

Renormalization Scrutinized

Sébastien Rivat*

Abstract

In this paper, I propose a general framework for understanding renormalization by drawing on the distinction between effective and continuum Quantum Field Theories (QFTs), and offer a comprehensive account of perturbative renormalization on this basis. My central claim is that the effective approach to renormalization provides a more physically perspicuous, conceptually coherent and widely applicable framework to construct perturbative QFTs than the continuum approach. I also show how a careful comparison between the two approaches: (i) helps to dispel the mystery surrounding the success of the renormalization procedure; (ii) clarifies the various notions of renormalizability; and (iii) gives reasons to temper Butterfield and Bouatta's claim that some continuum QFTs are ripe for metaphysical inquiry (Butterfield and Bouatta, 2014).

1 Introduction

Renormalization is one of those great success stories in physics that fly in the face of philosophers' ideals of scientific methodology. Quantum Field Theories (QFTs) have been known to be plagued by mathematical infinities since the 1930s and it was only in the late 1940s that physicists had their first significant victory by developing appropriate renormalization techniques. It could have been hoped that they would eventually construct a realistic QFT from first principles without using these techniques; but even after seventy years, this has not been the case. Our best QFTs are still constructed by means of conceptually odd and *ad hoc* renormalization techniques. One notable example is to isolate and cancel infinite quantities by shifting the dimension of space-time by some infinitesimal amount. Another one is to simply impose some arbitrary restriction on the range of distance scales of the theory.

Among the philosophers who take the formulation of QFT most widely adopted by physicists seriously, it has become standard to appeal to the Renormalization Group (RG) theory in order to explain the unlikely success of renormalization. For instance, Huggett and Weingard (1995, sec. 2) emphasize that the RG provides the appropriate tools for identifying the class of well-defined continuum QFTs and dispels the interpretative worries related to cancellations of infinities in perturbation theory. To give another example, although with a different understanding of QFT this time, Wallace (2006, pp. 48-50; 2011, sec. 4) relies

*Department of Philosophy, Columbia University. Email: sr3109@columbia.edu

on RG-based considerations to dispel the interpretative worries related to the crude and arbitrary implementation of a physically meaningful cut-off.

Those philosophers are right to emphasize the role and the importance of the RG in contemporary physics. But there are reasons to be dissatisfied. Of central importance is the failure to appreciate the existence of conceptually distinct modern formulations of renormalization, RG included. Consider for instance Huggett and Weingard’s attempt at clarifying renormalization in the case of continuum QFTs. If by ‘RG’ they mean the Gell-Mann & Low RG, then their account does not really dissolve the methodological worries that physicists had in the 1940s. The delicate fine-tuning of theories in the infinite cut-off limit is nothing but the old-fashioned cancellation of infinities in a different guise. On the other hand, if by ‘RG’ they mean the Wilsonian RG, then their account does not properly deal with continuum QFTs. At least as we traditionally understand it, the Wilsonian RG is built on the idea of integrating out high energy degrees of freedom and restricting the applicability of the resulting theories to sufficiently large distance scales (e.g., Weinberg, 1995, sec. 12.4; Schwartz, 2013, chap. 23).

To give another example, Cao and Schweber (1993) somewhat overstate the triumph of the modern Wilsonian renormalization programme. Many renormalization techniques conceptually akin to the approach of the late 1940s are still the “industry standard” in high energy physics, as Hollowood (2013, p. 3) felicitously puts it. These techniques include modern regularization methods such as dimensional regularization in standard QFTs and regularization by dimensional reduction in supersymmetric QFTs. More importantly perhaps, the Wilsonian RG does not fully dispel the traditional mathematical, conceptual and methodological worries associated with renormalization. With regard to methodology, for instance, one might be concerned about the infinite number of independent parameters typically required to compensate for the uncertainty associated with the exact value of a physically meaningful cut-off.

The main goal of this paper is to offer a more accurate and systematic way of understanding the overall conceptual structure of renormalization. For this purpose, I will distinguish between the “effective” and the “continuum” approach to renormalization and show that all the important features of perturbative renormalization can be understood along this distinction. The idea is simple: current working QFTs in high energy physics are understood and formulated either as continuum QFTs or as effective QFTs, and each of these two types of QFTs is associated with a specific methodology of theory-construction—or at least, given the diversity of renormalization techniques, each of them is most conceptually consistent with a specific methodology. In the effective approach, the domain of applicability of the theory is restricted by a physically meaningful short distance scale and the structure of the theory adjusted by including the appropriate sensitivity to the physics beyond this scale. Here, the goal is to focus on the appropriate low energy degrees of freedom. In the continuum approach, the theory is defined across all dis-

tance scales and its structure adjusted according to the physical scale of interest. Here, the goal is to define a putatively fundamental QFT and resist the suggestion that realistic QFTs are ultimately to be understood and formulated as phenomenological theories restricted to some limited range of distance scales.

The central claim of this paper is that the effective approach provides a more physically perspicuous, conceptually coherent and widely applicable framework to construct perturbative QFTs than the continuum approach. I will defend this claim by showing, in detail, how the steps underlying the perturbative construction of an effective QFT are physically justified and how the resulting parts of the theory are physically meaningful, unambiguously characterized and coherently related to one another—and this independently of the particular local QFT considered. And I will show how a careful comparison between the two approaches: (i) helps to dispel the mystery surrounding the success of the renormalization procedure discussed early on (e.g., Teller, 1988, 1989; Huggett and Weingard, 1995, 1996) but never fully dispelled in my sense, not even in the most recent literature (e.g., Butterfield and Bouatta, 2015; Crowther and Linnemann, 2017; Fraser J., 2017; 2018); (ii) helps to clarify the various notions of renormalizability; and (iii) gives reasons to temper Butterfield and Bouatta’s claim that some continuum QFTs are ripe for metaphysical inquiry (Butterfield and Bouatta, 2014; Butterfield, 2014).

The paper is organized as follows. Section 2 introduces the QFT framework and the problem of ultraviolet divergences. Section 3 compares the effective and the continuum approach to the renormalization procedure. Section 4 disentangles the effective and continuum notions of perturbative renormalizability. Sections 5 and 6 briefly compare the effective and the continuum approach to the RG and clarify the scope of the continuum approach.¹ Section 7 examines the implications of the discussion in sections 3-6 for Butterfield and Bouatta’s defense of continuum QFTs.

Three important clarifications before I begin. First, I do not think that the methodological superiority of the effective approach to renormalization offers a sufficient reason to take effective QFTs to be the correct way of understanding QFTs. It is a good step forward. But it needs to be supplemented with a careful analysis of the theoretical virtues of effective QFTs, and this goes beyond the scope of the present paper. Second, I do not mean to claim that the distinction between the effective and the continuum approach is absolutely perfect and exhaustive. All I aim to capture is a set of salient conceptual differences that do not reduce to mere practical differences (e.g., computational simplicity and efficiency). Third, unless otherwise indicated, I will follow Butterfield (2014, pp. 30-31) and understand ‘theory’ in its specific sense throughout the paper, that is to say as given by a particular action, a Lagrangian or a Hamiltonian.

¹For two recent and insightful reviews of the Wilsonian RG, see Williams (2018) and Fraser J. (2018).

2 Relativistic QFT and the Problem of Ultra-violet Divergences

Relativistic Quantum Field Theory (QFT) is the mathematical framework developed by physicists since the late 1920s to extend the tools of quantum mechanics to classical electromagnetism (and more) and to overcome the failure of quantum mechanics to account (among other phenomena) for the creation and annihilation of particles observed in decay experiments.

As its name suggests, a QFT describes the quantum analogue of classical fields, and the simplest way to think about a quantum field is to treat it as a continuous physical system composed of one individual quantum system at each space-time point. Each individual quantum system is associated with at least one independent variable quantity (a “degree of freedom”) determining the set of its possible states, and the possible states of the quantum field over space-time are obtained by combining the state spaces of these individual quantum systems together. From there, things work exactly as in quantum mechanics. A sum of states of the field (a “state superposition”) also defines a possible state of the field. Each state of the field is associated with a possible configuration or “history” of the field specifying a set of values that the field can take over space-time: for instance, one real number $\phi(x)$ at each space-time point for a simple scalar field. The probability for the quantum field to be found in the configuration state $|\phi(x)\rangle$ is given by the absolute square value of the wave functional $\psi[\phi(x)]$ (assuming that we could measure the whole state of the field). And the possible energy excitation states of the field are obtained by representing the possible configuration states of the field in momentum space. One odd thing, however, is that in this picture, a “particle” corresponds to a localized pattern of energy excitations.

Quantum fields are also dynamical physical systems. They vary smoothly over space-time and interact locally at space-time points with other fields and often with themselves too. Physicists typically describe the dynamics of fields by a Lagrangian functional density \mathcal{L} and the strength of interactions by coupling parameters g_i . I will take the ϕ^4 -theory as my main example in what follows:

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

with $\phi(x)$ an arbitrary field configuration of a scalar field, m a mass parameter, and λ a quartic self-interaction coupling (using the Euclidean metric for simplicity). Of crucial importance are the action $S[\phi] = \int d^4x\mathcal{L}$ and the path integral $\mathcal{Z} = \int d[\phi(x)]e^{S[\phi]}$ which give us the different weights $e^{S[\phi]}$ associated with each possible field configuration $\phi(x)$.²

Finally, the correlations between the states of the field at n different

²Of course, the difficulty is that we do not yet have a mathematically rigorous definition of the path integral for realistic continuum QFTs in four dimensions, but I will ignore this problem for now.

space-time points are given by n -point correlation functions $\langle \phi(x_1) \dots \phi(x_n) \rangle$. Roughly speaking, these correlation functions tell us the degree to which the different “parts” of the field are sensitive to one another, i.e., here, the probability (once these functions are squared) that the field is found in a certain state at some space-time points x_1, \dots, x_k given its state at other space-time points x_{k+1}, \dots, x_n ($1 \leq k \leq n-1$). We compute empirical predictions—say, about the probability that two incoming particles decay into two outgoing particles—by calculating the absolute square value of the scattering amplitude Γ between the appropriate asymptotic particle states of the field, with Γ obtained by taking into account all the possible correlations between these states.

These are the basic tools to define and test any QFT. Unfortunately, we face two immediate problems with this “naive” schematic construction if we want to make predictions. The least severe is that realistic QFTs are highly non-linear interacting theories and therefore not exactly solvable by current mathematical means. We can still work out approximate solutions and predictions thanks to perturbation theory: provided the (dimensionless) couplings are small (e.g., $\lambda \ll 1$), scattering amplitudes can be expanded perturbatively as follows:

$$\Gamma = \lambda + \lambda^2 \Gamma_2 + \lambda^3 \Gamma_3 + \dots \quad (2)$$

where each sub-amplitude $\lambda^n \Gamma_n$ represents field correlations between the incoming and outgoing particles given n possible interaction points.³

The most severe, the so-called problem of “ultraviolet” (UV) divergences, is that a large majority of the sub-amplitudes Γ_n actually diverge when we attempt to compute them.⁴ This is clearly a disaster (at least at this stage) since it means that most empirical predictions in QFT are infinite. If we keep all the other assumptions of the theory in place (e.g., four space-time dimensions and standard types of fields, symmetries and interactions), the problem naturally originates from what is known as the continuum assumption:

Continuum assumption: For any extended region of space-time no matter how small, quantum fields have infinitely many interacting degrees of freedom.

In practice, the continuum assumption forces us to take into account correlations over arbitrarily short distances (or, equivalently, over arbi-

³Here one might worry about two things. First, one should be wary not to interpret too quickly these perturbative terms as representing distinct real sub-processes (the so-called “virtual processes”) since they might be interpreted as mere mathematical artifacts of the decomposition of Γ . Let me briefly offer one reason to resist this worry: as we will see shortly, the renormalized coupling λ is a function of an arbitrary mass scale Λ which can be interpreted as the experimental energy E at which we probe the system. Since each $\lambda^n(E) \Gamma_n(E)$ does not vary with the same rate with respect to E , we can evaluate them separately by making successive measurements at different experimental energy scales E . If this succeeds, each term receives independent empirical confirmation. Second, the perturbative series diverges in realistic cases for arbitrarily small but non-zero λ (see Helling, 2012, pp. 1-13, and Duncan, 2012, chap. 11, for more details, and Miller, 2016, for a philosophical discussion). I will ignore this problem too.

⁴I will leave aside the problem of low energy or “infrared” (IR) divergences.

trarily high energies) when calculating a correlation function between any two states of the field. Consider for instance the scattering amplitude $\Gamma(p_1, \dots, p_4)$ in ϕ^4 -theory describing the scattering event of two incoming particles decaying into two outgoing particles. Then, for example, the second order perturbative term $\lambda^2 \Gamma_2$ describes a specific set of correlations which diverge logarithmically in the high energy domain of integration:

$$\Gamma_2 \approx \int^\infty d^4k/k^4 \quad (3)$$

with k a momentum variable. So the problem is that we have to take into account the correlations of the field over arbitrarily short distances and that the values of these correlations are small but sufficiently important once summed up to make Γ_n diverge.⁵ What does it mean physically? To give a rough analogy, it is as if two distinct macroscopic parts of a table were sensitive enough to the individual particles constituting the table for the slightest movement of a particle to significantly affect on average the distance between these two parts. The sensitivity is even more dramatic in the present case. The theory is not just empirically inadequate but also inconsistent as it predicts measurement outcomes with infinite probability (i.e., here, $|\Gamma(p_1, \dots, p_4)|^2$ diverges).

The claim that the problem of UV divergences originates from the continuum assumption is in fact controversial, and physicists have come up with three main types of responses which I will respectively call the “continuum”, the “effective” and the “axiomatic” approach to the problem of UV divergences. According to the continuum approach, the problem arises because we are not working with the correct type of QFT or because we have not appropriately parametrized the QFT at hand in the first place. The hope is that the continuum assumption holds for a specific class of QFTs and that all that needs to be done is to sensibly fine-tune their parameters with the tools of renormalization theory. According to the effective approach, the problem arises because the continuum assumption is false. The solution is to impose explicit restrictions on the domain of energy scales of QFTs and adjust the sensitivity to high energy phenomena with the tools of renormalization theory.⁶ According to the axiomatic approach, the problem arises because the mathematical structure of the QFT framework is ill-defined in the first place. The solution is to develop a rigorous mathematical formulation of QFTs with explicitly stated axioms—so that, if anything goes wrong, we can at least clearly identify the origin of the problem.⁷

⁵Note that the problem does not arise in the case of non-interacting theories since there is no non-trivial correlation between distinct states in this case (i.e., $\Gamma_n = 0$ for $n \geq 1$). Note as well that in typical interacting QFTs, some contributions to the perturbative expansion are finite (e.g., box diagram integrals).

