DOI: 10.1017/psa.2025.10132
This is a manuscript accepted for publication in *Philosophy of Science*.
This version may be subject to change during the production process.

CONTRIBUTED PAPER

A Problem of Scale

Laura Ruetsche[1]

[1]The University of Michigan
Corresponding author: Laura Ruetsche; ruetsche@umich.edu

Abstract
Subject to techniques of perturbative renormalization, the Standard Model makes empirical predictions that are stupendously successful. But also deeply mysterious. Not every quantum field theory (qft) is renormalizable. Indeed, most aren't. The mystery is: why should we be so lucky, that we live in a world governed by a renormalizable qft? I explicate this *Renormalizability Puzzle*, and explain why Renormalization Group (RG) approaches are widely thought to resolve it. Looking under the hood of the RG resolution, I identify a load-bearing element that might not be adequate to the explanatory burden the RG resolution places upon it.

1. Introduction

Our best quantum field theories (qfts), the interacting qfts making up the Standard Model of high energy physics, are *renormalizable*. When pressed into service to calculate empirical quantities, these theories threaten to yield meaningless infinities. However, a variety of *perturbative renormalization techniques*, developed in the mid-twentieth century, succeed in averting that threat.[1] Subject to techniques of perturbative renormalization, the Standard Model makes empirical predictions that are not only sensibly finite but also outlandishly (e.g., to one part in a trillion!) accurate. This stunning success numbers among the greatest triumphs of contemporary physics. It's also deeply mysterious. Not every qft is renormalizable. Indeed, most aren't. The mystery is: why should we be so lucky, that we live in a world governed by a qft with the rare property that its froward infinities can be tamed by perturbative renormalization techniques? This is *the Renormalizability Puzzle*.

It is now a catechism among high-energy physicists that a solution to the puzzle comes from Renormalization Group (RG) analyses, which "help explain why renormalizable theories play such an important role in physics" (Schwartz 2014, 443). The Standard Model consists of *effective* theories (the catechism goes). Effective theories aren't the fundamental truth, only a tolerable approximation to its implications for the limited domain of energies accessible in accelerator experiments. And Renormalization Group (RG) analyses teach us that effective theories are renormalizable. This *RG*

[1] (Schweber 2020) and (Blum forthcoming) are philosophically-minded accounts of the history.

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

Explanation's elegant resolution of the Renormalizability Puzzle has inspired philosophers to reconfigure the very terms of the scientific realism debate (Fraser 2016, 2018, 2020ab; Williams 2019, 2023, §5.2).

Although my brazenly iconoclastic aim here is to suggest that the RG Explanation might not be as satisfying as it initially appears, I also have aims that are more community-minded. One is to provide a technically undemanding account of aspects of the RG Explanation complementary to aspects foregrounded by a number of excellent recent treatments in the literature.[2] Another is to draw attention to methodological questions about applied mathematics dramatized by the RG Explanation. Chief among these is: is it reasonable to draw conclusions about the physical world on the basis of what's necessary for (or even merely conducive to) the success of certain strategies of mathematical analysis?

I proceed as follows. §2-§3 are stage-setting. §2 motivates the Renormalizability Puzzle. Sketching the RG Explanation's resolution, §3 highlights the role played therein by a "dimension-cancelling" maneuver undertaken to render the qfts under study amenable to convenient analysis. The stage thus set, the drama unfolds in §4. Conducting a very brief trial of the dimension-cancelling manueuver, §4 asks whether it's capable of supporting the explanatory burden the RG Explanation places on it.

2. The Renormalizability Puzzle

In qft, *lagrangians* are used to calculate quantities of empirical interest. The hitch is that, naively pursued, these calculations deliver infinities, rather than finite numbers to test against experiment. The conventional fix—renormalization— is to pursue more sophisticated calculations. Not every lagrangian can be cured by the conventional fix: not every lagrangian is *renormalizable*. Indeed, in some sense, most aren't. The good news is that Standard Model lagrangians are: they can be renormalized to yield finite predictions. The bad news is that the good news stuck us with the Renormalizability Puzzle, that "it seemed incredibly lucky that we could describe so much in particle physics using renormalizable QFTs" (Williams 2023, 45).[3] This section offers a minimally technical introduction to the Renormalizability Puzzle.

A lagrangian is a function of a collection of spacetime fields, where each field is map from points of a spacetime (which I'll assume to be four-dimensional) to a value, and their spacetime derivatives. Throughout, I'll adopt the drastically simplifying assumption, that a lagrangian involves only one field, a scalar-valued field ϕ, and its derivatives $\partial\phi$.[4] Thus a generic scalar field lagrangian $\mathscr{L}$ is a sum of mononomials of ϕ and $\partial\phi$, modified by coupling coefficients. A schematic example, shaped by some physical

[2] I aspire for my exposition to be accessible to anyone acquainted with the fact, that integrating the same function against different domains—$\int_0^a f(x)dx$ and $\int_0^b f(x)dx$ for $a \neq b$, say—can yield different answers.

