UNIVERSITY OF LEEDS

School of Philosophy, Religion
and History of Science

What is Quantum Field Theory?

Idealisation, Explanation and Realism in High Energy Physics

James Duncan Fraser

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy

July 2016

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

Acknowledgements

First and foremost, I must thank my supervisor Steven French for his generous guidance. His open mindedness and enthusiasm have made my project much less fraught, and much more fun, than it might otherwise have been. I would also like to extend my gratitude to all of the staff and postgraduate students who have made the school of Philosophy, Religion and History of Science a wonderful place to work, and especially to Juha Saatsi who has influenced my thinking a great deal. I have been lucky enough to meet and discuss my ideas with a number of fellow travellers in the philosophy of quantum field theory at various conferences and summer schools. Thanks in particular to Leif Hancox-Li, Michael Miller and Porter Williams for sharing their work with me. I have also greatly benefited from email correspondences with the physicists and mathematicians Jean Bricmont, Anthony Duncan, Aernout Van Enter, Klaas Landsman, Tom Lancaster, Peter Lepage and Hal Tasaki. Finally, I would like to thank my partner Ruth and my family for their much appreciated support.

Abstract

Quantum field theory (QFT) poses a challenge to the orthodox methodological framework of the philosophy of science. The first step of investigating a physical theory, it is commonly assumed, is to specify some mathematical structures which constitute its formalism. In the case of QFT however this is a highly non-trivial task. There are a variety of prima facie distinct formulations of the theory: most strikingly, there is a gulf between the axiomatic formulations of QFT developed by mathematical physicists and the formalisms employed in mainstream high energy physics. In recent years this has led to debate about which version of the theory those interested in the foundations of physics ought to be looking at. This thesis offers a response to this problem. I argue that we should abandon the search for a single canonical formulation of QFT and instead take a more pluralistic stance which allows that different strands of the QFT programme may be appropriate starting points for addressing different philosophical questions. The overarching claim of the thesis is that, while the axiomatic approach to QFT may be the right framework for addressing some internal question raised by the QFT programme, formulations of QFT which incorporate cutoffs on the allowed momentum states can express all of the claims that we have any reason to believe on the basis of the empirical successes of high energy physics. In the course of this discussion I offer new perspectives on foundational issues like the status of relativity and unitarily inequivalent representations in QFT. Furthermore, I suggest that QFT and high energy physics offer fruitful grounds for discussing broader issues in the philosophy of science, and in particular the challenge of formulating a viable version of scientific realism.

Contents

List of Figures

Abbreviations

The following abbreviations are used in this thesis.

QFT — quantum field theory.

QED — quantum electrodynamics

QCD — quantum chromodynamics

SSB — spontaneous symmetry breaking.

KMS state — Kubo-Martin-Schwinger state.

Chapter 1

Introduction

1.1 Quantum Field Theory and Philosophy

Quantum field theory (QFT) ought to be of great interest to philosophers. It certainly has all of the features which scientific realists take to warrant belief in a scientific theory: it is among the most empirically successful frameworks in the history of science and purports to explain many features of the physical world, from the short range of nuclear forces to the existence of anti-matter. It is also arguably the most fundamental physical theory we currently have at our disposal; metaphysicians of a naturalistic bent should be taking notice. Furthermore, high energy physics promises to provide fertile ground for discussions of a plethora of epistemic and methodological issues, from the status of symmetry principles to analogical reasoning.

Unfortunately there seem to be serious obstacles facing any attempt to philosophically engage with QFT. For one thing, there are numerous foundational puzzles that need to be addressed if we are to draw any substantive conclusions from the theory. The puzzles posed by unitarily inequivalent Hilbert space representations and notoriously suspect renormalisation techniques, for instance, both present difficulties with understanding what the theory is telling us. But QFT also poses another sort of problem for the philosopher of science. The first step of the philosophical analysis of a physical theory, according to standard lore amongst philosophy of physics, is to identify the class of mathematical structures that constitute its formalism. While there is widespread agreement about how to do this for well trodden theories like quantum mechanics and general relativity in the case of QFT this task is far from straightforward. We can find a plethora of prima facie distinct characterisations of QFT in the physics literature. Most notably, there is a gulf between the axiomatic formulations of QFT found in the mathematical physics literature and the approaches employed in high energy physics phenomenology. Identifying the mathematical structures which are constitutive of QFT becomes a highly non-trivial, and philosophically loaded, task.

This issue has recently come to a head in an exchange of papers between Doreen Fraser and David Wallace (Fraser, 2009, 2001; Wallace 2006, 2011). These authors have very different views about which version of QFT philosophers ought to take as their starting point. Fraser advocates the primacy of axiomatic characterisations of QFT originating in the mathematical physics literature. Conversely, Wallace argues that if we are interested in what we can learn about the world from the success of high energy physics, axiomatic QFT is largely irrelevant and philosophers ought to be looking at the 'conventional' formulation of QFT employed in mainstream physics.

This controversy over the formulation of QFT might seem to dash any hopes of drawing wide reaching philosophical conclusions from QFT. If we can't even agree on what the theory is how can we hope to learn anything about the metaphysics or epistemology of science from it? This thesis develops a response to this problem with the ultimate aim of clearing the way for more fruitful interactions between philosophy and high energy physics. This introductory chapter sets out the background of my discussion and provides an overview of what is to come. The next section serves as an introduction to QFT and the problem with pinning down its theoretical identity. I then provide a summary of the core arguments (and assumptions) of the thesis.

1.2 What is Quantum Field Theory?

Given the controversy surrounding the formulation of the theory it will be useful to have a neutral way of speaking collectively about the various approaches to QFT found in the scientific literature. To this end I will use the term 'QFT programme' in this thesis to refer to all of the theoretical material which can possibly be construed as falling under the QFT rubric. Understood in this maximally permissive way the QFT programme will not only include approaches which resemble orthodox philosophical characterisations of a scientific theory but also the kind of computational and approximation techniques that pervade the practice of high energy physics. This section surveys this theoretical landscape and discusses why it presents a challenge for the philosopher. §1.1 gives an informal overview of the QFT programme. §1.2 discusses the construction of QFT systems and distinguishes between cutoff and continuum models. §1.3 turns to how philosophers have approached the QFT programme and summarises the key points of contention in the debate between Fraser and Wallace.

1.2.1 Overview of the QFT Programme

Owing to its multifaceted nature, the QFT programme does not lend itself to neat summary. I focus here on the approaches to QFT which will play an important role in the overarching argument of the thesis. Largely for convenience of exposition, my discussion here follows the historical development of the QFT programme.[1]

QFT became a distinct theoretical enterprise shortly after the birth of ordinary quantum mechanics. One obvious motivation driving its development in these early days was the need to bring classical field theories into the quantum fold. In order to obtain a fully quantum description of electromagnetic interactions, in particular, a quantised version of classical electrodynamics seemed to be required. A more high minded project which also played a crucial role in the formation of the QFT programme was the goal of unifying quantum and relativistic physics. Prima facie this has nothing to do with fields, but attempts to formulate a fully relativistic version of quantum mechanics turned out to lead to the introduction of 'matter fields' corresponding to fermionic and bosonic particles.[2] The consensus which emerged in this early period was that quantum electrodynamics (QED) ought to be a theory that couples a fermionic matter field, associated with the electron, to a quantised electromagnetic field.

The formulation of such a theory was beset with theoretical problems however. While free field theories could be analysed using methods familiar from ordinary quantum mechanics, interacting theories, like QED, proved to be much more difficult to make sense of. The first signs of trouble came when physicists attempted to apply standard perturbative methods to QED scattering problems. Expanding scattering observables in powers of the fine structure constant gave rise to badly behaved series which contained divergent terms beyond the first order. It was not until after the second world war that any real progress was made with this problem. In the late 1940's Dyson, Feynman, Schwinger and Tomonaga and others showed that the infinities that appear in naive perturbative expansions can be systematically removed via a process which came to be known as renormalisation.

In many ways this method was a huge success. Renormalised QED perturbation theory yields results that are not only finite but also in very close agreement with experimental data. Most famously, perturbative evaluations of the anomalous

[1] I do not claim to give anything more than a caricature of the complex history of QFT here. Schweber (1984) is an account of the formation of the perturbative approach to QFT. Cao (1998) is an account of the history of QFT with an emphasis on conceptual issues. Brown (1993) is a collection of papers on the history of renormalisation theory.

[2] Cao (1998) chapter 7 discusses the conceptual development of the notion of matter fields and second quantisation.

magnetic moment of the electron are found to agree with the measured value up to more than 10 significant figures! Furthermore, the perturbative approach to QFT that was made viable by the invention of the renormalisation procedure continues to dominate the practice of high energy physics to this day. The perturbative formalism is not without its problems however. For one thing, the renormalisation procedure was viewed with suspicion by many, including some of its creators.[3] Even putting these worries aside perturbation theory remains an intrinsically approximate approach to QFT with significant limitations. It rests on the assumption that interactions are weak which immediately raised worries about how strong nuclear interactions could be handled in the QFT programme. More generally, the fact that perturbation theory only yields approximate evaluations of particular physical quantities seems to render it incapable of answering deeper structural questions about interacting QFTs.

Dissatisfaction with the perturbative approach led to the rise of two new branches of the QFT programme in the second half of the 20th century. On the one hand, there was the emergence of the axiomatic approach to QFT in the mathematical physics community.[4] The idea here was to put QFT on a firm non-perturbative footing by writing down a set of mathematically precise conditions that any QFT model could be expected to satisfy. Two main axiomatisations of QFT on Minkowski space-time resulted from this project. The Wightman axioms treat quantum fields as operator valued distributions which are required to have certain relativistic properties: the fields are said to covariant under Poincaré transformations, a so-called microcausality condition is imposed, and so on. The other main axiomatisation of QFT, known as the Haag-Kastler axioms, originates in the algebraic formulation of quantum physics. In this framework QFTs are characterised by nets of operator algebras over Minkowski space-time which are required to satisfy similar conditions inspired by special relativity.[5] Though these two frameworks have not been shown to be equivalent in general I will speak as though they are interchangeable in this thesis—the distinction between them won't play an important role in my discussion.

[3]Feynman (2006, 128), for instance, wrote: "The shell game that we play... is technically called 'renormalization'. But no matter how clever the word, it is still what I would call a dippy process... I suspect that renormalization is not mathematically legitimate." See Cao (1998, 203-207) for a survey of the views of the other founders of QFT.

[4]Wightman's accounts of the origin of the axiomatic approach often emphasise the need to go beyond the limitations of the perturbative approach: "The motivation for the creation of an axiomatic field theory was simple enough: it was frustration. Perturbative renormalisation theory is exceedingly complicated and was even more so in the early 1950's. Everyone agreed that it should be generalized to be non-perturbative. However a simple-minded person asking: what is the problem and what is to be regarded as an acceptable solution of it? could get no answer from the available account of quantum field theory" (Wightman, 1979, 1000).

[5]The standard reference for the Wightman axioms is Streater and Wightman (1964). Haag and Kastler (1964) contains the original statement of the Haag-Kastler axioms. Haag (1992) and Halvorson (2006) are reviews of the algebraic approach to QFT.

While some important general results, such as the spin statistics theorem, can be derived from these axiomatisations of QFT the major rub with this approach has been the difficulty of demonstrating the existence of concrete models of the axioms. Some toy models have been constructed in a reduced number of space-time dimensions, but showing that QED and other realistic theories employed in high energy physics admit formulations which satisfy the Wightman or Haag-Kastler axioms is presently a rather distant aspiration—I will have more to say about this problem in §1.2.

Meanwhile, in mainstream physics, a different response to the limitations of the perturbative approach emerged in the form of the renormalisation group framework.[6] Drawing on ideas from condensed matter and statistical physics, Kenneth Wilson and others developed a non-perturbative framework for studying the behaviour of QFT systems at different energy scales, the key idea being that it is possible to remove high momentum degrees of freedom from a system and encode their effect in the dynamics of a low energy 'effective' theory. This approach not only shed new light on the perturbative renormalisation procedure but also led to the formation of new QFT models and ultimately the rise of the standard model of particle physics, which still reigns supreme today. Perhaps most importantly, the renormalisation group framework played a crucial role in accommodating strong nuclear interactions into the QFT programme. The renormalisation group analysis of quantum chromodynamics (QCD) revealed that the theory's coupling becomes weaker as the energy scale increases—a property known as asymptotic freedom—meaning that perturbative methods could in fact be successfully applied at sufficiently high energies. Furthermore, the methodological shift precipitated by the renormalisation group approach was instrumental in the development of lattice formulations of QFT which are nowadays employed in numerical simulations of the strong coupling region of QCD.

Much of the rest of this thesis will be concerned with the teasing out the relationship between the perturbative, axiomatic and renormalisation group approaches to QFT just described. It is worth emphasising before we move on however that this brief sketch of the QFT programme really only skims its surface. A variety of other theoretical frameworks can be found in the physics literature, some of which have yet to receive much attention from philosophers. QFT has been developed on non-relativistic (Bain 2011) and curved space-time backgrounds (Wald, 1994). Another important development has been the fruitful exchange of ideas between QFT and statistical physics, which has led to the application of QFT models in condensed matter physics and the QFT at finite temperatures formal-

[6] The history of the renormalisation group is a complex subject in its own right. Still, it is clear from Wilson's discussion in his Nobel lecture (Wilson, 1983) that concerns about the limitations of perturbative QFT were one important motivation for his work on the renormalisation group.

ism (Zinn-Justin, 2002). The approaches introduced in this section are already sufficient to pose significant challenges for the philosopher of physics, as we shall soon see. First though, I will switch from a historical to a formal mode and say something more precise about the problems associated with the construction of QFT models.

1.2.2 Cutoff and Continuum Models

How does one go about writing down a QFT model? As in quantum mechanics the starting point is invariably a classical system—in this case a classical field theory. There are, of course, many kinds of classical field theory which we might try to quantise. The theories which seem to do a good job of describing our world, like QED and QCD, are gauge theories: their Lagrangians consist of fermionic fields coupled to vector fields in such a way that local internal symmetry groups are preserved. While these theories will be central to the narrative of this thesis the technical and conceptual issues raised by gauge and fermion fields are largely tangential to the thrust of my discussion. Most of the features of QFT models which will be important for our purposes already arise in the case of bosonic scalar fields and it will often be expedient to discuss the technical issues in this more straightforward context. In particular, I will often refer to the simplest interacting field theory we can write down involving a real scalar field $\phi(x)$ over Minkowski space-time; the aptly named ϕ^4 theory. The Lagrangian density[7] of this theory is given by:[8]

$$\mathcal{L} = (\partial_\mu \phi(x))^2 - m^2\phi(x)^2 - \lambda\phi(x)^4. \tag{1.1}$$

The first two terms describe a free scalar field with mass m, and the final term is a quartic self interaction, whose strength is parameterised by the coupling λ.

This theory makes perfectly good sense as a classical field theory: the question is how to obtain a quantum version. Aping the canonical approach to quantisation familiar from ordinary quantum mechanics, an obvious strategy is to elevate $\phi(x)$, and its conjugate momentum density $\pi(x) = \partial\mathcal{L}/\partial(\partial_0\phi(x))$, to Hilbert space

[7]The Lagrangian density is related to the standard Lagrangian familiar from classical mechanics via $L = \int d^3x\mathcal{L}$, so the action of a classical field theory can be written $S = \int d^4x\mathcal{L}$. From here on I will simply refer to the Lagrangian density as the Lagrangian, as per the standard sloppy terminology of mainstream physics.

[8]I use the short hand $(\partial_\mu\phi(x))^2 = \partial_\mu\phi(x)\partial^\mu\phi(x)$, where $\partial_\mu = \partial/\partial x^\mu = \partial/\partial t - \nabla$ is the four-derivative, and adopt the signature $(+, -, -, -)$ for the Minkowski metric. Note that in later chapters I will often drop the explicit dependence on the spacial coordinate and simply write ϕ for the scalar field.

operators and impose the commutation relations:[9]

$$[\phi(x), \phi(y)] = [\pi(x), \pi(y)] = 0, \tag{1.2}$$

$$[\phi(x), \pi(y)] = i\delta(x - y)I, \tag{1.3}$$

where I denotes the identity operator.

An alternative method often employed in both the mainstream high energy physics and mathematical physics literature is known as functional, or path integral, quantisation.[10] The key quantity in this approach is the partition function, Z, which is identified with the functional integral:

$$Z = \int \mathcal{D}\phi e^{-\int d^4x\mathcal{L}}. \tag{1.4}$$

Informally, $\mathcal{D}\phi$ indicates that a sum is being taken over all possible configurations of the field. The significance of the partition function is that all of the system correlation functions can be derived from it. Correlation functions are the vacuum state expectation values of the fields at disparate space-time points of the form $\langle\Omega|\,\phi(x_1)\phi(x_2)...\phi(x_n)\,|\Omega\rangle$ (where $|\Omega\rangle$ is the vacuum state) and are arguably the most fundamental properties of a QFT system. For one thing, they are closely related to the S-matrix and scattering observables, to be discussed more in chapter 3, but crucially for our current purposes we can actually construct QFT models from a complete set of correlation functions. This means, in effect, that supplying a well defined functional integral expression for the partition function is tantamount to supplying a well defined QFT model.[11]

This all sounds promising enough in the abstract but complications arise when we try to apply these strategies to particular classical field theories. In the case of free field theories the quantisation process goes through fairly smoothly. Some subtleties arise in the canonical quantisation of free field theories but, with a bit of care, a well defined quantum system can be obtained via this method (it turns out that free fields on Minkowski space-time cannot be simply be converted to operators, as I suggested above, and instead become operator valued distributions

[9]These are the commutation relations in the Schrodinger representation; in the Heisenberg representation they become the, perhaps more familiar, equal time commutation relations. Note that I have already set $\hbar = c = 1$ here—see the discussion of dimensional analysis in QFT below.

[10]Peskin and Schroeder (1995) chapter 9 is a standard introduction to the way that function methods are employed in mainstream physics. Glimm and Jaffe (1981) develops the constructive field theory programme from a function integral perspective.

[11]Correlation functions go by different names in different traditions, and exactly how they are defined depends on the formalism in which you are working—in the Wightman axiomatisation of QFT, for instance, they are often called Wightman functions. For a discussion of the construction of QFT models from a complete set of correlation functions in the Wightman framework see Streater and Wightman (1964).

in the more sophisticated treatment).[12] It is also possible to give precise meaning to the sum over field configurations needed to define the partition function of a free quantum field theory in the functional integral approach. Furthermore, the quantum systems produced by these methods can be shown to satisfy the Wightman and Haag-Kastler axioms of the axiomatic approach.[13]

Things get a lot murkier when it comes to interacting theories however. Whether interacting QFT models exist on Minkowski space-time at all remains one of the great unanswered questions in theoretical physics. What we do know is that well defined quantum systems with interacting Lagrangians can be obtained if the degrees of freedom associated with arbitrarily high and arbitrarily low momentum Fourier modes of the classical field are omitted in the quantum version. I will refer to the quantum systems which result from this kind of procedure as cutoff QFT models, and the thresholds imposed on the high and low momentum states as ultraviolet and infrared cutoffs respectively.[14] Dimensional analysis helps clarify what is going on here physically. The standard convention amongst field theorists is to work in units that set $\hbar = c = 1$ (so-called natural units). This has the welcome effect of dramatically simplifying dimensional analysis. In this convention the physical units of any quantity can be expressed as a power of the energy, often called its mass dimension (since energy, momentum and mass are dimensionally equivalent in these units). Crucially, the dimension of length is an inverse power of the energy in this scheme so we can often talk interchangably of length and energy scales. Removing the high momentum modes of a theory then amounts to freezing out variations in the fields on arbitrarily short length scales. Conversely, removing low momentum modes corresponds to modifying the long range behaviour of the theory; in particular, a sharp cutoff at low momentum is equivalent to putting the theory in a finite volume box.

The most straightforward way of writing down a cutoff model is to formulate the theory on a finite volume lattice over continuous space-time, the simplest being a hypercubic lattice with lattice spacing a and volume V. Field variables ϕ_x can then be associated with each point x on the lattice and derivatives replaced by discrete difference expressions which converge to ordinary derivatives in the limit $a \to 0$, the standard choice being $\partial_\mu \phi(x) \to (\phi_{x+a\mu} - \phi_x)/a$, where μ is a unit vector. The effect of the discrete lattice spacing is to eliminate any field modes with momenta over $\Lambda = 1/a$—as expected the discretisation of the theory implements an ultraviolet cutoff. Obtaining a quantum theory is then straightforward. We can impose canonical commutation relations in the manner described above

[12] See Wald (1994, 43-46) for the details.

[13] For a discussion of free models of the Wightman and Haag-Kastler axioms respectively see Streater and Wightman (1964, 103-106) and Araki (1999, 216-221).

[14] I use the term 'cutoff model' to refer to a system with both an ultraviolet and infrared cutoff by default throughout the thesis.

and it is also possible to give precise meaning to the path integral over field configurations. In the case of a hypercubic lattice we can simply define:

$$\int \mathcal{D}\phi \to \prod_{x\in L} \int d\phi_x \tag{1.5}$$

where L is the set of points on the lattice.[15]

These lattice QFTs are nowadays directly employed in high energy physics, especially in the study of strongly interacting theories like QCD. The most successful non-perturbative computational method in contemporary particle physics phenomenology is based on running numerical simulations on finite volume lattices. (More precisely, these simulations are done on lattice QFTs formulated on Euclidean space, the results then being analytically continued back to Minkowski space-time via the so-called Wick rotation—see Smit (2002) chapter 2.) There are also other ways of constructing cutoff QFT models however which are sometimes employed in mainstream and mathematical physics.[16] In the path integral approach it is possible to impose a 'smooth cutoff' by suppressing the high/low momentum modes via a quickly decaying function rather than imposing a sharp threshold on the allowed momenta. Though there are different ways of implementing the cutoffs however it will turn out that many of the important features of these systems are not sensitive to the details of how the cutoffs are implemented. In what follows I will tend to treat lattice QFTs as the exemplar of a cutoff model.

While these cutoff QFTs may be well defined quantum systems, they might not seem to be very satisfying quantisations of classical field theories. Classical fields take values on a continuous space and have an infinite number of degrees of freedom: these cutoff systems have neither of these features. Furthermore, cutoff models invariably violate relativistic properties like Poincaré covariance which are often taken to be constitutive of what a quantum field theory is. It goes without saying that cutoff models do not satisfy the Wightman or Haag-Kastler axioms, since these systems are based on the assumption that a quantum field should be defined on arbitrarily small spacial regions and fulfil the demands of special relativity. A truly relativistic QFT would be directly defined on Minkowski space-time and therefore take values on arbitrarily large and arbitrarily small length scales. I will refer to this kind of system as a continuum QFT model.

The quest to construct interacting continuum QFT models has spawned a whole

[15] Putting gauge and fermion fields on a lattice raises additional technical issues which will not be discussed here. See Smit (2002) for the details of formulating these theories on a lattice.

[16] There are also a variety of regularisation methods used in the perturbative approach to QFT. As I will argue in chapter 3 however, the cutoffs employed in perturbative calculations are importantly distinct from cutoff QFT models.

branch of mathematical physics known as constructive field theory. What researchers working on this problem invariably do when attempting to construct interacting continuum QFTs is start with the a cutoff model and try to take the continuum limit, in which the ultraviolet cutoff is removed, and the infinite volume limit, in which the infrared cutoff is removed.[17] This turns out to be very difficult to do in a mathematically controlled way however. For a select class of interacting theories, in a reduced number of space-time dimensions continuum field theory models have been constructed in this manner. In particular, models of the Wightman and Haag-Kastler axioms have been shown to exist for ϕ^4 theory on a two dimensional analogue of ordinary Minkowski space-time.[18] But when it comes to realistic theories like QED and QCD, and even ϕ^4 theory in four dimensional Minkowski space-time, there are considerable obstacles to taking these limits, and even reason to think that they do not exist. The renormalisation group analysis of QED and ϕ^4 theory apparently suggests that an interacting theory cannot be produced by taking the continuum limit—this issue will be discussed further in chapter 4. The infinite volume limit, on the other hand, is related to the infrared behaviour of a theory, which is notoriously poorly understood in the case of QCD. The construction of a continuum formulation of the quantum Yang-Mills theory, a close cousin of QCD, is currently among the unsolved millennium problems set by the Clay institute (Jaffe and Witten, 2006).

In sum then, we know how to write down cutoff versions of realistic theories like QED and QCD, but the constructive field theory programme, as it stands, is very much a work in progress, with the few models that have been shown to exist in the continuum being quite far removed from real physics. I want to conclude this section with a brief note on the question of mathematical rigour in the QFT programme. Fraser refers to what I have just dubbed continuum QFT as the "formal" variant of QFT and consistently emphasises the higher standard of mathematical rigour found in the axiomatical and constructive approach to QFT. It is certainly true that mathematical physicists working on QFT are much more careful with mathematical subtleties than mainstream physicists, though this is true of any branch of physics. When it comes to cutoff and continuum QFT systems however rigour is not the relevant distinguishing factor. There is nothing unrigorous about the cutoff QFT models described above, indeed, as I pointed out, these structures are often employed in the constructive field theory

[17] To be precise, this is usually done in Euclidean space. The Osterwalder-Schrader reconstruction theorem (Osterwalder and Schrader, 1973, 1975) is then appealed to establish the existence models of the Wightman axioms.

[18] Glimm and Jaffe's (1985) famous treatment of the ϕ^4 theory on two dimensional Minkowski space-time follows this approach. Rivasseau (2014) surveys results in the constructive QFT programme which have been obtained in this way. In fact, no gauge theories have been constructed in any number of space-time dimensions—see Douglas (2004) some discussion of gauge related issues in the constructive field theory programme.

programme. The motivation for seeking interacting QFT systems defined directly on Minkowski space-time is not simply about rigour, it is primarily a physical enterprise. For this reason, continuum QFT is a more apt name for these sought structures, in my view, than formal QFT. While mathematical rigour is one parameter which varies greatly among different traditions found within the QFT programme a theme of my later discussion will be that it is not really the key issue in the formulation debate.

1.2.3 The Formulation Debate

We have now seen enough to appreciate the force of the methodological challenge flagged at the outset. A range of sub-programmes and species of physical model fall under the jurisdiction of the QFT programme. How should philosophers approach this heterogeneous body of theory?

In the early days of the philosophy of quantum field theory literature this question simmered away in the background but was rarely explicitly discussed. Philosophers engaged with different approaches to QFT, sometimes drawing conflicting conclusions. Paul Teller's early work on QFT, for instance, focused on the methods found in mainstream textbooks and the practice of high energy physics (Teller, 1995). At the same time another approach to the philosophical study of QFT emerged based on the axiomatic approach to QFT, and in particular the algebraic formulation of the theory. Robert Clifton, Hans Halvorson and others derived a number of foundationally significant results from the Haag-Kastler axiomatisation of QFT.[19] These two starting points seemed to lead to very different perspectives on QFT however. Teller advocates an interpretation of QFT based on the existence of particle-like quanta, while results in the algebraic approach seem to show that no particle based interpretation of QFT is possible.[20]

The debate between Fraser and Wallace has thrown the dilemma facing the philosophical interpreter into sharp relief.[21] Each author marshals a complex set of considerations, often grounded in more general debates in the philosophy of science, in support of formulations of QFT which they both take to be theoretically distinct. Many of these issues will be taken up in later chapters but I will summarise the key points of contention in the debate here to set the scene for what

[19] Many of Clifton and Halvorson's seminal papers on algebraic QFT are collected in Clifton (2004).

[20] One problem for particle interpretations of QFT raised in this tradition comes from non-localisability results originating in Malament (1995) and made more precise in the algebraic QFT context by Halvorson and Clifton (2002). Another challenge, championed by Fraser (2008), comes from the non-existence of Fock space representations in interacting algebraic QFTs.

[21] MacKinnon (2008) and Kuhlmann (2010) are other contributions to the recent debate about which version of QFT ought to be the basis of philosophical enquiry.

is to come.

Fraser (2009) distinguishes three formulations of QFT: the "infinitely renormalised", "cutoff" and "formal" variants. The cutoff and formal variants correspond to what I have called cutoff and continuum QFT models above. The infinitely renormalised variant refers to the characterisation of QFT that seems to be implicated by the perturbative renormalisation procedure. I will have more to say about this supposed variant in chapter 3 but, since it is not central to the core debate, I put it aside for now. Limiting herself to the ϕ^4 theory on two dimensional Minkowski space-time, which, as I mentioned above, has been shown to admit a continuum formulation, Fraser claims that the choice between these three formulations of QFT is underdetermined by empirical evidence. Rather than drawing an anti-realist moral however, Fraser argues that theoretical considerations break the underdetermination in favour of the continuum formulation. The key idea here is that the project of unifying quantum and relativistic physics is integral to QFT, and while continuum models fulfil this commitment, cutoff models do not.[22] For Fraser then it is the axiomatic approach to QFT, and its constructive field theory offshoot, which philosophers ought to be engaging with.

Wallace takes a very different view of the situation facing the philosopher of physics. According to him we ought to be focusing on the approach to QFT found in mainstream physics which he calls, variously, "Lagrangian" (Wallace, 2006) and "conventional" (Wallace, 2011) QFT and associates with what I called cutoff QFT models.[23] There is something inapt about both of these terms in my view. The status of the Lagrangian representation of the dynamics is not a crucial issue here and the suggestion that it is a cutoff formulation of the standard model that is directly employed in high energy physics is somewhat misleading. Most of the empirical predictions of the QFT programme do not come directly from cutoff QFT structures, and those that do are based on simulations on Euclidean lattices, as briefly described above. What Wallace really wants to say, I think, is that the empirical and explanatory successes of contemporary high energy physics can be attributed to cutoff QFT models, a claim which he bases on results gleaned from the renormalisation group framework. Given these successes, and the relative failure of the axiomatic/constructive programme to produce empirically viable models, the claim is that it is cutoff formulations of the standard model, and other realistic QFTs, which the scientific realist should be focusing on. Crucially

[22] Fraser (2009) frames this argument in the context of the apparent threat to the consistency of QFT posed by Haag's theorem. The axiomatic approach is said to supply a principled response to this problem, while imposing cutoffs avoids the problem in an ad hoc manner.

[23] More precisely, Wallace focuses on systems with an ultraviolet cutoff. He has a different position regarding the status of the infrared cutoff, as will be discussed in more detail in chapter 6.

though, Wallace does not take cutoff QFT models to be exact descriptions of reality. On the contrary, on his account QFT is an "intrinsically approximate" theory which makes no claims about the way the world is at arbitrarily small length scales (Wallace, 2006, 46).

There are a number of factors underlying this dispute; I will briefly touch on four key issues here which will be picked up on, in one way or another, in later chapters. One of the key points of contention is the significance of renormalisation group methods. According to Wallace, the renormalisation group approach has largely solved the problems which the axiomatic QFT project was originally trying to address. Recall that axiomatic QFT originated, at least in part, in an attempt to provide a non-perturbative framework which could shore up the inadequacies of the perturbative approach. Wallace claims that the renormalisation group analysis has resolved the mathematical and conceptual puzzles with the perturbative approach and thus left axiomatic QFT without a tangible motivation. Fraser expounds a radically different perspective on the renormalisation group. In her view these methods represent an important breakthrough in articulating the "empirical content" of QFT—that is, they facilitate the extraction of empirical predictions—but they do not tell us about the theoretical content of QFT—they do not help to answer the question of which version of QFT should be taken to be canonical for the purposes of foundational research.

Another issue over which they disagree is the force of Fraser's underdetermination argument, just sketched. Wallace points out that, while Fraser's underdetermination scenario might hold in the case of toy models like ϕ^4 theory in two dimensions, it does not hold for realistic theories QFTs like QED and QCD because they have no known continuum formulation. Fraser counters this line by stressing that the continuum field theory project is a work in progress and that the mere possibility of continuum versions of realistic QFT is sufficient for her arguments in support of the axiomatic approach to go through (Fraser, 2011). A third issue, which plays into their diverging assessments of this argument, concerns the status of special relativity within QFT. As Fraser makes clear in her presentation of her underdetermination argument, she takes the project of unifying relativity and quantum theory to be integral to QFT. By contrast, Wallace does not take a commitment to relativistic space-time structure because on his view QFT makes no claims about the fundamental nature of space-time at all. He claims that cutoff QFT's are approximately Poincaré covariant and this is all we can reasonably demand given their status as effective theories.

Finally, besides these issues which are entangled with the details of the QFT programme, we can also detect broader methodological differences underlying the dispute. As Fraser makes explicit, she and Wallace have different understandings of the whole project of interpreting QFT:

> By ‘interpretation’ I mean the activity of giving an answer to the following hypothetical question: ‘If QFT were true, what would reality be like?’ In contrast, the interpretive question that Wallace focuses on is ‘Given that QFT is approximately true, what is reality (approximately) like?’ (Fraser 2009, 558).

While Fraser makes clear that approximate notions are not permitted in her approach to interpretation, the thrust of Wallace claims about QFT is that the theory furnishes us with partial representations of the world. Yet while both authors seem to recognise that these background methodological assumption play an important role in the debate very little is said in support their preferred construals of the philosophy of QFT project.

In sum, we have two radically different visions of what QFT is, and how the project of philosophically engaging with it should proceed. On the face of it, this disagreement over the formulation of QFT needs to be resolved before any philosophical conclusions which might interest the general philosopher of science or the metaphysician can be drawn from high energy physics.

1.3 Overview of the Thesis

This thesis develops a response to the challenge posed by the theoretical diversity of the QFT programme. I suggest that we should move away from the idea that there is a canonical formulation of QFT that needs to be determined before philosophical engagement can begin. Instead, I put forward a more pluralist methodological approach to the QFT programme which allows that different theoretical frameworks may be the right starting points for tackling different philosophical issues. The central claim of the thesis is that while the axiomatic formulations of QFT advocated by Doreen Fraser represent a relevant framework for addressing important questions about the internal relationship between theoretical principles, and in particular the issue of whether quantum theory can be unified with special relativity in a deep sense, cutoff QFT models have the capacity to express all of the claims about the world we have any right to believe on the basis of the QFT programme. The perspective advocated here is, therefore, perhaps closer to Wallace’s side of the debate in its spirit; indeed, much of the discussion can be understood as clarifying and building upon ideas found in Wallace’s work. But I will suggest that we can consistently maintain an important role for the axiomatic and constructive approach to QFT within this framework. My answer to the question of whether philosophers should be engaging with axiomatic or cutoff characterisations of QFT is ultimately ‘both’.

Chapter 2 sets out this methodological stance in the abstract. Picking up on the aforementioned disagreement between Fraser and Wallace regarding the objectives of the interpretive project, I subject some orthodox ideas about the methodology of the philosophy of physics to sustained criticism and put forward a broader framework in their place. Specifically, I argue that investigations of a particular physical theory ought to include consideration of a theory’s relationship to the actual world and not merely its truth conditional semantics. Furthermore, drawing on comparisons with similar debates in which stalemates arise over the existence of multiple inequivalent formulations of a theory, I argue that the right response is to relativise our choice of formalism to a particular line of philosophical enquiry. This leads to a more piecemeal approach to the different stands of the QFT programme. Rather than asking which approach to the theory is ‘correct’ we can simply ask how each bears on particular philosophical issues.

Chapter 3 applies this methodological stance by offering an analysis of the perturbative approach to QFT. Drawing on resources from the philosophical literature on scientific modelling I put forward an analysis of perturbative QFT according to which it is in the business of producing approximations rather than explicitly constructing quantum systems. This leads to a reassessment of the foundational standing of the perturbative approach. The real problem with QFT perturbation theory, and perturbative renormalisation, on this account is not that it is mathematically unrigorous or inconsistent, as philosophers have often thought, but that it leaves us without a physical explanation for its success.

Chapter 4 turns to the renormalisation group approach to QFT. Contra Fraser, and others who advocate an instrumental reading of the role of the renormalisation group in the QFT programme, I argue that this approach has important philosophical implications. On the one hand, I discuss how the renormalisation group framework bears on the question of when continuum QFT models exist. More importantly however, I claim that it has ramifications for the kind of representational success which ought to be attributed to QFT models. I sketch how the renormalisation group framework solves the aforementioned problem with the perturbative approach by providing a justification for the perturbative renormalisation procedure and argue that this story should lead us to take cutoff QFTs to accurately describe reality on scales far removed from the cutoffs. What emerges is a view of QFT models as coarse-grained representations of the world.

The later chapters of the thesis are concerned with fleshing out this perspective on the representational success of QFT models and defending it against possible objections. One worry about taking cutoff QFT models seriously in this way is that they violate Poincaré covariance and therefore appear to be in conflict with relativity. Chapter 5 examines this issue in depth. I suggest that all the evidence is that cutoff QFT models are approximately Poincaré covariant in a sense which can be made fairly precise. A more subtle issue about the theoretical status of relativistic principles in the QFT programme remains however. Applying the aforementioned pluralist approach to variant formulations of a theory, I distinguish the question of whether QFTs can be constructed on four dimensional Minkowski space-time from the question of what we ought to believe about the world given the success of the standard model. In the later context I claim that there is no license to demand Poincaré covariance down to arbitrarily small length scales.

Another potentially worrying feature of cutoff models is that, if both an ultraviolet and infrared cutoff are imposed, the resulting system has a finite number of degrees of freedom. This rules out the possibility of unitarily inequivalent Hilbert space representations which some philosophers of physics take to play a role in high energy physics. In particular, Ruetsche (2003, 2006, 2011) has argued that these resources afforded by the limit of infinite degrees of freedom are indispensably involved in the explanation of spontaneous symmetry breaking phenomena in quantum theory and consequently need to be taken physically seriously. I take this challenge head on by questioning whether the role played by novel properties of infinite systems in this context really motivates taking them representationally seriously. I put forward an approach to understanding how standard theoretical accounts of spontaneous symmetry breaking in statistical mechanics and high energy physics are successful which avoids the need to reify these properties.

Finally, chapter 7 reconnects with more general issues in the philosophy of science and tackles the question of what a realist stance towards QFT amounts to. Rejecting claims about the pragmatic status of cutoff QFTs, I claim that scientific realism is actually easier to make sense of in this context. Realists have long struggled to spell out the sense in which our best physical theories are 'approximately true'. I argue that the renormalisation group helps address this problem by providing a framework for identifying which features of QFT models latch onto the world in the absence of detailed information about future physics. I also address Doreen Fraser's claim that underdetermination between cutoff and continuum QFT models stands against a realist reading of the latter. I suggest that the kind of underdetermination we find in the QFT programme actually strengthens my claims about the partial representational success of QFT models.

Before we enter into the substance of the discussion I would like to conclude this chapter by making explicit some of the background assumptions driving the approach taken in this thesis. First of all, I will be assuming, rather than arguing for, a broadly realist view of science and, no doubt, beg the question against the anti-realist in many places. In general, I have conceived this project in a constructive spirit—my goal has been to develop a well motivate perspective on the QFT programme, not to establish that it is the only viable view of it. My guiding question has not been should we be realists about QFT, but what might a realist view of QFT look like.

In terms of my general perspective on the philosophy of science, I have tended to err on the side of caution, taking the more conservative route wherever possible. This is not to say that QFT cannot form the basis for revisionary ideas regarding the structure and semantics of scientific theories, emergence and reduction, explanation, and so on—I will flag the possibility of taking a less conventional route at various junctures. An overriding theme of this thesis however will be that we do not need to reinvent the philosophy of science in order to make sense of QFT. We can make progress in analysing the QFT programme using fairly orthodox philosophical resources. I will tend to speak in the language of the dominant semantic approach to the structure of scientific theories, though I will be fairly fast and loose with the terms 'theory' and 'model'—I will, for instance, speak interchangeably of the ϕ^4 theory and the ϕ^4 model. Both of these terms, of course, have been taken to mean a great many things by different philosophers, and are often used in conflicting ways in the scientific literature. For the most part, subtleties about exactly how these terms should be understood will not impinge on my discussion and I have opted for sloppiness over pedantry.[24]

One final point I want to touch on concerns the term 'effective field theory', which has perhaps been conspicuous by its absence in the discussion so far. While the perspective taken in this thesis is in many respects inspired by the effective field theory turn in high energy physics since the 1970s, I will make sparse use of this term. This is more a choice of convenience than of principle. The effective field theory language has connotations which are not necessarily helpful. In the physics literature effective field theory is often used to refer to a specific set of ideas and techniques within the perturbative approach to QFT that I will not be discussing in detail here, and is typically not applied to lattice QFT, which I will be discussing. On the other hand, the term can be used in a much broader sense to refer to any theory which has a limited domain of descriptive adequacy. On this interpretation however, the meaning seems to me to be, at best, a minor precisification of what philosophers of science usually mean when they say that a theory is approximately true (I expand on this connection in chapter 7). Indeed the thrust of this thesis will be that taking QFTs to be effective field theories in this sense is not a radical departure from previous ideas in physics and

[24]Though I will have a few words to say about the term 'model' in chapter 2.

the philosophy of science; it is similar to the relationship which, say, Newtonian gravity bears to the world. Furthermore, in the philosophical literature, the notion of an effective field theory has been associated with controversial ideas about anti-reductionism (Bain, 2013a) and anti-foundationalist metaphysics (Cao and Schweber, 1993) which are best avoided in my view. While I will be developing an effective field theory view of QFT then, I have, for the most part, chosen to express these ideas in more mundane philosophical terminology.

Chapter 2

Formulation and Interpretation

2.1 In Search of a Methodological Framework

This chapter is concerned with how philosophical engagement with theoretical physics in general ought to be understood. We have already seen some reasons for thinking that the methodological framework of the philosophy of physics needs to be reassessed in light of the challenges posed by the QFT programme. It is usually assumed that the formalism of a theory must be fixed at the outset of foundational investigation, but, as I described in the previous chapter, in the case of QFT philosophical problems already seem to enter at this stage. Furthermore, one of the factors I identified as underlying Fraser and Wallace's dispute over the formulation of QFT in the previous chapter was that they disagree about the nature of the philosophical project itself. For Fraser interpreting QFT is a matter of spelling out what the world would be like if the theory were true; for Wallace the key philosophical issue is what we can say about the actual world given that it is approximately described by QFT. It seems inevitable then that questions about the broader goals of the philosophical engagement with the physical sciences will need to be broached if we are to make progress in resolving the formulation problem.

I start by examining a common story about the task of the philosopher of physics, sometimes called the standard account of interpretation, which takes philosophical engagement with a physical theory to be centrally concerned with the possible worlds in which it is true. Whatever the merits of this doctrine, I argue that the full range of philosophical issues that physical theories actually provoke cannot be expressed in these terms (§2.2). On the one hand, I make the case that our understanding of a physical theory turns on epistemic judgements about its representational success, which are not included in a specification of the possible worlds it exactly describes. §2.3 sketches a broader framework for regimenting philosophical investigations of physical theories which incorporates claims about how the content of the theory relates to the physics of the actual world. On the other hand, I claim the standard account does not provide a mechanism for addressing the existence of variant formulations of a theory. We know that this is a problem in the case of QFT, but I suggest that this issue also occurs in many other branches of physics. In §2.4 I argue that the best response to the existence of inequivalent formalisms is to simply abandon the idea that there is a single canonical formulation of a physical theory. Instead we should look to justify our interest in one formulation over another within the context of a particular philosophical project. §2.5 considers the ramifications of these ideas for approaching the QFT programme and sketches how they will be employed later on in the thesis. Throughout I make use of examples drawn from classical physics to illustrate and motivate my claims.

2.2 The Standard Account and its Deficiencies

The philosophy of physics is often said to be about interpreting physical theories, but what does this mean exactly? Much of the time nobody worries too much about this question. Foundational issues raised by particular physical theories often seem to be adequately stated in purely local terms. On the occasions that philosophers of physics do address the nature of the interpretive project in the abstract however, they typically endorse a view which Ruetsche (2011) calls the standard account. On this view, an interpretation of a physical theory is a specification of the class of possible worlds in which it is true.[1]

This notion of interpretation, I suggest, is at base a semantic one. Philosophers of physics, the advocate of the standard account will say, are concerned with what physical theories say about the world—that is, what they mean. A popular view in the philosophy of language is that the semantic content of a proposition is given by the set of possible worlds in which it is true. Furthermore, this approach to content fits well with, and is, in fact, incorporated within, many versions of the dominant semantic view of scientific theories. Within the purview of the standard account, the process of stating a physical theory can be thought of as a two step process.[2] In the first, structural phase, a mathematical formalism is specified, which for proponents of the semantic view means a class of mathematical

[1]Explicit endorsements of the standard account can be found in Belot (1998), Rickles (2008) van Fraassen (1991), and Ruetsche (2011). Often this view is tacitly assumed: the talk of classical and relativistic worlds in Earman (1986) being one example among many.

[2]I follow Ruetsche (2011) here. Van Fraassen often tells a similar story, and endorses a characterisation of the second semantic phase in terms of possible worlds, in his presentations of the semantic view—see, for instance, van Fraassen (1980, 1991).

structures. In the case of general relativity, for instance, the relevant structures are typically taken to be of the form $(M, g_{\mu\nu}, T_{\mu\nu})$, where M is a manifold, and $g_{\mu\nu}$ and $T_{\mu\nu}$ are tensor fields defined on it that satisfy the Einstein equations. At this stage all we have are mathematical objects. For the theory to make claims about the physical world, as general relativity surely does, we need to add some prescription for assigning physical content to these structures. According to the standard account, this second semantic phase boils down to providing a characterisation of the possible worlds which the theory exactly describes, where these worlds are taken to be physical, as opposed to merely mathematical, objects.[3] The operative question then is the one Fraser endorses in her discussion of QFT: what would the world be like if the theory were true. Indeed, it seems clear that the standard account of interpretation guides Fraser's approach to QFT.

Of course, if an interpretation is an assignment of meaning to a theory's formalism then it is surely not only of interest to philosophers. Physicists are not doing pure mathematics, and their theories clearly come equipped with some prescriptions about what the relevant mathematical structures represent—the models of general relativity are taken to describe space-time and the distribution of matter-energy over it. If the advocate of the standard account is being careful then, they ought to say that the job of the philosopher of physics is to make the physical content of already partially interpreted theories more precise—which for them just means saying something more precise about the ways the world could be according to the theory. The debate surrounding the hole argument, for instance, is quite naturally cast in these terms: assignments of physical content to general relativity which treat diffeomorphic models as representing distinct possible worlds are shown to be problematic, with the residual puzzle being how best to characterise the worlds actually described by the theory.

I do not want to reject the standard account outright here. There are interesting motivations for moving away from a possible world semantics for physical theories which I am not unsympathetic with.[4] I am not convinced that we need to take this revisionary path to make sense of the QFT programme however; we can make progress in understanding QFT and high energy physics while accepting the orthodox conception of physical content embodied in the standard account. What I do want to object to however is the claim, often loosely implied in presentations of the standard account, that *all* philosophical issues posed by physical

[3]How the notion of possible worlds should be understood is, of course, a matter of debate amongst metaphysicians and philosophers of language, but I take it nothing I have to say here is sensitive to these issues.

