

Renormalization Group Methods

Porter Williams

It is a truism, in physics if not in philosophy, that in order to study physical behavior at a particular scale one is best served by using degrees of freedom defined at (or suitably near) that scale. If one wants to study the classical behavior of a simple fluid on relatively long length scales ($\gg 10^{-10}\text{m}$), for example, the appropriate degrees of freedom to use to describe the fluid are two continuous fields, the mass density $\rho(x)$ and the velocity field $\mathbf{v}(x)$. This despite the fact that the fluid itself is constituted by a large collection of discrete, atomic constituents, which in turn have their own subatomic structure. Remarkably, one can do a tremendous amount of empirically successful fluid dynamics ignoring almost entirely the more “fundamental,” discrete description of the fluid and modeling its long-distance behavior using “effective” continuum degrees of freedom.

Even in purportedly fundamental physics the truism holds. For example, consider quantum chromodynamics, the theory that describes the physics of quarks and gluons interacting via the strong force. Although hadrons (such as protons and neutrons and pions) interact via the strong force and are (in a rather complicated sense) “composed” of quarks and gluons, a considerable amount of low-energy nuclear physics proceeds perfectly well by modeling nuclear processes in ways that *ignore* those “fundamental” quark and gluon degrees of freedom entirely, modeling physical processes in terms of “effective” hadronic degrees of freedom. If one wants to model the low-energy scattering of pions off of a fixed nucleon target, for example, one is better off eschewing any attempt to describe the scattering in terms of the

“fundamental” quark and gluon degrees of freedom in favor of an “effective” description in terms of pions and nucleons. On the other hand, if one begins to scatter the pions off the nucleon target at energies so high that the pions begin to scatter inelastically (roughly, to probe the internal structure of the nucleon), then the theory involving only hadronic degrees of freedom becomes empirically and explanatorily inadequate and a set of degrees of freedom that include the quark and gluon fields becomes appropriate for modeling physics at that new energy scale (See Jansson (this volume) for further discussion of the relationship between scale and explanation).

Theories that take into account only the degrees of freedom that are necessary for characterizing physical processes at some particular scale, but which (i) break down when pushed to scales beyond their limited domain of applicability and (ii) incorporate this inevitable breakdown into their mathematical framework, are called *effective theories*. Wallace (this volume) has given a clarifying description of how this approach to physical theorizing plays out in quantum field theory (QFT), and the reader is strongly encouraged to read Wallace’s chapter in conjunction with this one. The basic requirement for the modeling strategy that underlies the effective theory approach to work might be called the “autonomy of scales”: that physical processes at long distances must exhibit minimal dependence on the shorter-distance physics that the effective theory’s description leaves out.¹

In particle physics and solid state physics, the main set of tools for determining how physical processes at longer scales depend on short-distance physics, and thus for determining which degrees of freedom are optimal for characterizing physical processes at a particular scale, are renormalization group (RG) methods. As Steven Weinberg puts it,

I think that this in the end is what the renormalization group is all about. It’s a way of satisfying the Third Law of Progress

¹For a philosophical discussion of the issues raised by multi-scale modeling, see (Batterman, 2013).

in Theoretical Physics, which is that *you may use any degrees of freedom you like to describe a physical system, but if you use the wrong ones, you'll be sorry.* (Weinberg, 1983, p. 16)

Despite this seemingly sensible motivation, the process of renormalization in QFT has historically been treated as an ill-founded but necessary evil: a collection of distasteful tricks, of dubious mathematical and physical standing, employed by practicing particle physicists in order to extract empirical predictions from the theory. This was the predominant attitude in the particle physics community from the 1930s through much of the 1960s, and this attitude contributed significantly to a widespread distrust in the entire framework of quantum field theory that took hold shortly after World War II and held sway until the early 1970s.²

A major shift in the particle physics community's attitude toward QFT coincided with the invention and refinement of RG methods by Murray Gell-Mann and Francis Low, Leo Kadanoff, Michael Fisher, Kenneth Wilson, and others from the mid-1950s through the early 1970s. The RG was widely viewed as having put the earlier renormalization methods on secure physical foundations. It also provided a new set of mathematical tools that proved to be indispensable both to practicing particle physicists, who were engaged in extracting empirical predictions from particular quantum field-theoretic models, and to mathematical physicists, whose aim was to put the theoretical framework itself on secure mathematical footing. The development of the RG ultimately ushered in a transformation of the way physicists understand the

²One finds this description scattered throughout particle physicists' descriptions of the period; already in 1965, Kenneth Wilson lamented that "The Hamiltonian formulation of quantum mechanics has been essentially abandoned in investigations of the interactions of π mesons, nucleons, and strange particles" in favor of approaches using dispersion relations (Wilson, 1965, p. 445), while Steven Weinberg refers to the "the general distrust of quantum field theory that set in soon after the brilliant successes of quantum electrodynamics in the late 1940s" (Weinberg, 1983, p. 7), and David Gross begins his description of particle physics in the 1960s by recalling that "Field theory was in disgrace; S-Matrix theory was in full bloom" (Gross, 2004, p.16). (Cushing, 1990) is a book-length treatment of the history of S-Matrix Theory, the main competitor framework to QFT in the 1960s.

conceptual foundations of quantum field theory itself: once thought to be a candidate fundamental theory, QFT is now widely understood to be just another effective theory.³

This chapter proceeds as follows. I begin with a brief introduction to path integral methods, since the path integral formulation of quantum field theory is a natural home for examining the generality of RG methods. I then introduce in some detail the machinery of early renormalization theory before turning to RG method; my aim is to illustrate the sense in which RG methods changed the way that we think about renormalization in QFT. I then turn to a major conceptual shift in the way that we think about QFT, brought on by the development of RG methods. The development of the RG has led naturally to a perspective in which individual quantum field theories live in a vast “space” of quantum field theories, and that one of the tasks of the theoretical physicist is to provide a map of theory space. I will conclude by considering the benefits that thinking in terms of theory space has for the task of defining and classifying QFTs.

Finally, a note on conventions: I use so-called “natural units” in which $\hbar = c = 1$. This means that energy and mass have the same dimensions, and length and time have dimensions of mass⁻¹. This makes possible two things that will prove useful going forward: we can characterize all physical quantities as having a positive or negative dimension of mass, and we can identify high energy scales with short distance scales. I will also use a “mostly minus” signature for the Minkowski metric, so that $\eta_{\mu\nu}x^\mu x^\nu = (x^0)^2 - (x_i x^i)$ and e.g. $(\partial_\mu\phi)^2 = (\partial_t\phi)^2 - (\nabla\phi)^2$.

³There is, however, an active research program aimed at determining whether a QFT describing gravitation may be well-defined at arbitrarily high energies; see (Niedermaier and Reuter, 2006) for a review. If the asymptotic safety program is successful, that opens the door to the possibility that QFT may be a fundamental framework after all.