[3] The mystery abates if renormalizability is a *consistency constraint*. In the 1970s, it was widely regarded to be just that (Weinberg 1977). "There is something puzzling about demanding renormalizability a priori in this way however. What licenses us to assume that the world is amenable to perturbative analysis?" (Fraser 2020b, 399). To elevate renormalizability to a consistency constraint is to make the sufficiency of an *approximation technique we've ironed out* to calculate *empirical quantities we care about* criterial for a theory's consistency. But a theory could fail to satisfy the criterion through no logical fault of its own, if either our approximation methods or our notion of empirical content are inadequate.

[4] This simplifying assumption sets to one side a host of further foundational questions; see (Miller 2021) for a sense of some of them.

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

background knowledge and a few additional simplifying assumptions[5]:

$$\mathscr{L} = (\partial\phi)^2 + m^2\phi^2 + g_1\phi^4 + g_2\phi^6 + g_3\phi^6 + \ldots \tag{1}$$

The physical background knowledge is that

$$\mathscr{L}_0 = (\partial\phi)^2 + m^2\phi^2$$

is a cracker-jack lagrangian, *the free-field lagrangian*, under complete analytic control. We can draw upon analogies with more familiar theories to recognize its *kinetic term* $(\partial\phi)^2$ as describing the kinetic energy of the free field, and its mass term $m^2\phi^2$ as describing a potential energy associated with the mass m of the field. And that's not all! We can explicitly characterize solutions to the classical equations of motion $\mathscr{L}_0$ generates, and use the classical solution space to construct a Hilbert space undergirding a *quantum* field theory based on that lagrangian (see (Wald 1994) for more on the construction).

Alas, the qft thus obtained is deathly boring. Free fields don't interact; they don't scatter; they pass unaltered through our particle accelerators. For the sake of characterizing interesting physics, we need to add more terms to $\mathscr{L}_0$, *interaction terms* third-order or higher in ϕ. A sum of $\mathscr{L}_0$ and a host of interaction terms, (1)'s generic lagrangian $\mathscr{L}$ amends the free theory $\mathscr{L}_0$ by adding *interactions* ϕ^4, ϕ^6, ϕ^8 … multiplied by *coupling coefficients* g_1, g_2, g_3 … that indicate the strength of those interactions. Think of $\mathscr{L}$ a residing in space $\mathscr{G}$ of lagrangians; each lagrangian in $\mathscr{G}$ is specified by a *coupling vector* listing the coefficients[6] $g_k, g_m, g_1, g_2, g_3, \ldots$ modifying the monomials it depends on.

Now we've got an interacting qft that might describe interesting physics! Only there's a hitch. Even the simplest interacting qft, one that adds just a ϕ^4 interaction to $\mathscr{L}_0$,

$$\mathscr{L}_{\phi^4} = \mathscr{L}_0 + g_1\phi^4 \tag{2}$$

eludes our analytical reach. We don't know how to construct an Hilbert space from its classical solution space. So we pursue quantum theoretic results by other means. We undertake to calculate quantities of empirical interest *perturbatively*. That is, we assume that the coupling g_1 is small, so that $\mathscr{L}_{\phi^4}$ represents a small perturbation to $\mathscr{L}_0$. And we approximate a quantity $\mathscr{M}$ of empirical interest—a scattering amplitude, say—by performing a perturbative expansion, in orders of g_1, around the (usually very dull) precise value $\mathscr{M}_0$ the free lagrangian predicts:

$$\mathscr{M} = \mathscr{M}_0 + g_1\mathscr{M}_1 + (g_1)^2\mathscr{M}_2 + (g_1)^3\mathscr{M}_3 \ldots \tag{3}$$

The $\mathscr{M}_n$ in (3) are n^{th}*-order perturbative corrections*. Perturbation theory is tasked with directing us in their calculation. Assuming g_1 is small, we hope that truncating the perturbative expansion after a few orders yields an approximation to $\mathscr{M}$ accurate enough to test in our particle accelerators.

Path-integrals are a standard method for extracting empirical content from qfts. To calculate a quantity $\mathscr{M}$ of empirical interest, one evaluates a path-integral whose

[5]I'm assuming $\mathscr{L}$ is symmetric under interchange of ϕ and $-\phi$, and pretending spacetime derivatives are absent from interaction terms. I'm also writing things like factors of $\frac{1}{2}$ and minus signs in invisible ink, and setting $c = \hbar = 1$.

[6]where g_k and g_m are coefficients of the kinetic and mass term respectively.

integrand depends on the qft lagrangian.[7] Schematically

$$\mathscr{M} = \int D[\phi]\,\{\ldots \mathscr{L} \ldots\} \tag{4}$$

The $D[\phi]$ denotes a functional integration over field configurations ϕ whose energy and momentum depend on ϕ's frequency. The path integral defining $\mathscr{M}$ integrates over field configurations of all energies/momenta/frequencies, including arbitrarily high ones.

