

Renormalization Group Methods and the Epistemology of Effective Field Theories

Adam Koberinski* † and Doreen Fraser‡ §

Abstract

The effective field theory (EFT) perspective on particle physics has yielded insight into the Standard Model. This paper investigates the epistemic consequences of the use of different variants of renormalization group (RG) methods as part of the EFT perspective on particle physics. RG methods are a family of formal techniques. While the semi-group variant of the RG has played a prominent role in condensed matter physics, the full-group variant has become the most widely applicable formalism in particle physics. We survey different construction techniques for EFTs in particle physics and analyze the role that semi-group and full-group variants of the RG play in each. We argue that the full-group variant is best suited to answering structural questions about relationships among EFTs at different scales, as well as explanatory questions, such as why the Standard Model has been empirically successful at low energy scales and why renormalizability was a successful criterion for constructing the Standard Model. We also present an account of EFTs in particle physics that is based on the full-RG. Our conclusion about the advantages of the full-RG is restricted to the particle physics case. We argue that a domain-specific approach to interpreting EFTs and RG methods is needed. Formal variations and flexibility in physical interpretation enable RG methods to support different explanatory strategies in condensed matter and particle physics. In particular, it is consistent to maintain that coarse-graining is an essential component of explanations in condensed matter physics, but not in particle physics.

1 Introduction

Since the late 1970s, there has been a growing consensus in particle physics that quantum field theories (QFTs) are best understood as effective field theories (EFTs), constructed with an explicit reference to the energy scale at which the theory no longer applies. Historically, we can understand this shift to thinking of QFTs as EFTs as a way to (dis)solve the mystery of renormalization in particle physics (Butterfield and Bouatta 2016), or as a means to construct a theory space that goes beyond the Standard Model (Weinberg 1979). The renormalization group (RG) played a major role in understanding some aspects of renormalization, and in analogy with condensed matter physics one could then see QFTs as an effective description of physics relative to some reference energy or length scale.

From a foundational standpoint, philosophers have begun to focus on EFTs for a few reasons. First and perhaps most obvious, philosophers of physics should be exploring the foundations of our current best theories as they are currently understood (Wallace 2006; Wallace 2011). Since a majority of physicists now view the Standard Model as an EFT, philosophers ought to make sure that they understand the structure of the EFT framework. Second, and growing from this, many philosophers have noticed that EFTs do not fit neatly into standard philosophical understandings of theories, models, and the relationships between the two. Some have argued that EFTs lead to a hierarchy of merely effective descriptions in nature, with no bottoming out at the fundamental level (Cao and Schweber 1993; Bain 2013). Arguing against this position, recent work has taken the EFT framework and the RG to give rise to a new prospective realism and a foundation for emergence as separate from reduction (Crowther 2015; J. D. Fraser 2018; Williams 2019; Wallace 2019). Meanwhile other work has focused on the assumptions needed to set up the EFT

*Center for Philosophy of Science, University of Pittsburgh, Pittsburgh, PA, USA

†adam.koberinski@pitt.edu

‡Department of Philosophy, University of Waterloo, Waterloo, ON N2L 3G1, Canada

§dlfraser@uwaterloo.ca

framework (Williams 2015; Rivat 2020; Koberinski and Smeenk 2022). We note here that there is sometimes a tendency to slip between precise technical notions grounded in physics—such as EFTs, coarse-graining, and RG relations—and more general concepts broadly relevant to philosophy of science—such as theories with limited scope, changing degrees of freedom, and general intertheoretic relations, respectively—using the same terms to refer to both the narrow and broad concepts. While we wholeheartedly agree that the former have important consequences for understanding the latter, the technical notions are much more limited in scope and entail specific mathematical relationships not necessarily entailed by the more general concepts. We will exclusively use the terms ‘EFT’, ‘coarse-graining’, and ‘RG’ to refer to the precise concepts from QFT and condensed matter physics.

In this paper, we will focus our attention on the epistemic dimensions of the construction and use of EFTs in particle physics. Our conclusions about the interpretation of EFTs and the explanations that they support may have implications for metaphysical debates about emergence and reduction, but those are outside of the scope of the present discussion. What many philosophical approaches to EFTs in particle physics have in common is a heavy dependence on analogies with condensed matter physics, resulting in an emphasis on coarse-graining and the semi-group variant of the RG as a basis for robust theoretical relations such as emergence, universality, or independence. Since EFTs and RG methods are used in both disciplines, it is tempting to understand particle physics—where physical interpretation is further removed from ordinary experience—in terms of condensed matter physics—where one has a better intuitive grasp on the physical situation and ontology. Due to the historical origins of RG methods, and the formal analogies with condensed matter physics, the dominant understanding of EFTs in particle physics utilizes the renormalization semigroup (semi-RG). This means that one often thinks of constructing an EFT by starting with a higher-energy theory, and coarse-graining by introducing a means to screen off effects of that theory beyond some specified energy scale. Using the semi-RG in condensed matter physics allows one to systematically transform from a theory that includes all fine-grained details relevant at one scale to different theories which exclude fine-grained details that are negligible at larger distance scales.

We argue that EFTs in particle physics are distinct from EFTs in condensed matter physics with respect to both epistemology and interpretation. This distinctness manifests in the different uses to which RG methods are put in particle physics. We begin in Sec. 2 by arguing that the group of mathematical techniques known as RG methods are diverse, and often loosely related. The family resemblance among RG methods is that all are mathematical tools meant to deal with problems where a wide range of weakly coupled scales are all important to the dynamics. However, because these problems arise across a variety of domains, one should expect to find that RG methods—and the physical interpretation attached to them—will likewise differ. We focus here on the semi-RG and the full renormalization group (full-RG) in particular, since these are the most relevant to current applications in particle physics. The individual formalisms that we will investigate do not wear their physical interpretations on their sleeves, but some physical interpretations do not fit naturally with some formalisms. For example, the full-RG does not fit with the physical interpretation of the RG transformation as representing a coarse-graining in which a change of variables leads to the discarding of fine-grained details because the full-RG transformation is reversible.

In Sec. 3, we review methods for the construction of EFTs in particle physics. The semi-RG plays a central role in the abstract top-down construction procedure, exemplified by reparameterizing the Wilson action for different values of the separation scale Λ . Concrete realizations of the top-down procedure do not involve iterative application of the semi-RG, though the (full-)RG equations are used in fixing low-energy couplings and calculating with the resulting EFT. However, neither the abstract nor concrete top-down procedures are readily applied to all EFTs in particle physics. The top-down procedure is relational, starting with a high-energy theory and integrating out the high-energy degrees of freedom. But for EFTs like chiral perturbation theory and soft-collinear effective theory, this perspective doesn’t apply. Instead, we should think of these EFTs from the bottom-up perspective. From this perspective, EFTs are constructed as independent theories, whose degrees of freedom are determined relative to the energy scales for which they will be used. These EFTs include a separation scale and an infinite set of coupling terms consistent with the symmetries and fields chosen for construction. Further, the bottom-up perspective can be applied to *all* EFTs in particle physics. We argue that, if one wants to pick a procedure with which to understand EFT construction in particle physics at a general level, then the bottom-up perspective is the only viable option. In constructing EFTs from the bottom-up, RG methods play less of a direct role. Only the full-RG plays a role, and this is in calculating amplitudes and determining scaling behaviour.

Next we examine the role of RG methods in *using* EFTs. The full-RG is needed for understanding the construction of EFTs in particle physics, but are there aspects of their use or structure that can only be captured by the semi-RG? In Sec. 4 we argue that the bottom-up perspective using the full-RG is able to account for the structure and use of EFTs in particle physics. In particular, we need never use the semi-RG as a coarse-graining procedure in particle physics. Whether one aims to explore the asymptotic scaling behaviour of a given EFT, understand why the renormalizable sector of the Standard Model EFT continues to be so successful, or explore the insensitivity of low-energy physics against some variations in high-energy EFTs, one can stick with the full-RG and prioritize the bottom-up perspective for all explanatory purposes. The top-down perspective can be useful in explicitly considering the relationship between classes of EFTs with differing couplings and cutoff scales, but this can also be accomplished with the full-RG; the semi-RG and the accompanying interpretation of coarse-graining is never needed. Sec. 4.1 focuses on the uses of RG methods for local empirical questions regarding a given EFT, while Sec. 4.2 examines explanatory questions dealing with the global EFT theory space. Sec. 4.3 provides a philosophical analysis of the explanatory value of semi-RG and full-RG characterizations of EFT theory space. We conclude that the full-RG perspective is best suited to understanding the frontiers of particle physics, and draw implications for other debates in the foundations of particle physics.

Finally, we turn back to the motivating analogy between EFTs in particle physics and condensed matter physics in Sec. 5. We argue that the formal analogy between the two disciplines actually highlights an important physical disanalogy, blocking the straightforward carryover of interpretation. The disanalogy involves the different quantities that vary using the semi-RG in the limit where the separation scale becomes unimportant.

We argue that significant interpretive choices are being made in the standard reading of EFTs and the RG. Making these interpretive choices explicit will serve to highlight the landscape of possible interpretations of EFTs and the RG. In particular, much stock is currently placed in lattice regularization and the Wilson action approach to RG flow. We believe that the differences between particle physics and condensed matter physics extend to the understanding of EFTs and the RG; an interpretation of the formalism should be discipline-specific. We aim to make this point clear by looking at how EFTs are constructed and used in particle physics, noting in particular the ways that RG methods come into play. These examples highlight our broader claim that interpretations of any formalism must be individualized to the physical context in which they are used. At the very least, this supports the point that one cannot read an interpretation directly off of mathematical formalism (Bokulich 2020).

To fix terminology at the outset: we take an EFT to be defined in terms of a Lagrangian with some set of coupling terms—in general infinite in number, but if a semi-RG is applied then finite—between relevant fields, subject to symmetry constraints, with coupling parameters fixed at some energy scale \mathcal{E} . Built into the coupling terms is a (potentially unfixed) separation scale, which we denote as Λ .

2 Renormalization group methods as a family of mathematical techniques

RG methods are a family of applied mathematical techniques. The diversity among formal techniques that fall into the category of RG methods is evident in both contemporary and past applications. For example, there are a variety of approaches to the use of perturbative approximations. Historically, Bogoliubov and Shirkov (1955) applied the term *renormalization group* to describe a technique that involved “renormalization-invariant improvement of perturbative results” (Shirkov 1999, p. 17). One of Wilson’s contributions was the introduction of a fundamentally non-perturbative formulation of a (different) renormalization group: the RG transformation can be defined non-perturbatively, though of course perturbation series are often used to evaluate the resulting expressions. More recently, different formulations of the RG have been introduced in the setting of algebraic QFT to facilitate the construction of models (e.g., the “exact” renormalization group and perturbative AQFT) (Hancox-Li 2015; Rejzner 2015). Our focus in this paper will be the uses of semi-group and full-group variants of RG methods in particle physics. Before turning to analysis of that particular case, we will offer some reflections on what it means for RG methods to be a family of formal techniques.

For our purposes, the characterization of RG methods that Wilson offered back in 1975 remains a useful

starting point for explaining how RG methods are a loosely related set of formalisms rather than a single, univocal formalism. The family resemblance is in both the type of the mathematical problem solving strategy and the character of the mathematical problem that the techniques are designed to solve. Wilson presents the renormalization group as the analogue of the derivative, and the statistical continuum limit as underlying the renormalization group in the same manner in which continuum limits underlie the derivative (Wilson 1975, pp. 773–774). A statistical continuum limit arises when limits are taken of functions of continuous variables and the values of the functions at each point in the continuum are treated as independent variables. Such a limit is statistical because the independent variables are typically fields that fluctuate; consequently, the end goal is to calculate averages or expectation values of products of the fields. For example, limits of this type are used to specify correlation functions in statistical mechanics or expressions for vacuum expectation values in QFTs. RG methods solve problems that directly or indirectly involve evaluating statistical continuum limits.

Wilson illustrates the statistical continuum limit and RG methods with examples of applications in QFT and statistical mechanics (Wilson 1975, pp. 773–774). This is a theme that will be picked up in this paper, so it is worth entering into some detail. In a QFT, the goal is to calculate vacuum expectation values (VEVs), or equivalently propagators or n -point functions. To make the problem tractable, the fields are initially regularized by imposing a high energy-momentum or low space(time) (i.e., lattice) cutoff. The statistical continuum limit arises when the cutoff is removed: either the high energy-momentum cutoff is taken to infinity or the space(time) lattice spacing is taken to zero. In a statistical mechanical model of a magnet, there is an atomic lattice that remains fixed. The statistical continuum limit increases the spatial extent of the region under consideration. At the critical point of the magnet, the correlation length becomes infinite and it is natural to take the infinite limit of the spatial extent of the region. In both cases, RG equations relate descriptions of the system at different scales. Scaling properties encapsulated in the RG equations can be leveraged to calculate the quantities of interest (VEVs or correlation functions). The salient family resemblances between these (and other) applications of RG methods are that the problem involves limits of functions (representing fields) that are independent variables, that the calculated quantities are averages or expectation values of these functions, and that the RG equations relate descriptions at different scales (e.g., length, energy).¹ These family resemblances are generally accompanied by many differences in physical interpretation. In the examples sketched in the preceding paragraph, the limits used to set up the problems and the scales have different physical interpretations. The starting point for the analysis in this paper will not be these differences in physical interpretation, but formal differences between implementations of the RG. Differences in physical interpretation between condensed matter and particle physics will be examined in Sec. 5.