1 Path integrals

The path integral representation of a QFT provides a particularly natural home for RG methods. Although path integrals in QFT are notoriously plagued by mathematical difficulties, they are perfectly well-defined mathematical objects in the *nonrelativistic* quantum mechanics of finitely-many particles. Since the physical ideas are largely the same in both cases, I will introduce path integrals in the latter, mathematically well-behaved setting before extending the discussion to QFT.

I will borrow a lovely example from Richard Feynman to motivate the notion of the path integral (Feynman and Hibbs, 1965/2010). Consider a double-slit experiment, in which one has (i) a source emitting individual particles, (ii) a detection screen, and (iii) between the source and the detector, a wall with two small holes labeled 1 and 2. The aim is to calculate the probability that the particle will be detected at some particular point on the screen. The particle must go through either hole 1 or hole 2 in order to reach the screen, and so the total amplitude for the particle being detected at any particular point x on the screen is given by the sum of two amplitudes: one for the particle to pass through hole 1 and arrive at x and another for the particle to pass through hole 2 and arrive at x :

$$\psi_{12}(x) = \frac{1}{\sqrt{2}}(\psi_1(x) + \psi_2(x))$$

with the probability that the particle is detected at x given by $|\psi_{12}(x)|^2$.

Now suppose that between our source and our original wall we put two more walls, which we label A and B, and drill several holes in each of them. Now there are a number of paths that the particle can take from the source to the detection screen: it could go from the source to A_1 , then to B_3 , and then through hole 2; it could go from the source to A_5 , then to B_2 , and then through hole 1; and so on. Each of these paths has its own amplitude, and the total amplitude for the particle being detected at the point x on the

detection screen is the sum of all of them.

Now imagine inserting more and more walls between the source and the detection screen, and drilling holes in each of the walls until there is nothing left. Eventually there will be infinitely many walls, each with infinitely many holes, and thus infinitely many possible paths that will take the particle from the source to the point x on the detection screen. In order to determine the total amplitude for the particle being detected at x we will have to sum over the amplitudes for *all* of them.

This idea of summing the amplitudes for each of the possible ways that a physical system in one state could transition into another state is what underlies the path integral formalism. Though we motivated the idea of the path integral as a method for determining the probability for detecting a particle emitted from a source at some particular point x on a detection screen, the path integral formalism is more general: it gives an efficient means for calculating the probability that a physical system in a state $|\alpha\rangle$ at time $t = t_0$ will transition into a state $|\beta\rangle$ at some later time $t = T$.

More formally, suppose we would like to calculate the amplitude that a system in state $|\alpha\rangle$ will transition into $|\beta\rangle$ between $t = 0$ and $t = T$, where the dynamics governing the time evolution are contained in the Hamiltonian $\hat{H} = \frac{\hat{p}^2}{2m} + V(\hat{x})$:

$$\langle\beta, T|e^{-i\hat{H}T}|\alpha, 0\rangle$$

We want to sum over all of the “paths” that the system could take from $|\alpha\rangle$ at $t = 0$ to $|\beta\rangle$ at $t = T$. We can do this in three steps:

1. Chop up the time interval $[0, T]$ into N equal time-steps dt , with $\frac{T}{N} = dt$. This allows us to re-write the time-evolution operator $e^{-i\hat{H}t}$ in the form $e^{-i\hat{H}t} = \underbrace{e^{-i\hat{H}dt} \times e^{-i\hat{H}dt} \times \dots \times e^{-i\hat{H}dt}}_{N \text{ times}} = \prod_{i=1}^N e^{-i\hat{H}dt}$.
2. After each time-step, insert a resolution of the identity operator $\mathbb{I} =$

$\sum^n |\gamma_n\rangle\langle\gamma_n|$ in some appropriate basis. For example, in our motivating example above one would insert resolutions of the identity in the position basis $\mathbb{I} = \int dx_i |x_i\rangle\langle x_i|$. In the latter case, each insertion of the identity plays the role of inserting a new “wall” in which infinitely many holes have been drilled and through which the particle must pass after every dt . This allows us to write $\langle\beta, T|e^{-i\hat{H}T}|\alpha, 0\rangle$ in the form

$$\prod_{i=0}^N \left(\int dx_i \right) \langle\beta|e^{-i\hat{H}dt}|x_N\rangle \times \langle x_N|e^{-i\hat{H}dt}|x_{N-1}\rangle \times \cdots \times \langle x_2|e^{-i\hat{H}dt}|x_1\rangle \times \langle x_1|e^{-i\hat{H}dt}|\alpha\rangle$$

3. Finally, let the time-steps dt get arbitrarily small by letting $N \rightarrow \infty$. What one finds (though it should not be obvious from what we’ve done here; see e.g. (Zee, 2010, §I.2)) is that

$$\langle\beta, T|e^{-i\hat{H}T}|\alpha, 0\rangle = \int Dx(t) e^{i \int_0^T dt \int d^3x [\frac{1}{2}m\dot{x}^2 - V(x)]}$$

where the “measure” is given by $\int Dx(t) = \lim_{N \rightarrow \infty} \prod_{i=0}^N dx_i$. The quantity $\mathcal{L}[x(t), \dot{x}(t)] = \frac{1}{2}m\dot{x}^2 - V(x)$ is called the *Lagrangian density*, and the integral of the Lagrangian density $\mathcal{S} = \int_0^T dt \int d^3x \mathcal{L}$ is called the *action*.

So far we have been dealing with a nonrelativistic quantum system, but all of this carries over directly to quantum field theory.⁴ One modification of presentation is that in the field-theoretic case, the Lagrangian density (more precisely, the action) takes center stage: one typically specifies a quantum field theory by writing down a classical action and then “quantizing” it through one or another method of quantization. If one has in hand the action for a physical system then one can obtain the classical equations of motion

⁴Modulo serious issues with its mathematical sensibility; see (Glimm and Jaffe, 1987). Most notably, the “measure” $Dx(t)$ typically does not exist in realistic models of QFT, leaving the path integral ill-defined.

for that system as the solution to the Euler-Lagrange equation; for this reason one often says that the action contains all the dynamical information about the physical system.

To connect this with QFT, consider a theory of a real scalar field with mass m and a self-interaction given by a ϕ^4 term:

$$\mathcal{L} = \frac{1}{2} \partial_\mu \phi(x) \partial^\mu \phi(x) - \frac{1}{2} m^2 \phi(x)^2 - \frac{\lambda}{4!} \phi(x)^4$$

Using the equation derived above, we can use the action to calculate the amplitude that a quantum field will transition from one state $|\phi_0(x)\rangle$ to another state $|\phi_1(x)\rangle$ over a particular temporal interval:

$$\langle \phi_1(x) | e^{-i\hat{H}t} | \phi_0(x) \rangle = \int D\phi e^{i \int_0^T d^4x \mathcal{L}}$$

where $\int_0^T d^4x = \int_0^T dt \int d^3x$ is an integration over all of space and the temporal interval $[0, T]$, and the field $\phi(x)$ over whose histories we are integrating is constrained to take on the values $\phi(x) = \phi_0(x)$ at $t = 0$ and $\phi(x) = \phi_1(x)$ at $t = T$. The essential idea developed in the nonrelativistic case holds equally well here, although at a somewhat more abstract level: instead of possible “paths” of a quantum particle between two possible states of that particle, we are now integrating over all possible histories of states of a quantum field that would carry it between the states $|\phi_0(x)\rangle$ and $|\phi_1(x)\rangle$.