If $\mathscr{L}$ is a free lagrangian, we can execute the path integral defining $\mathscr{M}$. If $\mathscr{L}$ contains interaction terms, however, we can't. Instead, we resort to the perturbative approximation (3). Feynman diagrams guide our calculation of perturbative corrections $\mathscr{M}_1, \mathscr{M}_2 \ldots$. Starting with the "loop-order" correction $\mathscr{M}_2$, Feynman diagrams instruct us to calculate corrections by evaluating integrals that conspicuously and tragically diverge. Brought on by integrating over arbitrarily large momenta/frequencies, these catastrophes are known as *ultraviolet divergences*. Rather than the small adjustments to empirical quantities presupposed by our perturbative approximation, corrections plagued by ultraviolet divergence contribute infinite ones. Ultraviolet divergences render perturbative approximation meaningless.

Renormalization, the conventional way to address the catastrophic divergences, is a recipe for an order-by-order doctoring of a lagrangian's couplings that eventuates, at each order, in finite perturbative corrections. An essential ingredient in the recipe is a finite set of measurements that serve to constrain the doctoring at all orders. Other ingredients struck even the framers of perturbative renormalization approaches as suspiciously ad hoc, even desperate.[8] Notoriously, renormalization restores empirical sanity to perturbative approximations by subtracting one infinity from another to obtain a finite answer.

And yet, for some lagrangians, renormalization techniques succeed, stunningly well. Call lagrangians that surrender finite predictions when subject to renormalization techniques *renomalizable*. Dyson articulated a criterion for identifying such amiable lagrangians. The criterion hinges on the notion of *mass dimension*. It's standard operating procedure in high-energy physics to set Planck's constant equal to the speed of light equal to the dimensionless number 1. This convention means we need only one physical dimension, which particle physicists declare to be $[mass]$.[9] Scalar fields and their derivatives have mass dimension 1, as—unsurprisingly!—does the mass m. Thus each term in $\mathscr{L}_0$ has mass dimension 4, as does the ϕ^4 interaction. Higher order interactions—ϕ^6, ϕ^8 $\ldots$— have mass dimension > 4.

The notion of an action supplies a nomic scaffolding for path-integral approaches. (4) implicates $\mathscr{L}$ in empirical calculations by using it to define an action via $S = exp(-\int \mathscr{L} d^4x)$. The definition requires $\int \mathscr{L} d^4x$ to be dimensionless. Because $d^4x = (\frac{1}{\partial})^4$ has mass dimension -4, this obliges each term in $\mathscr{L}$ to have mass dimension 4.

[7]Details: What I'm schematizing is the calculation of n-point functions from the partition function $\mathscr{Z}(J) = \int D[\phi] exp(-i \int d^4x[\mathscr{L}(\phi) + J])$ defining a qft. For a nice introduction to path integrals and functional integration, see (Hall 2013). But don't fret if these are alien concepts. A feeling for the basic mechanics of ordinary integration will be sufficient to follow the plot of this essay.

[8]Feynman infamously called renormalization "a dippy process" (1995, 128).

[9]Because $c = 1$, $[length] = [time]$. Because $h = 1$, $[mass][length]^2 = [time]$, which implies that $[length] = [mass]^{-1}$. So every physical dimension can be expressed as mass dimension: $[mass] = [mass]$ (obviously!), $[length] = [time] = [mass]^{-1}$.

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

Thus an interaction term with mass dimension d must be multiplied by a coupling g of mass dimension $4 - d$.

The reward for this slog through the niceties of mass dimension is that Dyson's renormalizability criterion is simply stated. A lagrangian is renormalizable only if it contains no interactions of mass dimension > 4, and an interaction is renormalizable only if it has mass dimension ≤ 4. But the vast majority of interactions appearing in our generic lagrangian (1) violate this criterion! Renormalizable lagrangians are confined to a tiny corner of theory space $\mathscr{G}$: the measley three-dimensional subspace $\mathscr{G}^{PR}$ spanned by the handful of terms with mass dimension ≤ 4.

Hence the Renormalizability Puzzle. It somehow just so happens that the lagrangians figuring in the Standard Model—not to mention other useful lagrangians, such as the ϕ^4 lagrangian—*are* renormalizable. They're residents of that tiny region $\mathscr{G}^{PR}$. Particularly given how suspicious renormalization techniques seemed even to their framers, this seems absolutely extraordinary. Why should it be that the lagrangians that govern real-world physics providentially belong to the small but exclusive subspace $\mathscr{G}^{PR}$?

3. The RG Explanation

Viewed through the lens of RG approaches, the space $\mathscr{G}$ of lagrangians is pierced by a flow of *RG trajectories*; the RG Explanation uses the structure of this flow to explain why so many useful lagrangians are renormalizable. Here's how.