Our primary focus will be on tracing the implications of one of the formal differences between different techniques that fall under the umbrella concept of RG methods: the uses of a semi-group or a full-group of RG transformations that arise in practice. The semi-group formulation of RG methods has a distinguished history in condensed matter physics. Kadanoff’s block-scaling transformation was the prototype for setting up the problem in condensed matter physics that RG methods were designed to solve. Consider an Ising model that represents a ferromagnet. A spin value is assigned to each atomic lattice site. The lattice is divided into blocks of lattice sites, and the block-scaling transformation averages the spins in each block and rescales the Hamiltonian representing the interactions among the spins accordingly. The averaging is a paradigmatic example of a *coarse-graining* operation: there is an iterative procedure in which one set of variables (e.g., spin variables) is replaced by another set of variables (e.g., average spin variables) and this results in the loss of fine-detail information about the initial variables (e.g., array of spin values).² Iterative application of the block-scaling transformation is useful near the critical point, when the correlation length of the system becomes large and the fine-detail, micro-level properties of the system are not relevant for determining certain macro-level properties.

¹Wilson stresses that other conditions need to obtain in order for RG methods to be applicable, including that the fluctuations over a large range of scales contribute to the values of the VEVs or correlation functions, that there are no characteristic scales that dominate the calculation, and that scales are locally coupled (e.g., 1000-2000Å wavelengths primarily affected by nearby wavelengths 500-1000Å and 2000-4000Å).

²As we will discuss below in the context of the Wilsonian semi-RG that applies to special cases in particle physics, integrating out high-energy degrees of freedom also counts as a form of coarse-graining, as it instantiates another iterative procedure for systematically introducing new field variables in a way that excludes fine-grained information about the original variables from a given theory or model.

The full-RG has a historical start in particle physics, in the context of renormalization and scaling. For example, Gell-Mann and Low (1954) considered the scaling of the electric charge in QED as a function of the energy scale at which the electron is probed.³ While (as J. D. Fraser (2021) points out) it would be a mistake to read the contemporary physical interpretation of such equations into Gell-Mann and Low (1954) (e.g., the concept of a renormalization scale), their scaling transformation is written as a differential equation, and is reversible. Bogoliubov and Shirkov picked up on Gell-Mann and Low’s formulation and further developed it. J. D. Fraser (2021) emphasizes the diversity in formal structures and physical interpretations that fall into the category of “the” renormalization group in this formative period. The central goal of this paper is to investigate the continued diversity of formalism and physical interpretation in the subsequent development of “the” RG.

Wilson took inspiration from both condensed matter and particle physics in constructing the modern Wilsonian RG.⁴ The “Wilsonian RG” is commonly taken to be a semi-RG transformation describing the scaling properties of a field system. Formally, the averaging operation is represented by a semi-group of transformations because it is irreversible. In condensed matter physics applications of RG methods, the semi-RG is naturally interpreted as a coarse-graining operation. Clearly, mathematical formalisms do not wear their interpretations on their sleeves. However, the full-group variant of the RG transformation is reversible, and thus does not count as a coarse-graining operation. As we shall see, the full-RG is widely applicable in particle physics. This raises a question: Is the full-RG sufficient for the applications and explanatory purposes to which the RG is put in particle physics, or is the semi-RG also necessary? In Sec. 3 and 4 we will argue that the full-RG is sufficient for all purposes in particle physics, and that it also has virtues that give it an advantage over the semi-RG. To make this case, we will carefully examine the applications and explanations afforded by the EFT perspective on theories in particle physics.

Back in 1975, Wilson’s impression was that the disparate nature of RG methods was more obvious than their common features. He hypothesized that “the very general nature of the renormalization group has been less apparent than the general nature of the derivative” because “the problems that one studies with the renormalization group are rarely formulated explicitly in terms of continuum limits” (p. 774). He also notes that “[t]he renormalization group is at a much more primitive stage than the derivative” because “[t]here is only a small subset of problems involving the statistical continuum limit that have been solved so far, and to solve these problems a large amount of labor and theoretical artifice is required” (p. 774). For example, Wilson notes that his solution of the Kondo problem in condensed matter physics involves “many special tricks which help to make the calculation practical” (p. 777). Today, after decades of research in which a wide range of problems in physics and other fields have been successfully solved using RG methods, it might seem that we should have also discovered a common formalism underlying these applications. However, RG methods have continued to resist a unified general formal presentation with respect to which particular techniques are instantiations. Jona-Lasinio’s attempts to find a more general framework for RG methods illustrate some of the difficulties. Having published (with DiCastro) a precursor to Wilson’s renormalization group equations in 1969, Jona-Lasinio was well aware of the variety of mathematical formulations of RG methods and also that the relationships between the different formal presentations were not clear. Motivated by the desire to find the “mathematically most faithful implementation” of Kadanoff’s idea that statistical mechanical systems have self-similar descriptions near a critical point, Jona-Lasinio (2001) uses a generalization of probability theory to underpin a more general framework for RG methods. This framework does illuminate the relationship between multiplicative formulations of the RG and Kadanoff’s idea (Jona-Lasinio 2010). However, the framework has only been extended as far as Gaussian fields (i.e., free fields), and Jona-Lasinio believes that extension beyond Gaussian fields would be very difficult (Jona-Lasinio 2001, Sec. 3.3). Furthermore, this generalized framework is designed to capture the physical ideas underlying the application of RG methods in condensed matter physics, which do not necessarily carry over to other applications of RG methods. In any case, in the absence of a general, unifying framework for RG methods, the best that we can do is to recognize that in practice RG methods are a cluster of mathematical techniques, and to do our best to understand how and why particular formal techniques are used in particular

³Cf. Koberinski (2021a, Sec. 2.2, 3.2) for a discussion of the conceptual move from the Gell-Mann and Low formalism to the modern Wilsonian RG, and J. D. Fraser (2021) for a broader account of the parallel developments of the Gell-Mann Low scaling and the Stueckelberg renormalization group in the 1950s.

⁴As Rivat (2021) argues, the analogies with condensed matter physics form only a part of the story for Wilson’s work leading to the renormalization group. His earlier work in meson physics also played an important role.

contexts. We turn our attention to an important context for particle physics: the construction of effective field theories.

3 Constructing effective field theories

On standard presentations of the construction of EFTs, semi-RG methods are often used to provide a high-level justification of the effectiveness of top-down constructions. (Recall that “top-down” refers to starting with high energies; “bottom-up” constructions begin with low energies.) One can also construct new EFTs from the bottom-up, where the role of the RG is indirect. Beyond the construction of EFTs, RG methods also allow one to explore the structure of either particular EFTs or the relationship between different EFTs via transformations in theory space. In this section we start by highlighting two different construction procedures for EFTs: the top-down and bottom-up approaches. Within the top-down approach, there is an abstract characterization in terms of semi-RG transformations between Wilson actions, and a concrete characterization where matching conditions are used between an EFT and its successor theory. In both cases, one thinks of the EFT as explicitly defined with reference to a more complete theory. After discussing the top-down approach (Sec. 3.1), we turn in Sec. 3.2 to the bottom-up EFT construction procedure to better understand how the Standard Model can be treated as an EFT. Constructing an EFT from the bottom-up contrasts with the top-down approach in that there is no longer explicit reference to any successor theory. We therefore find the bottom-up approach better suited to cases where a successor theory is unknown, or where the abstract notion of “integrating out” high-energy degrees of freedom from the successor does not apply.

There are three main takeaways from this section. First, the top-down construction procedure is not universally applicable in practice: the abstract construction is severely limited, while the concrete construction still misses important cases. Second, the concrete top-down procedure requires one to know what sort of low energy EFT they want to end up with; we don’t just start from the high-energy theory and arrive at a unique EFT. This complicates the abstract picture of *deriving* the EFT from its successor. Like many interesting cases in science, the concrete top-down perspective is actually a combination of top-down and bottom-up construction, meeting in the middle at an EFT with fixed couplings. Importantly for the discussion of uses of EFTs in Sec. 4, concrete top-down constructions do not utilize the semi-RG. Finally, the bottom-up procedure, in which one constructs an EFT independently from its relationship to a high-energy successor, is the most general, modern, and widely applicable perspective. This bottom-up construction does not lend itself to an interpretation in terms of coarse-graining transformations across theory space. The bottom-up perspective is also not dependent on the assumption that the top-down perspective applies. Instead, as we will argue in Sec. 4, the full-RG provides the full set of tools needed to explore the structure of EFTs.

3.1 Top-down constructions: EFTs as relational

From the top-down perspective, we construct EFTs in relation to some high-energy successor at energy scale Λ , and use the RG to compare the EFT to the low-energy limit of the successor at some energy scale $\mathcal{E} \ll \Lambda$. The abstract form of the top-down construction procedure is currently widely used to motivate a strategy for interpreting our current best theories—i.e., the Standard Model—as EFTs (cf. Weinberg 1979; Williams 2015; Wallace 2022). We start by outlining the abstract version of the top-down approach. This uses the Wilsonian semi-RG, which one can think of as successive iterations of “integrating-out” the high-energy modes of the original theory to arrive at a low-energy EFT. We then note that, in practice, the Wilsonian picture is rarely realized. In most concrete applications of top-down EFT construction, one uses dimensional regularization and a matching procedure for S-matrix elements. When used at all, one does not repeatedly apply the RG transformation as a coarse-graining procedure, as this is computationally intractable.

The Wilsonian semi-RG is a transformation from a generating functional $Z[\phi_i]$ for a local QFT with fields $\{\phi_i\} = \{\phi_1, \phi_2, \dots, \phi_n\}$

$$Z[\phi_i] = \int \mathcal{D}\phi_i \exp[S[\phi_i]] = \int \mathcal{D}\phi_i \exp \left[\int d^d x \mathcal{L}[\phi_i] \right], \quad (1)$$

to another with field modes of momentum less than some separation scale Λ “integrated out”:

$$\mathcal{Z}[\phi_i^l, \Lambda] = \int \mathcal{D}\phi_i^l \int \mathcal{D}\phi_i^h \exp [S[\phi_i^l, \phi_i^h]] \quad (2)$$

$$= \int \mathcal{D}\phi_i^l \exp [S'[\phi_i^l, \Lambda]], \quad (3)$$

where the field modes are split into high-energy (ϕ_i^h) and low-energy (ϕ_i^l) relative to the Λ by the momentum space relationship $p^2 + m^2 > \Lambda$ and $p^2 + m^2 < \Lambda$, respectively. Here $S[\phi]$ is the (classical) action for the field ϕ , \mathcal{L} the Lagrangian, and $\mathcal{D}\phi_i$ the measure for variations of all possible field modes. When we integrate out the high-energy modes, we are left with a generating functional containing an action S of similar form, though the dependence on high-energy modes has been replaced by the possible addition of new local interaction terms containing the $\{\phi_i^l\}$ and Λ . This transformation can be iterated, by lowering the separation scale $\Lambda' < \Lambda$ and integrating out the newly defined high-energy modes with $\Lambda' < p^2 + m^2 < \Lambda$.⁵ If we take the limit of infinitesimal decreases in the cutoff, i.e. $\Lambda' = \Lambda - \delta\Lambda$, we arrive at a semi-group transformation on the actions. It is a semi-group since the information regarding high-energy modes is irreversibly encoded in the low-energy dependence on Λ , changes in the coupling constants, and the addition of new local interaction terms. Once we have moved from $\Lambda \rightarrow \Lambda - \delta\Lambda$, we cannot go back. By Fourier transforming from energy-momentum to position variables, the semi-RG is formally analogous to the coarse-graining block-spin transformations introduced by Kadanoff for Ising models of Ferromagnets. Like the block-spin transformations, one can go from $\Lambda \rightarrow \Lambda'$, but not back again.

In practice, this sort of top-down construction is rarely used, as the process of integrating out high-energy fields at the level of the generating functional is intractable for most realistic QFTs. One notable exception is the ϕ^4 theory, for which this process has been carried out explicitly (Wilson and Kogut 1974; Polchinski 1984). The aim of this abstract top-down EFT construction is to illustrate the way in which information about the high-energy degrees of freedom is lost in moving from the full theory to the EFT. In practice, top-down constructions are typically far more subtle, and involve recognizing the (full or approximate) symmetries and degrees of freedom relevant to the low-energy theory, and where possible using the high-energy theory to construct matching conditions (cf. Burgess 2007; Manohar 2020). Semi-RG methods do not play a direct role in the construction of most empirically useful EFTs; they provide an in-principle understanding of how high-energy modes could possibly be removed from one theory to arrive at the EFT. If we compare the concrete top-down approach with the abstract one, the EFT is constructed with one major coarse-graining: the single transformation of fields from the high- to low-energy theory. But this is not accomplished using any form of the RG. Instead, one constructs the EFT by focusing on the low-energy degrees of freedom relevant to that domain, and neglects the contributions of high-energy fields. We will illustrate different concrete approaches with two well-known examples: Fermi theory and chiral perturbation theory.