⁶Here I will ignore the specific technicalities of Effective Field Theories (EFTs) and lattice QFTs and regroup them under the category of effective QFTs together with cut-off QFTs (see, e.g., Bain, 2013; Williams, 2015; Franklin, 2018 for recent philosophical discussions about effective theories). Note, however, that lattice QFT is often understood as a specific non-perturbative regularization framework and, in this context, the goal is usually to take the continuum limit.

⁷In this context, the problem of UV divergences is usually associated with the fact that the product of distributions at a point is ill-defined (see, e.g., Steinmann,

The crucial point is that physicists have only been able to formulate empirically successful and realistic QFTs by making extensive—if not indispensable—use of renormalization theory. It is beyond the scope of this paper to examine the axiomatic approach, but it is worth noting here that, even after seven decades, there has not yet been any finite, exact and mathematically rigorous formulation of a realistic continuum QFT in four dimensional space-time. If we want to understand the structure of our current best theories, a natural starting point is to look carefully at the details of renormalization.

Before delving into the details, it is instructive to start with the general idea of renormalization. Originally, renormalization was introduced as a set of techniques in high energy physics to isolate the divergent parts of scattering amplitudes and make them disappear from the final predictions by absorbing them into the formal expression of the couplings of the theory. In practice, the mathematical trick works because we never directly measure the value of couplings and we can already see a similar trick at work in the simpler and more vivid case of classical electromagnetism.

Consider for instance the standard example of an electrostatic field produced by an infinitely long and straight wire with uniform charge density λ (per unit length), lying along the z axis of a three dimensional Euclidean space. The measurable value of the field at some distance $r > 0$ from the wire in the xy plane orthogonal to the z direction is finite ($E \propto \lambda/r$). In contrast, the potential $V(r)$ obtained by summing up the contributions from each infinitesimal part of the wire diverges logarithmically ($V(r) \propto \lambda \int_{-\infty}^{+\infty} dz/\sqrt{z^2 + r^2}$). But since we only measure differences in the values of the potential (e.g., the field $\vec{E}(x) = -\vec{\nabla}V(x)$), it makes no physical difference to subtract or add some infinite quantity in the formal expression of the potential and work with the resulting finite “renormalized” expression. One way to make this precise and well-defined is to limit ourselves to a finite portion of the wire of arbitrary length L_0 ($V_{L_0}(r) \propto \lambda \int_{-L_0/2}^{+L_0/2} dz/\sqrt{z^2 + r^2}$). Subtracting the value of $V_{L_0}(L)$ for some fixed constant L to $V_{L_0}(r)$ leaves us with the finite function $-\lambda \ln(r/L)/2\pi$ and a finite residue depending on L_0 which vanishes if we take L_0 to infinity. The resulting renormalized expression of the potential is given by $V_R(r) = \lim_{L_0 \rightarrow \infty} V_{L_0}(r) - V_{L_0}(L)$.

More generally, the term ‘renormalization’ designates a set of techniques used to transform the kinetic and the interacting structure of theories. On the more practical side, one finds (among others) the renormalization procedure where the main goal is to generate finite and accurate predictions. On the more theoretical side, one finds the Renormalization Group (RG) theory where the main goal is to analyze the scale-dependent structure of QFTs. As we will see in section 5, it is also

2000, p. 73). The axiomatic criterion of rigor typically demands that the theory satisfies Wightman’s axioms (equivalently, Osterwalder-Schrader axioms) or Haag-Kastler axioms. For instance, the former includes assumptions of relativistic quantum theory, locally smeared fields, micro-causality, cyclicity and the existence of a unique vacuum state (see Streater and Wightman, 2000). And these axioms are usually considered to be significant to the extent that they are satisfied by toy-model theories.

useful to distinguish between perturbative and non-perturbative renormalization methods, even if renormalization theory is, in large part, a set of techniques specified and used in the context of perturbation theory.⁸ And, finally, other areas in physics have specific renormalization techniques that I will not discuss here, such as: (i) the discretized versions of the RG in condensed matter physics and (ii) the holographic RG in the context of gauge/gravity dualities.

3 Understanding the Renormalization Procedure

I argue in this section that the effective approach to renormalization offers a more physically perspicuous and conceptually coherent framework for constructing perturbative QFTs. By ‘physically perspicuous’ and ‘conceptually coherent’, I mean that the steps involved in the perturbative construction of the theory are physically justified, that the parts of the theory have a clear physical meaning and that they are coherently related to one another. I will focus on the renormalization procedure since the main differences between the two approaches are most clearly visible at this level. The upshot is, I believe, considerable: the contrast helps dissolve the much-discussed mystery of renormalization, i.e., the issue of explaining the unlikely success of the renormalization procedure (e.g., Teller, 1988; 1989; Huggett and Weingard, 1995; 1996; Fraser J., 2017). Here, again, I should emphasize that there are many different ways to implement the renormalization procedure. I will present the steps that are most conceptually consistent with the appropriate type of perturbative QFT in each case.

3.1 The Effective Approach

The effective approach to the renormalization procedure is a two-step maneuver.

(i) One first *regularizes* the divergent sub-amplitudes Γ_n by introducing a limiting high energy scale Λ_0 , called the “cut-off” or “regulator”. If we look at Eq. 3 and disregard potential trouble in the IR (i.e., $k \rightarrow 0$), $\Gamma_2(\Lambda_0) \approx \int^{\Lambda_0} d^4k/k^4$ is now a mathematically well-defined and manipulable finite quantity. But one might worry about the arbitrary choice of cut-off. A sharp cut-off separates low energy and high energy scales in a crude way, and we do not have enough information at this stage to decide whether an exponentially decreasing cut-off (e.g., $\int^{+\infty} d^4k e^{-k/\Lambda_0}/k^4$), a Gaussian cut-off (e.g., $\int^{+\infty} d^4k e^{-k^2/\Lambda_0^2}/k^4$) or what have you is the appropriate regulator.

(ii) The renormalization step compensates for this lack of constraint: one *renormalizes* the sub-amplitudes $\Gamma_n(\Lambda_0)$ by analyzing the relevant sensitivity to high energies and absorbing it into the couplings. The best way of implementing this idea is to include contributions to Γ_n

⁸For (various versions of) the Non-Perturbative Renormalization Group (NPRG) Theory, see Bagnuls and Bervillier (2001), Polonyi (2003) and Delamotte (2012) for introductory reviews. Note also the existence of axiomatic renormalization methods (e.g., Scharf, 1995).

from a specific layer of energy scales $[\Lambda, \Lambda_0]$ into a low energy theory defined only up to Λ . Call the initial regularized theory the “bare” theory $\mathcal{L}_0(\Lambda_0)$ and its parameters the “bare” parameters λ_0 and m_0 . The cut-off scale Λ_0 is the physical scale at which the theory is believed to become inapplicable and the “renormalization scale” Λ is the energy scale specifying the physics of interest, with $\Lambda \ll \Lambda_0$. In the example above, the contribution from $[\Lambda, \Lambda_0]$ is equal to

$$\begin{aligned}\lambda_0^2 \Gamma_2(\Lambda, \Lambda_0) &= \frac{3}{2} \lambda_0^2 \int_{\Lambda}^{\Lambda_0} \frac{d^4 k}{(2\pi)^4} \frac{1}{(k^2 + m_0^2)^2} \\ &\approx \frac{3}{16\pi^2} \lambda_0^2 \ln\left(\frac{\Lambda_0}{\Lambda}\right)\end{aligned}\quad (4)$$

assuming that the bare parameters are small ($\lambda_0, m_0/\Lambda \ll 1$).⁹

The essential point now is that we can simulate the effect of this high energy contribution in the expression of the bare theory $\mathcal{L}_0(\Lambda)$ restricted to the energy scale Λ (see Fig. 1).¹⁰ For that, we just need to include a new interaction term $\delta\mathcal{L}_0(\Lambda, \Lambda_0, \lambda_0) := -\frac{\lambda_{ct}}{4!} \phi^4$, called a “counter-term”, with $\lambda_{ct} = -\frac{3}{16\pi^2} \lambda_0^2 \ln(\frac{\Lambda_0}{\Lambda})$. Given Eq. 1, this amounts to shifting the value of λ_0 to $\lambda_0 + \lambda_{ct}$, i.e., to absorbing the contributions from $[\Lambda, \Lambda_0]$ into the parameter of the theory $\mathcal{L}_0(\Lambda)$. If we keep the details explicit and restrict ourselves to the second order, the new “renormalized” scattering amplitude derived from $\mathcal{L}_0(\Lambda) + \delta\mathcal{L}_0(\Lambda, \Lambda_0, \lambda_0)$ takes the form (cf. Eq. 2):

$$\begin{aligned}\Gamma_R(\Lambda) &= -(\lambda_0 + \lambda_{ct}) + \frac{3}{16\pi^2} (\lambda_0 + \lambda_{ct})^2 \left(\ln\left(\frac{\Lambda}{m_0}\right) - \frac{1}{2}\right) + \dots \\ &= -\lambda_0 + \frac{3}{16\pi^2} \lambda_0^2 \left(\ln\left(\frac{\Lambda_0}{m_0}\right) - \frac{1}{2}\right) + O(\lambda_0^3)\end{aligned}\quad (5)$$

The renormalized *effective* theory $\mathcal{L}_R(\Lambda) := \mathcal{L}_0(\Lambda) + \delta\mathcal{L}_0(\Lambda, \Lambda_0, \lambda_0)$ defined up to Λ is obtained by defining “renormalized” parameters up to the relevant order in perturbation theory:

$$\lambda_R(\Lambda) := \lambda_0 + \lambda_{ct} = \lambda_0 - \frac{3}{16\pi^2} \lambda_0^2 \ln\left(\frac{\Lambda_0}{\Lambda}\right)\quad (6)$$

This calls for two comments. First, the regularization step violates the continuum assumption only if we take the cut-off to eliminate high energy states in the state space of the original theory. Note, however, that there is a difference between restricting the possible states of a quantum field and assuming that the quantum field is a discrete physical system composed of one individual quantum system at each point of a space-time lattice. One way to see this is to look at the following toy-model. Consider the infinite set of oscillating field configurations $\phi_a(x) = \exp(iax)$ parametrized by $a > 0$ over a one dimensional con-

⁹For simplicity, I will ignore terms in $O(1/\Lambda)$ and $O(1/\Lambda_0)$ and any field renormalization of the bare field ϕ_0 (“wavefunction renormalization factors”).

¹⁰At this stage, one way of understanding the dependence of the Lagrangian functional density on the parameter Λ (or Λ_0) is to take it to refer to a restriction imposed on the Feynman rules used to compute scattering amplitudes in momentum space.

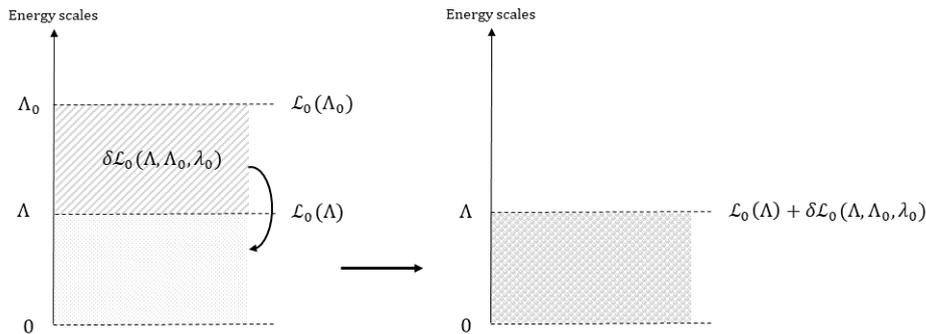


Fig. 1: Schematic representation of the effective approach to the renormalization procedure.

tinuous space and the corresponding infinite set of energy excitations $\tilde{\phi}_a(k) = \delta(k - \frac{a}{2\pi})$ obtained by Fourier transform. Suppose that the state space of the theory is reduced by multiplying the energy excitations by a step-function parametrized by a cut-off Λ_0 :

$$\tilde{\phi}_{a,\Lambda_0}(k) = \delta(k - \frac{a}{2\pi})\theta(\Lambda_0 - k) \quad (7)$$

with $\theta(\Lambda_0 - k) = 1$ if $k \leq \Lambda_0$ and 0 otherwise. For $a/2\pi \leq \Lambda_0$ (i.e., for sufficiently long wavelength oscillations), the function $\theta(\Lambda_0 - k)$ does not affect the value of $\tilde{\phi}_{a,\Lambda_0}(k)$ and we obtain the original oscillating function $\phi_{a,\Lambda_0}(x) = \exp(iax)$. Otherwise, $\phi_{a,\Lambda_0}(x)$ vanishes for $a > 2\pi\Lambda_0$. So this toy-model shows that restricting the state space of the theory by a sharp high energy cut-off implies that the possible field configurations have a minimal periodicity pattern (of wavelength $1/\Lambda_0$ here)—but it does not necessarily imply that the quantum field is discrete. To give a classical analogy, it is as if we had ignored all the possible little ripples of characteristic size smaller than $1/\Lambda_0$ in the ocean and restricted ourselves to large enough waves in order to evaluate the correlations between the oscillations of two corks floating at some macroscopic distance $1/\Lambda$ from each other.