Recall the schematic (4): Integrating over all energies, we can predict amplitudes for scattering experiments conducted at all energies. Only we haven't got the wherewithal to conduct arbitrarily high-energy experiments. What we have the wherewithal to conduct are experiments at energies we can achieve in the particle accelerators we've managed to construct. Introduce an *effective energy scale* Λ that sets an upper bound on the energies at which we can test amplitudes extracted from our qft.

Now indulge in what I'll call the *autonomy-of-scales hope*. This is the hope there is some way to describe scale Λ effective physics that prescinds from the details of ultra-high energy physics—that to generate predictions empirically adequate at scale Λ, we need only survey scale Λ fields: there is a lagrangian $\mathscr{L}$ such that we can describe amplitudes for scattering experiments conducted at energies up to Λ, a la (5), by integrating $\mathscr{L}$ against a domain of field configurations that's cut off at Λ. $(\mathscr{L}, \Lambda)$ is a *scale-*Λ *effective theory*: its domain, both of definition and of application, is limited, but it succeeds in that domain.

Revised to implement the scale-Λ effective theory inspired by our autonomy-of-scales, (4) becomes

$$\mathscr{M} = \int_{\phi<\Lambda} D[\phi]\,\{\ldots \mathscr{L} \ldots\} \tag{5}$$

where the subscript $\phi < \Lambda$ flags the cutoff we've imposed on the domain of integration. Evaluating (5) requires performing a path integral over field configurations of energies up to—but not beyond—the cutoff energy Λ. Issuing amplitudes for scattering experiments conducted at energies up to (but not above) Λ, (5) is a recipe for surveying fields at scale Λ in order to derive predictions concerning phenomena at that scale. Comparing (4) to (5) makes manifest that *qft lagrangians do not describe phenomena on their own,*

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

but only in conjunction with a range of field energies over which they're integrated to yield amplitudes.

We are ready to embark on an RG analysis of the space $\mathscr{G}$ of lagrangians.[10] Our first move is to pick a very high energy Λ_{uv}. Take a theory defined by a lagrangian $\mathscr{L}$ and the UV cutoff energy Λ_{uv}. We're being studiously ecumenical about what form UV physics might take: even lagrangiansthat aren't renormalizable might play the role of $\mathscr{L}$.

The theory $(\mathscr{L}, \Lambda_{uv})$ predicts values for low (scale Λ) energy amplitudes. So augment our autonomy-of-scales hope: posit a scale Λ effective theory $(\mathscr{L}', \Lambda)$ that reproduces the low-energy amplitudes $(\mathscr{L}, \Lambda_{uv})$ predicts:

$$\int_{\phi<\Lambda_{uv}} D[\phi]\,\{\ldots\mathscr{L}\ldots\} =_{<\Lambda} \int_{\phi<\Lambda} D[\phi]\,\{\ldots\mathscr{L}'\ldots\} \tag{6}$$

where $=_{<\Lambda}$ denotes agreement about low energy amplitudes.

RG equations identify theories $(\mathscr{L}, \Lambda_{uv})$ and $(\mathscr{L}', \Lambda)$ that satisfy the hopes just expressed. Given a high-energy (scale Λ_{uv}) lagrangian $\mathscr{L}$, RG equations identify a low-energy (scale Λ) lagrangian $\mathscr{L}'$ that mimics $\mathscr{L}$'s low-energy empirical implications—where (6) explicates the mimicry relations. RG equations induce a map $R_{\Lambda\Lambda_{uv}}$ on $\mathscr{G}$: $\mathscr{L}' = R_{\Lambda\Lambda_{uv}}\mathscr{L}$. It's important to keep the a pair of energy scales indexing the RG map in view: although the map acts on the space $\mathscr{G}$ of lagrangians, it's establishing an association (explicated by (6)) between theories $(\mathscr{L}, \Lambda_{uv})$ and $(\mathscr{L}', \Lambda)$.

RG equations are devised by starting with the high energy theory $(\mathscr{L}, \Lambda_{uv})$ and "integrating out" field configurations with energies between Λ_{uv} and the new cutoff Λ—roughly speaking, blurring $(\mathscr{L}, \Lambda_{uv})$'s high-energy content by replacing individual high energy field configurations with averages over the collection of high-energy fields—to identify the low-energy theory that reproduces $(\mathscr{L}, \Lambda_{uv})$ low-energy amplitudes.[11] Because low-energy amplitudes are approximated perturbatively, so are RG equations: a first-order RG equation identifies low- and high-energy theories whose low-energy amplitudes, calculated to tree order, agree; a second- or loop order RG equation identifies theories whose low-energy amplitudes, calculated to loop order, agree, and so on.