Though the Fermi theory was historically constructed before the Standard Model, one can proceed top-down by starting with the Standard Model, constructing the Fermi theory, and using perturbative—i.e., order-by order—matching in the domain where $\mathcal{E}/M_{W,Z} \ll 1$ (cf. Manohar 2020, Sec. 4.8). In the limit of massless neutrinos and dealing with neutrino flavour eigenstates, the W boson mediates muon decay $\mu \rightarrow e + \nu_\mu + \bar{\nu}_e$ via the weak neutral current

$$j_W^\alpha = (\bar{\nu}_\mu \gamma^\alpha P_L \mu) + (\bar{e} \gamma^\alpha P_L \nu_e), \quad (4)$$

where ν_μ, ν_e are the muon and electron neutrinos, μ is the muon and e the electron, γ^α are the Dirac matrices, and $P_L = (\mathbf{1} - \gamma^5)/2$ is the left-handed projection operator. At the tree level, the amplitude for this decay is

$$\mathcal{A} = \left(\frac{-ig}{\sqrt{2}} \right)^2 (\bar{\nu}_\mu \gamma^\alpha P_L \mu) (\bar{e} \gamma^\beta P_L \nu_e) \left(\frac{-i\eta_{\alpha\beta}}{p^2 - M_W^2} \right), \quad (5)$$

where $g/\sqrt{2}$ is the W boson coupling. In the EFT regime, $\mathcal{E} \ll M_W$, so the W boson propagator can be

⁵Instead of starting with an action defined without a separation scale, the iterative application of the semi-RG can start with a theory which itself has a cutoff. If we think of our best QFTs as EFTs, then from the abstract top-down perspective one can arrive at a picture of a tower of EFTs for successively higher and higher energy scales (Cao and Schweber 1993; Bain 2013).

expanded to

$$\frac{-i\eta_{\alpha\beta}}{p^2 - M_W^2} = \frac{i\eta_{\alpha\beta}}{M_W^2} \left(1 + \frac{p^2}{M_W^2} + \frac{p^4}{M_W^4} + \dots \right). \quad (6)$$

If we retain only the first order term, the amplitude is

$$\mathcal{A} \approx \frac{i\eta_{\alpha\beta}}{M_W^2} \left(\frac{-ig}{\sqrt{2}} \right)^2 (\bar{\nu}_\mu \gamma^\alpha P_L \mu) (\bar{e} \gamma^\beta P_L \nu_e), \quad (7)$$

which is the same result one would obtain if we started with the Fermi Lagrangian

$$\mathcal{L} = -\frac{4G_F}{\sqrt{2}} (\bar{\nu}_\mu \gamma^\alpha P_L \mu) (\bar{e} \gamma^\beta P_L \nu_e), \quad (8)$$

provided that the matching condition $G_F/2 = g^2/(8M_W)$ holds. Thus the Fermi EFT, containing four-Fermion operators of mass dimension six, reproduces the low-energy Standard Model effects to first order. To higher orders in \mathcal{E}/M_W , one may need more complicated matching conditions, and additional higher dimensional terms in the Lagrangian. In general, an effective Lagrangian should contain an infinite set of local interaction terms, constrained only by the chosen degrees of freedom and their symmetries. The values of the couplings will in general depend on the energy scale at which they are probed, as well as the order at which matching conditions are set. The effective Lagrangian can then be truncated after all terms of dimension N are included, where N is the minimum dimension needed to ensure matching up to n^{th} order in the mass-scale expansion, within the specified error bounds.

It should be clear from the above example that practical top-down constructions, even when matching conditions can be determined exactly, often look quite different from the Wilson action coarse-graining procedure. The Fermi EFT construction can be thought of as a one-time coarse-graining; by constructing an effective Lagrangian that only contains fermion interactions, we have “integrated out” the effects of the W and Z bosons. But in other examples, it’s not clear that the EFT should be thought of in terms of integrating out the high-energy degrees of freedom at all. Take the example of chiral perturbation theory as a low-energy EFT of quantum chromodynamics (QCD). Here the high-energy theory is QCD, which describes the strong interactions between quarks and gluons. Chiral perturbation theory is a low-energy effective theory describing bound quark-gluon states—nucleons and pions—and their interactions. Like the Fermi theory, chiral perturbation theory was first developed by Nambu and Jona-Lasinio (1961) and refined through the 1960s, while QCD wasn’t developed until the early 1970s. But given that we now know the high-energy theory, we can try to reconstruct the EFT in a top-down manner. However, even though the high-energy theory is known, one cannot analytically compute matching conditions with the EFT. This is because, at low energies, QCD displays strong coupling, and perturbation theory is invalid. Further, one cannot think of chiral perturbation theory as the EFT resulting from integrating out the high-energy degrees of freedom of QCD. First, the two theories are written in terms of different fields. While the QCD Lagrangian contains quark and gluon fields, chiral perturbation theory contains meson and baryon fields. Second, the lightest fields in chiral perturbation theory are much heavier than the lightest fields in QCD. Gluons are presumed to be massless, while the up and down quarks have masses of $\mathcal{O}(\text{MeV})$. Meanwhile, pions have mass 139MeV, while the proton and neutron are $\mathcal{O}(\text{GeV})$ (Scherer 2003). There is no straightforward way in which one can integrate out the fundamental quark fields to arrive at nucleons and pions. Nor can one interpret the effective theory as a coarse-graining of QCD, since the simple idea of baryons being composed of constituent quarks is too simplistic a picture. Quark composites are constantly changing, and the effects of the gluons dominate the effects of individual quarks. Since the relationship between the two domains is complicated, matching in regimes where QCD and chiral perturbation theory both apply is done strictly numerically. This has given insight into the relationship between the strong force and its residual effects on nucleons.

In sum: in special cases, the top-down semi-RG construction of an EFT can be carried out. The concrete top-down construction is applicable to more cases, but does not use the semi-RG. In general, it is often not possible to think of EFTs as constructed using the top-down semi-RG, and so we are motivated to consider a different approach to understanding EFTs. We turn to this approach next.

3.2 Bottom-up construction: EFTs as standalone theories

EFTs like chiral perturbation theory, that are relatively autonomous from their high-energy counterparts, suggest a different way to construct and understand EFTs. Though not constructed from another theory by “integrating out” high-energy modes, these still fit the earlier definition of an EFT: they include an explicit separation of energy scales Λ —of which the EFT is only concerned with low energies—and a collection of local fields with interactions indexed by energy scale \mathcal{E} . In cases where the corresponding high-energy theory isn’t known, the exact domain of applicability is unfixed for the EFT, though the presence of a separating scale is still important for understanding the theory as an EFT. We note the shift in focus away from a relational understanding of EFTs to treating them as independent theories from the bottom-up perspective.

To construct an EFT from the bottom-up, one must write down the most general possible Lagrangian representing local interactions between fields describing the effective degrees of freedom, consistent with the assumed symmetries. This Lagrangian will contain an infinite number of terms, and must be truncated at finite order to be useful for predictions. Except where symmetries or field redefinitions allow, even terms with a coupling at energy \mathcal{E} of $g_i(\mathcal{E}) = 0$ must be included, as the full-RG will in general act to transform g_i if $\mathcal{E} \partial/\partial \mathcal{E} g_i \neq 0$. Taking the lesson from the top-down constructions, where an energy scale from the high-energy theory comes into play, we introduce a separation scale into the Lagrangian in the local coupling terms. The separation scale Λ is used to ensure that couplings for each interaction term have the correct dimension to fit into the Lagrangian. In four spacetime dimensions the Lagrangian must have mass dimension 4, which we denote as $[\mathcal{L}] = 4$. The general form of an effective Lagrangian is

$$\mathcal{L}[\{\phi_j(x)\}] = \sum_i g_i O_i[\{\phi_j(x)\}], \quad (9)$$

where the $\{g_i\}$ are coupling constants determining the strength of the local Lorentz invariant interaction operators $O_i[\{\phi_j(x)\}]$, constructed as polynomials of the sets of fields $\{\phi_j(x)\}$ and their derivatives. If we write the dimension of the local interaction operators as D , then the coefficients $\{g_i\}$ must have dimension $4 - D$. By mass dimensional analysis we can group together all operators with the same mass dimension D , and rewrite the coupling constants as $g_i = \alpha_i/\Lambda^{(D-4)}$.⁶ Relevant terms—those interaction terms whose magnitude grows under full-RG flow down to lower energies—are those terms with mass dimension < 4 . Irrelevant terms become negligible at low energies, and these are terms with mass dimension > 4 . Marginal terms do not vary under RG flow, and have mass dimension 4. The EFT Lagrangian can then be written as a sum over mass dimensions

$$\mathcal{L}_{\text{EFT}} = \sum_{D \geq 0} \frac{\mathcal{L}_D}{\Lambda^{D-4}} = \sum_{D \geq 0} \sum_i \frac{\alpha_i O_i^{(D)}}{\Lambda^{D-4}}, \quad (10)$$

where $O_i^{(D)}$ are the possible local operators of dimension D . In effect, the EFT Lagrangian is an expansion in negative powers of Λ . The order at which one truncates the expansion depends on the degree of precision required of the EFT. It is a proven feature of local QFTs that the contribution to a given scattering amplitude (normalized to be dimensionless) from an insertion of higher dimension operators into *any* graph leads to an additional contribution

$$\mathcal{A} \sim \left(\frac{\mathcal{E}}{\Lambda}\right)^{\sum_i (D_i - 4)}, \quad (11)$$

where the sum is over all inserted operators, and \mathcal{E} is the characteristic energy/momentum scale for the external particles. This is the power-counting formula (cf. Manohar 2020, Sec. 2.4), which tells one how to organize calculations in the EFT. For corrections to the renormalizable part of \mathcal{L}_{EFT} of order $(\mathcal{E}/\Lambda)^n$, one computes graphs with a single insertion of \mathcal{L}_{n+4} , two insertions of \mathcal{L}_{n+3} , three insertions of \mathcal{L}_{n+2} , and so on until one gets to \mathcal{L}_5 .

⁶Physicists typically expect, for reasons of naturalness, that the dimensionless constants α_i are of $\mathcal{O}(1)$. For terms suppressed by powers of Λ^n , $n > 1$, this assumption can be relaxed to a certain degree, so long as $\alpha_i/\Lambda^n \ll 1$. For operators relevant at low energies, the assumption of naturalness is harder to justify, and the Higgs mass and cosmological constant problems highlight the apparent failure of naturalness (cf. Williams (2015) and Rosaler and Harlander (2019) for discussion of the Higgs mass, and Schneider (2020a), Schneider (2020b), Koberinski (2021b), Koberinski (2021c), and Koberinski and Smeenk (2022) for philosophical discussion of the cosmological constant problem).

The key difference between EFTs and renormalizable QFTs is that terms of mass dimension > 4 require counterterms of higher and higher dimension in order to generate finite scattering amplitudes, while terms of mass dimension ≤ 4 only generate new counterterms of mass dimension ≤ 4 . Renormalizability is the requirement that only a finite set of counterterms are needed to make all graphs finite; this is only possible in the case where no terms from $\mathcal{L}_{\geq 5}$ appear in the theory. In cases where one is interested in corrections up to some finite power of $\sum_i (D_i - 4)$ in Equation (11), only a finite number of operators are needed and the bottom-up EFT is predictive once all free couplings are fixed at some scale \mathcal{E} . We can think of all truncated EFTs—including the standard renormalizable QFTs—as the first few terms in a general EFT containing an infinite number of coupling terms. This is the sense in which one may think of the Standard Model as an EFT; implicit is a separation scale Λ and an infinite set of interaction terms compatible with the symmetries of the quarks, leptons, and gauge bosons.

RG methods do not enter at the construction stage for bottom-up EFTs. However, once an EFT has been constructed in this manner, the RG equations are used for examining the scaling behaviour of the theory. RG methods also underwrite the justification for constructing an EFT in this bottom-up fashion. The requirement that the theory’s scattering amplitudes be independent of the arbitrary energy scale \mathcal{E} leads to the full-RG equations describing the running of coupling constants with respect to \mathcal{E} . The semi-RG is not used here because the objective is the invariance of the scattering amplitudes with respect to \mathcal{E} . There is no sense in which one is coarse-graining: one is not averaging, integrating out, nor using any other systematic means of iteratively replacing the fine-grained variables of the model to obtain the EFT. The bottom-up construction procedure leads naturally to thinking of EFTs as full-fledged theories in their own right, albeit ones with a built-in separation of scales. This leads to the definition of EFTs we outlined at the start of Sec. 3.

4 The structure of effective field theories

So far we have made a distinction between the full- and semi-RG, and argued that the semi-RG only plays a role in the abstract top-down construction of EFTs. As a result, the notion of coarse-graining—an iterative procedure in which one set field variables is replaced by another set of field variables by “integrating out” high energy fields and this results in the loss of fine-detail information about the high energy fields—only applies in abstract top-down constructions. We turn now to explore the applications of the RG for elucidating structural relations among EFTs. RG methods—in particular the full-RG—are essential for the application of EFTs in particle physics. After surveying some of the uses of RG methods in exploring the structure of EFTs, we conclude that the bottom-up approach employing the full-RG is sufficient for all major explanatory purposes. While a top-down point of view can be useful, the semi-RG can always be replaced with the full-RG without loss of explanatory power. The converse is not true; simple scaling properties of a given EFT require the full-RG. In particular, it is never necessary to think of the RG transformations in particle physics as an iterative coarse-graining procedure. There are conceptual distinctions in understanding EFTs via the full- and semi-RG, but we will see that for most practical uses the two perspectives yield very similar formal relationships among EFTs (e.g., the distinction between relevant and irrelevant terms). We argue that these formal relationships provide a basis for answering the structural and explanatory questions that arise in particle physics.

