

Effective Theories and Infinite Idealizations: A Challenge for Scientific Realism

Sébastien Rivat*

September 3, 2020

Abstract

Williams and J. Fraser have recently argued that effective field theory methods enable scientific realists to make more reliable ontological commitments in Quantum Field Theory (QFT) than those commonly made. In this paper, I show that the interpretative relevance of these methods extends beyond the specific context of QFT by identifying common structural features shared by effective theories across physics. In particular, I argue that effective theories are best characterized by the fact that they contain intrinsic empirical limitations, and I extract from their structure one central interpretative constraint for making more reliable ontological commitments in different subfields of physics. While this is in principle good news, this constraint still raises a challenge for scientific realists in some contexts, and I bring the point home by focusing on Williams's and J. Fraser's defense of selective realism in QFT.

1 Introduction

There is a deeply entrenched strategy in philosophy of physics about how to interpret our best theories in realist terms. Philosophers usually start by pretending that the theory at stake is complete, true and final, even if it is known not to be true in all respects. Then, they eliminate its redundant parts by implementing sophisticated constraints on its structure. And eventually, they draw from the resulting theory some putatively complete picture of the world. The goal, ultimately, is to identify a definite set of unobservable entities or structures, whether they are fundamental or not, and thereby lay the ground for explaining the success of the theory in realist terms.¹

As it turns out, this strategy somewhat falls apart in the case of our most fundamental and empirically successful theories. We do not yet know whether realistic Quantum Field Theories (QFTs) can be consistently defined across all scales and therefore whether we can even consistently speculate about the possible worlds in which these theories are exactly true. Wallace (2006, esp. sec. 3.3; 2011), Williams (2019b), and J. Fraser (2018; 2020) have proposed a more modest and cautious strategy in response, which is also better suited to the limited success of current and past theories. They enjoin philosophers to

*Max Planck Institute for the History of Science, Berlin. Email: srivat@mpiwg-berlin.mpg.de

¹For a critical discussion of this traditional interpretative strategy, including references in the literature, see Ruetsche (2011, chap. 1) and Williams (2019b).

identify the ontological commitments necessary to explain the success of our best QFTs in the limited regimes where they are known to be reliable and not in the regimes where they are likely to break down.

The crucial part of Wallace, Williams and J. Fraser’s proposal resides in the set of techniques they employ to implement this new strategy, namely, Effective Field Theory (EFT) methods (including the Wilsonian Renormalization Group). Broadly speaking, these methods have been developed in QFT to treat phenomena at different scales separately, and they became popular in physics in large part because of their remarkable heuristic, computational and predictive power. More crucially for interpreters, the QFTs constructed by using these methods, i.e., EFTs, are intrinsically restricted to some limited range of distance scales. The physics within this range can even be shown in typical cases to be largely independent of the specific details of the short-distance physics. And this has led Williams and J. Fraser, in particular, to argue that EFTs provide a more perspicuous and reliable interpretative standpoint to identify unobservable entities or structures in the appropriate regimes, even if realistic QFTs are ultimately shown to be consistent across all scales.

This paper has two closely related aims. The first is to show that the interpretative relevance of EFT methods extends beyond the specific context of QFT. Given that most if not all known physical systems exhibit distinct scales in most circumstances, it should come as no surprise that the EFT paradigm has been successfully implemented in most areas of contemporary physics during the last decades.² Yet, we might still wonder whether the theories constructed by using EFT methods share distinctive structural features that might help us make more reliable ontological commitments in different subfields of physics. I will first argue that effective theories are best characterized in general by the fact that they contain intrinsic empirical limitations, i.e., their structure incorporates a robust specification of the scales at which they are likely to be empirically inaccurate before we probe these scales in experiments. This contrasts with the usual situation where the empirical limitations of a theory are found only by a direct confrontation with experimental data obtained at the relevant scale. Then, I will briefly present and justify the realist account of effective theories which follows the most naturally from this characterization. I will call it the “Standard Effective Account” and show that the structure of an effective theory forces us to restrict our commitments to entities or structures which can be specified within the limited range where the theory is likely to remain empirically reliable.

The second aim is to assess whether Wallace, Williams and J. Fraser’s strategy enables scientific realists to fulfill their explanatory duties. Starting with the traditional form of scientific realism (cf. Psillos, 1999, pp. xvii-xix), I will first give a concrete example of the restrictions we face if we treat our best current theories as effective theories.³ We may think, for instance, that we

²For references to the extension of EFT methods outside condensed matter and particle physics, see, e.g., Endlich et al. (2011), Dubovsky et al. (2012), and Gripaos and Sutherland (2015) for fluid dynamics; Donoghue (1995) and Burgess (2004) for general relativity; Goldberger and Rothstein (2006) and Porto (2016) for post-Newtonian gravitation; Baumann and McAllister (2015, chap. 2) and Burgess (2017) for inflationary cosmology; Polchinski and Strominger (1991) and Hellerman et al. (2014) for low-energy string theories; Baumann and Green (2012) and Kaplan (2016, esp. sec. 8.4.3) for advanced topics relevant to quantum gravity.

³Of course, this requires assuming that we do not yet have some decisive evidence that

have good reasons to take the descriptions of continuum fields in the effective versions of the Standard Model of particle physics and General Relativity to be approximately true and therefore to commit to the existence of those entities, i.e., of continuous systems with an infinite number of degrees of freedom. I will argue that on the Standard Effective Account, we cannot reliably make such ontological commitments. And my point here is not so much to claim that infinite physical systems are beyond our ken—in a way, we have known this for a long time—but rather to illustrate how the structure of effective theories imposes clear-cut restrictions on one’s ontological commitments.

I will then argue that, in some specific theoretical contexts including classical and quantum field theory, these restrictions still raise a challenge for more refined forms of scientific realism. To bring the point home, I will focus on Williams’s (2019b) and J. Fraser’s (2018; 2020) defense of selective realism in QFT and, expanding on Ruetsche’s (2018; 2020) discussion, show that the candidates which look at first sight the most appealing for making ontological commitments in the appropriate regimes—namely, correlations, particles, and lattice fields—fail in other important respects. The best candidates that do not suffer from the same issues appear to be continuum fields, with the proviso that they are approximately similar to large distance scale features of the world. But, again, selective realists cannot take the descriptions of continuum fields to be approximately true *simpliciter*, which leaves them with no obvious candidate for offering a genuine defense of the realist cause. I will conclude briefly with a more radical suggestion to circumvent this issue: namely, to modify the standard semantic tenet of scientific realism endorsed by selective realists (e.g., Psillos, 1999; Chakravartty, 2007) and index (approximate) truth to physical scales.