Second, the specific counter-term $\delta\mathcal{L}_0$ leaves the theory empirically invariant, in the sense that $\mathcal{L}_0(\Lambda_0)$ and $\mathcal{L}_R(\Lambda) := \mathcal{L}_0(\Lambda) + \delta\mathcal{L}_0(\Lambda, \Lambda_0, \lambda_0)$ generate the same scattering amplitudes. The high energy contributions to $\Gamma(\Lambda_0)$ are just parceled out among the lower order terms of $\Gamma_R(\Lambda)$ (see Eq. 5). Had we chosen a different counter-term, say, $\delta\mathcal{L}_0(\Lambda, \Lambda_0, \lambda_0) + C$ with C some finite quantity, the original and modified renormalized theories would still be empirically equivalent since we only measure variations of the same renormalized scattering amplitude at different energies (e.g., $\Gamma_R(E') - \Gamma_R(E) \propto \ln(E'/E)$).¹¹ So the renormalization step is really a matter of reformulating the regularized theory $\mathcal{L}_0(\Lambda_0)$ in an epistemically safer way, i.e., around the scales $\Lambda \ll \Lambda_0$ where we can trust its physical content. Inversely, if we fix the value of the renormalized pa-

¹¹This is a particular case of “renormalization scheme dependence”. I will not discuss this issue here.

parameters at some specific scale, Eqs. 5 and 6 show that variations in the value of the cut-off Λ_0 can be absorbed by adjusting the value of the bare parameters, at least for a finite range of energy scales.

3.2 The Continuum Approach

Let us now turn to the continuum approach. It is standard in this case too to impose a regulator and split the initial regulator-dependent bare Lagrangian into a renormalized and a counter-term Lagrangian.¹² I will proceed somewhat differently by subtracting counter-terms to the physical Lagrangian. The two methods are equivalent and, most importantly, the conclusion that the continuum approach is physically ambiguous and conceptually incoherent remains the same whether we use one method or the other. The main reason for choosing the second method is that it makes the conceptual differences between the effective and the continuum approach more explicit and allows us to follow more closely the original motivation of the continuum approach.

The natural starting point, then, is to think that the original theory \mathcal{L} in Eq. 1 corresponds to the physical theory and that its parameters are fixed by experiments. Upon finding that \mathcal{L} yields divergent amplitudes, we introduce a cut-off Λ_0 (regularization) and the goal of the renormalization procedure under the continuum approach is to eliminate the problematic Λ_0 -dependent terms and take $\Lambda_0 \rightarrow \infty$ at the end. So, contrary to the effective approach, the physical theory of interest is the regularized theory $\mathcal{L}(\Lambda_0)$ with fixed physical parameters λ and m and not a low energy effective theory defined only up to Λ . Likewise, the problematic Λ_0 -dependent terms derived from $\mathcal{L}(\Lambda_0)$ are cancelled by adding counter-terms to *that* theory and not by adding them to some low energy theory $\mathcal{L}_0(\Lambda)$ as defined above. This means that the counter-terms depend on λ instead of λ_0 and that the bare theory $\mathcal{L}_0(\Lambda_0) := \mathcal{L}(\Lambda_0) - \delta\mathcal{L}(\Lambda, \Lambda_0, \lambda)$, defined up to Λ_0 as well, is an intermediary construct under the continuum approach (see Fig. 2).¹³ Finally, the parameter Λ is an arbitrary mass scale introduced to ensure that the physical expressions in the theory have a correct physical dimension, and it parametrizes the particular choice of counter-term: e.g., $\delta\mathcal{L}(\Lambda, \Lambda_0, \lambda) = \delta\mathcal{L}(\Lambda', \Lambda_0, \lambda) + C$ for C some finite quantity and Λ' a new definition of the arbitrary mass scale. I will label the renormalization scale μ instead of Λ in the continuum approach in order to keep track of the difference of interpretation.

Once all the divergent terms have been removed up to some order n in the original expression of Γ , we can stop the renormalization procedure and safely take the limit $\Lambda_0 \rightarrow \infty$ in the renormalized expression of Γ . By assumption, the value of the physical parameters λ and m is fixed

¹²The method is often called “renormalized perturbation theory” because the perturbative analysis is done in terms of the physical renormalized parameters (e.g., Peskin and Schroeder, 1995, p. 326). See, e.g., Collins (1986, sec. 2.3; 2009, sec. 2) and Schwartz (2013, part III) for different ways of implementing the renormalization procedure.

¹³In renormalized perturbation theory, the bare Lagrangian corresponds to the initial Lagrangian with the “wrong” parameters, i.e., with the parameters that we split into a finite and an infinite part in order to cancel divergences.

(e.g., to their experimental value measured at some energy scale E). So, by taking the limit $\Lambda_0 \rightarrow \infty$, we are required to take the limit of the bare parameters too. In our example, λ_0 diverges:

$$\lim_{\Lambda_0 \rightarrow \infty} \lambda_0 = \lim_{\Lambda_0 \rightarrow \infty} \left(\lambda + \frac{3}{16\pi^2} \lambda^2 \ln\left(\frac{\Lambda_0}{\mu}\right) \right) = +\infty \quad (8)$$

In principle, the original scattering amplitude Γ can be made finite at any order by repeating the procedure. And if we know the experimental values of λ and m at the scale E , we can directly compute the quantum corrections obtained at some higher energy scale E' .

Complications arise once we realize that the formal expression of the finite renormalized scattering amplitude Γ_R still depends on the arbitrary value of the mass scale μ . Since this amplitude is supposed to be a physical amplitude, we have to assume that its formal expression does not depend on some arbitrary choice of μ . This has interesting consequences.¹⁴ First, the value of the bare parameters does not depend on μ while the value of the original parameters depends on μ , as it can be easily seen from the expression of Γ_R :

$$\begin{aligned} \Gamma_R &= -\lambda + \frac{3}{16\pi^2} \lambda^2 \left(\ln\left(\frac{\mu}{m}\right) - \frac{1}{2} \right) + \dots \\ &= -\lambda_0 + \frac{3}{16\pi^2} \lambda_0^2 \left(\ln\left(\frac{\Lambda_0}{m}\right) - \frac{1}{2} \right) + O(\lambda_0^3) \end{aligned} \quad (9)$$

This means that the original theory is a particular case of a more general *renormalized* theory $\mathcal{L}_R(\mu)$, defined in terms of renormalized parameters $\lambda(\mu)$ and $m(\mu)$. Second, in the absence of experimental measurement, we can give an explicit perturbative definition of the renormalized parameters by redefining them order by order in terms of the “fixed” (i.e., μ -independent) bare parameters (i.e., $\lambda_0 = \lambda(\mu) + O(\lambda^2(\mu)) \rightarrow \lambda(\mu) = \lambda_0 - O(\lambda^2(\mu))$). As a result, the general renormalized theory is defined perturbatively by fine-tuning the expression of the Λ_0 -dependent bare theory with the help of counter-terms:

$$\mathcal{L}_R(\mu) := \lim_{\Lambda_0 \rightarrow \infty} \mathcal{L}_0(\Lambda_0) + \delta\mathcal{L}(\mu, \Lambda_0, \lambda_0) \quad (10)$$

Note that the correction $\delta\mathcal{L}$ takes the form of the original counter-term as defined in the effective approach in this simple case.

Let me make one crucial comment. The finite renormalized amplitude $\Gamma_R := \lim_{\Lambda_0 \rightarrow \infty} (\Gamma - \Gamma_{ct})$ obtained by subtracting the appropriate Λ_0 -dependent terms Γ_{ct} in the original expansion Γ is derived from the expression of the bare Lagrangian $\mathcal{L}_0(\Lambda_0)$.¹⁵ Both the original and the general renormalized theory yield divergent amplitudes $\Gamma(\Lambda_0, \lambda)$ and $\Gamma(\Lambda_0, \lambda(\mu))$ in the limit $\Lambda_0 \rightarrow \infty$ if we do not restrict the state space of the theory. Similarly, in the method where the bare Lagrangian is split

¹⁴Note that if we choose to fix the value of μ at some energy scale E , the resulting amplitude still depends on the arbitrary choice of counter-term, which we can parametrize by introducing some μ' .

¹⁵To see this, note that the expression $\Gamma - \Gamma_{ct}$ at second order in λ is obtained from Γ_0 by expressing λ_0 in terms of λ at each order.

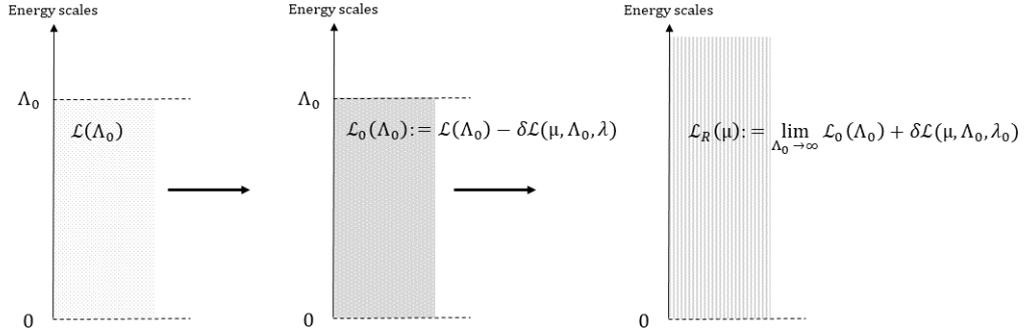


Fig. 2: Schematic representation of the continuum approach to the renormalization procedure.

into a renormalized and a counter-term Lagrangian, the finite renormalized amplitude is derived from the initial bare Lagrangian with the wrong parameters, and the physical renormalized theory yields divergent amplitudes in the limit $\Lambda_0 \rightarrow \infty$.

3.3 Comparing the Effective and the Continuum Approach

Let me now explain why the effective approach offers a more physically perspicuous and conceptually coherent formulation of renormalization. To begin with, in the somewhat naive approach taken so far, the bare theory of most QFTs makes no physical sense under the continuum approach. The reason is that most QFTs, including Quantum Chromodynamics (QCD) and Quantum Electrodynamics (QED), are plagued with UV divergences and these divergences are cancelled by choosing the bare couplings to diverge exactly in the same way. Even in QCD, the naive perturbative expression of the bare coupling parameter between quarks and gluons takes the form of a series in the physical coupling parameter with increasingly divergent Λ_0 -dependent terms at each order (see, e.g., Collins, 2011, sec. 3.3). So if we take the limit $\Lambda_0 \rightarrow \infty$ at this level, the bare coupling diverges and the resulting bare Lagrangian is ill-defined (e.g., as in our example, $\lim_{\Lambda_0 \rightarrow \infty} \lambda_0 = \pm\infty$).¹⁶ This means that the bare theory used in the renormalization procedure under the continuum approach is nothing more than a physically meaningless intermediary mathematical tool to generate finite renormalized scattering amplitudes. Therefore, we cannot explain the empirical success of the renormalized amplitudes by pointing at their successful derivation by means of some more general law and additional conditions since the bare Lagrangian, i.e., what plays the role of the general law here, has no physical meaning. Appealing to the renormalized Lagrangian does not help either since it generates divergent amplitudes.

A somewhat less naive approach is to realize that the perturbative expression of the bare parameters does not depend on the renormalization scale μ . If we take $\mu = \Lambda_0$ before taking the infinite cut-off limit,

¹⁶More generally, all the cases where the perturbative assumption $\lambda_0 \ll 1$ breaks down are pathological.

the bare parameters are equal to the renormalized parameters defined at the cut-off scale Λ_0 . As we will see in sections 5 and 6, the appropriate perturbative expression of the renormalized parameters is obtained by means of RG methods. And we will see that even from the perspective of the RG, it turns out that there are still many QFTs for which the bare parameters diverge. There are even QFTs for which it is impossible to take the infinite cut-off limit without affecting the expression of the renormalized parameters (because of the existence of a so-called “Landau pole”, see section 6). At any rate, all of these cases leave us in exactly the same situation as above. But perhaps the continuum approach offers a physically perspicuous and conceptually coherent formulation of renormalization only in well-behaved cases (i.e., $0 \leq \lim_{\Lambda_0 \rightarrow \infty} \lambda_0 \ll 1$). For instance, when we take $\mu = \Lambda_0$, the expression of the bare coupling in QCD converges to zero in the infinite cut-off limit and so the bare theory does not seem to be plagued with the same issues as the bare theory in the ϕ^4 -theory example.

Still, it turns out that the continuum approach faces important interpretative difficulties and suffers from severe conceptual ambiguities even in well-behaved cases. First, the renormalized theory yields divergent perturbative amplitudes if we do not restrict the state space of the theory. This should at the very least refrain us from taking this theory at face value too quickly (see section 7). Second, the conceptual status of counter-terms is ambiguous under the continuum approach, and this is independent of the value of the bare parameters. Recall that, whether we add the counter-terms to the original theory or obtain them by splitting the initial bare parameters into two pieces, the main role of the counter-terms is to make the original amplitude finite. We might attempt to clarify their conceptual status in two different ways.

- (i) Counter-terms correspond to surplus component parts of the bare theory which cancel out with other divergent parts of the bare theory when we calculate amplitudes. That is, by adding counter-terms to the physical theory, we simply re-arrange the structure of the bare theory in such a way that its superfluous divergent parts cancel each other.
- (ii) Counter-terms correspond to scaling factors relating the parameters of the bare and physical theories. That is, by adding counter-terms to the physical theory, we simply reparametrize the original parameters in such a way that the resulting theory, i.e., the bare theory, yields finite predictions. In both cases, however, the counter-terms cancel out precisely because we choose the component parts of the bare theory or the scaling factors of the physical theory in such a way that they cancel the divergent parts of the original amplitudes. That is, in both cases, it seems difficult to escape the conclusion that counter-terms are introduced just for the purpose of canceling divergences, which makes the whole renormalization procedure look *ad-hoc*. Moreover, it is hard to see how one could possibly interpret the counter-terms in physical terms, including those relating the original field variable and the bare field variable (which I ignored for simplicity), and therefore to make sense of the relationship between the bare and the renormalized theory. The counter-terms are, as it were, intrinsically meaningless formal tools

to derive finite predictions.