For a given EFT, there are a number of questions one might want to ask about its structure, for which RG methods can be useful. We start with a focus on more empirical explanations one might seek from an individual EFT, and then generalize to global features of the space of EFTs. The local questions are germane to particle physics because they involve calculational or empirical tractability. This is because an important purpose of RG methods is renormalization. The global questions also have an empirical focus: structural properties of EFTs are used to explain the empirical success of the Standard Model and the success of renormalizability as a criterion for constructing an empirically successful model.

One important distinction is that of relevant versus irrelevant operators under RG flow. In four spacetime dimensions, the couplings for terms in the Lagrangian with mass dimension > 4 are *irrelevant* under RG flow at low energies. This means that the strength of the coupling diminishes as one probes the system at lower and lower energies. By contrast, *relevant* terms have couplings that increase at low energies. *Marginal* terms are those whose couplings do not change under RG flow. A coupling that is marginal to

first order in perturbation theory may be either relevant or irrelevant at higher-orders; these are termed *marginally relevant* or *marginally irrelevant*, respectively (cf. Williams (2022) for further discussion of these distinctions in QFT). Importantly, these distinctions apply to either full- or semi-RG transformations. The irrelevance of nonrenormalizable terms at low energies plays a key explanatory role in the understanding of EFTs. Neglecting terms that make a negligible contribution to the calculation of a quantity of interest is a standard approximation procedure in physics. Neglecting irrelevant terms at low energies is an application of this standard approximation procedure. In contrast, coarse-graining is a more specialized procedure that involves iterative replacement of one set of fine-grained (or high energy) variables by another set of coarser-grained (or lower energy) variables and results in the loss of fine-grained (or high energy) information. While coarse-graining is one type of approximation, it is important to recognize that coarse-graining is not the type of approximation procedure used when applying the full-RG. Figure 1 provides a simplified picture of RG flow in the (a) semi-RG and (b) full-RG pictures. This will be especially useful for the discussion in Sec. 4.2.

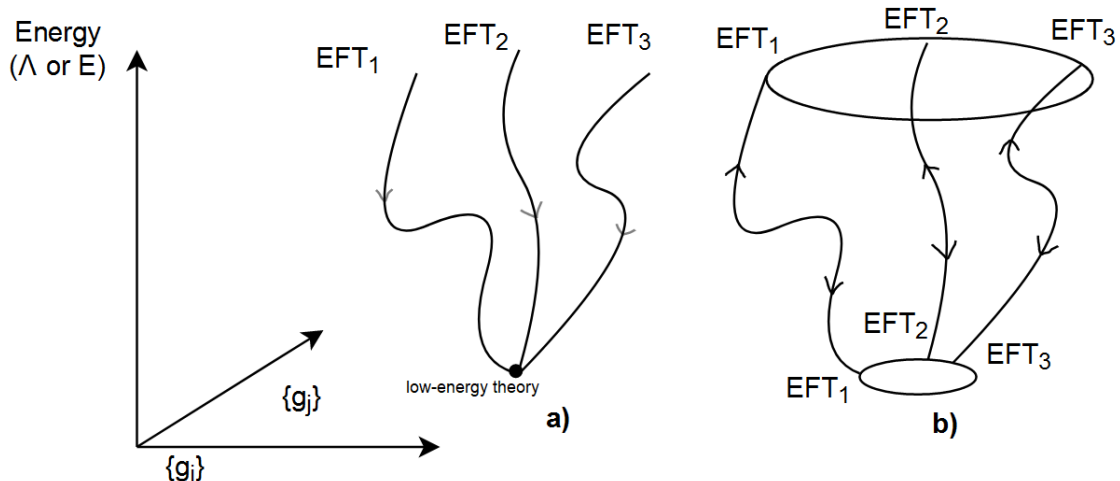


Figure 1: Renormalization group flow in EFT space in the semi-RG (a) and the full-RG (b) picture. Note that the infinite dimensional coupling space has been suppressed to two dimensions. 1a) For the semi-RG (top-down) perspective, we start with a class of different EFTs specified by values of their couplings at a reference energy. The application of the semi-RG induces a directed flow down to the renormalizable subspace (here a point). With each application of the semi-RG, we transform from one high-energy EFT down to a different, lower-energy EFT. As an example, the point could represent the Standard Model, and each high-energy EFT a different possible SMEFT. 1b) For the full-RG (bottom-up) perspective, we start at low energies with known couplings and known bounds on the first few nonrenormalizable couplings (plus uncertainties on both). This defines a fuzzy, higher-dimensional subregion of theory space, within which all possible EFT completions lie. Using the full-RG, one can reversibly flow up to higher or down to lower energies. We then see that what starts out as a larger subsurface at high energies shrinks to a surface described by the initial low-energy EFT (including uncertainties). Within this region, we can idealize a single trajectory line as a precisely specified EFT. Again, we can think of the Standard Model and various possible SMEFTs as a concrete example.

4.1 Empirical questions about local structure

First, one might wonder whether perturbative methods (i.e., expansion in powers of the coupling strength⁷) are valid for a given EFT, and in what domains one should expect them to work. Similarly, one may also wonder how values assigned to couplings in the EFT change with changing energy scale. These questions are answered in the same way. Starting with a given EFT, we can examine the scaling behaviour of the coupling constants using the full-RG to determine the domains in which the values of the couplings are small enough

⁷Note that this is a different expansion than the EFT expansion in powers of \mathcal{E}/Λ . For large values of coupling constants, standard perturbative QFT methods break down, though one may still be able to organize EFT terms in the \mathcal{E}/Λ expansion.

for perturbative QFT methods to apply. From a given reference scale, the full-RG allows one to explore energies higher and lower than that scale. Here, since we are taking an initial EFT as given, the bottom-up understanding is the most applicable. We have an EFT with a separation scale Λ (not necessarily precisely known, but fixed at some relatively high value), whose couplings have been renormalized and contain a dependence on a second lower energy scale \mathcal{E} . The action of the full-RG is on \mathcal{E} , where the imposition of invariant n -point functions under changes in \mathcal{E} results in β -functions determining the scaling behaviour of couplings for the theory.⁸ Standard perturbative QFT is applicable when couplings are less than one, and breaks down otherwise. For theories like QED, perturbative methods are valid at low energies, but break down at higher energies, while QCD has the opposite behaviour, such that perturbative QCD works at high energies but not at low energies.

Note that the semi-RG cannot be used to understand scaling of couplings in a reversible manner. For these local questions, we start with an EFT given by specifying the couplings at some energy scale, and explore the scaling of couplings. Since β -functions are meant to determine scaling with respect to both increase and decrease of energy, the action of the full-RG is needed. In particular, the semi-RG cannot be used to determine the scaling behaviour, as coarse-graining is a unidirectional operation from high to lower energies.

Now that we understand how the full-RG acts on a given EFT, we could ask whether that theory is well-defined to all energy scales, that is, whether its n -point functions are finite at all energies. For an EFT constructed with a built-in separation scale Λ , the answer to that question is always no. All terms with mass dimension > 4 (in 4D spacetimes) contribute factors of $(\mathcal{E}/\Lambda)^n$, $n > 0$ to any n -point functions, with higher powers of n coming from terms with higher mass dimension. An EFT can only give sensible predictions when $\mathcal{E} < \Lambda$, since one can then truncate the interaction terms when $(\mathcal{E}/\Lambda)^N$ is negligible for some $n = N$. At the point where $\mathcal{E} \geq \Lambda$ the EFT becomes ill-defined, since an expansion in \mathcal{E}/Λ no longer converges. Though we might not know what the exact value of Λ is, we can say with certainty that the EFT is not well-defined to all energy scales.

However, a more interesting question is whether a given renormalizable QFT (equivalently, the renormalizable sector of a given EFT) is well-defined to all energy scales. A renormalizable QFT can be written without an explicit dependence on the separation scale Λ , and such a theory is predictive with only a small finite number of coefficients fixed by experiment. Here we use the same independent EFT approach, and employ the full-RG to determine scaling behaviour in the asymptotic limits of $\mathcal{E} \rightarrow 0$ and $\mathcal{E} \rightarrow \infty$. A theory is said to be UV (resp., IR) safe if the β -functions for the set of couplings $\mathbf{g} = (g_1, g_2, \dots, g_n)$ all reach zero at some finite fixed point \mathbf{g}_* as $\mathcal{E} \rightarrow \infty$ (resp. $\mathcal{E} \rightarrow 0$), meaning that $\partial \mathbf{g} / \partial \mathcal{E} |_{\mathbf{g}=\mathbf{g}_*} = 0$. The special case where the couplings $\mathbf{g}_* = \mathbf{0}$ as $\mathcal{E} \rightarrow \infty$ is called asymptotic freedom. Theories that are UV and IR safe are well-defined to all energy scales, and therefore one can use them to calculate n -point functions up to arbitrarily high or low energies. The caveat is that perturbative methods may break down even if the couplings remain finite, so while the tools used to determine UV/IR safety are the same as those used to determine where perturbative methods are applicable, the domains of applicability may come apart. It is a separate question whether a given EFT is empirically adequate to all energy scales, though this question can of course only be answered by comparing its predictions to experiment.

4.2 Explanatory questions about global structure

Moving beyond particular features of a single EFT, we can use RG methods to ask more general questions about the relationship between a renormalizable QFT and its EFT generalization, or even about the structure of the space of possible local Lagrangian field theories. We highlight three major global explananda in particle physics that are explained by appeal to RG methods. All three explananda can be given similar explanations, but the phenomena to be explained are distinct.

First, we deal with the concept of renormalizability. Above, we needed to restrict ourselves to the renormalizable sector of an EFT to meaningfully inquire if that theory is well-defined to all energy scales. Why is this permissible? By definition, a renormalizable theory is one that can be written without an explicit

⁸In practice, one cannot typically determine β -functions analytically, and they must instead be determined in a local neighbourhood of a well-behaved sector of the theory. This is typically a region where the couplings are small or zero, or involves studying asymptotic scaling behaviour. So in practice, the explanation we get from RG methods is more about the ways we *should expect* scaling behaviour to change away from regions where perturbative methods are well-defined.

dependence on the separation scale Λ , and such a theory is predictive with only a small finite number of coefficients fixed by experiment. But what of the relationship between an EFT and its renormalizable sector? The focus on the renormalizable sector of a given EFT further requires that one renormalizes that portion of the theory, i.e., removes the explicit Λ dependence by subtracting counterterms from the regularized Lagrangian. Early in the development of the Standard Model, renormalizability was a necessary criterion for model construction, as non-renormalizable theories yielded divergent predictions (Koberinski 2021a, Sec. 3.1). Later on, RG methods helped explain how renormalization works (Huggett and Weingard 1995), and the EFT framework now allows for nonrenormalizable terms in any effective theory (Weinberg 1979; Williams 2022).⁹ But given the EFT perspective, one might wonder why the requirement of renormalizability led to successful theories that continue to hold up today. From the methodological perspective, we might rephrase the explanandum as follows:

Explanandum #1: *Renormalizability was a successful methodological criterion for constructing the empirically successful Standard Model despite the fact that the modern EFT perspective does not privilege renormalizable Lagrangians.*

This can be explained in two slightly different ways, using either the full-RG or semi-RG. Both rely on insights gained after the fact via the EFT framework. Think generally about the structure of an EFT. In four spacetime dimensions, any terms in an effective Lagrangian with mass dimension > 4 will become irrelevant at lower and lower energies. In the limit that $\mathcal{E}/\Lambda \rightarrow 0$, only the marginal and relevant terms will contribute to the effective Lagrangian, and for interactions described by the Standard Model these are the renormalizable terms. In the regime that $\mathcal{E} \ll \Lambda$, we will therefore only see small corrections from the nonrenormalizable terms to predictions made using only the renormalizable sector. From our current perspective that the Standard Model is actually the first few terms of a Standard Model EFT (SMEFT); the restriction to focus on renormalizable theories was methodologically successful because these match the predictions of an EFT to a high level of precision—provided that one is probing physics at energy scales well below the separation scale Λ .

The next explanandum has a similar explanation, but focuses on the continued empirical success of the Standard Model:

Explanandum #2: *The Standard Model continues to boast precise empirical success today without taking into account any terms of an SMEFT (specified at the energy scale of empirical tests).*

Consider the Standard Model. As constructed in the 1960s and 1970s, the Standard Model is a renormalizable Yang-Mills type QFT, defined by the symmetry groups $SU(3) \times SU(2) \times U(1)$. It has served as the theoretical foundation of particle physics for decades, and boasts some of the most precisely confirmed predictions in the history of science. But if we think of all QFTs in particle physics as EFTs, then the Standard Model is actually just the common renormalizable sector of a more general class of SMEFTs. Why has the Standard Model been so empirically successful if it only contains the first few terms of some more general SMEFT? If a SMEFT is the better approximation to some future theory, why haven't we noticed deficiencies in the Standard Model now or earlier?