The paper is organized as follows. Section 2 presents two distinct examples of effective theories. Section 3 argues on the basis of these examples that effective theories are best characterized by the fact that they contain intrinsic empirical limitations. Section 4 presents the Standard Effective Account. Section 5 shows that traditional scientific realists cannot, as a matter of principle, commit to the existence of the infinite systems specified by a literal interpretation of our best current effective theories. Section 6 extends the discussion to Williams’s and J. Fraser’s defense of selective realism.

2 Two Examples of Effective Theories

Philosophers have not paid much attention to the diversity of effective theories across physics (e.g., Cao and Schweber, 1993; Hartmann, 2001; Bain, 2013); and when they treat the particular case of EFTs in particle and condensed matter physics as a new paradigm for understanding physical theories, they often remain too elusive or attribute too much importance to parochial fea-

we have hit a true, final and complete theory in physics or some complete theory providing an approximately true description of the world in all respects. We also need to assume that effective theories display sufficiently many theoretical virtues to be even considered candidates for making approximately true claims about the world (see Wells, 2012, chap. 5, for a discussion related to this point). We do not need, however, to deny the existence of a final theory, which is implicit in the traditional scenario of an infinite “tower” of EFTs, where each theory of an endless series of EFTs describes phenomena within a limited range of energy scales.

tures absent in other types of effective theories. For instance, it is common to characterize effective theories as theories that directly incorporate into their mathematical structure the imprint of their breakdown at some non-trivial finite physical scale (e.g., Bain, 2013, p. 1; Williams, 2019a, p. 2; 2019b, pp. 6-7, 9-10, 13). But seldom is it specified whether, in the general case, effective theories display some mathematical singularity, become physically meaningless, make inconsistent predictions, or become merely empirically inaccurate at that scale.⁴ In order to give a sufficiently comprehensive and informative characterization, I will thus first present two different kinds of effective theories and examine, in particular, the way in which they “break down” at some scale.⁵

Example 1: Consider first the mathematically most simple formulation of the Newtonian gravitational theory for a body of mass m_1 interacting with another body of mass m_2 :

$$m_1 \frac{d^2 r}{dt^2} = -m_1 \frac{m_2 G}{r^2} \quad (1)$$

with r the relative distance between the centers of mass of the two bodies and G the Gravitational constant.

There are two distinct ways to construct an effective version of this theory. Since we already know its closest successor, i.e., classical General Relativity, we can simply follow the “top-down” strategy: namely, we appropriately restrict the range of parameters of the more comprehensive theory and eliminate its theoretical constituents which do not contribute significantly to predictions within this range. For instance, we can derive Eq. 1 with additional correction terms encoding relativistic effects by implementing weak-gravity and low-velocity restrictions on the simplest solutions to the equations of classical General Relativity (see, e.g., Poisson and Will, 2014, for more details).

We can also pretend that we do not yet know the more comprehensive theory and follow the “bottom-up” strategy. We first identify a limited range where we think that the theory provides reliable information. For instance, we may suspect from the infinite value of $m_1 m_2 G / r^2$ in the limit $r \rightarrow 0$ that Eq. 1 becomes mathematically inadequate for describing the gravitational interaction between arbitrarily small bodies moving arbitrarily close to one another. Or we may have already found that the theory makes slightly inaccurate predictions when the gravitational force $m_1 m_2 G / r^2$ becomes too strong. Then, we restrict the range of the theory by introducing some arbitrary limiting scale, namely, a short-distance scale r_0 in this case. And finally, we include all the possible terms depending on r_0 / r which are allowed by the symmetries of the theory, with one arbitrary coefficient for each new term. As we perform these steps, we do not need to know anything about the value or the underlying meaning of the limiting scale, namely, that r_0 turns out to be the Schwarzschild radius $2m_2 G / c^2$ of the body of mass m_2 , with c the speed of

⁴Other overly broad characterizations include “approximate theories” (e.g., Castellani, 2002, p. 263; Ruetsche, 2020, p. 298), “non-fundamental theories” (e.g., Egg et al., 2017, p. 455), and “phenomenological theories” (e.g., Huggett and Weingard, 1995, p. 189; Butterfield and Bouatta, 2014, p. 65).

⁵For simplicity, I will understand ‘theory’ in its specific sense throughout the paper, that is to say, as given by a specific action, a Lagrangian or a Hamiltonian—or even more simply by equations of motion.

light. The value of the additional coefficients and r_0 is ultimately determined by means of experimental inputs, at least for a finite number of them.⁶

Now, whether we follow the top-down or the bottom-up strategy, the resulting effective theory takes the following form:

$$m_1 \frac{d^2 r}{dt^2} = -m_1 \frac{m_2 G}{r^2} \left(1 + a_1 \frac{r_0}{r} + a_2 \left(\frac{r_0}{r} \right)^2 + a_3 \left(\frac{r_0}{r} \right)^3 + \dots \right) \quad (2)$$

with a_1, a_2, a_3 , etc. some arbitrary coefficients. The most complete version of Eq. 2 includes an infinite number of terms which depend on r_0/r and leave the equation invariant under Galilean symmetry transformations (i.e., translations in space and time, spatial rotations, and velocity boosts). We can also define an effective theory by means of a finite number of terms and fix the value of their coefficients by means of experiments.⁷

How should we interpret the scale r_0 if we take the structure of these effective theories at face value? Suppose for the sake of the argument that we are interested in predicting the value of the acceleration $d^2 r/dt^2$ in Eq. 2. The first thing to note is that the contributions of higher-order terms $(r_0/r)^n$ to predictions are negligible for $r \gg r_0$ and very large for $r \ll r_0$. If we include increasingly many higher-order terms in Eq. 2, the predictions remain overall the same for $r \gg r_0$ and become increasingly large around and below r_0 . And if we include an infinite number of terms, the resulting expansion $\sum_i a_i (r_0/r)^i$ takes an infinite value for $r_0/r \geq 1$. Hence, if we simply look at the mathematical structure of the family of effective theories associated with Eq. 2, we find that their predictions display a sharp pattern of variation around the characteristic scale r_0 , which remains robust as we add or remove higher-order terms.