The contrast with the effective approach is striking. First, recall that on this approach, we start with the assumption that the bare theory breaks down at some physically meaningful scale Λ_0 . The structure of the theory may give us very good internal reasons to believe that it becomes inconsistent at some point beyond this scale, or we may have very good external reasons to believe that the theory starts to make empirically inaccurate predictions around this scale. Either way, we take the domain of applicability of the theory to be restricted by some limited range of energy scales. On this assumption, $\mathcal{L}_0(\Lambda_0)$ is naturally interpreted as the most fundamental formulation of the theory, i.e., the theory defined up to the scale Λ_0 where it is supposed to break down. The renormalized theory \mathcal{L}_R is naturally interpreted as a more physically reliable low energy effective version of the bare theory. If we take $\Lambda = \Lambda_0$, the physical renormalized theory $\mathcal{L}_R(\Lambda)$ simply corresponds to the bare theory $\mathcal{L}_0(\Lambda_0)$ for some appropriate choice of counter-terms. And insofar as Λ_0 is kept fixed, both theories yield finite predictions, are mathematically well-defined (at least according to physicists' standards), and even yield exactly the same scattering amplitudes if we choose the counter-terms appropriately. Second, the effective approach offers a physically salient interpretation of counter-terms: whether we fix the parameters of the bare theory or those of the renormalized theory, the counter-terms are naturally interpreted as standing for high energy effects described by the bare theory. Moreover, the introduction of counter-terms is physically justified on the grounds that the low energy scales are not completely insensitive to the high energy ones.

All of this should help to clarify the mystery surrounding the renormalization procedure discussed in the literature (e.g., Teller, 1988; 1989; Huggett and Weingard, 1995; 1996; Fraser J., 2017). The mystery, if anything, is a mystery about the continuum approach: it arises because the meaning and the status of the bare theory, the renormalized theory, and the counter-terms are ambiguous, and because the method used for deriving the renormalized theory and the finite renormalized scattering amplitudes is physically unjustified. By contrast, the effective approach relies on well-specified physical concepts and offers a clear physical picture of inter-scale sensitivity. The effective approach also offers a better rationale for each step of the renormalization procedure: while there are good reasons to expect a physical theory to break down at short distances (regularization step), it does not mean that it automatically fails to provide physically relevant and empirically accurate descriptions at larger distances if the relevant sensitivity to short distances is taken into account (renormalization step).

Now, the mystery surrounding the continuum approach is not as mysterious as it might seem, at least in this simple case.¹⁷ It is a standard principle in physics that physical expressions must have the same physical dimension upon mathematical transformation for them to remain physically meaningful. This principle requires us to introduce the

¹⁷More generally, the following explanation works for QFTs displaying logarithmic divergences (e.g., QED, QCD).

new arbitrary parameter μ with the introduction of the regulator Λ_0 (e.g., to use $\lambda^2 \ln(\Lambda_0/\mu)$ instead of $\lambda^2 \ln(\Lambda_0)$ as a counter-term). In this specific example, this principle also ensures that the arbitrary parameter μ captures exactly the sensitivity to high energies as parametrized by the regulator Λ_0 , as it can be seen from the expression of the counter-term. The continuum approach therefore successfully offers a measure of the sensitivity of the low energy physics to the high energy physics, and this is in fact all that is needed to explain the empirical success of the theory. Had we chosen a different counter-term, say, $\lambda^2 \ln(\Lambda_0/\mu) + C$, with C some finite quantity, the same sensitivity would be captured by some appropriate redefinition of μ . Hence, even if the continuum approach offers a highly formalistic and instrumental framework, it remains at least possible to identify the reasons for its empirical success. Needless to say, the effective approach offers a more physically perspicuous and conceptually clear explanation.

Let me conclude this section by responding to two potential concerns. First, taking the limit $\Lambda_0 \rightarrow \infty$ under the effective approach does not turn the situation around. Agreed, there is nothing problematic if the goal is to probe the mathematical structure of the theory, or if we add by hand a high energy cut-off afterwards. But, strictly speaking, taking the limit $\Lambda_0 \rightarrow \infty$ is conceptually incoherent since the introduction of the cut-off Λ_0 is justified on the grounds that it marks the physical scale at which the theory is supposed to break down. Another option is that, by taking $\Lambda_0 \rightarrow \infty$, we are actually making the approximation that the low energy physics is largely insensitive to the high energy physics beyond Λ_0 . But in this case, it is implicitly assumed that the theory is restricted to low energies and that it should not be used to make predictions at arbitrarily high energies.

Second, the distinction between the effective and the continuum approach does not crucially depend on the specific regularization method we choose and on the specific way we subtract divergences or absorb the appropriate sensitivity to high energies (although I will emphasize in section 5 that each approach is most conceptually consistent with its own specific type of regularization and renormalization method). In particular, the distinction between the effective and the continuum approach does not reduce to the distinction that Georgi (1992, 1993) and Bain (2013) draw between Wilsonian and continuum EFTs. This distinction is mainly based on whether the split between the low energy and high energy physics depends on the mass parameter of the theory (see Georgi, 1993, sec. 1.2; Bain, 2013, sec. 4). And continuum EFTs are called “continuum” because the most famous mass-independent regularization method, namely, dimensional regularization, does not eliminate high energy states in the state space of the theory. This, however, does not mean that continuum EFTs are meant to be used to make predictions across all energy scales. In particular, they are restricted by the energy scale characterizing the matching conditions.

4 (Perturbative) Renormalizability Yes... But Which One?

We have seen that the continuum approach to the renormalization procedure offers a highly formalistic and instrumental perturbative framework to derive consistent and empirically relevant predictions. It turns out that the situation is even worse for the continuum approach since the procedure only works at every order in perturbation theory for the restricted class of “perturbatively renormalizable” QFTs. After distinguishing between two distinct notions, one for the continuum approach and the other for the effective approach, I argue in this section that the continuum approach is all the less attractive as it fails to apply to a large number of successful and physically significant theories. We will see in sections 5 and 6 that the RG does not substantially affect this claim.

First consider the continuum approach. Here the notion of perturbative renormalizability is best introduced by noting that the ϕ^4 -theory example used so far is extremely fortunate. All the divergent terms depending on positive powers of Λ_0 or $\log(\Lambda_0)$ that appear in the perturbative expansion of $\Gamma(p_1, \dots, p_4)$ can be absorbed by introducing counterterms that depend only on the coupling λ . All the finite terms depending on positive powers of $1/\Lambda_0$ vanish as we take $\Lambda_0 \rightarrow \infty$. More generally, all the divergences that appear in any sub-amplitude can be cancelled by using only λ and m in the ϕ^4 -theory. There are many theories, however, for which infinitely many new couplings need to be introduced—the 4-Fermi theory is one such example (see, e.g., Schwartz, 2013, chap. 22)—and the difference between this example and the ϕ^4 -theory can be captured as follows:

A theory is perturbatively renormalizable iff we only need to introduce a finite number of independent couplings in order to eliminate divergences and define $\mathcal{L}_R(\mu)$ at any order in perturbation theory in the limit $\Lambda_0 \rightarrow \infty$.

A theory is perturbatively non-renormalizable iff we need to introduce an infinite number of independent couplings.

This characterization is of course somewhat superficial. According to Dyson’s criterion, what makes a theory perturbatively non-renormalizable is that it contains *at least* one “non-renormalizable” individual interaction term, i.e., an interaction term parametrized by a coupling g_i with strictly negative mass dimension Δ_i .¹⁸ These types of interactions generate an infinite number of sub-amplitudes with an increasing degree of divergence, and each of the resulting types of divergent quantities usually requires the introduction of a new counter-term. In contrast, the so-called “renormalizable” ($\Delta_i = 0$) and “super-renormalizable” ($\Delta_i > 0$) interaction terms generate only a finite number of different types of divergences.¹⁹ Having said that, perturbative non-renormalizability is

¹⁸The mass dimension Δ of a physical quantity is the power of that quantity expressed in terms of some energy variable (i.e., energy ^{Δ}) with natural units $c = \hbar = 1$.

¹⁹The longer explanation is based on the so-called “power-counting” argument (e.g.,

not a dead end. In general, both perturbatively renormalizable and non-renormalizable theories are “renormalizable” in the sense that the structure of the theory is such that it is possible to construct a counter-term to cancel any type of divergence at any order in perturbation theory. I will call this notion “renormalizability_{RT}” to avoid confusion as it is sometimes referred to as the “Renormalization Theorem” (e.g., Osborn, 2016, sec. 4.3).²⁰ We can even take the limit $\Lambda_0 \rightarrow \infty$ in a number of expressions obtained from perturbatively non-renormalizable theories at each finite order in perturbation theory if the theory is not too exotic (i.e., if $\lim_{\Lambda_0 \rightarrow \infty} g_0$ formally exists for each bare coupling g_0 given some fixed finite order in perturbation theory).

At first sight, it seems that the distinction between perturbatively renormalizable and non-renormalizable theories captures the amount of work needed in order to renormalize a theory—and the amount is of course infinite if we want to define the parameters of a non-renormalizable theory at every order in perturbation theory. In fact, the notion of perturbative renormalizability provides a much stronger criterion of theory-selection under the continuum approach. If the perturbative expression of a non-renormalizable theory is defined by introducing an infinite number of new parameters, it means that quantum corrections to scattering amplitudes depend on the specification of an infinite number of constants and that we therefore need an infinite number of experiments in order to fix their value. Since this is impossible in practice, the perturbative formulation of non-renormalizable theories obtained by applying the renormalization procedure at every order in perturbation theory turns out to be empirically useless. We should therefore restrict the class of empirically relevant theories to perturbatively renormalizable theories under the continuum approach.

So far, the analysis only applies to the continuum approach and one might wonder whether there is any equivalent notion of perturbative

Weinberg, 1995, sec. 12.1). A divergent integral $I = \int^\infty dk k^{D-1}$ is characterized by the value of its superficial degree of divergence D (the integral diverges in the UV if $D \geq 0$) and D can be expressed in terms of the mass dimensions Δ_i of the interactions involved in the scattering process described by I : schematically, $D =$ positive number $-\sum_i n_i \Delta_i$, with n_i the number of times we need to use the interaction i to define the integral. Then, if there is at least one non-renormalizable interaction in the theory ($\Delta_i < 0$), it is possible to find infinitely many different types of divergent integrals ($D \geq 0$) by considering more and more complex sub-amplitudes at higher orders in the perturbative expansion. By contrast, D has a positive upper bound for (super-) renormalizable theories, i.e., there is only a finite number of different types of divergent integrals. Note, however, that the superficial degree of divergence is not always reliable: there are cases where $D < 0$ and the integral diverges (notably because of the so-called “sub-divergence” problem), and cases where $D \geq 0$ and the integral is finite (usually the divergence cancels because of symmetry constraints). Perturbatively renormalizable theories are sometimes called “renormalizable in the power-counting sense” or “renormalizable in Dyson’s sense”.

²⁰In contrast, the term ‘non-renormalization theorem’ usually refers to a specific result to the effect that a parameter or an interaction term does not need to be renormalized at all at any order in perturbation theory, as it is common in supersymmetric QFTs (see, e.g., Weinberg, 2000, sec. 27.6). Of course, in practice, the interesting question is whether a theory is renormalizable_{RT} given a set of constraints imposed on the construction of counter-terms (e.g., that they leave the resulting Lagrangian invariant under the action of a given symmetry group).

renormalizability under the effective approach and, if so, whether it plays the same role. Let me suggest the following notion of “perturbative renormalizability_E”, to be distinguished from the traditional notion and the notion of renormalizability_{RT}:

A theory is perturbatively renormalizable_E iff for any $p \in \mathbb{Z}$, all the possible contributions to predictions up to order $O((\Lambda/\Lambda_0)^p)$ obtained from $\mathcal{L}_0(\Lambda_0)$ can be absorbed in $\mathcal{L}_R(\Lambda)$ by introducing only a finite number of new parameters. (*mutatis mutandis* for perturbatively non-renormalizable_E.)

The basic idea is the following: a theory is perturbatively renormalizable_E if we can always simulate high energy effects up to a specific accuracy ϵ with only a finite number of couplings, and perturbatively non-renormalizable_E if we cannot. It is not entirely clear what perturbatively non-renormalizable_E theories would look like. Presumably, these types of theories would have to include exotic interaction terms such that the contributions of these terms vary too rapidly between Λ_0 and Λ to be approximated by the contributions of a finite number of independent polynomial interaction terms in the field variables and their derivatives given some accuracy ϵ . For instance, we can imagine a theory with an exotic non-local interaction term including some non-analytic function $F(\phi(x), \phi(y))$ of field variables specified at distinct space-time points x and y such that the contributions of F vary too abruptly between Λ_0 and Λ to be approximated by the contributions of a finite number of independent polynomial interaction terms.

Be that as it may, the notion of perturbative renormalizability_E is much less constraining than the traditional notion of perturbative renormalizability. Perturbative renormalizability_E is satisfied if the interaction terms of the theory are local polynomials in the field variables and their derivatives and if the theory has a finite number of independent interaction terms with the same dimension Δ_i . Most crucially, the notion of perturbative renormalizability_E does not prevent the theory from including non-renormalizable interaction terms. Quite the contrary: under the effective approach, we often need to introduce non-renormalizable terms into the effective theory if we want to absorb contributions in $O((\Lambda/\Lambda_0)^p)$ ($p > 0$) obtained—say—from the renormalizable interaction terms of the bare theory.²¹ There is no specific reason to worry about these contributions in the continuum approach since they cancel out in the limit $\Lambda_0 \rightarrow \infty$. But to the extent that we keep the cut-off fixed, we usually need to include non-renormalizable terms in the effective theory if we want to maximize the match between the effective and the bare theory.

As a corollary, if we keep the cut-off fixed, perturbatively non-renormalizable theories remain perfectly predictive and empirically relevant. Typically, it is sufficient to consider interaction terms with dimension larger or

²¹Lepage (1989, sec. 2.3) provides a concise explanation of this pattern. If we consider again the superficial degree of divergence of integrals (see footnote 19), it is possible for any renormalizable interaction to generate infinitely different types of finite integrals with negative superficial degrees of divergences, i.e., with contributions in $O((\Lambda/\Lambda_0)^p)$ ($p > 0$). For some examples, see Schwartz (2013, chap. 21).

equal to $\Delta_\epsilon = -\ln(\epsilon)/\ln(\Lambda/\Lambda_0)$ in order to make predictions at the energy scale Λ with accuracy ϵ (e.g., Georgi, 1993, p. 214). The total number of interaction terms with $\Delta \geq \Delta_\epsilon$ is finite in standard QFTs and we can increase the empirical accuracy of the theory by adding non-renormalizable interaction terms with $\Delta < \Delta_\epsilon$ (keeping in mind that the mass dimension of non-renormalizable interaction terms is negative). In general, the most empirically successful and physically informative version of an effective theory (the so-called “Wilsonian” effective Lagrangian) includes all the possible interaction terms compatible with the assumption of the theory—in particular, its symmetries.²² To give an example, the effective Lagrangian \mathcal{L}_W generalizing the ϕ^4 -theory takes the following form:

$$\begin{aligned} \mathcal{L}_W = & -\frac{1}{2}(\partial\phi)^2 - \frac{m^2}{2}\phi^2(x) - \frac{\lambda}{4!}\phi^4(x) - \sum_{n \geq 3} g_{2n}\phi^{2n} \\ & - \sum_{n \geq 2} g'_{2n}(\partial\phi)^{2n} - \sum_{n,m \geq 1} g''_{2n,2m}(\partial\phi)^{2n}\phi^{2m} \end{aligned} \quad (11)$$

Non-renormalizable interaction terms are those associated with the couplings g_{2n} with $n \geq 3$, g'_{2n} with $n \geq 2$ and $g''_{2n,2m}$ with $n, m \geq 1$ in this example.