Again, either the semi-RG or full-RG can be used to supply an explanation. Start by thinking about the class of SMEFTs. Each contains an infinite number of coupling terms consistent with the fields and symmetries included in the Standard Model, and each member of this class is distinguished by different values of the couplings at a reference scale \mathcal{E}^* . Each also includes a separation scale Λ . If we relate each of these SMEFTs to a successor theory, we can think of each SMEFT at \mathcal{E}^* as giving a competing effective, low-energy account of the physics of the future theory. Then the upper bound for the separation scale is the mass/energy at which qualitatively new physics arises in the successor theory, Λ_{new} . We must assume that either a) the successor theory *just is* another local Lagrangian field theory, or b) that the successor is well-approximated by a SMEFT at energies $< \Lambda_{new}$.

The top-down, semi-RG explanation for the empirical success of the Standard Model can be summarized as follows. Suppose a successor theory is well-approximated by some SMEFT at \mathcal{E}^* —for example, in the

⁹The details of renormalization are outside the scope of this paper, but see Huggett (2002), Butterfield and Bouatta (2016), and Rivat (2019) for some contrasting philosophical accounts.

sense that one obtains that SMEFT by “integrating out” new physics at some separation scale Λ_{new} .¹⁰ Then we must assume that the energy scales we are concerned with in the Standard Model are far smaller than the energy scale at which new physics comes in, such that the Standard Model applies at energies $\ll \Lambda_{new}$. As we move from the SMEFT defined with a separation scale of Λ_{new} under a semi-RG flow down to new theories with lower separation scales, these new effective theories look increasingly like the Standard Model, since high mass dimension terms are irrelevant and get discarded. The notion of equivalence here is at the level of Lagrangians. An EFT with nonrenormalizable terms and a separation scale $\Lambda_{new} \gg \Lambda_{SM}$ will be equivalent to the SM after iterative coarse-graining if $\mathcal{L}_\Lambda = \mathcal{L}_{SM} + \mathcal{O}(\frac{\Lambda_{SM}^2}{\Lambda_{new}^2})$, since coarse-graining involves throwing out the additional terms. Terms in the SMEFT with mass dimension $\gg 4$ are coarse-grained away quickly, and in the limit of $\Lambda_{SM}/\Lambda_{new} \rightarrow 0$ the *only* remaining terms are the renormalizable terms found in the Standard Model (Fig. 1(a)). While we know that $\Lambda_{SM}/\Lambda_{new} \neq 0$, the Standard Model is a better approximation to the successor EFTs the closer this ratio gets to 0. So, the Standard Model has been—and continues to be—so successful because it approximates the dominant effects of a variety of possible SMEFT extensions. The Standard Model subspace is an IR fixed point for these EFT extensions: the couplings transform under RG flow such that only Standard Model terms remain. Assuming that one of the SMEFTs is a good approximation to whatever the future theory is, the accuracy and precision with which one can use the Standard Model predictions in place of the EFT depends on the ratio of $\Lambda_{SM}/\Lambda_{new}$. Over the last fifty or so years, we have pushed the energy scale Λ_{SM} to higher and higher levels, and thus far the Standard Model has still been sufficient to account for empirical results, indicating that Λ_{new} must likewise be very large.

The top-down semi-RG explanation for the success of the Standard Model in light of the EFT framework is often taken to be one of the major philosophical insights gained from the application of RG methods to particle physics. However, an explanans for Explanandum #2 can be constructed making similar—though strictly weaker—assumptions in the context of the bottom-up approach to EFTs. The main difference is that one should think of the *full*-RG acting on \mathcal{E} as a transformation of parameters within a single EFT, instead of the semi-RG acting on Λ as a transformation *between* theories. As discussed in Sec. 3.2, the separation scale Λ in a bottom-up EFT is treated as an unknown, used to sort calculations into powers of \mathcal{E}/Λ . We can keep Λ fixed at some large but unknown value, and find diminishing contributions from the nonrenormalizable sector of the SMEFT as we move down to energy scales for which $\mathcal{E}/\Lambda \ll 1$.

The bottom-up full-RG explanation for the success of the Standard Model starts, like the top-down explanation, by assuming that some SMEFT is a good approximation to whatever theory succeeds the Standard Model (quantum gravity, grand unified theory, etc.) at low enough energies, though here the nature of the approximation remains open. In particular, we need not assume any fixed relationship between an SMEFT and a successor theory, such as that the SMEFT is obtained by “integrating out” the high energy degrees of freedom of the successor theory. The explanation for the success of the Standard Model is cast in terms of prediction. The Standard Model makes *the same* predictions—up to some pre-specified error bounds—as the SMEFT. As we move to lower energies under the (full-)RG flow, we can neglect more and more of the irrelevant terms in the SMEFT, and we move closer and closer to the Standard Model fixed point (Fig. 1(b)). For most current predictions, up to the possible sensitivity of experimental apparatus, we will find that one can truncate the SMEFT to include only terms of order $(\mathcal{E}/\Lambda)^0$ and still obtain the correct predictions. We can thus think of comparing the empirical equivalence of different degrees of expansion in powers of \mathcal{E}/Λ from the *same* theory, namely the SMEFT. In Fig. 1(b) this would entail picking a point in the low-energy region for which all nonrenormalizable terms are exactly 0. From this point of view, in using the Standard Model we have actually just been truncating a SMEFT to include only the renormalizable terms. Additionally, if and when discrepancies arise requiring new terms from the SMEFT, careful analysis of the magnitude of the discrepancy can provide indirect evidence for the separation scale Λ at which new physics arises. Given the high degree of precision to which some predictions in the Standard Model have been confirmed, one should expect the separation scale Λ to be rather high.

Instead of thinking of relationships between various different possible EFTs and the Standard Model, the bottom-up explanation takes the perspective that the Standard Model just is the first part of some SMEFT. Further, due to our epistemic position, we could be dealing with any of a class of possible SMEFTs, consistent

¹⁰Though not strictly required for the explanation of Explanandum #2, this stronger relationship between a SMEFT and the successor theory is essential for grounding a notion of theoretical equivalence that is stronger than simply empirical equivalence. We discuss this further in Sec. 4.3 below.

up to experimental uncertainty with the finite number of known or bounded couplings (see Fig. 1(b)). We get away with truncating the SMEFT at the renormalizable sector because, up to this point, the precision of experiment δ has been such that $\delta > \mathcal{O}(\mathcal{E}/\Lambda)$.¹¹ Thus all nonrenormalizable terms are negligible at energies probed thus far. Future results for which $\delta \approx \mathcal{O}(\mathcal{E}/\Lambda)$ would force us to include higher mass dimension terms in the SMEFT, and would provide indirect evidence for an upper bound on the domain of applicability of the SMEFT, since one could thereby place bounds on Λ . As we use the full-RG to take the class of compatible SMEFTs up to higher energies, the region compatible with low-energy empirical constraints takes up larger portions of the theory space, since terms irrelevant at low energies become larger at high energies. Thus, from the bottom-up perspective, we see that the known Standard Model terms are compatible with a wide range of EFT completions.

Both the semi-RG and full-RG explanations for Explanandum #2 assume that some SMEFT makes predictions in close agreement with the successor theory at energies $< \Lambda_{new}$. Further, both require that the energies probed in the best tests of the Standard Model are far lower than Λ_{new} , such that all nonrenormalizable terms are negligible. The semi-RG explanation further assumes that the successor can be transformed into the SMEFT via something like the Wilsonian action transformation, allowing one to integrate out degrees of freedom successively. Coarse-graining involves iteratively discarding high-energy terms, which results in a relationship between different EFTs. This raises the possibility that the semi-RG establishes a notion of equivalence between theories that is stronger than mere empirical equivalence. The full-RG explanation does not make the assumption about the applicability of coarse-graining, and leads to a notion of empirical equivalence between sectors of the same theory (conceived as an EFT).

We can also use RG methods to elucidate the relationship between an EFT and its possible successors:

Explanandum #3: *Low-energy EFTs make predictions that are empirically robust with respect to a set of possible variations in the high-energy successor.*

This explanandum is of great interest to both physicists pursuing a successor theory and philosophers who deploy the explanans to articulate either a variant of scientific realism (e.g., Williams 2019, J. D. Fraser 2018) or empiricism (e.g., Ruetsche 2020). In fact, Explanandum #3 is weaker than the desired explanandum; the stronger version of this explanandum invokes the renormalizability of the low-energy EFT (as in Explanandum #1) and includes a principled specification of a large class of variations in the high-energy theory that correspond to a given low-energy EFT. For our purposes in this paper, analysis of the weaker Explanandum #3 is sufficient to support our conclusion that either the semi-RG or the full-RG can be used to ground this type of explanation. Below we will also comment on the role of the full-RG in explanations of the stronger version of this explanandum.

The explanation of Explanandum #3 proceeds as follows in the context of the abstract Wilsonian action and semi-RG.¹² The semi-RG acting on Λ defines a flow on the space of possible local Lagrangian theories at high-energy Λ . We obtain an EFT by lowering Λ and considering physics at energy scales $\mathcal{E} \ll \Lambda$. We know that the resulting EFT is insensitive to changes in the high-energy theory when we can consider “nearby” theories in the space of local Lagrangians (e.g., those with different values of couplings, or additional higher-order terms). This will be the case when the variations in the high-energy Lagrangian affect only irrelevant terms (or, perhaps, marginal terms). These nearby theories, when acted upon with the semi-RG, flow down to the same space of EFTs, with the only dependence being confined to small differences in the coupling constants of the EFT. Note that this is a sort of counterfactual robustness: we say that a set of possible high-energy theories yields empirically equivalent EFTs at low energies as the RG flow converges down to

¹¹There are two ways to try to push the Standard Model such that $\delta \approx \mathcal{O}(\mathcal{E}/\Lambda)$. The first is to raise the energy scales at which we probe, i.e., increasing \mathcal{E} . This is done by using particle accelerators at very high energies, and has been the primary mode for testing the Standard Model. Unfortunately, however, experiments are generally less precise at higher energies, and an increase in \mathcal{E} also results in an increase in δ . The second way is to perform precision tests of the Standard Model, thereby lowering δ . The most precise tests to date are low-energy tests of the QED sector (cf. Koberinski and Smeenk 2020; Parker et al. 2018; Hanneke, Hoogerheide, and Gabrielse 2008 for recent measurements and discussion). This also involves lowering \mathcal{E} , but the increase in precision is often worth the cost. Current precision measurements of the fine-structure constant (Parker et al. 2018) even place bounds on the presence and strength of interaction of beyond Standard Model particles.

¹²When the concrete top-down construction procedure applies (i.e., we can construct an EFT by neglecting the heavy particles from a higher-energy theory and matching conditions are used to ensure n -point functions agree), a limited explanation is also available. The functional dependence of the EFT couplings on quantities from the high-energy theory can be calculated perturbatively via the matching conditions, which allows the effects of variations in high-energy couplings to be determined by substitution. However, the principled determination of the size of the class of variations discussed below relies on the RG.

a finite dimensional subspace, though at most one of these possible theories is a candidate description of physics in the actual world. The abstract top-down approach thus provides information about the robustness of EFTs against (perhaps unknown) details of the high-energy successor theory.

A similar story holds using the full-RG, from the bottom-up perspective. Though we cannot determine the explicit form of dependence of the low-energy couplings on those of the successor—from this perspective, the successor theory is unknown—we still fix the low-energy couplings via renormalization and measurement. We can show robustness of the expectation values in the low-energy EFT by considering neighbouring EFTs whose couplings differ from the EFT under consideration only for terms above some mass dimension d . The explanation of the robustness of the empirical predictions is that at low energies \mathcal{E} , these terms are irrelevant, meaning that initially distinct EFTs in theory space yield the same low-energy n -point functions, up to some small uncertainty.

Our main conclusion is that either the semi-RG or full-RG can be used to explain Explanandum #3. As noted above, Explanandum #3 can be strengthened by making the connection to renormalizability of the terms in the EFT at low energy scales, as in Explanandum #1, and supplying a principled specification of a suitably large set of possible variations in the high-energy theory. The formally analogous explanation in condensed matter physics appeals to the universality class of systems at small-distance scales to explain universal behaviour at large-distance scales. In particle physics, details about the particular class of EFTs under consideration will matter for establishing an explanation for the strengthened Explanandum #3. This is important, but the point that we want to draw attention to here is that the template for the general explanation can invoke either the abstract Wilsonian semi-group or the full-group. The stronger version of Explanandum #3 requires detailed study of a particular class of EFTs, and this may require selective application of either the semi-group or full-group variant. For example, Polchinski (1984) treats the ϕ^4 case by using the Wilsonian formalism, but with a full-group transformation acting on the separation scale Λ . By focusing on a specific class of theories with pre-specified symmetries and fields, one can more explicitly characterize the class of EFTs that flow down to the renormalizable subspace. This is a consideration that one could use as part of an argument for convincing oneself that the set of successor theories captured by this explanation is suitably large.

While detailed discussion of Polchinski (1984) is beyond the scope of this paper, it is interesting to note that the analysis uses the full-RG. Since this paper is regarded as supplying the most detailed proof of renormalizability and robustness under RG flow in a concrete case, this is another consideration in support of the full-RG being the more useful formalism in particle physics (cf. Ruetsche 2020, 301 fn8).