At first sight, this predictive pattern does not appear to tell us much about r_0 since the expansions $\sum_i^N a_i (r_0/r)^i$ for finite N are mathematically well-defined across all distance scales (except for the trivial scale $r = 0$). Yet, if we consider these finite expansions in relation to one another, we learn that we can always add small correction terms of increasing order in r_0/r in any given expansion and adjust their coefficients if we want to improve its predictive accuracy for $r \gg r_0$. And if we consider these finite expansions in relation to the limiting case of the infinite expansion, we also learn that they ultimately become mathematically ill-defined at r_0 when we add increasingly many such terms. In short, if we try to make any of these finite expansions as predictively accurate as possible for $r \gg r_0$, we end up with theories making infinite predictions at r_0 and below, i.e., with theories which, as a matter of principle, cannot be empirically accurate for $0 < r \leq r_0$. And this, in turn, provides at least preliminary reasons to believe that the pattern of variation around r_0 does not simply reflect some notable qualitative physical change but rather signals that these finite expansions are likely to become unreliable around r_0 .

⁶In general, we also need to assume that the dimensionless constants of the theory are of order 1 to get a first estimate of the limiting scale, i.e., we need to endorse the “naturalness” principle $a_i = O(1)$ in Eq. 2 below.

⁷For more details about the first-order relativistic and quantum corrections to the non-relativistic gravitational potential, see, e.g., Donoghue (1995), Burgess (2004), and Blanchet (2014). Note that, in some cases, existing empirical measurements (or some other reason) may require us to break some of the symmetries of the original equation.

Now, this interpretation is grounded in the experimental profile of existing theories displaying the same predictive pattern. If, for simplicity, we use Eq. 2 as an example, the experimental pattern takes the following form. We start with some effective theory defined by means of a finite expansion and fix its parameters by means of experiments at large distance scales r . At shorter distance scales, however, we find small experimental discrepancies and decide to add new terms to compensate for them. Yet, as we probe even shorter distance scales, the effective theory with the additional terms becomes all the more quickly empirically inaccurate and we need, at least in principle, to introduce new terms if we want to maintain its predictive power and accuracy. In practice, physicists directly look for a new theory in situations like this. If we were to keep up with the original theory and probe phenomena closer and closer to r_0 , however, we would need to introduce an infinite number of terms. Since all these terms are equally important at r_0 , we would not be able to select a finite number of them in order to make approximate predictions. And since we cannot in practice make an infinite number of measurements to fix the value of an infinite number of arbitrary coefficients, the theory would lose its predictive power. Hence, according to this pattern, r_0 corresponds to the maximal predictive limit of the family of effective theories associated with Eq. 2. For the infinite expansion, r_0 corresponds both to a characteristic scale where the theory becomes mathematically ill-defined and predictively powerless. For the finite expansions, the demarcation is not as vivid and sharp; but, overall, the corresponding effective theories make empirically accurate predictions for $r \gg r_0$ and empirically inaccurate ones for $r \ll r_0$.

Note that the same argument does not apply to the original Newtonian theory in Eq. 1 despite its divergent behavior at $r = 0$. If we leave aside the apparent physical impossibility of the situation characterized by $r = 0$, we still face the issue that the limiting scale $r = 0$ is experimentally trivial from the perspective of classical Newtonian gravitation. Even if we can, in principle, probe the system down to arbitrarily short distances in this context, we can only gain experimental information about finite size effects resulting from the gravitational interaction between two bodies at some finite distance from one another. In the case of effective theories, the situation is different because there is no physical principle or experimental constraint which indicates that the regime specified by $r \leq r_0$ is either experimentally inaccessible or trivial. The infinite expansion becomes deficient at r_0 . But nothing in the theory suggests that we cannot use bodies to probe distance scales within $0 < r \leq r_0$ compared, say, to string theory where we cannot use strings in scattering processes to probe distances shorter than the string scale (see, e.g., Hossenfelder, 2013, sec. 3.2, for a discussion).⁸

Example 2: Consider now a standard example of QFT, the ϕ^4 -theory. The theory describes a simple quantum field, i.e., a continuum of smoothly coupled individual quantum systems over space-time with each system characterized by only one degree of freedom. In a somewhat analogous way as in Eq. 1, the original dynamical equation is given by:

$$\partial_\mu \partial^\mu \phi(x) + m^2 \phi(x) = -\lambda \phi^3(x) \quad (3)$$

⁸Note, moreover, that we cannot define some non-trivial limiting distance scale r_0 by using only m_1 , m_2 , and G in Eq. 1. By dimensional analysis, we would need to introduce a new arbitrary velocity scale c and therefore modify the structure of the original theory.

where $\phi(x)$ is a real-valued variable describing a possible configuration of the field over space-time, ∂_μ the analog of d/dt in the four-dimensional Minkowski space-time, m a mass parameter, and λ a self-coupling parameter. This equation contains no explicit intrinsic limitation, which suggests that there is *a priori* no reason to believe that the theory fails to apply at arbitrarily large and short distances (or, equivalently, at arbitrarily low and high energies). The trouble comes when we try to compute predictions. Typically, in QFT, this is done by evaluating the correlations between some initial and final field configuration states characterizing some scattering process, where these states describe, roughly speaking, the particles that we prepare and detect in experiments. Calculating these correlations requires, in turn, including the contributions from all the possible transitions between these states and therefore summing over all the possible intermediary field configuration states. If we do that, however, the high-energy configurations of the field, i.e., the configurations which vary rapidly over short-distance scales, give rise to infinite probabilistic predictions, which is inconsistent.

As of today, the only way to solve this issue in realistic QFTs is to modify the structure of the theory by means of “renormalization” methods.⁹ In the case of the ϕ^4 -theory, for instance, we can smoothly lower the contributions of the high-energy field configurations $\tilde{\phi}(k)$ over some high-energy cut-off Λ by using a new field variable $\phi_\Lambda(x)$ with exponentially decreasing contributions above Λ :

$$\phi_\Lambda(x) \propto \int d^4k e^{ikx} (e^{-k/\Lambda} \tilde{\phi}(k)) \quad (4)$$

Similarly to Example 1, the value of the limiting scale Λ is not fixed at this stage. Yet, the QFT case is special. If we keep a finite cut-off, we can make the predictions of the theory Λ -independent by absorbing Λ -dependent terms into its parameters, at least for a finite range of values of Λ . But this requires including all the possible interaction terms allowed by the symmetries of the theory:

$$\partial_\mu \partial^\mu \phi_\Lambda(x) + m^2(\Lambda) \phi_\Lambda(x) = -\lambda(\Lambda) \phi_\Lambda^3(x) - g_5(\Lambda) \phi_\Lambda^5(x) - g_7(\Lambda) \phi_\Lambda^7(x) - \dots \quad (5)$$

where the g_i 's are new arbitrary coupling parameters depending on Λ . If we have appropriate experimental inputs, we can define an effective theory by means of a finite number of interaction terms, fix their parameters, and estimate the value of Λ (as in Example 1).