In the space $\mathscr{G}$, a lagrangian is coded by the coupling vector giving the coefficients of its kinetic, mass, and interaction terms. The central dogma of the RG approach is that "changing [the cutoff scale Λ] and demanding the the physics be the same (since Λ is arbitrary) means the the couplings in the theory …must depend on Λ" (Schwartz 2014, 418). Poetically put, RG equations describe the running of the couplings. More prosaically, for each coupling g, they determine a function $g(\Lambda)$ that tracks how that coupling evolves as the cutoff scale is lowered from Λ_{uv} to Λ.[12]

So how do the couplings run? A celebrated analysis of RG flows in the space of scalar field theories due to Polchinski supports a striking answer. Fix a UV cutoff scale Λ_{uv}. Pick a lagrangian $\mathscr{L}$ and start the RG flow. At first the news is bad. "Although the [cutoff scale Λ_{uv}] lagrangian might start with a simple form, at lower scales it

[10]The variety of RG analysis sketched here is a Wilsonian RG analysis. There are other varieties, with the conceptual distinctions between them non-trivial. See (Fraser, ms) for a splendid discussion.

[11]See (Peskin and Schroeder 1995, §12.1) for a less impressionistic account.

[12]Not essential to the story told is the use of *beta functions* to describe the running of the couplings under the RG. A coupling's beta function gives its infinitesimal variation with respect to the logarithm of the cutoff scale Λ.

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

becomes complicated" (1984, 277), as interactions of all sorts, not just renormalizable ones, scramble to encapsulate the empirical implications of high energy modes integrated out of $(\mathscr{L}, \Lambda_{uv})$. But good news is coming: "At scales far below Λ_{uv} …a great simplification will occur. …no matter what initial lagrangian we start with (within limits), the lagrangian will be strongly attracted toward a three-dimensional submanifold in the infinite-dimensional space of possible lagrangians" (1984, 277). For $\Lambda << \Lambda_{uv}$, $\mathscr{L}' = R_{\Lambda\Lambda_{uv}}\mathscr{L}$ will be well-approximated by a lagrangian featuring just three terms.[13] Additional terms in the high-energy lagrangian $\mathscr{L}$ are suppressed—multiplied by factors of $\frac{\Lambda}{\Lambda_{uv}}$—by the RG flow. Although they won't be strictly speaking absent from $\mathscr{L}'$, these additional terms have so little empirical impact that we might as well write them in invisible ink. As far as effective physics within experimental tolerances is concerned, it'll be described by a lagrangian in a three-dimensional subspace of $\mathscr{G}$.

What's more, according to Polchinski's analysis, *the exponentially suppressed terms are exactly the interactions declared non-renormalizable by Dyson's criterion*. In other words, to all intents and purposes, lagrangians participating in effective physics reside in $\mathscr{G}^{PR}$. "In light of the above RG analysis," Williams remarks, "it no longer seems like incredible luck that so many useful QFT [lagrangians] are renormalizable" (2023, 54).

And that's the RG Explanation. Useful physics is low-energy effective physics, and the RG flow suppresses the contributions of non-renormalizable interactions to low-energy effective lagrangians. Polchinski tells us that this happy outcome "follows in a very general way when dimensional analysis is applied to the RG equation for an effective lagrangian" (1984, 274). Where g_i is a coupling vector specifying a scale Λ_{uv} lagrangian, RG analysis yields an account $g_i(\Lambda)$ of how each coupling runs as high-energy modes are integrated out of the high-energy theory. We expect that couplings are going to mix under the action of the RG—that after a RG transformation, the coupling for each interaction will typically be a function of the complete set of pre-transformation couplings. If the couplings $\{g_j(\Lambda)\}$ have different mass dimensions, what sense can we make of the mass dimension of their combination? An easy way to avert this awkwardness is to have the RG act on *dimensionless* couplings. "It is," after all, as Weinberg remarks, "convenient to work with dimensionless parameters" (1995, 525).

To facilitate this convenience, take a coupling g that modifies a lagrangian term whose mass dimension $d \neq 4$. For reasons rehearsed in §2, g must have mass dimension $4-d$. To transmute g to a dimensionless coupling $\tilde{g}$, multiply it by the appropriate power of the cutoff Λ (which, recall, has mass dimension 1) collaborating with the lagrangian to define the effective theory. Thus

$$\tilde{g} := (\Lambda)^{d-4} g \tag{7}$$

$\tilde{g}$ is a duly dimensionless coupling whose running will be amenable to RG analysis. Because the expedient $(\Lambda)^{d-4}$ plays a major role in what follows, I'll give it a name: the *dimension-cancelling pre-factor*.

For a brutalist approximation to how $\tilde{g}$ runs under the RG, suppose that the "physical" coupling g itself carries no scale-dependence. Then only the dimension-cancelling prefactor varies as the cutoff scale changes—assuming, that is, that the Λ in the dimension-cancelling prefactor slides along with the scale at which integrals like (5)

[13]Supposing $\mathscr{L}$ represents only a small perturbation to the gaussian fixed point lagrangian whose only non-zero term is its kinetic term. This restriction lies behind Polchinki's "within reason" qualification. See (Wallace 2019) for a discussion of the restriction.