From a pragmatic perspective, the point of developing this machinery is to enable us to compute the amplitudes for transitions between different field states. Typically one is interested in computing the amplitude that a specified state $|\phi_{\text{in}}\rangle$ of the field at $t = -\infty$ will transition into some other specified state $|\phi_{\text{out}}\rangle$ at $t = +\infty$, given a particular dynamics specifying how the excitations of the fields – roughly, particles – will interact. This information is contained in mathematical objects called correlation functions (equivalently, n -point functions). These correlation functions are of physical interest because they can be turned into amplitudes for the outcomes of scattering processes (S-

matrix elements) via the Lehmann-Symanzik-Zimmerman (LSZ) reduction formula. The fact that the main situation of interest involves transitions between asymptotic states of the field at over the interval $[-\infty, +\infty]$ means that the path integral is typically an integration over all of space *and* time. We will therefore write the path integral as

$$\langle \phi_1(x) | e^{-i\hat{H}t} | \phi_0(x) \rangle = \int D\phi e^{i \int d^4x \mathcal{L}}$$

where the temporal interval over which we integrate is now $[-\infty, +\infty]$.

The path integral above describes the amplitude for a quantum field to transition between two generic states $|\phi_1(x)\rangle$ and $|\phi_0(x)\rangle$. Typically in quantum field theory one is interested in $|\phi_1(x)\rangle = |\phi_0(x)\rangle = |\Omega\rangle$; speaking loosely of particles (the notion of a “particle” in QFT is subtle; see Fraser (this volume)), this gives the amplitude that the field begins in the vacuum state, out of which the field is “excited” in some spacetime region to create a particle that propagates for some period of time, and then the particle annihilates at some other spacetime region to return the field to its vacuum state. Thus one is usually interested in correlation functions of the form $\langle \Omega | \phi(x_1) \phi(x_2) \cdots \phi(x_n) | \Omega \rangle$ which capture the correlation between the field values of the field ϕ at distinct spacetime points.⁵ The simplest correlation function in quantum field theory is the *2-point function* $\langle \Omega | \phi(x_1) \phi(x_2) | \Omega \rangle$ which is called the *propagator*. Again speaking loosely, this gives the amplitude for an excitation of ϕ at x_1 to propagate to x_2 before annihilating and returning the field to the vacuum state $|\Omega\rangle$.

With these facts about correlation functions in hand, we are finally in a position to write down the path integral representation of a generic n -point function:

⁵These are called *vacuum expectation values*, or VEVs. They are just correlation functions that describe correlations of field values in one particular state of the field.

$$\langle \Omega | \phi(x_1) \phi(x_2) \cdots \phi(x_n) | \Omega \rangle = \mathcal{N} \int D\phi \phi(x_1) \phi(x_2) \cdots \phi(x_n) e^{i \int d^4x \mathcal{L}}$$

where \mathcal{N} is a normalization factor. The obvious thing to do at this point is to try to *compute* one of these correlation functions. This will lead us naturally to the need for renormalization and, eventually, for RG methods.

2 Renormalization

One might naively think that computing correlation functions should be straightforward: just do the appropriate path integral! It turns out that things are not so simple, and that in physically interesting models – such as models describing quantum fields in four-dimensional spacetimes and in which the field excitations are allowed to interact – one is almost always forced to resort to various approximation methods.⁶

The most common of these is perturbation theory. When doing perturbation theory, one splits the path integral up into two parts: an integral over the “free” part of the dynamics $\mathcal{L}_{\text{free}} = \frac{1}{2}(\partial\phi)^2$ and an integral over the “interaction” part $\mathcal{L}_{\text{int}} = -\frac{1}{2}m^2\phi^2 - \frac{\lambda}{4!}\phi^4$. The free part of the dynamics is trivial, and one can do the path integral directly. The interacting part of the theory is where the difficulties arise. We cannot do the path integral over \mathcal{L}_{int} directly, but one profitable approximation method is to treat \mathcal{L}_{int} as a small perturbation of $\mathcal{L}_{\text{free}}$; this amounts to assuming that the *coupling* λ which determines how strongly the field excitations interact is sufficiently small that we can treat the full dynamics of the theory $\mathcal{L} = \mathcal{L}_{\text{free}} + \mathcal{L}_{\text{int}}$ as “close” to the theory devoid of interactions described by $\mathcal{L}_{\text{free}}$, and expand the path integral over \mathcal{L}_{int} as a series approximation.

⁶Many of the difficulties one encounters can be traced back to the fact mentioned in fn. 4 above, namely that the “measure” in the path integral is typically ill-defined for realistic models of QFT.

Slightly more explicitly, when doing perturbation theory one attempts to calculate (for example) the four-point function $G^{(4)}$ (speaking loosely, the amplitude for two excitations of the scalar field in some initial state to interact and scatter into another pair of excitations in some other final state) as follows:

$$\begin{aligned} & \langle \Omega | \phi(x_1) \phi(x_2) \phi(x_3) \phi(x_4) | \Omega \rangle \\ &= \mathcal{N} \int D\phi \phi(x_1) \phi(x_2) \phi(x_3) \phi(x_4) e^{i \int d^4x \mathcal{L}} \approx \sum^n \lambda^n \mathcal{I}_n \end{aligned}$$

where the \mathcal{I}_n are integrals over the momenta of the particles and, since λ was assumed to be small at the outset, one can reasonably expect that successive terms in the series will get smaller.⁷ Consider the structure of the integrals \mathcal{I}_n appearing in our perturbative calculation of $G^{(4)}$. It is well-known that many of the integrals appearing in the series expansion $\sum^n \lambda^n \mathcal{I}_n$ are divergent, and that these divergences arise from including field excitations of arbitrarily high momenta in our calculations. We will begin with a schematic discussion of these divergences and their elimination, and then go through a concrete example.

The process of eliminating the divergences can be split into two steps: *regularization* and *renormalization*. One first adopts a method of *regularizing* the divergent integrals. The central idea of a regularization method is the following. Given a divergent integral \mathcal{I} in the perturbative expansion of the path integral, one makes the integral a function of new parameter θ – the regulator – so that the integral now becomes $\mathcal{I}(\theta)$. In order to count as a successful regulator, there are two constraints that θ must satisfy: (i) finite values of the regulator θ must render $\mathcal{I}(\theta)$ finite, and (ii) if one were to remove the regulator by taking $\theta \rightarrow \infty$, the result must be the original

⁷In standard textbook presentations, one would use Feynman diagrams to keep track of the integrals \mathcal{I}_n at each order of perturbation theory.

divergent integral \mathcal{I} . A third, weaker constraint is that it is desirable to choose a regularization method that leaves the mathematical structure of the theory as unchanged as possible; we will see below that *some* modification of that structure is inevitable, and which aspects of that structure one is willing to modify will often depend on the physical processes one is trying to understand.