The predictive pattern in this example is overall similar to the one displayed in the previous example. Once we fix the parameters of the theory, we can show that the higher-order interaction terms $g_i(\Lambda) \phi_\Lambda^i$ in Eq. 5 contribute to predictions by increasing powers of (E/Λ) , with E the characteristic energy scale of the scattering process considered. Yet, there is one crucial difference: the predictions of the theory typically become inconsistent for energies E close to and above Λ whether we include a finite or an infinite number of interaction terms in Eq. 5. Hence, if we take the structure of effective QFTs at face value, Λ is naturally interpreted as the scale at which the theory is likely to make inconsistent and *a fortiori* empirically inaccurate predictions.¹⁰

⁹See, e.g., Butterfield and Bouatta (2015), Williams (2019a), and Rivat (2019) for introductory discussions.

¹⁰As it turns out, the ϕ^4 -theory is even more special: the perturbatively renormalized

3 What is an Effective Theory?

Now that we are equipped with two different examples, let us look at some options for characterizing what is so distinctive about effective theories. I will argue that the structure of an effective theory is best characterized by the fact that it incorporates a robust specification of the scales at which it is likely to be empirically inaccurate (assuming, in particular, that we have appropriate experimental inputs at some other scales to fix its free parameters).

Characterization 1: A first option is to characterize an effective theory as a low-energy limit of a more complete theory—even if this more complete theory is not fully known, which means that an effective theory is a particular realization of a given theory over a restricted range of energy scales. This relational characterization fits well with high-energy physicists’ general description of EFTs (e.g., Burgess and Moore, 2006, p. xi; p. 456) and with the top-down Wilsonian procedure for deriving an EFT by eliminating high-energy field configurations.

To give a concrete example, suppose that the ϕ^4 -theory is a low-energy realization of a more complete theory including a light scalar field $\phi(x)$ of mass m and a heavy scalar field $\psi(x)$ of mass M , with $m \ll M$. We can derive effective theories as follows. First, we eliminate, or “integrate out”, the heavy field variable $\psi(x)$ in the high-energy theory (or, more precisely, in its functional path integral \mathcal{Z}). This gives rise to exotic terms depending on the variable $\phi(x)$ such as $\phi(x)(-\partial_\mu\partial^\mu + M^2)^{-1}\phi(x)$. Assuming that the characteristic energy E of the scattering processes of interest is much smaller than the mass of the heavy field, i.e., $E \ll M$, we can expand these exotic terms into an infinite series of polynomial terms depending only on the variable $\phi(x)$, its derivatives, and some inverse power of M . Schematically,

$$\begin{aligned} \mathcal{Z} &= \int \mathcal{D}[\phi]\mathcal{D}[\psi] e^{i \int d^4x \left[\frac{1}{2}(\partial_\mu\phi)^2 - \frac{m^2}{2}\phi^2 - \frac{\lambda}{4!}\phi^4 + \frac{1}{2}(\partial_\mu\psi)^2 - \frac{M^2}{2}\psi^2 - \frac{g}{4}\phi^2\psi^2 \right]} \\ \implies \mathcal{Z} &= \int \mathcal{D}[\phi] e^{i \int d^4x \left[\frac{1}{2}(\partial_\mu\phi)^2 - \frac{m^2}{2}\phi^2 - \frac{\lambda}{4!}\phi^4 - \frac{g_6}{M^2}\phi^6 - \frac{g_8}{M^4}\phi^8 - \dots \right]} \end{aligned} \quad (6)$$

with the appropriate coupling parameters g and g_i .¹¹ The structure of the effective theory is fully specified by the restrictions imposed on the high-energy theory with the appropriate low-energy assumption. In particular, since the contributions of the interaction terms $(g_i/M^{i-4})\phi^i$ give rise to inconsistent predictions close to M , the high-energy theory provides a natural high-energy cut-off for renormalizing the effective theory, namely, the mass of the heavy field. We can also define effective theories by restricting the series to some finite order in $1/M$ and obtain the original ϕ^4 -theory by taking the limit $M \rightarrow \infty$.

The main issue with Characterization 1 is that it is either too broad or too narrow depending on how we understand it. If we take it to apply to any theory which is, in principle, derivable from a more complete theory in its low-energy limit, even indirectly, we may have reasons to suspect that it applies to all empirically successful theories built up so far. However, if we do not

coupling $\lambda(\Lambda)$ diverges at some finite high-energy scale, i.e., it displays a “Landau pole” singularity.

¹¹See, e.g., Baumann and McAllister (2015, sec. 2.1.1) and Petrov and Blechman (2016, sec. 4.1) for more details.

specify the structure of the high-energy theory or provide the specific details of the derivation, we will be left with a characterization which is overly vague and which, in particular, does not help us to circumscribe specific structural features common to effective theories. And to make the matter even worse, some standard cases of EFTs do not seem to have any high-energy completion and therefore to be even derivable, as a matter of principle, from a high-energy theory (see, e.g., Adams et al., 2006, for a discussion).

Inversely, if we take this characterization to apply only to theories which are explicitly related to a more comprehensive theory by means of some energy parameter or mass scale, as in Eq. 6, we will leave out many standard cases of effective theories, including Example 1. In general, the types of limiting scales and power counting schemes underlying the structure of effective theories, i.e., the rules for evaluating how the contributions to predictions of the different parts of an effective theory vary with some parameter, can be extremely diverse. Examples 1-2 illustrate this variety of scales. Example 1 provides a simple velocity power counting scheme when applied to a system of two bodies with the same mass m_1 and orbital radius r . In the non-relativistic regime, the virial theorem holds ($v^2 \sim Gm_1/r \sim r_0/r$), which means that the interaction terms in Eq. 2 contribute to predictions by increasing powers of the characteristic velocity v of the system. And it is more appropriate in this case to speak of a low-velocity realization of a more complete theory.

Characterization 2: A more promising strategy might be to look for some abstract feature internal to the mathematical structure of an effective theory.¹² Suppose for instance that we take an effective theory to be a theory which, while remaining mathematically well-defined over some limited range of parameters, becomes ill-defined at some non-trivial finite scale. This characterization fits well with the most complete versions of the effective theories presented in Examples 1-2 (e.g., $\sum_i a_i (r_0/r)^i$). It also fits well with the attitude sometimes expressed in the philosophical literature according to which the framework of EFTs provides a general, efficient, and “opportunistic” way of solving the mathematical issues of QFTs (see, e.g., Butterfield, 2014, sec. V.2.2; Butterfield and Bouatta, 2015, sec. 3.1.3). Indeed, the very idea of introducing and keeping a finite cut-off is vindicated by the pathological behavior of QFTs at high energies (cf. Example 2). And even if we attempt to cure QFTs of their mathematical difficulties with renormalization methods, some paradigmatic cases like the ϕ^4 -theory and Quantum Electrodynamics (QED), the quantum theory of the electromagnetic force, are likely to remain mathematically ill-defined at some large yet finite energy, i.e., to display a Landau pole singularity. If we want to define these pathological cases of QFTs consistently, they leave us with no choice but to restrict their range of parameters, and this suggests that EFT methods were meant to be applied to these sorts of theories.