[4]Frisch (2005) and Curiel (2011) both reject the standard account and a possible world semantics for physical theories. Mark Wilson's work on the semantics of scientific theories also seems to have similar morals (Wilson, 2006). My hope is that the key claims of this thesis will be recoverable, in one way or another, within revisionary approaches to theory semantics.

theories boil down to questions about the class of possible worlds in which they are true.

There are two challenges to viewing the standard account as a comprehensive methodological framework in this sense that I want to focus on here. The first arises in the initial structural phase. On the naive picture just sketched, one starts a philosophical investigation of a theory by writing down mathematical structures which constitute its formalism, perhaps copying from a nearby textbook. As we have seen, in the case of QFT matters are not this straightforward; there are a number of mathematical frameworks that physicists and philosophers take to fall under the QFT umbrella, and the question of which set of models constitutes the canonical formulation of QFT is itself a matter of considerable controversy. Really though matters are not this straightforward for any physical theory: there are Newtonian, Lagrangian and Hamiltonian formulations of classical mechanics; flat space and curved space formulations of Newtonian gravitation; Boltzmanian and Gibbsian formulations of statistical mechanics; and so the list goes on. Before we can engage with a theory in the way that the standard account prescribes there is often a philosophically non-trivial question about which formalism we ought to start with.

There are, of course, well trodden ways of responding to the proliferation of alternative formulations of a theory. Many of the apparently distinct formulations of physical theories we find in scientific practice are commonly thought to be equivalents. Newtonian, Lagrangian and Hamiltonian mechanics are often taken to be equivalent rather than rival formulations of classical mechanics, for instance.[5] But while appealing to some notion of theoretical equivalence may abate the problem of variant formalisms somewhat I will argue in §2.4 that it does not completely resolve the issue. For many branches of physics there are multiple prima facie reasonable ways of choosing a collection of mathematical models that cannot plausibly be viewed as theoretically equivalent—QFT being one, but by no means the only, such case.

The second, and perhaps more fundamental problem I see with the standard account as a putatively comprehensive methodological framework is that it is simply not the case that all philosophical issues raised by physical theories concern their semantics. The standard account gives us a way of understanding the physical content of a theory but it does not say anything about how that content relates to the actual world. This is often presented as a matter to be resolved by very general epistemic considerations. The scientific realist believes that the theory is true,

[5]It is difficult to find uncontroversial examples of equivalent formalisms (we shall see some possible reasons why in §2.4). Many philosophers do not agree that the Newtonian, Lagrangian and Hamiltonian frameworks are theoretically equivalent—see the debate between North (2009) and Curiel (2014) over which version should be viewed as fundamental.

and therefore that the actual world is among the class of possible worlds admitted by the theory, while the constructive empiricist remains completely agnostic about the theory's truth. It seems to me, however, that presenting the question as an all or nothing affair, settled entirely in the abstract, does not leave room for the fine grained claims that philosophers, and scientists themselves, actually make about the representational success of particular physical theories.

Consider Newtonian gravitation theory for instance (putting aside issues raised by the existence of variant formulations for now). The standard account asks us to delimit a class of worlds which are exactly described by the theory. These will presumably be worlds in which there are absolute (i.e. frame independent) facts about the temporal duration between events, which, in the wake of relativistic physics, we do not believe actually obtain. It is not just that we know that our world is not a Newtonian one then, we also make more selective judgements about the features of the theory's interpreted models which misrepresent real physical systems. Spelling out what Newtonian gravity models get right about the world also seems to be a thoroughly theory specific enterprise. There is a complex story to tell about which features of Newtonian gravity models are embedded within models of general relativity, for instance, which has been explored in detail in foundational research—see, in particular, Malament (1986). This suggests that the question of how a physical theory latches onto the world hangs on the peculiarities of its structure, its relationship to other theories, and the kind of empirical success it enjoys, and not just on the general epistemic issues discussed in the traditional scientific realism debate. Furthermore, I think that these local claims about which aspects of a theory's models faithfully represent are, in an important sense, part of our understanding of the theory which is missed if we limit ourselves to the resources of the standard account. Motivating this view in the context of the QFT programme will be one of my overarching aims in the chapters that follow.

The next section develops this critique of the standard account and puts forward my preferred way of incorporating theory-world relations into debates about particular physical theories. With this framework in place, I come back to the problems posed by variant formalisms in §2.4.

2.3 Accommodating Theory-World Relations

I am not the first to think that the absence of any mention of the actual world in the standard account leads us into trouble. Gorden Belot, for one, raises this point in his discussion of the Aharonov-Bohm effect (Belot, 1998). The Aharanov-Bohm effect is often taken to undermine the traditional view of classical electrodynamics as a theory about the electromagnetic field. In one way or another, we seem to need to attribute additional physical significance to the vector potential to accommodate the relevant physics. As Belot points out, this is difficult to make sense of on the standard account. Classical electrodynamics and quantum mechanics presumably describe disjoint sets of possible worlds so why should a quantum phenomenon tell us anything about the content of electrodynamics? Furthermore, since classical electrodynamics is a false theory, it is not clear what we learn from this kind of interpretive revision. Aren't we simply making claims about possible worlds we know to be distinct from our own on the standard account? What does this have to do with the physics of the Aharanov-Bohm effect?

Belot's solution is to add another layer to the standard account. We are not only interested in the possible worlds which can be associated with a theory's formalism in Belot's expanded framework, we also need to specify an additional structure representing "our evaluation of the relative merits of each of the possible interpretations" (Belot, 1998, 551). Belot takes this to be an ordering relation on the space of possible worlds encoding our assessment of their closeness to actuality. With this new apparatus in place he claims we can do a better job of assimilating the philosophical ramifications of the Aharanov-Bohm effect. What this quantum phenomenon is telling us, on this account, is that possible worlds consisting of electric field configurations and assignments of holonomies to closed curves in space are closer to the actual world than worlds consisting of classical electromagnetic fields. This is supposed to explain how the revaluation process triggered by the Aharanov-Bohm effect, and the philosophical investigation of false theories more generally, can tell us something about our world. In so far as claims about the closeness of other possible worlds constrain where the actual world might sit in the space of all possible worlds they are claims about actuality.

Belot is on the right track here; we do need some way of talking about the relationship a theory bears to the actual world if we are to do justice to the full range of philosophical issues raised by theoretical physics. I doubt that a closeness relation on the space of possible worlds is the right way go however. The problem I see with Belot's approach is that the relation of closeness holds between worlds in toto—saying that one world is close to ours, or is closer than some other world, does not tell us in what specific respects it resembles actuality. Notice that the Newtonian gravity case introduced in the previous section displays many of the features of the Aharanov-Bohm effect case, admittedly in a prima facie less puzzling form. It is similarly unclear on the standard account why the advent of general relativity ought to affect our understanding of Newtonian gravity at all, since again these theories deal with distinct sets of possible worlds. The lesson I

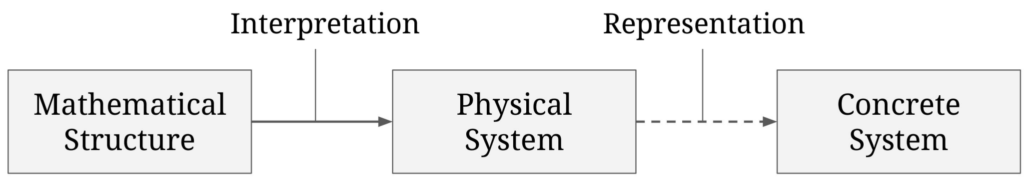

Figure 2.1: Schematic of the relationship between mathematical structures, physical models, and concrete systems out in the world.

drew from this case, however, was that we need to be able to accommodate more selective claims about which parts of an interpreted theory's content accurately describe reality. The blanket claim that Newtonian worlds are close to ours does not capture these more fine-grained judgments about the features of its models that faithfully represent.

I will now sketch a framework which tries to do a better job of incorporating the partial nature of theory-world relations. As in Belot's approach, I suggest we think of the process of stating a physical theory as dividing into three parts. The first two are those already covered by the standard account: a class of mathematical structures is delimited and some physical content is assigned to these structures (again, I am suppressing the issues raised by variant formalisms until the following section). What we have at this stage is a class of physical systems, associated with non-actual possible worlds on the standard account, whose properties can be investigated independently of any putative relation they bear to our world. The final step, which I call the epistemic phase, amounts to providing some prescriptions about the doxastic attitude we ought to take towards these physical systems as representations of the actual world. In particular, and in contrast to Belot's proposal, for me this comes down to identifying specific features of a theory's interpreted physical systems that should be taken to get things right. (Figure 2.1 illustrates the relationship between a theory's formalism, physically interpreted systems and the world I have in mind here.)[6]

Exactly how the representation relation between scientific models and concrete physical systems should be understood is an issue which has provoked much de-

[6]This framework helps to clarify a common confusion regarding the use of the word 'model' in the philosophy of science. Proponents of the semantic view of theories often use the term in, roughly, its logical sense, to mean a structure which satisfies a set of axioms. In much of the literature on scientific modelling however, the term model is used to mean a representation of another system. The picture I have sketched here supports Thomson-Jones' (2006) claim that these two meanings are, in principle, distinct. In the first stage of stating a physical theory we have a collection of mathematical models, in the logical sense, which via a physical interpretation become physical models, again in the logical sense. Only in the last phase do these systems become models of real systems in the representational sense. As outlined in the chapter 1 however, I will tend to use the term model in a deliberately loose way when these kind of distinctions are not philosophically important.

bate in recent years.[7] I take it, however, that the fact that scientific models get some features of their targets right while getting others wrong is something all sides agree on and are ultimately trying to capture. Rather than taking sides in this debate then I will sketch how I understand the process of carrying out the epistemic phase in schematic terms which can presumably be fleshed out and made more precise in the language of any viable approach to scientific representation. When an interpreted mathematical structure is used to represent a concrete physical system, I suggest, there are three attitudes we can take towards its properties. We can take a property to be representationally faithful; that is we can hold that the target system actually has that property. We can take a property of the model to misrepresent; following the terminology Jones (2005) I will refer to properties of a model which are not believed to be instantiated by the target system as idealisations. Finally we can remain agnostic about whether the target system has or does not have the property in question.[8] Carrying out the epistemic phase of analysing a physical theory, on this simple minded approach, just means saying something about which properties of its models ought to be treated as descriptively accurate, idealisations or neither.

With this machinery in hand we can make sense of what is going on in the Newtonian gravity case (and, I hope, the impact of the Aharanov-Bohm effect on classical electrodynamics as well, though discussing the subtleties of this case is a story for another time). What has changed with the rise of general relativity is that much of the content of Newtonian gravity theory is now thought to misdescribe the world. Newtonian gravity models, equipped with a suitable semantics, make claims, not just about absolute durations, but also about gravitational effects at long distances and in the presence of very large masses, which contemporary physicists believe to be false. As representations of actual gravitating systems then these properties of Newtonian models are viewed as idealisations, in the terminology just introduced. Yet other features of Newtonian gravity theory are preserved in relativistic gravitational theory and can, at least provisionally, be treated as descriptively accurate. Examining this issue in detail is a task which calls for sophisticated engagement with the intertheoretic relations that hold between the two theories. As Malament (1986) and Barrett (2006) describe

[7] Model-theoretic (French, 2003), similarity (Geire, 2004) and inferential (Saurez, 2004) accounts of scientific representation have been put forward, to name but a few of the more prominent approaches.

[8] A similar tripartite distinction is found in the partial structures approach to scientific representation developed in da Costa and French's (2003). Their approach is certainly one way that the output of the epistemic phase might be formalised. The Bayesian formalism may also provide a useful framework here. In this context we could generalise from belief, disbelief and agnosticism to a continuum of attitudes towards particular aspects of the theory's content. I leave the possibility of clarifying the claims of this thesis within these formalisms as a topic for future investigation.

in detail, the sense in which Newtonian gravity latches onto the world is most clearly illuminated in the geometric formulation of the theory, where the embedding of its content within general relativity can be made mathematically precise. In my view spelling out how a physical theory relates to the world in this way represents an important contribution of our understanding of the theory and is certainly a legitimate task for the philosophy of physics.

If we enlarge our methodological framework in this way the question which immediately arises is how we go about making judgements about the representational success of different features of a physical theory. As with its semantics the kind of representational success a theory enjoys is not a mere philosophical afterthought. The development of a theory in scientific practice invariably involves prescriptions about how it describes and distorts the world, either explicitly or tacitly. Indeed, these assumptions are often crucial in guiding how scientific models are applied in practice. The philosophers job can, again, be viewed as a matter of precisification, in this case the project being to pin down the kind of qualified representational success which can rationally be attributed to a particular theory. As I have already urged, a variety of local factors are important in carrying out this task. A theory's relationship to other, typically more fundamental, theories often plays an important role in identifying which of its features faithfully represent and which are idealisations, the relationship between Newtonian and relativistic gravitation theory we have already discussed being a prime example of this kind of reasoning. In other contexts direct empirical information about the target system can play a role in determining the representational status we assign to a model. As will be discussed further in later chapters, physicists often employ infinite volume systems, generated by the so-called thermodynamic limit, to model the properties of concrete systems in statistical mechanics. We don't need any overarching theory to tell us that this is an idealisation—the pieces of iron whose behaviour we are trying to describe clearly have a finite volume.

What about cases in which we do not have access to a more fundamental theory, or relevant information about the target system, which can inform our assessment of a theory's representational success however? After all, this is more or less the situation we find ourselves in with respect to QFT—extant quantum gravity theories are arguably too speculative at present to provide any reliable guidance about which aspects of current QFT models will be preserved in future physics. There is a general principle which the realist philosopher of science can appeal to here. Namely, that we ought to believe in those parts of a physical theory which underwrite its predictive and explanatory successes. This is not a new idea. The claim that we ought to reserve our optimism for the parts of a theory which drive its success has become the dominant way of understanding the realists' epistemic commitments in the face of challenges posed by theory change, classic exposition of this sort of doctrine being Kitcher's (1993) distinction between idle and working posits and Psillos's (1999) "divide et impera" strategy. The motivation for demarcating the belief worthy content of a theory in this way is a local form of the no miracles intuition, the thought being that we need to take the aspects of a theory which essentially contribute to it accurate predictions to latch onto the world in order to explain these successes.[9]

In order to apply this criterion in practice though we need to say something relatively precise about what it means for a constituent of a theory to underwrite, or essentially contribute to, its empirical success (I focus on predictive success here, but I will have more to say about explanatory success in chapter 6). One standard approach, found in Psillos (1999) for instance, is to say that something contributes to a prediction if it plays an indispensable role in its derivation—that is if there are no alternative ways of deriving the result from the theory which do not make use of this property. For our purposes, however, a slightly different modal characterisation of what it means for something to underwrite a prediction will be useful (though I won't be defending it as a universally applicable definition here). The idea is that a property of a model contributes to a particular prediction if varying it necessarily spoils its empirical accuracy in this respect. To use terminology which has recently gained currency in the scientific explanation literature, we might say that it makes a difference to the relevant observable property.[10] Conversely, a property of model does not contribute to an empirical success, and consequently is not supported by it, if it can be varied without affecting its predictive accuracy. A simple example may help to illustrate what I have in mind here. The simple pendulum model, which neglects air resistive forces correctly predicts that the time period of a pendulum is approximately proportional to the square root of its length. Adding a small non-zero air resistive force to the model does not affect this prediction and consequently the assumption that air resistance is zero does not contribute to this empirical success. I will be plying this criterion into service in my discussion of the representational success of QFT models in chapter 4.

[9]Of course, the anti-realist will balk at this. Van Fraassen holds that the empirical success of particular theory calls for no explanation while the empirical adequacy of science as a whole can be accounted for without invoking truth or accurate representation (van Fraassen, 1980, 40). As I set out in chapter 1 however, I will be flagrantly begging the question against the anti-realist .

[10]Where Psillos's deductive characterisation of what it means for something to 'essentially contribute' to a prediction is perhaps allied to deductivist approach to explanation, my characterisation here is closely related to modal accounts of explanation defended in Woodward (2005) and Strevens (2008). There are subtle issues relating to how the notion of difference making should be understood which have been glossed over here—exploring the connections with the broader debate surrounding scientific explanation is an interesting topic for further discussion. The criterion suggested here is also inspired, in part, by the discussion of idealised models in Saatsi (2016a).

I close this section with a few clarificatory remarks about the methodological system which has emerged from the discussion so far. First, a terminological note. Should the expanded framework just described be understood as a new account of what it means to interpret a physical theory, or as an injunction that interpretive questions do not exhaust the purview of the philosophy of physics? This is ultimately just a matter of how we choose to use the term interpretation. There is a case to be made for keeping interpretation as a purely semantic notion however. For one thing, this maintains the connection with interpretation in the standard logical sense and avoids potential misunderstandings with advocates of the standard account. But, perhaps more importantly for my purposes, using interpretation to refer to the full three stage process as a whole will likely lead to a blurring of semantic and epistemic issues which should be recognised as separate in my view. To keep this distinction clear I will continue to use interpretation in its standard account sense and treat the specification of a theory's representational status as an extra-interpretive matter.

I should also make clear that I am not claiming that all philosophical issues raised by physical theories concern their relationship to the world. As far as I am concerned, we would do best to avoid general proclamations about what is, and is not, a legitimate topic for philosophical enquiry and interesting issues arise at all three levels of analysis just distinguished. Philosophers of physics certainly sometimes investigate physical theories in a purely internal mode, as it were, remaining at the semantic level and bracketing any questions about how the theory relates to the world. We can investigate the status of space-time in string theory, for instance, while remaining completely agnostic about its claims about the world—I will have more to say about this kind of enquiry in the next section. The point here is simply that epistemic issues are raised in the analysis of particular physical theories and we need resources to deal with them in our philosophy of physics toolbox. In fact, I suspect that epistemic questions do play a more important role in extant debates in the philosophy of physics than they are often given credit for, though supporting this claim in detail would require a thorough survey of the field. In any case, I will be arguing that they certainly play an important role in understanding the QFT programme in the chapters that follow.

Finally, I should stress again that my primary objective here has been to marshal sufficient resources for the task at hand not to put forward a stand alone account of the structure, semantics and epistemic status of scientific theories. Accordingly, the framework I have put forward here should be understood in a provisional spirit. I have already said, I am open to the possibility of moving away from a possible world semantics for physical theories, and I am similarly open minded about the details of how theory-world relations are made precise. My hope is that the claims about the QFT programme put forward in later chapters are not overly sensitive to the way these issues are handled and should be recoverable within many approaches to theory semantics and scientific representation. To reiterate once more, the most important point here is that philosophers of physics need to take seriously both semantic and epistemic issues in their investigations of physical theories and that this necessarily means going beyond the standard account. The following section returns to the other problem I identified with taking the standard account to provide a comprehensive guide to the philosophy of physics: the existence of variant formalisms.

2.4 Accommodating Variant Formalisms

On both the standard account and the expanded framework put forward in the previous section the first step of stating a physical theory is to supply a mathematical formalism. As I pointed out however, there are many alternative formulations, not just of QFT, but of almost any physical theory you care to mention. How should we proceed in the face of this kind of diversity at the formal level? One reason this is not always perceived as a pressing issue is that many of the formulations of physical theories found in the scientific literature are assumed to be equivalent. If Newtonian, Lagrangian and Hamiltonian mechanics are simply different ways of stating the same hypotheses about the world, for instance, then the dilemma about choosing between them disappears. But while identifying equivalent formulations of a theory in this way is surely part of the story here, it does not completely resolve the issues raised by the existence of variant formalisms, as I will now argue.

The most immediate challenge facing any attempt to diffuse the problem of variant formalisms along these lines is the need for a principled way of identifying equivalent formulations of a theory: that is, an account of theoretical equivalence. Unfortunately this has proven to be a contentious philosophical project in its own right. I will briefly give my take on the extant debate surrounding this issue before going on to claim that no plausible account of theoretical equivalence identifies all prima facie reasonable formulations of a physical theory—in many contexts the question of how to accommodate variant formalisms inevitably remains.

Many accounts of theoretical equivalence take it to be a purely formal relation which holds between uninterpreted mathematical formalisms. On the syntactic view of theories propagated by the logical positivists, which took physical theories to be axiomatic systems, it was natural to try to analyse theoretical equivalence in terms of logical equivalence. It was soon realised however that logical equivalence is too stringent a condition as it automatically distinguish theories formulated in different languages. Various attempts were made to develop more nuanced definitions of theoretical equivalence based on the possibility of translating between theories expressed in different languages, the account given in Quine (1975) being a prominent example. Owing in part to the lack of syntactic axiomatisations of interesting physical theories the concrete implications of these definitions were difficult to assess and no syntactic account of theoretical equivalence gained widespread acceptance.

Indeed, the difficulty with capturing the identity conditions of scientific theories in the syntactic framework was one of the original motivations driving the rise of the semantic view, which takes a theory's formalism to consist of a class mathematical structures rather than an axiomatic system.[11] In this framework isomorphism becomes the natural vehicle for an analysis of theoretical equivalence. The definition which immediately suggests itself is that two theories are equivalent if each mathematical model of the first is isomorphic to a mathematical model of the second. Again however, this doesn't seem to be quite right. Halvorson (2012) has recently raised technical objections to analysing theoretical equivalence in terms of isomorphisms between mathematical models alone. Drawing on various examples from pure mathematics, Halvorson argues that isomorphism on its own is both too strong, in the sense that it equates theories which ought to be distinct, and too weak, in the sense that it distinguishes theories which ought to be equivalent. One potentially counterintuitive consequence of the definition given above, for instance, is that Lagrangian and Hamiltonian mechanics turn out to not be theoretically equivalent because their models are structurally distinct—the configuration space of the Lagrangian approach is a tangent bundle while the phase space of Hamiltonian mechanics is a cotangent bundle.[12]

Halvorson seems to be optimistic that the failure of previous formal definitions of theoretical equivalence is due to a lack of mathematical sophistication and points to category theory as a possible framework for a more promising approach. While attempts to analyse theoretical equivalence in category theoretic terms are in their infancy there are general reasons to doubt that this programme will yield a universally adequate formal definition of theoretical equivalence. Coffey (2014) has recently argued, to my mind convincingly, that theoretical equivalence is not a relation which holds between uninterpreted formalisms at all. One simple minded argument to this effect is that we do not intuitively identify theories

[11] Suppe often presents the semantic view as superior in part because it does a better job of capturing the identity criteria of scientific theories. He writes, for instance: "As actually employed by working scientists, theories admit of a number of alternative linguistic formulations... As such, scientific theories cannot be identified with their linguistic formulations." (Suppe, 1989, 82).

[12] Having said this some authors take the structural distinctness of Hamiltonian and Lagrangian mechanics to indicate that they are genuinely distinct theories—see North (2008) and Curiel (2014).

which employ the same mathematical formalism in radically different scientific domains. To borrow an example from van Fraassen (2014), though the same field equation—the diffusion/heat equation—is used to describe both the diffusion of a gas and spread of heat throughout a system, it seems highly unnatural to suggest that these models are one and the same. Another point raised by Coffey is that it is difficult to make sense of the widespread disagreement among philosophers about which theories are in fact equivalent if it is a purely formal question. Quine (1975) takes theories in euclidean space and 'Poincaré theories', which mimic the effects of euclidean geometry by positing compensating universal forces, to be paradigm cases of equivalent theories, while Psillos (1999), and others, take them to be obvious inequivalents. It seems clear that the dissagreement here is not about the formal relations between the two theories but rather springs from different interpretive judgements about their physical content.

In opposition to previous formal approaches Coffey (2014) advances a view which takes the relata of the theoretical equivalence relation to be completely interpreted theories.[13] On this account, two theories are equivalent if they share the same physical content, which on the standard account's analysis of physical content simply means that they pick out the same set of possible worlds. In my view this is the most plausible approach to theoretical equivalence currently on the market (though my preference for this account will not play a crucial role in the claims of this section, or the broader argument of this thesis). It captures the key intuition that two theories are equivalents if they say the same thing about the world, since, in the framework I have been developing here, a theory's mathematical formalism on its own says nothing about the world at all until it has been furnished with a physical semantics. A potentially disheartening feature of this account however is that it makes judgements about theoretical equivalence hostage to difficult questions about the semantic interpretation of particular physical theories.

What is the upshot of this debate for the challenge posed by variant formalisms? For one thing, it suggests that appealing to the notion of theoretical equivalence is not going to be a quick fix. On either of the approaches just discussed deciding whether two formulations of a theory are equivalent is a non-trivial task. For believers in a formal approach careful mathematical work will presumably be needed to determine whether the formalisms employed in scientific practice can, in fact, be identified, while on Coffey's account equivalence claims become

[13]Note that rejecting formal accounts of theoretical equivalence certainly does not mean that the mathematical relations between theoretical formalisms are irrelevant to questions of equivalence. The existence of isomorphisms between formulations of a theory, for instance, will influence how we understand their content. In this way, bringing category theory and other resources from pure mathematics to bear on physical theories may well inform our judgements about which particular theories are equivalent, in so far as this programme impacts on our understanding of their physical content.

a function of semantic interpretation, meaning that the formal and semantic phases of specifying a theory will have to be explored in tandem for a number of variant formalisms. Either way, the result is a more sophisticated picture than the naive reading of the standard account originally described above, on which questions about a theory's formalism are simply decided by consulting a physics textbook.

More importantly for our purposes however, we are now in a position to see that none of these approaches to theoretical equivalence, or any plausible future account, will identify all of the variant formulations of physical theories as equivalents. There are genuinely inequivalent formulations of many (perhaps all) physical theories which are on equal footing as far as the mathematics is concerned. The QFT programme offers a clear example in which variant formulations of a theory cannot reasonably be identified. Though philosophers disagree about the nature of theoretical equivalence there are certain general conditions which everyone agrees equivalent theories must satisfy. One of them is empirical equivalence—they must be indistinguishable by any possible observation. Cutoff and continuum formulations of QFT do not seem to pass this test;[14] as we shall see in later chapters, while cutoff theories may be in very close empirical agreement with continuum models in many contexts they always assign different values to observable quantities. Furthermore, cutoff and continuum QFT differ with respect to fundamental theoretical properties: continuum QFTs have infinite degrees of freedom, while QFT models with an ultraviolet and infrared cutoff violate Poincaré covariance and have finite degrees of freedom. On any viable physical interpretation of these structures then, they surely cannot be saying the same thing about the world.

While QFT is a particularly contentious and complex case in which inequivalent formulations of a theory arise, it is certainly not the only such case. In fact, the need to accept the existence of multiple inequivalent formulations of a theory has been noted in a number of recent debates in the philosophy of physics. As Vickers (2013) discusses in depth, this is a recurring theme, and frequent source of miscommunication, in debates about inconsistency in science. To take just one example, Frisch (2005) notoriously argues that classical electrodynamics is inconsistent, while Belot (2007) and Muller (2007) contest this. As Vickers (2008, 2013) points out however, these authors make their respective claims about different, and clearly inequivalent formulations of the theory—crucially, Frisch and his detractors employ different versions of the Lorentz force law. Consequently, there is a sense in which the participants in this debate are talking past one another: Frisch, Belot and Muller give different verdicts on the consistency of

[14]This point is discussed further in §7.4.

classical electrodynamics, but they also take 'classical electrodynamics' to refer to different things.[15]

Another context in which this kind of ambiguity arises is the recent debate surrounding determinism in Newtonian mechanics, and in particular the 'Norton dome' model introduced by Norton (2008). This simple indeterministic system was designed to show that Newtonian mechanics is not, in fact, a deterministic theory. Some philosophers have responded by arguing that the dome model, as Norton describes it, is not a legitimate Newtonian system at all. Zinkernagel (2010), for instance, argues that, on a proper interpretation of Newton's first law, the indeterministic solutions appealed to by Norton are ruled out. As Fletcher (2012) describes in detail, the dome is an example of a system which admits indeterministic solutions because the relevant force function violates the property of Lipschitz continuity. There are various ways of demanding Lipschitz continuity during the formulation of Newtonian mechanics which would eliminate the indeterminism here and there may well be good motivations to construe the theory in this way—indeed, if Zinkernagel is right, this may be a more accurate formalisation of Newton's historical conception of the theory. Again, the question of whether Newtonian mechanics is deterministic or not comes down to which version of the theory you chose, with different philosophers in the debate adopting different inequivalent formulations.

Naively applying the standard account to these debates seems to lead to a stalemate as each side is talking about inequivalent theories, and therefore different sets of possible worlds. After a detour via the theoretical equivalence debate then, we have essentially arrived back at the question we started out with: how should philosophers negotiate the diverse range of genuinely distinct formulations of physical theories?

The response advanced by Vickers and Fletcher in the context of their respective debates, which I would like to endorse as a general methodological stance here, is that we move away from the idea that we need to identify one particular formulation of a physical theory as canonical. Instead, we should adopt a more pluralist perspective according to which theory names like 'classical electrodynamics', 'Newtonian mechanics', and so on, are not univocal—there are many inequivalent formalism which legitimately fall under these headings, none of which has any a priori claim to be the fundamental version of the theory.[16] Ac-

[15] A subtlety in this case is that the question of how inconsistent theories should be characterised within accounts of the structure of scientific theories is controversial. Frisch (2005) actually takes the inconsistencies he identifies in classical electrodynamics to be an argument against the semantic view of theories, while da Costa and French (2003) claim that inconsistent theories can be accommodated with a model-theoretic framework.

[16] In fact, Vickers (2014) goes further than this, urging that we should stop using theory names, and even the general theory concept, entirely, an approach he calls theory eliminativism.

knowledging the viability of many distinct formalism as objects of philosophical study leads to a realignment of the traditional methodological picture previously sketched. Rather than thinking of the mathematical formalism of a theory as being fixed before philosophical engagement begins we should instead look to justify our interest in one of a number of possible formulations of a theory in the context of a particular philosophical project.

In the case of the classical electrodynamic debate, for instance, focusing on the question of whether 'the theory' is inconsistent simply leads to a purely semantic dispute about how classical electrodynamics should be defined. Theory names are ultimately just labels however, and which set of assumptions we choose to call classical electrodynamics has no intrinsic interest. As Vickers urges, the kind of question we should be focusing on is whether the inconsistent set of assumptions identified by Frisch tells us anything important about, say, the nature of scientific theorising, and framing the issue as one about the essence of classical electrodynamics distracts attention away from these issues. Similarly, the question of how QFT should be defined lacks substance on its own and ultimately leads us down dead ends if we attempt to answer it in isolation. The key issue should instead be how the various approaches to QFT bear on the diverse philosophical issues raised by the QFT programme.

I won't try to give a comprehensive characterisation of the many possible lines of philosophical enquiry that might lead us to focus on one precisification of a theory's formalism over another, but there is one distinction I would like to touch on, as it will become important later on.[17] On the one hand, philosophers of physics often pursue what we might call internal questions. These are queries about the our theories themselves—their properties, the relations between them, and so on—that do not hang on the representational relationship they bear to the real world. Norton's presentation of the indeterminism debate in classical mechanics is arguably an example of this kind of project. He is interested in whether classical mechanics is an deterministic theory, not in the causal structure of the actual world—as Norton (2008) makes explicit, his dome is not intended to be an accurate representation of any physically realisable system. In terms of the methodological framework I have advocated, this kind of internal enquiry takes

While I am sympathetic to the motivations for this methodological proposal I do not follow this doctrine to the letter in this thesis. I continue to use the terms 'cutoff QFT', 'continuum QFT' and 'the QFT programme', but I as Vickers says that these names are simply labels and don't carry any "conceptual weight" on their own (Vickers, 2014, 120).

[17] To mention another sort of investigation, you might be interested in the conceptual issues which impacted on a particular episode in the history of science and therefore want to focus on the formal structures actually employed by historical scientists. In the context of this kind of project it could be perfectly legitimate to object to Norton's dome on the grounds that it violates the original meaning of Newton's first law, for instance, as Zinkernagel (2010) seems to do.

place in the structure-specifying and semantic phases of analysing a physical theory. On the other hand, philosophers ask what we might contrastively call external questions about how a theoretical model latches onto the world, or what we ought to believe about the world given the successes of a particular theory. This clearly takes place in the epistemic phase. If we take the pluralist route I am endorsing there is no particular reason to think that all of the internal and external questions raised by a branch physics should be directed at the same set of mathematical structures; indeed, I will be suggesting in the coming chapters that these issues come apart in the case of QFT.

The take home message of this section then is that instead of trying to eliminate the formal diversity found in many areas of physics we need to learn to accept it. I have tried to motivate the thought that this is not really as problematic as it might initially seem. There is ultimately no good reason to think that there should be a single fundamental formulation of a physical theory which is the right starting point for all philosophical enquiries into a particular field of science and giving up on this idea helps to focus the dialectic on the issues which really matter. Of course, this shift in thinking gives us a new perspective on the formalism problem with QFT set out in chapter 1. I close by summarising the implications of the claims of this chapter for the broader project of the thesis.

2.5 A Look Ahead

To summarise the conclusions of this chapter, I have argued that an adequate framework for regimenting work in the philosophy of physics needs to go beyond the standard account in at least two ways. First, we need to acknowledge that philosophical debate about particular physical theories concerns not only the semantic interpretation of the theory but also its epistemic status as a representation of the actual world. Specifically, I argued that we must recognise the partial nature of the representational success of scientific theories and put forward a rough framework for capturing this. Second, we need a more sophisticated approach to dealing with variant formalisms than a naive reading of the standard account seems to suggest. Instead of assuming that we can identify a canonical, universally applicable, formulation of a physical theory we should allow that there can be many prima facie viable choices of formalism to investigate. The central question is then which set of structures is the right vehicle for exploring a particular philosophical issue.

As I mentioned at the outset, philosophers of physics often address the foundational puzzles raised by particular physical theories in a local fashion without explicitly embedding their discussion within a broader programme. Why worry about these general methodological issues then? While I think that the approach put forward in this chapter may be helpful for clarifying other debates in the philosophy of physics, my central goal here was to lay the groundwork for approaching the QFT programme, and the conclusions endorsed here certainly have important implications for this project. For one thing, we are now better placed to make sense of Fraser's and Wallace's conflicting statements about the proper goal of the philosophy of QFT. Recall that for Fraser the operative question was what the world would be like if QFT were true, while Wallace focuses on what the actual world is like given that QFT is approximately true. In terms of the framework I have presented here, we can see these are not really rival accounts of the interpretive project but rather different aspects of the analysis of a physical theory: Fraser's question is a semantic one, while Wallace's is an epistemic one. This raises the possibility that at least some of what Wallace and Fraser say about QFT is mutually consistent.

More generally, the pluralist approach to alternative formulations of physical theories set out in the previous section leads to a new perspective on the problem raised by the many approaches to QFT found in the scientific literature. The challenge posed by the diversity of the QFT programme should no longer be viewed as one of picking out the correct formulation of the theory against a backdrop of distracting impostors. Rather, the task is to tease out how the various approaches to QFT bear on broader philosophical issues. This suggests a more piecemeal and constructive approach to the project. Rather than asking, is this approach to QFT the right framework for philosophical work, we can simply ask how this approach should be understood, and how it bears on particular philosophical issues raised by the QFT programme, without worrying about whether it is the 'correct' formulation of the theory. I take the analysis of perturbative QFT put forward in the next chapter to exemplify this approach. This methodological shift also opens up the possibility that different theoretical approaches to QFT are the appropriate frameworks for answering different kinds of philosophical questions. In fact, I will be endorsing this kind of claim in what follows. In brief, a key claim of the thesis will be that the axiomatic approach to QFT is a suitable starting point for various internal questions raised by the QFT programme, in the sense discussed in the previous section, while cutoff QFT models are appropriate objects of study for addressing the external issue of what we ought to believe about the world on the basis of the empirical successes of high energy physics.

Chapter 3

The Real Problem with Perturbative Quantum Field Theory

3.1 Three Worries about the Perturbative Approach to Quantum Field Theory

This chapter initiates our study of the QFT programme by examining the perturbative approach to QFT. There are good motivations for starting here. As I described in the introduction, the perturbative approach played a crucial role in the development of the QFT programme. Both the axiomatic and renormalisation group approaches to QFT developed, at least in part, as a response to perceived limitation with the perturbative formalism. It seems plausible, therefore, that a better understanding of the perturbative approach will also help us to understand the significance of these later developments. There is also a more simple minded reason for philosophers to be interested in perturbative QFT: it remains the main source of empirical predictions within the QFT programme. The famously accurate prediction of the anomalous magnetic moment of the electron is just one among many perturbative results which form the backbone of particle physics phenomenology. Consequently, when it comes to the question of representational success, and the kind of ontological claims which are warranted by the QFT programme, perturbation theory will surely have to be addressed in one way or another.

Philosophers have often shied away from perturbative methods in their discussions of QFT however.[1] To some extent this reticence is understandable. The

[1] There are, of course, exceptions. Teller (1989, 1995) and Huggett and Weingard (1995) examine perturbative renormalisation from a philosophical perspective. There has also been discussion of virtual particles (Weingard, 1988; Redhead 1988) and Feynman diagrams (Meynel 2008; Wüthrich 2012) in the philosophical literature, which necessarily engages with perturbative

renormalisation method at the heart of the perturbative approach has long been viewed with suspicion by physicists and philosophers alike. We can distinguish three main problems sometimes raised in the literature that seem to present a barrier to philosophical engagement.

i) *The rigour problem.* One concern about the perturbative approach to QFT is that it lacks mathematical rigour. Amongst philosophers this kind of critique is typically rooted in the assumption that a certain standard of mathematical rigour is necessary (or at least strongly desired) before interpretive work can get off the ground. I suspect this attitude towards perturbative methods is behind Halvorson's comments about the lack of a "mathematically intelligible description of QFT" underlying the practice of mainstream high energy physics (Halvorson, 2006, 731). Fraser (2009) raises this kind of objection more explicitly in her discussion of what she calls the "infinitely renormalised" variant of QFT, which she associated with the practice of taking a momentum cutoff to infinity in the course of perturbative calculations. Fraser argues that this procedure is ill defined and complains that "the standard criticism levelled against unrigorous theories—that they are difficult to analyse and interpret—certainly applies in this case" (Fraser, 2009, 543). One putative obstacle to engaging with perturbative QFT then is that it is on too flimsy ground mathematically for foundational issues to be properly addressed.

ii) *The consistency problem.* A prima facie more severe charge sometimes brought against perturbative QFT is that it runs foul of inconsistency. A notorious theorem due to Haag, Hall and Wightman (henceforth Haag's theorem) seems to show that standard perturbative calculations rest on an inconsistent set of assumptions. While Earman and Fraser (2006, 306) want to allay worries about the consistency of interacting QFTs in general in light of Haag's theorem, they claim that the result does "pose problems for some of the techniques used in textbook physics for extracting physical prediction from the theory", and perturbation theory in particular. Fraser even suggests that physicists may be tacitly employing inferential restrictions which prevent contradictions from being derived when they perform perturbative calculations (Fraser, 2009, 551). The threat of inconsistency might provide further reason to be wary of taking perturbative QFT seriously from a foundational perspective, and especially of drawing ontological conclusions from the perturbative formalism.

iii) *The justification problem.* Even putting these concerns about the internal coherence of perturbative QFT aside however, a further puzzle seems to re-

QFT.

main. In physics textbooks perturbative renormalisation is often presented as a mathematical trick for removing divergent terms in the naive perturbation series. But why are these infinities there in the first place, and what justifies the procedure used to remove them? The renormalisation procedure has often been seen as lacking in physical motivation. Wallace is pointing to this sort of worry when he suggests that (at least when it was first developed) perturbative renormalisation made little "physical sense" (Wallace 2011, 117). Similar complaints about the lack of a physical picture underlying the renormalisation procedure can be founded dotted throughout the literature.[2] A final reason for trepidation about engaging with the perturbative approach to QFT then is that it appears to be problematically ad hoc.

In light of these problems one might conclude that philosophers should simply remain silent about the use of perturbative methods in QFT, at least until physicists and mathematicians have developed a coherent picture of what is going on. This chapter puts forward a different response. Drawing on Norton's (2012) discussion of the distinction between idealisation and approximations I argue that perturbative calculations should be understood as producing functions which approximate physical quantities, rather than mathematical structures semantically interpreted as physical systems. Taking perturbative QFT to be in the business of producing approximations, in Norton's sense, leads to a reassessment of all three of the worries just identified. The rigour and inconsistency problems, in particular, lose much of their bite. Perturbative renormalisation does not furnish a class of continuum QFT models on this view not because of mathematical imprecision but because this is not the objective of the procedure. In fact, I suggest that there is nothing problematically unrigorous about the perturbative expansions found in QFT textbooks. Furthermore, this understanding of QFT perturbative expansions provides resources for an effective response to Haag's theorem. In brief, the result does not undermine standard perturbative calculations because they do not posit the existence of a model satisfying the relevant set of inconsistent assumptions.

What this analysis does not do however is resolve the justification problem. What ends up being the really salient puzzle about QFT perturbation theory is why it is so successful—why, that is, the approximations produced by the perturbative approach are in fact good ones. I will suggest, however, that this is a issue which philosophers can contribute to rather than a reason to eschew perturbative

[2] A typical example of this kind of attitude is found in McMullin (1985). After arguing that scientists are generally dissatisfied with modifications to a model which lack physical justification he comments: "The renormalisation techniques required in quantum field theory because of the assumption introduced by the original idealization of the electron as a point particle would be a case in point" (McMullin, 1985, 261).

methods, and later chapters take up this challenge. The plan for this chapter is as follows. §3.2 sketches the basic features of perturbative QFT. §3.3 sets out the distinction between idealised models and approximations in general terms. §3.4 and §3.5 make the case for understanding QFT perturbation theory as a method for producing approximations and discusses the ramifications of this view for the three putative problems just outlined.

3.2 The Perturbative Framework

Perturbative QFT is a huge subject in its own right. I focus here on the general aspects of the approach which are most relevant to assessing the aforementioned worries about its foundational respectability.[3] I first describe how perturbative expansions of QFT observables are set up, stressing the role of the interaction picture and the challenge posed by Haag's theorem. I then discuss the divergences which appear in naive implementations of the expansion and the renormalisation procedure which is used to tame them.

3.2.1 Expanding the S-matrix

I claimed in §1.2.2 that the really fundamental quantities of a QFT model are its partition function and the correlation functions it generates. The key object we are trying to get at in the perturbative approach to QFT however, at least in the first instance, is the so-called S-matrix. In scattering theory scattering events are represented as transitions from localised initial and final states at asymptotic times. The S-matrix is the operator which maps incoming states $|\alpha\rangle_{in}$ at $t \to -\infty$ onto outgoing states $|\beta\rangle_{out}$ at $t \to \infty$:

$$S_{\beta\alpha} = {}_{out}\langle\beta|S|\alpha\rangle_{in}. \tag{3.1}$$

The S-matrix is of paramount importance to particle physics phenomenology because S-matrix elements associated with particular classes of in and out states are closely related to scattering cross sections, the quantities measured in collider experiments. But the S-matrix also plays an indispensable role in perturbative evaluations of other quantities of theoretical interest. In particular, once we have the perturbative expansion of the S-matrix in hand it can be used to evaluate correlations functions as well.[4]

[3] Peskin and Schroeder (1995) is a classic text which focuses on the perturbative approach to QFT. Useful discussions are also found in Duncan (2012) and Lancaster and Blundell (2014).

[4] This means that perturbative QFT is rooted in scattering theory in a way in which other computational approaches are not. Monte Carlo simulations in lattice QFT, for instance, can

There are two major obstacles to getting at the S-matrices of the realistic QFTs studied in high energy physics. On the one hand, there is the issue of giving concrete meaning to the S-matrix in the first place. To substantiate the above definition a precise characterisation of the in and out states, and the space in which they live, is needed. As we have seen however, identifying mathematical structures which are constitutive of interacting QFTs in four dimensions is a task plagued by technical and conceptual hurdles—this is, after all, what the formulation debate is all about. Besides this question however, interacting field theories inevitably engender more practical concerns about computational tractability. Interaction terms lead to non-linear equations of motion, and finding exact solutions to empirically successful QFTs are typically out of the question. However they are characterised structurally then, there is going to be a difficulty with computing the S-matrix elements of theories like QED and QCD in practice.

In the face of these problems the perturbative strategy is to use what we already know about free QFTs to generate expressions for the S-matrix elements of weakly interacting theories. The construction of free QFT models is much better understood than their interacting counterparts—we know how to write down continuum models in this case. Furthermore, the resulting models can be exactly solved; we can determine the system's spectrum, get explicit Fock space representations of the field operators and obtain the S-matrix analytically (unsurprisingly, the S-matrix of a free field theory is trivial). The idea then is to split the Hamiltonian of an interacting theory into a free part, H_0, and an interaction potential, V, parameterised by a 'coupling' g (it will be convenient to use the Hamiltonian representation of the dynamics here):

$$H = H_0 + gV. \tag{3.2}$$

We then construct power series in g for the theories S-matrix elements whose coefficients can be computed using the explicit representations of the fields afforded by the exact solution of the free model associated with H_0. If g is sufficiently small the hope is that the first few terms of this series yield accurate predictions of experimentally observed scattering cross sections.

In order to get these expansions up and running though we need to invoke the so-called interaction picture. As is familiar from quantum mechanics there are multiple ways to represent the time evolution of a quantum system. In the Schrödinger picture the state evolves in time while the operators remain constant; in the Heisenberg picture the operators take on the time dependence and the states are constant. The idea behind the interaction picture is to use this freedom to isolate the free representations of the fields. If A_S and $\psi_S(t)$ are operators and states

be used to evaluate correlation functions directly without referring to the S-matrix at all.

in the Schrödinger picture the corresponding operators and states in the interaction picture are given by $A_I(t) = e^{iH_{0,s}t} A_S e^{-iH_{0,s}t}$ and $|\psi_I(t)\rangle = e^{iH_{0,s}t} |\psi_S(t)\rangle$. This means that the field operators are governed by H_0 while the remaining time dependence due to the interaction potential is shifted into the states.

The time evolution operator in the interaction picture, $|\psi_I(t)\rangle = U_I(t, t_0) |\psi_I(t_0)\rangle$, is then used to set up perturbation series for particular S-matrix elements. We can expand $U_I(t, t_0)$ in a power series in g of the form:

$$U_I(t, t_0) = \sum_{n=0}^{\infty} \frac{(-ig)^n}{n!} \int_{t_0}^{t} dt_1 ... \int_{t_o}^{t} dt_n T[V_I(t_1)...V_I(t_n)], \tag{3.3}$$

where T denotes the time order product, which rearranges operators in descending order with respect to their time arguments. This expansion can then be plugged into the definition of the S-matrix:

$$S_{\beta\alpha} = \lim_{t_i \to -\infty} \lim_{t_f \to \infty} \langle \beta | U_I(t_f, t_i) | \alpha \rangle. \tag{3.4}$$

Taking the in and out states to be eigenstates of H_0, on the grounds that the system is effectively isolated before and after the scattering event, we end up with a string of terms made up of Fock space operators acting on the free field vacuum state—expressions that we know how to compute, at least in principle.