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

are cut off. (This is an assumption we'll revisit.) It follows that

$$\frac{\partial \tilde{g}}{\partial \Lambda} = g\frac{\partial(\Lambda^{d-4})}{\partial \Lambda} \qquad (\text{NAIVE}) \tag{8}$$

Ignoring the fundamental moral of RG approaches, that physical couplings change as scales change, (NAIVE) gives a first-order approximation to the RG equations.[14] Higher-order approximations to the RG equations, approximations sensitive to the details of how qft lagrangians underwrite predictions at higher orders of perturbation theory, could afflict corrections on (NAIVE).

But let this not distract us from the story (NAIVE) has to tell! According to (NAIVE), if $\tilde{g}(\Lambda_{uv}) = (\Lambda_{uv})^{d-4}g$ is a coupling's value in a UV cutoff lagrangian $\mathscr{L}$, and we lower the cutoff from Λ_{uv} to Λ, then the Λ scale coupling $\tilde{g}(\Lambda)$ becomes $(\Lambda)^{d-4}g$. But then

$$\tilde{g}(\Lambda) = (\frac{\Lambda}{\Lambda_{uv}})^{d-4}\tilde{g}(\Lambda_{uv}) \tag{9}$$

d is the dimension of the interaction $\tilde{g}$ modifies. Dyson's criterion tells us that if $d > 4$, that interaction is non-renormalizable. (9) tells us that if $d > 4$, the RG flow suppresses—multiplies by positive powers of $\frac{\Lambda}{\Lambda_{uv}}$, where Λ_{uv} is *a lot* bigger than Λ—that interaction's coupling. So (9) is the happy result we're after. (9) attests that under the RG flow toward the infrared, couplings of non-renormalizable interactions are dramatically suppressed.

They're dramatically suppressed, at least, according to the (NAIVE) approximation to the RG equations. It could be that higher-order approximations complicate this pleasing upshot, by revealing $\tilde{g}$—and by extension, the non-renormalizable interaction it modifies—to matter more to low-energy physics than (9) allows. This menace, that the higher-order corrections "overwhelm" (1984, 273) (NAIVE)'s account of the RG flow, is ruled out by Polchinski's analysis. Approximating RG equations to all orders, he shows that (NAIVE) gets the basic story basically right. The brutalist approximation is a good guide to how couplings run under the RG flow. And that approximation tells us what we want to hear: couplings of non-renormalizable interactions are dramatically suppressed.

4. Some Dimensions of the Problem

This criminally brief concluding section asks: how happy should we be with the RG Explanation? What's doing almost all of the work in the version of the explanation just sketched is the "convenient" dimension-cancelling pre-factor (7) introduces to ensure that all the couplings, whose running we treat via RG methods, are dimensionless. The pre-factor condemns non-renormalizable interactions to suppression under the RG flow. And that condemnation resolves the Puzzle. But can the dimension-cancelling prefactor really bear the explanatory weight the RG Explanation places on it?

Here, in rapid-fire, are a few reasons to think the answer isn't obvious.

Start with a difference between the instances of dimensional analysis that lead to (i) the verdict that interactions of mass dimension d must be accompanied in a lagrangian by couplings of mass dimension $(4 - d)$, and (ii) the dimension-cancelling

[14] Impressionistically: at tree order, loop integrations over high energy fields don't figure in the perturbative calculation, relieving g of the obligation to adapt to offset changes in the domain of integration.

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

pre-factor in (7). For the first verdict, a nomic commodity—the action $S := \int d^4x\mathscr{L}$—guides the conclusion reached by dimensional analysis. That analysis reveals how couplings must behave in order for the action, and the swathe of physical theorizing based on the action, to make sense. For verdict (ii), by contrast, something arguably less nomic guides dimensional analysis—the maneuver, of using dimension-cancelling pre-factors, to facilitate easy management of functions describing the running of the couplings. Whereas in case (i), dimensional analysis reveals how couplings must behave for the sake of nomic compliance, in case (ii), dimensional analysis reveals how dimension-cancelling pre-factors must behave for the sake of ... convenience. It's not obvious that we should trust conclusions reached through dimensional analysis guided by considerations of convenience as much as we trust those guided by nomic considerations.

But grant, for the sake of argument, that functions describing how couplings run should act on dimensionless couplings.[15] Why use powers of a *sliding* cutoff Λ as the dimension-cancelling pre-factor? The result (NAIVE) anchoring the RG explanation follows only if the dimension-cancelling pre-factor varies with scale—indeed (NAIVE) is an account of *how* the dimension-cancelling pre-factor varies with scale! But a fixed cutoff— Λ_{uv}, for instance—would serve just as well in a pre-factor that cancels $g(\Lambda)$'s dimension as Λ does. It's the mass dimension of the pre-factor, not its value, that effects the cancellation. The problem, for the RG Explanation, is that if we use a fixed value of the cutoff in the dimension-cancelling pre-factor, dimensionless couplings inherit *no* cutoff-dependence from that pre-factor, eradicating the exponential suppression that's the crux of the explanation.