Once again, however, this characterization excludes simple cases of effective

¹²Appealing to a particular mathematical structure does not seem to give an adequate trade-off between generality and informativeness. The closest we can probably get to Examples 1-2 and standard cases of effective theories is to characterize the structure of effective theories in terms of Taylor (or Laurent) series in some parameter (or truncations thereof). But even then, this solution excludes exotic cases of effective theories with non-polynomial interaction terms in the field variables (see, e.g., Gripaios, 2015, sec. 5, for some models including such terms).

theories and therefore appears to be too restrictive. For instance, the effective theories defined by means of a finite number of terms in Example 1 remain mathematically well-defined across all distance scales (except at the trivial scale $r = 0$) and therefore do not fall under Characterization 2. Agreed, being mathematically ill-defined at some non-trivial finite scale is presumably a sufficient condition for a theory to be characterized as effective (provided we introduce some cut-off); but these simple examples of classical point-particle effective theories show that this condition is not necessary.

Characterization 3: A third option, the one I favor, is to characterize effective theories by the fact that they contain intrinsic empirical limitations. Namely: an effective theory incorporates into its structure a robust specification of the ranges of scales where it is likely to be empirically inaccurate. There are four essential ingredients here:

1. The mathematical structure of the theory contains some non-trivial finite scale (“intrinsic limiting scale” or “cut-off”);
2. It is possible to include increasingly many terms depending on this limiting scale which are consistent with the core principles governing the structure of the theory, with one arbitrary coefficient for each new term introduced;
3. These terms are systematically organized according to the importance of their contributions to predictions below and above the limiting scale (“power counting scheme”);
4. As we include increasingly many such terms, the predictions derived from the theory remain approximately the same, say, below the limiting scale and become increasingly large around and above this scale (“robustness”).

The predictive pattern is well illustrated by Examples 1-2, although it does not essentially depend on the particular details of their mathematical formulation, and, in general, the interpretation in terms of intrinsic empirical limitations is grounded in the experimental profile of existing theories displaying the same predictive pattern. Note as well that Characterization 3 does not imply that the mathematical structure of an effective theory delineates *by itself* the scales at which its predictions are likely to break down. We usually need to have experimental inputs in some accessible regime and assume that the dimensionless constants of the theory are of order one if we want to estimate the value of the limiting scale. Similarly, adding a list of provisos of the form ‘For velocities much smaller than the speed of light’ or ‘ $r \gg r_0$ ’ in the preamble of the theory is not sufficient: Characterization 3 requires the theory to have the imprint of its probable predictive failure directly written in its mathematical structure.

Now, the advantage of this option is twofold. First, Characterization 3 is neither too restrictive nor too permissive. In particular, it applies to Examples 1-2 and standard cases of classical and quantum effective theories. It also excludes standard cases of theories putatively applicable across all scales such as the Newtonian theory defined in Eq. 1 and the perturbatively renormalizable version of Quantum Chromodynamics (QCD), the quantum theory of the

strong force.¹³ As explained in section 2, if we take such theories at face value, their structure does not explicitly delineate non-trivial experimental regimes where their predictions are likely to break down. Of course, we may impose a finite cut-off on the perturbatively renormalizable version of QCD because we suspect that QCD is likely to be empirically inaccurate at very high energies, and include higher-order interaction terms into the theory. We may also exploit the hierarchy of scales exhibited by the different masses of the quarks in QCD and define a low-energy theory of the light quarks u , d and s with some cut-off because we suspect that it is easier to compute low-energy predictions if we eliminate irrelevant high-energy degrees of freedom. In both cases, however, we will be dealing with a different kind of theory, strictly speaking: namely, an effective theory which falls under Characterization 3.

Second, the characterization is also informative. Most remarkably, it offers a sharp distinction between two kinds of theories (or models): (i) theories with intrinsic empirical limitations, i.e., which already contain in their structure information about where they are likely to make inaccurate predictions before we probe the relevant scales in experiments; and (ii) theories with extrinsic empirical limitations, i.e., which are found to make inaccurate predictions only by a direct confrontation with experimental data obtained at the relevant scale. As we will see in the next section, the structure of an effective theory also gives good reasons to believe that it provides reliable ontological guidance only within a limited part of the world.

4 The Standard Effective Account

So far, I have argued that effective theories are best characterized by the fact that they contain intrinsic empirical limitations, but I have not said anything yet about their representational achievements. Suppose then that some effective theory is found to make accurate predictions within some regime and that its predictions are likely to break down at some scale beyond this regime. The most straightforward realist explanation in this case is to take the theory to accurately represent a limited part of the world and misrepresent, or fail to represent, other parts. Since this explanation fits well with the set of commitments shared by philosophers who explicitly defend a realist interpretation of EFTs, I will be relatively brief in this section. I will clarify the idea that the domain of applicability of effective theories is intrinsically limited by means of four common claims made about EFTs, briefly justify them by relying on general features of effective theories, call the resulting account the “Standard Effective Account”, and extract one central interpretative constraint from it. This is, of course, not to say that these philosophers agree on everything. There are indeed substantive interpretative disagreements in the literature on EFTs. But I will ignore those disagreements and restrict myself to extending the four common claims beyond the context of QFT.

The first difficulty here is that the term ‘domain of applicability’ is ambiguous. We could arguably take it to refer to the universe of discourse or interpretation of the theory, to the set of phenomena accounted for by the theory, to the range of variables specifying the possible physical states of the

¹³This supposes that we set aside potential trouble at low energies and assume that the theory is sufficiently mathematically well-defined at arbitrarily high energies.

system described by the theory, or perhaps even to the range over which the theory is mathematically well-defined. If we keep in mind that the target of the theory is the actual world, the following notions should be sufficiently neutral and adequate for clarifying the Standard Effective Account. (i) The “domain of applicability” of a theory is the set of concrete physical objects—entities, structures, properties, quantities, states, phenomena, dispositions, and so on—that the theory accurately represents. The domain of applicability of a theory is not necessarily identical to its putative domain of applicability, i.e., to the set of putative physical objects specified by a literal interpretation of the theory.¹⁴ (ii) The “domain of empirical validity” of a theory is the range of physical parameters over which its predictions are likely to remain accurate. If we have good reasons to believe that we have found a final theory, this domain ranges over all physically possible scales. Otherwise, if we do not have any means to estimate the empirical limitations of the theory in advance as in the case of effective theories or any evidence that the theory will remain empirically accurate in new regimes, this domain reduces to the range over which the theory has been found to be empirically accurate.