What we typically need to do to work out the coefficients of the series at each order in g is evaluate a set of integrals over momentum space. Feynman diagrams bring some order to the proceedings, but the number and complexity of these integrals grows rapidly as the series proceeds and the best we can hope to do in practice is calculate the first few terms. As is now well known however, the integrals which result from naively following the prescription just described typically diverge. The next section describes the renormalisation procedure that is needed to deal with these divergences.

Before we move on to discuss the renormalisation procedure however there is another issue with the expansion technique as we have described it thus far: the apparent clash with Haag's theorem. What Haag's theorem shows is that, given certain assumptions common to all of the main axiomatic formulations of QFT, the interaction picture does not exist.[5] The root of the problem is that the continuum QFT model associated with H_0 and the continuum QFT model associated with the full interacting Hamiltonian, H, cannot be formulated on the same Hilbert space. As will be discussed in detail in chapter 6, systems

[5] The original result can be found in Haag (1956). It is Hall and Wightman's (1957) generalisation however which is nowadays commonly referred to as Haag's theorem. The assumptions needed to prove it are discussed in detail in Earman and Fraser (2006) and Miller (2015).

with infinite many degrees of freedom admit unitarily inequivalent Hilbert space representations, and the Haag result flows from the fact that the ground state of the free and interacting theories live in unitarily inequivalent Hilbert spaces. This means that there cannot be a global unitary transformation connecting the states and field operators of a free and interacting theory, which the interaction picture is clearly predicated on. Given the role that the interaction picture plays in setting up the power series expansion of the S-matrix, Haag's theorem has seemed to some to point to a fundamental inconsistency in the perturbative method—this is, of course, the root of the consistency problem flagged in §3.1.

3.2.2 Perturbative Renormalisation

To make things concrete let's consider applying the method just described to the ϕ^4 theory introduced in §1.2.2. The Lagrangian of this theory was:

$$\mathcal{L} = (\partial_\mu \phi)^2 - m^2\phi^2 - \lambda\phi^4. \tag{3.5}$$

In the absence of exact solutions the hope is that expanding the S-matrix of the theory in powers of λ will at least allow us to probe the region of the parameter space in which λ is small and the interaction is weak. When we follow the prescriptions set out in the previous section however troublesome terms like the following start to appear at second order in λ:[6]

$$\lambda^2 \int \frac{d^4k}{(2\pi)^4} \frac{1}{(k^2 - m^2 + i\epsilon)((p-k)^2 - m^2 + i\epsilon)} \tag{3.6}$$

This integral diverges, leaving the expansion ill defined and predictively useless. The same kind of divergent integrals arise when the perturbative method is applied to realistic gauge theories like QED and QCD.

This apparently disastrous behaviour can be seen as resulting from a poor choice of expansion parameter however. The renormalisation procedure used to obtain a sensible power series divides into three steps. The first is to 'regularise' the offending integrals, that is render them convergent. In the case of the integral above this can be done by taking the upper limit of integration to be some finite constant Λ, known as an ultraviolet perturbative cutoff. It can then be explicitly evaluated as a function of Λ which goes to infinity as $\Lambda \to \infty$; in our example there turns out to be an explicit dependence on $\ln \Lambda$. Terms of this kind are said to be ultraviolet divergent. The perturbative expansions of theories with mass-

[6] k here is an the internal 'loop momentum' which is being integrated over, p is an external momentum factor associated with the in and out scattering states and the standard $i\epsilon$ prescription has been used to deal with the poles on the real axis.

less particles, like QED photons, also contain infrared divergent integrals that blow up as momentum variables are integrated down to zero. In this case the integral can be regularised by taking the lower limit of integration to be some small but non-zero constant—an infrared perturbative cutoff. While imposing cutoffs on the momentum in this manner often plays a privileged conceptual role in discussions of renormalisation, a variety of other regularisation techniques are used in perturbative calculations in high energy physics. One method which is often more convenient in the treatment of gauge theories like QED and QCD, and has also sometimes been attributed special foundational significance in the philosophical literature,[7] is dimensional regularisation. In this approach the integration measure is modified so as to range over a fractional number of dimensions, $4-\epsilon$. For ultraviolet divergent integrals this leads to finite results for $\epsilon > 0$, with divergences manifesting as poles at $\epsilon = 0$, and infrared divergences can also be regularised in this way.[8]

Once the coefficients of the perturbation series can be manipulated as finite expressions the next step is to redefine the coupling in which we are expanding so as to remove the singular dependence on the cutoff—it is to this process that the term renormalisation properly refers. In the case of ϕ^4 theory we can do this by reparameterising the Lagrangian as follows:

$$\mathcal{L} = (1+\delta Z)(\partial_\mu \phi_r)^2 - (m_r + \delta m)^2 \phi_r^2 - (\lambda_r + \delta\lambda)\phi_r^4. \tag{3.7}$$

Where $\phi_r = (1+\delta Z)^{-1/2}\phi$ is the renormalised field, m_r and λ_r are the renormalised mass and charge, and $\delta Z, \delta m, \delta\lambda$ parameterise so-called counterterms. Note that nothing mathematically or physically dubious is going on here. The Lagrangian has simply been rewritten in terms of different variables. The value of the counterterm parameters make no difference to the dynamics described by the Lagrangian. It does make a difference to the terms of the perturbation series at each order however: both the renormalised expansion parameter and coefficients of the series depend on the choice of counterterms. It turns out to be possible to choose these factors in such a way that the part of the series coefficient at second order which diverges as $\Lambda \to \infty$ is completely removed. In fact, this can be done at each order in perturbation theory.

Theories whose ultraviolet divergences can be systematically eliminated at all or-

[7] Bain (2013b) makes a distinction between what he calls Wilsonian effective field theories, which employ an explicit cutoff, and continuum effective field theories, which employ dimensional regularisation. One upshot of this approach to perturbative QFT set out in this chapter is that claims Bain makes about the significance of this distinction end up being somewhat misleading. I discuss this point in more detail in §5.1.

[8] A number of technical issues with the dimensional regularisation approach, which are not germane to the thrust of my discussion, have been glossed over here. Collins (1984) provides a thorough treatment of this regularisation method.

ders of perturbation theory via a redefinition of a finite number of parameters in the Lagrangian in the manner just sketched are said to be renormalisable (showing that infrared divergences can also be removed from the perturbation series is another matter which will not be discussed in detail here). There turns out to be a simple dimensional criterion for determining whether a theory has this property: a theory is renormalisable if all of the coupling parameters associated with terms in the Lagrangian have zero or positive mass dimension.[9] The mass dimension of the ϕ^4 coupling, λ, for instance is zero (in four space-time dimensions) so this interaction term satisfies the requirement. The standard model Lagrangian also satisfies this criterion, and consequently the ultraviolet divergences in the perturbation series of QED and QCD can be completely removed at each order. In fact, the renormalisability of these theories is not a coincidence. Traditionally renormalisability was viewed as a necessary condition for a viable perturbative analysis of a theory and was often treated as a theory selection principle in the formulation of the standard model.[10] There is something puzzling about demanding renormalisability a priori in this way however. What licenses us to assume that the world is amenable to perturbative analysis? As we shall see in the following chapter, the renormalisation group leads to a new perspective on the renormalisability of empirically successful QFTs.

Even assuming renormalisability however, there might seem to be something odd about the redefinition prescription just described. How can a change of variables, which by construction leaves the underlying physics unchanged, transform the physical respectability and predictive power of the expansion? While there is a puzzle to be addressed here, that renormalising the coupling in this way can radically alter the properties of the series should not be surprising in itself. In general, the quality of a power series approximation can be very sensitive to the choice of expansion parameter. Suppose we are want to expand $\ln(1+x)$ as a power series, for instance. Expanding in x gives rise to a series which converges for $|x| < 1$, while making a change of variables and expanding in $x' = x/(x+2)$ produces a series which converges for any positive value of x, and typically converges much more rapidly. Moreover, the practice of renormalising the coupling parameter is often employed in the perturbative treatment of models in ordinary quantum mechanics which have no problems with divergences. In

[9]The term mass dimension, and dimensional analysis in natural units, was introduced in §1.2.2.

[10]In fact non-renormalisable theories can also be treated with the perturbative framework. What happens in this case is that new non-renormalisable interactions need to be added at each order in perturbative theory in order to absorb the ultraviolet divergences. This means that in principle an infinite number of parameters would need to be fixed by experiment in order to give meaning to the expansion at all orders. As long as we stick to a finite order of perturbative theory however, the renormalisation procedure can essentially be carried out as described here for non-renormalisable theories. For the purposes of this chapter I focus on the renormalisable case.

the case of the quantum anharmonic oscillator for instance, the most obvious choice of coupling leads to a divergent series, but by redefining the coupling, in essentially the same manner just sketched for interacting QFTs, a convergent series can be obtained.[11] The need to carefully define an appropriate coupling is not intrinsically problematic then. What is puzzling about this redefinition procedure just described however is that it seems to be entirely ad hoc. No reason has been given for thinking that removing the divergent part of the perturbative coefficients should lead to accurate predictions.

The final step of standard perturbative calculations is to take the cutoffs to infinity, or, more generally, remove the regulator. If the redefinition of the coupling has done its job this limit is well-defined and we obtain finite expressions from truncations of the series at each order. This is sometimes presented as returning us to the realms of continuum field theory after a detour through strange regulated theories with cutoffs on the momentum or in a fractional number of dimensions. As I will shortly be arguing, this way of speaking is quite misleading—removing the regulator on integrals in the perturbation series does not amount to constructing a continuum QFT model. It is worth bearing in mind however that there are also good practical motivations for taking the cutoff to infinity in perturbative calculations. The most convenient way of performing the rescaling process in practice is to stipulate the value of a small number of S-matrix elements (containing so-called 'primitive divergences') at a particular energy in the limit $\Lambda \to \infty$. It is this step which allows us to use multiple regularisation methods, often crucial to practical computations, and be guaranteed to get equivalent results. Furthermore, it is generally much easier to deal with integrals over all of momentum space than cutoff expressions. Again, however, the physical justification for removing the regulator remains unclear.

After all of these steps have been completed truncations of the resulting series can be in extremely good agreement with experimental data. As has already been stressed, particle physics phenomenology is largely based on computations of this kind. There is a final, potentially unnerving, feature of renormalised QFT perturbation series however which is worth mentioning here: they are thought not to converge. That is, even after each term has been rendered finite by the renormalisation procedure the sum of the series at all orders is not. Dyson (1952) was the first to give heuristic arguments to the effect that QED perturbation series have zero radius of convergence and this conclusion has since been borne out by studies of the large order behaviour of the perturbation series of realistic QFTs. In the case of the interacting QFTs which have been constructed as models of the Wightman axioms on lower dimensional space-times the divergences of the

[11] Delamotte (2004) and Neumaier (2011) give useful discussions of renormalisation in this more general context. The $\ln(1+x)$ example is borrowed from Neumaier (2011).

perturbation series can, in fact, be rigorously demonstrated.[12] Of course, the apparent convergence of QED perturbation theory in the first few orders, and their excellent agreement with experiment, assure us that the series is, at least, an asymptotic expansion. But the fact that it divergences prevents us from directly identifying its sum with the S-matrix elements of the theory. This perhaps is one of the features of QFT perturbative expansions which contribute to their reputation for a lack of mathematical rigour.

In sum then, our survey of the perturbative QFT has identified prima facie motivations for the rigour, consistency and justification problems identified in §3.1.[13] In order to assess the substance of these worries however a more careful analysis of what is going on in the perturbative method is needed. The next section lays the groundwork for such an analysis by introducing some notions drawn from philosophical work on scientific modelling which illuminates the perturbative method, namely the distinction between an approximation and an idealised model.

3.3 Approximation and Idealisation Distinguished

The best way to present the distinction between approximations and idealisation I want to make use of here is with reference to examples. Take a familiar classical system: a unit mass falling in a uniform gravitational field while acted upon by a linear air resistive force.[14] The equation of motion of this system can be written in terms of the velocity:

$$dv/dt = g - kv, \tag{3.8}$$

[12] See Strocchi (2013, 35-39) for results pertaining to perturbation series of ϕ^4 theory on two and three dimensional space-times.

[13] There are also other puzzles raised by perturbative formalism which have not been discussed here. One issue is the apparent ambiguity of perturbative calculations owing to the dependence of truncations of the series on the so-called renormalisation scheme. As I explained, it is the fact that the terms in the perturbation series depend on the counterterms invoked in the renormalisation process which allows us to remove the divergent dependence on the regulator. By the same token however, we can also add finite contributions to the perturbative coefficients. This means that there are multiple ways of removing the divergences in the coefficients, known as renormalisation schemes, which give different results when the series is truncated. What happens in practice is that a small number of conventional schemes are used, but the justification for choosing these schemes over possible alternatives is a matter of contention in the physics literature—see James Fraser (2012) for a review of this debate. How we can make sense of the success of the perturbative method in light of this apparent ambiguity is another issue which philosophers of science have yet to tackle.

[14] This example is taken from Norton (2012). A clarificatory note here: Norton uses the term idealisation in a different way from my preferred definition set out in the previous chapter. For Norton an idealisation is a system which is used to represent a distinct target system—i.e. he uses the term in the way that many authors, myself included, use model. I use the term idealisation in a more fine grained way to refer to features of a model which misrepresents the target system. This is simply a matter of terminology and does not make a difference to the substance of the distinction I am after here.

where g is the gravitational acceleration and k parameterises the frictional force. Assuming the body is initially at rest the solution of this equation is:

$$v(t) = g/k(1 - e^{kt}). \tag{3.9}$$

Suppose however that air resistance is not very important for the aspects of the system's behaviour we are interested in describing. Perhaps k is small relative to g and we are only interested in the system's evolution over comparatively short time periods.

There are two paths one might take in this situation. On the one hand, we could consider a new classical system, with a simplified equation of motion $dv/dt = g$, and treat it as a representation of our original system. This model clearly misrepresents its target in some respects; in the terminology of the previous chapter it contains idealisations. It falsely represents k as being zero and has quite different asymptotic behaviour as $t \to \infty$, since it lacks a terminal velocity. But it does get some of the features of the target system right—the velocity function of the idealised model will stay within some error bound of the target system's velocity over some finite time period. There is another way to proceed however. Suppose that instead of moving to a new simplified system we take the limit $k \to 0$ of the velocity function of the original system:

$$\lim_{k \to 0} v(t) = gt \tag{3.10}$$

This produces a new function which is, again, within some error bound of the target system's velocity over sufficiently short periods of time. It can therefore be viewed as an inexact description of this property of the system. Following Norton (2012), I will call the production of such a function, without referring to a new physical system, an approximation.

This distinction between representing a system with an idealised model and providing an inexact description of one of its properties directly has long been recognised in the literature on scientific modelling. Its significance is less agreed upon however. Redhead (1980) claims that the two approaches are always interchangeable, in the sense that, for any function which approximates a quantity of a target system we can construct a system to which it is an exact solution. This certainly does seem to be the case in the above example, since the approximation obtained by taking $k \to 0$ in the velocity function is identical to the velocity of the corresponding idealised model. Following examples like these, philosophers of science have tended to see models as the fundamental unit of analysis and approximations as ultimately derived from them, a doctrine which probably also owes something to the dominance of the semantic view of theories.

Norton (2012) has recently put forward a different view of the distinction which I want to endorse here, namely that approximation and idealised representation are not, in general, interchangeable, and that distinguishing between them can play an important role in resolving some puzzles in the philosophy of physics. Norton presents a number of counterexamples to the idea that every approximation piggy backs on a corresponding idealised model. Take, for instance, an ellipsoid with semi-major axis of length a. The ratio of the surface area and volume of this geometric object can be expressed as a function of a. If we take the limit $a \to \infty$ of this function it converges to $3\pi/4$; this might be an appropriate approximation if a is very large. If we take the semi-major axis of the ellipsoid itself to infinity however we get an infinite cylinder whose surface area to volume ratio is 2 (see figure 3.1). In this case applying an operation to a function defined on a model and applying the analogous operation to the model itself gives different results and we clearly cannot say that the former is underwritten by the latter.

A second class of examples which speak against Redhead's equivalence claim are cases in which no model exists corresponding to an operation used to implement an approximation. Norton discusses the example of a sphere of radius r. The ratio of surface area to volume of a sphere is $3/r$, which clearly has a finite limit as $r \to \infty$. But, at least on standard definitions of a sphere in Euclidean geometry, there is no such thing as a sphere with infinite radius. While it may be appropriate to treat the area to volume ratio of a large sphere as vanishing, there is no geometric object which actually has this property. A case of this kind closer to our area of interest is the use of the thermodynamic limit in statistical mechanics. A common modelling strategy in statistical mechanics is to take the limit in which the volume and number of degrees of freedom, N, go to infinity (this is closely related to the infinite volume limit of a cutoff QFT model discussed in §1.2.2). Once again, there is some ambiguity about how this practice should be understood. One could take this limit of a particular thermodynamic function, or actually construct an infinitely extended system by taking the boundary of a finite system to spatial infinity in a controlled way. The latter task turns out to be difficult to do in practice: the limit system typically has to be established on

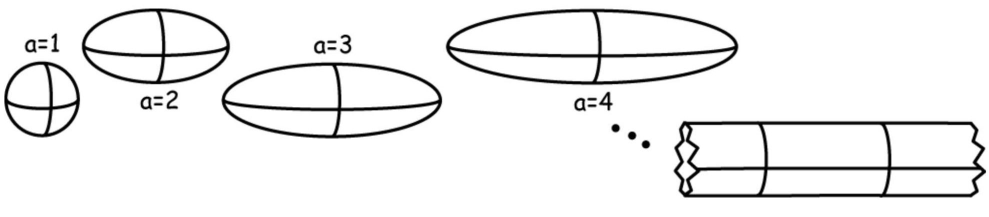

Figure 3.1: Taking the semi-major axis of an ellipsoid to infinity. Taken from Norton (2012).

a case by case basis and is known not to exist for gravitating systems.[15] Again, it might be that the $N \to \infty$ limit is well-defined for some thermodynamic quantities though there is no infinite volume limit system.

Apart from counterexamples of this kind, another motivation for taking approximations and models to be distinct kinds of theoretical output is that it often leads to a more natural interpretation of scientific practice. As the previous examples attest to, it is typically more difficult to establish the existence of a physical system than a single function. As Norton points out, we often find that physicists do not do this extra work; he argues, for instance, that in many applications the $N \to \infty$ limit is taken of particular thermodynamic functions without addressing the question of whether an infinite system can be generated in this way. Those who hold that approximations are underwritten by the presence of a corresponding model will be forced to reproach physicists for their sloppiness in such cases. But if we give up this idea, as Norton's examples suggest we should anyway, a more sympathetic reading becomes available. On this view, producing a function which approximates some property of the target system, and producing a model which resembles it in certain respects are simply different things—the former does not rest on the latter. Consequently, there is nothing wrong with using approximate expressions for physical quantities without embedding them within idealised models.

It is this capacity of a robust notion of approximation to illuminate scientific practice which I want to draw on here. In the context of the QFT programme, the existence of QFT models is a delicate issue and consequently the construction of approximations and models naturally come apart. This, I think, is the key to understanding what is going on in perturbative QFT.

3.4 Perturbative Theory Produces Approximations, Not Models

The key claim of this chapter is that perturbative QFT ought to be understood as a method for producing approximations in the sense just elaborated, rather than picking out QFT models. The truncated power series obtained by following the prescriptions set out in §3.2 can approximate physical quantities like scattering cross sections, but the various steps involved in getting to these expressions should not be interpreted as an attempt to provide a structural characterisation of an interacting QFT. The central motivation for this view is that it makes sense

[15] Ruelle (1969) is a classic reference for rigorous results about the existence of the thermodynamic limit of particular systems. Callender (2011) gives a discussion of the non-existence of the thermodynamic limit in the case of gravitating systems.

of what physicists actually do, or perhaps more importantly do not do, when they treat QFTs perturbatively. As was discussed in chapter 1, constructing interacting QFT models in four dimensions is a project fraught with technical difficulties. Looking for a solution to this problem in the standard perturbative approach leaves us disappointed. At each juncture in the perturbative method we find that the work needed to construct a physical model is missing.

Consider, for instance, the regularisation of the momentum space integrals making up the perturbative coefficients. Imposing a high momentum cutoff on such an integral is sometimes described, in both the physics and philosophy literature, as taking us to a cutoff theory. In light of the distinction drawn in the previous section however, we can see that putting a cutoff on a single integral and constructing a quantum system which lacks high momentum degrees of freedom is not the same thing. In fact, early applications of renormalised perturbation theory made no attempt to verify the existence of such systems, or explore their properties. The details of how to formulate QFTs on a lattice have since been worked out, and other ways of implementing a 'non-perturbative' cutoff have been developed in the constructive field theory programme (as I described in §1.2.2). But this came decades after the original perturbative treatment of QED. The lack of any attempt to construct a quantum system corresponding to a regularised integral is even more evident in the case of dimensional regularisation. Recall that this method works by continuing the integration measure to a non-integer number of dimensions. Again, one sometimes finds talk of quantum theories with a fractional number of dimensions in this context. Yet (to my knowledge) no attempt has been made to actually construct QFT models of this kind, and dimensional regularisation techniques do not play any role in extant work in the constructive field theory programme.[16]

Another place where the perturbative approach is apt to disappoint those looking for a class of mathematical structures to identify with realistic QFTs is the removal of the regulator. I have already suggested that it is misleading to talk about taking a continuum limit here, and we can now see why. All that is happening when an ultraviolet cutoff is taken to infinity in a typical perturbative calculation is that a limit is being taken of a particular truncated power series expression. That a great deal more is needed to verify the existence of a continuum QFT model is evident when we look at the work that goes into constructing models of the Wightman or Haag-Kastler axioms. As I described in §1.2.2, what mathematical physicists working on this problem typically do is start with a cutoff QFT model, with non-perturbative cutoffs at high and low momentum, and attempt to take the continuum and infinite volume limits of this structure. The difficul-

[16]Rivasseau (2014, 7) comments that dimensional regularisation "cannot be used up to now in a constructive non-perturbative program".

ties associated with taking these limits in a controlled way outstrip, and are in fact largely independent of, the problem of ridding perturbative expansions of divergences. As will be discussed further in the following chapter, the continuum limit is best understood within the context of the renormalisation group framework. For now it is enough to say that the existence of the continuum limit is commonly taken to depend on the existence of a so-called ultraviolet fixed point, and this property is independent of the property of perturbative renormalisability. Moreover, the treatment of infrared divergences in the perturbation series of gauge theories like QCD does not solve the problems associated with bringing its infrared behaviour under mathematical control in the non-perturbative context. In sum then, removing the cutoffs on perturbative approximants has very little to do with the project of constructing continuum QFT models.

A final blow for those searching for explicit constructions of QFT models in perturbative framework is the divergence of the expansion. If renormalised perturbation series converged we could use their sums to define S-matrix elements and correlations functions and construct QFT models from there. Alas, all indications are that they do not. This may not rule out the possibility of extracting a structural characterisation of interacting QFTs from renormalised perturbation series entirely—there are, after all, a number of ways of defining sums for divergent series.[17] It does, at least, show that it is a non-trivial undertaking. And again, is not a project which we find physicists working with perturbative methods engaging in. In most practical calculations particle physics phenomenologists simply truncate the perturbation series after the first few terms and do not concern themselves with the fate of the sum to all orders. Once again, no attention is being given to the existence of QFT models.

This all makes sense if we take the intended output of the perturbative approach to be approximations rather than physical systems. On this reading the strategy underlying the perturbative treatment of interacting QFTs is to dodge rather than solve the problem of how to characterise the theory in terms of mathematical structures. While philosophers of a realist persuasion in particular will find this troubling—a point I will come back to shortly—there is nothing manifestly incoherent about it if we accept the conclusions of the previous section. As I argued there, it is a mistake to think that approximations must be embedded within a model in order to be meaningful. We should not berate particle physics

[17] It turns out that the perturbation series of interacting theories which have been constructed as models of the Wightman axioms in a reduced number of space-time dimensions have a unique Borel sum which is identical to the exact values of the corresponding S-matrix elements—see Magnen and Seneor (1977) and Riviseau (2014). There are good reasons to think that this behaviour does not generalise to the case of realistic QFTs however, for a discussion of this problem see Duncan (2012) chapter 11. Thanks to Michael Miller for discussions about this issue.

phenomenologists then for failing to tell us what an interacting QFT is. Obtaining approximate expressions for scattering cross sections and constructing QFT models are different objectives; in principle the former can be pursued without addressing the latter.

A great advantage of this reading is that it opens the way for more constructive philosophical engagement with perturbative QFT. The analysis I have developed here sheds new light on the three putative problems with the perturbative approach I distinguished in §3.1 and ultimately leads to a less pessimistic assessment of the conceptual respectability of the perturbative framework.

One of the complaints about perturbative QFT, which I called the rigour problem, was that its lack of mathematical rigour makes it impossible to engage with from a foundational perspective. This worry is not usually stated in very explicit terms—we rarely find philosophers pointing to specific aspects of the perturbative method which they find problematically unrigorous. But I suspect that what is often driving this sort of objection is the absence of an explicit specification of a class of mathematical structures underlying the approach. As I discussed in the previous section, philosophers of science have often viewed approximations as derived from, and consequently underwritten by, physical models. Given what has been said already about perturbative QFT it is easy to see how a zealous follower of this doctrine might perceive it as hopelessly sloppy, and attribute this to an unrigorous treatment of the relevant mathematics. As I mentioned in §1.2.3, in addition to cutoff and continuum models, Fraser identifies an "infinitely renormalised" variant of QFT which is associated with the process of imposing and removing a cutoff in the perturbative renormalisation procedure. According to her this amounts to adding infinite counterterms to the Lagrangian and leads to a mathematically ill defined system.

Rather than reading the perturbative approach as a botched attempt at constructing QFT models however, I have been arguing that it is much more natural to interpret it as never taking up this project in the first place. Understood as a method for producing approximations, the perturbative approach is more difficult to dismiss as mathematically unsound. This is not to say that standard perturbative computations are paragons of mathematical precision. The treatment of the convergence of momentum space integrals in the physics literature, for instance, often falls short of the standards of rigour upheld by mathematicians. But this is the kind of everyday imprecision that is ubiquitous in applied mathematics—if we reject the perturbative approach on these grounds we will also be jettisoning much of physical science. The perturbative approach has been seen as mathematically problematic in a more radical sense, I think, because of the lack of attention payed to the existence of QFT models. On the view I have been developing here however this has nothing to do with mathematical rigour and instead reflects the more modest objective of the perturbative methodology. Whether "infinitely renormalised" QFTs make sense as constructive mathematical objects or not they are not deployed in any way in conventional perturbative calculations.

The apparent inconsistency problem posed by Haag theorem also turns out to be less threatening than it initially appeared. The crucial point here is that Haag's result pertains to the properties of QFT models. Roughly speaking, it tells us that the time evolution of models of the Wightman and Haag-Kastler axioms cannot be carved up in the manner prescribed by the interaction picture. Why doesn't this undermine perturbative evaluations of scattering cross sections? The short answer is that since the perturbative method is not in the business of providing a structural characterisation of QFT there cannot possibly be a conflict here. Perturbative evaluations of S-matrix elements do not posit the existence of models satisfying the assumptions shown to be inconsistent by Haag's theorem because they do not pick out any physical models at all.

In fact, the perspective on perturbation theory advocated here fits well with a recent response to Haag's theorem due to Miller (2015), which helps to clarify how perturbative calculations avoid inconsistency. As Miller points out, the methods used to regularise ultraviolet and infrared divergences in the perturbative expansion invariably cut against the assumptions needed to prove Haag's theorem. Consider, in particular, the imposition of high and low momentum cutoffs on the relevant integrals. As I argued above this is not the same thing as constructing a cutoff model, but we now know how to write down such systems. If we put a QFT on a finite volume lattice, for instance, the resulting model is not touched by Haag's theorem because it violates the assumption needed to prove it. In fact, since the number of degrees of freedom is finite, such as system does not admit unitarily inequivalent Hilbert space representations at all. This means that the interaction picture exists and the steps involved with setting up the perturbative expansion of the S-matrix can be concretely implemented on such a structure. Crucially, when we remove the cutoffs at the end of the calculation this should be understood as taking a limit of a particular function. Just as taking the radius to infinity of particular quantities defined on a sphere does not require one to posit the existence of an infinite sphere, in Norton's example, removing the cutoffs does not rest on the assumption that the interaction picture can be implemented in the continuum limit. The perturbative method simply does not assert an inconsistent set of claims about QFT models.[18]

[18]This is not to say that Haag's theorem has no foundational significance at all, merely that it does not directly undermine standard perturbative QFT.

The rigour and inconsistency problem turn out to be red herrings on the reading of the perturbative approach I have developed then. Worries about the internal coherence of the QFT perturbation theory have largely sprung from a misunderstanding of its aims in my view. Still, there clearly is something puzzling about the perturbative approach, as I have described it. It is the justification problem, I think, which ends up being the really fundamental conceptual issue with perturbative QFT.

3.5 The Real Problem

The method described in §3.2 may provide a sound method for generating well defined functions but there remains a mystery about why the results provide such good approximations of their experimentally measured values. The renormalisation procedure, in particular, seems to be entirely ad hoc. The process of regularising integrals, redefining the coupling and removing the regulator in the standard renormalisation procedure all seem to proceed in the absence of any physical argument underlying each step.

While my analysis of the perturbative approach does not solve this problem it does help to provide a diagnosis of the root cause. The success of the perturbative approach is mysterious, I suggest, precisely because it dodges the question of what an interacting QFT is. Many of the seemingly peculiar features of QFT perturbation series are actually found in many applications of perturbative methods elsewhere in physics. As I mentioned in §3.2.2, the process of redefining the expansion parameter in the renormalisation procedure is often employed in perturbative calculations in classical and quantum theory even when there is no problem with divergences. Furthermore, divergent perturbative expansions are commonplace in many branches of physics. In most of these cases however, we typically know what the physical systems whose properties we are trying to approximate looks like. Although we cannot solve the standard quantum mechanical models of the Helium atom, for instance, we know how to write down its Hamiltonian and Hilbert space. What sets the original perturbative treatment of QED apart from these cases is the absence of any non-perturbative characterisation of the system of interest. While I have argued that this does not render perturbative QFT incoherent it undercuts the possibility of telling a physical story which could explain its success.

This appears to be bad news for the scientific realist, who wants to say that scientific predictions succeed because they are derived from theories which accurately represent the way the world is. But this situation also seems to be problematic for the best developed forms of anti-realism. Semantically interpreted models are still the basic unit of analysis for the constructive empiricist; they define their epistemic commitments and notion of empirical adequacy with respect to these structures. If I am right then the perturbative approach to QFT cannot be read in these terms either: it does not provide us with physical models at all, empirically adequate or not. If an anti-realist reading of the perturbative approach as it stands can be maintained then, it must be a fairly radical instrumentalist position, according to which the perturbative method is, as Buchholz suggests, merely "an efficient algorithm for the theoretical treatment of certain specific problems in high-energy physics" (Buchholz, 2000, 1-2). Whether we can make sense of this stance is a question for another time, but familiar worries about instrumentalism are likely to rear their head again in this context.[19]

Besides these connections with broader debates in the philosophy of science however, the ad hoc character of perturbative renormalisation can also be viewed as an local problem for the QFT programme. Indeed, as I gestured at in §1.2.1, it was perceived as such by physicists in the wake of the success of the initial success of the perturbative approach. The emergence of the axiomatic and renormalisation group approaches to QFT in the second half of the 20th century can both be understood as responding to inadequacies with the perturbative formalism, the search for a physical justification for the renormalisation procedure being at least one of motivation for these developments. We should not move too quickly to dismiss the perturbative approach to QFT as conceptually unsound on the basis of the justification problem then. The analysis put forward here suggests that the perturbative formalism lacks the resources to answer this challenge on its own, but in the modern context we have non-perturbative approaches to QFT on the table which may yet provide a solution to this puzzle. The following chapter takes up this possibility by examining the impact of renormalisation group methods on this issue.

[19] My analysis of perturbative QFT can thus be understood as posing a challenge to the idea that theories and models are the key objects for understanding the structure of science. Incidentally, Kaiser (2009) endorses a similar conclusion in his study of the historical development of Feynman diagram methods in high energy physics. There is scope for a more radical departure from the traditional picture of the philosophy of science here then. In keeping with the conservative stance set out in the introduction however, I will be exploring the possibility that we can, in fact, make sense of what is going on here within a fairly conventional realist framework.

Chapter 4

Lessons From the Renormalisation Group: Structures and Representations

4.1 The Contested Status of the Renormalisation Group

The application of renormalisation group methods to QFT was undoubtedly a major breakthrough in 20th century physics; as I described in chapter 1, the renormalisation group approach paved the way for the development of QCD and ultimately the standard model. Its philosophical significance is more controversial however. The status of the renormalisation group is a central point of contention in Fraser and Wallace's dispute over the theoretical identity of QFT. Wallace takes renormalisation group methods to have important ramifications for the philosophy of QFT and appeals to them frequently in his arguments in support of cutoff formulations of the theory. By contrast, Fraser advocates a more instrumentalist reading:

> [Renormalisation group (RG)] methods illuminate the empirical structure of QFT. That is, they furnish a more perspicuous procedure for deriving empirical predictions... However, RG methods do not shed light on the theoretical content of QFT. For this reason, appeal to RG methods does not decide the question of which set of theoretical principles are appropriate for QFT. (Fraser, 2011, 131)

For Fraser it seems the advent of the renormalisation group approach to QFT represents progress of a pragmatic kind but does not amount to a substantial improvement in our understanding of the foundations of QFT.

This chapter aims to clarify the relevance of the renormalisation group to the philosophy of QFT. I argue that the renormalisation group approach does in fact impact on a number of philosophical issues raised by the QFT programme.

After describing the basics of the renormalisation group framework (§4.2), with particular emphasis on the non-perturbative momentum space approach, I argue that this formalism impacts on the project of constructing continuum QFTs, an issue which Fraser herself clearly takes to be of considerable foundational importance (§4.3). My real focus here, however, is on how the renormalisation group approach informs judgements about the kind of representational success which can reasonably be attributed to QFT models. Picking up where we left off in the previous chapter, I claim that the renormalisation group furnishes a non-perturbative explanation of the success of perturbative QFT (§4.4). A key element of this story is that the renormalisation group approach reveals the insensitivity of physical quantities at relatively low energies to the details of the dynamics at high energies, and in particular to the value of an ultraviolet cutoff. This leads me to argue that cutoff QFT models ought to be viewed as successful representations of coarse-grained features of the actual world, a claim which I spend most of the rest of the thesis clarifying and defending (§4.5). I conclude by returning to the question of how the foundational status of the renormalisation group ought to be understood (§4.6). The best way to view the renormalisation group, I suggest, is as a powerful tool for obtaining information about the properties of QFT models—contra Fraser however, I do not think its reach is limited to their empirical properties.[1]

4.2 The Renormalisation Group in Quantum Field Theory

The term 'renormalisation group' might give the impression that we are dealing with a single, monolithic, mathematical formalism. In reality a plethora of distinct methods, which differ in potentially conceptually significant respects, fall under this rubric.[2] Disentangling the various strands of the renormalisation group concept, and their application across different branches of physics, is not my project here (though this is certainly an issue which calls for further investigation by historians and philosophers of science). This section focuses on the incarnation of the renormalisation group which I take to be most important for

[1] My discussion in this chapter owes a great deal to previous work on the renormalisation group approach to QFT in the philosophical literature, and especially to Wallace (2006, 2011) and Hancox-Li (2015a, 2015b).

[2] There are at least three distinct renormalisation group traditions, which historically arose more or less independently. First, there was the early work of Stueckelberg and Petermann (1951), Gell-Mann and Low (1954) and others which introduced some of the key ideas of the renormalisation group in the context of perturbative QFT. Second, there is the real space or blocking approach originating in the work of Kadanoff (1966) in the context of condensed matter physics (touched on briefly below). Third there is the momentum space approach developed by Kenneth Wilson, which I focus on in this chapter. Even within these traditions however, there is a great deal of variation in the methods which are employed.

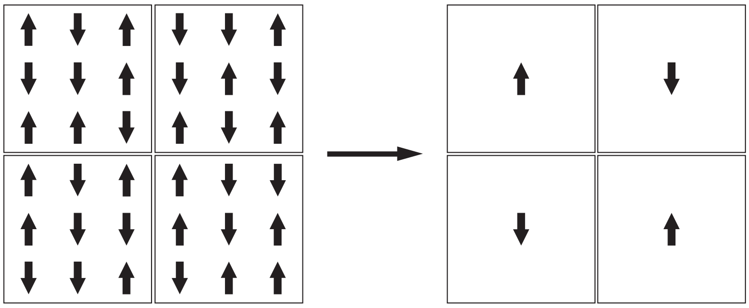

Figure 4.1: A ‘majority rule’ blocking procedure on a two valued spin system.

the philosophy of QFT; namely the non-perturbative momentum space approach developed by Kenneth Wilson and others in the 1970s.[3]

The core idea underlying the modern conception of the renormalisation group, in all of its guises, is the study of coarse-graining transformations. That is, operations which take us from an initial system of interest to a new one which lacks some of the degrees of freedom associated with variations at small length scales but agrees with its large scale properties. One reason that renormalisation 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 4.1). In some cases it may be possible to invert the transformation, forming a group structure—hence the name. But this 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 renormalisation 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’ in a space of possible models. Studying this flow gives us information about the way that systems behave at different length/energy scales.

Rather than working in real space, as in the blocking approach, Wilson pioneered the idea of setting up a coarse-graining transformation in momentum space. In the QFT context this is most naturally carried out within the path integral approach to the theory. The key quantity here is the partition function, Z, introduced in §1.2.2. For a single scalar field, this quantity is associated with the path integral expression:

$$Z = \int \mathcal{D}\phi e^{-\int d^4x\mathcal{L}}. \tag{4.1}$$

[3]Wilson and Kogut (1974) is a classic review of this approach. Binney et al (1992) provides a useful account of the Wilsonian renormalisation group in the context of condensed matter physics. Duncan (2012) describes its application to QFT.

As we saw, there are grave difficulties associated with precisely defining this integral on continuum space-times—we will see in the next section that the Wilsonian renormalisation group turns out to be relevant to this project. It is possible to give concrete meaning to the path integral however if cutoffs are introduced. The Wilsonian renormalisation group is based on setting up a coarse-graining transformation on these cutoff QFT models.

The key idea is that, rather than evaluating the whole path integral at once, we can isolate the contribution due to high momentum field configurations, whose Fourier transforms have support above some value μ. This part of the path integral can then be computed separately and absorbed into a shift in the initial Lagrangian. In symbols:

$$\int_{|p|\leq\mu} \mathcal{D}\phi \int_{\mu\leq|p|\leq\Lambda} \mathcal{D}\phi e^{-\int d^4x\mathcal{L}} = \int_{|p|\leq\mu} \mathcal{D}\phi e^{-\int d^4x(\mathcal{L}+\delta\mathcal{L})} \tag{4.2}$$

This defines a transformation which takes us from an initial cutoff QFT model to a new one, which has a lower ultraviolet cutoff and a modified dynamics but behaves like the original.[4]

This transformation will not only alter the values of coupling parameters in the initial Lagrangian 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-renormalisable interactions (which, you will recall, are interaction terms with couplings that have zero or positive mass dimension). For scalar fields this means going beyond the ϕ^4 theory to a broader class of Lagrangians of the form:[5]

$$\mathcal{L} = (\partial_\mu\phi)^2 - m^2\phi^2 - \sum_{n=2}^{\infty} \lambda_{2n}\phi^{2n}. \tag{4.3}$$

The renormalisation group transformation can then be seen as inducing a flow in a space of possible Lagrangians, spanned by the couplings (m, $\lambda_4, \lambda_6, ...$). The way that these parameters change as μ is lowered, and more and more high momentum degrees of freedom are 'integrated out', is conveniently described by the so-called beta functions:

$$\beta_i = \mu\frac{\partial}{\partial\mu}\lambda_i. \tag{4.4}$$

[4]There are some subtleties here when it comes to precisely defining this kind of coarse-graining transformation. In particular, there is a difficulty with controlling the contributions of configurations in which the magnitude of the field diverges in a space-time region—known as the large field problem (Rivasseau, 2014). See Hancox-Li (2015a) for a discussion of the work that has been done to make the Wilsonian renormalisation group more rigorous in the mathematical physics literature.

[5]Assuming a $\phi \to -\phi$ symmetry to ensure that the energy is bounded from below.

The most important points in this space for understanding the behaviour of the renormalisation group flow are those at which all of the beta functions vanish, known as fixed points. We will have more to say about their significance shortly.

One important point to note about the presentation of the renormalisation group approach so far is that perturbation theory has not been mentioned at any point. In practice, the path integral over large momentum configurations cannot be evaluated exactly for physically interesting models and some kind of approximation method is needed to make computations tractable. Perturbative methods are often employed in this context, but other kinds of approximation are also invoked.[6] As I argued in the previous chapter however, the problem with perturbative QFT is not the use of perturbative approximations itself but the lack of a non-perturbative characterisation of the underlying physics which explains their success. I will argue in §4.4 that the renormalisation group approach fills this lacuna, and provides a way of justifying perturbative renormalisation.

Another significant property of this kind of momentum space renormalisation group transformation is that it can be inverted—so we really are talking about a group here. This means that we can explore what happens when we take the renormalisation group flow in the other direction, as it were, and increase the value of μ beyond the initial value of the cutoff. As I describe in the next section, it is this feature of the Wilsonian approach which makes it relevant to the project of constructing continuum QFT models.

4.3 Ultraviolet Fixed Points and the Continuum Limit

As we have discussed already, what mathematical physicists working in the constructive field theory programme invariably do when they attempt to establish the existence of interacting QFTs on continuum space-times is start with a cutoff QFT model and take the continuum and infinite volume limits of this structure. The renormalisation group provides a natural framework for attacking the former, ultraviolet, part of this problem, as I will now explain. Hancox-Li (2015b) has recently argued that this speaks against Fraser's view, quoted in §4.1, that the renormalisation group merely facilitates the extraction of empirical predictions within the QFT programme. I reiterate this case here, paying particular attention to Fraser's discussion of this issue.

For simplicity's sake, let's consider a single interaction term in the Lagrangian

[6] This is another reason why renormalisation group methods are so diverse; even when we have settled on a coarse-graining transformation to study a range of different approximations can be employed to gain information about the resulting renormalisation group flow.

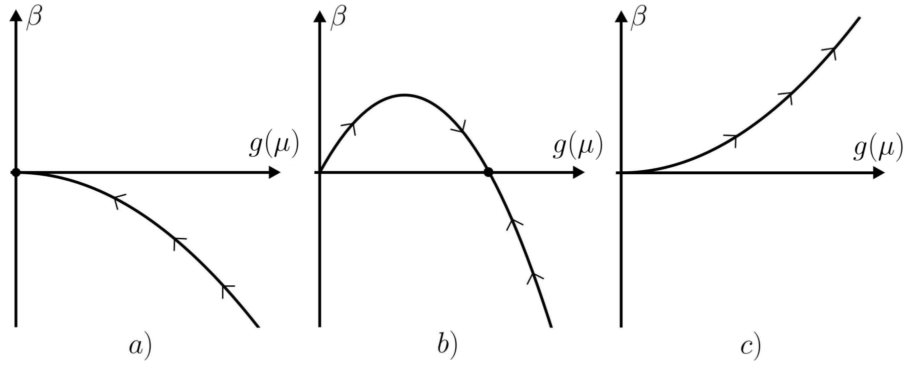

Figure 4.2: Sketch of the beta function in the case of a) asymptotic freedom, b) asymptotic safety and c) no fixed point.

of a cutoff model, parameterised by a coupling g. As we saw in the last section, the momentum space renormalisation group transformation can be inverted and μ can be taken beyond the initial value of the cutoff. Essentially what we are doing here is adding more degrees of freedom associated with variations at smaller length scales and modifying the dynamics to compensate. Suppose we start with g at some non-zero value and take $\mu \to \infty$: what happens to g? One possibility is that the corresponding beta function goes to zero and the renormalisation group terminates at a so-called stable ultraviolet fixed point (henceforth simply ultraviolet fixed point). An important special case is asymptotic freedom in which $g \to 0$ as $\mu \to \infty$. All of the evidence from perturbative calculations suggests that this is what happens in the case of QCD. The coupling can also converge to non-zero value, a scenario sometimes known as asymptotic safety. The possibility that QFTs obtained from a naive quantisation of the Einstein equations have this kind of ultraviolet behaviour is currently a subject of intense research in the quantum gravity programme.[7] The other possibility is that there is no fixed point and the coupling does not converge to a finite value as $\mu \to \infty$. All the indications are that this is how the interaction couplings of QED and the ϕ^4 theory in four space-time dimensions behave. (See figure 4.2 for an illustration of these scenarios.)

The relevance of the renormalisation group to the constructive field theory project should now be clear. We need the Lagrangian to converge as the cutoff goes to infinity if we are to obtain an interacting theory in this way—we need an ultraviolet fixed point. Theories which have no ultraviolet fixed point are often said to be trivial because the only way to obtain a well defined continuum system in this case is to take $g \to 0$, i.e. completely 'turn off' interactions, as the cutoff is removed. Whether realistic theories like QED and QCD admit continuum

[7]Niedermaier and Reuter (2006) is a review of this programme. Note that this substantiates my claim from §3.3 that perturbative renormalisability is not a necessary condition for the existence of a well-defined continuum QFT. In fact, rigorous results about the continuum limit of the non-renormalisable Gross-Neveu model in three dimensions have been obtained (Rivasseau, 2014).

formulations is clearly a question which many philosophers of physics, and not least Doreen Fraser herself, take to be of paramount importance. I will have more to say about my take on the philosophical significance of this issue later on in this chapter. For now though, the key point is that we have identified one way in which the renormalisation group apparatus impacts on an issue of foundational import.

This clearly cuts against the view that these methods only illuminates the empirical content of QFT. Perhaps unsurprisingly then, Fraser takes a deflationary view of the relationship between the renormalisation group treatment of the ultraviolet behaviour of a theory and the existence of continuum QFT models. She writes:

> ...[T]he fact that QED is not asymptotically free does not entail that there are no non-trivial models because the formulation of QED using the RG formalism relies on particular approximation techniques; different approximation techniques could be compatible with the existence of non-trivial rigorous models.[8] (Fraser, 2011, 131)

While these comments point to some important provisos about what can be legitimately inferred from the renormalisation group analysis of realistic QFTs as it currently stands, they do not undermine the relevance of the general approach for the project of constructing continuum QFT models, as I will now argue.