For example, one historically important regularization method – the so-called Pauli-Villars method – involves introducing new fictitious “ghost” particles of mass θ into the Lagrangian density. Pauli-Villars regularization does render divergent integrals finite while allowing one to recover the original divergent integral if the mass of the fictitious particles $\theta \rightarrow \infty$, and thus satisfies the two constraints. However, the method also requires the fictitious particles to violate the spin-statistics theorem, introduces a violation of the assumed unitarity of the time evolution operator in quantum theory, and typically breaks any gauge symmetries the theory might possess. From the perspective of the third desirable constraint, it is thus a less than ideal regularization method. Another very useful regularization method involves replacing Minkowski spacetime with a Euclidean hypercubic lattice with lattice spacing $\frac{1}{\theta}$. This satisfies the two constraints by restricting integrals $\mathcal{I}(\theta)$ to include only momenta k such that $-\frac{\pi}{\theta} \leq k \leq \frac{\pi}{\theta}$, and sending the lattice spacing to zero then returns the originally divergent integral. It also preserves any gauge invariance the theory may have, but violates Poincaré invariance. One typically has to make such choices about which parts of the mathematical structure of the theory can be altered when choosing a regularization method. Ideally the original mathematical structure of the theory will be restored when the regulator is eliminated, but in many cases this is a subtle question.⁸

The second step is renormalization. The renormalization procedure iden-

⁸The Nielsen-Ninomiya theorem, for example, captures a sense in which a lattice regulation introduces mathematical artifacts that cannot be eliminated by naively letting the lattice spacing $\frac{1}{\theta} \rightarrow 0$.

tifies the part of $\mathcal{I}(\theta)$ which causes the divergence as $\theta \rightarrow \infty$, then shifts the couplings in the action in such a way as to eliminate the divergent pieces of $\mathcal{I}(\theta)$. The modified action is called the *renormalized* action, while the original action is called *bare* and the shifts of the couplings are called *counterterms*. Schematically, the couplings are shifted from $\lambda_0 \rightarrow \lambda_R = \lambda_0 - \delta\lambda$ where $\delta\lambda$ is a counterterm chosen to cancel a divergent contribution the integrals $\mathcal{I}(\theta)$ appearing in the perturbative calculation of $G^{(4)}$. If one is lucky, then shifting finitely many couplings will eliminate *all* of the divergences that might appear in the perturbative calculation of *all* of the n -point functions of the theory (not only for the calculation of $G^{(4)}$!); in that case one says the theory is *renormalizable*.

Once one has shifted the couplings in the action to their renormalized values, one can then safely remove the θ -dependence of the integrals \mathcal{I}_n by letting $\theta \rightarrow \infty$ *without* reintroducing the original divergence. After renormalization, all of the resulting integrals appearing in any perturbative calculation of any n -point function will be finite and independent of the regulator.⁹ While the values of the masses and couplings in the bare Lagrangian are arbitrarily specifiable parameters, the values of the couplings in the renormalized Lagrangian are the *physical* values of those couplings and must be extracted from experiment. Once those couplings have been measured at one scale for one physical process (e.g. $2 \rightarrow 2$ particle scattering at the scale $\mu \sim 100$ GeV), they can be used to calculate amplitudes for any other physical process at any other scale.

This two-step procedure of regularization and renormalization afforded a classification scheme for quantum field theories: if one can eliminate all of the divergences appearing in $\sum^n \lambda^n \mathcal{I}_n$ by making *finitely* many shifts of the bare couplings in the action, then the theory is *perturbatively renormalizable*. The theory is *perturbatively nonrenormalizable* if each higher order term in $\sum^n \lambda^n \mathcal{I}_n$ contains a divergence with a novel structure that, to be eliminated,

⁹Assuming that the theory is renormalizable, of course.

would require a corresponding novel interaction to be added to the bare Lagrangian. Nonrenormalizable theories require an action containing infinitely many interactions to cancel all of the divergences in perturbative calculations of quantities of empirical interest, such as scattering amplitudes. It was long thought that nonrenormalizable QFTs were physically nonsensical, and the requirement that any physically meaningful QFT be perturbatively renormalizable was a guiding principle in particle physics through the 1970s.

In order to make this concrete, consider the simple ϕ^4 theory and imagine computing a particular scattering amplitude: the four-point function $G^{(4)}$ representing (again, loosely) the amplitude for two excitations of a scalar field to scatter off one another. The integrals \mathcal{I}_n appearing in our series expansion will be integrals over allowed momenta of field excitations, and will generically take the form

$$\mathcal{I}_n \propto \lambda^n \int^\infty \frac{d^4k}{(2\pi)^4} \frac{i}{f(k^2)}$$

Consider only the first two terms of the series expansion of $G^{(4)} = \lambda\mathcal{I}_1 + \lambda^2\mathcal{I}_2 + \mathcal{O}(\lambda^3)$. The first term in this series (the so-called “tree-level” term) is finite, and the integral is equal to $-i\lambda$. The second term – so-called “one-loop” level – is in fact a sum of three integrals, each of which has the form

$$\mathcal{I}_2 \propto (-i\lambda)^2 \int^\infty \frac{d^4k}{(2\pi)^4} \frac{i}{k^2 - m^2} \frac{i}{(C - k)^2 - m^2}$$

where the constant C represents one of three Mandelstam variables s , t , and u which are (roughly) a function of the square of the energy at which the particles are scattered. For sufficiently large values of the momentum k , the mass and Mandelstam variables can be neglected. In that high-momentum regime, the integrals \mathcal{I}_2 behave like $\int^\infty \frac{d^4k}{k^4}$, which diverges logarithmically (like $\lim_{k \rightarrow \infty} \ln(k)$) when one allows for the possibility that the field excitations might possess arbitrarily high momenta at some intermediate stage of the scattering process. Since the divergence comes from the high-momentum

behavior of $G^{(4)}$, these are called *ultraviolet divergences*.¹⁰

In order to render this integral mathematically sensible, one first must regularize it. The most physically transparent regularization simply ignores contributions of high-momenta field excitations to our scattering amplitude by sharply cutting off the integral at some large, but finite, value Λ . This makes \mathcal{I}_2 a function of Λ ; explicitly, we get that $\mathcal{I}_2 = 2i \ln\left(\frac{\Lambda^2}{C^2}\right)$. Note that, as required above, this regularization reproduces the logarithmically divergent integral if we let $\Lambda \rightarrow \infty$. After regularization, the amplitude $G^{(4)}$ to one-loop reads:

$$G^{(4)} = -i\lambda + -i\lambda^2 \left[\ln\left(\frac{\Lambda^2}{s}\right) + \ln\left(\frac{\Lambda^2}{t}\right) + \ln\left(\frac{\Lambda^2}{u}\right) \right] + \mathcal{O}(\lambda^3)$$

where I have replaced the stand-in variable C with the Mandelstam variables for which it was a placeholder. What is important to note about this amplitude is that it now depends explicitly on the arbitrarily chosen value Λ for the cutoff. This should be alarming; what if one had chosen Λ^2 ? What about some $\Lambda' \ll \Lambda$? The n -point functions like $G^{(4)}$ contain all of the empirical content of QFT, and that content should not be a function of arbitrary choices of theorists. Ultimately it will be RG methods that provide satisfying answers to these questions.