Perturbatively non-renormalizable theories have been much appraised in the recent physical and philosophical literature (e.g., Lepage, 1989; Cao and Schweber, 1993; Butterfield and Bouatta, 2015; Williams, 2018). I do not have much to add here, except the following four important points. First, for any approach, the restriction to a finite number of independent couplings is necessary in practice if we want to make empirical predictions. Second, the effective approach provides a clear physical justification for the introduction of an infinite number of additional non-renormalizable interaction terms: they capture the full sensitivity of the low energy physics to high energies, even the most insignificant parts of it. Third, perturbative renormalizability *remains* a decisive criterion of theory-selection for the perturbative formulation of *continuum* theories insofar as it is possible to define (at least formally) perturbatively non-renormalizable theories at every order in perturbation theory in the $\Lambda_0 \rightarrow \infty$ limit. Fourth, the notion of perturbative renormalizability_E under the effective approach offers a highly inclusive criterion of theory-selection and, as far as I can tell, all the traditional perturbatively renormalizable and non-renormalizable QFTs are perturbatively renormalizable_E. In a way, perturbative renormalizability_E is as non-constraining as the notion of renormalizability_{RT} discussed above (but less general, though).

Now, it is a matter of fact that perturbatively non-renormalizable theories have proven to be extremely useful in deriving successful empirical predictions and describing physically relevant patterns at different energy scales, from low energy effective phenomenological models to

²²A complication comes from anomalies: i.e., the renormalization procedure might require new terms which explicitly break the symmetries of the theory. This is called “anomalous” or “quantum” symmetry breaking, but I will ignore this problem here.

extensions of QFTs beyond the Standard Model. This success, however, requires us to explicitly restrict the domain of applicability of these theories by means of some finite cut-off. For if we attempt to define the perturbative formulation of these theories across all scales and derive exact predictions without making any approximation, we will find that these theories lose their predictive power and empirical relevance. Of course, if we have empirical inputs and restrict ourselves to some finite order in perturbation theory, we may take the limit $\Lambda_0 \rightarrow \infty$ at this order and use the perturbatively non-renormalizable theory to make predictions. For instance, if we know the value of the Fermi constant $G_F \sim 10^{-5} \text{ GeV}^{-2}$, we can use the 4-Fermi theory to make tree-level predictions. However, if we endorse the continuum approach and intend to renormalize theories at every order in perturbation theory, we will be forced to rule out a large class of empirically and physically relevant theories. And so insofar as we want to praise a framework for constructing perturbative QFTs that proves to be (sufficiently) universal, the effective approach looks more attractive than the continuum approach.

5 The Renormalization Group Theory

What has been at the center stage of the renormalization procedure so far is the attempt to address the problem of UV divergences:

- (1) How can finite and accurate predictions be obtained if the original theory is inconsistent?

We have seen that in both the effective and the continuum approach, the introduction of an arbitrary mass scale Λ (or μ) is forced upon us if we want to derive the expression of renormalized quantities. The genius of the physicists who developed the Renormalization Group (RG) theory was to use this seemingly idle and arbitrary parameter as a lever to address the (new) questions:

- (2) What is the scaling behavior of the theory?
- (3) Does the theory make consistent predictions in the continuum limit?

The goal of this section is to show how the RG theory clarifies the notion of renormalizability and therefore complicates the argument of section 4. Of crucial importance is the possibility that a theory both includes non-renormalizable interaction terms and makes consistent predictions in the continuum limit. At the same time, some perturbatively renormalizable theories such as the Standard Model of particle physics are likely to make inconsistent predictions at very high energies. This suggests that the scope of the continuum approach might not be as restricted as initially thought—and yet still restricted in important ways.

5.1 The Effective and the Continuum RG

What, exactly, is the RG? Strictly speaking, the RG refers to the structure of invariance of theories under rescaling by the renormalization

scale Λ (or μ). It is helpful, though, to distinguish between three types of RG equations. First, at the level of theories, the RG describes how the path integral, the action and the Lagrangian transform under rescaling. In a way, the renormalization procedure already gives us a rudimentary RG transformation: e.g., in the effective approach, $\mathcal{L}_0(\Lambda_0) \rightarrow \mathcal{L}_R(\Lambda) = \mathcal{L}_0(\Lambda) + \delta\mathcal{L}_0(\Lambda, \Lambda_0)$ for $\Lambda_0 \rightarrow \Lambda$. Second, at the level of scattering amplitudes and correlation functions, the RG describes the specific trade-off between the kinetic and interacting parts of the theory required for the scattering amplitudes to remain invariant under rescaling. The so-called ‘‘Callan-Symanzik’’ equation for a N -particle amplitude with one renormalized coupling g is given by:

$$\left(\Lambda \frac{\partial}{\partial \Lambda} + \beta \frac{\partial}{\partial g} + N\gamma_\phi\right) \Gamma_R(p_1, \dots, p_N; g(\Lambda)) = 0 \quad (12)$$

where $\beta(g) = \Lambda \frac{\partial g}{\partial \Lambda}$ is the ‘‘beta-function’’ of the coupling g and γ_ϕ is the ‘‘anomalous dimension’’ of the field. Eq. 12 describes how much we need to shift the value of the coupling ($\beta \frac{\partial}{\partial g}$) and the size of the field configurations ($N\gamma_\phi$) in order to absorb an infinitesimal rescaling ($\Lambda \frac{\partial}{\partial \Lambda}$) and leave the amplitude Γ_R invariant. Third, at the level of couplings and local operators, the RG describes how the strength of an interaction varies across scales in accordance with the sign of its beta-function. For instance, the quartic interaction in the ϕ^4 -theory becomes increasingly strong at high energies:

$$\Lambda \frac{\partial \lambda_R}{\partial \Lambda} = \beta(\lambda_R) = \frac{3}{16\pi^2} \lambda_R^2 + \mathcal{O}(\lambda_R^3) \quad (13)$$

Note, however, that this perturbative RG equation remains only valid for $\lambda_R \ll 1$.

The effective (or Wilsonian) RG and the continuum (or Gell-Mann & Low) RG have a relatively similar formal structure overall. But again, there are significant conceptual differences between the two.²³ Most crucially, the effective renormalized theory is obtained by integrating out high energy field configurations in the path integral, while the continuum renormalized theory is obtained by fine-tuning the expression of the bare theory. Schematically,

$$\begin{aligned} \text{Effective theory:} & \quad \int d[\phi_{<\Lambda}] e^{S_{\text{eff}}(\Lambda, \Lambda_0)} = \int d[\phi_{<\Lambda_0}] e^{S_0(\Lambda_0)} \\ \text{Continuum theory:} & \quad \int d[\phi] e^{S(\mu)} = \lim_{\Lambda_0 \rightarrow \infty} \int d[\phi] e^{S_0(\Lambda_0) + \delta S(\mu, \Lambda_0)} \end{aligned} \quad (14)$$

with the same conventions as before ($\phi_{<\Lambda}$ refers to field configurations with energy lower than Λ). The effective RG transformation obtained by decreasing Λ is irreversible since it eliminates high energy degrees of freedom, while the continuum RG transformation obtained by varying

²³For more technical details, see, e.g., Weinberg (1995, sec. 12.4; 1996, chap. 18); Zinn-Justin (2007); Schwartz (2013, chap. 23). Here I rely on the standard understanding of the Wilsonian RG. For a formal interpretation, see Rosaler and Harlander (2019).

μ is reversible since it merely amounts to subtracting or adding some finite quantity in the action (i.e., to imposing a different renormalization condition).²⁴

Next, the most conceptually consistent interpretation of the cut-off and of the renormalization scale is not the same in the two cases. In the effective approach, the idea of integrating out *all* the high energy degrees of freedom makes sense only if we use a sharp cut-off (e.g., a lattice or a momentum cut-off). If we use a smooth cut-off, the path integral measure is still defined by summing over arbitrarily high energy states. Similarly, if we have good reasons to think that the high energy states close to the sharp cut-off Λ_0 misrepresent, in some way or another, the correct state of matter, we should make sure that we exclude them. One conceptually simple and clear way of ignoring these high energy states is to integrate out *all* the high energy degrees of freedom between Λ and Λ_0 . In contrast, the continuum approach is based on the idea that all the degrees of freedom in the original theory \mathcal{L} are relevant in some respect. One way of making sure that the continuum assumption holds is to use a regularization method that gives weight to the physical states of interest without eliminating the others. In the method of dimensional regularization, for instance, the divergences are analyzed by shifting the dimension of space-time by $\pm\epsilon$, and the state space of the theory is smoothly distorted in the UV in a way that keeps all the possible energy states but significantly lowers the weight of the states above the scale μ .²⁵

Agreed, one important lesson of the modern understanding of renormalization is that the specific value of the cut-off and the specific details of the regularization method do not really matter. They can always be absorbed in the formal expression of the renormalized parameters and, overall, the predictions obtained with different regularization methods are empirically equivalent.²⁶ Yet, this does not mean that all regularization methods are on the same footing. If the goal is to define a theory across all energy scales, for instance, it appears somewhat conceptually inconsistent to construct the theory by first eliminating all the high energy degrees of freedom beyond some fixed scale. Similarly, if the goal is to offer a restricted description of low energy degrees of freedom, it appears somewhat conceptually inconsistent to include the contributions from arbitrarily high energy degrees of freedom when calculating low energy predictions. At the very least, some regularization methods make

²⁴Note the difference between active and passive transformation in both cases. An active RG transformation corresponds to a genuine change of scales (and hence to integrating out degrees of freedom in the effective case). A passive RG transformation corresponds to a conventional redefinition of energy units (in which case no degree of freedom is integrated out in the effective case).

²⁵In more detail, if we take $d = 4 \pm \epsilon$, we have to rescale each coupling by some power of the renormalization scale μ for dimensional consistency (e.g., replace λ by $\lambda\mu^\epsilon$) and modify the dimension of the divergent integrals. Integrating out the extra ϵ dimension in those integrals leaves an additional damping factor in the integrand that depends on both ϵ and μ . If we ignore potential trouble in the IR, this damping factor smoothly vanishes for momenta much larger than μ with ϵ small and μ fixed (Georgi, 1992, p. 4; 1993, sec. 2.4, provides a very helpful and concise explanation).

²⁶I would like to thank a referee for pressing me on this point.

these specific goals more explicit and provide a conceptually simpler and clearer way of achieving those goals. In the case of the effective approach, for instance, a deformation that eliminates all the high energy degrees of freedom appears to be more natural than a deformation that merely lowers the weight of UV contributions. For if we believe that the theory literally breaks down at high energies, we should rather avoid using those high energy degrees of freedom altogether instead misrepresenting their properties and using them to compute low energy predictions. Likewise, a sharp cut-off introduces a conceptually simple and clear classification of low energy and high energy field configurations, while a smooth cut-off makes the boundary between them somewhat fuzzier. Agreed again, both a sharp and a smooth cut-off offer a highly idealized representation of the boundary between the low energy and high energy regimes of the theory. But we do not need to take the exact form of the cut-off to be physically meaningful in order to grant that those differences between a sharp and a smooth cut-off regularization method are significant. And of course, if our primary goal is simply to compute low energy predictions, we should probably select the regularization method which allows us to achieve this goal in the simplest, most efficient and appropriate way.

5.2 RG and Renormalizability

Now, let us look at the implications of the RG for the notion of renormalizability and for the scope of the continuum approach. Before we do that, it is necessary to spend some time clarifying: (i) the notion of RG space or theory space; (ii) the notion of fixed point; and (iii) the different types of local behaviors in the vicinity of fixed points.²⁷

(i) Consider first the infinite set of renormalized couplings $g_n(\Lambda)$, including both couplings from renormalizable and non-renormalizable interactions, which can be used to define any sort of renormalized (local) QFT in four dimensions for a specific set of fields and symmetries. The infinite set of RG equations span an infinite dimensional space, the so-called “RG space”, where each coupling stands for an independent coordinate and where each point in the space represents a theory defined at some energy scale Λ (see Fig. 3). The RG transformations of the couplings induce trajectories in this space, the so-called “RG flows”, either towards the IR or the UV as we respectively decrease or increase the value of Λ .²⁸

(ii) Fixed points g^* are defined by the points in theory space where the RG flow terminates. The fixed point is either IR or UV depending on whether the RG flow converges to the fixed point in the low energy or high energy limit. In each case, the β -function $\beta(g_n^*)$ of each coupling vanishes at the fixed point, which means that each coupling $g_n(\Lambda) = g_n^*$ remains constant whether the value of Λ is increased or decreased and that the theory specified by the couplings g_n^* is scale-invariant. It turns

²⁷The analysis does not depend on the type of RG used since the effective and the continuum RG are formally equivalent at the level of couplings.