Of course, if we use a fixed value of the cutoff in dimensional-cancelling pre-factors, it would be nice to have a reason for the value we choose. There's one choice for which a reason comes ready-made with RG approaches. That choice is Λ_{uv}. Standard accounts of RG approaches cast Λ_{uv} as the energy at which lagrangian field theories break down. The only "natural" energy in sight, Λ_{uv} is like (or may even be) the Planck length, which physicists are accustomed to wield in the cause of dimensional analyses in other contexts.[16] Given that we can define dimensionless couplings using a fixed cutoff, and that it's hardly anathema to RG approaches to identify Λ_{uv} as the fixed cutoff, we needn't accept the result (NAIVE) that secures the RG explanation.

There is an even shorter route to the suppression of non-renormalizable interactions than the one §3 charted. Forget about dimensionless couplings. Adopt another perspective on (6). We've been using (6) to answer the question: which low-energy theory

[15] The concession incorporates another one. Granting that beta functions should act on couplings of the same dimension, why should that dimension be 0? Note that if it's different from 0, the pleasing alignment—between non-renormalizable interactions and interactions whose couplings are exponentially suppressed by the RG flow —breaks down. I've conceded that that dimension should be 0. Parts of the story this essay leaves off-stage—including the pivotal role of the gaussian fixed point in the sorts of RG analysis discussed here—make this concession easy.

[16] A further consideration: Qft calculations are customarily prosecuted using Euclidean path integrals, which replace the time coordinate t with an imaginary time coordinate it. This Wick rotation simplifies calculations by transmogrifying unwieldy oscillating exponentials e^{-it} into tractable damped ones e^{t}. It also transforms a qft with cutoff Λ_{uv} to a statistical mechanical theory of a lattice system with a finite lattice spacing related to Λ_{uv}. And the lattice spacing of a physical lattice is plausibly a physically relevant parameter. Whether this plausibility transfers to Λ_{uv} understood as a cutoff for a real-time qft is very non-trivial additional question. See (Fraser 2020) for more.

https://doi.org/10.1017/psa.2025.10132 Published online by Cambridge University Press

$(\mathscr{L}', \Lambda)$ mimics the low-energy amplitudes defined by a high-energy theory $(\mathscr{L}, \Lambda_{uv})$? This is a question about theories with different (but overlapping!) domains. But let's ask a different question, one about theories with the same domain: *how do $\mathscr{L}$ and $\mathscr{L}'$, both understood as governing physics upto the UV cutoff Λ_{uv}, compare?* In condensed matter physics, asking the counterpart of this question has huge payoffs: if the answer is, "they're the same!," you've found a fixed point of the RG flow and are well on your way to calculating critical exponents and explaining universality.

To compare $\mathscr{L}$ and $\mathscr{L}'$ as lagrangians for the same, high-energy domain, we need to *rescale*. $\mathscr{L}'$ as it appears in (6) governs physics only upto a scale Λ. To broaden its coverage so that its domain coincides with the domain of $\mathscr{L}$, multiply all the couplings that appear on the r.h.s. of (6) with mass (=energy) dimension d by a scale factor $(\frac{\Lambda_{uv}}{\Lambda})^d$. After this rescaling, $\mathscr{L}'_{rescaled}$ has the same domain of validity as $\mathscr{L}$. It's $\mathscr{L}'_{rescaled}$ we should consider alongside $\mathscr{L}$ if our aim is to compare lagrangians with a common domain. But soft! If $\mathscr{L}$ includes an interaction of mass dimension $d > 4$—a non-renormalizable interaction—that interaction will be modified by a coupling of *negative* mass dimension. This coupling will appear in $\mathscr{L}'_{rescaled}$ multiplied by a *positive* powers of $\frac{\Lambda}{\Lambda_{uv}}$. This is the exponential suppression key to the RG explanation—here exhibited as a consequence of the rescaling step of an RG transformation.[17]

But why in the context of thinking, fueled by the autonomy-of-scales hopes, about effective lagrangians, should we take the rescaling step? The question—which low-energy theory $(\mathscr{L}', \Lambda)$ mimics the low-energy amplitudes defined by a high-energy theory $(\mathscr{L}, \Lambda_{uv})$?—is the right one to ask in this context. The adjacent question—what high energy theory would $\mathscr{L}'$ be, if $\mathscr{L}'$ were a high energy theory?—isn't. It's $\mathscr{L}'$ as participant in the effective theory $(\mathscr{L}', \Lambda)$ we care about. And how we've expressed that care makes rescaling $\mathscr{L}'$ to extend its domain up to the ultra-high energy cutoff Λ_{uv} weird. If physics breaks down at Λ_{uv}, what (after rescaling) are we to make of the high-energy modes we (before rescaling) integrated out of $\mathscr{L}$ to obtain $\mathscr{L}'$? This is hardly to deny what the previous paragraph acknowledged: that there are plenty of contexts where it's illuminating to pursue the question about how $\mathscr{L}$ and $\mathscr{L}'$, considered as theories with the same domain, compare. It is, however, to suggest that the present context might not be one of them. And if it isn't one of them, the RG Explanation that emerges from the pursuit rings hollow.