What does Fraser mean exactly by "approximation techniques" in this passage? If we take approximation to mean an inexact description of the properties of a target system, as in the previous chapter, we can certainly find many approximations being employed in the renormalisation group approach to high energy physics. As has already been mentioned, the Wilsonian renormalisation group transformation rests on path integral expressions which cannot be evaluated analytically for physically interesting models. Approximations of one kind or another need to be employed, which give us only partial and imperfect information about the renormalisation group flow. Perturbative approximations of the beta function, in particular, are intrinsically limited by the fact that they can only be expected to be accurate in the region in which the coupling is small. It is possible that perturbative calculations indicate that a model is trivial in the continuum limit when in fact an ultraviolet fixed point exists at large g in a region of the space which cannot be accessed via perturbative theory. This is one sense in which the use of approximations within the renormalisation group approach could give

[8]Two quick comments on this passage: Fraser refers to asymptotic freedom here but we can safely substitute the more general notion of having an ultraviolet fixed point sketched above. Furthermore, by a 'rigorous' model I think Fraser really means a continuum model—as I said in §1.2.2, the question of mathematical rigour is not the most salient issue in the distinction between cutoff and continuum models.

rise to misleading impressions about the non-existence of the continuum models. In the case of QED and ϕ^4 theory a number of different perturbative and non-perturbative approximation methods all point to the non-existence of an ultraviolet fixed point, but we still lack a rigorous proof of this result.[9]

This clearly does not mean that the renormalisation group has no bearing on the constructive field theory project however; it merely serves to remind us that more work is needed to decisively determine whether particular QFT models have ultraviolet fixed points.[10] In fact, this has been one of the main focuses of the constructive field theory programme in recent years. As Hancox-Li (2015b) describes in some depth, mathematicians working in this field have attempted to probe the renormalisation group flow of QFT models using non-perturbative approximations with well controlled error bounds, the hope being that rigorous results about the existence or non-existence of ultraviolet fixed points can be established.

I suspect, therefore, that Fraser is getting at something else here. It is the attempt to construct continuum QFT models by means of the continuum limit of a cutoff theory itself which Fraser is referring to as an "approximation technique".[11] While I have been, somewhat sloppily, speaking about taking *the* continuum limit there are really number of ways of implementing this sort of constructive strategy: in part because there are many ways to write down cutoff models, but also because questions of convergence depend on the topological properties of the background space, which can be handled in different ways. The fact that a standard construal of the continuum limit does not converge does not necessarily mean that a continuum QFT model cannot be obtained by some other means. This is especially true if we are working with an open ended conception of what a continuum QFT is—that is if we are willing to weaken the Wightman axioms in the face of new developments as Fraser attests to.

While this suggests that it will be a difficult task indeed to prove that a particular QFT model does not exist on continuum space-time I do not think that this should push us towards the instrumentalist reading of the renormalisation group

[9] For a non-perturbative treatment of the renormalisation group flow of QED see Gies, H. and Jaeckel, J. (2004). Lüscher and Weisz (1987) is a standard reference for evidence of the triviality of ϕ^4 theory.

[10] Note that this point applies just as forcefully to cases in which the ultraviolet behaviour of the theory seems to be well behaved, such as QCD. The asymptotic freedom of QCD has not been rigorously demonstrated, and this would surely be needed if a continuum version of the theory were to be concretely constructed.

[11] On this interpretation 'approximation' is being used to refer to a limiting process used to define a mathematical structure. This is clearly a different meaning from my use of the term in the previous chapter. In her discussion of this issue Fraser (2011) cites a passage in Wightman (1984) which refers to the construction of a continuum model by means of a series of lattice models as a "lattice approximation", which corroborates this reading.

which Fraser advocates. Speaking of the apparent lack of an ultraviolet fixed point in the renormalisation group flow of the ϕ^4 theory Rivasseau (2014, 271) remarks:

> Mathematicians often cure this defect by reformulating the flow equations in another space... It is perhaps in this direction that one might reach in the future some positive results concerning [ϕ^4 theory], although for the moment we have no geometrical insight on how to change the ordinary space on which the parameters such as the coupling constant live.

Rivasseau is clearly optimistic about the existence of a continuum ϕ^4 model on Minkowski space-time here, but he is not dismissing the relevance of previous work on the renormalisation group analysis of its ultraviolet behaviour. What he is proposing is that it may be possible to modify the space in which the renormalisation group analysis is being conducted so as to obtain a convergent continuum limit. This seems to me to be the appropriate response to these results; the challenge theories like ϕ^4 theory and QED pose for the constructive field theory project calls for more engagement with the renormalisation group framework, not less.

One moral of the foregoing discussion then is that very little can be said with certainty about the existence or non-existence of continuum QFT models on the basis of extant work within the renormalisation group approach. Claims that QED has been shown to admit no continuum formulation, not uncommonly found in the physics literature, are premature. Nevertheless gaining control over the ultraviolet behaviour of four dimensional QFTs is a crucial challenge facing the constructive field theory project and despite these qualifications the renormalisation group framework clearly remains relevant to this problem. This is one respect in which renormalisation group methods amount to more than a "procedure for deriving empirical predictions" within the QFT programme.

The next section returns to the unsolved puzzle about the perturbative approach to QFT discussed in the previous chapter. In this context too, the renormalisation group has important ramifications which are difficult to make sense of on an instrumentalist reading.

4.4 Explaining the Success of Perturbative Quantum Field Theory

The renormalisation group approach is often said to have put perturbative QFT, and the perturbative renormalisation procedure in particular, on a firmer conceptual footing.[12] This section offers a diagnosis of why this is in fact the case. The crucial point is that the renormalisation group approach largely solves the key problem with the QFT perturbation theory which I identified in the previous chapter: the lack of explanation for its success.

Where the previous section looked at what happens when we take the renormalisation group transformations in the ultraviolet direction, the crucial issue here is how the renormalisation group flow behaves as we decrease the value of the cutoff. As was discussed in §4.2, this kind of coarse-graining operation inevitably gives rise to non-renormalisable, as well as renormalisable, interaction terms in the new Lagrangian and the full renormalisation group flow lives in an infinite dimensional space of couplings. In the perturbative context the distinction between renormalisable and non-renormalisable interactions was based on whether ultraviolet divergences in the expansion could be absorbed into a finite number of parameters in the Lagrangian—this turned out not possible for terms whose coupling had negative mass dimension. The mass dimension of the coupling in the Lagrangian also has important implications for how it behaves under renormalisation group transformations. Though the equations governing the renormalisation group flow are highly non-linear and intractable, we can say something surprisingly general about how the renormalisation group flow behaves in the low energy regime: as μ decreases it is attracted to a finite dimensional surface spanned by the parameters associated with renormalisable terms in the Lagrangian.[13] In the case of scalar field theory, for instance, non-renormalisable interaction terms, like $\lambda_6\phi^6$, $\lambda_8\phi^8$, and so on, have very weak effects at low energies—at sufficiently large length scales almost any scalar QFT will look like the familiar ϕ^4 model (see figure 4.3 for an illustration).

Some indication of this behaviour is provided by naive dimensional analysis. As-

[12] See Lepage (2005). Similar discussions can be found in many modern QFT textbooks. Huggett and Weingard (1995), Butterfield and Bouatta (2014) and Hancox-Li (2015b), chapter 4, all advocate the significance of the renormalisation group for understanding perturbative renormalisation.

[13] The distinction between renormalisable and non-renormalisable couplings directly corresponds to the distinction between relevant, irrelevant and marginal couplings in condensed matter physics. As this terminology suggests, couplings with negative mass dimension are dubbed irrelevant because they have very weak affects on the critical behaviour of condensed matter systems. For more on the connections between the application of the renormalisation group in QFT and condensed matter theory, see Binney et al (1992), chapter 14.

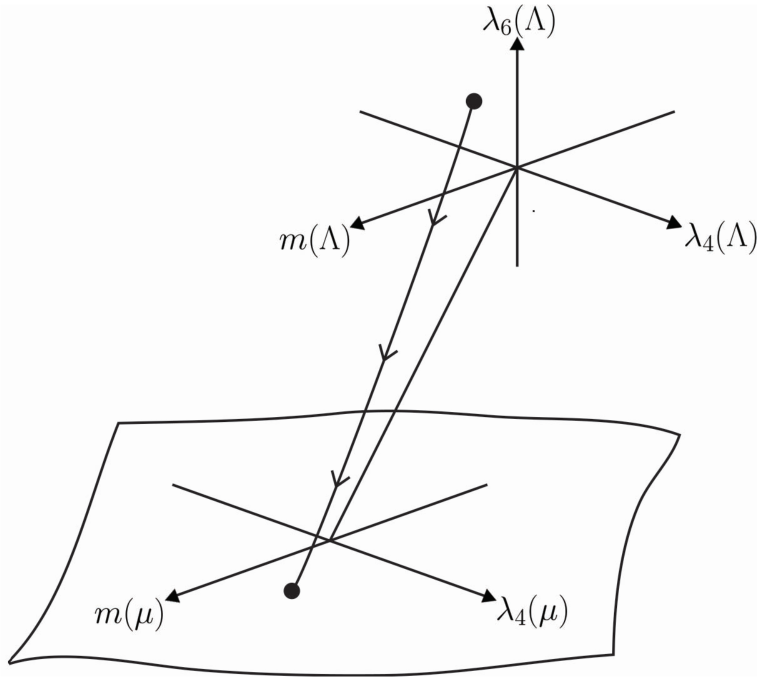

Figure 4.3: The renormalisation group flow of scalar field theories to a surface spanned by renormalisable parameters.

suming that the only dimensional quantity at play is the ultraviolet cutoff, we can rewrite the couplings of non-renormalisable interactions in terms of a dimensionless factor multiplied by an inverse power of Λ. In the case of scalar field theory for instance, we can define, $\lambda_{2n} = c_{2n}\Lambda^{4-2n}$, where $\{c_{2n}\}$ is a new set of dimensionless parameters. The contribution of a particular interaction term to observables at some energy scale E can thus be estimated to be of the order $c_{2n}(E/\Lambda)^{4-2n}$. If $E \ll \Lambda$ then we can expect the affect of non-renormalisation terms (with $n > 2$) to be heavily suppressed. This is not a foolproof argument however because the new couplings c_{2n}, though dimensionless, can still depend implicitly on the energy scale, and could, in principle compensate for the inverse powers of Λ. A more rigorous, and genuinely non-perturbative demonstration of this behaviour, originally due to Polchinski (1984), is based on linearising the renormalisation group equations by treating the relevant couplings as infinitesimals. If we do this it is possible to show that the non-renormalisable couplings at an initial cutoff scale Λ are almost completely absorbed into variations in renormalisable couplings as we coarse-grain to a much lower cutoff scale $\mu \ll \Lambda$, up to powers of $(\mu/\Lambda)^{|d|}$, where d is the mass dimension of the non-renormalisable coupling.[14] This analysis only tells us about the renormalisation group flow for small values of the coupling (a limitation I will return to in the following section) but since perturbation theory rests on this assumption it will be sufficient for present purposes.

One upshot of this result is that it allows us to understand why it is that the

[14]See Duncan (2012), chapter 17, for an especially clear presentation of this result.

standard model, and its component theories, are renormalisable. As I mentioned in the previous chapter, renormalisability was traditionally treated as a theory selection principle, but there was always something mysterious about this way of thinking. Why should the world be structured in such a way as to facilitate perturbation approximations? Without the renormalisation group results just described the fact that the Lagrangians of empirically successful theories like QED and the standard model are renormalisable would seem to be a lucky coincidence. We can now see, however, that limiting one's attention to renormalisable interaction terms is a very reasonable thing to do. Suppose we take Λ to be at a very high energy scale at which new fields, or effects beyond the scope of QFT entirely, come into play—in the most extreme case this could be the Planck scale at which quantum gravity effects are expected to become important. On the basis of the above argument, we should then expect physics at currently accessible energies to be very well described by a renormalisable Lagrangian, as the effects of non-renormalisable interaction terms will be heavily suppressed by the large ultraviolet cutoff.

While this new perspective on renormalisability has been emphasised in the nascent philosophical literature on renormalisation theory,[15] this does not get to the heart of how the renormalisation group illuminates the QFT perturbation theory in my view. After all, the fundamental puzzle about perturbative renormalisation identified in the previous chapter was why it works at all, even granting the renormalisability of the interactions under consideration. I argued there that the perturbative approach is incapable of answering this question on its own because it lacks a non-perturbative characterisation of the underlying physics. The renormalisation group framework allows us to respond to this challenge however. In fact, all of the machinery needed to explain the success of the standard perturbative renormalisation procedure in physical terms has already been introduced.

You will recall that the perturbative renormalisation procedure can be divided into three steps. In the first, regularisation phase, divergent momentum space integrals are replaced by convergent expressions. There are many ways to do this but the basic approach in the case of ultraviolet divergences is to put a finite upper bound on the integral (I'll come back to infrared divergences shortly). The renormalisation group framework provides a physical interpretation of what is going on here which was lacking in the original perturbative method. 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 legitimise the perturbative cutoff on momentum space integrals in the same

[15] Butterfield and Bouatta (2014) and Hancox-Li (2015a) both stress how the renormalisation group impacts on the concept of renormalisability.

terms, on the understanding that the effects of physics beyond the cutoff can be absorbed into the Lagrangian. Of course, if we are viewing the perturbative cutoff in this way then we should be taking non-renormalisable interactions into account, which is not done in the standard perturbative method—this will become important shortly.

The second step of the renormalisation procedure is to redefine the expansion parameter so as to remove the diverging dependence on the cutoff in the perturbative coefficients. But what licenses us to do this? The best answer we can give within the perturbative method itself is that it is required if the expansion is to be mathematically well behaved, and yield sensible predictions, as the cutoff is removed. But this is a purely instrumental rationale. The fact that this prescription produces finite results does not give us any reason to think that they ought to be good approximations to physical quantities like scattering cross sections. It is this manoeuvre, more than anything else, which gives perturbative renormalisation its ad hoc flavour. The renormalisation group analysis of the low energy regime provides a physical justification for removing the hypersensitivity to the cutoff however. It turns out that physical quantities like scattering cross sections actually are very weakly dependent on the cutoff at low energies. Choosing the expansion parameter so as to minimise 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 it is supposed to approximate.

We have seen that the low energy behaviour of a cutoff theory can be almost entirely parameterised by the renormalisable couplings. All that varying the value of the cutoff at high energies can do to the renormalisation group flow then is move us around on the surface of renormalisable couplings. This means that fixing the values of the renormalisable couplings at low energies by measuring the values of a finite number of scattering cross sections absorbs almost all of the cutoff dependence of physical quantities. The upshot of the Polchinski (1984) argument outlined above is that the correlation functions of a cutoff model only depend on the cutoff through powers of E/Λ, for scattering process taking place at kinetic energy E.[16] This amounts to a non-perturbative demonstration that

[16]These renormalisation group results are often described as indicated a general insensativity of low energy behaviour to physics at much higher energy, also associated with the decoupling theorem (Appelquist and Carazzone, 1975), and I will sometimes slip into this way of talking in what follows. As Williams (2015) stresses however, we need to be somewhat careful here. While the renormalisation group indicates that the cutoff dependence can be absorbed into the low energy renormalisable couplings if there are parameters in the theories Lagrangian with a positive mass dimension—corresponding to so-called relevant operators—this masks an implicit hyper-sensitivity to the value of the cutoff. This is the root of the naturalness problem with the Higgs sector of the standard model. I won't have anything more to say about this here though it will ultimately be an important issue for investigating the status of particular QFT models.

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 renormalised in the manner described in the previous chapter the only dependence on the cutoff that remains in perturbative approximants takes the correct form of powers of E/Λ. The perturbative renormalisation procedure then is fundamentally about ensuring that our approximations have the right scaling behaviour on the view I am advocating here, not about ensuring that they are mathematically well behaved as the cutoff is removed.

We can also provide a natural justification for the removal of the cutoff in the final stage of the perturbative renormalisation procedure within this framework. Assuming that the cutoff scale is much higher than the energy scale we are trying to describe, the E/Λ cutoff dependence of renormalised perturbative approximants, 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 neglected. What we are doing when we take the cutoff to infinity in perturbative calculations is setting $E/\Lambda \to 0$, and completely eliminating the dependence on the cutoff. As I mentioned in the previous chapter, removing the cutoff has significant computational benefits, and since the renormalisation group gives us a handle on the kind of errors that result from doing so it is pragmatically justified.

This reinforces the claim I made in the previous chapter that taking the cutoff to infinity in perturbative calculations and taking the continuum limit of a cutoff QFT model are very different things. As we saw in §4.3, when it comes to the continuum limit it is the behaviour of the renormalisation group flow in the ultraviolet direction which is of interest. In contrast, the justification for removing the cutoff on perturbative approximations I have just described is based entirely on the behaviour of the renormalisation group flow at low energies, and the fact that physical quantities are almost entirely independent of the cutoff in this regime. In fact, nothing we have said in this section hangs on the existence or non-existence of a continuum limit at all. This explains why the worries about the continuum limit of QED raised by its apparent lack of an ultraviolet fixed point do not lead to doubt about the reliability of QED perturbation theory.

What about the infrared divergences that appear in the perturbation series of theories with massless particles? In this case the renormalisation group does not help us and other theoretical resources are needed to understand what is going on. But while the details of the treatment of infrared divergences are different, and will not be discussed in detail here, there is reason to believe that the basic explanatory story here is essentially the same. As in the ultraviolet case, there is strong evidence that infrared divergences are unphysical artefacts of the perturbative approximation scheme. For one thing, there are rigorous results, admittedly still within the context of perturbative QFT itself, which tell us that the infrared divergences in QED and QCD perturbation series ultimately cancel when in and out states are treated appropriately and do not manifest in physical observables.[17] Furthermore, we have strong non-perturbative evidence that correlation functions over relatively short length scales are very insensitive to the details of a long distance cutoff. In the context of lattice QCD the impact of finite size effects on correlation functions can be estimated quite precisely, using methods like chiral perturbation theory.[18] In general, the result is that correlation functions at energy scales E depend on the infrared cutoff through powers of $\exp(-EL)$, where L is the length of the finite volume lattice. As in the ultraviolet case then, the singular dependence on the infrared cutoff is unphysical and systematically removing infrared divergences can again be understood as a matter of ensuring that our perturbative expressions share the relevant behaviour of the quantities they are supposed to approximate.

While I certainly would not claim that all of the conceptual puzzles with the perturbative approach have been solved in the preceding discussion, we have certainly improved the situation as we left it in the previous chapter.[19] The perturbative renormalisation procedure, in particular, has been shown to have a sound physical motivation and grounding in a non-perturbative characterisation of QFT models. This paves the way for further philosophical engagement with perturbative QFT since the final outstanding problem—the justification problem—has been at least partially addressed. It also represents another context in which renormalisation group methods have genuine philosophical import. As I argue in the next section, the explanation of the success of perturbative renormalisation which has just been sketched also leads a more general shift in our understanding of QFT; it motivates a realist view of cutoff QFT models.

[17]The Bloch–Nordsieck theorem is the classic result for the cancellation of infrared divergences in QED, see Duncan (2012) for an exposition. New subtleties arise in the context of QCD however. Though general cancellation results can be obtained isolating the infrared safe part of perturbative approximants is often non-trivial. See Muta (2010) for an in-depth discussion of infrared divergences in QCD perturbation theory.

[18]See Golterman (2011) for a discussion of the estimation of finite size effects in lattice QCD.

[19]One issue which I have not addressed here, for instance, is the scheme dependence problem mentioned in chapter 3. There is also a great deal more to be said about how renormalisation group methods are employed in the context of contemporary perturbative QFT, and the way that non-renormalisable theories are dealt with perturbatively in the modern effective field theory methodology. I take my discussion here to be a first step in the broader project of philosophically engaging with perturbative QFT.

4.5 Quantum Field Theories as Coarse-Grained Representations

As I stressed in chapter 3, one of the key features of perturbative QFT, in its original incarnation, that made it problematically different from perturbative methods employed elsewhere in physics was the lack of a physical characterisation of the system whose properties are being approximated. The previous section detailed non-perturbative results which, I claimed, help to rationalise the standard perturbative method. It may not be obvious from the discussion so far however what the physical picture underlying this account is supposed to be. This section addresses this issue explicitly and argues that the explanatory story developed in the previous section is, in fact, of the sort that should satisfy the scientific realist. Properly understood, the renormalisation group analysis just described furnishes an explanation of the empirical successes of the QFT programme in terms of the way the world is. Spelling this out will lead to a broader thesis about how QFT models relate to reality. The key claim is that cutoff QFT models ought to be taken seriously as faithful representations of the physics of our world at length scales far removed from the ultraviolet and infrared cutoffs, though, as we'll see, this view of QFT systems as coarse-grained models will also apply to continuum formulations of realistic QFTs, if they exist.

Recall that in the methodological framework set out in chapter 2 the process of stating a physical theory was divided into three parts. In the first structural phase mathematical structures which make up the theory's formalism are supplied. In the second semantic phase some physical content is assigned to these structures—this is what is usually meant by providing an interpretation of the theory, and the so-called standard account, which I was happy to provisionally accept, took this step to involve a specification of a set of possible worlds in which the theory is true. In the final epistemic phase, which I argued also needs to be taken into account if we are to make sense of the full range of conceptual issues raised by physical theories, something is said about how these interpreted structures ought to be understood as representations of the actual world. Specifically, on the minimal framework I put forward, we need to identify which properties of a theory's models we should take to faithfully represent, which we should take to misrepresent, and which we should simply remain agnostic about. My approach here will be to examine how each of these steps play out in the case of cutoff QFTs in light of renormalisation group results just outlined (I'll return to continuum models at the close of the section). What will end up emerging from the discussion is an explanation of the extant empirical successes of the QFT programme in terms of the representational success of these theories.

When it comes to the first phase of writing down mathematical structures, we saw in §1.2.2 that ultraviolet and infrared cutoffs can be implemented in a number of different ways, leading to a diverse collection of mathematical models which I have been referring to collectively as cutoff QFTs. For the moment though keep in mind the canonical case of a finite volume hypercubic lattice QFT—we will see that focusing on these systems is not as ad hoc as it might initially seem shortly. As I discussed in chapter 1, lattice QFT models are very well understood mathematically and can be quantised using both the canonical and functional integral approaches. The resulting structure is a quantum system with a finite number of degrees of freedom that is highly analogous to the lattice systems used in condensed matter physics to describe a quantum crystal.

How should we go about assigning physical content to these structures? On the one hand, QFT models (cutoff and continuum) will inherit puzzles from ordinary quantum mechanics here. Owing in large part to the enduring controversy surrounding the measurement problem, there is very little agreement about what any quantum theory says about the world. In particular, there has been much debate in recent years about the ontological status of quantum states and these issues will inevitably spill over into QFT, though so far they have been much less discussed in this context.[20] Additional semantic issues have also been raised in the philosophy of QFT literature. A key focus of philosophical work on QFT has been the status of particles and fields in QFT systems, and specifically whether they can be counted as part of the fundamental furniture of possible worlds that are exactly described by QFT models. Furthermore, in the case of continuum QFT models the existence of unitarily inequivalent representations poses puzzles for understanding what these theories say about the world (as will be discussed further in chapter 6).

I don't have anything new to say about these issues here. What I do want to stress however is that the imposition of ultraviolet and infrared cutoffs does not lead to any new problems at the semantic level. In fact, since systems with finite degrees of freedom do not admit unitarily inequivalent representations cutoff systems are easier to make sense of than continuum QFTs (again, this feature of cutoff models will be discussed further in chapter 6). The issues which remain here are essentially those that plague quantum physics in general. Since this point is not uncontroversial I will pause to defend it.

It is sometimes suggested that cutoff QFTs are incapable of completely describing a possible world at all. This would mean that the standard account cannot be applied in this context and there is a serious puzzle with making sense of their

[20]Recent papers discussing the fundamental ontology of ordinary quantum mechanics are collected in Ney and Albert (2013). Wallace and Timpson (2010) and Arntzenius (2014) chapter 3 extend these discussions to the context of QFT systems.

content. Wallace (2006) describes cutoff QFT as "intrinsically approximate", for instance, and the claim that QFTs which incorporate an ultraviolet cutoff 'break down' at high energies is not uncommonly found in the scientific and philosophical literature on effective field theories.[21] In my view this is another area where we need to be careful to distinguish perturbative and non-perturbative issues. It is true that if a finite ultraviolet cutoff is maintained in the perturbative framework, for instance when non-renormalisable theories like the Fermi model of weak interactions are treated perturbatively, the resulting approximations become badly behaved at energies that exceed the cutoff—in particular, unitarity is violated. In the non-perturbative context however there is no sense in which a lattice theory 'breaks down' at high energies—the system simply does not admit states with momenta in excess of the ultraviolet cutoff. As far as I can tell it makes perfectly good sense to talk about the possible worlds in which a lattice QFT model is true: these will presumably be worlds in which space-time actually has a lattice structure at the fundamental level.[22]

The really interesting issues with cutoff QFTs, I think, arise in the final epistemic phase when the question of representational success is broached. How do these systems relate to the physics of the actual world? Before the rise of the renormalisation group approach the default answer was probably not at all. As we have now discussed in some detail, cutoff models have long been employed in the constructive field theory programme as an intermediary step towards obtaining continuum formulations of realistic QFTs, and in this context can simply be viewed as an instrument in the search for continuum models. A similarly dismissive attitude towards the physical significance of the cutoff can be found in the high energy physics literature prior to the influential work of Wilson. The success of the renormalisation group approach has transformed physicists' attitudes towards the cutoff however. Nowadays lattice QFT models, and perturbative calculations which maintain a finite cutoff, are taken seriously as descriptions of actual physics.

This transformation should, I think, be understood as taking place at the epistemic rather than semantic level. The shift in late 20th century high energy physics towards thinking of QFTs as so-called effective field theories was not a revolution in how QFT systems are formulated but was rather a reassessment of their representational status. How did this come about? The first step is to

[21] See McKenzie (2012) chapter 6 and Williams (2016) for discussion of the break down of effective field theories in perturbation theory at the ultraviolet cutoff scale.

[22] This is not, of course, to say that the violations of unitarity found in perturbation theory have no significance whatsoever. In many contexts the appearance of these effects can be used to estimate the scale of a physically significant cutoff at which new physics kicks in—in the case of the fermi model of weak interactions, for instance, the natural physical cutoff is the mass of the electroweak W bosons.

recognise that the results discussed in the previous section show that cutoff QFT models are empirically successful in their own right. We have seen that the values of observable quantities associated with a cutoff model at energies much lower than the ultraviolet cutoff, for instance, are thought to differ from standard perturbative approximations via powers of E/Λ, which, if Λ is taken to be sufficiently large, can obviously be ensured to be much smaller than the experimental error associated with the measurement of the quantities in question. This means, in effect, that the myriad predictive successes of QFT perturbation theory accrue to cutoff versions of realistic QFT models.

Furthermore, in the wake of the success of the renormalisation group approach, more evidence of the predictive capacities of cutoff models has come to light, with the success of non-perturbative approximations generated by numerical simulations of lattice QFT systems. These numerical methods have proven to be our most predictively powerful probe of the low energy behaviour of QCD, where interactions are strong and perturbative approximations are no longer useful. I mentioned in the previous section that the Polchinski (1984) renormalisation group analysis of the low energy regime is based on the assumption that the coupling is small and consequently cannot be trusted to give us accurate estimates of the cutoff dependence when the coupling is large. In the low energy region of QCD we expect the appearance of so-called intrinsically non-perturbative contributions to correlation functions proportional to factors like $\exp(-1/g)$ (where g is the QCD coupling) that are 'invisible' to perturbation theory. It turns out to be possible to estimate the dependence on the lattice spacing a in this context however—the basic result being that it vanishes like powers of a as $a \to 0$.[23] Again, this means the dependence on the ultraviolet cutoff takes the form of inverse powers of $\Lambda = 1/a$. As I briefly mentioned in the previous section, finite size effects in lattice QFT simulations are also exponentially suppressed. Thus the basic rationalisation of the success of lattice QFT simulations is essentially the same as that of perturbative approximants: this method produces good approximations because the observable quantities we are trying to predict are very insensitive to the values of the lattice spacing/ultraviolet cutoff and the volume/infrared cutoff.

Should we take the empirical success of cutoff formulations of realistic QFTs to confirm their claims about the physical world? I argued in chapter 2 that we would do best to move away from the idea that the content of any physical theory should be accepted en bloc. The kind of representational success which can reasonably be attributed to a theory invariably involves taking some parts to be veridical and others to misrepresent. As I discussed there, there are numerous

[23] See Duncan (2012), chapter 9, for a discussion of intrinsically non-perturbative effects and Weisz (2011) for a detailed discussion of the estimation of lattice defects in lattice QCD.

factors which can influence these epistemic judgements. But I also put forward a master criterion for identifying the belief worthy content of a theory based on the familiar realist idea that we ought to reserve our optimism for the parts of the theory that are really responsible for its predictive achievements. I offered a counterfactual characterisation of what this means according to which a property of a model contributes to a prediction, and therefore is supported by it, if varying it spoils the relevant empirical success. Conversely, if a model's empirical adequacy is robust under variations in one of its properties then it does not essentially contribute to its success and should not be taken representationally seriously.

Coupling this approach with what we have already learned about the scaling behaviour of cutoff QFTs leads to a general thesis about the representational status of these models. The renormalisation group analysis of cutoff QFTs tells us that, once the quantities of a model have been reparameterised in terms of couplings fixed by low energy measurements, its large scale behaviour is very insensitive to the value of the ultraviolet cutoff and the details of the dynamics at the cutoff scale. Less systematic, but still compelling, evidence suggests that the empirical consequences of cutoff QFT models are also robust under variations in the value of the infrared cutoff and the degrees of freedom at arbitrarily large distances. What this means is that the very small and very large scale properties of these models can be varied without undermining its empirical adequacy, which is exactly the criterion I put forward for identifying features of a model that do not contribute to its success. The upshot then is that we should not trust these features of cutoff models to accurately represent anything in reality.

While there are speculative motivations coming from quantum gravity and cosmology for thinking that the world actually has some kind of short distance cutoff at the fundamental level, and really is finite in volume, these are, as far as I know, unanswered questions. As far as the empirical success of the QFT programme goes then, I suggest we ought to remain agnostic about these issues. On the other hand, we have good reason to believe that any particular implementation of the cutoffs misrepresents the way the world is. For one thing, the fact that there are many ways of writing down a cutoff version of a realistic QFT which make no difference to its relevant observable consequences suggests that none of these cutoff schemes has a reasonable chance of describing reality as it is. But there are also compelling reasons external to the QFT programme for viewing the claims that our current QFT models make about the small and large scale structures of the world to be false—due to the clash with general relativity we expect radically new physics to come into play at both of these scales. The model's claims about the world at the cutoff scales then should be viewed as idealisations.

If the details at the cutoff scale are not responsible for its predictive accuracy though, what is? The answer which emerges from the discussion of this chapter is that it is the system's correlation functions over length scales much longer than the ultraviolet cutoff and much shorter than the infrared cutoff which really underwrite its success. These are the fundamental properties of QFT models that remain almost entirely invariant under variations in the cutoffs. They are also modally connected to directly observable quantities like scattering cross sections, in the sense that QFT models with different correlation functions at the energy scales we are currently probing in collider experiments will necessarily lead to discernibly different predictions. Following through on the idea that the realist commits themselves to the aspects of scientific theories that fuel its success then, we should take these quantities to faithfully represent coarse-grained aspects of reality. This is, of course, closely allied to the stance towards cutoff QFTs set out in Wallace (2006, 2011). Unlike Wallace however, I have advocated a very similar attitude towards both the ultraviolet and infrared cutoff (more on this point in chapter 6); we will see another difference between Wallace's approach and mine shortly.

Some clarificatory remarks are in order. First of all, it may be helpful to point out that the idea that a model captures coarse-grained features of its target while misrepresenting it at other scales is far from a novel one in the broader landscape of scientific modelling. We are very familiar with the use of continuum models to describe the properties of systems ultimately composed of a finite number of particles in fluid mechanics and condensed matter physics, for instance. In statistical mechanics so-called lattice gas models, which discretise the states of a system of particles so that they can only take positions on the sites of a spacial lattice, are used to describe fluid to solid phase transitions. Models like this clearly idealise the short scale structure of the concrete systems they aim to describe but nevertheless accurately capture many of their large scale properties. We can also find more mundane instances of this kind of coarse-grained representational success in non-scientific contexts. Consider a digital photograph, for instance. The discrete changes in colour at the scale of individual pixels clearly do not reflect real variations in the target but, again, digital images accurately represent some features of their subjects on length scales much larger than the pixel length. The kind of representational success I am attributing to realistic cutoff QFTs can be thought of as broadly analogous to these more tangible cases.

I should also stress here that my claim that cutoff models get coarse-grained features of the world right is not to be confused with the claim that they merely capture observables, in the philosopher's sense. In particular, the view I am putting forward here is not to be assimilated to the constructive empiricist claim that these models are merely empirically adequate. For the empiricists an entity or property posited by a theory is said to be observable if it is a possible object of unmediated human perception. In the QFT context, scattering cross sections are perhaps the obvious candidate for quantities that might satisfy this definition. But the correlation functions I am taking to be the locus of our epistemic commitments here are not, in general, observable in this sense—they are, of course, supported by empirical evidence but they are not directly measured in scattering experiments. It is worth pointing out that when I say coarse-grained here I mean coarse-grained relative to a suitable ultraviolet cutoff, which in the case of the standard model and its sub-theories will be very small distance scales indeed. In the case of QCD, for instance, I am suggesting that correlation function on scales much smaller than the diameter of the proton should be taken representationally seriously; these properties will clearly not be observable by the empiricist's lights.[24]

A final issue which needs to be addressed is the plethora of distinct cutoff formulations of any realistic QFT. At first blush, the fact there are many different ways of freezing out the degrees of freedom associated with arbitrarily high and low momenta might seem to stand against taking these systems seriously from a foundational perspective. We can now see that this is not as problematic as it initially appears however. The existence of many cutoff versions of QED, that are all essentially indistinguishable at sufficiently low energies, might be a problem if we took them to be rival hypotheses about the fundamental nature of the world. But, as I have argued, cutoff QFTs should not be viewed in this way. Lattice QFTs and systems which employ a smooth ultraviolet cutoff (as briefly as described in §1.2.2) disagree on their short scale structure, but we should not take these claims representationally seriously in any case on the view I have developed here. In contrast, these systems agree about everything we have any right to believe, namely the correlation functions or scales intermediate between the ultraviolet and infrared cutoffs. There is no conflict in the claims about the world that are warranted by the success of these theories then.[25] We can, I think,

[24]There is a potential concern here about how the representational content of these correlation functions should be understood however, given the controversy surrounding the semantics of quantum theory in general. Correlation functions are expectation values of products of field operators and consequently how we interpret quantum operators and states will determine how their physical significance is spelled out. Consequently, the precise characterisation of the coarse-grained features captured by cutoff QFT models will depend, to some extent, on one's broader view of quantum theory. This should not be too surprising however; familiar interpretive issues with quantum mechanics are bound to reoccur in the QFT context, as I have already mentioned. This need not mean that the philosophy of QFT is held hostage to the measurement problem. We can arguably deal with a certain amount of semantic ambiguity here; as long as we maintain a broadly ontological view of quantum states, steer clear of epistemic and pragmatic views, such as that of Healey (2012), correlation functions will end being robust physical quantities.

[25]I develop this point further in §7.4 in the context of concerns about underdetermination among QFT models.

understand the ambiguity associated with the imposition of the cutoffs in a similar way to idealisations elsewhere in physics. In the case of the simple pendulum model, briefly invoked in §2.3, what justifies sending the air resistive force to zero is that its precise value makes no difference to the time period (to the required level of accuracy). Similarly, since the very small/large scale properties of QFT models make no difference to their relevant empirical predictions are free to deal with them in whatever way we find convenient. This explains why it is legitimate to focus on lattice QFTs, despite the existence of other ways of imposing the cutoffs.

Now we have dealt with the cutoff case in some detail, I will close this section with some comments about continuum QFTs, as promised. What does all this mean for the continuum theories sought in the constructive field theory programme? If we had continuum formulations of realistic QFTs the basic morals of the analysis of cutoff systems just detailed would also apply to them. Suppose that there is a breakthrough in mathematical physics and continuum versions of a theory like QED or QCD are constructed. As I mentioned above continuum systems raise new semantic puzzles, but putting these issues aside, what impact would this have on the epistemic question of what we ought to believe about the world on the basis of the QFT programme? In my view, essentially none at all (though I will add a caveat to this claim in §6.5). Though these continuum theories would purport to describe the world on arbitrarily large and small length scales these are not claims which we are licensed to believe for the same reasons that we have just run through in the case of cutoff systems. We have seen that physics at the energy scales we currently have access to in collider experiments, is effectively isolated from what goes on in these domains and consequently is not supported by the empirical success of QED and QCD. Furthermore, we know that these theories do not describe the world in these domains for external theoretical reasons, the need for a consistent theory of quantum gravity being the most immediate (I will have more to say on this point in the next chapter). Does this mean that the axiomatic and constructive approaches to QFT are philosophically barren? I think not. In the chapters that follow I will be developing the idea that studying continuum QFT systems can still be philosophically rewarding without furnishing new claims about the way the world is.

In contrast to Wallace's presentation of the dialectic, for me the lack of continuum formulations of realistic QFTs is not a key motivation for focusing on cutoff QFTs. What comes out of the discussion of this section is a general picture about how QFT models relate to the world; they capture coarse-grained features of the world. This is, I think, a precisification of the familiar claim in the physics literature that the standard model is an effective field theory. Constructing a continuum formulation of this theory will do nothing to change this and the cutoff formulation we have already has all that we need to articulate the claims about the world that it warrants. Furthermore, I have suggested that the coarse-grained representational success of QFT models furnishes a physical explanation of the empirical successes of the QFT programme. Indeed, I think that the view which emerges from the discussion of this section is a genuinely realist one—a claim which I defend in detail in chapter 7.

4.6 Understanding the Significance of the Renormalisation Group

I want to conclude this chapter by returning to the debate over the significance of the renormalisation group approach to the philosophy of QFT we started out with in §4.1. I have argued that the renormalisation group approach does indeed have important implications for philosophical engagement with the QFT programme. In terms of the framework set out in chapter 2, we can see the renormalisation group results as coming into play in both the structure-specification and epistemic phases of analysing QFT. In the former case, I have argued that this framework is relevant to the question of whether interesting continuum QFT models exist; in the latter, I claimed that it informs epistemic judgements about which aspects of QFT models ought to be taken to get things right about the world.

There is room for clarification about how the conceptual significance of renormalisation group methods ought to be understood however. Indeed, there might seem to be a puzzle about how they can illuminate the foundations of physics at all. To caricature somewhat, there sometimes seems to be an assumption amongst philosophers of physics that all that really matters about a physical theory, for the purposes of foundational and interpretive study, is its fundamental structure and the theoretical principles it invokes. The way that results are derived from this base in scientific practice is a pragmatic concern which does not impact on the philosophical analysis of the theory. The renormalisation group does not furnish us with a novel theoretical characterisation of QFT models however; it is essentially a method for studying behaviour of physical systems at different scales. I suspect that this way of thinking about the interpretive project is a key factor behind Fraser's characterisation of the renormalisation group as an essentially pragmatic device with limited bearing on the foundational issues addressed in the philosophy of QFT.

This austere methodological stance does not seem to do justice to the kind of arguments we actually find in the cut and thrust of debate about particular physical theories however. Decoherence theory, for instance, does not invoke new physical principles or offer a novel reformulation of ordinary quantum mechanics—it provides a framework for studying the dynamical behaviour of quantum systems coupled to an external environment. Still, decoherence theory is often taken to have important implications for the philosophy of quantum mechanics. The dynamical behaviour of quantum systems it helps to capture, which is certainly not obvious from a causal inspection of standard axiomatic presentations of the theory, is invoked in various approaches to the measurement problem. Despite not making an intervention at the level of fundamental theoretical principles then, decoherence theory gives us information about properties of quantum systems that is at least putatively relevant to foundational questions about quantum theory.[26]

The philosophical import of the renormalisation group should be understood in similar terms. The renormalisation group framework is not primarily concerned with the fundamental structure of physical theories, but this does not mean it only deals with their empirical content. In fact, the renormalisation group framework as I have described it does not invoke the observable/unobservable distinction in any way. What it gives us is a way of studying how *any* property of a model, observable or not, varies with scale. Indeed, the quantities at the heart of the renormalisation group approach to QFT are correlations functions which, as I have stressed, are not directly observable. The real significance of the renormalisation group framework for the QFT programme, I think, is that it gives us a way of probing the modal, rather than empirical, properties of QFT models; it tells us how the properties of QFT models depend, and do not depend, on its small scale structure. This is crucial for gaining mathematical control over the continuum limit. But more importantly, counterfactual information of this kind is precisely what should be informing our judgements about which parts of a physical theory faithfully represent. I have argued that the renormalisation group, in conjunction with other theoretical resources, motivates a general epistemic attitude towards successful QFT models.

Much of the rest of this thesis will be concerned with developing and defending this epistemic stance. The following two chapters deal with objections coming from within the QFT programme, as it were. Chapter 5 addresses the fact that imposing a cutoff on the momentum violates key relativistic properties, most importantly Poincaré covariance. Given the central role that special relativity has played in the QFT programme this immediately raises worries about cutoff QFT models as objects of foundational study. Chapter 6 addresses another concern

[26]Of course, one might end up concluding that decoherence theory has no decisive impact on the debate surrounding the measurement problem. My point here is that it is not a priori obvious that it does not simply because it is not concerned with the fundamental structure of the theory.

about cutoff systems, namely that they are explanatorily inadequate. As we have seen, a system with both an ultraviolet and infrared cutoff has a finite number of degrees of freedom and consequently does not admit unitarily inequivalent Hilbert space representations. But unitarily inequivalent representations have been thought to play an essential role in explaining crucial physical phenomena like spontaneous symmetry breaking. Chapter 7 considers objections coming from the general philosophy of science and defends my claim that the view of QFT put forward in this chapter is a realist one.

Chapter 5

Emergent Relativity

5.1 An Apparent Clash with Relativity

The key claim of the last chapter was that cutoff QFT models provide accurate, if partial, representations of the actual world. One feature of these systems which urgently needs to be addressed is their non-relativistic properties. Cutoff models invariably violate the requisites of special relativity. This is particularly transparent in the case of lattice QFT models. Any lattice we can put on Minkowski space-time will not be invariant under rotations, translations or Lorentz boosts—at best it will respect a finite subgroup of these transformations. Furthermore, the dynamics of a lattice model will clearly not be local: influences propagate instantaneously over the cutoff length scale manifesting as non-zero commutators between field operators at space-like separated points.[1] And though there are other ways of constructing cutoff models, and perhaps further approaches which have not yet been discovered, there are, at least, heuristic reasons for thinking that relativity will always be compromised in one way or another; any way of suppressing high momentum modes is going to introduce a privileged length scale that will not be a relativistic invariant. While there are approaches to quantum gravity which attempt to discretise space-time in a fully relativistic way, such as loop quantum gravity and causal set theory, this inevitably involves substantial deviations from the QFT programme as we know it.[2]

Why should this raise concerns about the representational capacities of these models? On the one hand, there might seem to be a straightforward empirical

[1] I won't have much to say about non-locality in this chapter but my hope is that any worries the non-local nature of lattice systems generate can be addressed in an analogous way to my discussion of Poincaré symmetry.

[2] The possibility of a 'fundamental length' and the relationship of this idea to Poincaré covariance is a topic of much debate in the quantum gravity literature. Hagar (2009) is a recent philosophical discussion of this issue.

worry here. Special relativity is very well confirmed experimentally.[3] This, one might think, gives us good reason to believe that it is true, and therefore that any theory which is inconsistent with it is false. Shouldn't this be enough to rule out any kind of realism about cutoff QFTs? There is also a more subtle theoretical worry about the non-relativistic properties of cutoff models. Bringing together relativistic and quantum theory has long been a key motivation driving the QFT programme. In both axiomatic formulations of QFT and the more informal discussions of physics textbooks, relativistic properties like Poincaré covariance are often presented as constitutive of what a QFT is. Consequently, the violations of relativity engendered by a momentum cutoff might be thought to undermine the theoretical assumptions underlying the QFT programme in an unacceptable way.

This second line of objection comes to the fore in Fraser and Wallace's debate about the formulation of QFT. While both authors agree that there is no empirical problem with cutoff QFT models, they take very different views about the theoretical status of relativity principles in QFT. Fraser consistently stresses the centrality of special relativity to her understanding of QFT:

> Quantum field theory is *by definition* the theory that best unifies quantum theory (QT) and the special theory of relativity (SR)... Since QFT = QT + SR, the project of formulating quantum field theory cannot be considered successful until either a consistent theory that incorporates both relativistic and quantum principles has been obtained or a convincing argument has been made that such a theory is not possible. (Fraser 2009, 550, her italics).

As passages like this one make clear, Fraser takes QFT to incorporate a commitment to Minkowski space-time structure and Poincaré covariance. In fact, this forms the basis for her central argument for rejecting cutoff QFTs in favour of the axiomatic approach. The Wightman and Haag-Kastler axioms, which incorporate Poincaré covariance explicitly, amount to a principled unification of relativistic and quantum principles according to Fraser, while cutoff QFTs compromise this project in an ad hoc way.

By contrast, Wallace is unfazed by the non-covariance of cutoff QFT models:

> It is true that those cut-off schemes we can actually concretely implement violate Poincaré covariance... But provided that the real cut-off is actually imposed by some Poincaré covariant theory, this will just be an artefact of the particular cut-off schemes we are using at the moment... Having said this, if Poincaré covariance turns out

[3]See Mattingly (2005) for a survey of phenomenological tests of relativistic space-time symmetries.

to be only phenomenological, so what? (Wallace 2011, 123).

Notice that while Wallace advocates the primacy of cutoff formulations of QFT he is clearly not wanting to assert that the world really is non-relativistic. Rather, his stance seems to be that QFT, as an approximate description of reality, does not make any claims about the fundamental nature of space-time at all. He does maintain, however, that cutoff models are approximately covariant in a sense which is sufficient to accommodate impressive empirical support for relativistic space-time symmetries (Wallace, 2006, 50-52).

This chapter takes up these issues raised by the apparent clash between relativity and cutoff QFTs. §5.2 examines the sense in which cutoff models can be said to be approximately Poincaré covariant, something Wallace says relatively little about. I sketch an approach to making the notion of an approximate space-time symmetry more precise and argue that there is reason to believe that, at energies far below the cutoff, cutoff formulations of realistic models like QED and QCD are approximately Poincaré covariant in a sense which is sufficient to placate worries about their empirical adequacy. §5.3 takes on the question of the theoretical status of relativity in QFT. At first blush, it is difficult to see how to negotiate Fraser and Wallace's dispute over this question. We seem to be faced with different conceptions of what QFT is: Fraser takes a commitment to Minkowksi space-time structure to be integral to QFT while Wallace does not. Returning to themes introduced in chapter 2, I argue that we should move away from the idea that certain theoretical claims about space-time structure are constitutive of QFT and instead question in what philosophical contexts we ought to entertain these assumptions. I distinguish between the internal question of whether, and to what extent, quantum theory and relativity can be consistently combined from the external question of what we ought to believe about the world given the successes of high energy physics. In the latter context, I claim, we have good reason to remain agnostic about fundamental space-time structure.