Then, the Standard Effective Account can be spelled out in terms of the four following claims:

1. The domain restriction claim: The domain of applicability of an effective theory is restricted by the limits of its domain of empirical validity (cf., e.g., Cao and Schweber, 1993, p. 76; Castellani, 2002, p. 260; Wallace, 2006, sec. 3.2-3.3; Schweber, 2015, p. 60; J. Fraser 2018, p. 1173; Williams, 2019b, p. 13).

To take the simplest case of physical object, the domain restriction claim states that an effective theory accurately represents some concrete entity only if its core properties can be specified within the limited range where the theory is likely to remain empirically accurate. By ‘core property’ I mean that the property is constitutive of the identity of the entity (e.g., an infinite number of degrees of freedom for a continuum field). Now recall that if we have appropriate experimental inputs, say, at large distances, we can estimate the value of the limiting scale of an effective theory, say, a short-distance cut-off scale. And even if we have not yet probed phenomena close to this scale in experiments, the structure of the theory already gives us good reasons to believe that its predictions are inaccurate beyond this scale. As a realist, it is standard to assume that if a theory accurately represents the entities characterizing a specific domain, it also makes accurate predictions in this domain. Hence, the standard realist explanation of the probable predictive failure of an effective theory beyond its limiting scale is that the theory is likely to misrepresent, or fail to represent, the entities characterizing the corresponding domain (assuming here that there are such entities). And this means that the structure of an effective theory prevents us from remaining agnostic about its putative representational success beyond its limiting scale. We also have good reasons to think that the theory provides unreliable information about physical properties beyond this scale and therefore fails to give an accurate picture of the entities which are individuated by such properties.

¹⁴By ‘literal’ I mean that the physically meaningful descriptions of the theory are understood in their standard sense and taken to be either true or false.

In Example 2, for instance, the imposition of the smooth cut-off in Eq. 4 does not eliminate any degree of freedom in the original theory. On the face of it, then, the effective theory represents a putative continuum field with one degree of freedom at every point of space-time and therefore attributes core properties to its target system within any arbitrarily small region of space-time. At the same time, the pathological predictions of the theory around and beyond Λ also give very good reasons to believe that the theory misrepresents the structure of matter at arbitrarily short distances and therefore that it does not accurately represent a putative continuum field, strictly speaking. According to the domain restriction claim, however, it is perfectly possible for the theory to accurately represent, say, a real physical pattern of characteristic size larger than $1/\Lambda$ (see sections 5-6 for a discussion).

2. The new physics claim: The structure of an effective theory strongly suggests that the theory misrepresents or fails to represent some putative physical objects (cf., e.g., Robinson, 1992, p. 394; Cao and Schweber, 1993, p. 76; Wallace, 2006, sec. 3.2-3.3; J. Fraser 2018, p. 1173; Williams, 2019b).

This claim is best supported by examining the relation between successive effective theories, or even the relation between an effective theory and some putatively fundamental theory. If we take effective theories in isolation, however, we can still give some support to this claim by relying on their structure. Consider Example 2 again. The effective version of the ϕ^4 -theory with a smooth cut-off is mathematically well-defined at any point of space-time (at least according to physicists' standards) and does not contain any physical principle or constraint implying that the range beyond Λ is physically forbidden. To take again the simplest case of physical object, the theory thus appears to allow for the existence of concrete entities at arbitrarily short distances. Yet, as already emphasized, the theory also makes inconsistent predictions beyond Λ . Taken together, these two features strongly suggest that the theory is deficient in some way or another rather than that the world contains some physical limit at the scale Λ . And the best realist explanation, in this case, is that the theory does not include the appropriate theoretical constituents which would give rise to consistent predictions at short-distance scales and therefore that the theory either misrepresents or fails to represent putative entities at these scales instead of specifying, say, the fundamental graininess of space-time.¹⁵

3. The approximate truth claim: Effective theories offer approximately accurate representations in their domain of empirical validity (cf., e.g., Castellani, 2002, p. 260; J. Fraser 2018, p. 1173; Williams, 2019b, sec. 3).

The approximate truth claim states that an effective theory provides some accurate representations of unobservable physical objects specifiable within

¹⁵Note that the scale at which the predictions of an effective theory break down does not need to be *exactly* the same as the scale at which the new physics kicks in. For a discussion about the intricate link between violations of perturbative unitarity and the onset of new physics in the context of QFT, see, e.g., Aydemir et al. (2012) and Calmet and Casadio (2014).

the limited range where the theory is likely to remain empirically accurate—or, at least, that we can construct such representations by modifying the original structure of the theory.¹⁶ Again, the argument is relatively standard for the realist: (i) the best explanation for the predictive success of the theory within some regime is that the theory is approximately true; (ii) the probable predictive failure of the theory beyond its limiting scale gives good reasons to take only the descriptions below this scale to be approximately true. In Example 2, for instance, we should expect the descriptions of the dynamical properties of the field to be approximately true if they are restricted to scales lower than Λ . We can also impose limits at large distances by introducing a low-energy cut-off. And one way to construct a model satisfying this restricted set of descriptions is to replace the standard Minkowski space-time with a space-time lattice of finite extent (a sharp low-energy cut-off) and non-zero spacing (a sharp high-energy cut-off) and represent the quantum field in terms of a lattice field defined by assigning a variable $\phi(x)$ to each point of the space-time lattice. As we will see in section 5, the approximate truth claim does not mean that, in its standard formulation, an effective theory always accurately represents the putative objects specified by a literal interpretation of its core descriptions. And in section 6, we will see that the approximate truth claim sits in tension with other realist requirements in the context of QFT.

4. The stability claim: The representations of an effective theory specified within its domain of empirical validity are likely to remain approximately accurate under theory-change (cf., e.g., Cao and Schweber, 1993, sec. 4.1; sec. 4.3; Wallace, 2006, sec. 3.2-.3.3; J. Fraser 2018, sec. 3-4; Williams, 2019b, sec. 3).