The foregoing are reasons, internal to the RG Explanation, to question its bona fides. I'll close with an external reason. It's that the RG Explanation explains too much. The RG explanation primes us to expect effective theories to be renormalizable. But paradigm examples of effective theories—the 4-Fermi theory of beta decay (which implements the electroweak theory at low energies), the chiral lagrangian (which implements QCD at low energies)— aren't renormalizable. This is adamantly not to deny that such theories play important roles, rewarding of methodological and conceptual analysis, in working physics, or that a broad effective theory philosophy helps to make sense of those roles. It is rather to observe that non-renormalizable effective theories sit awkwardly alongside a very specific line of thought about effective theories: the line of

[17] There is collateral evidence for (NAIVE) that takes the following form: we can characterize the WRG flow directly, by calculating what effective lagrangian results when we integrate higher energy modes out of the higher energy theory. (NAIVE) results. The exhibitions of such evidence I'm aware of (e.g. Melo 2019, §3) rely on the rescaling described here.

thought traced by the RG Explanation. As sketchy as reservations I've raised about that explanation are, perhaps this awkwardness is a reason to take them seriously.

References

[1] Blum, Alexander (forthcoming). *The decline and fall of QED*. Cambridge: Cambridge University Press.

[2] Feynman, Richard (1995). *QED: The strange theory of light and matter*. Princeton: Princeton University Press.

[3] Fraser, Doreen (2020). "The development of renormalization group methods for particle physics: Formal analogies between classical statistical mechanics and quantum field theory." *Synthese 197*: 3027-3063. https://doi-org.proxy.lib.umich.edu/10.1007/s11229-018-1862-0

[4] Fraser, James Duncan (2016). "What is Quantum Field Theory? Idealisation, Explanation and Realism in High Energy Physics." Ph D Diss., University of Leeds.

[5] Fraser, James D. (2018). "Renormalization and the formulation of scientific realism." *Philosophy of Science 85*): 1164-1175. https://doi.org/10.1086/699722

[6] Fraser, James D. (2020a). "Toward a realist view of quantum field theory." In *Scientific Realism and the Quantum*, ed. Steven French and Juha Saatsi, 276-292. Oxford: Oxford University Press.

[7] Fraser, James D. (2020b). "The real problem with perturbative quantum field theory." *The British Journal for the Philosophy of Science 71*: 391-413. https://doi-org.proxy.lib.umich.edu/10.1093/bjps/axx042

[8] Fraser, James D. (ms). "Classifying Renormalization Groups."

[9] Hall, Brian (2013). *Quantum mechanics for mathematicians*. Berlin: Springer-Verlag.

[10] Melo, Joao (2019). "Introduction to renormalisation." arXiv:1919.11099.

[11] Miller, Michael (2021)."Infrared cancellation and measurement." *Philosophy of Science 88*: 1125 - 1136. https://doi.org/10.1086/714707

[12] Peskin, Michael and Schroeder, Daniel (1995). *An Introduction to Quantum Field Theory*. Boulder, CO: Addison Wesley.

[13] Polchinski, Joseph (1984). "Renormalization and effective Lagrangians." *Nuclear Physics B 231*: 269-295. https://doi.org/10.1016/0550-3213(84)90287-6

[14] Schwartz, Matthew D. (2014). *Quantum field theory and the standard model*. Cambridge: Cambridge University Press.

[15] Schweber, Silvan S. (2020). *QED and the men who made it: Dyson, Feynman, Schwinger, and Tomonaga*. Princeton: Princeton University Press.

[16] Wald, Robert M. (1994). *Quantum field theory in curved spacetime and black hole thermodynamics*. Chicago: University of Chicago Press.

[17] Wallace, David (2019). "Naturalness and emergence." *The Monist 102*: 499-524. https://doi-org.proxy.lib.umich.edu/10.1093/monist/onz022

[18] Weinberg, Steven (1977). "The search for unity: Notes for a history of quantum field theory." *Daedalus* : 17-35. http://www.jstor.org/stable/20024506

[19] Weinberg, Steven (1995). *The Quantum Theory of Fields*, vol. I. Cambridge: Cambridge University Press.

[20] Williams, Porter (2019). "Scientific realism made effective." *The British Journal for the Philosophy of Science 70*: 209-237. https://doi-org.proxy.lib.umich.edu/10.1093/bjps/axx043

[21] Williams, Porter (2021). "Renormalization Group Methods." In *The Routledge Companion to Philosophy of Physics*, ed. Eleanor Knox and Alistair Wilson, 296-310. New York: Routledge.

[22] Williams, Porter (2023). *The Philosophy of Particle Physics*. Cambridge: Cambridge University Press.

Acknowledgements.

For feedback on earlier versions, I am grateful to Dave Baker, Gordon Belot, and John Earman.