4.3 Analysis

To recap, in Sec. 3 we distinguished top-down and bottom-up construction procedures for EFTs in particle physics. The bottom-up, full-RG strategy is more widely applicable and is based on less substantial assumptions. Furthermore, in Sec. 4.1 we saw that the local structural properties of a given EFT are probed using the full-RG, not the semi-RG. But these would not be compelling rationales for basing the interpretation of EFTs in particle physics on the bottom-up, full-RG procedure if the top-down procedure were necessary for some uses of EFTs in particle physics. Is the abstract Wilsonian semi-RG and the associated interpretation as a coarse-graining transformation needed for the deeper explanations of global structural relations among EFTs? Sec. 4.2 argued that both the semi-RG and the full-RG can ground sufficient explanations for the empirical adequacy of the Standard Model and the successful use of renormalizability as a methodological criterion for its construction. More generally, either the full-RG or semi-RG supplies a basis for explaining the empirical robustness of an EFT with respect to possible successor theories. In this sub-section we will offer analysis of why the full-RG variant is sufficient for these explanatory purposes and argue that the full-RG actually supplies better explanations than the semi-RG. We will also assess how the full-RG informs the conception of an EFT in particle physics and compare our account of EFTs in particle physics to other recent philosophical accounts.

Why does the full-RG suffice for addressing the three explananda in Sec. 4.2—explaining the empirical adequacy of the Standard Model at low energy scales, the success of renormalizability as a criterion for constructing the Standard Model, and the empirical robustness of the EFTs (including the SMEFT) with respect to variations in the high-energy successor theory? One reason has to do with the nature of the explananda. What requires an explanation is the empirical success of the (renormalizable) Standard Model

(or other EFTs) at the relatively low energy scales that we have been able to probe. In particle physics, collider experiments typically test the predicted values of scattering amplitudes. In all of the explananda empirical content is interpreted in terms of scattering amplitudes. Either the full-RG or the semi-RG can be used to determine the scattering amplitudes at some scale \mathcal{E} to some degree of approximation. The second reason is that the ingredient that does most of the work in the explanation for this explanandum is the behaviours of relevant and irrelevant operators under an RG transformation. As long as the non-renormalizable terms at some high energy scale Λ are irrelevant with respect to the RG transformation, the only terms in the effective low-energy Lagrangian that will make non-negligible contributions to scattering amplitude calculations at low \mathcal{E} will be the renormalizable, Standard Model terms. Non-renormalizable terms are, formally, irrelevant with respect to both the full-RG and the abstract semi-RG flow. Irrelevant terms can be neglected at low energies as a standard matter of invoking an approximation procedure. The physical interpretation of an RG transformation as a coarse-graining operation that induces a loss of fine-grained information contained in the high energy description of the system is not needed to supply this explanation.

Thus, both the semi-RG and full-RG suffice for explaining global structural features of EFTs in particle physics. However, the analysis in this paper supports the stronger conclusion that the full-RG supplies a better framework than the semi-RG for understanding the role of EFTs in particle physics. One line of argument for this conclusion is that the full-RG is more widely applicable in particle physics than the semi-RG. The idea of successive coarse-graining underpinning the semi-RG is not readily applicable to many realistic EFTs—in particular those effective to QCD. In these cases we know the successor theory, but still find that the low-energy EFTs are not obtained by integrating out high-energy degrees of freedom, nor by any other form of successive coarse-graining. Thus the semi-RG plays no role in their construction. Using EFTs for making predictions requires use of the full-RG, but not the semi-RG. Sec. 4.1 outlined local explanations of the scaling behaviour of particular EFTs that depend on the full-RG.¹³

The second line of argument for the conclusion that the full-RG provides a better framework for the understanding of EFTs in particle physics is that the full-RG is better suited to our epistemic situation in particle physics than the semi-RG. The semi-RG must be coupled with a top-down perspective, and assumptions must be made about the high-energy theory to apply the semi-RG. Consider the Standard Model. Applying the top-down, semi-RG perspective, we get from some SMEFT specified just below energy Λ_{new} where one expects new physics to arise down to a low-energy renormalizable theory by arbitrarily picking values of couplings just below energies Λ_{new} . We have no justification for picking any particular values of the coupling constants at Λ_{new} . In contrast, the full-RG can be applied from the bottom up. The starting point is the assignment of empirically-well confirmed values for the couplings in the low-energy SMEFT—and uncertainties about the precise empirical values of the couplings can also be taken into account. Either the full-RG or the semi-RG can be used to show that the assumptions about the values of the irrelevant couplings in the SMEFT at Λ_{new} that are made in the course of applying the semi-RG are inconsequential to low-energy predictions. The difference is that the application of the full-RG does not require arbitrary posits about high-energy physics. Instead, the full-RG more transparently represents our current epistemic situation by starting with the empirically accessible low-energy theory and extrapolating to higher energies.

The full-RG supports a particular understanding of EFTs in particle physics. At the abstract level, an EFT is (incompletely) specified in terms of the fields chosen to represent relevant degrees of freedom, and values of couplings at some reference energy scale \mathcal{E}^* . The full-RG acting on couplings in the theory relates their values at other possible empirical reference energies \mathcal{E} , with the constraint that the n -point functions for the theory remain invariant under this scaling transformation. One can think of this as a type of invariance requirement on the theory; though the couplings vary with energy scale, they vary in such a way that the EFT yields invariant predictions. A separation scale Λ is included, with the interpretation that physics at or above that scale falls outside the domain of applicability of the EFT. From the perspective of the full-RG, Λ can remain relatively unfixed; the only constraint is that $\Lambda \gg \mathcal{E}$. In principle, one could place bounds on Λ by observing slight deviations from predictions made using the renormalizable sector of the EFT. The main interpretative point is that Λ has epistemological significance in this full-RG-based understanding of EFTs; Λ is not given a physical interpretation in the EFT framework. Perhaps a successor theory will supply a physical interpretation of Λ_{new} as part of its account of new physics above Λ_{new} , but this would be extrinsic to the particular EFT and is not pertinent to how EFTs are used now. Since we can only ever fix

¹³A potential addition to this list is the application of the full-RG in Polchinski (1984) mentioned at the end of Sec. 4.2.

a finite number of the infinite couplings in the EFT, we actually work with an equivalence class of different EFTs, specified by the precision to which we can empirically determine known couplings, and the degree to which irrelevant couplings are suppressed. We can treat all QFTs in particle physics as EFTs in this sense, noting that for the special case of the renormalizable sector of an EFT, the separation scale need not appear explicitly.

This understanding of EFTs based on the full-RG is reflected in current attitudes towards EFTs in particle physics. In particular, the Standard Model, on this view, just is an incompletely specified SMEFT. In the more general case, this view emphasizes the fact that EFTs are standalone theories. As Manohar (2020, p. 65) argues,

It is much better to think of EFTs in terms of the physical problem you are trying to solve, rather than as the limit of some other theory. The EFT is then constructed out of the dynamical degrees of freedom (fields) that are relevant for the problem. The focus should be on what you want [in the EFT], not on what you don't [i.e., integrating out high-energy modes].

Thus, our view of EFTs using the full-RG fits better with the idea that EFTs can be treated as independent theories, rather than limits of some other theory.

Next, we will highlight some differences between our account of EFTs in particle physics and other recent accounts. We agree with some elements of Rivat's (2020) recent characterization of an effective theory as being defined in terms of an explicit separation of scales within the theory. However, we disagree that this amounts to a "robust specification of the ranges of scales where it is likely to be empirically inaccurate" incorporated within the structure of the EFT (p.10). In top-down constructions, the actual value of the separation parameter is only known for reasons *extrinsic* to the EFT, i.e., from the high-energy theory. In bottom-up constructions, there may be no known value for the separation scale, especially if the future theory is unknown. We claim instead that effective theories should be understood as constructed with a separation of scales, and the value—or range of values—of the separation scale is *external* to the EFT. For example, the Standard Model does not tell us that it must break down beyond the Planck scale; the arguments supporting that upper bound of validity come from dimensional analysis relating quantum theory and general relativity. It could also turn out to be the case that the separation scale for the Standard Model is actually much lower than the Planck scale. This would in turn be dictated by some currently unknown successor theory, not by the Standard Model itself.¹⁴

Our view of the primacy of the full-RG for EFTs—noting the equivalence of Lagrangians with couplings scaled by the full-RG acting on \mathcal{E} —leads to an understanding of EFTs that is similar in spirit to that of Rosaler and Harlander (2019) and Harlander and Rosaler (2019). In particular, we agree with their argument against privileging any one set of parameters (e.g., bare parameters) as more fundamental. Under the action of the full-RG, changes to the scale induce corresponding changes in the couplings to ensure that the n -point functions remain unchanged, and therefore that these parameters at different \mathcal{E} values yield the same predictions. Thus, values of parameters at different \mathcal{E} can be considered empirically equivalent in this sense. However, Rosaler and Harlander take all reparameterizations of the couplings under the full-RG to be physically equivalent, and they conclude that the scale parameter is unphysical. We disagree, and maintain that the parameter \mathcal{E} has a clear physical interpretation as the energy scale at which the system is probed. That the full-RG construction ensures that the theory's predictions are constrained to be invariant under changes to \mathcal{E} does not undermine this interpretation of \mathcal{E} .

Thinking of EFTs in particle physics in terms of the full-RG has implications for views on reduction and emergence. There is an extensive literature on the implications of EFTs and the RG for intertheoretic relations (Bain 2013; Batterman 2011; Batterman 2019; Butterfield and Bouatta 2016; Castellani 2002; Morrison 2012), and much of it turns on the coarse-graining procedure defined by the semi-RG. Bain (2013), for instance, argues that emergence holds between EFTs due to the elimination of degrees of freedom: this is most naturally interpreted as a coarse-graining procedure. Palacios (2022) similarly argues for a version of emergence based on coarse-graining. Recognizing that the full-RG is the primary form of the RG applied in particle physics thus undermines the applicability of this account in this domain because the full-RG

¹⁴In cases of QFTs with a Landau pole, one can place an (extremely high) upper bound on the domain of applicability of the theory. But this is neither particularly informative nor a general feature of EFTs. For QED, the Landau pole is estimated to be $\mathcal{O}(10^{286})\text{eV}$, far above and Planck scale, and other interesting QFTs like QCD are asymptotically free, and therefore definable to all energy scales.

does not admit an interpretation in terms of coarse-graining. Of course, there is the potential for accounts of emergence such as Bain’s to be adapted to the context of the full-RG in particle physics by adopting a characterization of the requisite intertheoretic relations that is less substantial than coarse-graining. The main point that we want to stress is that any account according to which EFTs in particle physics are emergent in virtue of coarse-graining is undermined by our arguments to favour the full-RG. We agree with Castellani (2002), for example, who argues only that the relationships between EFTs as defined by the RG pose a novel form of intertheoretic relationship, and serve as a fruitful ground for issues of emergence and reduction. We insist only that the full-RG should be the basis for analyzing the relationships between EFTs in particle physics. A full investigation into the compatibility of the full-RG formalism with different notions of emergence and reduction will be the subject of future work.

Many other accounts of emergence related to EFTs and the RG are developed in the context of condensed matter physics (Batterman 2011; Butterfield and Bouatta 2016; Morrison 2012). We have argued in Secs. 1 and 2 that any interpretational work on EFTs and the RG—including their implications for emergence and reduction—must be domain specific. As a consequence, our analysis here has no bearing on these issues as they pertain to condensed matter physics. We will further elaborate on the physical differences between domains in the next section.

Finally, we comment on possible implications of our arguments in favour basing our understanding of EFTs in particle physics on the full-RG for recent proposals in Williams (2019) and J. D. Fraser (2018) for Effective Realism—that is, a brand of realism based on the EFT perspective in particle physics. As Ruetsche (2020) convincingly argues, Effective Realism must include a specification of how EFTs describe the world in order to distinguish the position from mere empiricism. How our arguments in this paper relate to Effective Realism depends on the details of how this realist commitment is spelled out. If establishing the approximate truth of the low-energy EFT as judged from the perspective of a high-energy theory requires the semi-group and associated coarse-graining, then our arguments in support of the full-RG would put pressure on Effective Realism.

Effective realism also invokes an assumption about the relationship between the high-energy EFT just below Λ_{new} and the (currently unknown) successor theory for physics above Λ_{new} . For our current EFTs to be deemed approximately true as judged from the perspective of a successor theory, the relationship between the EFT just below Λ_{new} and the successor theory must be stronger than mere empirical equivalence. One must assume that there is a principled way to construct the SMEFT at Λ_{new} from the successor theory, and that this construction proceeds via a semi-RG style coarse graining. In the strongest form, this requires that the successor theory has at least one formulation as a local Lagrangian field theory (cf. Ruetsche (2018) and Koberinski and Smeenk (2022) for arguments as to why this is implausible). At minimum, the unknown theory must be amenable to some generalization of the semi-RG formalism, such that this relationship can hold. In contrast, application of the full-RG formalism does not require any assumptions about the relationship between the successor theory and the EFT just below Λ_{new} . As we have emphasized, the full-RG implements the requirement that EFTs for different \mathcal{E} yield the same n -point functions, which secures empirical equivalence between EFTs at different energy scales. The fact that the full-RG requires assumptions about high-energy physics that are more minimal in these respects is, in our view, a point in favour of the full-RG formalism.