²⁸Typically, an effective theory is defined by a finite segment of the RG flow with an upper bound while a (well-defined) continuum theory is defined by a complete segment specified by the values of $g_n(\mu)$ for any μ .

out that we can distinguish between two kinds of fixed points. A Gaussian fixed point $g^* = 0$ defines a non-interacting theory, and theories converging towards a Gaussian fixed point are called “asymptotically free”. A Wilson-Fisher fixed point $g^* \neq 0$ defines a non-trivial scale-invariant dynamics, and theories converging towards a Wilson-Fisher fixed point are called “asymptotically safe”. As we can already anticipate, the existence of a UV fixed point indicates that the corresponding continuum theory behaves well at high energies, i.e., that the value of its couplings remains finite at high energies.²⁹

(iii) The infinite set of RG equations determine local topological properties of the RG flow on theory space. To see that, we need to examine first the behavior of couplings in the vicinity of a fixed point. In the simple case of a Gaussian fixed point, the perturbative RG equation for a coupling g looks like:

$$\Lambda \frac{\partial g}{\partial \Lambda} = \beta(g) \approx (-\Delta + \gamma)g + bg^2 + cg^3 + O(g^4) \quad (15)$$

where γ , b and c are constants. Assuming that γ is negligible, the solution at lowest order is given by:

$$g(\Lambda) = \left(\frac{\Lambda}{\Lambda_0}\right)^{-\Delta} g(\Lambda_0) \quad (16)$$

Three types of behaviors can be distinguished from this elementary perturbative equation, and each of them clarifies the scale-dependence of the non-renormalizable, renormalizable and super-renormalizable individual interaction terms defined in section 4 (here with the flow directed towards the IR):

- (a) Super-renormalizable interaction: If $\Delta > 0$, the coupling $g(\Lambda)$ becomes large at small scales $\Lambda \ll \Lambda_0$ and negligible near the cut-off $\Lambda \lesssim \Lambda_0$. The coupling and the corresponding operator are said to be “relevant” at low energies.
- (b) Non-renormalizable interaction: If $\Delta < 0$, the coupling $g(\Lambda)$ becomes negligible at small scales $\Lambda \ll \Lambda_0$ and large near the cut-off $\Lambda \lesssim \Lambda_0$. The coupling and the corresponding operator are said to be “irrelevant” at low energies.
- (c) Renormalizable interaction: If $\Delta = 0$, dimensional analysis is ineffective. The sign of the next dominant term in Eq. 15 determines whether the coupling is “marginally” relevant or irrelevant. For instance, the (dimensionless) λ_R coupling in ϕ^4 -theory is marginally irrelevant (see Eq. 13).

With these three properties in hand, we can specify the distinct local topological features of the RG flow in theory space (see Fig. 3).

²⁹More precisely, there are three conditions for asymptotic safety/freedom in both the IR and the UV case: (a) the vanishing of the β -function; (b) the existence of a finite-dimensional surface in the vicinity of the fixed point if we want the theory to be predictive; and (c) the existence of a well-behaved RG flow on the way to and at the fixed point. In the IR case, for instance, condition (c) obtains if $\sup_{\Lambda \leq \Lambda'} g(\Lambda) < \infty$ for some Λ' and $\lim_{\Lambda \rightarrow 0} g(\Lambda) = g^* < \infty$.

The “critical surface” is defined by the set of couplings whose trajectory ends on the fixed point and the “unstable manifold” is defined by the set of couplings whose trajectory departs from the fixed point. In general, trajectories can be extremely varied: the flow might seemingly converge toward a fixed point and quickly diverge away as it comes too close to it, or the flow might seemingly diverge away from a fixed point and suddenly converge extremely fast towards it. Some RG flows even display periodic behaviors (see Wilson, 1971, and Bulycheva and Gorsky, 2014, for a discussion and examples of cyclic flows). In typical cases, the critical surface corresponds to the subspace spanned by the set of irrelevant couplings while the unstable surface corresponds to the subspace spanned by the set of relevant couplings, and most trajectories converge towards the fixed point and suddenly diverge away as the relevant couplings become too important (as depicted in Fig. 3). The analysis applies both to IR and UV fixed points, and we may speak similarly of IR/UV relevant, irrelevant and marginal operators.

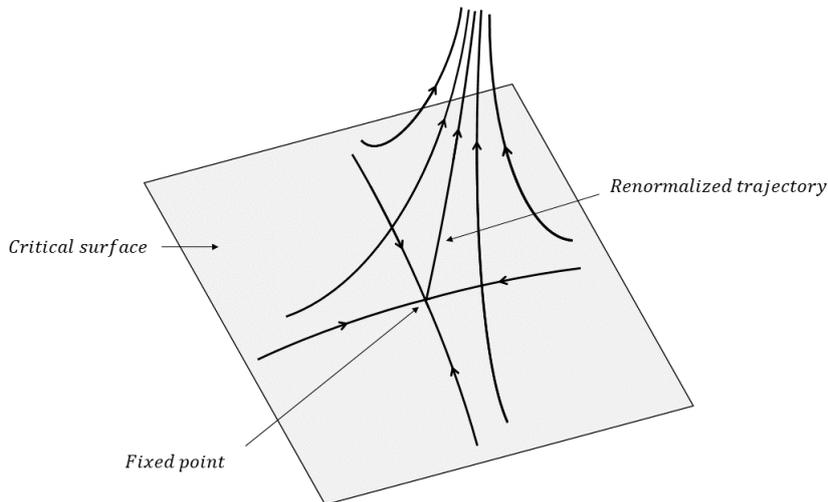


Fig. 3: Theory space in three dimensions, with a two dimensional critical surface and a one dimensional unstable manifold. The possible trajectories towards the IR are denoted by the lines with arrows.

This analysis has three important implications.³⁰ First, it shows that the pathological high energy behavior of non-renormalizable interactions (i.e., IR-irrelevant/UV-relevant interactions) is closely linked to the fact that they generate increasingly divergent integrals in perturbation theory. Consider for instance a scalar theory in four dimensions with one non-renormalizable interaction term $g_6\phi^6$. The 6-particle physical amplitude $\Gamma(p_1, \dots, p_6)$ is just equal to g_6 at first order. On dimensional grounds, the total amplitude Γ and any of the higher order sub-amplitudes $g_6^n\Gamma_n$ ($n > 1$) have mass dimension -2 . As I briefly

³⁰Note that this analysis also explains why typical realistic QFTs are likely to be only approximately perturbatively renormalizable since they might contain IR irrelevant interactions that we have not detected yet (e.g., Butterfield and Bouatta, 2015; Williams, 2018).

indicated above, the amplitudes at some order n have in general the schematic form $g_6^n \int dk k^{D_n-1}$. On dimensional grounds, we can therefore infer from the mass dimension of g_6^n (namely, $-2n$) that the number D_n increases with n ($D_n = 2n - 2$), which shows that the sensitivity of non-renormalizable interactions to high energies is closely linked to the pathological perturbative behavior of the theory.

Second, the RG theory suggests a general non-perturbative characterization of renormalizability. In the continuum approach, the notion can be defined as follows (I drop the label “non-perturbative” for simplicity):

A theory is renormalizable_{RG} iff there is some μ' such that the RG flow remains on the same finite dimensional surface as the theory is rescaled toward the UV (i.e., for any $\mu > \mu'$). (*mutatis mutandis* for non-renormalizable_{RG}.)

In other words, the theory is renormalizable_{RG} if it can be expressed in terms of a finite number of independent couplings as the theory flows towards high energies and non-renormalizable_{RG} if it is impossible to constrain the RG flow to stay on a finite subspace. If we add the additional requirement that the theory converges toward a UV fixed point g^* as μ is increased, we obtain Weinberg’s characterization of renormalizability as asymptotic safety (Weinberg, 1979, p. 802). The couplings of a theory satisfying this more sophisticated criterion of renormalizability, call it renormalizability_{AS}, remain finite at arbitrarily high energies. As we will see in section 6, this is a good sign that the predictions of the theory remain under good mathematical control at high energies.³¹

Third, the RG theory implies that the scope of the continuum approach is not as restricted as initially considered. The definition of relevant, irrelevant and marginal operators by means of dimensional analysis in Eq. 16 is only valid near a fixed point. In general, this is a good rule of thumb. But it is perfectly possible that some non-perturbative effects come into play either at low or high energies. In particular, it is perfectly possible that a coupling which looks UV-relevant at low energies actually happens to be well-behaved at high energies because of fortuitous non-perturbative quantum corrections. That is, a theory can both converge towards a UV Wilson-Fisher fixed point and include non-renormalizable interactions, as some physicists expect for the quantum field theoretic approach to quantum gravity.³² Likewise, we might attempt to tame the pathological UV behavior of a given theory by embedding it into a larger theory displaying a UV fixed point. Overall, this suggests that the continuum approach is suitable for a larger class of physically relevant theories than initially expected. Yet, there still remains a large number of theories ill-handled under this approach, namely those which fail to converge towards a UV fixed point. As we will see in the next section,

³¹We can also speak of degrees of renormalizability_{AS} or “approximate” renormalizability_{AS} in the case of a renormalizable_{AS} theory with additional UV-relevant interaction terms diverging at high energies if the contributions of these UV-relevant interactions are negligible compared to the contributions of the other interactions at low energies.

³²See Niedermaier and Reuter (2006) for a review of the asymptotic safety scenario.

if we take the current perturbatively renormalizable formulation of the Standard Model by itself for instance, there are very good reasons to believe that it exhibits a Landau pole singularity and therefore makes inconsistent predictions at very high energies.

6 The infinite cut-off limit and the continuum limit

The goal of this section is to distinguish between different types of QFTs on the basis of their behavior in the continuum limit. For the sake of the argument, I will assume that the theory at hand has been renormalized under the continuum approach, except that we have kept the parameter Λ_0 fixed and not yet attempted to take the limit $\Lambda_0 \rightarrow \infty$. I will also assume that the analysis applies both to the perturbative and the non-perturbative case (with specific provisos when needed and with the speculative assumption that the non-perturbative notion of continuum QFT makes sense in realistic cases).

The first thing to note is that the notion of “continuum limit” is ambiguous. It may refer either to the *infinite cut-off limit* ($\Lambda_0 \rightarrow \infty$) or to the *continuum limit* ($\mu \rightarrow \infty$), properly speaking.³³ Note, moreover, that the distinction is robust under different regularization methods. For instance, the infinite cut-off and continuum limits correspond respectively to $\epsilon \rightarrow 0$ and $\mu \rightarrow +\infty$ for dimensional regularization. Likewise, using this terminology, taking the lattice spacing to zero in a lattice QFT amounts to taking the infinite cut-off limit, except in cases where the lattice spacing also plays the role of the renormalization scale (in which case there is only one type of limit).

Accordingly, the notion of “good behavior” in the limit should be understood in two distinct ways:

- (1) The low energy predictions of the theory at μ are unaffected by taking the infinite cut-off limit $\Lambda_0 \rightarrow \infty$.
- (2) The theory makes consistent predictions at arbitrarily high energies ($\mu \rightarrow \infty$).

(1) corresponds to cases where the low energy physics described by the theory is sufficiently insensitive to the high energy physics described by the theory, while (2) corresponds to cases where the high energy predictions of the theory do not violate typical assumptions such as unitarity.³⁴ So (2) is not about empirical adequacy, properly speaking. After all, the theory might turn out to be empirically inaccurate at very high energies. But we should at least require that it makes consistent predictions, say, by making sure that the values of the couplings remain sufficiently small for unitarity to hold.

³³For a slightly different understanding of this distinction emphasizing the difference between the removal of a perturbative regulator and the removal of a non-perturbative regulator, see Delamotte (2012, sec. 2.6; esp. 2.6.3).

³⁴Unitarity is the assumption that the total sum of probabilities for the possible measurement outcomes of some specific physical process add up to unity.

How should we discriminate between well-behaved and ill-behaved theories in the two cases then? Consider first the case of the infinite cut-off limit. Let us bracket any issue about the violation of the perturbative assumption in the case of the bare theory, and simplify the discussion by looking at the following toy-model theory $\mathcal{L}_R(\mu)$ with two renormalized couplings (g_0 and g'_0 correspond to the bare couplings):

$$\begin{cases} g(\mu) = \left(\frac{\mu}{\Lambda_0}\right)^{-\Delta} g_0(\Lambda_0) \\ g'(\mu) = \left(\frac{\mu}{\Lambda_0}\right)^{\Delta} g'_0(\Lambda_0) \\ \Delta > 0 \end{cases} \quad (17)$$

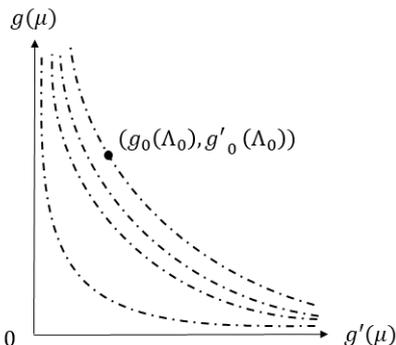


Fig. 4: Theory space in two dimensions with one IR-relevant/UV-irrelevant coupling g and one IR-irrelevant/UV-relevant coupling g' .

The condition $\Delta > 0$ implies that g is super-renormalizable and g' non-renormalizable or, equivalently, that g and g' are respectively relevant and irrelevant near the IR Gaussian fixed point, and irrelevant and relevant near the UV Gaussian fixed point. In general, RG equations define families of solutions parametrized by different boundary conditions and, in the present case, each solution $(g(\mu), g'(\mu))$ of the two dimensional RG equation is uniquely determined by specifying a single point $(g_0(\Lambda_0), g'_0(\Lambda_0))$ for some Λ_0 (see Fig. 4). Inversely, if the value of the renormalized couplings at μ is fixed by means of experiments, we can analyze the behavior of the bare theory in the infinite cut-off limit.

This toy-model is interesting because it displays two common types of behaviors. (i) $g'(\mu) = 0$ and $g(\mu) \neq 0$: that is, we do not detect non-renormalizable effects at low energies and take the liberty to fine-tune $g'_0(\Lambda_0)$ to zero, which implies $g'(\mu) = 0$. In this case, the RG flow lies on what is called the “renormalized trajectory” (see Fig. 3 and the $g(\mu)$ axis in Fig. 4) and we can take the infinite cut-off limit by assuming that the relevant bare coupling appropriately vanishes at infinity (i.e., such that $\lim_{\Lambda_0 \rightarrow \infty} \Lambda_0^\Delta g_0(\Lambda_0)$ is finite). It is therefore possible to take the infinite cut-off limit without affecting the low energy predictions of the theory.