Here the challenge is that a future higher-level or same-level theory might undermine the putative representational achievements of our best effective theories. As we will briefly see in section 6, Williams (2019b) and J. Fraser (2018; 2020) rely on the machinery of EFTs, including Wilsonian Renormalization Group (RG) methods, to defend the stability claim in the context of QFT. If we move outside of this context, we can still gain some support for this claim by focusing on the role of higher-order terms in effective theories.

Consider Example 1 and suppose that the predictions of the effective Newtonian theory with a few lowest-order terms are accurate at large distances $r \gg r_0$. If we discover a radically new and more comprehensive theory that makes slightly better predictions than the effective theory at large distances, we can always add higher-order terms to compensate for these empirical discrepancies. This move is, of course, largely *ad hoc*. But it shows that the higher-order terms can be used to encode the contributions of new physics at large distances according to their relevance and thus suggests that these terms do not simply correspond to arbitrary modifications of the theory, with no physical significance whatsoever. The ability of higher-order terms to stand for fine-grained features of new physics is also supported by explicit derivations of effective theories from more comprehensive ones (see, e.g., Eq. 6

¹⁶I will set aside issues related to the nature of scientific representation and use interchangeably “approximately accurate representation” and “approximately true description”, assuming that a description is approximately true relative to the actual world if it is satisfied by some model that provides an approximately accurate representation of some actual target system.

above). And, in general, the structure of an effective theory is such that we can parametrize the contributions of any type of new physics at large distances up to an arbitrarily high degree of precision by adding increasingly many terms depending only on the degrees of freedom of the original theory. In the Newtonian case, we can even include such terms by preserving *all* the core principles of the original theory (e.g., Galilean invariance).

Now, the crucial point is that the contributions of the higher-order terms become increasingly negligible at large distances $r \gg r_0$, no matter what the new physics looks like. And insofar as these higher-order terms are assumed to stand for fine-grained features of new physics, this shows that the descriptions of the effective theory which are relevant at large distances are largely insensitive to the particular details of the new physics. This new physics affects at most the value of the parameters of the lowest-order terms. At the scale r_0 , by contrast, the core principles of the effective theory do not even allow us to give an approximately true description of the dynamical behavior of the system and we have no choice but to look for a new theory.

Of course, this argument is far from fully ensuring that the theoretical content of some effective theory will not be found to be radically incompatible with the theoretical content of some future theory, even within its domain of empirical validity (see Ruetsche, 2018; J. Fraser 2020, p. 288, for a similar worry). One might also raise legitimate doubts about the ability of the higher-order terms to adequately encode the entirety of the new physics relevant at large distances. Giving a full response to these worries goes beyond the scope of this paper. If we leave them aside, the previous argument still goes some way toward giving us confidence in the robustness of the theoretical content of the effective theory within its domain of empirical validity.

To summarize, the Standard Effective Account takes effective theories to make approximately true and stable claims about a limited part of the world beyond which it is reasonable to expect to discover (or beyond which we have already discovered) new entities or structures. Although more work needs to be done in order to give a full defense of these features, they suggest nonetheless that effective theories provide us with a reliable epistemic standpoint to identify unobservable entities or structures in the regimes where our best theories are known to be successful. This extends Williams and J. Fraser's recent claim beyond the context of QFT and provides a further response to philosophers who deem EFTs unfit for interpretative purposes (e.g., D. Fraser, 2009; 2011; Kuhlmann, 2010). And if we are to interpret effective theories in realist terms, their structure provides us with one central constraint for making more reliable ontological commitments than those commonly made across physics: namely, we should only commit to the existence of concrete physical objects—entities, structures, properties, quantities, states, phenomena, dispositions, and so on—specifiable within the domain of empirical validity of the theory. Beyond this domain, the structure of effective theories gives us good reasons to believe that they fail to represent, or misrepresent, physical objects.

5 A Challenge for the Traditional Realist

I will now illustrate how effective theories force the traditional scientific realist to be more selective about her ontological commitments than she might think

she has good reasons to be.

Suppose for the sake of the argument that our realist feels unmoved by the traditional constructive empiricist concerns about unobservables and underdetermination (van Fraassen, 1980), the pessimistic meta-induction argument (Laudan, 1981), and the more recent problem of unconceived alternatives (Stanford, 2006). Yet, impressed by the new dogma of effective theories, our realist concedes that our best current theories are best understood and formulated as effective theories and agrees to endorse the account developed in sections 3-4. She examines the standard formulation of our best effective theories (e.g., the Standard Model Effective Field Theory), either eliminates or disregards their artifactual mathematical structures (e.g., gauge redundancies), and, after interpreting the remaining core theoretical descriptions in their literal sense as she has always done, finds out that our best effective theories represent putative infinite entities and structures, including continuum quantum fields and their infinitary symmetry structure. She also takes the remarkable empirical and explanatory success of these theories to be a good enough reason to commit to those entities and structures. But of course, knowing that our best effective theories might be superseded one day, perhaps by some advanced type of effective string theory or maybe even by some final theory, she is ready to grant that these putative entities and structures are only approximately similar to more fundamental ones.

I will argue in what follows that, on the Standard Effective Account, our realist is actually not warranted in taking the representations of these putative entities and structures to be even approximately accurate and cannot, therefore, reliably commit to their existence. Three important remarks are in order. (i) For simplicity, and in line with the traditional form of scientific realism, I will restrict myself to concrete entities, i.e., continuum fields in this case. I should emphasize, however, that a similar argument could be made for their infinitary symmetry structure and more generally for physical objects whose core features are specified well-beyond the limits of empirical validity of the effective theory of interest. (ii) I will assume that standard mathematical means of comparison (e.g., measure, cardinality, isomorphisms, etc.) provide reliable standards of relative similarity and accuracy as it is usually assumed in the literature (e.g., da Costa and French, 2003; Weisberg, 2013, chap. 8). So, for instance, I take two distinct finite sets of degrees of freedom of the same type to be much more similar to one another than either of them is to an infinite set of degrees of freedom of the same type. (iii) I will first rely on a general notion of similarity and then use the model-theoretic account of similarity to make the argument more concrete.¹⁷

¹⁷I doubt that the argument actually depends on one's favored account of similarity if we assess whether the representation of a continuum field itself (and not some finite representation thereof) is similar to the representation of a lattice field of finite extent. In the contrast-account, for instance, we need to evaluate the amount of properties shared by two representations and subtract the properties that differ between them, with specific weights assigned depending on whether the property is deemed more or less relevant (see, e.g., Weisberg, 2013, chap. 8, for a recent defense of this account). If we want to compare different fields themselves (and not simply their configurations), the number and type of their degrees of freedom appear to be essential, which means that, according to the contrast-account, two lattice fields of the same type with different spacing are much more similar to one another than either of them is to the corresponding continuum field (cf. below). I would like to thank an anonymous referee for pressing me on this point.