5 Interpretations of EFTs are domain-specific

To this point, we have focused on EFTs and RG methods in particle physics. We have argued that the full-RG formulation best serves the purposes for which RG methods are used in particle physics. Furthermore, we contend that the proper understanding of what an EFT is depends on the full-RG. As a result, coarse-graining is not essential to explanations supported by EFTs in particle physics because the full-RG is reversible—no fine-grained information is lost, nor explicitly excluded via some transformation. In contrast, both the physical concept of coarse-graining and the semi-RG have played an important role in condensed matter physics. The role that coarse-graining plays in RG-supported explanations of critical phenomena in condensed matter physics is a matter of debate (Batterman 2019; Reutlinger 2014; Sullivan 2019). We do not aim to defend any particular interpretation of the RG and EFTs in condensed matter physics in this paper, but we do defend the position that the interpretation of EFTs in condensed matter physics is largely

independent of their interpretation in particle physics. Formal differences between the applications of RG methods have been emphasized to this point in the paper. In this section we will consider differences in the goals of applying RG methods in condensed matter and particle physics and explain why formal analogies between the two domains are compatible with differences in physical interpretation.¹⁵

Consider the use of an Ising model for investigating the large distance-scale magnetic properties of an iron bar. Near a critical point the correlation length approaches infinity, indicating that there are large distance-scale correlations in the system. The iron bar is represented by a d -dimensional lattice of atoms with spacing a . Each of the atoms is assigned a spin. Kadanoff’s block-scaling transformation divides the lattice into blocks of n^d sites, assigns each block the mean spin of each of the atoms in the block, and scales the Hamiltonian representing the interactions among the spins accordingly. The averaging is a coarse-graining operation that explicitly involves loss of fine-detail information about the initial array of spins, yielding a semi-group of transformations. Furthermore, averaging iteratively introduces new variables as spins are replaced by mean spins. Iterative application of the block-scaling transformation is useful for probing the large distance behaviour of the system when the micro-level details that are averaged out are irrelevant to the large distance-scale behaviour, such as near a critical point. The block-scaling transformation is the basis for an explanation of the universality of critical behaviour: there is a universality class of different Hamiltonians at the lattice scale that, after repeated applications of the block-scaling transformation, all yield approximately the same effective Hamiltonian at large distance scales. Underlying this explanation is the same mathematical structure presented in Sec. 4.2: the effective Hamiltonian reflects terms that are irrelevant and relevant with respect to the RG transformation and a fixed point. The particle physics analogues of this explanation of universal behaviour are also those surveyed in Sec. 4.2, namely the explanations of the empirical success of the Standard Model at low energies and the success of renormalizability as a model construction criterion.

The Wilsonian semi-RG transformation is the direct formal analogue of the block-scaling transformation, with energy-momentum scale the analogue of distance scale. The goal of describing low energy behaviour of the particle physics system is analogous to the goal of describing large distance behaviour of the condensed matter system. “Integrating out” the high energy fields is formally analogous to the block-spin averaging transformation. The initial Lagrangian $\mathcal{L}[\phi_i]$ in Eq. (1) is the analogue of the initial Hamiltonian for the atomic lattice and the Lagrangian $\mathcal{L}[\phi_i^l, \Lambda]$ that results from integrating out the high energy field modes is the analogue of the Hamiltonian that results from applying the block-scaling transformation.

An immediately apparent difference in physical interpretation is that the Wilsonian semi-group transformation is not interpreted as an averaging operation. The interpretation of the block-scaling transformation as an averaging operation is intimately related to the coarse-graining in condensed matter physics.

There are also important differences in elements of the formalism associated with the Kadanoff and Wilsonian semi-RG scaling transformations. The physical interpretations of the base scales a and Λ_{new} differ in important respects. In the Ising model, the original “bare” value of a has physical significance as atomic lattice spacing and the spins at this scale are interpreted as the actual spins. In Sec. 4.3 we argued that the bottom-up, full-RG accords Λ epistemological significance, and does not physically privilege a particular value of Λ . In the top-down perspective of the Wilsonian semi-RG, the separation scale Λ_{new} cannot be interpreted as physically privileged in the same way as the initial lattice spacing either. As in the full-RG case, the choice of a particular value for Λ is arbitrary, as long as $\Lambda \gg \mathcal{E}$. The value of the lattice spacing can be determined empirically, whereas determining a possible value or range of possible values for Λ_{new} where our current EFTs become empirically inadequate requires the appeal to a theory external to the Standard Model (e.g., quantum gravity) or, in principle, the bottom-up application of the full-RG methods. As a result, there is a physically well-determined starting point for initiating coarse-graining in the condensed matter case, but not in the particle physics case.

RG methods underwrite a disparate set of explanations in condensed matter and particle physics, with prominent examples being the the explanation of universal behaviour and the explanations of the empirical

¹⁵It has been argued elsewhere that formal analogies between particle physics and condensed matter do not support the carryover of physical interpretation (D. Fraser and Koberinski 2016; D. Fraser 2020a; D. Fraser 2020b). This section supplies additional arguments to support the same conclusion. The analysis here is complementary to that in D. Fraser 2020a, which analyzes the early use of formal analogies between critical phenomena and renormalization in QFT in Wilson and Kogut 1974. We focus on the later development of the EFT perspective in particle physics. More specifically, for Wilson and Kogut 1974 the $\Lambda \rightarrow \infty$ limit in Eq. (13) is of central importance, whereas the $\mathcal{E} \rightarrow 0$ limit in Eq. (13) is the most important one for investigating EFTs.

success of the Standard Model laid out in Sec. 4.2. These explanations are based on fixed points of the pertinent RG transformation and irrelevant terms of the Hamiltonian or Lagrangian tending to zero in a physically applicable limit. RG methods are capable of supporting such a diverse set of explanations in virtue of flexibility in the application of the formalism. The limits that are used in these explanations are the following (cf. discussion of Wilson on statistical continuum limits in Sec. 2):

$$\text{Condensed matter (distance scales): } \frac{a}{\xi} \rightarrow 0 \quad \text{as } \xi \rightarrow \infty, \quad (12)$$

$$\text{Particle physics (energy scales): } \frac{\mathcal{E}}{\Lambda} \rightarrow 0 \quad \text{as } \mathcal{E} \rightarrow 0 \text{ or } \Lambda \rightarrow \infty, \quad (13)$$

where ξ is the correlation length for the condensed matter system. Abstractly, the ratio of the upper and lower scales that bound the application of the RG transformation must go to zero. The concrete interpretation of the scales (e.g., distance or energy) and which term in the ratio is varied is application-dependent. In particle physics from the abstract Wilsonian top-down perspective, it is possible to fix Λ at some arbitrarily high energy and take the $\mathcal{E} \rightarrow 0$ limit in which the low energy scale goes to zero. This limit is of course an idealization; in the condensed matter application, the critical limit $\xi \rightarrow \infty$ is also an idealization. In particle physics, the context of renormalization and regularization and the arbitrariness of Λ generates another question that does not arise in the condensed matter case: can the imposed cutoff Λ be removed without affecting the set of scattering amplitudes? That is, is it possible to take the $\mathcal{E}/\Lambda \rightarrow 0$ limit by taking $\Lambda \rightarrow \infty$? (Of course, an affirmative answer to this question requires the existence of a suitable fixed point.) There is no physical analogue of this question in the iron bar case. For a given theory, the atomic lattice spacing a is assumed to remain fixed for physical reasons. This is another illustration of how the RG formalism detaches from physical interpretation. The flexibility in how the formalism is applied (i.e., how the limits are taken) along with differences in physical interpretation partly accounts for the wide scope of applicability of RG methods. The salient physical differences undermine a physical analogy.

In this section we have focused on the Wilsonian semi-group variant of RG methods because this is the closest formal analogue to the Kadanoff block-scaling transformation, but in the rest of this paper we have advocated for basing foundational conclusions on the full-group variant in particle physics. The full-RG makes it even clearer that coarse-graining is not an essential concept in particle physics. For particle physics EFTs, the separation scale is externally imposed, or remains unfixed at some high value such that we can assume $\mathcal{E}/\Lambda \ll 1$. We then care about scaling as \mathcal{E} varies, such that the n -point functions remain invariant. All of the theory's predictions are therefore largely insensitive to this scale. The full-RG equations for the system are a direct result of this constraint, which gives rise to the compensating scaling of couplings for the theory.

The Wilsonian abstract semi-RG is a formal analogue of the block-spin transformation, but the applications of these formalisms and their physical representations diverge. As a result, the dissociation of RG-transformations and coarse-graining in particle physics does not imply that this is also the appropriate interpretation of scaling transformations in condensed matter physics. Of course, the Kadanoff block-scaling transformation and intuitions about coarse-graining from condensed matter physics played an inspirational role in the development of RG methods in particle physics by Wilson and others. In retrospect, these intuitions supported the use of formal analogies to develop new formal frameworks for particle physics, but do not provide physical insight into particle physics (D. Fraser 2020a). The physical interpretation of RG methods and associated EFTs is domain-specific.

6 Conclusions

We have argued that understanding the role of RG methods for EFTs in particle physics requires close attention to their use *in this domain*. Our main topic of focus has been on the explanatory role played by the full-RG for understanding relationships between EFTs. As with any physical interpretation of theories, the local, contextual features of the target domain are of central importance, and one should be careful not to read physical significance straight off of the formalism. This paper marks a first step in a philosophical analysis of the full-RG, by clarifying the distinctions between the semi-RG and full-RG as applied to EFTs in particle physics. We examined the construction procedure for EFTs, as well as the role of the (full-)RG

in explanations about the local features of EFTs and their role in explaining physical phenomena. Our focus has been on the types of explanation provided by the full-RG in the context of the Standard Model of particle physics. Consistent with our thesis about domain-specific interpretation, we understand that further work will be required to elucidate further philosophically fruitful features of this understanding of EFTs in particle physics. This future work must be within the proper context of particle physics and (relativistic) QFT.

With this in mind, the upshot of this paper is as follows. First, RG methods are a loose family of formal mathematical techniques, the details of which differ depending on the discipline and formalism to which they are applied. For the purposes of understanding the structure of EFTs in particle physics, we focused on two methods: the semi-RG and the full-RG. The semi-RG provides the justification for the Wilson action picture of EFTs, where one starts with a high-energy theory and successively “integrates out” the high-energy degrees of freedom to construct a low-energy EFT. This coarse-graining picture is abstract, and has occasional uses in particle physics. The top-down approach is also used in concrete cases, such as to relate the Fermi theory to the Standard Model. However there are important examples of EFTs that cannot be constructed in a top-down fashion using the semi-RG or concrete methods. This motivates consideration of the bottom-up perspective using the full-RG, which we have argued in this paper is better suited to the construction of EFTs in general and also to the applications of EFTs in particle physics.

The main conclusion of our analysis of the structure of EFTs is that the full-RG is sufficient for understanding the structure of a given EFT, and for exploring theory space more generally. Though the top-down perspective is sometimes useful, coarse-graining and the semi-RG are not needed, and in fact inapplicable for many of the more mundane uses of RG methods, like scaling and renormalization. In particular, we focused attention on questions of global structure on the space of EFTs, and argued that the full-RG can ground these more ambitious explanations as well. Not only is the full-RG more generally applicable in particle physics, but it grounds a more secure explanation of global structure as well. In Sec. 4.3 we argued that the full-RG explanations are not only adequate, but furnish better explanations of global structure in particle physics EFTs. This is due to the greater generality of the full-RG, and to its weaker assumptions, more closely grounded in our epistemic situation based on the empirical success of known EFTs.

Our rejection of the semi-RG in particle physics may surprise some readers. Much of the existing philosophical literature treats EFTs and the RG in particle physics almost exclusively in terms of a top-down perspective, using the semi-RG. We believe that the prevalence of this view has two sources: (1) from looking only at the abstract semi-RG within particle physics, and generalizing from there; and (2) from a mistaken physical analogy with condensed matter physics. While Sec. 3 and 4 provide arguments to convince those who adopt the semi-RG view because of (1), in Sec. 5 we turned to (2). We argued that the analogy between the semi-RG in condensed matter physics and the semi-RG in particle physics is purely formal, and that the physical differences between the domains prevent a clear carryover of interpretation. In particular, the coarse-graining associated with the semi-RG in particle physics is not an averaging procedure, while it is in condensed matter. The limits that figure in RG-based explanations have both different formal implementations and different physical interpretations. This should not come as a surprise in light of the fact that RG methods are heterogeneous and bear a loose family resemblance (Sec. 2). By contrasting the goals of using RG methods in each domain, we gain some insight into why the full-RG is better suited to particle physics and can appreciate why interpretations of EFTs need to be domain-specific.