(ii) The most likely case is that both $g(\mu)$ and $g'(\mu)$ are non-zero, i.e., that the theory contains UV-relevant couplings. The toy-model indicates that the constraints we need to impose on these couplings are relatively minimal: the infinite cut-off limit leaves the low energy predictions intact with the appropriate limits $g_0(\infty) = 0$ and $g'_0(\infty) = +\infty$. Of course, in general, taking the limit $\mathcal{L}_R(\mu) = \lim_{\Lambda_0 \rightarrow \infty} (\mathcal{L}_0 + \delta\mathcal{L})$ might require some delicate fine-tuning with the bare theory; and, as already emphasized, the perturbative assumption is explicitly violated in the case of the bare theory. But, overall, the RG theory is highly permissive since it is possible to take the infinite cut-off limit (at least formally) even if the theory contains pathological UV-relevant couplings. As we will see shortly, this fails to be the case if the renormalized coupling diverges at some finite energy scale Λ_M on the way to the limit.

Consider now the case of the continuum limit. RG flows towards the UV fall under four main types (e.g., Weinberg, 1996, sec. 18.3).³⁵ (i) Asymptotic freedom ($g^* = 0$) and (ii) asymptotic safety ($g^* \neq 0$) are the best case scenarios. In both cases, the values of the renormalized couplings remain finite in the continuum limit, which is a good sign that the theory makes consistent predictions at high energies since the main source of violations of (perturbative) unitarity comes from arbitrarily large values of the renormalized couplings in the expression of the scattering amplitudes. Of course, in those two cases as much as in the two cases below, our confidence in the behavior of couplings across energy scales depends on the reliability of the methods used to derive their expression. Asymptotic freedom is a special case in that respect. It is firmly based on perturbation theory like many of the results usually obtained from renormalization theory. But the fact that the values of the couplings become arbitrarily small at very high energies justifies the use of perturbation theory in the first place and suggests that non-perturbative results do not spoil the asymptotic behavior of the theory (e.g., a non-perturbative contribution to a scattering amplitude depending on some factor $e^{-1/g(\mu)}$ becomes arbitrarily negligible for $\mu \rightarrow \infty$ if $\lim_{\mu \rightarrow \infty} g(\mu) = 0$).

(iii) Let me call “asymptotic UV instability” the type of limiting behavior characteristic of theories containing divergent UV-relevant interactions as μ tends to ∞ (e.g., g' in the toy-model above). This case is problematic because, in general, these divergent UV-relevant interactions lead to violations of (perturbative) unitarity at high energies. These interactions even contain explicit information about the energy scale where those violations of unitarity arise. That being said, it is still possible to define the perturbative expression of the renormalized theory in the infinite cut-off limit $\Lambda_0 \rightarrow \infty$ if we restrict the range of the parameter μ to low energies. Moreover, if we use a smooth regularization method, the renormalized theory still includes negligible contributions from arbitrarily high energy excitation states (compared to contributions from low energy states). Hence, even though the theory behaves badly in the continuum limit, the continuum assumption holds in this

³⁵Other cases include scale-invariance ($\beta(g) = 0$), in which case the RG does not flow, strictly speaking, and the cyclic behaviors mentioned above.

case for processes probed at sufficiently low energies.

(iv) Let me call “finite UV instability” the type of limiting behavior characteristic of theories containing a Landau pole, i.e., a finite energy scale Λ_M at which at least one of the couplings $g(\Lambda_M)$ diverges. As the ϕ^4 -theory shows, finite UV instability is the worst case scenario. The solution to the perturbative RG equation of the quartic coupling $\lambda_R(\mu)$ is given by (cf. Eq. 13):

$$\lambda_R(\mu) = \frac{\lambda_R(\mu')}{1 - \frac{3\lambda_R(\mu')}{16\pi^2} \ln\left(\frac{\mu}{\mu'}\right)} \quad (18)$$

Given a fixed experimental value $\lambda_R(E)$ at the energy scale E , the coupling λ_R diverges at $\Lambda_M = E \exp(16\pi^2/3\lambda_R(E))$. Similarly, if we evaluate the expression of the bare coupling at the scale $\mu = \Lambda_0$, the bare coupling diverges at the same finite scale Λ_M . And if we do not make any low energy measurement and decide to take $\lambda_R(\mu) = 0$, we have to give up the initial assumption that the theory is an interacting theory. So, overall, the theory cannot be consistently defined in the infinite cut-off and continuum limits.³⁶ Now, the framework of the continuum approach is such that it is possible to take the infinite cut-off limit at the level of perturbative scattering amplitudes if we restrict ourselves to the first few orders in perturbation theory (see Eq. 8). However, the RG reveals that the partial perturbative relationship between the bare and the renormalized coupling obtained from the renormalization procedure is only superficially well-defined in the infinite cut-off limit: if we include the leading logarithms at higher orders in perturbation theory (as the derivation of the RG automatically does), we find a Landau pole.

In sum, continuum QFTs are likely to make consistent predictions at high energies when they are known with confidence to have a fixed point. The most reliable property of QFTs that we can typically find by means of perturbative methods is asymptotic freedom. And, for the large majority of continuum QFTs, there are good reasons to believe that they are not only conceptually incoherent and physically dubious but also that they make inconsistent predictions at high energies—or, at the very least, that standard perturbative techniques cannot be used in those cases. Table 1 below summarizes the main interpretative differences between the effective and the continuum approach, including the results from section 6.

³⁶Of course, it might be the case that the Landau pole turns out to be an artifact of the perturbative analysis. For a discussion about the existence of a Landau pole and triviality in the case of QED, see, e.g., Gockeler et al. (1998a,b); Gies and Jaeckel (2004).

	Effective approach	Continuum approach
The continuum assumption	False	True
Goal	Select the appropriate low energy degrees of freedom	Define the theory across all length scales
Bare theory	Physical theory	Intermediary/initial mathematical tool
Renormalized theory	Effective theory	Physical theory
Regulator Λ_0	The scale at which the theory breaks down	Intermediary mathematical tool
Regularization and renormalization method (most conceptually consistent)	Sharp cut-off (Λ)	Smooth cut-off (μ)
Infinite cut-off limit	Physically irrelevant	Mandatory
Continuum limit	Physically irrelevant	Consistent for a restricted class of well-behaved theories
Perturbative renormalizability	Perturbative predictions within ϵ with a finite number of parameters	Exact perturbative predictions with a finite number of parameters
Non-perturbative renormalizability	Finite dimensional RG surface within ϵ + IR fixed point	Finite dimensional RG surface + UV fixed point

Table 1: Main interpretative differences between the effective and the continuum approach.

7 Butterfield and Bouatta on continuum QFTs

An advocate of the axiomatic approach might raise the following objection at this point: why should we take the differences between the effective and the continuum approach seriously if both fail to meet satisfying standards of mathematical rigor in the first place? And why should we attach any importance to the good behavior of asymptotically safe QFTs as opposed to finitely unstable QFTs if there is a chance that they are both mathematically inconsistent and *a fortiori* physically incoherent?³⁷ Wallace (2006, sec. 3-4; 2011, sec. 6-9) has rightly argued, I believe, that effective Lagrangian QFTs are as well-defined as any of the past theories that we usually take to be mathematically well-defined, and therefore should be considered fit for foundational and philosophical scrutiny. Butterfield and Bouatta (2014) recently extended this claim to continuum QFTs (see also Butterfield, 2014, pp. 8-9, sec. II.2-3, p. 31). They argue that even if the path integral formulation of realistic continuum QFTs has not yet received a precise mathematical definition

³⁷One might see these objections as particular cases of the general objection that the conventional mathematical apparatus of QFTs is ill-defined (e.g., Halvorson, 2007, p. 731; Fraser D., 2008, p. 550; Kuhlmann, 2010; Baker, 2016, p. 5; Summers, 2016).

according to the standards exhibited in the axiomatic, algebraic and constructive programs, some of these theories appear to be sufficiently mathematically well-defined according to physicists' common standards to be fit for philosophical scrutiny. Hence, by endorsing less stringent criteria of mathematical rigor, they claim, we should feel confident to draw world pictures for the heuristic formulation of *some* continuum QFTs. I will argue that the methodological and conceptual differences between the effective and the continuum approach discussed in sections 3-6 suggest reasons to temper Butterfield and Bouatta's claim.

Let me begin by making two friendly amendments to their discussion of continuum QFTs. (i) They contend that the contrast between theories likely to be (A) mathematically well-defined and (B) mathematically ill-defined depend, broadly speaking, "on the type of fields in the theory concerned" (Butterfield and Bouatta, 2014, p. 65). Agreed: as Butterfield and Bouatta rightly emphasize, in four dimensions, QFTs including only non-abelian gauge fields fall under case (A) while QFTs including only scalar or fermionic fields typically fall under case (B). In general, however, the field content of a QFT does not provide a reliable guide to assess whether the QFT is mathematically well-defined or not. Examples of asymptotically free scalar and fermionic QFTs in two and three dimensions show that the mathematical well-definedness of a QFT is not simply determined by the type of its quantum field operators.³⁸ In contrast, the scaling behaviors of QFTs exhibited by means of RG methods offer a more systematic way of distinguishing between (A) and (B), and Butterfield and Bouatta's diagnosis somewhat obscures the remarkable fact that this criterion does not depend on the content of the theory. Agreed, the definition of a particular RG space depends on the specification of a set of couplings and therefore on the specification of a set of (local) interactions—which, in turn, depends on the specification of a set of fields (e.g., scalar, fermionic, gauge, etc.), symmetries, and a space-time dimension. However, the possible types of RG trajectories, i.e., the possible types of behaviors of theories across energy scales, do not depend on these constraints. And so what it means for a theory to be mathematically well-defined is independent of the specific QFT model considered.

(ii) Butterfield and Bouatta's classification of QFTs under (A) and (B) is also incomplete. They argue that we should group asymptotically free, safe and conformal theories under (A) and theories presenting a Landau pole under (B). Agreed, this provides a good rule of thumb for the high energy limit of continuum theories; and, for the perturbative theories we have so far, there are, in general, good reasons to expect that Landau poles in the IR ("infrared slavery") are perturbative artifacts, as it seems for perturbative QCD. However, it is worth being more systematic here since the non-perturbative definition of a theory might display, say, a Landau pole in the IR and asymptotic freedom in

³⁸See, e.g., Weinberg (1996, sec. 18.3); Gross (1999, lecture 3, sec. 3.2; lecture 4). Examples of asymptotically safe theories in lower dimensions involving scalar or fermionic fields include: the Gross-Neveu model, the nonlinear σ -model with dimension $2 < d < 4$, and the 2d sine-Gordon model.

the UV. I distinguished in section 6 between (a) finite instability (i.e., existence of a Landau pole), (b) asymptotic instability (i.e., asymptotically divergent couplings), (c) asymptotic freedom (i.e., convergence to a zero fixed point), and (d) asymptotic safety (i.e., convergence to a non-zero fixed point), to which we might add the two additional cases of (e) non-convergent cyclic scaling behavior (i.e., non-convergent oscillating couplings) and (f) scale-invariant theories (i.e., theories defined at a fixed point). It is perfectly possible that the non-perturbative definition of a theory displays two properties out of the five (a)-(e).

We should therefore only include under (A) theories defined by a continuous RG flow between two distinct fixed points and theories defined at a fixed point (ignoring (e)). The first class of theories corresponds to the class of IR/UV asymptotically safe/free theories, i.e., theories that continuously connect two conformal theories in the RG space.³⁹ For instance, the RG equation $\mu dg/d\mu = Ag - Bg^2$ for some coupling g with $A, B > 0$ describes the behavior of a theory asymptotically free in the IR (flowing towards the fixed point $g^* = 0$) and asymptotically safe in the UV (flowing towards the fixed point $g^* = A/B$). The second class of theories corresponds to the class of scale-invariant theories (i.e., $dg/d\mu = 0$). Although this has not been proven for models in dimension $d > 2$, these scale-invariant theories can typically be formulated as conformal field theories (CFTs).⁴⁰ Moreover, since our confidence in the existence of the properties (c), (d) and (f) is usually based on perturbative methods (as Butterfield and Bouatta rightly recognize), we should add the further constraint $g^* \ll 1$ for perturbation theory to be reliable.⁴¹

Let us now turn to Butterfield and Bouatta's claim that some continuum QFTs are ripe for metaphysical inquiry. At least as I understand them, Butterfield and Bouatta's claim relies on two key ideas. First, physics exhibits various standards of mathematical well-definedness and mathematical existence, and the heuristic standard commonly used in physics' practice provides a perfectly reasonable standard for interpretative purposes. In the context of QFT, the heuristic standard requires the theory to have a finite UV scaling behavior. By contrast, a theory is mathematically well-defined according to the axiomatic standard if it is axiomatizable and has a consistent model (Butterfield and Bouatta, 2014, p. 69). Second, the current perturbative formulation of some realistic continuum Lagrangian QFTs displays a UV fixed point and therefore satisfies the heuristic standard. QCD is one such example. Of course, the finite behavior of the theory at high energies does not mean that the functional integral resulting from the path integral quantization of the classical Lagrangian density is mathematically well-defined according to more stringent criteria of rigor. But the lack of a mathematically well-defined formulation should not prevent us from interpreting

³⁹For reference to the existence of well-defined and non-trivial RG flows from IR to UV fixed points, see, e.g., Caswell (1974); Banks and Zaks (1982); Bond and Litim (2017).

⁴⁰For references and discussions, see Polchinski (1988); Dymarsky et al. (2015).

⁴¹Here it is worth mentioning the efforts made to formulate non-perturbative theories in the asymptotic safety scenario programme briefly mentioned in section 5.2.

the heuristic formulation of our best continuum QFTs (Butterfield, 2014, p. 15). I take it that when Butterfield and Bouatta speak of the “heuristic” formulation of QFTs (Butterfield and Bouatta, 2014, p. 64, p. 68; Butterfield, 2014, p. 15), they refer to the current perturbative formulation that we have of these theories. And by ‘perturbative formulation’ I mean the formal expression of the path integral and the perturbative expression of the renormalized action and Lagrangian together with the set of perturbative techniques used to compute correlation functions.

Now, even if we accept to endorse less stringent criteria of mathematical rigor and philosophically assess the heuristic formulation of some continuum QFTs, it does not mean that we are warranted in attempting to draw “ontological claims” or “world-pictures” for these continuum QFTs (e.g., Butterfield and Bouatta, 2014, p. 68). It was central to the argument of section 3 that the structure of a physical theory does not only need to be under good mathematical control but also needs to make physical sense. Even if a QFT has a finite behavior at all energy scales, it is no indication that the theory has a physically coherent interpretation. Agreed, we do not need to demand that all the component parts of the theory make physical sense in order to dive into the metaphysical interpretation of a theory. But we should at least require that the core component parts of the theory do. Section 3 suggests that the perturbative formulation of our best continuum QFTs does not even meet this requirement in contrast to effective QFTs.