One natural way of eliminating the Λ -dependence of $G^{(4)}$ is to go out and extract from experiment the value of the coupling λ by scattering particles at some particular energy scale corresponding to particular values of the Mandelstam variables s , t , and u . In fact, for simplicity one can set $s = t = u = \mu^2$ and imagine extracting the value of λ from a scattering experiment at the scale μ . Call this value of the coupling $\lambda_R(\mu)$ (R for “renormalized”).

It turns out that one can replace the Λ -dependent coupling that we have

¹⁰It is also true that divergences arise from the arbitrarily *low* momentum behavior of the integral. These are called infrared divergences; while they raise their own set of conceptual questions, I will neglect them in this chapter.

been using – the bare coupling λ – with the measured, renormalized coupling $\lambda_R(\mu)$ in the calculation of the scattering amplitude $G^{(4)}$. What one finds is that

$$G^{(4)} = -i\lambda_R + i\lambda_R^2 \left[\ln\left(\frac{\mu^2}{s}\right) + \ln\left(\frac{\mu^2}{t}\right) + \ln\left(\frac{\mu^2}{u}\right) \right] + \mathcal{O}(\lambda_R^3)$$

The dependence on Λ has dropped out! One can safely take $\Lambda \rightarrow \infty$ without reintroducing divergences into our calculation, completing the second stage of the regularization and renormalization procedure sketched above.

However, the scattering amplitude now depends on the ratio between the energy scale at which the value of λ_R was measured and the energy scale at which we are scattering our particles, captured (roughly) by the Mandelstam variables s , t , and u . What’s so special about μ ? What would have happened if we chose to measure the value of the coupling at some other scale μ' and carried out our calculation with $\lambda_R(\mu')$ instead? A skeptic might think that while the introduction of $\lambda_R(\mu)$ eliminated the dependence of $G^{(4)}$ on any arbitrarily cutoff Λ , it replaced it with an equally arbitrary dependence on the scale μ at which we chose to measure the value of $\lambda_R(\mu)$. Attempting to answer this skeptic leads naturally to the renormalization group.

3 The renormalization group

Suppose that, unsettled by the skeptic, one wanted to see how the ϕ^4 theory with its coupling defined by its measured value at the scale μ – henceforth “the theory defined at μ ” – relates to the theory defined at μ' . First consider the expression for the four-point correlation function in terms of $\lambda_R(\mu)$:

$$G^{(4)} = -i\lambda_R(\mu) + i\lambda_R^2(\mu) \left[\ln\left(\frac{\mu^2}{s}\right) + \ln\left(\frac{\mu^2}{t}\right) + \ln\left(\frac{\mu^2}{u}\right) \right] + \mathcal{O}(\lambda_R^3)$$

Suppose that one chose instead to use $\lambda_R(\mu')$, the value of the coupling measured at μ' . This amounts to making the simple replacement

$$G^{(4)} = -i\lambda_R(\mu') + i\lambda_R^2(\mu') \left[\ln\left(\frac{\mu'^2}{s}\right) + \ln\left(\frac{\mu'^2}{t}\right) + \ln\left(\frac{\mu'^2}{u}\right) \right] + \mathcal{O}(\lambda_R^3)$$

In order to see how these two quantities relate, subtract the latter from the former to express $\lambda_R(\mu')$ in terms of $\lambda_R(\mu)$:

$$\lambda_R(\mu') = \lambda_R(\mu) + 3\lambda_R(\mu) \ln\left(\frac{\mu'^2}{\mu^2}\right) + \mathcal{O}(\lambda_R^3)$$

For $\mu' = \mu - d\mu$, this can be expressed as a differential equation for the way the coupling changes with scale:

$$\mu \frac{d}{d\mu} \lambda_R(\mu) = 6\lambda_R(\mu)^2 + \mathcal{O}(\lambda_R^3)$$

With a few lines of algebra, we have derived the *renormalization group equation* for the coupling $\lambda_R(\mu)$. The right-hand side of this equation, often written $\beta(\lambda)$, is called the *beta function* for the coupling λ_R , and captures how the value of the coupling changes as a function of the energy scale at which it is measured. One thing that becomes immediately apparent is that the notion of a “coupling constant” in QFT is a misnomer: the physical couplings that appear in the renormalized Lagrangian are functions of an energy scale. This may seem a startling result, but it is empirically quite well-confirmed. The fine-structure constant, for instance, is rather famously said to have the value $\alpha = \frac{1}{137}$. However, the renormalization group teaches us that it has the value *only at the relatively low energies at which we were historically able to measure it*. More recent experiments have confirmed the scale-dependent variation of α , with the value of α at the energy scale $\mu \sim 80$ GeV measured to be $\alpha(\mu) \approx \frac{1}{128}$ (Patrignani and al., 2016 and 2017 update).

Though the path to its derivation was relatively simple, the implications

of this scale-dependence for our understanding of QFT are vast. I will first develop these implications in a somewhat restricted context, then pair them explicitly with the path integral formulation of QFT that I said above was a natural home of renormalization group methods.

The first thing to note is that *all* of the couplings in a QFT will vary with scale according to a renormalization group equation. For example, suppose that one added a (nonrenormalizable) interaction $\lambda_6\phi^6$ to the action of our scalar field theory; the coupling λ_6 would have its own beta function that would capture how that coupling changed as a function of μ . Similarly, there is also a renormalization group equation determining how the (effective) mass m of the scalar field varies with scale. In general, the beta function for any individual coupling g_i is a function of *all* of the couplings in the theory:

$$\mu \frac{d}{d\mu} g_i(\mu) = \beta_i(g_1, \dots, g_n)$$

Of particular interest is the behavior of a theory's couplings at very high energies. It provides some guidance about whether the theory could have a well-defined continuum limit; an interesting mathematical question, albeit one of limited physical import.¹¹ The renormalization group equation tells us that there are three possible behaviors that a theory can have in that asymptotic regime:

1. The renormalization group flow hits a point $g^* \equiv (g_1^*, \dots, g_n^*)$ at which the beta functions β_i for *all* of the couplings are zero. This is called a *fixed point*, and the QFT becomes *scale-invariant* at this point.
2. The beta function for at least one of the couplings in the theory is positive, entailing that the coupling *increases* with higher energies. If the coupling does not hit a fixed point, it eventually becomes infinite at some very large, but finite, energy scale μ_{Landau} and the theory is

¹¹See (Li, 2015) for a philosophical discussion of the uses of RG analyses in the mathematically rigorous setting of constructive QFT.

ill-defined; one says that its RG equations encounter a *Landau pole*. QFTs with this asymptotic behavior are called *trivial*, since they can only be well-defined if the values of the couplings are zero. Conversely, as $\mu \rightarrow 0$ the couplings go to zero, and these theories are *infrared free*.