One product of our arguments is an account of EFTs in particle physics that is informed by the full-RG formalism, which we reiterate. The bottom-up perspective encourages us to think of an EFT as defined in terms of a Lagrangian with an infinite number of coupling terms, and a built-in (but perhaps unfixed) separation scale Λ . Importantly, the EFT *on its own* does not fix the value of Λ . The fields and dynamical symmetries are chosen to reflect the degrees of freedom at the energy scale of concern. Couplings are likewise fixed at some reference energy scale, and their scaling behaviour is determined by the action of the full-RG. In principle, one EFT can be distinguished from another by containing different fields,¹⁶ different symmetry constraints, or different values of *any* couplings at a given energy scale. Since in practice we can only ever empirically determine some finite set of couplings at any energy scale, we are in the epistemic position of

¹⁶With the major caveat that EFTs are quite permissive in the transformations one can make between different field representations. One can ensure that n -point functions are invariant under general transformations of the form $\phi \rightarrow f(\phi, \phi')$, where f contains combinations of polynomials of its arguments. These transformations are compensated by including higher-order terms from the EFT in n -point functions.

working with a *set* of EFTs, at least in instances where the successor theory is unknown. Precise empirical tests of the EFT allow us to narrow down the set of possible EFTs.

We can use the full-RG to justify this picture, so long as we acknowledge the finite precision with which we can make predictions and measurements in particle physics. Assuming a fixed separation scale Λ , the action of the full-RG on a class of EFTs specified at some high-energy scale that differ only in the values of a subset of irrelevant couplings as we lower the energy scale \mathcal{E} is to make predictions converge (cf. Sec. 4.2). Our predictions and measurements are never infinitely precise, so idealizing a theory as partially defined by a set of exact values will always lead to imprecision with respect to our use of these theories.¹⁷ The additional imprecision introduced here is that EFTs require an infinite set of values to be specified to exactly delimit the theory. We don't think that this additional imprecision poses any problems for working with or using EFTs, as long as we acknowledge that we never have exactly one theory in the class to compare with experiment.

There are open questions that we have not addressed here. Our focus has been on the epistemology and interpretation of EFTs in particle physics, and we have not considered the important metaphysical issues of emergence and reduction as they pertain to relationships between different EFTs. The adoption of the full-RG interpretation of EFTs that we have advocated has implications for the relationship between EFTs at different scales in virtue of the inadmissibility of an interpretation of the full-RG in terms of coarse-graining. A negative point is that any view of emergence (or reduction) in particle physics that hinges on coarse-graining would be incompatible with our interpretation of EFTs. We briefly discussed this issue in Sec 4.3; however, unpacking the implications of the full-RG for emergence and reduction in particle physics will require more careful consideration that we leave for future work.

Finally, though our discussion has been focused on the empirical success of EFTs, we note that the bottom-up, full-RG perspective is neutral regarding debates on realism versus instrumentalism. Earlier work on understanding the RG and EFTs suggested that thinking of our best theories as EFTs forces an instrumentalist perspective, as argued in Cao and Schweber (1993) and Butterfield and Bouatta (2016). More recent work, which has made heavy use of the semi-RG and the top-down perspective, argues to the contrary that EFTs support a novel prospective realism (J. D. Fraser 2018; Williams 2019). Though we have made a few skeptical remarks about the prospects of this form of realism, we want to stress that while the bottom-up, full-RG perspective foregrounds empirical equivalence, it does not automatically lead to instrumentalism. As we have stressed throughout this paper, we believe that no formalism comes with a built-in interpretation. This holds true for the bottom-up perspective of EFTs as well. We have instead argued that an interpretation of EFTs in particle physics based on the full-RG formalism clarifies the empirical and epistemological roles that the EFT framework plays in this domain.

¹⁷Alternatively, one could accept a fundamental imprecision as characteristic of physical facts, and define theories with this imprecision built in, as in Miller (2020).

References

- Bain, Jonathan (2013). “Emergence in effective field theories”. In: *European Journal for Philosophy of Science* 3.3, pp. 257–273.
- Batterman, Robert W. (2011). “Emergence, singularities, and symmetry breaking”. In: *Foundations of Physics* 41.6, pp. 1031–1050.
- Batterman, Robert W. (2019). “Universality and RG Explanations”. In: *Perspectives on Science* 27.1, pp. 26–47.
- Bogoliubov, N.N. and D.V. Shirkov (1955). “Application of the Renormalization Group to Improve the Formulae of Perturbation Theory”. In: *Dokl. Akad. Nauk SSSR* 103, pp. 391–394.
- Bokulich, Alisa (2020). “Losing Sight of the Forest for the Ψ ”. In: *Scientific Realism and the Quantum*. Ed. by Steven French and Juha Saatsi. Oxford University Press, USA, pp. 185–211.
- Burgess, Cliff P. (2007). “An introduction to effective field theory”. In: *Annual Review of Nuclear and Particle Science* 57, pp. 329–362.
- Butterfield, Jeremy and Nazim Bouatta (2016). “Renormalization for philosophers”. In: *Metaphysics in Contemporary Physics*. Ed. by Tomasz Bigaj and Christian Wüthrich. Brill Rodopi, pp. 437–485.
- Cao, Tian Yu and Silvan S Schweber (1993). “The conceptual foundations and the philosophical aspects of renormalization theory”. In: *Synthese* 97.1, pp. 33–108.
- Castellani, Elena (2002). “Reductionism, emergence, and effective field theories”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 33.2, pp. 251–267.
- Crowther, Karen (2015). “Decoupling emergence and reduction in physics”. In: *European Journal for Philosophy of Science* 5.3, pp. 419–445.
- Fraser, Doreen (2020a). “The development of renormalization group methods for particle physics: Formal analogies between classical statistical mechanics and quantum field theory”. In: *Synthese*, pp. 1–37.
- Fraser, Doreen (2020b). “The non-miraculous success of formal analogies in quantum theories”. In: *Scientific Realism and the Quantum*. Ed. by Steven French and Juha Saatsi. Oxford University Press, USA, p. 255.
- Fraser, Doreen and Adam Koberinski (2016). “The Higgs mechanism and superconductivity: A case study of formal analogies”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 55, pp. 72–91.
- Fraser, James D. (2018). “Renormalization and the formulation of scientific realism”. In: *Philosophy of Science* 85.5, pp. 1164–1175.
- Fraser, James D. (2021). “The Twin Origins of Renormalization Group Concepts”. In: *Studies in History and Philosophy of Science* 89, pp. 114–128.
- Gell-Mann, Murray and Francis E Low (1954). “Quantum electrodynamics at small distances”. In: *Physical Review* 95.5, p. 1300.
- Hancox-Li, Leif (2015). “Moving Beyond “Theory T””. Available at <http://d-scholarship.pitt.edu/24801/>. PhD thesis. University of Pittsburgh.
- Hanneke, D., S. Fogwell Hoogerheide, and Gerald Gabrielse (2008). “New measurement of the electron magnetic moment and the fine structure constant”. In: *Physical Review Letters* 100.12, p. 120801.
- Harlander, Robert and Joshua Rosaler (2019). “Higgs Naturalness and Renormalized Parameters”. In: *Foundations of Physics* 49.9, pp. 879–897.
- Huggett, Nick (2002). “Renormalization and the Disunity of Science”. In: *Ontological Aspects of Quantum Field Theory*. Ed. by Meinard Kuhlman, Holger Lyre, and Andrew Wayne. World Scientific Publishing Co. Pte. Ltd, pp. 255–277.
- Huggett, Nick and Robert Weingard (1995). “The renormalisation group and effective field theories”. In: *Synthese* 102.1, pp. 171–194.
- Jona-Lasinio, Giovanni (2001). “Renormalization group and probability theory”. In: *Physics Reports* 352.4-6, pp. 439–458.
- Jona-Lasinio, Giovanni (2010). “Analogies in theoretical physics”. In: *Progress of Theoretical Physics Supplement* 184, pp. 1–15.
- Koberinski, Adam (2021a). “Mathematical developments in the rise of Yang–Mills gauge theories”. In: *Synthese* 198, pp. 3747–3777.

- Koberinski, Adam (2021b). “Problems with the cosmological constant problem”. In: *Philosophy Beyond Spacetime*. Ed. by Christian Wüthrich, Baptiste Le Bihan, and Nick Hugget. Oxford University Press, pp. 260–280.
- Koberinski, Adam (2021c). “Regularizing (away) vacuum energy”. In: *Foundations of Physics* 51.1, pp. 1–22.
- Koberinski, Adam and Chris Smeenk (2020). “Q.E.D., QED”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 71, pp. 1–13.
- Koberinski, Adam and Chris Smeenk (2022). “ Λ and the limits of effective field theory”. In: *Philosophy of Science* doi:10.1017/psa.2022.16, pp. 1–26.
- Manohar, Aneesh V (2020). “Introduction to Effective Field Theories”. In: *Effective Field Theory in Particle Physics and Cosmology: Lecture Notes of the Les Houches Summer School: Volume 108, July 2017*. Ed. by Sacha Davidson et al. Vol. 108. Oxford University Press, p. 47.
- Miller, Michael E (2020). “Worldly imprecision”. In: *Philosophical Studies*, pp. 1–17.
- Morrison, Margaret (2012). “Emergent physics and micro-ontology”. In: *Philosophy of Science* 79.1, pp. 141–166.
- Nambu, Yoichiro and Giovanni Jona-Lasinio (1961). “Dynamical model of elementary particles based on an analogy with superconductivity. II”. In: *Physical Review* 124.1, p. 246.
- Palacios, Patricia (2022). *Emergence and Reduction in Physics*. Elements in the Philosophy of Physics. Cambridge University Press.
- Parker, Richard H et al. (2018). “Measurement of the fine-structure constant as a test of the Standard Model”. In: *Science* 360.6385, pp. 191–195.
- Polchinski, Joseph (1984). “Renormalization and effective Lagrangians”. In: *Nuclear Physics B* 231.2, pp. 269–295.
- Rejzner, Kasia (2015). *Perturbative Algebraic Quantum Field Theory*. New York, NY: Springer Berlin Heidelberg. ISBN: 978-3-319-25899-7.
- Reutlinger, Alexander (Dec. 2014). “Why Is There Universal Macrobehavior? Renormalization Group Explanation as Noncausal Explanation”. In: *Philosophy of Science* 81.5, pp. 1157–1170. DOI: 10.1086/677887.
- Rivat, Sébastien (2019). “Renormalization scrutinized”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 68, pp. 23–39.
- Rivat, Sébastien (2020). “Effective theories and infinite idealizations: a challenge for scientific realism”. In: *Synthese*, pp. 1–30.
- Rivat, Sébastien (2021). “Drawing scales apart: The origins of Wilson’s conception of effective field theories”. In: *Studies in History and Philosophy of Science Part A* 90, pp. 321–338.
- Rosaler, Joshua and Robert Harlander (2019). “Naturalness, Wilsonian renormalization, and “fundamental parameters” in quantum field theory”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 66, pp. 118–134.
- Ruetsche, Laura (2018). “Renormalization group realism: The ascent of pessimism”. In: *Philosophy of Science* 85.5, pp. 1176–1189.
- Ruetsche, Laura (2020). “Perturbing realism”. In: *Scientific Realism and the Quantum*. Ed. by Steven French and Juha Saatsi. Oxford University Press.
- Scherer, Stefan (2003). “Introduction to chiral perturbation theory”. In: *Advances in Nuclear Physics* 27, pp. 277–538.
- Schneider, Mike D (2020a). “Betting on future physics”. In: *The British Journal for the Philosophy of Science*. DOI: 10.1093/bjps/axz040.
- Schneider, Mike D (2020b). “What’s the Problem with the Cosmological Constant?” In: *Philosophy of Science* 87.1, pp. 1–20.
- Shirkov, D.V. (1999). *Evolution of the Bogolubov Renormalization Group*. arXiv: hep-th/9909024. URL: <http://arxiv.org/abs/hep-th/9909024>.
- Sullivan, Emily (2019). “Universality Caused: The Case of Renormalization Group Explanation”. In: *European Journal for Philosophy of Science* 9.3, p. 36. DOI: 10.1007/s13194-019-0260-x.
- Wallace, David (2006). “In defence of naïveté: The conceptual status of Lagrangian quantum field theory”. In: *Synthese* 151.1, pp. 33–80.
- Wallace, David (2011). “Taking particle physics seriously: A critique of the algebraic approach to quantum field theory”. In: *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics* 42.2, pp. 116–125.

- Wallace, David (2019). “Naturalness and emergence”. In: *The Monist* 102.4, pp. 499–524.
- Wallace, David (2022). “The Quantum Theory of Fields”. In: *The Routledge Companion to the Philosophy of Physics*. Ed. by Eleanor Knox and Alistair Wilson. Routledge, pp. 275–295.
- Weinberg, Steven (1979). “Phenomenological Lagrangians”. In: *Physica a* 96.1-2, pp. 327–340.
- Williams, Porter (2015). “Naturalness, the autonomy of scales, and the 125 GeV Higgs”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 51, pp. 82–96.
- Williams, Porter (2019). “Scientific realism made effective”. In: *The British Journal for the Philosophy of Science* 70.1, pp. 209–237.
- Williams, Porter (2022). “Renormalization Group Methods”. In: *The Routledge Companion to the Philosophy of Physics*. Ed. by Eleanor Knox and Alistair Wilson. Routledge, pp. 296–310.
- Wilson, Kenneth G. (1975). “The renormalization group: Critical phenomena and the Kondo problem”. In: *Reviews of Modern Physics* 47.4, p. 773.
- Wilson, Kenneth G. and John Kogut (1974). “The renormalization group and the ϵ expansion”. In: *Physics Reports* 12.2, pp. 75–199.