Before launching into these issues I want to briefly comment on a recent intervention in the debate surrounding relativity in QFT due to Bain (2013b). In the context of his discussion of the effective field theory approach to QFT Bain suggests that there is in fact a way of implementing a high energy cutoff which does not compromise relativity. He points out that much of the renormalisation group apparatus described in the previous chapter can be implemented in the context of perturbation theory without explicitly freezing out large momentum degrees of freedom. Dimensional regularisation methods, briefly described in §3.1.2, can be employed and the decoupling of high and low energy degrees of freedom can be inserted by hand using so-called matching conditions. According to Bain this approach, which following Georgii (1993) he calls "continuum effective field theory", amounts to an alternative cutoff variant of QFT, not considered by Fraser and Wallace, that does not violate Poincaré symmetry.

This claim is misguided however. I argued in chapter 3 that we ought to distinguish between the regularisation of divergent integrals in perturbation theory and the construction of cutoff models. It is true that the dimensional regularisation method ensures that no explicit Poincaré violations occur in any step of a perturbative calculation, indeed this is one of the main practical advantages of this approach. This has nothing to do with the non-perturbative project of constructing continuum QFT models on Minkowski space-time however—as I pointed out, dimensional regularisation has so far found no application in this context. What Bain calls continuum effective field theory is a purely perturbative approach to QFT and does not contribute to the issue Fraser and Wallace are arguing over: which non-perturbative mathematical structures should be identified as the correct formulation of QFT. I assume in what follows that anything worthy of the name cutoff QFT violates Poincaré covariance and continue to treat lattice QFTs as the exemplar of such a system.[4]

5.2 Approximate Space-Time Symmetries

Wallace defends cutoff QFTs against concerns about their non-relativistic properties by claiming that they are, nevertheless, approximately covariant. It is not immediately obvious what this means however, and one might even worry about whether it makes sense at all. Part of the problem here is that the notion of approximate symmetry itself stands in need of clarification. The burgeoning philosophical literature on symmetries in physics has almost exclusively focused on exact symmetries. Yet there are clearly many contexts in which inexact symmetries play an important role in physical science. In high energy physics, isospin, and the more general flavour symmetries which have been found to hold approximately in QCD, are prominent examples. What will turn out to be a more instructive case for our discussion here is the emergence of approximate rotational symmetry in condensed matter systems. Crystals are not rotationally invariant but their large scale behaviour, such as the propagation of sound waves, can be effectively isotropic and approximate continuum symmetries of this kind are often invoked in model building in condensed matter physics. While the existence of approximate symmetries is sometimes mentioned in the philosophical literature however there has been very little discussion about how they ought to

[4]Thanks to Michael Miller for helpful discussion about the notion of effective Poincaré symmetry in QFT and quantum gravity.

be understood.[5]

A comprehensive analysis of approximate symmetry in physics is beyond the scope of this thesis. Still, we need to say something about this issue if we are to assess the severity of the problem raised by the non-covariance of cutoff QFTs. In what follows I set out a simple-minded approach to understanding how a symmetry can hold approximately that goes some way towards clarifying the status of relativistic space-time symmetries in cutoff models. Whether these ideas have more general application is a question for another time.[6]

Start with the familiar characterisation of exact symmetries. An exact symmetry of a geometric figure, function, physical theory, or other object, is a transformation that leaves its properties (or some specified subset of its properties) invariant. The complete set of these transformations forms a group structure. The Poincaré group, for instance, is the set of transformations which preserve intervals on Minkowski space-time, consisting of spacial rotations, Lorentz boosts and translations. An important class of physical symmetries are transformations which leave the dynamics of a physical theory or model invariant. Much of the philosophical debate about exact symmetries in recent years has focused on how to understand these dynamical symmetries, and how they bear on the semantic content of a theory.[7] Formally at least, what we mean when we say that a QFT model has a symmetry is that its partition function is invariant under the transformation in question.[8] It is worth mentioning that we can also talk about the symmetries of particular states of a physical theory in the same terms—these will, again, be transformations that map the state in question to itself. Symmetries (and asymmetries) of states will become important in the discussion of spontaneously broken symmetry in the following chapter, but for our current purposes it is symmetries of theories which will be our primary concern.

If an exact symmetry is a transformation that gives you back what you started with an obvious suggestion about what it would mean for a symmetry to be approximate is that the transformation in question gives you back something similar. An object which has an inexact but approximate symmetry, on this line

[5]Kosso (2000) and Castellani (2003) contain discussions of approximate symmetry in physics.

[6]I do not claim that the approach taken here is the only way of discussing approximate symmetries. Castellani (2003) defines an approximate symmetry as one which is "valid under certain conditions". This suggests a different understanding of approximate symmetry akin to the notion of a ceteris paribus law, or non-universal generalisation. While I don't explore this possibility here, it may be that the literature on ceteris paribus laws can help to clarify the status of approximate symmetries in some contexts.

[7]See Belot (2013), Dasgupta (2015) and Caulton (2016) for recent work on these issues.

[8]We also often talk about symmetries of the Lagrangian in the QFT context, however symmetries of the classical Lagrangian can be broken by so-called anomalies (briefly described in Castellani, 2003). The partition function, therefore provides a more general characterisation of the symmetries of a QFT system.

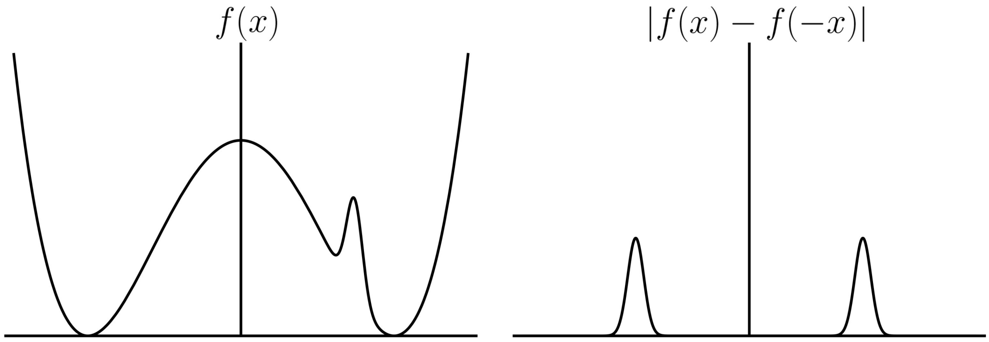

Figure 5.1: Function with approximate reflection symmetry.

of thought, will not be invariant under the corresponding transformation group; instead the action of the group will generate a set of objects which are 'close' to each other. To make this precise we need to spell out the intended sense of closeness. At the very least this will mean specifying a topology on the space of relevant object, but if we want to talk about the amount that a symmetry is broken in quantitative terms we will also need a metric. At this point things might seem to get worryingly messy. We are interested in the symmetries of many different kinds of objects, and there will likely be multiple ways of defining an appropriate topology in each case. It is unlikely then that this route will lead to a neat general definition of approximate symmetry.

This is not a conceptual disaster however. The existence of many possible topologies simply reflects the fact that there are many ways that two objects can be similar to each other; approximate notions are always infected with this kind of ambiguity. Furthermore, in many important contexts there are standard ways of defining a metric which will do the job for us here. If we are talking about the approximate symmetries of real functions, for instance, the natural measure to use is simply the absolute value of the difference between pre- and post-transformed functions. Take a double well potential function which has a $x \to -x$ reflection symmetry and add a small Gaussian perturbation at positive x (depicted in figure 5.1). This new function, $f(x)$, will no longer have an exact reflection symmetry, but the difference between the initial and reflected function, $|f(x) - f(-x)|$, will be small, especially for values of x far away from the asymmetric perturbation. In this way we can pin down the sense in which the function has an approximate reflection symmetry and even quantify the amount of symmetry breaking incurred by the perturbation. One lesson to be drawn from this example is that, while exact symmetries are necessarily exact everywhere, an inexact symmetry can be broken to different extents in different places. Just as one function can provide a good approximation to another in some regions of its domain but not in others, it is possible that a symmetry can hold to a good degree of approximation in one region but not in another. This local nature of the quality of an approximate symmetry will be important shortly.

We now have the resources in hand to clarify the sense in which cutoff QFT models may be said to be approximately covariant. Note, first of all, how continuum space-time symmetries are employed in the model building process in high energy physics. Typically Poincaré symmetry is invoked at the level of classical field theory when it comes to writing down a Lagrangian for the relevant interactions the theory is supposed to describe. Possible interaction terms which explicitly violate Poincaré covariance are discarded a priori. The Poincaré violations of cutoff QFT models obtained via these classical Lagrangians then entirely originate in the imposition of the cutoffs in the course of writing down a corresponding quantum system. Asking how a quantity defined on a cutoff QFT model behaves under Poincaré transformations then is tantamount to asking how it depends on the cutoffs.

As we saw in the previous chapter however, we have good reason to believe that many of the properties of cutoff QFT models are highly insensitive to the exact value of both the ultraviolet and infrared cutoff. To gloss over some subtleties discussed there, the basic picture to emerge from the renormalisation group analysis of QFT models was that, after quantities like correlation functions and S-matrix elements have been reparameterised in terms of renormalised couplings fixed via measurements at low energies, the residual dependence on the ultraviolet cutoff takes the form of powers of E/Λ, which are heavily suppressed at energies $E \ll \Lambda$. The imposition of an infrared cutoff will also give rise to violations of Poincaré symmetry, but again, there is strong evidence that correlation functions over length scales which are short relative to the infrared cutoff depend only very weakly on its precise value. In the context of lattice QFT models the contribution of finite size effects is expected to fall off at least as fast as $\exp(-EL)$, where L is the length of the lattice. As we saw, this means that correlation functions on scales intermediate between the ultraviolet and infrared cutoffs are very nearly independent of the details of how they are imposed.

Cutoff QFTs are not invariant under Poincaré transformations then—applying the Poincaré group to a lattice QFT generates a set of cutoff systems which differ with respect to their partition functions. But the above argument suggests that the correlation functions of this collection of systems at energy scales suitably far away from the cutoffs will be very similar indeed. The difference between correlation functions pre- and post-transformation due to the ultraviolet cutoff will be of the order E/Λ, for instance, and therefore vanishingly small at low energies. Using this difference to quantify how 'close' Poincaré transformed cutoff models are we can give a precise meaning to the claim that Poincaré symmetry holds approximately at low energies in these systems. Just as the quality of the reflection symmetry in the double well example depends on the value of x, the amount of Poincaré breaking incurred by the cutoffs depends on the energy scale at which we are probing the system. At the cutoff scale itself Poincaré transformations give rise to large changes in physical properties of the model, while at coarse-grained scales the system's behaviour will be essentially invariant under these transformations.[9]

What this means, in effect, is that special relativity will have the status of an effective, low energy, theory in a possible world exactly described by realistic cutoff QFTs.[10] We might say that Poincaré symmetry is 'emergent' in this context, in an innoculous, reduction compatible, sense of the term.[11] Furthermore, we can make an immediate connection with empirical tests of special relativity on the approach I have put forward here. The empirical effect of Poincaré breaking in experimentally observed quantities, like scattering cross sections, will also be supressed by inverse/exponential powers of the cutoffs and can therefore be expected to fall well within the relevant experimental error if the cutoffs scales are far away from the scales at which we are measuring these quantities. In this way, we can allay worries about cutoff models based on the lack of any detectable violation of Poincaré symmetry in extant experimental tests.

Formally at least, this is all highly analogous to the way that continuum symmetries emerge in condensed matter systems. At the scale of the crystal spacing rotational symmetry is badly broken in these systems but correlations between widely separated lattice sites can be approximately invariant under rotations giv-

[9]There is a different way of cashing out the sense in which a cutoff system has approximate continuum space-time symmetries at low energies which is often invoked in the lattice QFT literature. The so-called operator product expansion is used to express non-local combinations of lattice operators as series of local continuum operators. The standard lattice kinetic term on a Euclidean lattice, for instance, is expanded like:

$$\sum_{\mu}\Big[\frac{\phi_{x+a\hat{\mu}}-\phi_x}{a}\Big]^2 = (\partial_\mu\phi(x))^2 + C_1\phi(x)\partial_\mu^4\phi(x) + ...$$

Eliminating terms in the expansion which do not respect the discrete symmetries of a hypercubic lattice, it turns out that the terms beyond the familiar continuum kinetic term are non-renormalisable and consequently can be expected to diminish in importance at low energies. Consequently, the claim is that the effective Lagrangian governing the low energy degrees of freedom will exhibit standard continuum symmetries. This argument is discussed in Moore (2003) and Williams (2016). I take this approach to be complimentary to my discussion here, though how to spell out the precise sense in which continuum space-time symmetries are approximate in this context is a topic for further discussion.

[10]Exactly how one spells this out will, I suspect, depend on one's broader views about space-time theory. As a side note, on Brown's (2005) dynamical view of relativity theory, the story seems to be quite simple. For Brown Minkowski space-time is just an encoding of a symmetry of the dynamics, namely Poincaré symmetry. Consequently, if we can pin down the sense in which this symmetry is approximate at low energies, and I am suggesting that we can, we have already done much of the work needed to spell out the sense in which Minkowski space-time might be said to be emergent.

[11]As in the sense of emergence put forward in Butterfield (2011).

ing rise to isotropic bulk behaviour. There is an important difference between the representational status of the cutoffs in condensed matter and QFT however, which impacts on how the symmetry violations should be understood in the latter case. The molecules in a solid really do form a lattice structure, and there really is a threshold on the allowed momentum modes of the system. In this case then rotation symmetry breaking effects are physically real and can be detected by probing the system at sufficiently short length scales. In the case of cutoff QFTs however I have argued that the cutoff ought to be viewed as an idealisation, and we should not believe what the model says about the world at the cutoff scales. My claim that cutoff QFTs capture coarse-grained aspects of reality does not entail that the world is non-relativistic at the fundamental level then. Rather, since the Poincaré breaking effects originate in the imposition of the cutoffs they too should be viewed as unphysical idealisations, and the belief worthy content of the theory will respect Poincaré symmetry to the level of approximation which it is actually supported by current evidence.

Pressing questions remain about the story just sketched. The possibility of emergent Poincaré symmetry has recently been the subject of some controversy in quantum gravity research. In the literature on the causal set theory approach to quantum gravity, for instance, one often finds claims that a regular discrete structure to space-time, like a lattice, cannot recover approximate relativistic symmetries at low energies. This is taken to motivate the need to 'sprinkle' discrete space time points in a Lorentz invariant way in a fundamental quantum theory of space-time (Dowker et al, 2004). Renormalisation group studies of the effects of various kinds of high energy Poincaré violations on low energy phenomenology have also motivated various claims that seem to conflict with picture I have put forward here—see, for instance, Polchinski (2012), and references therein. Addressing these arguments is clearly an urgent task for defenders of the representational significance of cutoff QFTs, as prima facie they seem to suggest that these systems are not empirically adequate after all. There remains a great deal more to be said about the emergence of continuum space-time symmetries in discrete systems then.[12]

Still, I take it that the above discussion has shown that the notion of approximate Poincaré symmetry makes sense. We now have a programme at least for reconciling my previous claims about the representation success of cutoff QFTs with the empirical success of special relativity, as well as its role in model building in high energy physics. As I have already flagged however, Fraser's main objection to the Poincaré violations incurred by cutoff theories is an extra-empirical one.

[12]There are also potential connections here with broader debate surrounding the notion of emergent space-time in quantum gravity theories. See Huggett and Wüthrich (2013) for a recent discussion of these issues.

The next section turns to these theoretical concerns about the status of relativity in the QFT programme.

5.3 Relativity and the Quantum Field Theory Programme: Internal and External Questions

Fraser grants that cutoff QFTs are empirically adequate, indeed the claim that the choice between cutoff and continuum models is underdetermined by empirical evidence is central to her presentation of the formulation debate. For Fraser the violations of relativity engendered by the cutoffs are problematic not for empirical or pragmatic reasons but because they compromises the project of unifying quantum and relativity theory which is essential to the theoretical identity of QFT.

It is worth pausing to unpack this claim as unification is a notoriously nebulous notion whose philosophical significance is controversial.[13] After all, if the picture sketched in the previous section is correct, a cutoff formulation of the standard model may qualify as a unification of quantum theory and special relativity in a weak sense of the word. What physicists sometimes seem to mean when they say that two theories have been unified is simply that they can be recovered, perhaps in some limit, from a third more fundamental theory. This is the sense in which some ‘theories of everything’, such as the string theory programme, propose to unify the standard model and general relativity—these theories are supposed to emerge as effective theories from a multidimensional fundamental theory. If cutoff QFT models are approximately Poincaré covariant in the manner described above then, as I pointed out, special relativity should be an effective low energy theory in this sense. Fraser clearly has something stronger than this in mind. The objective of the unification project she focuses on, it seems, is to consistently combine quantum and relativistic principles at the fundamental level, and not merely recovering them as approximations. Haag theorem is presented as a major obstacle to achieving this and the Wightman and Haag-Kaiser axioms are held up as exemplars of the kind of theory that is sought. Though Fraser allows that extant axiomatic frameworks may need to be modified to accommodate realistic QFTs, the idea seems to be that an adequate formulation of QFT must marry quantum structures and dynamics with a commitment to fundamental Minkowksi

[13] Morrison (2007) argues that unification in physics is a diverse phenomenon which cannot be captured by general definitions. While other authors are less pessimistic than Morrison about the possibility of saying something general about unification, it is usually admitted that multiple senses of unification can be distinguished, which have potentially quite different roles and significances. Indeed, we find many distinct characterisations of unification in the literature on scientific explanation and the foundations of physics.

space-time and Poincaré covariance.[14]

Note however that the kind of unification we find in the axiomatic approach to QFT is still of a somewhat shallow kind. What these frameworks aim to do is put together quantum and relativistic axioms in a consistent way, but as Maudlin (1996) points out, a great deal more than mere consistency is needed for a deep physical unification. To use one of Maudlin's examples, general relativity enacts a genuinely physical unification of gravitational and inertial mass, which are only coincidentally connected in Newtonian gravitation theory, by showing that both arise from a common source, namely the metrical structure of space-time. Another example of a robustly physical unification is Einstein's light-quanta hypothesis which explained the success of Planck's black body radiation equation, the photoelectric effect, and various other phenomena, in terms of a single posit. Putting aside the question of whether they apply to realistic QFTs, the Wightman axioms do not seem to offer this kind of explanatory unification of quantum and relativistic physics; they codify rather than account for the fact that the world is very well described by Minkowski geometry and quantum theory.

I will come back to the significance of unification shortly, but for now the key point is that Fraser takes QFT to incorporate a commitment to fundamental Minkowski space-time structure. Wallace clearly does not accept this characterisation of the theoretical content of QFT however. For Wallace QFT is not in the business of describing the fundamental structure of space-time; this is the goal of the quantum gravity programme. QFTs are effective theories which describe the world at relatively large length scales and consequently the kind of emergent relativity discussed in the previous section is all that we can reasonably demand. Again, while Wallace endorses a realist stance towards cutoff QFTs he does not want to say that the non-relativistic effects engendered by the cutoff are physically real. Instead, we should simply remain neutral about whether Poincaré covariance holds down to arbitrarily small length scales as far as the QFT programme is concerned.

In sum then, Fraser and Wallace take quite different views of the status of special relativistic constraints within QFT. This seems to lead to a stalemate situation; it is not obvious what considerations we could bring to bear to determine which assumptions about space-time structure should be taken to be part of the theory. There are close parallels here, I think, with the philosophical debates involving multiple inequivalent formulations of a theory we discussed in §2.4. One case we considered there, for instance, was the controversy surrounding the consistency of classical electrodynamics. Frisch (2005) and his detractors present their dispute

[14]Fraser (2009, 557) comments that there is "a certain amount of latitude in deciding what counts as a relativistic principle and what counts as a quantum principle", but it seems clear that, for her, core relativistic properties like Poincaré covariance are non-negotiable.

as one about whether classical electrodynamics is consistent, but, on closer inspection, philosophers on each side of the debate take different sets of theoretical assumptions to constitute the theory. As in this case, once we recognise that Fraser and Wallace are delineating the content of the theory in different ways the debate threatens to degenerate into a purely semantic dispute about which version is worthy of the name QFT.

I argued in chapter 2 that the right way to respond to this situation is to move the dialectical emphasis away from the question of what the 'correct' version of the theory is. Theory names are just labels at the end of the day and which collection of models and theoretical assumptions we decide to call QFT has no philosophical interest on its own; a particular delimitation of theoretical content only becomes interesting via its relation to broader foundational issues. Rather than thinking of the theoretical structures under consideration as being set in stone before philosophical engagement begins we should instead look to justify focusing on one formulation of a theory over another within the context of a particular philosophical project. Making this move in the present debate leads to a different perspective on the issue regarding the status of relativity in QFT. The question now is not whether fundamental Poincaré covariance is part of the definition of QFT but whether we ought to be demanding this property when addressing the disparate philosophical issues raised by the QFT programme. This opens up the possibility that different approaches to QFT, which incorporate contrary commitments about space-time structure, might be the appropriate starting points for different lines of philosophical enquiry. I will now argue that this is in fact the case: it is appropriate to assume exact Poincaré symmetry when addressing some salient internal questions raised by the QFT programme but there is no case for demanding it when it comes to the question of what QFT tells us about the world.

On the one hand, there clearly are contexts in which we do want to demand fundamental Poincaré covariance. A central issue raised by the QFT programme is whether, and to what extent, quantum theory and special relativity can be consistently combined. As I sketched in §1.2.1, bringing these theories together was one of the key motivations driving the development of QFT in its early days, and while I don't think we should define QFT as a fundamental unification of quantum and relativistic principles, as Fraser does, this project is surely a legitimate strand of the QFT programme. Furthermore, the relationship between quantum and relativity theory has long been a topic of scrutiny and debate amongst philosophy of physics, especially in the wake of Bell's theorem. If we are interested in the question of how quantum phenomena can be consistent with fundamental relativity there is an obvious sense in which cutoff formulations of the standard model are besides the point. Not only are cutoff QFT models not covariant or local they don't tell us anything about whether a theory which actually has these properties is possible. By contrast, the axiomatic and constructive approaches to QFT provide a natural starting point for addressing this issue. Axiomatic formulations of QFT on continuum space-times give us a powerful framework for exploring questions about quantum theories on Minkowski space-time in the abstract, while the constructive field theory programme's raison d'être is to establish if, and when, interacting QFTs on Minkowski space-time exist.

The question of how quantum and relativity relate to one another is ultimately an internal one however, in the sense introduced in §2.4—it is a question about the way our theories are rather than the way the world is. As a point of comparison, consider the debate between Norton (1995, 2002) and Malament (1995) over the question of whether Newtonian cosmology is a consistent theory. Norton argues that it isn't, while Malament claims that using modern geometric formulations of Newtonian gravitation we can make sense of it as a consistent cosmological model. This question is interesting for a number of reasons: it illuminates the structure of gravitational theory and may be a useful case study for understanding how scientists can subscribe to an inconsistent set of assumptions without falling foul of the principle of explosion.[15] But answering it does not directly tell us anything about the world. If it is possible to formulate a consistent Newtonian cosmological model we will obviously not be in a position to believe any of its content—we now know that many of the core assumptions of pre-20th century cosmology are radically false. Similarly, answering the question of whether interacting QFTs exist on continuum Minkowksi space-time does not, on its own at least, amount a discovery about the way the world is. How such a theory relates to the actual world is a separate question, and I have urged in this thesis that it is often a highly non-trivial one.

As far as I am concerned, investigating the internal relations between quantum theory and special relativity is interesting and important in its own right. As I have already stressed, philosophers of physics often remain at the semantic level in their analyses of physical theories and busy themselves with questions about their content. This kind of enquiry improves our understanding of the theoretical landscape of modern physics and can thus be hoped to indirectly contribute to our knowledge of the world, either by clarifying how a theory's epistemically warranted claims should be spelled out, or by paving the way for future theoretical developments. I take it that work on the philosophy of quantum gravity, for instance, is typically conducted in this spirit. Philosophical investigations into string theory, canonical and loop quantum gravity are not predicated on the assumption that these frameworks accurately describe the world in their current

[15]See Vickers (2013) chapter 5 for a discussion of why the inconsistencies in Newtonian cosmology were not noticed until the end of the 19th century.

state; the hope is rather that elucidating the conceptual issues raised by these approaches will ease progress in the quantum gravity project. Clarifying the relationship between quantum and relativistic principles in an internal mode can be viewed in a similar light. In my view, much of the philosophical work on the axiomatic approach to QFT that has been conducted in recent years can be understood as contributing to this kind of project, a prime example being Clifton and Halvorson's (2000, 2001) work on how distinctively quantum phenomena like entanglement can be made sense of the fully relativistic context of axiomatic QFT.

There are, in fact, some indications in Fraser's discussion of the formalism problem that her vision of the philosophy of QFT is an internal one. For one thing, she consistently stresses that her focus is the question of what the world would be like if QFT were true—which for her, of course, means what the world would be like if it were quantum and perfectly relativistic. This amounts to an endorsement of the adequacy of the standard account of interpretation for her purposes, but I argued in chapter 2 that this definition of interpretation is a purely semantic one and is not equipped to deal with the epistemic question of how a theory's content relates to the actual world. She also dismisses quantum gravity research as irrelevant to her project (Fraser, 2009, 552). This makes perfect sense if we are addressing internal issues, but when it comes to assessing how QFT models relate to the world the difficulties associated with bringing gravity into the quantum fold are surely relevant. (By way of analogy, modern cosmology does not impact on the question of whether a consistent Newtonian cosmology can be formulated but is clearly crucial for assessing the epistemic status such a theory would have.) If Fraser's claim is that the axiomatic/constructive approach to QFT is the right starting point for addressing the internal quantum-relativity relationship within the QFT programme she makes a good case for it.[16]

This thesis is wholly compatible with the view of the representational success of cutoff QFT models put forward in the previous chapter however. There is no tension in advocating the study of continuum QFT systems for insights into quantum-relativity relations while maintaining that we have no reason to believe that they exactly describe the world. The external question of how QFT models relate to the world and what we ought to believe about reality given the empirical successes of the QFT programme is, in principle at least, a separate issue. Should we be positing fundamental Poincaré symmetry as a theoretical constraint in this context then? The answer, I think, is no.

Why might we think that fundamental Poincaré covariance is a warranted as-

[16]Though 'defining' QFT as a unification of these theories does not play any role in motivating this conclusion.

sumption when we come to the epistemic question of what the QFT programme tells us about reality? It should be clear by now that the idea that QFTs are Poincaré covariant by definition does not count for much here, and we saw in the previous section that the extant empirical support for relativity theory does not establish that Poincaré symmetry is actually exact. While Fraser sometimes seems to style her philosophical project as an internal issue, we also find an argument for viewing continuum QFT models and the axiomatic approach as epistemically superior in her discussion, which goes something like this.[17] First of all, the claim is made that the relationship between cutoff and continuum formulations of QFT is a case of underdetermination of theory by evidence. Cutoff and continuum versions of a QFT say different things about the world but they are equally well supported by empirical evidence.[18] Fraser suggests that we should not draw an anti-realist moral however because we can appeal to a familiar realist strategy to break the underdetermination. Continuum QFT systems have extra-empirical theoretical virtues which give us reason to believe their claims over those of the corresponding cutoff theory. Namely, they are more unified; they enact a principled unification of quantum and relativity theory while cutoff models violate fundamental Poincaré covariance in an ad hoc way.

This argument mitigates against the view of the epistemic status of QFT models put forward in the previous chapter. I argued there that if we had continuum formulations of realistic QFTs their claims about the large and small scale structure of the world would not be belief worthy because they don't make a difference to the predictive successes of the QFT programme. The suggestion here, however, is that there are extra-empirical reasons for believing a continuum QFT is representationally accurate down to arbitrarily small length scales; it is only by accepting these commitments that we adequately unify quantum and relativistic principles at the fundamental level.

How should we evaluate to this claim? Wallace's reply has been to deny the premise that there is an underdetermination scenario here in the first place (Wallace, 2011, 121-122). As we have seen, we do not currently have continuum formulations of realistic QFTs. We aren't in a position of choosing between the claims of a cutoff and continuum version of QED then, only the former is on the table. And we surely cannot believe in a theory we have not yet formulated no matter how theoretically virtuous it supposed to be! Fraser's response to this obvious worry is to emphasise the incompleteness of the constructive field theory programme. The suggestion seems to be that, though we aren't in a position

[17] I say a good deal more about extra-empirical virtues here than we find in Fraser (2009). My motivation has been to state the argument in what I take to be its most plausible and persuasive form rather than to offer a faithful exegesis of what Fraser actually says.

[18] It is not obvious that this is the case given that the cutoffs do lead to in principle observable effects—I will have more to say about this in §7.4.

to believe in continuum QFT models at the moment, the argument from superior theoretical virtues still motivates a positive epistemic attitude towards the constructive field theory programme as a work in progress and we should be optimistic that further work on this problem will produce belief worthy theories at some point in the future.[19]

Whether you find this way of responding to the unsubstantiated character of Fraser's underdetermination scenario compelling or not, there are more basic problems with this line of thought in my view. As I made clear in the previous chapter, I am not convinced that the question of whether realistic continuum models exist is as crucial as Wallace and Fraser's discussion suggests, and I think this perspective can also be defended here. In order to separate out the relevant issues it will be useful to engage in a couple of thought experiments. First of all, imagine that we had a continuum formulation of an interacting QFT in four space-time dimensions. Further suppose that this theory is well supported by the available empirical evidence and there is no reason to think that it was not a complete description of reality—there are no gravitational interactions that stubbornly resist quantisation, or any other indications that the theory will break down in some, as yet unprobed, domain in this fantasy.

Would we be justified in believing the claims this theory makes about the large and small scale structure of the world over those of an empirically indistinguishable cutoff alternative? It seems to me that this is a more subtle issue than Fraser and Wallace give it credit for. While the idea that extra-empirical virtues can provide some kind of evidential support is relatively mainstream amongst scientific realists there is very little agreement about how the details of this story are supposed to play out in the context of confirmation theory.[20] Simply saying that unification is a theoretical virtue will not do here, we have the right to demand a more detailed story about why the unification enacted by the continuum theory makes it more worthy of our belief than its cutoff competitor. This challenge is exacerbated by the fact that, as I mentioned earlier, unification is an ambiguous notion. A common way of cashing out how extra-empirical virtues play a role in confirming the claims of a theory is that they contribute to its explanatory

[19]There is, perhaps, a connection here with the axiological characterisation of scientific realism favoured by van Fraassen, according to which the core doctrine is that science aims at truth, rather than that our current theories are truthlike (van Fraassen, 1980, 6-7). A realist stance towards the constructive field theory programme could be construed along these lines as a commitment to it being a path to true theories while admitting that, as of yet, it has not achieved this goal. While I think this makes sense as an epistemic stance it is much less clear to me that Fraser's arguments support it. Even if we grant that a continuum version of QED would be epistemically superior to a cutoff versions this does not, in itself, provide us with any reason for thinking that the continuum theory actually exists—Wallace (2011, 121-122) makes the same point.

[20]Specifically, there is disagreement over how the epistemic relevance of unification should be understood. For some recent proposals see Myrvold (2003), Sober (2003) and Lange (2004).

power, so that we can justify believing its claims over its rivals via an inference to the best explanation. As I mentioned above however, it does not seem that the kind of unification we find in axiomatic formulations of QFT is of an explanatory variety. It seems plausible that Einstein's light-quantum hypothesis is supported by the fact that it neatly accounts for a number of disparate phenomena in terms of a single hypothesis, but it is not clear that the same could be said of the Wightman axioms even if they were empirically adequate. At the very least then, more work is needed on the part of Fraser sympathisers to shore up the case that the extra-empirical virtues of continuum QFTs would support their claims about the arbitrarily large and small.[21]

Ultimately though, whatever we end up concluding about this counterfactual scenario, there are more staightfoward reasons why this appeal to extra-empirical virtues does not go through when we move closer to the epistemic situation we actually find ourselves in. This time imagine that we had continuum formulations of QED, QCD and the standard model. Would the extra-empirical virtues of these theories give us grounds to believe what they say about arbitrarily large and small length scales? Clearly not. In fact, we know with certainty that the standard model does not furnish us with an accurate description of the world at all energy scales. For one thing, it does not incorporate gravitational interactions, but there are also a plethora of other issues which lead many physicists to think that it will break down long before the Planck scale, the naturalness problem with the Higgs sector being perhaps the most significant. Furthermore, the quantum gravity project provides us with compelling reason to think that the QFT programme itself is not up to the task of representing the fundamental structure of reality. On all of the major approaches to quantum gravity Minkowski space-time ultimately ends up being a sort of idealisation—eventually space-time must become dynamical and quantum in nature. It is not just that the large and small scale structures of QFT models are not supported by empirical evidence then, they are directly undermined by broader theoretical considerations. And these external factors surely trump any arguments for trusting continuum QFTs at all length scales based on extra-empirical virtues.

[21] Another point here is that it is not at all clear that demanding exact Poincaré symmetry eliminates the underdetermination problem here at all. As Fraser (2011) points out, and I will discuss further in §7.4, the renormalisation group framework points to a more general kind of underdetermination in high energy physics. It suggests that very different high energy theories can have indistinguishable low energy physics. But this generalised underdetermination challenge would apply just as strongly to continuum formulations of realistic QFTs if we had them. It is possible that multiple continuum QFTs, which do an equally good job of unifying quantum theory and relativity by Fraser's lights, flow to the same effective theory under renormalisation group transformations. One can make the case then that the renormalisation group undermines any reason to trust our theorys' claims about energy scales beyond our experimental reach independently of the considerations discussed in the next paragraph.

To sum up, I have suggested that there is an ambiguity in Fraser's advocacy of the philosophical significance of the axiomatic approach to QFT. If we understand her project to be an investigation of internal issues raised by the QFT programme, and specifically into the relationship between quantum theory and relativity, then, to my mind, her arguments for focusing on the axiomatic/constructive approach go through. I claimed, however, that this is perfectly compatible with the claims about the representational status of QFT models put forward in the previous chapter. When it comes to what we ought to believe about the world given the empirical successes of QFT programme I have argued that we have no epistemic warrant to demand fundamental Poincaré covariance. Exact Poincaré symmetry is not established by extant empirical tests and even if we did have grounds to believe that Poincaré covariance is fundamental this would not mean that we should trust what a formulation of the standard model on Minkowski space-time says about large and small scale structure of the world. The master argument against appeals to the extra-empirical virtues of an exactly Poincaré covariant QFT model is that we know that current QFTs cannot be exact descriptions of reality at all scales because of their rivalry with general relativity. My claim that QFT models, both cutoff and continuum, are, at best, coarse-grained representations of reality can be defended against objections based on tensions with relativity theory then.

One upshot of the discussion of this chapter is that we have managed to reconcile at least some of what Fraser and Wallace say about the QFT programme. I have argued that we can accept that continuum QFT systems are worthy of philosophical study while maintaining that cutoff models capture all that we have any right to believe about the world. I will be developing this idea further in later chapters by expanding on the ways that the investigation of the axiomatic approach might be justified without taking the axioms themselves to be epistemically justified.

Chapter 6

Spontaneous Symmetry Breaking and the Limits of Infinite Degrees of Freedom

6.1 An Explanatory Challenge

I have been using the term cutoff QFT model, by default, to denote a system which lacks both arbitrarily high and arbitrarily low momentum modes—i.e. one which has both an ultraviolet and infrared cutoff. An important feature of this kind of doubly cutoff QFT system is that it has a finite number of degrees of freedom. My claim that these systems capture all that is worth believing in high energy physics is somewhat heterodox as many philosophers of physics have viewed the presence of infinitely many degrees of freedom as being the distinctive feature of QFT models which makes them interesting (and challenging) from a foundational perspective. Quantum systems with infinite degrees of freedom admit unitarily inequivalent Hilbert space representations and a great deal of work in recent years has focused on the implications of this fact for the interpretation of continuum QFT systems.[1] Furthermore, we find claims in this literature which appear to be in direct conflict with the view of cutoff models set out in chapter 4. Most notably, Laura Ruetsche (2003, 2006, 2011) has argued that the unitarily inequivalent representations afforded by the limit of infinite degrees of freedom play an indispensable role in explaining important physical phenomena like phase transitions and spontaneous symmetry breaking (SSB) in quantum theory, and consequently need to be taken physically seriously. This suggests a new explanatory challenge: one might admit that cutoff QFTs are empirically successful while maintaining that they lack the explanatory resources needed to really account for important features of the physical world.

[1] Fraser (2008) and Baker (2009) both make crucial use of the existence of unitarily inequivalent representations in their discussions of the ontological status of particles and fields in QFT, for instance.

Fraser (2009, 560) pushes this sort of objection in her debate with Wallace:

> ...[T]he 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.

The implication being that doubly cutoff QFTs cannot accommodate an apparently crucial phenomenon for high energy physics, namely SSB.[2]

Wallace has tended to reply to this line of attack by distinguishing between the ultraviolet and infrared cutoffs and advocating a different stance to the latter:

> 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. (Wallace, 2011, 119)

What Wallace seems to be saying here is that the QFT systems he advocates a realist reading of are those with an ultraviolet cutoff but no infrared cutoff—e.g. a QFT defined on an infinite volume lattice. Since these systems do have infinite degrees of freedom the claim that they lack the relevant representational resources of continuum QFT models misses its mark.

There are some tensions in this response however. Wallace (2006, 55) dismisses unitarily inequivalent representations which appear due to the presence of infinitely many degrees of freedom on arbitrarily short length scales as unphysical for two main reasons: i) there are mathematical difficulties associated with defining QFT models on a continuum, and ii) the renormalisation group tells us that the properties associated with the very small scale structure of a QFT model do not affect its large scale physics. Both of these points also seem to apply to the infrared cutoff however (and I have de-emphasised the significance of the former issue). Taking the infinite volume limit of QFT models is by no means mathematically trivial. We know that there are realistic physical systems for which the closely related thermodynamic limit does not exist, and proving that it exists in the case of particular models is often a subtle mathematical problem. In the case of QCD, in particular, it is the task of bringing the infrared, rather than

[2]Toader (2016) also explicitly raises this objection against Wallace.

ultraviolet, physics under mathematical control which is the main source of difficulty from the constructive field theory perspective.[3] Furthermore, as I pointed out in §4.5, while our knowledge here is much less systematic than in the case of the ultraviolet cutoff, we have good reason to believe that physics at finite length scales is very weakly affected by the imposition of an infrared cutoff. Prima facie then, there is a great deal of symmetry between the status of the ultraviolet and infrared cutoff.

Furthermore, the discussion of the unitarily inequivalent representations afforded by the infinite volume limit found in Wallace (2006) does not address the kind of explanationist arguments put forward by Ruetsche. Wallace focuses on the question of how we can discover which representation we are in, as it were, which he sees as a challenge to scientific realism. He argues that the representational ambiguity associated with infrared unitarily inequivalent representations simply codifies our ignorance of global facts about the universe, such as the total mass and charge. But Ruetsche's claim is that we actually need to take multiple unitarily inequivalent representations seriously in order to furnish adequate scientific explanations of phenomena as mundane as the boiling of a kettle, and this arguably poses a challenge to a realist view of QFT which cuts across Wallace's discussion.[4]

In this chapter I explore a more direct response to the explanatory challenge posed by the finiteness of doubly cutoff QFTs. I suggest that the case has not in fact been made that the novel properties afforded by the limit of infinite degrees of freedom need to be taken to represent features of the world. In effect, I will simply be doubling down on my previous claim that cutoff QFT models can represent all of the features of the world which we have any reason to believe in on the basis of current high energy physics.

I situate my discussion within a broader debate in the recent philosophy of physics literature about the explanatory and representational role of the limit of infinite degrees of freedom in both classical and quantum theory, and especially in statistical physics. In statistical mechanics, phase transitions, such as the boiling of water, are associated with non-analyticities in the free energy which only occur in the infinite volume, or in this context thermodynamic, limit. Though this limit is clearly an idealisation in this case, Batterman (2005) claims that the

[3]As we saw in §4.3 all indications are that the ultraviolet behaviour of QCD is well behaved. It's worth pointing out here that the Clay institute quantum Yang-Mills problem explicitly refers to the need to demonstrate the existence of a mass gap and confinement (Jaffe and Witten, 2006). These are infrared issues which indicates that the infrared aspects of the problem are taken to be the most challenging and interesting for mathematical physicists.

[4]See, in particular, Ruetsche (2011) chapter 15 for a discussion of the connection between her rejection of 'pristine' interpretations of quantum theories with infinite degrees of freedom and the scientific realism debate.

non-analyticities that occur in the thermodynamic limit are needed to successfully represent and explain phase transitions, a position which is closely allied to Ruetsche's claims about the limit of infinite degrees of freedom in quantum theory. Batterman's analysis of the thermodynamic limit has been contested by Butterfield (2011), Norton (2011) and Callender and Menon (2013) however. According to these authors, the thermodynamic limit is predictively and explanatorily successful in phase transition theory because of features infinite systems share with large finite systems—specifically, the non-analytic functions found in the limit are said to provide a good approximation to the values of macroscopic quantities in realistic, finite volume, models. Following this line of thought I suggest that it is not clear that resources afforded by unitarily inequivalent representations really are explanatorily indispensable in high energy physics in an ontologically committing sense.

Showing that the limit of infinite degrees of freedom is not representationally indispensable *in general* is a tall order, given the range of uses to which it is put. I focus on the case of SSB for two reasons: first, SSB raises new challenges for the kind of deflationary programme set out by Butterfield and others, and second, because the notion of SSB plays a prominent role in high energy physics and is the case which Fraser and Ruetsche focuses on in this context.[5] I start in §6.2 by setting out the role that novel properties of infinite systems play in theoretical accounts of SSB in both classical and quantum physics. §6.3 discusses the prospects of deflating the physical status of these properties with reference to the recent debate surrounding phase transitions. In §6.4 I argue that the deflationary reading of the limit of infinite degrees of freedom can be extended to the case of SSB. I conclude in §6.5 by sketching some avenues for further investigation into the epistemic and representational status of the limit of infinite degrees of freedom in the QFT programme.

6.2 Modelling Spontaneous Symmetry Breaking

What is SSB? A great deal is going to hang on this question so we need to be somewhat careful here. Before rushing into formal definitions let's start with a qualitative description of the kind of worldly phenomena we are trying to account for. The archetype is ferromagnetism. The interactions between the magnetic moments associated with electron spins inside a ferromagnet make it energetically favourable for neighbouring spins to align together. These interactions have a

[5]In fact, phase transitions are also dealt with in the QFT programme, specifically in the QFT at finite temperatures formalism and in particle physics cosmology, but this is a story for another time.

rotation symmetry and at high temperatures the system behaves in a spherically symmetric way: spins are uncorrelated and the magnetisation is effectively zero. If the temperature is lowered to some critical value however, spins suddenly line up in a particular direction, giving rise to a net magnetisation. Moreover, the spins remain frozen into this aligned state, unless an external magnetic field is applied or the temperature is raised above the critical temperature once more.

It will be useful in what follows to distinguish two aspects of the theoretical notion of SSB. On the one hand, there is the 'symmetry breaking' part. Informally speaking, the system displays 'stable' states which are non-invariant under some transformation group. In the case of the ferromagnet, the aligned spin states which manifest at low temperatures are clearly not rotationally invariant. On the other, there is the 'spontaneous' part. Naively this might be taken to mean that the system has, in some sense, an equal chance of choosing a number of asymmetric configurations; the spins inside the ferromagnet can align in any direction at the critical temperature.[6] In practice the sense in which the symmetry is broken spontaneously is often spelled out via a contrast with so-called explicit symmetry breaking, in which asymmetric behaviour is produced by asymmetries in the underlying interactions at play.[7] Applying an external magnetic field to the ferromagnet produces aligned states, for instance, but in this case the asymmetric behaviour arises due to violations of rotation symmetry in the relevant dynamical equations. Giving a theoretical characterisation of SSB comes down to spelling out more precisely what the relevant stable asymmetric states are and in what sense the asymmetric behaviour is produced spontaneously. One theme of this chapter will be that there is not a single way of doing this. As we shall see however, the novel properties of infinite systems play an important role in standard theoretical accounts of SSB in both classical and quantum systems. I start by discussing how SSB behaviour is typically described in classical theory, before moving on to the quantum case and the status of SSB in the QFT programme.

6.2.1 Classical Systems

One often finds SSB defined as a situation in which a system (classical or quantum) has multiple ground states which are mapped into each other under the action of an exact symmetry group of the dynamics. Why? The thought is that a system should eventually collapse into (or, at least, close to) a ground state,

[6] This idea becomes problematic in the case of gauge symmetries, but as I touch on below, many philosophers view the notion of spontaneously broken gauge symmetry as intrinsically problematic.

[7] See Castellani (2003) for a discussion of the distinction between spontaneous and explicit symmetry breaking.

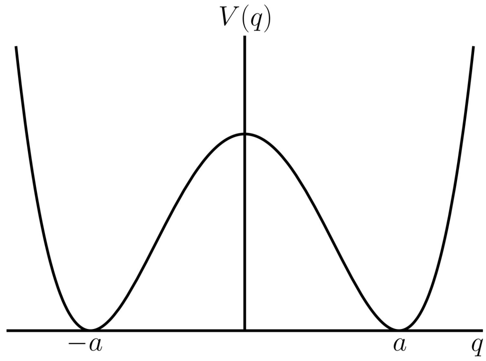

Figure 6.1: A double well potential.

so these are the appropriate states to identify as the locus of stable symmetry breaking behaviour. As I mentioned, the sense in which the symmetry is produced spontaneously is typically cashed out via a contrast with explicit symmetry breaking in which there are non-invariant terms in the system's dynamical equations. If we assume that there are no such terms, and the symmetry in question is exact, the only way to have a non-invariant ground state is if it is non-unique—a unique ground state is necessarily mapped into itself under all of the exact symmetries of the dynamics.

A little thought shows that the existence of multiple ground states is not really sufficient to ensure that SSB type behaviour occurs however (later I will argue that, in quantum systems, it is not necessary either). A simple classical system which motivates this conclusion is a particle on a two-dimensional plane. Each point on the plane is a ground state, but the system does not display the kind of stable asymmetric configurations that are characteristic of SSB (Strocchi, 2008, 4). In order for interesting asymmetric behaviour to emerge there needs be an energy cost associated with moving from one ground state to another. A toy model which is often discussed in presentations of SSB is a classical system with a double well potential (figure 6.1). In this system there are two ground states which are not parity invariant and are separated by a potential barrier. As long as the system does not have sufficient energy to traverse the potential barrier it can be expected to 'choose' one side of the potential.[8] Finite classical systems can display the characteristics of SSB phenomena then, but the sense in which we have stable asymmetric behaviour here is inevitably a matter of degree; there is no natural answer to the question of how high the potential barrier between degenerate ground states must be for interesting symmetry breaking behaviour to arise.

Furthermore, there are contexts in which the existence of a potential barrier of any

[8]For a discussion of SSB in this kind of simple mechanical model see Lui (2003) and Earman (2004).

height does not seem to capture the appropriate notion of stability. The classical Ising model in one spacial dimension has two asymmetric ground states separated by a potential barrier but it is usually said *not* to display SSB. Explaining why takes us into the account of SSB found in classical statistical mechanics which rests on the limit of infinite volume. As we shall see, infinite systems yield a much more clear cut characterisation of SSB than can be provided in finite systems.