3. The beta functions for all of the couplings in the theory are negative, entailing that the couplings *decrease* with higher energies. Eventually all of the couplings go to zero as $\mu \rightarrow \infty$, and the QFT hits a particular fixed point: the point $g^* = 0$. In this case, the QFT is said to be *asymptotically free*. Conversely, at low energies the couplings get large, eventually diverging at some long, but finite, energy scale; this is responsible for the phenomenon of *confinement* in quantum chromodynamics.¹²

Analysis of the high-energy behavior of a theory's RG flow allows for a more general notion of renormalizability than was described above. Recall that a theory was perturbatively renormalizable if all divergences arising in the perturbative calculations of its n -point functions could be eliminated by shifting the values of finitely many of its couplings. What made these QFTs noteworthy was that they seemingly remained consistent up to arbitrarily high energies, in the sense that one could consistently let the cutoff $\Lambda \rightarrow \infty$ at the end of calculations. What the RG analysis of the possible high-energy behaviors of a QFT offers is a sense of *nonperturbative renormalizability*: a theory is said to be nonperturbatively renormalizable if its RG flow hits a fixed point when one examines the structure of the theory as $\mu \rightarrow \infty$. Interestingly, this is a sense of renormalizability that can apply to QFTs that are *not* perturbatively renormalizable; the Gross-Neveu model in 3 spacetime dimensions, for example, is perturbatively nonrenormalizable but

¹²Strictly speaking, asymptotic freedom is just the special case of asymptotic safety where the fixed point $g^* = 0$. Its significance in modern particle physics makes it a special case worth highlighting.

asymptotes to a fixed point at high energies and has a perfectly sensible structure at arbitrarily high energies (Braun et al., 2011).

As I will discuss in more detail below, RG methods also naturally lead to a picture of a *space of QFTs* \mathcal{A} that is coordinatized by the values of the couplings g_1, \dots, g_n . The RG equations generate a “dynamics” on this theory space, determining how a theory – a point in \mathcal{A} defined by a particular set of couplings defined at a particular energy scale – will “flow” through theory space as one changes the energy scale at which one is studying the structure of the QFT.¹³ The fixed point(s) of the theory play a particularly important role in making sense of this notion of a space of QFTs. In order to appreciate the power of this picture, let us return to the path integral framework introduced earlier and discuss a particular approach to the RG, the so-called “Wilsonian RG”.¹⁴

Recall that in the perturbative calculation of $G^{(4)}$ above, we encountered integrals in the series expansion that were divergent. Those integrals were regularized by sharply cutting off our range of integration at some high but finite momentum Λ . At the level of the path integral, in which one integrates over entire histories of possible configurations of the field, this regularization procedure amounts to treating histories of the field $\phi(x)$ which may differ on length scales $L \leq \frac{1}{\Lambda}$ as equivalent. In the action, this amounts to replacing $\phi(x) \rightarrow \phi_\Lambda(x)$ and replacing the “bare” couplings g with the the scale dependent, renormalized coupling(s) $g_R(\Lambda)$, emphasizing the fact that one is now dealing with a theory defined only up to the scale Λ . This automatically regularizes the divergent integrals that appear in perturbative calculations by excluding the high momenta that gave rise to the divergences. The resulting theory is defined only up to the scale of the cutoff Λ , but inevitably becomes

¹³This notion of a “space of theories” is, in general, not mathematically well-defined. See (Douglas, 2013) for an interesting discussion of theory space and the current state of our understanding of it.

¹⁴The *locus classicus* for this formulation is (Wilson and Kogut, 1974); although it is beyond the scope of this chapter, there are interesting relationships between several different formulations of the RG. See (Schwartz, 2014, chapter 23) for a textbook discussion.

inapplicable at that scale; such a QFT is called an *effective field theory* and the action that defines such a theory is called the *Wilsonian effective action* \mathcal{S}_W . As (Wallace, this volume) describes, that QFTs are effective field theories is now the dominant attitude in the physics community toward the foundations of QFT. Although it should be clear from the preceding paragraph, perhaps it is worth emphasizing that in an effective field theory there is no problem of divergences whatsoever, and the renormalization of the couplings has nothing to do with the elimination of any infinities.

In the context of perturbative renormalization, the requirement that a theory's action not include any nonrenormalizable interaction terms made sense: one must include infinitely many renormalized parameters to cancel all divergences in such a theory and since values of renormalized parameters must be taken from experiment, these theories were useless for making predictions: one would need to perform infinitely many measurements before calculating anything at all. Steven Weinberg, writing just before the understanding of QFTs as effective field theories became widespread, exemplified this attitude:

Throughout this history I have put great emphasis on the condition of renormalizability, the requirement that it should be possible to eliminate all infinities in a quantum field theory by a redefinition of a small number of physical parameters...it has always seemed to me that the requirement of renormalizability has just the kind of restrictiveness that we need in a fundamental physical theory. (Weinberg, 1977, p. 33)

In the effective field theory approach, there is no longer any justification for including only renormalizable terms in the action. As a result, a Wilsonian action will generally include *infinitely* many terms. For example, the action for the ϕ^4 theory we have been considering becomes

$$\begin{aligned}
S_W &= \int d^4x \frac{1}{2} (\partial\phi_\Lambda)^2 + m^2\phi_\Lambda^2 + \lambda_4\phi_\Lambda^4 + \lambda_6\phi_\Lambda^6 + \dots \\
&= \frac{1}{2} (\partial\phi_\Lambda)^2 + \sum_{n \geq 0} \lambda_n \phi_\Lambda^{2+n} + \sum_{n \geq 0} \lambda'_n (\partial\phi_\Lambda)^2 \phi_\Lambda^n + \dots = \mathcal{L}_\Lambda(\phi, \partial\phi, \dots)
\end{aligned}$$

One might think that this is sufficiently unwieldy to be useless. However, as we will see below, one of the virtues of RG methods is that they demonstrate that only a finite set of interactions – precisely the renormalizable ones! – contribute meaningfully to the low energy physics to which we have experimental access.