The argument goes as follows. To simplify the discussion and as already emphasized, I will follow Butterfield’s usage of the term ‘theory’ in its specific sense and identify the perturbative formulation of a QFT with the perturbative formulation of its Lagrangian (Butterfield, 2014, p. 31). Then, we may either interpret the renormalized Lagrangian or the bare Lagrangian (or both) in order to extract dynamical information. Consider first the renormalized Lagrangian. However we construct it, the renormalized Lagrangian together with the standard rules for deriving amplitudes yields divergent quantities if we do not restrict the state space of the theory. Hence, if the goal is to interpret empirically successful theories, we have no reason to even attempt to draw a world picture out of the renormalized theory or to take the renormalized Lagrangian to give us reliable dynamical information about the target system. At the very least, we should show some degree of caution.

Let us look at the bare Lagrangian. In the least naive perturbative construction of a renormalized continuum QFT, we start with some initial bare Lagrangian with the “wrong” parameters and we split it into a physical Lagrangian and a counter-term Lagrangian. The split is made in such a way that the counter-terms cancel the original divergences in the scattering amplitudes derived from the bare Lagrangian. And, by re-expressing the parameters of the bare Lagrangian, we find that the original bare amplitude is actually finite. The problem, however, is that the parameters of the bare Lagrangian diverge in the infinite cut-off limit. We precisely use the freedom that we have in defining the original bare parameters to absorb the UV divergences that we find in the original perturbative expansion. So, at least at this level, the

original expression of the bare theory makes little physical sense. How about the “true” bare theory, i.e., the theory defined by the renormalized parameters evaluated at the cut-off Λ_0 (see section 3.2)? As already emphasized in section 3.3, there are concrete examples of theories where these bare parameters converge to a finite value in the infinite cut-off limit. However, if we choose to identify the bare parameters in this way, the resulting bare Lagrangian yields, once again, divergent predictions. Finite amplitudes are always derived from the theory which has the “wrong” parameters, as it were, since we always need to re-express the original couplings of the divergent amplitudes in order to absorb the divergences. And so whether we look at the renormalized or the bare Lagrangian, it does not appear that we can justifiably draw a world picture out of the perturbative formulation of a continuum QFT constructed under the continuum approach.

8 Conclusion

The aim of this paper has been twofold: (i) to propose a general conceptual framework to understand the various aspects of renormalization theory based on the distinction between effective and continuum QFTs; and (ii) to show that the effective approach to renormalization offers a more physically perspicuous, conceptually coherent and widely applicable way to construct perturbative QFTs in comparison to the continuum approach. The oddities of the continuum approach are best illustrated by the absence of physical justification for the introduction of counterterms, the instrumental status of the bare theory, and the fact that, strictly speaking, the renormalized theory yields divergent amplitudes if we do not restrict the state space of the theory. Evaluating the limiting behavior of continuum QFTs also provides important conceptual and classificatory insights into the scope of the continuum approach: only asymptotically safe and free theories are likely to make consistent predictions at high energies in contrast to asymptotically and finitely unstable theories. In comparison, the effective approach is applicable to any local QFT model (as far as I am aware). The paper concluded with some lessons for the debate over the interpretation of QFTs in response to Butterfield and Bouatta’s paper (2014): the distinction between the effective and the continuum approach gives reasons to doubt that perturbative continuum QFTs are yet ripe for metaphysical analysis.

Acknowledgments

I would like to thank David Albert, Jonathan Bain, Alexander Blum, Jeremy Butterfield, Bertrand Delamotte, Michael Miller, David Wallace, Porter Williams, as well as audiences at Oxford and Utrecht, for many helpful discussions and/or correspondences on the materials discussed in this paper. David Albert, Jonathan Bain, Jeremy Butterfield, Michael Miller, David Wallace, Porter Williams, as well as three anonymous referees, provided invaluable comments, suggestions and/or criticisms on previous drafts; the final product has greatly benefited from their generosity. .

References

- Bagnuls, C. and C. Bervillier (2001). Exact Renormalization Group Equations. An Introductory Review. *Physics Reports* 348(1-2), 91–157.
- Bain, J. (2013). Effective Field Theories. In R. Batterman (Ed.), *The Oxford Handbook of Philosophy of Physics*, pp. 224–254. Oxford: Oxford University Press.
- Baker, D. J. (2016). The Philosophy of Quantum Field Theory. *Oxford Handbooks Online*. Philsci preprint: <http://philsci-archive.pitt.edu/12678/>.
- Banks, T. and A. Zaks (1982). On the Phase Structure of Vector-like Gauge Theories with Massless Fermions. *Nuclear Physics B* 196(2), 189–204.
- Bond, A. D. and D. F. Litim (2017). Theorems for Asymptotic Safety of Gauge Theories. *The European Physical Journal C* 77(6), 429. ArXiv:1608.00519.
- Bulycheva, K. and A. Gorsky (2014). RG Limit Cycles. ArXiv:1402.2431.
- Butterfield, J. (2014). Reduction, Emergence and Renormalization. *Journal of Philosophy* 111(1), 5–49.
- Butterfield, J. and N. Bouatta (2014). On Emergence in Gauge Theories at the 't Hooft Limit. *European Journal for Philosophy of Science* 5(1), 55–87.
- Butterfield, J. and N. Bouatta (2015). Renormalization for Philosophers. In T. Bigaj and C. Wuthrich (Eds.), *Metaphysics in Contemporary Physics*, pp. 437–485. Leiden: Brill Rodopi. ArXiv:1406.4532.
- Cao, T. Y. and S. S. Schweber (1993). The Conceptual Foundations and the Philosophical Aspects of Renormalization Theory. *Synthese* 97(1), 33–108.
- Caswell, W. E. (1974). Asymptotic Behavior of Non-Abelian Gauge Theories to Two-Loop Order. *Physical Review Letters* 33(4), 244–246.
- Collins, J. (1986). *Renormalization: An Introduction to Renormalization, the Renormalization Group and the Operator-Product Expansion*. Cambridge: Cambridge University Press.
- Collins, J. (2009). The Problem Of Scales: Renormalization and All That. ArXiv:hep-ph/9510276.
- Collins, J. (2011). *Foundations of Perturbative QCD*. Cambridge: Cambridge University Press.

-
- Crowther, K. and N. Linnemann (2017). Renormalizability, Fundamentality, and a Final Theory: The Role of UV-Completion in the Search for Quantum Gravity. *The British Journal for the Philosophy of Science* *axx052*. Philsci preprint: <http://philsci-archive.pitt.edu/13219/>.
- Delamotte, B. (2012). An Introduction to the Nonperturbative Renormalization Group. In A. Schwenk and J. Polonyi (Eds.), *Renormalization Group and Effective Field Theory Approaches to Many-Body Systems*, Lecture Notes in Physics, pp. 49–132. Springer-Verlag Berlin Heidelberg.
- Duncan, A. (2012). *The Conceptual Framework of Quantum Field Theory*. Oxford: Oxford University Press.
- Dymarsky, A., Z. Komargodski, A. Schwimmer, and S. Theisen (2015). On Scale and Conformal Invariance in Four Dimensions. *Journal of High Energy Physics* (10), 171. ArXiv:1309.2921.
- Franklin, A. (2018). Whence the Effectiveness of Effective Field Theories? *The British Journal for the Philosophy of Science* *axy050*. Philsci preprint: <http://philsci-archive.pitt.edu/14731/>.
- Fraser, D. (2008). The Fate of 'Particles' in Quantum Field Theories with Interactions. *Studies in History and Philosophy of Science Part B* *39*(4), 841–859.
- Fraser, J. D. (2017). The Real Problem with Perturbative Quantum Field Theory. *The British Journal for the Philosophy of Science* *axx042*. Philsci preprint: <http://philsci-archive.pitt.edu/13348/>.
- Fraser, J. D. (2018). Towards a Realist View of Quantum Field Theory. Philsci preprint: <http://philsci-archive.pitt.edu/14716/>.
- Georgi, H. (1992). Thoughts on Effective Field Theory. *Nuclear Physics B - Proceedings Supplements* *29*(2), 1–10.
- Georgi, H. (1993). Effective Field Theory. *Annual Review of Nuclear and Particle Science* *43*(1), 209–252.
- Gies, H. and J. Jaeckel (2004). Renormalization Flow of QED. *Physical Review Letters* *93*(11).
- Gockeler, M., R. Horsley, V. Linke, P. Rakow, G. Schierholz, and H. Stuben (1998a). Is There a Landau Pole Problem in QED? *Physical Review Letters* *80*(19), 4119–4122.
- Gockeler, M., R. Horsley, V. Linke, P. Rakow, G. Schierholz, and H. Stuben (1998b). Resolution of the Landau Pole Problem in QED. *Nuclear Physics B - Proceedings Supplements* *63*(1), 694–696.
- Gross, D. J. (1999). Renormalization Groups. In P. Deligne, P. Etingof, D. S. Freed, L. C. Jeffrey, D. Kazhdan, J. W. Morgan, D. R. Morrison, and E. Witten (Eds.), *Quantum Fields and Strings: A Course for Mathematicians. Vol. 1*, pp. 551–596. Providence: American Mathematical Society.

-
- Halvorson, H. (2007). Algebraic Quantum Field Theory. In J. Butterfield and J. Earman (Eds.), *Handbook of the philosophy of physics, part A*, pp. 731–864. North Holland.
- Helling, R. C. (2012). How I Learned to Stop Worrying and Love QFT. *ArXiv:1201.2714*.
- Hollowood, T. J. (2013). *Renormalization Group and Fixed Points*. SpringerBriefs in Physics. Springer.
- Huggett, N. and R. Weingard (1995). The Renormalisation Group and Effective Field Theories. *Synthese* 102(1), 171–194.
- Huggett, N. and R. Weingard (1996). Exposing the Machinery of Infinite Renormalization. *Philosophy of Science* 63(3), S159–S167.
- Kuhlmann, M. (2010). Why Conceptual Rigour Matters to Philosophy: On the Ontological Significance of Algebraic Quantum Field Theory. *Foundations of Physics* 40(9-10), 1625–1637.
- Lepage, G. P. (1989). What is Renormalization? In T. DeGrand and D. Toussaint (Eds.), *From Actions to Answers, Proceedings of the 1989 Theoretical Study Institute in Elementary Particle Physics*, pp. 483–509. World Scientific.
- Miller, M. (2016). Mathematical Structure and Empirical Content. Philsci preprint: <http://philsci-archive.pitt.edu/12678/>.
- Niedermaier, M. and M. Reuter (2006). The Asymptotic Safety Scenario in Quantum Gravity. *Living Rev. Relativity* 9.
- Osborn, H. (2016). Advanced Quantum Field Theory Lecture Notes. Available at <http://www.damtp.cam.ac.uk/user/ho/AQFTNotes.pdf>.
- Peskin, M. E. and D. V. Schroeder (1995). *An Introduction to Quantum Field Theory*. Chicago: Westview Press.
- Polchinski, J. (1988). Scale and Conformal Invariance in Quantum Field Theory. *Nuclear Physics B* 303(2), 226–236.
- Polonyi, J. (2003). Lectures on the Functional Renormalization Group Method. *Open Physics* 1(1), 1–71.
- Rosaler, J. and R. Harlander (2019). Naturalness, Wilsonian Renormalization, and "Fundamental Parameters" in Quantum Field Theory. *Studies in History and Philosophy of Modern Physics*. Philsci preprint: <http://philsci-archive.pitt.edu/15810/>.
- Scharf, G. (1995). *Finite Quantum Electrodynamics: The Causal Approach* (2nd ed.). Springer-Verlag Berlin Heidelberg.
- Schwartz, M. D. (2013). *Quantum Field Theory and the Standard Model*. Cambridge: Cambridge University Press.
- Steinmann, O. (2000). *Perturbative Quantum Electrodynamics and Axiomatic Field Theory*. Springer-Verlag Berlin Heidelberg.

-
- Streater, R. F. and A. S. Wightman (2000). *PCT, Spin and Statistics, and All That*. Princeton University Press.
- Summers, S. J. (2016). A Perspective on Constructive Quantum Field Theory. ArXiv:1203.3991.
- Teller, P. (1988). Three Problems of Renormalization. In H. R. Brown and R. Harre (Eds.), *Philosophical Foundations of Quantum Field Theory*, pp. 73–89. Oxford University Press.
- Teller, P. (1989). Infinite Renormalization. *Philosophy of Science* 56(2), 238–257.
- Wallace, D. (2006). In Defence of Naivete: The Conceptual Status of Lagrangian Quantum Field Theory. *Synthese* 151(1), 33–80.
- Wallace, D. (2011). Taking Particle Physics Seriously: A Critique of the Algebraic Approach to Quantum Field Theory. *Studies in History and Philosophy of Science Part B* 42(2), 116–125.
- Weinberg, S. (1979). Ultraviolet Divergences in Quantum Theories of Gravitation. In S. W. Hawking and W. Israel (Eds.), *General Relativity: An Einstein centenary survey*, pp. 790–831. Cambridge: Cambridge University Press.
- Weinberg, S. (1995). *The Quantum Theory of Fields, Vol. 1: Foundations*. Cambridge: Cambridge University Press.
- Weinberg, S. (1996). *The Quantum Theory of Fields, Vol. 2: Modern Applications*. Cambridge: Cambridge University Press.
- Weinberg, S. (2000). *The Quantum Theory of Fields, Vol. 3: Supersymmetry*. Cambridge: Cambridge University Press.
- Williams, P. (2015). Naturalness, the Autonomy of Scales, and the 125 GeV Higgs. *Studies in History and Philosophy of Science Part B* 51, 82–96.
- Williams, P. (2018). Renormalization Group Methods. In E. Knox and A. Wilson (Eds.), *The Routledge Companion to the Philosophy of Physics*. Routledge. Philsci preprint: <http://philsci-archive.pitt.edu/15346/>.
- Wilson, K. G. (1971). Renormalization Group and Strong Interactions. *Physical Review D* 3(8), 1818–1846.
- Zinn-Justin, J. (2007). *Phase Transitions and Renormalization Group*. Oxford: Oxford University Press.