The Ising model is the quintessential toy model of a ferromagnet. In this model, one takes a d dimensional (cubic) lattice, with N sites and volume V, and associates with each site i a variable $s_i = \pm 1$. The values of s_i can be interpreted as a naive representation of the orientation of spins inside a ferromagnetic crystal: $s_i = 1$ corresponding to a spin aligning 'up'; $s_i = -1$ corresponding to a spin aligning 'down'. In keeping with this picture, we can identify the quantity, $m = \sum_i s_i/V$, with the (volume averaged) magnetisation of the system.[9] The basic Hamiltonian consists of a single term representing short ranged interactions that make it energetically favourable for spins to align parallel with one another:

$$H = -J \sum_{i,j} s_i s_j, \tag{6.1}$$

where J is a positive constant parameterising the interaction strength and the sum is taken over 'nearest neighbour' spin pairs.

This Hamiltonian is invariant under a 'spin flip' transformation $s_i \to -s_i$ and clearly also has two ground states which are mapped into each other under this transformation: one in which all spins point up, and one in which all point down. Naively then, we might expect the kind of SSB behaviour found in the ferromagnet—spins should align either all up or down breaking the spin flip symmetry. This turns out not to happen in the case of one spacial dimension however. Though multiple ground states exist they are not stable under thermal excitations: any non-zero temperature gives rise to purely symmetric behaviour.

In statistical mechanics the natural candidate for identifying stable asymmetric states are not grounds states but equilibrium macrostates. The equilibrium state of the classical Ising model is represented by a probability distribution over its microstates, $S_V = (s_1, ..., s_N)$; in particular, the canonical ensemble takes the form:

$$P(S_V) = \frac{\mathrm{e}^{-\beta H(S_V)}}{Z}, \tag{6.2}$$

[9] A note on terminology here. In statistical mechanics the magnetisation is often identified with the expectation value of the quantity I have called m. Part of the reason for this is that $\langle m \rangle$ and m effectively coincide in the thermodynamic limit, as fluctuations vanish. It will be crucial to my discussion in §6.4 to keep these quantities distinct however, and I refer to the latter quantity as the magnetisation here.

where $\beta = 1/k_B T$ is the inverse temperature and Z is the partition function (a close cousin of the QFT partition function discussed in chapter 1). This distribution is clearly unique and, in fact, the canonical ensemble assigns a unique equilibrium state to any system with finite degrees of freedom. For the same reason discussed above in the case of the ground state, this means that it must respect all of the symmetries of the dynamics.

Things become more interesting in the thermodynamic limit however. Taking $N, V \to \infty$, with N/V held constant, generates an infinite lattice with a divergent total energy so (2) is no longer a well defined probability distribution. In the classical context the canonical ensemble is typically generalised to such infinite systems by means of the Gibbs measure formalism.[10] Roughly speaking, a Gibbs measure is a probability measure over a system's microstates which behaves locally like the canonical ensemble—in the case of the Ising model, any finite sub-region of the lattice is stipulated to have a probability distribution defined over its associated microstates which takes the form of (2). Stated precisely, this definition agrees with the canonical ensemble for systems with finite degrees of freedom, but infinite systems also admit Gibbs measures in this sense.

What happens in the thermodynamic limit of the Ising model depends on the dimensionality. The d=1 case was solved exactly by Ising himself and was found to have a single symmetric Gibbs measure at all temperatures—his verdict was that no ferromagnetic behaviour occurs. For two dimensions and higher however, multiple Gibbs measures appear. An elegant argument (originally due to Peierls) illustrates how this is possible.[11] Strictly speaking, the Hamiltonian of a finite Ising model should include an additional term representing the effects of external spins at the boundary of the lattice. If we take all of the spins at this boundary to be 'up' it can be shown that, below some critical temperature, the expectation value of s_i at the centre of the lattice is positive no matter how large V is. This means that when we take $V \to \infty$ the boundary goes to spacial infinity while the expected magnetisation of the system converges to a positive value. Running the same argument with 'down' spins at the boundary, we find that there are two distinct Gibbs measures, P_+ and P_-, on the infinite lattice which assign positive and negative values to the magnetisation respectively. Normalised linear combinations of these measures are also Gibbs measures of the system. If we take the thermodynamic limit of a finite lattice with boundary conditions that are invariant under spin flips,[12] for instance, the Gibbs measure converges

[10]Presentations of the Gibbs measure formalism, and its application to SSB and phase transitions, can be found in Georgii (1988) and Lebowitz (1999).

[11]See Kindermann and Snell (1980), and references therein, for the details. Note that rigorous results are, for the most part, not available in three or more dimensions.

[12]In practice this typically means free or periodic boundary conditions, see Kinderman and Snell (1980, 34-35).

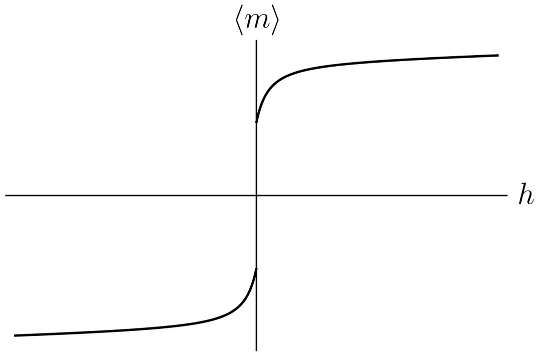

Figure 6.2: Expected magnetisation of the two dimensional classical Ising model in the thermodynamic limit.

to $^1\!/\!_2 P_+ + {^1\!/\!_2} P_-$. Only P_+ and P_- satisfy the properties required of genuine equilibrium states however: in general, systems with infinite degrees of freedom can have a convex set of Gibbs measures whose extremal elements correspond to macroscopically distinct equilibrium states of the system.

This leads to the standard definition of SSB commonly employed in statistical mechanics: SSB occurs if there are multiple equilibrium macrostates in the thermodynamic limit which are mapped into each other under an exact symmetry of the dynamics. Interpretive puzzles immediately arise with this approach to modelling SSB. Real ferromagnets are clearly finite in volume so how do these infinite systems relate to what is going on in concrete systems? As a piece of first order science however its success is undeniable.[13] The asymmetric equilibrium states which appear in the thermodynamic limit give an accurate quantitive description of the macroscopic properties of many concrete systems. And this characterisation of SSB is intimately tied up with a powerful framework for describing phase transitions in statistical mechanics. As has already been mentioned, phase transitions are associated with non-analyticities in macroscopic observables, and ultimately thermodynamic potentials, which only occur in the thermodynamic limit. The expected magnetisation of the infinite volume classical Ising model, for instance, is a discontinuous function of the external magnetic field, h, below the critical temperature, corresponding to the first order phase transition observed in real ferromagnets when the direction of an external field is reversed (figure 6.2). The appearance of these non-analyticities is closely related to the non-uniqueness of the system's equilibrium state at the phase boundary. In the case of the Ising model, it is the existence of the P_+ and P_- equilibrium states that allows the expected magnetisation to change discontinuously at $h = 0$. The interesting question for the philosopher of science, I think, is not whether the infinite systems obtained in the thermodynamic limit are predictive and explanatory

[13]Its worth mentioning, however, that this account of SSB may have important limitations as well, in so far as the thermodynamic limit does not exist for all systems.

in this context, but why.

One lesson we can draw from this discussion of classical systems is that it is possible to characterise SSB in much more categorical terms in the limit of infinite degrees of freedom. In finite mechanical models there is always a certain amount of vagueness in the sense in which we have stable asymmetric behaviour. In the limit of infinite degrees of freedom this ambiguity disappears. The reason is that the solution space of classical systems with infinite degrees of freedom can split into multiple isolated 'islands' which cannot be traversed without expending an infinite amount of energy. It is the appearance of these islands which opens up the possibility of degenerate equilibrium states in the thermodynamic limit of classical statistical mechanics, but the same kind of structures are also found in classical field theories, where degenerate ground states become completely dynamically isolated from each other.[14] As we shall see, there are parallels here with what happens in infinite quantum systems. There are new features of the quantum case however, which have often been thought to provide further grounds to attach special physical significance to the limit of infinite degrees of freedom.

6.2.2 Quantum Systems

The most dramatic difference between the quantum and classical case is that, while classical systems with finite degrees of freedom can have multiple ground states, the corresponding quantum systems typically do not. As we have seen, the classical Ising model has up and down aligned ground states. But the quantum Ising model (with a transverse magnetic field)[15] has a single ground state, which is a superposition of the states ψ_+ and ψ_- corresponding to all spins aligned up and down respectively before the thermodynamic limit is taken. More generally, a unique quantum ground state can typically be constructed from a superposition of quantum states associated with classical ground states if the system has finite degrees of freedom. If degenerate ground states are required for SSB behaviour to occur then it appears that symmetries cannot be spontaneously broken in finite quantum systems.

In the limit of infinite degrees of freedom radically new structures emerge however.

[14]See Strocchi (2008) for more on this perspective on SSB in finite classical systems, and a detailed discussion of the situation in classical field theories.

[15]In order to obtain a non-trivial quantum version of the Ising model a transverse magnetic field term has to be added, so the basic Hamiltonian is:

$$H = -J \sum_{i,j} \sigma_i^z \sigma_j^z - h_x \sum_i \sigma_i^x,$$

where $\{\sigma^x, \sigma^y, \sigma^z\}$ are the Pauli matrices. Henceforth I refer to this model simply as the 'quantum Ising model'.

The Stone-von Neumann theorem famously shows (roughly speaking) that all possible quantisations of a finite classical system, obtained via the usual canonical commutation relations, are unitarily equivalent. What this means is that, while it is possible to write down a finite quantum system on a Hilbert space of sequences of complex numbers, as in Heisenberg's matrix formulation of quantum mechanics, and on a Hilbert space of complex valued functions, as in Schrodinger's wave-function formulation, there is an isomorphism between these structures, which has typically been taken to invite identifying them as theoretically equivalent.[16] This result breaks down in the limit of infinite degrees of freedom. The best way to see what is going on here is to invoke the algebraic formalisation of quantum theory. In this framework a quantum system is associated with an abstract $C \times$ *-algebra which can be instantiated in multiple Hilbert spaces. What happens in quantum theories with infinite degrees of freedom is that the corresponding C^*-algebra admits multiple unitarily inequivalent Hilbert space representations which cannot simply be viewed as equivalent encodings of the same physics.

This leads to puzzles with understanding the semantics of these theories. As Ruetsche (2011) chronicles in detail, the failure of the Stone-von Neumann theorem signals the splintering of previously equivalent strategies for assigning physical content to quantum systems. But the appearance of unitarily inequivalent representations also provides resources for representing SSB behaviour. It turns out that these unitarily inequivalent representations can support distinct ground states; the considerations sketched above for finite quantum systems no longer apply in this context because the superposition of states in unitarily inequivalent representations is blocked by so-called superselection rules. In the thermodynamic limit, ψ_+ and ψ_- converge to genuine ground states belonging to unitarily inequivalent representations of the infinite quantum Ising model. Furthermore, it is impossible from the system to transition from one of these ground states to the other, so the sense in which we can expect stable asymmetric behaviour is especially stark.

In quantum statistical mechanics the existence of unitarily inequivalent representations in the thermodynamic limit furnishes an approach to modelling SSB which is, in other respects, highly analogous to the classical case. The canonical ensemble of a finite quantum system is traditionally represented by a density operator,

$$\rho = \frac{\mathrm{e}^{-\beta H}}{Z}. \tag{6.3}$$

which, again, is not well defined in the thermodynamic limit. In the quantum

[16]Precisely: a Hilbert space and collection of operators $(\mathcal{H}, \mathcal{O}_i)$ is unitarily equivalent to another, $(\mathcal{H}', \mathcal{O}'_i)$, if there exists a unitary map $U : \mathcal{H} \to \mathcal{H}'$ such that $U^{-1}\mathcal{O}'_i U = \mathcal{O}_i$ for all i. See Ruetsche (2011) chapter 2 for a discussion of the connection between the unitary and theoretical equivalence of quantum theories.

context the algebraic formulation of quantum statistical mechanics provides a rigorous framework for generalising the canonical ensemble to infinite systems.[17] The analogue of Gibbs measures are so-called KMS states, which agree with (3) for finite systems, but can also exist in the thermodynamic limit. As in the classical case, the canonical ensemble of a finite system is unique, but a system with infinite degrees of freedom can have a convex set of KMS states whose extremal elements correspond to distinct symmetry breaking equilibrium states. In the quantum case these symmetry breaking equilibrium states live in unitarily inequivalent Hilbert space representations afforded by the limit of infinite volume. Again this approach to SSB has been extremely successful and connects up with a powerful apparatus for dealing with phase transitions in quantum systems. As in the classical case however there is a puzzle about why these infinite volume systems provide a good description of the behaviour of concrete lumps of iron which are patently finite in extent.

What about high energy physics and the QFT programme: what role does SSB play in this context and how is it represented? The picture here is complex and ultimately poorly understood from a foundational perspective. The most significant, but also most controversial, appeal to the notion of SSB in high energy physics takes place within the so-called Higgs mechanism, a central pillar of the standard model. The Higgs mechanism is essentially a method for adding masses to the gauge bosons associated with weak interactions without explicitly breaking the underlying gauge symmetries, and is standardly understood in terms of SSB. Coupling the electroweak sector of the standard model to a scalar field—the Higgs field—is said to spontaneously break the relevant gauge symmetries giving rise to massive gauge particles in the process. A gamut of technical and conceptual issues arise here however. For one thing, many philosophers of physics have objected to the very idea of spontaneously broken gauge symmetry. The gist of this critique is that, because gauge freedom is standardly interpreted as a kind of descriptive redundancy, it just doesn't make sense to talk about gauge violating ground and equilibrium states, much less for particles to acquire masses in this way—see Earman (2004), Healey (2007) and Lyre (2008) for complaints along these lines.

In addition to these more philosophical worries there are tangible formal barriers to understanding the Higgs mechanism as a case of quantum SSB. What often happens in high energy physics is that the Higgs mechanism, and the accompanying SSB argument, are presented at the level of classical field theory and perturbative methods are then used to calculate quantum corrections, on the assumption that the non-perturbative picture is more or less preserved when

[17] See Bratteli and Robinson (2003) for the details of this approach. Presentations aimed at philosophers can be found in Emch and Lui (2005) and Ruetsche (2011).

the theory is quantised. Unfortunately indications are that this is simply not the case. A result known as Elitzur's theorem seems to show that local gauge symmetries cannot be spontaneously broken in QFT models (Elitzur, 1975).[18] Remnant global gauge symmetries are not touched by this results and can be spontaneously broken, but all of the evidence suggests this has nothing to do with imparting mass to gauge particles. As a result of all these difficulties many working in the foundations of physics have concluded that the Higgs mechanism needs to be reconceptualised in terms which do not appeal to SSB at all.

Even if we abandon the idea that spontaneously broken gauge symmetry plays an important role in the standard model however, this does not mean jettisoning the notion of SSB from the QFT programme entirely. In fact, SSB plays an important role in less controversial contexts in high energy physics. The spontaneous breaking of chiral symmetry in QCD, for instance, is said to account for around 99% of the mass of nucleons (and therefore most of the mass of ordinary matter). This case is not touched by gauge related concerns. SSB is also invoked in scalar field theories which do not involve gauge fields in any way. The spontaneous breaking of internal rotational symmetries are studied in multi-component generalisations of the familiar ϕ^4 model known as $O(n)$ models (which are, in fact, closely related to the Higgs field).[19] We do not necessarily need to solve the riddles of the Higgs mechanism to make progress in understanding the explanatory status of SSB in contemporary high energy physics then, and I will largely bracket the difficulties raised by the peculiarities of local gauge symmetries here.

When it comes to giving a precise characterisation of SSB in QFT models there are two main approaches on the table. On the one hand, a framework for capturing SSB has been developed in the axiomatic approach to QFT, and especially in the algebraic formulation of continuum QFT models. Just as in the algebraic approach to statistical mechanics, this characterisation of SSB makes use of unitarily inequivalent representations afforded by the limit of infinite degrees of freedom. In particular, we can see the spontaneous breaking of symmetries in QFT as arising from the existence of unitarily inequivalent representations that are related by a symmetry of the underlying algebra and support distinct non-invariant ground (or in this context, vacuum) states.[20] The main drawback of

[18] Elitzur's theorem, and its philosophical ramifications are discussed in Freidrich (2013). Note that the result is proven in the context of lattice gauge theory, but there are heuristic indications that it should also apply in the continuum limit (if it exists).

[19] The $O(n)$ models are scalar field theories with the classical Lagrangian density:

$$\mathcal{L} = \partial_\mu \phi^\alpha \partial^\mu \phi^\alpha - m^2 \phi^\alpha \phi^\alpha - \lambda(\phi^\alpha \phi^\alpha)^2$$

where $\phi^\alpha(x)$, $\alpha = 0, n-1$, is a n-vector of scalar fields. For $n > 1$ there is a $O(n)$ rotation symmetry, in the internal space of this vector, which can be spontaneously broken if $m^2 < 0$. Clearly when $n = 1$ we get back the familiar ϕ^4 theory.

[20] A detailed discussion of this formalism can be found in Earman (2004) and Strocchi (2008).

this approach comes as no surprise—we currently do not have continuum formulations of realistic QFT models in the algebraic framework. While it might hope to illuminate the nature of SSB in QFT at a very high level of abstraction then, this formalism does not furnish concrete scientific explanations in the way that the thermodynamic limit of quantum statistical mechanics does.

SSB is studied directly in realistic, four dimensional, lattice QFT models however. There are strong parallels here with the way that SSB is modelled in the lattice spin systems already discussed. In fact, there is a direct correspondence between the aforementioned $O(n)$ scalar field models quantised on a cubic lattice and generalisations of the Ising model (Smit, 2002, 34-35). Assuming zero temperature, SSB is associated with the existence of degenerate non-invariant ground states in the infinite volume limit which presumably belong to unitarily inequivalent representations. The kind of rigorous results pertaining to the Ising model sketched above are not forthcoming in the case of infinite volume lattice QFT systems however. On the one hand, various approximation methods imported from statistical mechanics, such as mean field theory techniques, can be used to estimate some of the relevant properties that obtain in this limit. But numerical simulations based on finite lattices are also commonly used to study SSB. Since these systems have a finite number of degrees of freedom unitarily inequivalent representations and degenerate ground states do not occur, but various techniques are used to extrapolate numerical results to the infinite volume limit. One method which is often employed to estimate the value of expectation values of degenerate vacuum states in the infinite volume limit is to add a symmetry breaking perturbation to the Lagrangian and examine what happens as it is gradually removed as the lattice volume is increased.[21] In sum, the limit of infinite degrees of freedom continues to play a role, if a less direct one, in the way that SSB is modelled in lattice QFT.

6.3 Motivating the Deflationary View

We have seen that the way that SSB phenomena are modelled in both classical and quantum physics appeals to novel properties of infinite systems. Ruetsche (2011) takes the role that resources afforded by multiple unitarily inequivalent representations play in the quantum case to show that these properties need to be taken physically seriously—states afforded by multiple inequivalent representations must to be viewed as genuine possible states for the ferromagnet in order to account for the onset of symmetry breaking behaviour, for instance. As I have already alluded to, this perspective is closely allied to Batterman's (2005)

[21]See Smit (2002, 64-66) and references therein.

claims about phase transition theory. According to him the discontinuities in the thermodynamic potentials that occur in the thermodynamic limit are needed to represent physical discontinuities which manifest in concrete systems during a change of phase. Call the general idea that novel properties of infinite systems must be taken to faithfully represent in order to underwrite the explanatory capacity of these models the indispensablist view.

In the context of statistical mechanics this position immediately runs into puzzles I have already touched on. The infinite volume models obtained in the thermodynamic limit are highly idealised representations of the concrete systems which exhibit SSB phenomena—real ferromagnets clearly have a finite volume and number of constituents. But the indispensablist line cuts against a familiar view of how idealised models relate to their targets. In calculating the electric field produced by a long charged wire one often treats it as being infinitely long. In so far as this model is predictive and explanatory it is presumably because it captures relevant features of the target system that are preserved when this infinite length idealisation is corrected—we can easily verify in this case that the electric field associated with an infinite charged wire provides a good approximation to the field close to a long finite length wire. The doctrine that idealised models are successful because the behaviour they aim to describe is robust under 'de-idealisation' is most famously associated with McMullin's notion of Galilean idealisation (McMullin, 1985). If the indispenablist thesis is right, however, the role of the thermodynamic limit cannot be understood in these terms. Properties of infinite systems, which completely disappear when we move to a large finite volume, are explanatorily essential and representationally faithful on this account.

The indispensablist then owes us an alternative account of how the infinite models employed in this context manage to successfully predict and explain. There are two main routes which philosophers with indispensablist sympathies have taken here, both of which have their fair share of problems in my view. Batterman has championed the idea that the thermodynamic limit is an instance of a broader class of essential (or ineliminable, or uncontrollable) idealisations which contribute to the explanatory power of a model in a radically different way from familiar Galilean idealisations, the use of continuum descriptions in fluid mechanics being another putative example of this phenomenon. My central gripe with this approach is that, while Batterman has consistently maintained that his conception of essential idealisation cannot be assimilated to established accounts of scientific explanation (Batterman, 2002) or the applicability of mathematics (Batterman, 2010), the positive picture of how essential idealisations are actually supposed to work remains inchoate.[22] If essential idealisations are to provide a viable alternative to the Galilean paradigm what we need is an explicit account of why essentially idealised models are successful which does not appeal to their relation to de-idealised representations of the target system.[23] In the absence of such an account the worry is that calling the thermodynamic limit as essential idealisation merely gives a name to the problem facing the indispensablist.

While the doctrine of essential idealisation has often been associated with claims about the failure of reduction relations between models at different level of description, Earman (2004) puts forward a very different way of understanding the indispensablist thesis which is thoroughly reductionist in character. As we have seen, the mathematical structures found in the thermodynamic limit also occur in a field theoretic context. Assuming that the world is, at base, a continuum QFT we could take the novel properties that appear in the thermodynamic limit to correspond to features of this more fundamental physical description which are missed by models with finite degrees of freedom. On this account the limit of infinite degrees of freedom is not an idealisation after all! Unfortunately for Earman, whether we have grounds to take the novel features of continuum QFT models to faithfully represent is precisely what I am questioning in this chapter. Even if we take this for granted however, there is something ad hoc, or at least thoroughly programmatic, about this response. What seems to be needed to meet the idealisation challenge here is not just an assurance that degenerate equilibrium states are in fact possible but rather a detailed story about why they occur in the particular materials, and at the particular temperatures, they do. Prima facie, field theory has nothing to do with ferromagnetism and the boiling of kettles.[24] Ruetsche (2011, 336-339) suggests a less risky way of implementing Earman's reductionist strategy; the indispensablist can simply claim that the relevant properties found in the thermodynamic limit will be embedded within future physical theories in some way or other. Again though, the idea that advances in quantum gravity will shed light on the question of why the thermodynamic limit works seems far fetched and unmotivated.

In the recent debate about phase transition theory a number of philosophers have argued that a more conservative way of understanding the success of the

[22] Batterman and Rice (2014) is perhaps the most detailed discussion of how Batterman and others understand the explanatoriness of essential idealisations to date. See Lange (2015) for criticisms with which I am sympathetic. Furthermore, it is not clear how the minimal model account of explanation given there motivates, or responds to concerns about, the indispensablist claim that novel properties afforded by the thermodynamic limit faithfully represent.

[23] Shesh (2013) makes the same point, though he is more optimistic than I that this challenge can be met.

[24] Butterfield (2011, 1078) and Mainwood (2006, 228-231) raise similar objections to invoking field theory to justify the use of the thermodynamic limit.

thermodynamic limit is available which avoids much of this wrangling. As I understand them, Butterfield (2011), Norton (2011) and Callender and Menon (2013) want to say that, despite appearances, the thermodynamic limit is an ordinary idealisation, not fundamentally unlike the limit of infinite length in the charged wire case. Note that the fact that novel properties obtain in an idealised limit does not, in itself, establish that they are indispensable in any interesting sense. The electric field produced by an infinitely long charged wire has a translation symmetry that is not respected by any finite length wire, but nobody claims that this property is explanatorily essential, much less physically real. The reason being that the success of the infinite wire model is very plausibly accounted for by other properties which are robust when the infinite length idealisation is removed. Similarly, according to these authors, it is not the novel properties of infinite systems but those they share with large finite systems that are doing the real explanatory and representational work in applications of the thermodynamic limit. Call this the deflationary view.

The basic observation underlying this position in the phase transition debate has been that, while the non-analyticities in the free energy disappear when the thermodynamic limit is removed, this function can still approximate the changes in macroscopic properties of a finite system. Recall that, below the critical temperature, the expected magnetisation of the Ising model becomes a discontinuous function of the external magnetic field in the thermodynamic limit, signalling a first order phase transition (figure 6.2). General theoretical considerations, and evidence from numerical simulations, suggest that the expected magnetisation of a finite Ising model has an abrupt change in sign at $h = 0$, which becomes increasingly steep as N increases, so that, if the lattice is sufficiently large, it will be very well approximated by the discontinuous magnetisation function of an infinite Ising lattice. We seem to have a simple schema for explaining the success of the orthodox approach to phase transitions in terms of the behaviour of large finite systems: a real phase transition is a sharp but smooth change in a large system's macroscopic properties that is approximated by the non-analytic free energy found in the infinite volume limit to accuracies within acceptable error.

How does this debate about the thermodynamic limit bear on the question we are interested in: the status of novel properties of infinite systems in QFT? Ruetsche suggests that the indispensablist thesis goes through much more smoothly in this context. The thought being that, since QFTs actually are infinite systems, there is no worry about idealisations in this case and the appeal to novel properties afforded by the limit of infinite degrees of freedom can simply be taken at face value. There are problems with this line of thought however. In the statistical mechanics case it is external information about the volume of concrete ferromagnets which tells us that the infinite volume limit is an idealisation. In the QFT case we do not have this kind of information. Still, this clearly does not mean that the continuum and infinite volume limits should be read realistically by default. For one thing, whether the universe has a continuum or discrete structure, and whether it is spatially infinite or finite, seem to be open scientific questions.[25] Even if there were good reasons to believe the universe is infinite however this does not mean that we should trust what our present QFT models say about the world at spacial infinity. I have been arguing in this thesis that the claims that QFT models make about physics at arbitrarily large and small distance scales are not supported by the extant empirical successes of high energy physics. On Ruetsche's account, however, it seems that it is precisely this kind of structure which needs to be taken physically seriously to account for SSB phenomena in high energy physics. While there is a less flagrant conflict in the QFT case than the statistical mechanics case then, there is still something puzzling about the claim that novel properties afforded by the limit of infinite degrees of freedom need to be taken to faithfully represent in this context.

Furthermore, there is a tension in the suggestion that one might be a selective indispensablist, giving in to worries generated by the thermodynamic limit in the statistical mechanics case while maintaining that novel properties of infinite systems must be taken representationally seriously in order to account for SSB in QFT. The challenge posed by the thermodynamic limit's status as an idealisation, as I see it, is not that it undermines the explanatoriness of infinite systems—I take it as read that the standard approach to phase transitions and SSB in statistical mechanics is genuinely explanatory. The point is rather that it mitigates against the indispensablist's claim that the representational success of features only found in infinite systems underwrite their explanatory success. If we accept that we do not need to reify these properties to account for the success of the thermodynamic limit in statistical mechanics, as I take it Butterfield and others do, why should we think that novel features of infinite systems engender ontological commitment in high energy physics? To motivate this selective stance we would need to point to a relevant difference in the explanatory set up in these two cases. In fact however, philosophers of physics have tended to rely on analogies with the relatively well understood description of SSB in quantum statistical mechanics to inform their

[25]It is sometimes suggested that the fact that all present measurements of the energy density of the universe are consistent with zero space-time curvature gives us reason to believe that the universe is infinite. This is a simplification however; this inference only goes through if we assume, as standard cosmological models typically do for reasons of simplicity, that the topology of the universe is simply connected (Lachieze-Rey and Luminet, 1995). In fact, zero curvature is perfectly compatible with a finite volume space-time. To my knowledge, the question of whether the universe has a finite or infinite volume is an open, if heavily constrained, question.

understanding of the more controversial case of QFT.[26] This suggests a strategy for responding to the attack on doubly cutoff QFTs we started out with. If we take a deflationary view of the explanatory role of the thermodynamic limit in statistical mechanics this will undermine the aforementioned worries about the representational inadequacies of doubly cutoff QFTs.

There is a hitch here however. It is not at all obvious that the kind of treatment of phase transitions advocated by Butterfield and others can be extended to the case of SSB. Pointing to the fact that non-analytic functions can approximate analytic ones does not help in this case. The issue now is not why we can legitimately model a ferromagnet's macroscopic properties as changing discontinuously under variations in the external magnetic field but why, in the absence of an external magnetic field, the system can be successfully represented as having multiple asymmetric equilibrium states, despite the uniqueness of the equilibrium state of any finite system. As we have seen, if a system's ground and equilibrium states are unique they must be invariant under the same symmetries as the dynamics. In the statistical mechanics context, this means that the expected values of macroscopic observables at equilibrium must be consistent with these symmetries; the expected magnetisation of a finite Ising model, for instance, is always zero no matter how large the lattice is. Prima facie then, the description of SSB afforded by the infinite volume limit cannot approximate behaviour already found in large finite systems in any reasonable sense. SSB might seem to offer new support to the indispensablist view then. The next section argues that this is not in fact the case; there is scope for extending the deflationary view to the case of SSB.

6.4 Spontaneous Symmetry Breaking in Finite Systems

The kind of no-go argument just sketched is misleading. SSB behaviour only appears to be impossible in finite systems if the modelling assumptions implicit in standard characterisations of SSB via infinite systems are treated as sacrosanct. Relaxing the requirement that the system's ground/equilibrium state is asymmetric (the symmetry breaking part), or the demand that the broken symmetry is exact (the spontaneous part), allows us to see how systems with finite degrees of freedom can exhibit the kind of behaviour observed in real ferromagnets. §6.4.1 explores the former route. I point out that a classical system with a unique, symmetric, equilibrium state can still be expected to exhibit asymmetric states over long time periods. Similarly, in the quantum case a system with a unique ground state can nevertheless furnish stable symmetry breaking behaviour. §6.4.2

[26] See, for instance, Emch and Lui (2005). Fraser (2012) raises worries about applying this strategy uncritically.

explores the latter route. Landsman (2013) has developed an approach to SSB in finite quantum systems based on the instability of their ground and equilibrium states under asymmetric perturbations to the dynamics. The upshot is that, despite initial pessimism, there is a viable programme for accounting for the success of infinite systems in modelling SSB without reifying them.

6.4.1 Approach I: Long Lived Asymmetric States

The fact that a finite system's equilibrium state is invariant under the symmetries of the dynamics does not, in fact, rule out SSB type behaviour: an equilibrated finite system can still be expected to exhibit asymmetric states over very long time periods. Once again, the classical Ising model provides a useful reference point here. The canonical ensemble of the Ising model assigns a probability distribution, $P(m)$, to the possible magnetisations of the system. Provided the boundary conditions imposed on a finite lattice do not break the interaction's reflection symmetry, $P(m)$ will also respect this symmetry, which, as has already been noted, implies that the expected magnetisation is zero. This does not mean however that $P(m)$ must take the form of a normal distribution centred at $m = 0$; it can have maxima at positive and negative m. The system would then be expected to spend most of its time in magnetised microstates.

In the case of the Ising model on a two dimensional lattice it can be shown that this is exactly what happens below the critical temperature. As was mentioned in §6.2.1, if the boundary conditions are suitably symmetric the Gibbs measure of the system is an equal linear combination of the P_+ and P_- measures. These extremal measures assign probability one to $m = \pm a$ (for some positive a) in

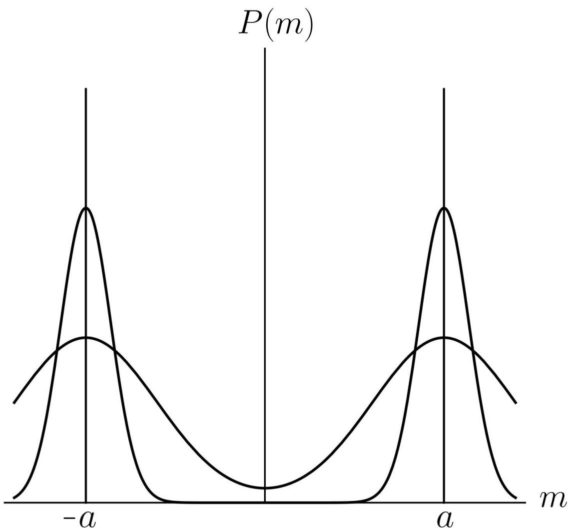

Figure 6.3: Sketch of $P(m)$ below the critical temperature for a classical Ising model with small, large and infinite N.

the thermodynamic limit, and can be shown to satisfy central limit theorems as $N \to \infty$.[27] This means that, below the critical temperature, $P(m)$ has two 'humps' at $m \pm a$, which become increasingly sharply peaked as N increases, finally converging to Dirac measures at $N = \infty$ (figure 6.3). Though rigorous results are less forthcoming for more realistic models there are general theoretical grounds for believing that analogous behaviour will hold whenever the Gibbs measure is non-unique in the thermodynamic limit. The generic result that statistical fluctuations scale like $1/\sqrt{N}$ (where N is now the number of constituents) implies that the canonical ensemble of a large finite system will have most of its mass in microstates which assign values to macrostate variables close to those associated with the degenerate equilibrium states that appear in the limit.

Consequently, we can expect such systems to spend long time periods in symmetry breaking states. A large finite Ising model below the critical temperature will enter one of the two regions of its state space in which $m \approx \pm a$, with equal probability, and remain there for a very long time. In fact, we can give heuristic arguments—ultimately based on artificial implementations of the system's dynamics, like the so-called Glauber dynamics—that a finite Ising model will transition between states of positive and negative magnetisation with a time period proportional to $\exp \sqrt{N}$.[28] In the thermodynamic limit the system will remain in an 'up' or 'down' aligned state forever—they become distinct equilibrium states. But if N is of the order of Avagadro's number the system can still be expected to exhibit a magnetisation in a particular direction for a very long time indeed: much longer than we can feasibly observe real ferromagnets for, and perhaps longer than the age of the universe!

This points to an explanation of the successful appeal to infinite systems in modelling SSB in classical statistical mechanics in terms of properties they share with large finite systems. On this view, what we actually observe in real ferromagnets at low temperatures are long lived asymmetric states whose macroscopic properties are well approximated by the degenerate equilibrium states found in the thermodynamic limit. The picture of SSB afforded by the thermodynamic limit, in which the system must choose between a set of symmetry breaking equilibrium states, thus provides an appropriate description of the behaviour of large systems over long, but finite, time periods.[29]

[27] See Kindermann and Snell (1980, 34-62) and references therein.

[28] See Lebowitz (1999, 353) and Kinderman and Snell (1980, 55-61).

[29] In fact, this way of understanding how the standard account of SSB relates to the behaviour of finite systems is sometimes found in the statistical physics literature. Commenting on the non-uniqueness of the Gibbs measure of the Ising model in the thermodynamic limit, Lebowitz writes:

> This means physically that when V is very large the system with "symmetric" [boundary conditions] will, with equal probability, be found in *either* the

Can this kind of story be applied to quantum systems as well? There are indications that it can. Finite quantum systems can support asymmetric states which are stable over very long time periods, despite the uniqueness of their ground states, as can be seen in the case of the quantum Ising model. If the potential of a finite system has multiple minima, the associated asymmetric quantum states can typically be expressed as superpositions of the ground and first excited states: in the quantum Ising model, $\psi_{\pm} = \frac{1}{\sqrt{2}}(\psi_0 \pm \psi_1)$. What happens as $N \to \infty$ is that the energy gap between these states vanishes. In the case of the quantum Ising model it can be shown that the energy gap between the ground and first excited states, ΔE, decays like $\exp(-N)$, and Koma and Tasaki (1994) prove analogous results for a broader class of lattice models which display degenerate ground states in the thermodynamic limit. This means that, if the system is large, these asymmetric states will have a very slow time evolution. A quantum Ising model which is prepared in the ψ_+ state will remain in a magnetised state for an arbitrarily long time, for arbitrarily large N.

In fact, the existence of these slowly precessing asymmetric states has long been put forward in the physics literature as a way of understanding how finite quantum systems can exhibit SSB, most notably in Anderson's famous discussion of the electric dipole moment of ammonia and sugar (Anderson 1972). The ammonia molecule is known to fluctuate very quickly between states of opposite polarity, which correspond to local minima in the potential separated by a small potential barrier. As Anderson describes however, when we move to larger and larger molecules the potential barrier between the states of opposite polarity increases and the gap between the ground and first excited states gets smaller. When we get to complex molecules the inversion occurs on such long time scales that the relevant symmetry breaking states can be treated as stationary for all practical purposes.

In the statistical mechanics context, this suggests that the macroscopic properties of a large quantum system can be well described by the degenerate ground and equilibrium states found in the thermodynamic limit. The details of this picture still need to be filled in however. There remains the question of how the system enters an asymmetric state in the first place, and why we do not see the true quantum ground state.[30] It may be that the canonical ensemble of a finite quantum

> "+ state" or the "− state". Of course as long as the system is finite it will "fluctuate" between these two pure phases, but the "relaxation times" for such fluctuations grows exponentially in V, so the either/or description correctly captures the behavior of macroscopic systems. (Lebowitz, 353)

A similar discussion is found in Binney et al. (1992, 48-51).

[30] Some authors have suggested that there is a close connection here with the measurement problem, see Landsman (2013) and Emch and Lui (2005). This is something of a red herring in my view however and I keep the question of how to understand quantum SSB separate

system which displays degenerate equilibrium states in the thermodynamic limit has most of its mass in symmetry breaking pure states at low temperatures, and the account sketched above for the classical Ising model can be straightforwardly extended to quantum systems. But it has been suggested by some authors that one must appeal to additional resources, and specifically the system's coupling to the environment, to account for the appearance of SSB behaviour in finite quantum systems. The next section examines an alternative approach to SSB in finite systems in this vein which has been explicitly developed in the quantum context.

6.4.2 Approach II: Instability Under Asymmetric Perturbations

Landsman (2013) has put forward a different approach to SSB in finite quantum systems. It turns out that the ground and equilibrium states of a quantum system can be very unstable under small asymmetric perturbations to the dynamics. If we relax the assumption that the symmetry which is broken at low temperatures is exact, this provides another way of understanding how SSB phenomena can manifest in finite systems.

It can be shown that adding a small perturbation to the Hamiltonian of a large quantum Ising model that favours 'up' aligning states causes the system's ground state to abruptly shift towards ψ_+, and converge to it as $N \to \infty$, with analogous statements holding for the system's equilibrium state below the critical temperature. One way to see why this occurs is to consider the effective Hamiltonian governing the ground and first excited states. In the $\psi_\pm$ basis,

$$H = \begin{pmatrix} 0 & -\Delta E \\ -\Delta E & 0 \end{pmatrix}, \tag{6.4}$$

where ΔE is, again, the energy difference between the ground and first excited state, which, as we have seen, decays exponentially in N. Consequently, any term added to the upper diagonal element, no matter how small, will dominate the effective Hamiltonian for sufficiently large N, shifting the ground state close to ψ_+. Of course, the same argument applies to perturbations favouring 'down' aligning spins. More generally, this type of instability under asymmetric perturbations can be expected to hold if the energy gap between ground and first excited states is vanishingly small, which, as we saw in the previous section has been shown to arise in a very wide range of systems with multiple minima in their potentials. Accordingly, Landsman provides a similar analysis of the quantum double well potential. What we find in this case is that under parity violating perturba-

from these broader interpretive considerations here.

tions the quantum ground state collapses towards one of the local minima of the potential. This is all spelled out in rigorous terms in Landsman (2013).

This provides another way of accounting for the success of the standard approach to modelling SSB in statistical mechanics by way of the behaviour of large finite systems. Modelling the symmetry which is broken at low temperatures as an exact invariance of the target system's dynamics is often itself an idealisation. In many concrete systems, small, otherwise empirically negligible, asymmetric perturbations are very likely to be physically instantiated. In the case of ferromagnetic crystals such perturbations might originate in the system's coupling to its environment, background magnetic fields, or defects in the lattice structure. At low temperatures these tiny asymmetries have a dramatic effect on the behaviour of the system, shifting it into an equilibrium state whose macroscopic properties are very well approximated by one of the degenerate equilibrium states found in the thermodynamic limit. Given our ignorance of the details of the perturbations affecting a real ferromagnet there is an obvious epistemic sense in which a ferromagnet has an equal chance of aligning in any particular direction, and it is appropriate to represent the system as having a number of available equilibrium states below the critical temperature—thus there remains a sense in which the symmetry breaking might be said to be spontaneous.

While Landsman's discussion is grounded in quantum theory, it is plausible that this approach to SSB in finite systems can also be applied to classical systems. If the expected magnetisation of a large classical Ising model below the critical temperature is well approximated by the discontinuous function found in the thermodynamic limit, as Butterfield and others argue, it will be very sensitive to changes in the external magnetic field. As a result, a tiny background magnetic field will shift the system's equilibrium state into a magnetised state which is well approximated by one of the symmetry breaking Gibbs measures found in the thermodynamic limit.[31] This is, of course, a much less general kind of perturbation to the dynamics than those considered by Landsman, but it does suggest that the equilibrium states of finite classical systems can display the same kind of instability properties as their quantum counterparts. SSB behaviour in finite classical systems may also be produced by small asymmetries arising from environmental effects then.

A puzzle might seem to arise here. We have now seen two, apparently quite different, ways of producing the phenomenological features of SSB phenomena in finite systems: the first based on long lived asymmetric states and the second on asymmetric perturbations to the dynamics. If both accounts apply to the

[31] Binney et al (1992, 48-51) gives a characterisation of SSB in finite classical systems along these lines.

same model, how are we to understand what actually happens in a concrete system which displays SSB? There is no real difficulty here however. Following Lui (2003) we can distinguish between the features of a model which make SSB possible and the physical mechanism which brings about asymmetric behaviour in a particular system that model is used to represent. A finite volume model can display both long lived asymmetric states and the kind of instability properties cited by Landsman, yet the way that SSB arises in different physical instantiations of the model may differ from case to case. Furthermore, there is also no reason to treat these two approaches to SSB in finite systems as in competition. Each can be viewed as getting around the apparent impossibility of SSB in finite systems by weakening the symmetry breaking and spontaneous part of the standard definition of SSB respectively, but it is, of course, also possible to weaken both. It is plausible that in many concrete systems environmental perturbations and the stability properties of asymmetric states both play a role in producing SSB type phenomena.

Another potential concern is that whether or not SSB occurs in finite systems does not have a precise answer on either of these approaches. On the first approach the relevant symmetry is not really broken in the system's full time evolution. While a finite classical Ising system can display magnetised states over long time periods, as $t \to \infty$ it will spend an equal amount of time in 'up' and 'down' aligned states, consistent with the fact that the expected magnetisation is zero. But the notion of a long lived asymmetric state is clearly a vague one—the difference between the broken and unbroken phase of the ferromagnet is a matter of degree rather than kind. The second approach, on the other hand, muddies the conventional distinction between spontaneous and explicit symmetry breaking. Explicit symmetry breaking, you will recall, refers to asymmetric behaviour produced by asymmetries in the dynamics, which is exactly what the perturbations appealed to in Landsman's approach are. While there is clearly a qualitative difference between the production of magnetised states by tiny environmental effects and via the application of an external magnetic field the distinction is not a sharp one. Admitting that both approaches can play a role in producing symmetry breaking behaviour further compounds this situation.

As I stressed in §6.2.1 however, SSB behaviour in finite systems is never clear cut. It is difficult to see how this could be used to object to the approach to SSB in finite systems developed here unless one has already taken the more categorical description of SSB afforded by novel features of the infinite volume limit to be representationally faithful. I should stress again that the deflationary view, as I understand it, need not have revisionary consequences for scientific practice, and certainly does not amount to an injunction that physics abandon standard definitions of SSB and cease their appeals to infinite systems.[32] There are instructive parallels here with Callender and Menon's (2013) discussion of extensive quantities in statistical mechanics. Roughly speaking, a quantity is said to be extensive if it behaves additively when the system is divided into sub-systems. Strictly speaking, no quantity satisfies this definition when the system is finite; boundary effects spoil the additivity of quantities like the entropy and the distinction between extensive and intensive quantities only really applies in the thermodynamic limit. But, though the notion of extensivity does not carve nature at its joints, as it were, there is clearly a sense in which it is epistemically beneficial to treat entropy and other observables as extensive in situations in which boundary effects are negligible. Similarly, if the approaches to SSB in finite systems put forward here is borne out it will justify the contemporary practice of statistical mechanics, not undermine it.

How do these ways of thinking about SSB in finite systems translate into the QFT context? A great deal more work is needed to answer this question with any precision but I will make some suggestive remarks here. The kind of rigorous demonstration of the vanishing of the energy gap between ground and first excited states with increasing volume provided by Landsman (2013) and Koma and Tasaki (1994) are typically not forthcoming in the case of interesting lattice QFT models. Nevertheless, numerical simulations do seem to bear out the expectation that large lattices exhibit very long lived symmetry breaking states—in many cases we can directly see the appearance of this kind of behaviour in the numerical data. It is possible that SSB phenomena in high energy physics can be understood as long lived but ultimately unstable along the same lines I sketched for the ferromagnet. This is not as radical a suggestion as it might initially seem. In fact, the possibility that the standard model 'vacuum' is not genuinely stationary and will eventually tunnel to some other minimum in the potential has long been a live, and much discussed, possibility in high energy physics.[33]

As for Landsman's instability approach, there appear to be close connections with the method of studying SSB in finite lattice QFTs briefly mentioned in §6.2.2, in which symmetry breaking perturbations are added to the Lagrangian of finite lattice QFT systems and removed as the volume is increased. Exploring the formal relations between Landsman's formalism and this approach to modelling SSB may also illuminate the question of the existence of a vanishing energy gap.

[32]It is important to distinguish the kind of deflationary programme advocated here from attempts to offer a new definition of a phase transition in terms of the resources of finite systems, surveyed in Callender and Menon (2013). I take the deflationary reading of the thermodynamic limit to be a claim about why infinite systems, and definitions which appeal to them, are successful in statistical mechanics. Whether an alternative definition of SSB of any practical use can be provided which only refers to finite systems is another question in my view.

[33]See Bednyakov et al (2015) for an up to date discussion of this scenario.

There is also potential for the physical picture underlying Landsman's approach to have currency in the high energy physics context. In the case of spontaneously broken chiral symmetry, there really are asymmetric perturbations in full QCD in the form of non-zero quark masses, for instance. On the other hand, the gauge symmetries of the standard model are supposed to be exact. How all of this bears on the Higgs mechanism and the embattled notion of spontaneously broken gauge symmetry I do not dare speculate.