It is part and parcel of the effective field theory approach that a theory is appropriate only for describing physical processes at energy scales $E \ll \Lambda$ and one doesn't trust the theory's description of physical processes at energy scales near the cutoff. However, an effective field theory defined at the cutoff Λ may still contain many high-energy degrees of freedom that are inappropriate for describing physics at some particular $E \ll \Lambda$. Suppose, for example, one wants to model the elastic scattering of protons at $E \sim 1$ GeV, but only has available to them quantum chromodynamics defined at (say) $\Lambda \sim 10^{15}$ GeV, which will contain quark and gluon fields as its fundamental degrees of freedom. This violates the truism with which I began this chapter. In order to achieve a more perspicuous description of the low energy nuclear physics of interest, one ought to remove some of these high-energy degrees of freedom from the path integral by lowering the scale of the cutoff.

Essentially, one achieves this by doing the path integral a little bit at a time. Wilson's approach involves splitting the the path integral of the “full-theory” into 2 parts, a “high-energy” component and a “low-energy” component. This is achieved by a splitting of the field $\phi(x)$ into $\phi(x) = \phi_H(x) + \phi_L(x)$. The high-energy component describes excitations of the field whose allowed energies E are contained in a momentum “shell” of width $d\Lambda$, i.e. $(\Lambda - d\Lambda) < E < \Lambda$. In order to lower the cutoff from $\Lambda \rightarrow \Lambda' = (\Lambda - d\Lambda)$,

one simply computes the high-energy part of the path integral and leaves the low-energy part alone.¹⁵ The resulting QFT contains only the low-energy components of the field and is defined only up to the new, lower cutoff scale Λ' ; one sometimes describes this process of integrating out the high-energy component of the action as “putting on blurry glasses” that prevent us from seeing excitations of the field on length scales shorter than $\frac{1}{\Lambda}$.

It is integral to the tenability of the effective field theory approach that the empirical consequences of the high-energy part of the path integral can be *entirely* incorporated into the new, low-energy theory via a modification of its couplings. The specific manner in which the couplings must be modified as one changes the scale of the cutoff is, perhaps unsurprisingly, given by the beta functions of the couplings. This procedure of “integrating out” high energy degrees of freedom can be iterated in the obvious way, with each change in the scale of the cutoff requiring a corresponding change in the value of the couplings of the theory, with the precise form of the change determined by the RG equations. The Wilsonian picture makes clear the physical import of the RG equations: the change of the couplings as a function of the energy scale should be understood as reflecting the way that the *physical effects* of field excitations at energies $\Lambda > \Lambda'$ are incorporated into a new Wilsonian action defined at a lower energy scale Λ' .

We can now join this description of the Wilsonian RG to our above discussion of the RG equations as defining a “flow” through a space of QFTs \mathcal{A} . In principle, one would like to begin by first pinpointing *all* of the fixed points of *all* QFTs. This is obviously impossible in practice – not only do we not know all the QFTs that could be consistently written down, but the mathematical techniques available to us are insufficient for identifying all of the fixed points even of the QFTs we *do* know how to consistently write down. Nevertheless, one can get the general idea by considering only free

¹⁵In general, this computation will be done using the perturbative approximation sketched above.

QFTs, which are the simplest possible fixed points. These will form an infinite set of isolated points in \mathcal{A} . This leads naturally to an interesting picture, somewhat common in the physics community, of what it means to *define* a quantum field theory.¹⁶ One begins with a scale-invariant QFT defined by an action at a fixed point and “perturbs” it by adding an infinite sum of local interactions to its action, which will shift the theory away from the fixed point.¹⁷ Following the reasoning sketched above, one then cuts the theory off at some energy scale Λ . This generates a Wilsonian action \mathcal{S}_W , defined at a particular scale. One can then use the RG to analyze how this theory “moves” through the space \mathcal{A} as one iteratively lowers the cutoff and changes the values of the couplings.

This definition of a QFT as a perturbation of a fixed-point action also offers a natural classification of the possible “perturbing” interactions appearing in the action into three categories: *relevant*, *marginal*, and *irrelevant*. The interactions are classified according to the RG behavior of their associated couplings in the neighborhood of the fixed point:

1. An interaction is *relevant* if its coupling *increases* at low energies, flowing *away* from the fixed point as one integrates high-energy degrees of freedom out of the path integral.
2. An interaction is *irrelevant* if its coupling *decreases* at low energies, flowing *back into* the fixed point.
3. An interaction is *marginal* if its coupling does not change with scale, and the interaction contributes equally to physical processes at all length scales.¹⁸

¹⁶See, for example, (Banks, 2008, chapter 9).

¹⁷Though it is beyond the scope of this review, this will also break the scale invariance of the theory, the study of which is an important part of contemporary QFT.

¹⁸Often a more sophisticated analysis will show a coupling to be marginal at the lowest-order of perturbation theory, but when higher-order corrections are included the coupling is shown to scale like a relevant or irrelevant interaction. In this case, the interaction is called *marginally relevant* or *marginally irrelevant*.

There are two important things to say at this point. The first answers the worry that although the Wilsonian action did not generate perturbative expansions that contained divergent integrals, this virtue was outweighed by the fact that it contained infinitely many interactions and so was too unwieldy to be useful. It also provides an *explanation* of why the criterion of renormalizability was successful as a constraint on theorizing through the 1970s.

In any given Wilsonian action \mathcal{S}_W , at most a finite number of interactions will be renormalizable, with the vast majority of the interactions in \mathcal{S}_W rendering the theory nonrenormalizable; for example, the only two renormalizable interactions one can write down in the ϕ^4 theory in four space-time dimensions are $m^2\phi^2$ and $\lambda\phi^4$ and in quantum electrodynamics, there is only one possible renormalizable interaction between the photon and electron fields. If the right way of understanding a QFT requires including infinitely many terms in the action, how could these theories have been so successful?

The answer goes as follows.¹⁹ Beginning with any arbitrary Wilsonian action, one can begin iteratively lowering the cutoff and studying the way that the couplings change. RG methods reveal that the couplings associated with all of the *nonrenormalizable* interactions will make contributions to low energy physics in only two ways: (i) through a rescaling of the couplings associated with *renormalizable* interactions and (ii) through contributions to scattering amplitudes which suppressed by inverse powers of the high-energy cutoff Λ , rendering them negligible for describing low-energy physics. In terms of an RG flow in theory space \mathcal{A} : one finds that a generic Wilsonian action, which lives at the point on some infinite-dimensional manifold in \mathcal{A} picked out by the values of the couplings associated with its infinitely-many interactions, flows at low energies to a finite-dimensional manifold defined by the values of only its *renormalizable* interactions. The result is that at the low energies for which QFT has proved empirically adequate, only a finite

¹⁹See (Duncan, 2012, chapter 17.4) for details.

number of interactions make non-negligible contributions to the computation of the low-energy scattering amplitudes that we have been able to test experimentally: precisely the renormalizable interactions! Thus what seemed like an inexplicable bit of good luck to physicists in the middle of the 20th century is rendered straightforwardly explicable by RG methods.