6.5 An Epistemic Role for the Limit of Infinite Degrees of Freedom

I certainly do not claim to have refuted the indispensablist view of the limit of infinite degrees of freedom in this chapter. As I flagged at the outset, novel features of infinite systems play an important role in many theoretical contexts which have not been considered here. A particularly glaring omission is the role of infinite systems in the renormalisation group framework, and especially in renormalisation group explanations of the universality of critical exponents in second order phase transitions. Batterman has long taken this case to be the key motivation for his claims about the explanatory indispensability of the thermodynamic limit.[34] Furthermore, the approach to SSB in finite systems put forward in the previous section is more of a research programme than a stable theoretical picture. In the statistical mechanics case, it is ultimately a conjecture that the key statements about large finite models can be generalised to the full range of systems to which the standard account of SSB is applied, and many of the details about how SSB behaviour manifests in concrete systems are open to further investigation. More importantly for the overarching dialectic of this thesis, many questions about SSB in the QFT programme remain unanswered.

With these caveats in mind however, I have at least developed a strategy for replying to the objection we started out with. The problem, remember, was that if novel properties afforded by the limit of infinite degrees of freedom need to be taken to faithfully represent in order to fulfil the explanatory duties of the QFT programme then the claim that doubly cutoff QFT systems are representationally adequate on the basis of present evidence will be undercut. The proceeding discussion takes much of the bite out of this critique. The idea that SSB behaviour cannot occur at all in finite quantum systems has turned out to be bogus and we

[34] Callender and Menon (2013) sketch a deflationist treatment of the thermodynamic limit in this case which again avoids the conclusion that the novel properties afforded by this limit should be taken to be physically real. I am sympathetic to their approach but there remains a great deal to be said about the status of idealisations in renormalisation group explanations.

have seen that there are prima facie viable ways of accounting for the success of the limit of infinite degrees of freedom in models of SSB in terms of the properties of finite systems. While it is possible that further investigation of the nature of SSB phenomena in high energy physics, and their representation in the QFT, will vindicate the indispensablist claim, once a prima facie viable programme for deflating the representational role of the limit of infinite degrees of is on the table the onus seems to be on the critics of cutoff QFT models to make this case.

I close with some clarificatory remarks about how denying the representational indispensability of the limit of infinite degrees of freedom should be understood in the QFT context. In statistical mechanics we (arguably!) have good reason to believe that the concrete systems we are trying to describe really do have a finite number of degrees of freedom. In this case then, the deflationist might legitimately claim that the symmetry breaking states exhibited by these systems are not truly stationary and/or are produced by otherwise negligible symmetry breaking effects. As I pointed out however, the representational status of the limit of infinite degrees of freedom in the QFT context is unknown. The claim is not that the universe actually has a finite number of degrees of freedom and SSB phenomena in high energy physics are ontologically on a par with what happens in a ferromagnet then. Rather, in this case the point is an epistemic one: we simply aren't in a position to know whether SSB phenomena in high energy physics are produced by the existence of degenerate ground states or one of the weaker characterisations of symmetry breaking behaviour furnished by finite systems. I am not claiming that unitarily inequivalent representations are categorically unphysical then, merely that everything we have licensed to believe on the basis of present evidence can be expressed within a single Hilbert space.

This does not rule out the possibility that the limit of infinite degrees of freedom plays a crucial epistemic role in the QFT programme however. Defenders of a deflationary view of the thermodynamic limit in statistical mechanics, such as Butterfield (2011, 1073), are generally happy to admit that it is "epistemically indispensable", in the sense that it facilitates derivations of the properties of large finite systems which would be difficult, or perhaps even strictly impossible, without it. It may be that we can say something similar about to the limit of infinite degrees of freedom in the QFT context. It is plausible (if speculative) that bringing the infrared behaviour of models like QCD under mathematical control, and obtaining a proper understanding of the elusive phenomenon of confinement, will essentially involve the limit of infinite volume, for instance. Though novel properties afforded by the limit of infinite degrees of freedom do not faithfully represent on my view, they can still play a role in articulating the belief worthy content of QFT models then. Applying this moral in the context of foundational study raises points to another place where Fraser's and Wallace's views of the

QFT programme can be brought closer together. We can maintain the view of QFTs as coarse-grained representations I have been defending in this thesis while allowing that the axiomatic approach to QFT provides a convenient, perhaps essential, framework for deriving foundationally important results about QFT systems. As with the thermodynamic limit however, if we are going to trust results obtained in this framework to describe the actual world we need to be sure that they apply, at least in some approximate sense, to cutoff QFT systems as well.

Suppose that, in light of further work on these issues, it turns out that we really do need to take novel properties afforded by the limit of infinite degrees of freedom to be physically real in order to meet the explanatory ambitions of high energy physics. Where will this leave us with respect to the perspective on QFT elaborated in this thesis? It may not be a death blow for realism about cutoff QFT models. There is still the possibility of going back to Wallace's original response to the challenge and advocate a realistic reading of the infinite volume limit while maintaining a different stance towards the ultraviolet cutoff. We could then simply accept Ruetsche's thesis that unitarily inequivalent representations play a thick, ontologically involved, role in explanations in high energy physics. As I mentioned in §6.1, Ruetsche claims that the way that unitarily inequivalent representations feature in explanations leads to a complications with spelling out the semantic content of infinite quantum systems which are inimical to scientific realism. Still, it is possible that these arguments can be countered in other ways from the strategy pursued in this chapter, or that some form of scientific realism can be defended while accepting Ruetsche's coalescence approach to the semantics of QFT.[35] I leave these possibilities as avenues for future investigation.

[35]See French (2012) for suggestions along these lines.

Chapter 7

A Realist View of Quantum Field Theory

7.1 The View from the General Philosophy of Science

The past two chapters addressed objections to taking cutoff QFT models seriously which were grounded in theory specific concerns. In this chapter I turn towards broader issues in the philosophy of science, and especially the scientific realism debate, which bear on our assessment of cutoff QFT models and the QFT programme in general.

A number of new challenges arise in this context. In chapter 4 I suggested that the view of cutoff QFT models which emerged from my examination of the renormalisation group approach constitutes a realist stance towards these theories. There might seem to be a tension here however: scientific realism is naively understood as the doctrine that our best theories tell us what reality is really like, but I claimed that we have good reason to take cutoff QFTs to misrepresent the small scale structure of the world. Indeed, many philosophers of physics see cutoff QFT models as being inimical to a realist view of science. Butterfield and Bouatta (2014) claim that imposing a high momentum cutoff amounts to taking an "opportunistic or instrumentalist attitude to quantum field theories" and suggest that a genuinely realist view is only viable for QFT models that admit continuum formulations.[1] Similar sentiments about the pragmatic character of cutoff models can be found dotted throughout the philosophical literature on

[1] According to Butterfield and Bouatta, it seems, we ought to take a selective attitude to QFT models, reserving realist commitments for theories which have ultraviolet fixed points, like QCD, and eschewing those that apparently don't, like QED, on the grounds that only the former admit continuum formulations (as we've seen, the question of whether continuum models exist is somewhat more complex than this). This view is put forward most explicitly in Butterfield and Bouatta (2012).

QFT—see, for instance, Huggett and Weingard (1995), Fraser (2009, 2011) and Crowther (2013). One challenge facing advocates of the physical significance of cutoff QFTs then is the apparent need to reconcile this stance with widely held intuitions about what realism about fundamental physics amounts to.

Another issue with roots in the scientific realism debate is the existence of a putative underdetermination problem within the QFT programme. We already discussed Doreen Fraser's claims about underdetermination between cutoff and continuum QFTs in our discussion of relativity in QFT in §5.3, but there are also more general issues raised by Fraser discussion in the context of the broader realism debate. According to her the underdetermination we find in the QFT programme, which is in fact aggravated by renormalisation group results, undermines any kind of no-miracles argument for taking cutoff QFT models to be approximately true and instead motivates an anti-realist view of these theories according to which they merely capture the empirical content of QFT (Fraser, 2011). Again, zooming out to consider broader issues in the philosophy of science yields new ammunition for opponents of the analysis of the QFT programme I have been expounding.

This chapter turns this assessment of the view from the general philosophy of science on its head. Not only does the view of QFT developed in previous chapters fit well with recent developments in the realism debate in my view, it even helps address some of the challenges facing contemporary formulations of realism. §7.2 surveys the persistent difficulties associated with formulating a viable version of scientific realism. I express sympathy with a recent trend in the realism debate which seeks to solve these problems on a local, theory-by-theory, basis. One worry with this approach, however, is that it leads to a very weak form of realism about our current most fundamental theories. In §7.3 I argue that the renormalisation group framework helps address this concern in the context of the QFT programme by providing grounds for identifying which features of current QFT models faithfully represent without making guesses about the details of future physical theories. With this machinery in hand, §7.4 addresses Fraser's claims about underdetermination in the QFT programme. I claim that the kind of underdetermination that exists between different QFT models actually supports my claims about the coarse-grained representational success of QFT models. I conclude, in §7.5, by clarifying a potentially worrying feature of the view which emerges from this discussion; QFT does not tell us how the world is at the fundamental level. While this leads to important issues for further discussion about the QFT programme, I claim that it does not undermine the realist credentials of the view put forward here.[2]

[2]The ideas put forward in this chapter bears a strong affinity to a recent approach to realism about QFTs developed independently by Porter Williams (Williams, 2016). I am grateful to

7.2 Formulating Scientific Realism

Many philosophers of science, and philosophers of physics in particular, take a broadly realist view of science. Despite the popularity of scientific realism however, pinning down what the position actually amounts to continues to be a matter of considerable controversy. Indeed, many of the most fierce disputes in the epistemology of science in recent years have been between those defending different conceptions of realism rather than between realists and anti-realists (though anti-realist concerns continue to play an important role in the dialectic, of course, as different versions of realism respond to them in different ways). This section discusses the overarching structure of this debate and puts forward my own take on the problem of formulating scientific realism.

We can think of the basic problem here as a dilemma triggered by the untenability of a naive statement of the realist position. On this reading, scientific realism is simply the claim that our best confirmed scientific theories are true. Adopting the orthodox view of the semantic content of theories embodied in the standard account, this means that to be realist about a theory is to place the actual world among those which are possible according to the theory.

It quickly becomes obvious that construed in this way realism is a highly implausible doctrine. One often cited reason for doubting that even our best scientific theories are true simpliciter is the challenge posed by the historical record. Laudan's (1981) famous pessimistic meta-induction trades on the fact that many defunct theories were successful in their day but turned out to be false, and while realists have resisted the claim that this necessarily leads to anti-realism they have typically admitted that the history of science gives us grounds to doubt that current scientific theories are right in all their details. Less discussed in the traditional realism debate, but equally important for our purposes, are the various indication within contemporary science itself, and theoretical physics in particular, that our present theories are strictly false. Most dramatically, our most fundamental physical theories, QFT and general relativity, furnish mutually inconsistent pictures of reality and there are a complex array of theoretical arguments which suggest that neither can completely characterise gravitational phenomena.[3] There are also internal issues within both of these frameworks that physicists take to point to their incompleteness: the naturalness problem with the standard model and the issue of space-time singularities in general relativity

Porter for sharing this work with me.

[3]Overviews of the conflict between QFT and general relativity can be found in Callender and Hugget (2001, 1-33) and Butterfield and Isham (2001)—I will have more to say about this issue and the search for a theory of quantum gravity in §7.5.

being two prominent examples.[4] Whatever one thinks about the debate surrounding the formulation of QFT then, being a realist about this theory clearly cannot simply mean taking it to exactly describe reality.

If this naive formulation is no good what should the would-be realist replace it with? What seems to be needed is an epistemic attitude towards suitably successful theories which is stronger than anti-realism, in that it takes them to latch onto unobservable features of the world in one way or another, but weaker than accepting their content wholesale. The obvious suggestion is that the realist should take suitably successful scientific theories to be approximately true, and this has indeed become the standard way of outlining the realist position in its broad brush strokes.[5] Ultimately though this answer only postpones the problem, as the notion of approximate truth is notoriously nebulous. If the realist position is to have any substance it seems reasonably to expect a more precise account of the way that scientific theories are purported to relate to the world.

One response to this challenge has been to attempt to formalise the notion of approximate truth, and in particular the idea that the truthlikeness or verisimilitude of a scientific theory is a matter of degree. Starting with Popper's (1972) famously flawed account of verisimilitude however, this project has been beset with technical difficulties and no approach has achieved anything like widespread acceptance. The best developed framework that have been put forward in this literature proceeds by defining a metric on the space of possible worlds (Oddie, 1986; Niiniluoto, 1987). Adopting a possible world semantics for scientific theories we can then give a precise, and in principle quantitative, meaning to the verisimilitude of a theory in terms of the closeness of its possible worlds to actuality. As might be expected however, there are many ways to set up this kind of framework, and the details quickly become contentious. One apparently worrying feature of the accounts which have been put forward along these lines is that this formal notion of approximate truth ends up being a language dependent one—the same theory formulated in different formal languages will, in general, be assigned different degrees of verisimilitude (Miller, 1976).[6]

[4] Another issue which pushes against the naive formulation of realism is the ubiquity of idealisation in scientific theorising. I will not discuss this point in depth here, except to say that I think that an important lesson of my analysis of the QFT programme is that realists need to take this aspect of scientific practice seriously, even in the context of fundamental physics.

[5] Though this is the basic conception of realism which most contemporary defenders of the doctrine take as a starting point it is not the only possible response to the failure of the naive formulation. Van Fraassen tends to characterise realism as a thesis about the basic aim of science, namely that it seeks truth, which is compatible with admitting that current theories do not meet this standard. I will have more to say about this axiological reading of realism in §7.5.

[6] Some proponents of the truthlikeness approach claim that the language dependence of these definitions is not a disaster and is in fact consistent with our pre-theoretic intuitions about approximate truth— see Niiniluoto (1999) and Psillos (1999) for contrasting perspectives.

Beside these internal issues with the verisimilitude programme however there is a more fundamental problem with using this kind of definition of approximate truth as the basis of a formulation of scientific realism. Notice that this approach to approximate truth is similar in spirit, if not in all its details, to Belot's (1998) approach to theory-world relations which I criticised in §2.3. The idea there was that we should broaden the methodological framework of the philosophy of physics by adding judgements about the relative closeness to actuality of possible worlds produced by different semantic interpretations of a theory. My basic worry with the vermisimilitude programme, as I understand it, is the same objection I raised against Belot's proposal. On the possible world approach we have been discussing verisimilitude is a property of the entire content of a theory. But, as I have been urging throughout this thesis, our understanding of particular scientific theories invariably involves more fine grained judgements about which parts of its content get things right. To return to an example I used in chapter 2, there is a detailed story to tell about which aspects of Newtonian gravitational theory are preserved in models of general relativity which should inform the attitude that the realist ought to take to this theory. The claim that Newtonian gravity worlds are close to ours, or indeed that the theory can be assigned a verisimilitude of, say 2/3, does not capture this complexity. A viable formulation of scientific realism, I suggest, needs to accommodate the partial nature of a theory's representational success, and extant formal approaches to verisimilitude do not seem to be up to the task.

In fact, this point seems to have been widely accepted by recent advocates of scientific realism. The orthodox realist response to the pessimistic meta-induction, in schematic terms, has been to claim that some aspects of successful scientific theories are preserved through theory change and this is what underlies their empirical success—this, in a nutshell, is the "divide et impera" move briefly mentioned in chapter 2. For this and other reasons, the idea that the realist should take a selective view of successful scientific theories, accepting some parts of their content while eschewing others, has become the dominant way of thinking about how the position should be stated. On this general approach to the formulation problem, the challenge to say something more precise about the sense that scientific theories are taken to approximately true is answered, at least in part, by identifying a subset of their claims about the unobservable which hit their mark—call formulations which fit this mold selective realisms. Accepting this much, the question is how do we identify the parts of a scientific theory which the realist should commit themselves to? This issue has been the focus of much debate over the past few decades, with myriad selective formulations of realism being put forward and defended against their rivals in the literature. Behind these disagreements over how the selective strategy should be developed in its details however, one can often detect an underlying assumption about the form that an acceptable account should take which has led defenders of realism astray in my view. It has typically been thought that what is needed here is a completely general characterisation of the belief worthy content of scientific theories which applies to all suitably successful theories—call formulations of this type global selective realisms.[7]

One brand of realism which has sometimes been understood in these terms is structural realism, and in particular Worrall's (1989) epistemic version of this doctrine. This formulation of realism was originally motivated by considerations of revolutionary theory change in physics, the key observation being that, while theoretical entites are often dropped in these episodes, structural continuities invariably exist between earlier theories and their successors. The transition from Newtonian gravity to general relativity seems to fit this account in its broad outline; the gravitational field is not retained in relativistic gravity theory, but mathematical connections exist between Newtonian and relativistic models which transcend their shared empirical content. Worrall, and others, take these cases to have wide-reaching implications for the kind of attitude which we ought to take towards scientific theories, the basic claim being that it is the 'structure' of a theory, or its structural claims, which should to be the locus of realist commitment.[8] But what is structure exactly? And here lies the rub, for if structuralism is understood as a global form of selective realism then what is presumably needed here is a definition of structure that captures what is preserved in *all* instances of theory change. Attempts to provide a precise characterisation of the structural content of a theory, using resources like Ramsey sentences, have arguably failed to lived up to this lofty ambition.

Much the same can be said of other formulations of realism which are global in spirit: no general recipe for demarcating the belief worthy content of successful theories have achieved anything close to general acceptance. One might conclude from this that the task is a difficult one and more work is needed to find the right version of global selective realism. In my view, however, seeking such a thing in the first place was a mistake. It is not just that theory-world relations are partial in nature, they are also local, in the sense that the kind of representational success which can reasonably be ascribed to a theory is sensative to the peculiarities of its structure, the surrounding scientific context, and perhaps even the specific

[7]Saatsi (2015) calls this idea "recipe realism" and my criticisms of this approach to the formulation problem follow his lead.

[8]Note that even putting aside the question of how to define 'structure' this strategy does not, in itself, solve the problem of characterising the sense in which our theories are supposed to be approximately true. Structuralist still talk about the structure of superseded theories being 'approximately retained' in their successors and there remains a challenge with making this precise (Saatsi, 2009, 256). The most developed framework which structural realists have appealed to here is da Costa and French's (2003) account of partial truth.

kind of evidential support it enjoys. One motivation for this view is simply the evident diversity of the sciences. Prima facie, scientific theories relate to the world in multifarious ways. As I mentioned, Worrall's version of structural realism is based on observations about particular instances of theory change in physics, but it is far from obvious that any lessons drawn in this context should carry over to theories in chemistry, biology and the social sciences. Physics alone provides plenty of grounds to be pessimistic about the idea that the representational status of successful theories can stated in a unified way. Can the complex stories about what is preserved in the transition from Newtonian gravity to general relativity and, say, thermodynamics to statistical mechanics, both be captured by a single definition of 'structure', or, for that matter, any general schema?[9] The foundational work which has gone into clarifying the relevant intertheoretic relations in each of these cases is highly theory specific in nature, and in my view the study of these details ought to inform the kind of representational success the realist attributes to these theories.

A more fundamental point which cuts against a global formulation of the realist's epistemic commitments comes from reconsidering the explanationist intuitions which typically motivate a realist stance towards science in the first place. The no-miracles argument is sometimes understood as an inference about science as a whole, which might seem to favour a global conception of realism; the rough idea being that, in order to explain the collective empirical successes of mature science (or the scientific method) we need to take it to be a reliable guide to the truth.[10] Whatever one thinks about this global formulation of the argument however, the same kind of explanationist considerations are typically also taken to have force at the level of particular scientific theories. We can ask, for instance, why it is that Newtonian gravity is empirically adequate in the domain that it is, and a philosopher with realist intuitions will want an answer which turns on how the theory latches onto the world. It is far from clear, however, that maximally general characterisations of the approximate truth of scientific theories, along either the formal or global selectivist lines I have just sketched, actually furnishes explanations of why particular theories are successful. In what sense does the fact that Newtonian gravity has a verimisilitude of 2/3, or gets the structure of the world right, where this is explicated in schematic terms, explain its success? Surely the relevant explanation of why Newtonian gravity is successful turns on which *specific* features of the world the theory gets right and how this underlies the *specific* predictive successes which can be obtained from the theory. The explanationist

[9]This is not to say that what is preserved in each case is not 'structural' if this term is construed sufficiently broadly.

[10]Boyd (1985) and Psillos (1999) are well known defences of global formulations of the no-miracles argument. Frost-Arnold (2010) and Fitzpatrick (2013) are recent critiques of this sort of defence of realism.

defence of realism then arguably leads us away from a global statement of the doctrine and towards the details of particular theories.

Where does this leave us with respect to the issue we started out with: the question of what the realist should say about the sense in which they take successful scientific theories to be approximately true. The moral of the discussion so far, I think, is that this challenge cannot be answered in one fell swoop. Accepting a broadly selective approach to the formulation problem, my suggestion is that, instead of making a blanket claim about which aspects of our theories get things right, the realist should address the question of which parts of their content faithfully represent on a theory-by-theory basis. The idea that the approximate truth of particular scientific theories should be spelled out in local terms has recently been advocated by Barrett (2008) and Saatsi (2016b). On this view, the question of which parts of Newtonian gravity get things right should be addressed via a close study of the theory itself, and its surrounding theoretical context. In particular, Barrett and Saatsi suggest that it is largely the theory specific story about how Newtonian gravity's content is embedded within general relativity which guides the kind of representational success that should assigned to the theory. Since general relativity is also strictly false this cannot be the whole story about how Newtonian gravity theory relates to the world, but it does at least constitute a precisification of the sense in which it latches onto the unobservable, and this is arguably the best we can hope to do in our present epistemic situation. The claim would be that the approximate truth of other scientific theories should be fleshed out along similarly local lines.

Combining this approach to theory-world relations with the core no-miracles intuition yields the following statement of the realist's position: scientific theories are empirically successful (at least for the most part) because of the way they latch onto unobservable aspects of the world, but which aspects of their content gets things right, and how this furnishes an explanation of their empirical accuracy, is a matter to be thrashed out in the context of particular scientific theories—call this formulation local selective realism. Thus construed realism is still a general attitude towards the whole of science, but the detailed truth claims which it engenders are cashed out in piecemeal terms. What makes a realist a realist on this view is not a global claim about what scientific theories get right, but rather a commitment to a global explanatory thesis about why the mature sciences are empirically successful—this marks an important difference between scientific realism and other sorts of philosophical realism, like moral realism and realism about universals.[11]

[11] This formulation of realism is consonant with a growing tendency to construe and defend realism in more local terms in recent years. Some papers which advocate a local approach to the realist debate, for a variety of different reasons, are Magnus and Callender (2004),

Fleshing out and defending this formulation of scientific realism in detail is not my project here, but I will touch on two potential objections by way of clarification (the second issue in particular will be discussed further in the next section when I return to the case of QFT). One worry is that going local in this way leaves the philosopher of science with nothing to do. If questions about the approximate truth of scientific theories are answered by the relevant science won't we simply be reciting first-order scientific evidence and arguments? Construed locally realism might seem to collapse into a something close to the philosophical quietism of Fine's natural ontological attitude (Fine, 1986). This worry is misplaced however. For one thing, the core thesis that a scientific theory's representational success explains its empirical success is clearly a second-order, and distinctively philosophical, claim about science, which needs to be defended against anti-realist concerns.[12] Moreover, the local realist stance instigates philosophical engagement with the relevant science. Making a theory's representational status precise will involve bringing general philosophical approaches to representation and explanation to bear on the theory, as well as considering how anti-realist challenges play out in its local context. The localist does not deny the relevance of the general philosophical issues discussed in the traditional realism debate then; the claim is simply that theory specific factors also play an important role in spelling out the realist's epistemic commitments.

A perhaps more troubling objection however, is that this formulation of realism ends up being too weak. Without a systematic account of the parts of our theories which accurately represent there is a worry that we end up with, at best, a very attenuated realist stance towards our current best science. Take a theory like general relativity, for instance. Realists of all stripes will want to say that the theory is latching onto the world in a way which accounts for its impressive empirical success, but which aspects of the theory should we take to be getting things right in this case? A global selective realist will have an answer here; the global structuralist can say they believe in the structural claims of the theory, whatever this is taken to mean exactly. But what can the local selective realist, who eschews this kind of formulaic response, say? The cases which Barret (2008) and Saatsi (2016b) point to as exemplars for the way that the approximate truth of theories can be cashed out locally invariably turn on their embedding within more fundamental theories. Consequently, realism about our current best theories might seem to end up amounting to little more than a promissory note: the local selective realist claims that general relativity latches onto reality in a way which explains its success, but exactly how we cannot yet say. Just as Newton had no idea which parts of his theory of gravity would be preserved in contemporary

Fitzpatrick (2013), Ruetsche (2014), Saatsi (2009, 2015), and Asay (2016).

[12]See Saatsi (2015, 12) for more on this point.

gravitational physics, we seem to be in a similar epistemic situation with respect to current fundamental physics.

This kind of worry is the essence of Stanford's (2003, 2006) so-called "trust argument". According to Stanford the central problem facing the realist is the need to tell us what current theories get right about the world now, and not merely in retrospect. I suspect that the perceived need to respond to this challenge is one reason why defenders of realism have sought to generalise from a few instances of theory change to a global claim about what is preserved in all instances of theory change. An alternative reaction to this line of attack however is to simply refuse the bait, and deny that the realist needs to state their epistemic commitments prospectively. After all, the claim about general relativity above, though schematic, seems to be a prima facie viable epistemic stance towards the theory that clearly goes beyond what the anti-realist would say—they would, of course, remain entirely agnostic about general relativity's claims about the unobservable and the deny the need for any kind of explanation of its empirical success. Saatsi (2016b) simply bites the bullet here and calls the resulting version of realism "minimal realism" in recognition of the fact that some intuitions about what a realist attitude towards current scientific theories amounts to are not necessarily borne out by this approach. While I think this response is perfectly defensible, I will suggest in the next section that the statement of the problem itself is somewhat overblown. In the case of at least some fundamental physical theories, we can make progress in identifying which parts of the theory faithfully represent without appealing to any global recipes or making guesses about future physics—my analysis of the QFT programme is the proof.

7.3 The Renormalisation Group as a Weapon for the Realist

I now return to the case of QFT and consider how the ideas put forward in previous chapters fit into the broader dialectic just outlined. My claim will be that the view of QFT I have been elaborating is a genuinely realist one—contra Butterfield and Bouatta, taking cutoff QFT models representationally seriously actually fits well with orthodox ideas about how realism should be characterised. But I will also suggest that my analysis of the QFT programme provides an instructive case study for the local selective formulation of realism. The role that the renormalisation group plays in identifying which features of QFT models ought to be taken to successfully represent exemplifies how resources found in a particular theoretical context can be crucial in articulating the relationship a scientific theory bears to the world. Furthermore, it demonstrates that, in at least some cases, we can make progress in precisfying the kind of approximate truth which can be attributed to our most fundamental theories in localist terms. The worries raised in the previous section were premature then; giving up on the ambition of a global characterisations of the realists commitments does not necessarily result in a problematically weak epistemic stance towards the forefront of science.

From a general philosophy of science perspective, the central claims of the last four chapters can be seen as flowing from an application of the no-miracles intuition to the case of QFT. How do we account for the spectacular predictive successes of the QFT programme? I suggested in chapter 3 that this explanatory challenge is especially pressing in this case because standard perturbative calculations, which are the most important source of empirical results in particle physics phenomenology, are prima facie difficult to make sense of in either realist or anti-realist terms. I argued in chapter 4 however that we can, in fact, provide a physical explanation of their success. In brief, the claim was that renormalised perturbative approximations work because many of the properties of QFT models are insensitive to the precise values of the cutoffs imposed at large and small momentum, a fact which is most systematically revealed by the renormalisation group analysis of these systems. One upshot of this account is that the empirical success of perturbation theory accrue to cutoff models, and non-perturbative numerical simulations, which constitute the other key source of successful empirical predictions in contemporary particle physics phenomenology, are also more directly attributable to these systems. This does not mean that we should accept the content of cutoff QFT models wholesale however because only some aspects of these systems really contribute to the relevant predictive successes. In particular, I argued that the long and short distance structure of QFT models is irrelevant to the empirical results we are trying to account for. Modulo worries about the details of this account discussed in chapters 5 and 6, the upshot is that to account for the extant predictive successes of the QFT programme, we need to take suitably coarse-grained properties of QFT models to latch onto the world, while many of their short and long scale features should be viewed as idealisations.

Looking back at the discussion of the previous section we see that this actually fits nicely with mainstream ideas about what a realist stance towards a scientific theory should look like. Putting aside the issue of global vs local construals of realism for the moment, we saw that most contemporary scientific realists defend a selective version of the doctrine, according to which being a realist about a theory means taking some, but not all, of its content to latch onto reality. This is exactly the kind of stance I have advocated taking towards empirically successful cutoff QFT models: some of their properties, namely those associated with length/momentum scales far away from the cutoffs are taken to faithfully represent while other aspects of its content are not. In particular, the claim that cutoff QFT models misrepresent the small scale structure of the world is not in conflict with a selective formulation of realism. After all, theories like Newtonian gravity and ordinary non-relativistic quantum mechanics certainly get this wrong too, yet I take it most contemporary realists would claim that these theories are approximately true, and do get something right about the way the physical world is despite their representational inaccuracies.

Furthermore, this view of QFT differentiates itself from an anti-realist stance in the two most important respects. First, it makes substantive epistemic commitments to unobservable aspects of reality. Anti-realist of all stripes disavow the claims that scientific theories make about unobservables. The constructive empiricist, for instance, remains agnostic about all a theory's content save its claims about observable matters, where an entity or property posited by a theory is said to be observable, roughly speaking, if it is a possible object of unmediated perception by human agents. As I stressed in chapter 4 however, my claim that QFT models get coarse-grained features of the world right should not be confused with the claim that it gets observable features of the world right. On my account, it is the correlation functions of a QFT model over relatively long length scales which we ought to take representationally seriously, and these quantities are not observable in this sense, or directly measured in collider experiments. The second feature of my account which distinguishes it from an anti-realist position is that it is based on an explanation of the success of the empirical predictions of the QFT programme in terms of the way the world is. The anti-realist, by contrast, takes the empirical success of a theory to be something primitive which does not stand in need of explanation (or at least not that kind of explanation).

In fact, I think that the best way to understand the account of the representational success of QFT models I have put forward is in terms of the local formulation of scientific realism I endorsed in the previous section. The key idea with this approach, remember, was that, instead of trying to give a completely general characterisation of the sense in which empirically successful scientific theories are approximately true, the realist should allow that the representational status that can reasonably be attributed to a theory is sensitive to theory specific factors. On the picture I have been sketching here, the renormalisation group framework plays a particularly important role in articulating the kind of representational success which can reasonably be attributed to QFT models. One way that the renormalisation group helps to articulate the approximate truth of cutoff QFT models is by identifying features of these theories which should not be taken to accurately describe reality. In general, the renormalisation group teaches us that we should not trust the claims that current QFT models make about very short length scale physics, but renormalisation group methods also allow us to identify more specific defects, produced by a particular implementation of the ultraviolet cutoff which should not be taken physically seriously. Chapter 5's discussion of Poincare symmetry in cutoff systems can be seen as a detailed case study of this point: we saw that the renormalisation group allows us to systematically identify the Poincare violating contributions to physical quantities induced by the ultraviolet cutoff which we have no reason to believe accurately represent.[13]

But as well as clarifying the relationship cutoff QFT models bear to reality in this negative sense, the renormalisation group arguably also gives us grounds to be confident in some aspects of empirically successful QFT models. I have thus far presented this point in explanationist terms: I claimed, for instance, that correlation functions over relatively coarse-grained length scales are modally connected to empirically observed scattering cross sections in the appropriate way, and consequently need to be taken representational seriously. But we can also think about how the renormalisation group helps the realist here from the perspective of theory change and intertheoretic relations.[14] We know that the standard model of particle physics is not an exact description of the world at all energy scales, and we do not really know what lies at higher energies or even at what energy new physics might kick in beyond the energies which have so far been probed at the LHC. Still, in demonstrating that the large scale properties of a QFT model are insensitive to the high energy dynamics, the renormalisation group is also telling us that these properties are essentially independent of the details of future physical theories which describe the dynamics of currently inaccessible high energy degrees of freedom. Thus the renormalisation group gives us a way of identifying properties of our present theories which will be embedded within future theories, in one way or another.[15]

As I see it, this is all good news for the local selective formulation of realism. On the one hand, the role that the renormalisation group plays in explicating

[13]Williams (2016) gives a similar characterisation of the role of the renormalisation group in identifying which features of cutoff QFT systems are the appropriate objects of realist commitment, and examines some other instances of 'lattice defects' which are illuminated by renormalisation group methods, such as the appearance of 'mirror fermion' degrees of freedom when fermionic fields are placed on a lattice.

[14]We need to be appropriately modest here though. While the renormalisation group helps to spell out the intertheoretic relations between QFT models at different scales other resources are also important. The notion of spontaneous symmetry breaking discussed in the previous chapter also plays an important role in explicating how low energy effective theories emerge from a more fundamental high energy description, for instance. The claim here is not that the renormalisation group tells us everything we might want to know about how QFT models relate to the world, but that it helps.

[15]The idea that renormalisation theory should be understood as isolating features of our present theories which are robust under future theory change was advocated in a prescient paper by Alexander Rueger (1990).

the relationship that realistic QFT models bear to the world exemplifies the local realist's key claim: that resources found in a particular theoretical context can play an important role in articulating the realist's epistemic commitments with respect to a particular theory. While it may be the case that the renormalisation group, and the ideas I have sketched here, can also help to explicate the sense in which theories in other branches of physics are approximately true (especially in condensed matter physics where very similar methods are employed) there is no claim here that the renormalisation group furnishes anything like a general approach to theory-world relations. The geometric resources that go into spelling out how Newtonian gravity theory relates to general relativity, for instance, are quite different in character from renormalisation group methods, and I suspect that the sense in which space-time theories like general relativity turn out to be approximately true will also differ in important respects from the story just sketched about QFT models.

The discussion also throws new light on the central objection to local selective realism raised in the previous section. The worry was that, without a general recipe for identifying the belief worthy content of any scientific theory, the local realist will only be able to make a highly tentative and provisional claim about the representational success of our current most fundamental theories. I take my discussion here to show that this need not be the case. The renormalisation group framework provides us with resources for identifying which aspects of empirically successful QFT models get things right without knowing what the quantum gravity programme has in store for us. The upshot seems to be that we are in a better epistemic situation with respect to the standard model of particle physics than Newton was with respect to Newtonian gravity. This is, in large part, due to the wealth of modal information about which features of QFT models can be varied without affecting their large scale physics provided by the renormalisation group. Consequently, the renormalisation group framework opens up the possibility of a more full bodied realist stance towards high energy physics, than Saatsi's (2016b) discussion of "minimal realism" would seem to suggest. At least part of the response to Stanford's trust argument then is that the demand for a prospective identification of the features of our theories which will be preserved through theory change can be met directly in some local scientific contexts.

If I am right, and the renormalisation group, and cutoff QFT models, help to furnish a realist view of high energy physics, then why have so many philosophers of physics associated them with an anti-realist attitude? The next two sections offer a diagnosis of this mistake: I first tackle Fraser's concerns about underdetermination before turning to a different kind of worry generated by the fact that, on the view I have been espousing QFT does not provide a complete or truly fundamental description of reality.

7.4 Underdetermination in the Quantum Field Theory Programme

Doreen Fraser's central argument for advocating an intrumentalist reading of cutoff QFTs is based on the putative existence of an underdetermination problem within the QFT programme. As we saw in §5.3, Fraser (2009) claims that the choice between cutoff and continuum formulations of QFT is underdetermined by empirical evidence, but argues that the theoretical virtues of the axiomatic approach to QFT break the underdetermination in favour of continuum models. I maintained that the argument for this positive claim does not go through, but whatever one thinks about this issue the presence of an underdetermination scenario itself might seem to be enough to threaten the realist view of QFT I have been developing here. Fraser (2011) suggests that the results of renormalisation group analysis further amplify this challenge. The renormalisation group framework shows how theories which are radically different in their short scale dynamics can nevertheless have very similar large scale behaviour. Consequently, the empirical success of a QFT model over a limited range of energy scales is compatible with many distinct theories describing the higher energy degrees of freedom beyond this domain.

The first step in assessing this objection is to get clear on what the problem is supposed to be. Prima facie the relationship between cutoff and continuum QFT models which Fraser points to does not resemble the kind of underdetermination scenario which has traditionally been focused on in the realism debate. It has usually been assumed that the really problematic situation for the realist would be one in which two rival theories are empirically equivalent, in the sense that they cannot be distinguished by any possible observations. But cutoff and continuum QFT models do not appear to be empirically equivalent. We have seen that, if a continuum QFT model exists, it will assign values to all observable quantities which differ from its cutoff counterpart; in many contexts we can estimate the relevant error as being bound by some power of E/Λ, where Λ is the ultraviolet cutoff and E is the energy scale at which the relevant correlation functions are being probed. This means that if the world actually has the structure of a cutoff QFT model—if space really does become grainy and discrete at some very small length, say—then this should, in principle, be empirically detectable. As we approach the physical cutoff scale we will begin to see effects which deviate from the expected continuum results. What we seem to have then is a case of transient or weak underdetermination, which has often been thought to be either unproblematic for the realist or reducible to more general philosophical challenges, like the problem of induction (Ladyman, 2001, 163-167).

There are two ways that Fraser, or an anti-realist who wants to co-opt her argument, might shore up their case here however. On the one hand, it might be pointed out that, while cutoff and continuum models are empirically distinguishable given the ability to probe arbitrarily high energies, assuming, as seems fairly plausible, that there is an in principle limit to the energy scales we can ever experimentally access we will never be able to tell whether there is a cutoff beyond this threshold or a continua of high energy degrees of freedom. In this scenario the cutoff and continuum hypothesis really will be underdetermined by all possible evidence, at least as far as the QFT programme is concerned. I suspect that it is something along these lines which Fraser has in mind. Another response is to simply deny that empirical equivalences is a necessary condition to get a pernicious underdetermination problem going. Stanford (2001, 2006) has argued that all that is needed is that the choice between two rival theories cannot be decided by *current* empirical evidence for the realist to be in trouble.

On either of these precisifications, Fraser's claim that the renormalisation group framework exacerbates the situation seems to be apt. From a renormalisation group perspective the indistinguishability of cutoff and continuum versions of a QFT models arises because they are drawn to the same region of the space of possible theories under the action of the renormalisation group transformation. But this is just one manifestation of the more general decoupling of high and low energy degrees of freedom captured by this framework. For one thing, different ways of implementing a high momentum cutoff will also be empirically indiscernible in the same way: whether the ultraviolet cutoff is imposed sharply or smoothly, and the precise value it takes, will not make a detectable difference to the low energy phenomenology. But we have also seen that theories with different interactions can exhibit essentially the same low energy behaviour. In particular, the effects of weak non-renormalisable interactions will invariably become completely undetectable at sufficiently low energies. As a consequence, there will be a great number of dynamically distinct QFT models which are empirically indistinguishable in a limited energy regime. If the kind of underdetermination revealed by the renormalisation group framework is a problem for the realist it is a very wide reaching one then, which goes well beyond the indistinguishability of cutoff and continuum QFT models that Fraser focuses on.

In fact, there is a further puzzle about this particular underdetermination claim which we discussed already in §5.3: we do not actually have continuum formulations of the empirically successful gauge theories which make up the standard model. It is this point which Wallace (2011) focuses on in his response to Fraser's argument. Wallace apparently concedes that there would be a problem here if models of the Wightman or Haag-Kastler axioms has been constructed in four dimensional Minkowksi space-time but stresses that, as it stands, the cutoff version of theories like QED and QCD are the only well defined mathematical systems we have at our disposal. There simply is no underdetermination scenario between a cutoff and continuum version of the standard model because, as far as we know, the latter does not exist.

But Fraser (2011, 133) argues that the fact that continuum formulations of realistic QFTs are in the offing is sufficient to undermine a realist stance towards the cutoff standard model. This claim is independent of her, I think, less plausible suggestion that we should take a realist stance towards the constructive field theory programme in its incomplete state. One way we might make this thought precise would be suggest that, since constructive field theorists are working towards constructing these continuum theories, and we have at least some idea of how they will be related to extant cutoff models from successes in lower dimensional space-times, there is contextual evidence that we have not yet exhausted the space of alternative theories which are supported by the relevant evidence, along the lines of Stanford's "unconceivable alternatives" formulation of the underdetermination argument against realism. Whether we find this compelling or not however, Wallace's response does not address the more general kind of underdetermination problem which seems to be ushered in by the renormalisation group framework. While we don't know whether a continuum version of QED exists it is straightforward to construct a cutoff model which is empirically indistinguishable from cutoff QED at low energies by adding non-renormalisable interactions to the Lagrangian. Besides the more ethereal underdetermination scenario Fraser bases her argument on then there are plenty of concrete cases which could also be used to push against a realist stance towards cutoff formulations of current QFT models.

There is, I think, a more general reason why the kind of underdetermination engendered by the renormalisation group is not as troubling for the realist as it seems however, which, once again, does not hang on the question of whether continuum formulations of realistic QFT models exist. To motivate the point I have in mind here it will help to reconsider why (and when) underdetermination is a problem for the realist in the first place. Roughly speaking, the challenge is generated by the thought that, if two theories are equally well supported by empirical evidence then there is no reason to believe one over the other. But if this situation is to motivate skepticism about *all* of the claims that a theory makes about unobservable aspects of reality, as the anti-realist wants, then the pair of indistinguishable theories must say completely different things about the world.[16] Equally well supported theories become much less troubling for the re-

[16] This ideal of empirically equivalent theories with completely disjoint content is often realised in the kind of algorthmic approach to generating underdetermined theories advocated in Kukla (1996). As Stanford (2001) points out however, this kind of underdetermination scenario

alist if their content is not disjoint in this sense. On the one hand, this point is what drives the common realist response when faced with an apparent case of underdetermination: identify the putative rivals as equivalent formulations of the same underlying theory.[17] If wave and matrix formulations of quantum mechanics are just different ways of expressing the same claims about the world, for instance, there is no dilemma about choosing between them which threatens realism. Another situation is also possible however; it may be that two indiscernible theories are genuinely inequivalent but there is nevertheless a great deal of overlap in their descriptive content.

Consider a classic example: two Newtonian gravity models $TN(0)$ and $TN(v)$, formulated on Newtonian space-time, which differ with respect to the absolute velocity of the universe (van Fraassen, 1980, 44-46). While these two theories are inequivalent by stipulation, they also share a great deal of content. In fact, they agree on everything save the absolute velocity of the universe—they assign the same values to gravitational forces acting on massive bodies, prescribe the same relative motions between them, and so on. Consequently, it seems perfectly possible for both $TN(0)$ and $TN(v)$ to be approximately true, which as we have seen is the only claim that a sane realist wants to make about well confirmed theories. Even Stanford, a prominent defender of underdetermination challenges to scientific realism, concedes this point (Stanford, 2001). The natural moral to draw from this case does not seem to be skepticism about all of Newtonian gravity's claims about the world, rather it mitigates against putting stock in the claims about the absolute velocity of the universe over which these two models disagree. Underdetermination of theory by evidence is not necessarily a bad thing for the realist then; it can act as a guide to locating which features of a model should be taken to faithfully represent. The way that this kind of partial underdetermination informs the realist's assessment of a theory's representational status is closely related the idea that we should not believe in the parts of a scientific theory which do not make a difference to its empirical success.

I claim that the underdetermination between QFT models revealed by the renormalisation group is like the absolute velocity case in this regard. QFT models which flow to the same region of the space of possible theories under the action of the renormalisation group transformation share a great deal of descriptive content in addition to their empirical predictions. Specifically, they agree about the values of correlation functions over relatively coarse-grained length scales, which, as I have stressed, are not directly observable quantities. What the underdetermi-

seems to simply be a restatement of Cartersian skeptical arguments.

[17]Of course, this strategy brings us back to the problem, discussed in §2.4, about how to characterise theoretical equivalence. See Magnus (2004) for a discussion of how the debate surrounding theoretical equivalence bears on the assessment of underdetermination arguments.

nation between QFT models described above motivates then is not anti-realism but rather a selective skepticism about the features which these models disagree on—namely the very small (and very long) distance structure. But this is exactly the attitude towards QFT models I have been defending! All this talk about underdetermination is really just another way of stating the core claim of this thesis: that the small and large scale structure of cutoff QFT models (and continuum models if they exist) do not contribute to the extant empirical success of the QFT programme and should not be taken to describe the world. The fact that a diverse range of currently unconfirmed high energy theories display the same low energy behaviour as the standard model should bolster our confidence in it as an accurate description of (relatively) low energy physics, not undermine it.

In sum, the kind of underdetermination we find in the QFT programme has the opposite moral to the one Fraser wants to draw in my view. It actually supports the realist reading of cutoff QFTs I have been elaborating.

7.5 Future Physics and Fundamentality

There is, I think, another intuition which drives some philosophers to associate cutoff QFTs with anti-realism. The unease springs from the fact that, on the view I have been defending, the successes of the QFT programme only licenses claims about the world in limited energy regimes and on relatively coarse-grained length scales: it does not provide us with a complete or truly fundamental picture of what the world is like. Isn't this what fundamental physics is supposed to do however, and isn't this what the scientific realist is after? It is this feature of the effective field theory view, I suspect, which leads Butterfield and Boautta (2014) to describe it as "opportunistic and instrumentalist".[18] Thus stated the problem is somewhat nebulous however. I will explore two ways of pinning down what the worry is supposed to be and argue that, while each raises important issues about the QFT programme that deserve further discussion, neither undermines the realist credentials of the view I have been defending.

One potential reason for thinking that viewing QFT models as effective theories is embarrassing for the realist is the perception that this runs counter to the idea that science is about finding the whole truth about the world. While the discussion of §7.2 focused on the attitude which the realist takes towards the current products of scientific theorising, there is another strand of thinking in the

[18] Another feature of Butterfield and Bouatta's (2014) discussion which perhaps leads them to this conclusion is the assumption that cutoff QFTs are purely heuristic and badly defined, while only continuum QFT models are sufficient rigorous to be called theories in the philosopher's sense. As I have made clear from the outset I do not agree with this characterisation of the situation.

realism debate which focuses on the goals and intended outcomes of the scientific enterprise. Van Fraassen, in particular, tends to characterise the difference between the realist and the anti-realist in terms of the fundamental aim of science rather than its current epistemic achievements: the realist takes it to be true theories, while the constructive empiricist seeks only empirically adequate ones (van Fraassen, 1980, 6-7). Imposing a cutoff on the momentum, and taking the resulting theory to be descriptively accurate in a limited domain, might seem to amount to giving up on the quest for a theory which exactly describes the world on all scales. This chimes with Fraser's complaint 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" (Fraser, 2009, 563), where "theory X" here refers to a currently unknown theory which exactly describes the world. The intuition here is that the realist should not be content with effective theories; they should be striving for this "theory X".

As I have argued, however, taking current empirically successful QFT models to capture salient, but ultimately incomplete, information about the physical world is on a par with the innocuous claim that Newtonian gravity and non-relativistic quantum mechanics are approximately true. I have not said that other theories cannot do a better job, or that we should not seek such theories—taking cutoff QFTs to provide successful coarse-grained representations of reality does not amount to "resting content" with these representations. If you think that scientific realism should entail an optimistic attitude towards future physics, and the possibility of finding a genuinely complete theory of the physical world, there is nothing in what I have said thus far in this chapter to contradict this.

Having said this, a great deal remains to be said about how the QFT programme as it stands bears on the project of constructing future physical theories, and the problem of quantum gravity in particular. This is another area where there is scope for reconciling Fraser's and Wallace's perspectives on QFT. I am inclined to agree with Fraser that the axiomatic and constructive approaches to QFT may yet have lessons for the project of constructing theories of quantum gravity. Recent developments such as the discovery of the AdS/CFT correspondence and the progress in asymptotic safety approach to quantum gravity suggest that getting a better understanding of the space of possible QFT models may be a crucial step in articulating a viable theory of quantum gravity. I see no reason why philosophical engagement with what is known about continuum QFT models might not also play a role here; just as when philosophers engage with speculative proposals in quantum gravity research we can hope that clarifying conceptual issues internal to the QFT programme will help pave the way for future theoretical progress.[19]

We have to realise however that when it comes to the construction of future theories, we leave behind the question of what we should believe about the world on the basis of present evidence at the heart of the realism debate. Underdetermination is not a problem in this context as we are not examining theories which are candidates for belief. In the context of discovery the more theoretical frameworks at your disposal the merrier; there is no tension in drawing on both the axiomatic and effective field theory traditions in the search for an adequate quantum theory of gravity. It is part and parcel of the decoupling of theories at different scales that there is unlikely to be a very direct connection between what we are licensed to believe about the world on the basis of present evidence and the details of physics at higher energies. To reiterate a claim I have made several times now, I do not think that constructing a continuum formulation of the standard model will license any new claims about the world which are not already captured by a cutoff formulation. But it may provide resources that aid the construction of future theories, thus contributing to the epistemic progress of science in a more indirect way. This is yet another way that engagement with the axiomatic approach to QFT can be justified in a way which is fully compatible with the realist view of cutoff QFTs I have defended in this thesis.

Another reason why viewing QFT models as effective theories may be unattractive to some however, is not so much that they don't capture the whole truth but that the truths they do capture are not fundamental.[20] There seems to be a pervasive, if rarely explicitly discussed, idea in the philosophy of physics that what we should primarily be concerned with is the fundamental ontology posited by physical theories. But it is precisely the fundamental features of QFT models (assuming these can be assimilated to their claims about the smallest length scales) which I have argued are not trustworthy as representations of reality. It is worth pointing out here that this metaphysical notion of fundamentality has not played much of a role in the traditional realism debate, and it certainly has not been viewed as part of the definition of realism that scientific theories capture fundamental truths.[21] The reason, presumably, is that realism is supposed to be

[19] Wallace (2011) suggests considerations of this kind only motivate the "first-order project" of constructing continuum QFT models, the implication being that the philosophy of physics has nothing to offer here. But this verdict seems to be based on the unnecessarily restrictive view that the *only* legitimate question in the philosophy of QFT is what we ought to believe about the world given the extant empirical successes of the standard model. I have tried to make room for internal philosophical questions about physical theories which may well bear on the construction of future theories.

[20] What philosophers mean by 'fundamental' is not an issue I will discuss much here, safe to say that it goes beyond what is meant by the more philosophically neutral term 'fundamental physics'—See Schaffer (2009), Sider (2012) and Barnes (2013) for contrasting views in the metaphysics literature.

[21] Fundamentality talk has increasingly played a role in the debate surrounding ontological

a philosophical thesis about the whole of science, including fields like biology and chemistry, as well as non-fundamental physical theories like Newtonian gravity, which most philosophers do not take to give us information about the fundamental ontological constituents of reality. I suspect, therefore, that the concern about the non-fundamentality of current QFT models is not an issue about scientific realism per se, but is rather triggered by additional assumptions which come into play when discussing high energy physics.

Why have philosophers of physics attached such importance to the fundamental structure of physical theories? One diagnosis of this tendency is that it springs from the influence of the standard account of interpretation. The standard account, remember, was the doctrine that an interpretation of a theory is a specification of the possible worlds in which it is true. It is plausible to think that characterising the worlds in which physical theory is true will involve spelling out their fundamental ontology. I argued in §4.5 that there is no special problem with treating cutoff QFTs in standard account terms—lattice QFTs will be exactly true in worlds in which space-time really has a lattice structure. But I also argued in chapter 2 that, in so far as the standard account is a plausible doctrine, what it provides is a purely semantic notion of interpretation. A standard account interpretation gives us an account of what a theory says, but the epistemic question of what we ought to believe about the world given the success of a theory, which is the concern of the realism debate, is a separate issue. Properly understood, the standard account has no concrete implications for how physical theories relate to the world, and certainly does not imply that their fundamental ontology must get things right.

A more interesting motivation for focusing on the fundamental ontology delineated by a physical theory is an adherence to a specific programme of naturalistic metaphysics. Many contemporary metaphysicians characterise, and in some cases even define, their discipline as the subject which concerns itself with the fundamental structure of reality.[22] If you take this view of metaphysics and also hold that metaphysics should be deeply informed by physics then it is easy to see why you would be disappointed with my claim that the QFT programme as it stands only licenses claims about the non-fundamental—the consequence would seem to be that QFT tells us nothing about metaphysics.

I think we simply have to bite this bullet: we cannot wish away the limitations of our current epistemic situation. This suggests that if we want to do metaphysics based on the QFT programme there are two ways we can go. On the one hand,

structural realism however, see McKenzie (2013).

[22]The idea that metaphysics is primarily concerned with the fundamental can be found in papers like Dorr (2005), Schaffer (2009), and Sider (2012). For an interesting critique of this characterisation of metaphysics see Barnes (2015).

we can engage in speculative metaphysics: we can address questions about the possible worlds in which various kinds of QFT models are perfectly fundamental. I don't see anything objectionable about this kind of work, indeed, on the standard account of interpretation this is closely related to the task of specifying a theory's physical semantics which is clearly an important task for the philosophy of physics. One simply has to be careful not to slip into drawing conclusions about the way the world actually is when working in this mode—in the terminology used in previous chapters, the metaphysical issues being addressed are internal to a particular theoretical framework. On the other hand, if you want to make claims about the metaphysics of the actual world based on the successes of the QFT programme I suggest that, in one way or another, non-fundamental ontology will have to be taken seriously. This may seem to be a revisionary message but it is worth pointing out here that the untenability of the naive formulation of scientific realism on its own seems to have a similar moral for the metaphysics of science. We have good reason to believe that no current theory describes reality at the most fundamental level so any metaphysical conclusions warranted by contemporary science must be about non-fundamental stuff.

I close with a quote from Polchinski (1999) which incapsulates the blunt response which must ultimately be given to the worries discussed in this section. After explicating the key results of the renormalisation group approach to QFT he recounts an imaginary dialogue:

> Q: Doesn't all this mean that quantum field theory, for all its successes, is an approximation that may have little to do with the underlying theory? And isn't renormalization a bad thing, since it implies that we can only probe the high energy theory through a small number of parameters?
>
> A: Nobody ever promised you a rose garden. (Polchinski 1999,10-11)

The decoupling of physics at different scales revealed by the renormalisation group is a mixed blessing: it gives us grounds to be confident in the partial representational success of our present QFT models, but it makes the task of finding a truly fundamental and comprehensive description of the physical world extremely difficult. Still, it is a fact about the way the world is structured not a negociable philosophical doctrine. Continuum formulations of current QFTs would not be any more capable of describing reality at its fundamental level than the cutoff structures we have at our disposal. The view of what we learn about the world from high energy physics elaborated in this thesis is appropriately modest then and realism enough.

Chapter 8

Conclusion

I started this thesis by pointing out a distinctive methodological challenge which QFT seems to pose for the philosopher. Perhaps more than any other branch of physics, the question of what QFT is at the theoretical level has vexed philosophers of physics. My hope is that the foregoing discussion has laid the groundwork for a response to this problem and, therefore, for a more productive interchange between high energy physics and philosophy. I conclude here by summarising my key claims via a comparison with Wallace and Fraser's approaches to the QFT programme and making some suggestions about avenues for future enquiry.

My basic strategy has been to dissolve rather than resolve the apparent dilemma raised by the existence of multiple theoretical approaches to QFT. The heterogeneity of the QFT programme only seems to be problematic if we cling to the idea that the answer to the question 'what is QFT?' should take the form of a precise set of mathematical axioms or structures. I suggested that this expectation is unfounded for quite general reasons. Rather than seeing the issue as one of identifying the 'correct' version of QFT my approach to answering this question has been to acknowledge the existence of the many strands of the QFT programme and start to flesh out the relationship they bear to each other and to broader questions of philosophical interest.

Like Wallace, I have defended the legitimacy and importance of cutoff QFT models as objects of philosophical study, especially when it comes to the epistemic question of what we ought to believe about the world on the basis of high energy physics. The view of cutoff QFTs I have defended here does differ from Wallace's in important respects however. In some ways my position is more modest. I have not claimed that cutoff QFTs are the only systems which the philosophy of QFT should be interested in. In fact, I think that Wallace's (2011) presentations of the axiomatic and renormalisation group approaches to QFTs as rival approaches to addressing the same issues has limited usefulness. While there are commonalities in their origins, in that they were both responding to problems with the original perturbative formalism, each has developed its own distinct agenda and we can make room for both in our philosophy of QFT.

In other respects my stance towards cutoff models is more radical than Wallace's. Wallace consistently emphasises the fact that we do not currently have continuum formulations of realistic QFTs—Wallace (2006, 54) even suggests that "it is only because that goal has not been achieved that it is interesting to explore alternatives [to axiomatic QFT]". In contrast, my central arguments for taking cutoff QFTs seriously have been largely independent of the mathematical difficulties associated with constructing continuum QFT models. For me, the epistemic status of cutoff formulations of QED and QCD will not change much if a proof of the existence of continuum formulations of these theories is obtained. My case for taking cutoff QFTs seriously as representations of reality has focused on what we do know about QFT models, and specifically on the insensitivity of their coarse-grained behaviour to their small and large scale structure, rather than what we do not know. This line of reasoning has led me to make stronger claims about the status of the infrared cutoff than Wallace: I have argued that QFT systems with both ultraviolet and infrared cutoffs are capable of representing all of the features of the world we have any right to believe on the basis of current experimental high energy physics.

As for Fraser, on the face of it our assessments of the formulation problem are fundamentally opposed. Indeed, I have rejected much of the machinery which Fraser uses to set up her presentation of the issue. I have resisted Fraser's claim that QFT should simply be defined as "the theory that best unifies quantum theory and the special theory of relativity" (Fraser's, 2009, 550), because I deny that QFT can be neatly defined in this way at all. Furthermore, presenting the whole issue as one of underdetermination of theory by evidence is, in my view, ultimately unhelpful. In any case, I argued that the kind of underdetermination between QFT models we do find in the QFT programme does not support the morals Fraser wants to draw.

Having said this, I have suggested that we can accommodate some of Fraser's motivations for focusing on the axiomatic approach to QFT. I identified a number of ways that philosophical engagement with the axiomatic QFT may illuminate issues of philosophical import without directly contributing to our picture of what the world is like. On the one hand, there is the question of the extent to which quantum theory is compatible with fundamental relativity, an issue which constructive field theory promises to tell us a great deal. As far as I am concerned this is an important issue in its own right, independently of whether the world actually is fundamentally relativistic, just as the status of determinism in classical mechanics is interesting though we do not live in a classical world. More than this though, I suggested in chapter 6 that study of the axiomatic and constructive approaches to QFT may still play an important epistemic role in the QFT programme in so far as they facilitate the derivation of results which also apply to cutoff QFT systems. Finally, there is scope for philosophical investigation of these approaches to inform the construction of future physical theories; as I sketched in chapter 7, there are multiple ways that learning more about continuum QFT systems, in both formal and conceptual terms, may impact on the quantum gravity programme. In this context however we are approaching axiomatic QFT in an exploratory mood rather than as objects of belief. There is scope to reconcile at least some of Fraser's and Wallace's claims about QFT within my approach then.

How might the arguments of this thesis be taken forward? Many loose ends remain from the discussions of previous chapters. Chapters 3 and 4 were primarily concerned with establishing the foundational respectability and interest of the perturbative and renormalisation group approaches to QFT. Plumbing the depths of these bodies of theory, which are themselves multifaceted and diverse, is another project in itself. There is also plenty of scope to sophisticate and challenge the view of the QFT programme I have developed here. As I emphasised in chapters 5 and 6 more work is needed on the subject of approximate space-time symmetries and spontaneous symmetry breaking in QFT, both of which being crucial issues for defending my claims about the coarse-grained representational success of QFT models. Only time will tell whether the perspective on QFT taken here will survive in the face of further investigation of these issues.

If my core claim about the partial representational veracity of QFT models is accepted however, it seems likely to have knock on implications for many previous debates in the philosophy of QFT. Much of the extant work in this literature has been concerned with the fundamental ontology of QFT, and specifically with the status of particles and fields. Furthermore, the dominant approach to these issues in recent years has been to address them within the algebraic axiomatisation of continuum QFT—Fraser (2008) and Baker's (2009) arguments against particle and field interpretations of QFT respectively being prime examples of this trend. My claim that the fundamental ontological claims of QFT models, and the novel properties afforded by the continuum limit, should not be taken to faithfully represent raises questions about the appropriate conclusions which can be drawn from these arguments. While I certainly do not want to say that the many impressive foundational theorems which have been derived within the algebraic formalism should simply be discarded as philosophically irrelevant, we do, I think, need to reassess the significance of these results. One theme of this thesis has been that philosophers of physics have tended to focus on semantic questions at the expense of epistemic ones. The key point here is that we need to carefully consider what claims about the world are licensed by results obtained in the algebraic framework. Re-examining previous foundational work on QFT from an epistemological stand point is a priority for the philosophy of QFT going forward then.

So much for the technical theory specific debates of the philosophy of physics. My original motivation for focusing on the formulation problem was that it seemed to be a key factor obstructing meaningful interactions between high energy physics and more wide-reaching philosophical issues. What lessons can the general philosopher of science and the metaphysician glean from the discussion of this thesis?

When it comes to the general philosophy of science the central take home message of this thesis has been not to panic. Worries that the QFT programme radically undermines a conventional, broadly realist, philosophy of science are, I think, largely unfounded. In a more constructive direction however, my hope is that my discussion has laid a foundation for further exchange between traditional debates in the philosophy of science and the details of the QFT programme. Chapter 7 claimed that there are morals to be drawn for the articulation of scientific realism and I think there are many other debates which can benefit from engagement with the QFT programme. To name just one, the scientific explanation debate, and specifically the question of how idealisation and abstraction ought to be understood in this context, is crying out to be brought into contact with high energy physics. A recent trend in this literature has been a focus on the notion of explanatory relevance and the idea that eliminating details which don't make a difference to the target explanandum can boost the explanatory power of a model—see especially Strevens (2008) kairetic account of explanation. The way that cutoff QFTs, understood as coarse-grained representations, abstract away information about high energy degrees of freedom which does not impact on low energy physics appears to exemplify this explanatory practice. My hunch is that the QFT programme may provides instructive case studies for honing our understanding of these issues.

What about the metaphysics of science? A key messages of chapter 7 was that if we are interested in addressing ontological questions about the actual world on the basis of the basis of the QFT programme then we will need to take non-fundamental ontology seriously. In my view, this is the way that the debate surrounding the status of particles in QFT should go. It seems clear, whichever approach to the theory you look at, that elementary particles are not fundamental entities in QFT. The key question going forward is how to spell out the kind of derivative status that they enjoy.

If we are to make progress in this direction however, we seem to need a framework for discussing non-fundamental ontology. Wallace has made use of Dennett's (1991) notion of "real patterns" in this context—see, in particular, Wallace (2012) chapter 2. It is notable, however, that this language has not been widely adopted by philosophers of physics, much less analytic metaphysicians. The worry, I suspect, is that while this approach may bolster intuitions about derivate entities and properties it does not appear to do much to make their ontological status precise. There are two ways we might go in search of clarification here. On the one hand, following the localist spirit of chapter 7, we could claim that the sense in which ontological constituents are emergent (in a weak reduction compatible sense) is largely substantiated by the details of the relevant physics. Spelling out the sense in which particles are non-fundamental entities, for instance, will come down to examining the details of how particle-like behaviour emerges in QFT systems.[1] On the other hand, recent work in analytic metaphysics, and specifically the burgeoning literatures on meta-ontology and fundamentality, has thrown up a variety of resources for articulating notions of ontological dependence.[2] Though there are certainly challenges with employing these ideas in the context of physical science recent work, such as McKenzie (2013), suggests it is not a fools errand. To my mind, pursuing both of these approaches in concert is the most promising way forward for metaphysical engagement with the QFT programme, and may lead to fruitful interactions between the philosophy of physics and analytic metaphysics.

The picture of how the different parts of the QFT programme presently hang together put forward in this thesis is ultimately intended in a provisional and explorative spirit. Perhaps more significant than these positive claims is the methodological approach I have advocated and attempted to exemplify. There is no reason for philosophical engagement with the QFT programme to be held hostage by the formulation problem; we can, and should, draw on the various theoretical approaches to QFT in a piecemeal fashion in order to enrich debate in the philosophy of science without getting caught up on the perceived need to provide a universally applicable formal answer to the question 'what is QFT?'.

[1] Bain (2000) and Wallace (2001) can be understood as contributions to this project.
[2] See Lowe and Tahko (2015) for a review of this literature.

References

Anderson, P. W. (1972), “More is Different”, Science, 177 (4047), 393-396.

Appelquist, T., and Carazzone, J. (1975), “Infrared Singularities and Massive Fields”, Physical Review D, 11(10), 2856–2861.

Araki, H. (1999), *Mathematical Theory of Quantum Fields*, Oxford University Press.

Arntzenius, F. (2014), *Space, Time, and Stuff*, Oxford University Press.

Asay, J. (2016), “Going Local: a Defense of Methodological Localism about Scientific Realism”, Synthese, 1-23.

Bain, J. (2000), “Against Particle/Field Duality: Asymptotic Particle States and Interpolating Fields in Interacting QFT (or: Who’s Afraid of Haag’s Theorem?)”, Erkenntnis, 53(3), 375-406.

—— (2011), “Quantum Field Theories in Classical Spacetimes and Particles”, Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 42(2), 98-106.

—— (2013a), “Emergence in Effective Field Theories”, European journal for philosophy of science, 3(3), 257-273.

—— (2013b), “Effective Field Theories”, In Batterman, B. (ed), *The Oxford Handbook of Philosophy of Physics*, Oxford University Press, 224-254.

Baker, D. J. (2009), “Against Field Interpretations of Quantum Field Theory”, The British Journal for the Philosophy of Science, 60(3), 585-609.

Barnes, E. (2012), “Emergence and Fundamentality”, Mind, 121(484), 873-901.

—— (2014), “Going Beyond the Fundamental: Feminism in Contemporary Metaphysics”, Proceedings of the Aristotelian Society 114: 335–51.

Barrett, J. A. (2008), “Approximate Truth and Descriptive Nesting”, Erkenntnis, 68(2), 213-224.

Batterman, R. W. (2002), *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, Oxford University Press.

—— (2005), "Critical Phenomena and Breaking Drops: Infinite Idealizations in Physics", Studies in History and Philosophy of Modern Physics, 36, 225-244.

—— (2010), "On the Explanatory Role of Mathematics in Empirical Science", British Journal for the Philosophy of Science, 61(1), 1-25.

Batterman, R. W. and Rice, C. C. (2014), "Minimal Model Explanations", Philosophy of Science, 81(3), 349-376.

Bednyakov, A. V., Kniehl, B. A., Pikelner, A. F., and Veretin, O. L. (2015), "Stability of the Electroweak Vacuum: Gauge Independence and Advanced Precision", Physical review letters, 115(20), 201802.

Belot, G. (1998), "Understanding Electromagnetism", The British Journal for the Philosophy of Science, 49, 531-555.

—— (2007), "Is Classical Electrodynamics an Inconsistent Theory?", Canadian Journal of Philosophy, 37(2), 263-282.

—— (2013), "Symmetry and Equivalence", In Batterman, B. (ed), *The Oxford Handbook of Philosophy of Physics*, Oxford University Press, 224-254.

Binney, J. J., Dowrick, N. J., Fisher, A. J., and Newman, M. (1992), *The Theory of Critical Phenomena: an Introduction to the Renormalization Group*, Oxford University Press.

Boyd, R. (1985), "Lex Orandi Est Lex Credendi", In Churchland, P. M. and Hooker, P. A. (eds), *Images of Science*, Chicago University Press, Chicago, 3-34.

Bratteli, O. and Robinson, D. (2003), *Operator Algebras and Quantum Statistical Mechanics, Vol II*, Springer.

Brown, L. M. (1993), *Renormalization: From Lorentz to Landau (and Beyond)*, Springer.

Brown, H. R. (2005), *Physical Relativity: Space-Time Structure From a Dynamical Perspective*, Oxford University Press.

Buchholz, D. (2000), "Algebraic Quantum Field Theory: A Status Report", arXiv preprint math-ph/0011044.

Butterfield, J. (2011), "Less is Different: Emergence and Reduction Reconciled", Foundations of Physics, 41 (6), 1065-1135.

Butterfield, J. and Bouatta, N. (2012), "Which Quantum Field Theories Exist?", talk presented at *Structuralism in Physics and Mathematics*, University of Bristol, September 21-23 2012.

—— (2014), “Renormalization for Philosophers”, arXiv:1406.4532.

Butterfield, J. and Christopher I. (2001), “Spacetime and the Philosophical Challenge of Quantum Gravity”, In Callender, C. and Huggett, N. (eds), *Physics Meets Philosophy at the Planck Scale*, 33-89.

Callender, C. (2011), “Hot and Heavy Matters in the Foundations of Statistical Mechanics”, Foundations of Physics, 41(6), 960-981.

Callender, C. and Huggett, N. (2001), *Physics Meets Philosophy at the Planck scale: Contemporary Theories in Quantum Gravity*, Cambridge University Press.

Callender, C. and Menon, T. (2011), “Turn and Face the Strange... Ch-ch-changes: Philosophical Questions Raised by Phase Transitions” in Batterman, R. (eds), *The Oxford Handbook of Philosophy of Physics*, Oxford University Press.

Cao, T. Y. (1998), *Conceptual Developments of 20th Century Field Theories*, Cambridge University Press.

Cao, T. Y. and Schweber, S. S. (1993), “The Conceptual Foundations and the Philosophical Aspects of Renormalization Theory”, Synthese, 97(1), 33-108.

Castellani, E. (2003), “On the Meaning of Symmetry Breaking”. in Brading, K. and Castellani, E. (eds), *Symmetries in physics: Philosophical reflections*, Oxford University Press.

Caulton, A. (2016), “The Role of Symmetry in the Interpretation of Physical Theories”, Forthcoming in Studies in the History and Philosophy of Science.

Clifton, R. (2004), *Quantum Entanglements: Selected Papers*, Butterfield, J. and Halvorson, H. (eds), Clarendon Press.

Clifton, R. and Halvorson, H. (2000), “Generic Bell Correlation Between Arbitrary Local Algebras in Quantum Field Theory”, Journal of Mathematical Physics, 41(4), 1711-1717.

—— (2001), “Entanglement and Open Systems in Algebraic Quantum Field Theory”, Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 32(1), 1-31.

Coffey, K. (2014), “Theoretical Equivalence as Interpretative Equivalence” The British Journal for the Philosophy of Science, 65(4): 821-844.

Collins, J. C. (1984), *Renormalization: an Introduction to Renormalization, the Renormalization Group and the Operator-Product Expansion*, Cambridge university press.

Crowther, K. (2013), “Emergent Spacetime According to Effective Field Theory:

From Top-Down and Bottom-Up" Studies in History and Philosophy of Modern Physics 44 (2013) 321–328.

Curiel, E. (2011), "On the Propriety of Physical Theories as a Basis for their Semantics", http://philsci-archive.pitt.edu/8702/

—— (2014), "Classical mechanics is Lagrangian; it is not Hamiltonian", The British Journal for the Philosophy of Science, 65(2), 269-321.

Da Costa, N. C., and French, S. (2003), *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*, Oxford University Press.

Dasgupta, S. (2015), "Symmetry as an Epistemic Notion (Twice Over)", The British Journal for the Philosophy of Science, doi:10.1093/bjps/axu049.

Delamotte, B. (2004), "A Hint of Renormalization", Amer. J Phys. 72, 170–184.

Dennett, D. C. (1991), "Real patterns", The journal of Philosophy, 88(1), 27-51.

Dorr, C. (2005), "What We Disagree About When We Disagree About Ontology", In Kalderon, M. E. (ed), *Fictionalism in Metaphysics*, Oxford University Press, 234-286.

Douglas, M. R. (2004), "Report on the Status of the Yang-Mills Millenium Prize Problem", http://www. claymath. org/millennium/Yang-Mills Theory/ym2.pdf.

Dowker, F., Henson, J. and Sorkin, R. D. (2004), "Quantum Gravity Phenomenology, Lorentz Invariance and Discreteness", Modern Physics Letters A, 19(24), 1829-1840.

Duncan, A. (2012), *The Conceptual Framework of Quantum Field Theory*, Oxford University Press.

Dyson, F. J. (1952), "Divergence of Perturbation Theory in Quantum Electrodynamics", Physical Review, 85(4), 631.

Earman, J. (2004), "Curie's Principle and Spontaneous Symmetry Breaking", International Studies in the Philosophy of Science, 18(2-3), 173-198.

Earman, J., and Fraser, D. (2006), "Haag's Theorem and its Implications for the Foundations of Quantum Field Theory", Erkenntnis, 64(3), 305-344.

Elitzur, S. (1975), "Impossibility of Spontaneously Breaking Local Symmetries", Physical Review D, 12(12), 3978.

Emch, G. G. and Lui, C. (2005), "Explaining Quantum Spontaneous Symmetry Breaking" Studies in the History and Philosophy of Modern Physics, 36, 137-163.

Feynman, R. P. (2006), *QED: The Strange Theory of Light and Matter* Princeton

University Press, Princeton.

Fine, A. (1986), *The Shaky Game: Einstein, Realism, and the Quantum Theory*, Chicago: University of Chicago Press.

Fitzpatrick, S. (2013), “Doing Away with the No Miracles Argument”, In Karakostas, V. and Dieks, D. (eds), *EPSA11 Perspectives and Foundational Problems in Philosophy of Science*, Springer, 141-151.

Fletcher, S. C. (2012), ‘What Counts as a Newtonian System? The View from Norton’s Dome.’ European Journal for Philosophy of Science, 2(3), 275-297.

Fraser, D. (2008), “The Fate of ‘Particles’ in Quantum Field Theories with Interactions”, Studies in History and Philosophy of Science Part B: 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

—— (2012), “Spontaneous Symmetry Breaking: Quantum Statistical Mechanics versus Quantum Field Theory”, Philosophy of Science, 79(5), 905-916.

Fraser, J. [James] (2012), “Renormalisation Scheme Dependence and the Scope of Perturbative Quantum Field Theory”, Msc. dissertation, University of Durham.

French, S. (2003), “A Model-Theoretic Account of Representation (Or, I Don’t Know Much about Art... but I Know It Involves Isomorphism)”, Philosophy of Science, 70(5), 1472-1483.

—— (2012), “Unitary Inequivalence as a Problem for Structural Realism”, Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 43(2), 121-136.

Friederich, S. (2013), “Gauge Symmetry Breaking in Gauge Theories—in Search of Clarification”, European journal for philosophy of science, 3(2), 157-182.

Frisch, M. (2005), *Inconsistency, Asymmetry, and Non-Locality*, Oxford University Press.

Frost-Arnold, G. (2010), “The No-Miracles Argument: Inference to an Unacceptable Explanation”, Philosophy of Science, 77(1), 35-58.

Gell-Mann, M and Low, F.E. (1954), “Quantum Electrodynamics at Small Distances”. Physical Review 95 (5): 1300–1312.

Georgi, H. (1993), “Effective Field Theory”, Annual Review of Nuclear and Particle Science 43: 209-52.

Georgii, H. (1988), *Gibbs Measures and Phase Transitions*, W. De Gruyter, Berlin.

Giere, R. N. (2004), “How Models are Used to Represent Reality”, Philosophy of science, 71(5), 742-752.

Gies, H. and Jaeckel, J. (2004), “Renormalization Flow of QED”. Physical Review Letters 93 (11): 110405.

Glimm, J. and Jaffe, A. (1981), *Quantum physics: a Functional Integral Point of View*, Springer-Verlag.

—— (1985), *Collected Papers: Constructive Quantum Field Theory Selected Papers.* Springer.

Golterman, M. (2011), “Applications of Chiral Perturbation Theory to Lattice QCD”, In Lellouch, L., Sommer, R., Svetitsky, B., Vladikas, A. and Cugliandolo, L. (eds), *Modern Perspectives in Lattice QCD*, Oxford University Press, 423-515.

Haag, R. (1955), “On Quantum Field Theories” Det Kongelige Danske Videnskabernes Selskab, Matematisk-fysiske Meddelelser 29(12), 1–37.

—— (1996), *Local Quantum Physics*, Springer-Verlag.

Haag, R. and Kastler, D. (1964), ”An Algebraic Approach to Quantum Field Theory”, Journal of Mathematical Physics 5: 848–861.

Hagar, A. (2009), “Minimal Length in Quantum Gravity and the Fate of Lorentz Invariance”. Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 40(3), 259-267.

Hall, D. and A. Wightman (1957), “A Theorem on Invariant Analytic Functions with Applications to Relativistic Quantum Field Theory” Det Kongelige Danske Videnskabernes Selskab, Matematisk-fysiske Meddelelser 31(5), 1–41.

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*, Elsevier.

—— (2012), “What Scientific Theories Could Not Be”, Philosophy of Science 79,183–206.

Halvorson, H. and Clifton, R. (2002), “No Place for Particles in Relativistic Quantum Theories?” Philosophy of Science, 69, 1-28.

Hancox-Li, L. (2015a), *Moving Beyond Theory “T”: The Case of Quantum Field Theory*, Doctoral dissertation, University of Pitsburgh.

—— (2015b), "Coarse-Graining as a Route to Microscopic Physics: The Renormalization Group in Quantum Field Theory", Philosophy of Science 82 (5), 1211-1223.

Healey, R. (2007), *Gauging What's Real*, Oxford University Press.

—— (2012), "Quantum Theory: a Pragmatist Approach", The British Journal for the Philosophy of Science, axr054.

Huggett, Nick and Robert Weingard (1995), "The Renormalisation Group and Effective Field Theories", Synthese 102 (1), 171-194.

Huggett, N. and Wüthrich, C. (2013), "Emergent Spacetime and Empirical (In)coherence", Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 44(3), 276-285.

Jaffe, A., and Witten, E. (2006), "Quantum Yang–Mills Theory", In *The Millennium Prize Problems*, Clay Mathematics Institute, 129–152.

Jones, M. R. (2005), "Idealization and Abstraction: A Framework", In Jones, M. R. and Cartright, N. (eds), *Idealization XII: Correcting the Model. Idealization and Abstraction in the Sciences*, 86, 173-217.

Kadanoff, L. P. (1966), "Scaling Laws for Ising Models Near T_c". Physics (2) 263-272.

Kaiser, D. (2009), *Drawing Theories Apart: The Dispersion of Feynman Diagrams in Postwar Physics*, University of Chicago Press.

Kindermann, R. and Snell, L. (1980), *Markov Random Fields and their Application*, American Mathematics Society.

Kitcher, P. (1993), *The Advancement of Science*, Oxford University Press.

Koma, T. and Tasaki, H. (1994), "Symmetry Breaking and Finite-Size Effects in Quantum Many-Body Systems", Journal of Statistical Physics 76.3-4 745-803.

Kosso, P. (2000), "Fundamental and Accidental Symmetries", International Studies in the Philosophy of Science, 14:2, 109-121.

Kuhlmann, M. (2010), "Why Conceptual Rigour Matters to Philosophy: On the Ontological Significance of Algebraic Quantum Field Theory", Foundations of Physics 40 (9):1625-1637.

Kukla, A. (1996), "Does Every Theory have Empirically Equivalent Rivals?", Erkenntnis, 44(2), 137-166.

Lachieze-Rey, M. and Luminet, J. P. (1995), "Cosmic Topology", Physics Reports, 254(3), 135-214.

Ladyman, J. (2002), *Understanding Philosophy of Science*, Routledge.

Lancaster, T. and Blundell, S. J. (2014), *Quantum Field Theory for the Gifted Amateur*, Oxford University Press.

Landsman, N. P. (2013), "Spontaneous Symmetry Breaking in Quantum Systems: Emergence or Reduction?", Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 44(4), 379-394.

Lange, M. (2004), "Bayesianism and Unification: A Reply to Wayne Myrvold", Philosophy of Science, 71(2), 205-215.

—— (2015), "On "Minimal Model Explanations": A Reply to Batterman and Rice", Philosophy of Science 82 (2):292-305.

Laudan, L. (1981), "A Confutation of Convergent Realism", Philosophy of Science, 48, 19-48.

Lebowitz, Joel L. (1999), "Statistical Mechanics: A Selective Reivew of Two Central Issues", Reviews of Modern Physics,. 71, 346-357.

Lepage, G. P. (2005), "What is Renormalization?", arXiv preprint hep-ph/0506 330.

Lowe, E. J. and Tahko, T. (2015), "Ontological Dependence", Stanford Encyclopedia of Philosophy.

Lui, C. (2003), "Spontaneous Symmetry Breaking and Chance in a Classical World", Philosophy of Science 70: 590-608.

Lüscher, M. and Weisz, P. (1987), "Scaling Laws and Triviality Bounds in the Lattice ϕ^4 Theory", Nucl. Phys. B290: 25.

Lyre, H. (2008), "Does the Higgs Mechanism Exist?", International Studies in the Philosophy of Science, 22(2), 119-133.

MacKinnon, E. (2008), "The Standard Model as a Philosophical Challenge", Philosophy of Science, 75: 447–457.

Magnen, J. and Seneor, R. (1977), "Phase Space Cell Expansion and Borel Summability for the Euclidean ϕ_3^4", Commun. Math. Phys., 56, 237–276.

Magnus, P. D. (2003), "Underdetermination and the Problem of Identical Rivals", Philosophy of Science, 70(5), 1256-1264.

Magnus, P. D., and Callender, C. (2004), "Realist Ennui and the Base Rate Fallacy", Philosophy of Science, 71(3), 320-338.

Mainwood, P. (2006) *Is More Different? Emergent Properties in Physics*, Doctoral dissertation, University of Oxford.

Malament, D. (1986), “Newtonian Gravity, Limits, and the Geometry of Space”, In Colodny, R. (ed), *From Quarks to Quasars*, University of Pittsburgh Press.

—— (1995), “Is Newtonian Cosmology Really Inconsistent?”, Philosophy of Science, 489-510.

—— (1996), “In Defense of Dogma: Why There Cannot be a Relativistic Quantum Mechanics of (Localizable) Particles”, In Clifton, R. (ed), *Perspectives on Quantum Reality*, Springer, 1-10.

Mattingly, D. (2005), “Modern Tests of Lorentz Invariance”, Living Rev. Rel, 8(5), 2003.

Mauldin, T. (1996), “On the Unification of Physics”, The Journal of Philosophy, 93(3), 129-144.

McKenzie, K. (2012), *Physics Without Fundamentality: A Study into Anti-Fundamentalism about Particles and Laws in High-Energy Physics*, Doctoral dissertation, University of Leeds.

—— (2013) “Priority and Particle Physics: Ontic Structural Realism as a Fundamentality Thesis”, The British Journal for the Philosophy of Science, axt017.

McMullin, E. (1985). “Galilean Idealization”, Studies in History and Philosophy of Science Part A, 16(3), 247-273.

Meynell, L. (2008) “Why Feynman Diagrams Represent”, International Studies in the Philosophy of Science 22(1): 39-59.

Miller, D. (1976), “Verisimilitude Redeflated”, The British Journal for the Philosophy of Science, 27(4), 363-381.

Miller, M. (2016), “Haag’s Theorem, Apparent Inconsistency, and the Empirical Adequacy of Quantum Field Theory”, Forthcoming in The British Journal for the Philosophy of Science.

Moore, G. (2003), “Informal Lectures on Lattice Gauge Theory”, Lecture notes, McGill University. http://www.physics.mcgill.ca/ guymoore/lattlectures.pdf.

Morrison, M. (2007), *Unifying Scientific Theories: Physical Concepts and Mathematical Structures*, Cambridge University Press.

Muller, F. A. (2007), “Inconsistency in Classical Electrodynamics?”, Philosophy of Science, 74(2), 253-277.

Muta, T. (2010), *Foundations of Quantum Chromodynamics: an Introduction to Perturbative Methods in Gauge Theories*, World Scientific.

Myrvold, W. C. (2003), “A Bayesian Account of the Virtue of Unification”, Philosophy of Science, 70(2), 399-423.

Neumaier, A. (2011), "Renormalization without Infinities – an Elementary Tutorial", Unpublished Manuscript, University of Vienna, http://www.mat.univie.ac.at/ neum/ms/ren.pdf

Ney, A. and Albert, D. Z. (eds), (2013), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford University Press.

Niedermaier, M., and Reuter, M. (2006), "The Asymptotic Safety Scenario in Quantum Gravity", Living Rev. Rel, 9(5), 173.

Niiniluoto, I. (1987), *Truthlikeness*, Reidel.

—— (1999), *Critical Scientific Realism*, Clarendon Press.

North, J. (2009), "The 'Structure' of Physics", The Journal of Philosophy 106 (2), 57–88.

Norton, J. D. (1992), "A Paradox in Newtonian Gravitation Theory", PSA: Proceedings of the biennial meeting of the philosophy of science association, 412-420.

—— (2002), "A Paradox in Newtonian Gravitation Theory II." In Meheus, J. (ed), *Inconsistency in Science*, Springer, 185-195.

—— (2008), "The Dome: An Unexpectedly Simple Failure of Determinism", Philosophy of Science, 75(5), 786-798.

Norton, John. D. (2012), "Approximation and Idealization: Why the Difference Matters", Philosophy of Science, 79(2), 207-232.

Oddie, G. (1986), *Likeness to Truth*, Reidel.

Osterwalder, K. and Schrader, R. (1973), "Axioms for Euclidean Green's functions I", Comm. Math. Phys. 31, 83–112.

—— (1975), "Axioms for Euclidean Green's functions II", Commun. Math. Phys. 42, 281-305.

Peskin, M. E., and Schroeder, D. V. (1995), *An Introduction to Quantum Field Theory*, Westview.

Polchinski, J. (1984), "Renormalization and Effective Lagrangians", Nuclear Physics B, 231(2), 269-295.

—— (1992), "Effective Field Theory and the Fermi Surface", arXiv preprint hep-th/9210046.

—— (2012), "Comment on 'Small Lorentz Violations in Quantum Gravity: Do they Lead to Unacceptably Large Effects?"', Classical and Quantum Gravity, 29(8), 088001.

Popper, K. (1976), "A Note on Verisimilitude", The British Journal for the Philosophy of Science, 27(2): 147–159.

Psillos, S. (1999), *Scientific Realism: How Science Tracks Truth*, Routledge.

Putnam, H. (1975), *Mathematics, Matter and Method*, Cambridge University Press.

Quine, W. (1975), "On Empirically Equivalent Systems of the World", Erkenntnis, 9, pp.313–28.

Redhead, M. (1980), "Models in Physics", *British Journal for the Philosophy of Science*, 31(2): 145-163.

—— (1988), "A Philosopher Looks at Quantum Field Theory", In Brown, H. and Harré (eds), *Philosophical Foundations of Quantum Field Theory*, Oxford University Press, 9-24.

Rivasseau, V. (2014), *From Perturbative to Constructive Renormalization*, Princeton University Press.

Rueger, A. (1990), "Independence from Future Theories: A Research Strategy in Quantum Theory", In Fine, A., Forbes, M. and Wessels, L. (eds), *PSA 1990*, East Lansing: PSA, 203–211.

Ruelle, D. (1969), *Statistical Mechanics: Rigorous Results*, W.A. Benjamin.

Ruetsche, L. (2003), "A Matter of Degree: Putting Unitary Inequivalence to Work", Philosophy of Science 70 (5):1329-1342.

—— (2006), "Johnny's so Long at the Ferromagnet", Philosophy of Science, 73(5), 473-486.

—— (2011), *Interpreting Quantum Theories*, Oxford University Press.

—— (2015), "The Shaky Game +25, or: on Locavoracity", Synthese, 192(11), 3425-3442.

Saatsi, J. (2009a), "Form vs. Content-driven Arguments for Realism", In Magnus, P. and Busch, J. (eds), *New Waves in Philosophy of Science*, Palgrave Macmillan, 8–28.

—— (2009b), "Whence Ontological Structural Realism?", Suárez, M., Mauro D. and Miklós R. (eds), *EPSA Epistemology and Methodology of Science* Springer, 255–65.

—— (2015), "Replacing Recipe Realism", Synthese, 1-12.

—— (2016a), "Models, Idealisation and Realism", http://philsci-archive.pitt.edu/11739/

—— (2016b), “What is Theoretical Progress of Science?”, Forthcoming in Synthese.

Schaffer, J. (2009) “On What Grounds What”, In Manley, D., Chalmers, D. J. and Wasserman, R. (eds), *Metametaphysics: New Essays on the Foundations of Ontology*, Oxford University Press, 347-383.

Schweber, S. S. (1994), *QED and the Men Who Made it: Dyson, Feynman, Schwinger, and Tomonaga*, Princeton University Press.

Shech, E. (2013), “What Is the Paradox of Phase Transitions?”, Philosophy of Science, 80(5), 1170-1181.

Sider, T. (2011), *Writing the Book of the World*, Oxford University Press.

Smit, J. (2002), *Introduction to Quantum Fields on a Lattice*, Cambridge University Press.

Sober, E. (2003), “Two Uses of Unification”, In Stadler, F. (ed), *Institute Vienna Circle Yearbook 2002*, Kluwer.

Stanford, P. K. (2001), “Refusing the Devil’s Bargain: What Kind of Underdetermination Should We Take Seriously?”, Philosophy of Science, 68(3), S1–S12.

—— (2003), “Pyrrhic Victories for Scientific Realism”, The journal of philosophy, 100(11), 553-572.

—— (2006), *Exceeding our Grasp: Science, History, and the Problem of Unconceived Alternatives*, Oxford University Press.

Streater, R. F., and Wightman, A. S. (1964), *PCT, Spin and Statistics, and All That*, Princeton University Press.

Strevens, M. (2008), *Depth: An Account of Scientific Explanation*, Harvard University Press.

Strocchi, F. (2008), *Symmetry Breaking*, Springer.

—— (2013), *An Introduction to Non-Perturbative Foundations of Quantum Field Theory*, Oxford University Press.

Stueckelberg, E.C.G. and Petermann, A. (1953), “La Renormalisation Des Constants Dans la Theorie de Quanta”, Helv. Phys. Acta 26, 499.

Suárez, M. (2004), “An Inferential Conception of Scientific Representation”, Philosophy of science, 71(5), 767-779.

Suppe, F. (1989), *The Semantic Conception of Theories and Scientific Realism*, University of Illinois Press.

Teller, P. (1989), “Infinite Renormalization”, Philosophy of Science, 238-257.

—— (1995), *An Interpretive Introduction to Quantum Field Theory*, Princeton University Press.

Thomson-Jones, M. (2006), "Models and the Semantic View", Philosophy of Science, 73, 524–535.

Toader, I. (2016), "Unitary Inequivalence and Inherent Modality", Unpublished Manuscript, University of Bucharest.

van Fraassen, B. (1980), *The Scientific Image*, Clarendon Press.

—— (1991), *An Empiricist View of Quantum Mechanics*, Clarendon Press.

—— (2014), "One or Two Gentle Remarks About Hans Halvorson's Critique of the Semantic View", Philosophy of Science, 81(2), 276-283.

Vickers, P. (2008), "Frisch, Muller, and Belot on an Inconsistency in Classical Electrodynamics", British Journal for the Philosophy of Science, 59(4), 767-792.

—— (2013), *Understanding Inconsistent Science*, Oxford University Press.

—— (2014), "Scientific Theory Eliminativism", Erkenntnis, 79(1), 111-126.

Wald, R. M. (1994), *Quantum Field Theory in Curved Spacetime and Black Hole Thermodynamics*, University of Chicago Press.

Wallace, D. (2001), "Emergence of Particles from Bosonic Quantum Field Theory", arXiv preprint quant-ph/0112149.

—— (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.

Wallace, D. and Timpson, C. G. (2010), "Quantum Mechanics on Spacetime I: Spacetime State Realism", The British Journal for the Philosophy of Science, 61(4), 697-727.

Weingard, R. (1988), "Virtual Particles and the Interpretation of Quantum Field Theory", In H. Brown and Harré, R. (eds), *Philosophical Foundations of Quantum Field Theory*, Oxford University Press. 43–58.

Weisz, P. (2011), "Renormalization and lattice artifacts", In Lellouch, L., Sommer, R., Svetitsky, B., Vladikas, A. and Cugliandolo, L. (eds), *Modern Perspectives in Lattice QCD*, Oxford University Press, 93-160.

Wightman, A. S. (1979), "Should We Believe in Quantum Field Theory?", In Zichichi, A. (ed), *The Whys of Subnuclear Physics*, Plenum Press, 983-1025.

—— (1986) "Some Lessons of Renormalization Theory", In de Boer, J. Dal, E. and Ulfbeck, O. (eds), *The Lesson of Quantum Theory*, Elsevier, 201–226.

Williams, P. (2015), "Naturalness, the Autonomy of Scales, and the 125 GeV Higgs", Studies in History and Philosophy of Modern Physics, 51, 82–96.

— (2016) "Scientific Realism Made Effective" Unpublished Manuscript, University of Pittsburg.

Wilson, K. G. (1983), "The Renormalization Group and Critical Phenomena", Reviews of Modern Physics, 55(3), 583.

Wilson, K. G., and Kogut, J. (1974), "The Renormalization Group and the ϵ Expansion", Physics Reports, 12(2), 75-199.

Wilson, M. (2006), *Wandering Significance: An Essay on Conceptual Behaviour*, Oxford University Press, Oxford.

Woodward, J. (2005), *Making Things Happen: A Theory of Causal Explanation*, Oxford University Press.

Worrall, J. (1989), "Structural Realism: The Best of Both Worlds?", Dialectica 43(1-2), 99-124.

Wüthrich, A. (2012), "Interpreting Feynman Diagrams as Visual Models", Spontaneous Generations, 6(1), 172-181.

Zinkernagel, H. (2010), "Causal Fundamentalism in Physics", In Suárez, M., Dorato, M. and Rédei M. (eds), *EPSA Philosophical Issues in the Sciences*, Springer.

Zinn-Justin, J. (2002), *Quantum Field Theory and Critical Phenomena*, Clarendon Press.