The second thing to say about the notion of theory space \mathcal{A} is that it affords us a means for partitioning the space of QFTs into classes of theories defined by their behavior under RG transformations. As (Batterman, this volume) explains, many different QFTs may flow to the *same* fixed point, and we say that these QFTs form a particular *universality class*; in principle, RG methods offer the possibility of partitioning the entirety of \mathcal{A} into distinct universality classes. Though we are very far from the attainment of anything like such a classification of all QFTs, the fact that RG methods make such a classification possible, even in principle, illustrates the far-reaching import of RG methods for our understanding of QFT. While renormalization may have begun as a mathematically and physically dubious procedure for eliminating infinities from perturbative calculations, RG methods are now at the very heart of virtually every aspect of the modern understanding of QFT. Since I can hardly emphasize this centrality better myself, I will conclude with an update to Eddington's famous words about the second law of thermodynamics proposed by Vincent Rivasseau:

The flow of the renormalization group holds, I think, the supreme position among the laws of Nature. If someone points out to you that your pet theory of the universe is in disagreement with Maxwell's equations - then so much the worse for Maxwell's equations. If it is found to be contradicted by observation - well, these experimentalists do bungle things sometimes. But if your theory is found to be against the flow of the renormalization group I can give you no hope; there is nothing for it but to collapse in deepest humiliation. (Rivasseau, 2012, p. 28)

4 Further Reading

RG methods will form a significant portion of any modern introduction to QFT. My recommendations for further QFT reading essentially mirror those of Wallace (this volume), with the addition of (Schwartz, 2014) which contains a helpful discussion of several topics related to renormalization and RG methods. There are also several book-length treatments of RG methods both in QFT and classical statistical physics; I can recommend (Collins, 1984), (Rivasseau, 1991), (Goldenfeld, 1992), (Cardy, 1996), and (Hollowood, 2013). A longer survey of renormalization explicitly aimed at philosophers is (Butterfield and Bouatta, 2015).

RG methods have also been put to considerable philosophical use, most notably in debates concerning emergence and reduction. A considerable literature has developed around the discussions of renormalization and intertheoretic reduction in (Batterman, 2000) and (Batterman, 2002); see e.g. (Butterfield, 2014) or (Reutlinger, 2014) and references therein. Recently it has also been suggested by (Fraser, 2016) and (Williams, 2017) that RG methods make possible a promising application of the *divide et impera* approach to scientific realism in QFT; see (Ruetsche, forthcoming) for a skeptical response.

References

- T. Banks. *Modern Quantum Field Theory*. Cambridge University Press, Cambridge, 2008.
- R. Batterman. Multiple realizability and universality. *The British Journal for the Philosophy of Science*, 51(1):115–145, 2000.
- R. Batterman. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press, Oxford, 2002.

- R. Batterman. The tyranny of scales. In R. Batterman, editor, *Oxford Handbook of Philosophy of Physics*, pages 255–286. Oxford University Press, Oxford, 2013.
- J. Braun, H. Gies, and D. Scherer. Asymptotic safety: a simple example. *Physical Review D*, 83(8), 2011.
- J. Butterfield. Reduction, emergence, and renormalization. *The Journal of Philosophy*, 111(1):5–49, 2014.
- J. Butterfield and N. Bouatta. Renormalization for philosophers. In *Metaphysics in Contemporary Physics*, volume 104 of *Poznan Studies in the Philosophy of the Sciences and the Humanities*, pages 437–485. 2015.
- J. Cardy. *Scaling and Renormalization in Statistical Physics*. Cambridge University Press, Cambridge, 1996.
- J. Collins. *Renormalization: an introduction to renormalization, the renormalization group and the operator-product expansion*. Cambridge university press, 1984.
- J. T. Cushing. *Theory construction and selection in modern physics: The S matrix*. Cambridge University Press, Cambridge, 1990.
- M. R. Douglas. Spaces of quantum field theories. In *Journal of Physics: Conference Series*, volume 462, 2013.
- A. Duncan. *The Conceptual Framework of Quantum Field Theory*. Oxford University Press, Oxford, 2012.
- R. Feynman and A. Hibbs. *Quantum Mechanics and Path Integrals*. Dover, New York, 1965/2010. edition emended by D. Styer.
- J. Fraser. *What is Quantum Field Theory? Idealisation, Explanation, and Realism in High Energy Physics*. PhD thesis, University of Leeds, 2016.

- J. Glimm and A. Jaffe. *Quantum Physics: A Functional Integral Point of View*. Springer, New York, 1987.
- N. Goldenfeld. *Lectures on Phase Transitions and the Renormalization Group*. Addison-Wesley, Reading, 1992.
- D. J. Gross. Asymptotic freedom and QCD—a historical perspective. *Nuclear Physics B-Proceedings Supplements*, 135:193–211, 2004.
- T. J. Hollowood. *Renormalization Group and Fixed Points in Quantum Field Theory*. Springer, London, 2013.
- B. Li. Coarse-graining as a route to microscopic physics: the renormalization group in quantum field theory. *Philosophy of Science*, 82(5):1211–1223, 2015.
- M. Niedermaier and M. Reuter. The asymptotic safety scenario in quantum gravity. *Living Reviews in Relativity*, 9(5), 2006.
- C. Patrignani and e. al. Review of particle physics. *Chin. Phys. C*, 40(10): 100001, 2016 and 2017 update.
- A. Reutlinger. Why is there universal macrobehavior? renormalization group explanation as noncausal explanation. *Philosophy of science*, 81(5):1157–1170, 2014.
- V. Rivasseau. *From Constructive to Perturbative Renormalization*. Princeton University Press, Princeton, 1991.
- V. Rivasseau. Quantum gravity and renormalization: the tensor track. In *AIP Conference Proceedings 8*, volume 1444, pages 18–29, 2012.
- L. Ruetsche. Renormalization group realism: the ascent of pessimism. *Philosophy of Science*, forthcoming.

- M. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, Cambridge, 2014.
- S. Weinberg. The search for unity: Notes for a history of quantum field theory. *Daedalus*, pages 17–35, 1977.
- S. Weinberg. Why the renormalization group is a good thing. In A. Guth, K. Huang, and R. L. Jaffe, editors, *Asymptotic Realms of Physics: Essays in Honor of Francis Low*, pages 1–19. The MIT Press, Cambridge, 1983.
- P. Williams. Scientific realism made effective. *The British Journal for the Philosophy of Science*, 2017.
- K. G. Wilson. Model hamiltonians for local quantum field theory. *Physical Review*, 140(2B):445–457, 1965.
- K. G. Wilson and J. Kogut. The renormalization group and the ϵ -expansion. *Physics Reports*, 12(2):75–199, 1974.
- A. Zee. *Quantum Field Theory in a Nutshell*. Princeton University Press, Princeton, 2010.