How should we evaluate the representation of a continuum field in a given effective theory then? Recall from the approximate truth claim that, for each effective theory, we can at least construct one accurate representation of its target system specified within its domain of empirical validity. In the ϕ^4 -theory case, for instance, we can represent the target system in terms of a lattice field defined by assigning one degree of freedom to each point of a space-time lattice of finite size and non-zero spacing. Of course, in the same way as we do not need to reduce a massive body to its point-like center of mass, we do not need to assume that the target system in the ϕ^4 -theory literally takes the form of a “grid”. A representation is approximately accurate if the putative entities specified by the representation are approximately similar to real ones. A representation which only ignores, omits, or abstracts away irrelevant features of the target system does not necessarily provide false information—the only thing we can be certain of is that it provides partial information about the target system.

Now, suppose that for the effective theory of interest, we are also able to represent the target system in terms of a continuum field. For instance, in Example 2, we can decrease the lattice spacing, increase the size of the lattice, and attribute a new degree of freedom to every newly added space-time point in the set specifying the elementary structure of the lattice. However, according to the domain restriction claim, the more we replace, add, or distort features of the target system in sufficiently small regions of space-time, i.e., the more we take into account descriptions assigning properties to the target system beyond the limits of empirical validity of the effective theory, the more the theory provides false information about the target system. In the limit, the lattice field is replaced by a continuum field with an infinite number of degrees of freedom, one at every point of space-time, and the resulting representation provides us with an infinite amount of false information about the target system in arbitrarily small regions of space-time compared to the original lattice representation. The Standard Effective Account thus does not only suggest that the representation of the putative continuum field is strictly inaccurate—and hence best understood as an infinite idealization. It also gives us principled reasons to believe that this representation is not even close to being approximately accurate.

We can make the argument more concrete by relying on a specific notion of similarity. According to the model-theoretic (or structuralist) account, for instance, two representations, or mathematical structures in this case, are similar to one another if they are isomorphic to one another, i.e., roughly speaking, if the two mathematical structures have the same number of elements and the same structural relations between their elements. Obviously, a mathematical structure with an infinite number of elements—an infinite representation in short—is not isomorphic to a finite one; but few philosophers actually think that the traditional notion of isomorphism provides an adequate standard of accuracy and the problem is to define an adequate notion of “approximate isomorphism”. da Costa and French (2003) suggest the notion of “partial isomorphism” (or “partial homomorphism”): briefly put, two mathematical structures M_1 and M_2 are partially isomorphic to one another if there is some mapping from the elements of M_1 to the elements of M_2 which preserves the substructures (and absence thereof) holding between the elements in M_1 and which does not say anything specific if we do not know whether some sub-

structure holds or not between the elements in M_1 (see, e.g., da Costa and French, 2003; Bueno and French, 2011, for more details).

Clearly, it is essential that the two representations have important chunks of substructures in common for them to be approximately similar to one another. In this case, two finite representations are always much more partially isomorphic (or homomorphic) to one another than either of them is to the corresponding infinite representation. It is non-trivial to give a precise account of degrees of partial isomorphism (or homomorphism) and I will restrict myself to giving an intuitive picture. In Fig. 1, for instance, the two lattice fields at the top have, respectively, 64 and 49 elements and share a large part of their spatial structure. We could also specify the substructures which are not preserved (e.g., the substructure associated with the local rotational symmetry transformations of the elements that leave the lattice invariant) and the substructures for which we do not know whether they are preserved (e.g., the substructure associated with some relations not apparent in the pictures). In contrast, the continuum field depicted in the top right-hand corner has infinitely many more elements than the two lattice fields and infinitely many spatial relations not reflected in the spatial structure of the two lattice fields. Agreed, the patterns of the continuum field might represent well some patterns of the lattice fields (see Fig. 1, bottom). But this does not affect the conclusion that the two lattice fields themselves are much more similar to one another than either of them is to the continuum field.

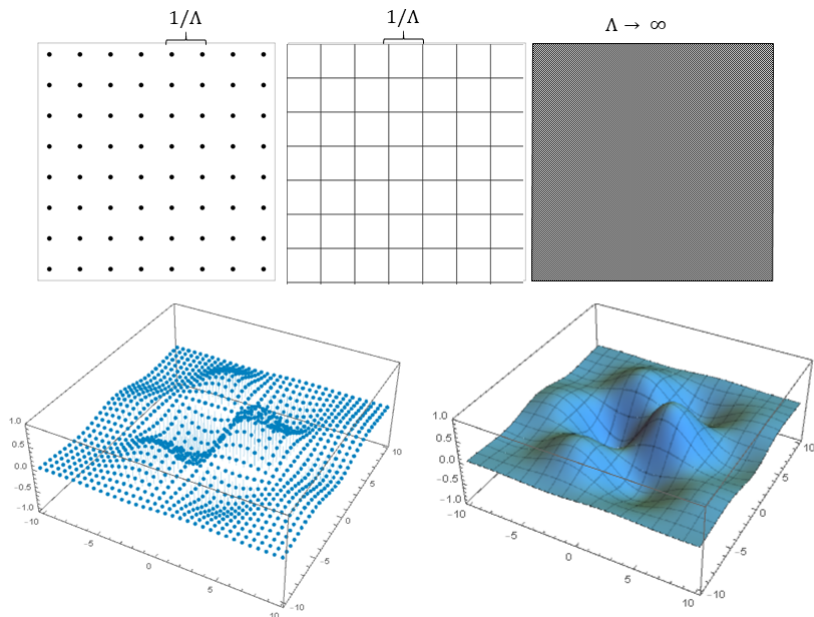


Figure 1: Schematic representations of a lattice field and a continuum field, with Λ a sharp cut-off. The two figures in the top left-hand corner represent, respectively, a finite set of points separated by a characteristic distance Λ and a finite set of blocks of characteristic size Λ . The figure in the top right-hand corner represents a continuum of points. The bottom figures represent, respectively, a lattice field configuration and its continuum counterpart.

Let me conclude this section with two comments before extending the dis-

cussion to Williams’s and J. Fraser’s defense of selective realism in QFT. First, the argument above applies to the standard formulation of our best effective theories, and therefore offers a concrete challenge to the traditional scientific realist insofar as he is willing to make ontological commitments by interpreting the central parts of our most successful theories in their literal sense. Second, the argument crucially relies on the structure of effective theories. If we have external reasons to believe that our best theories at a given time are likely to be empirically inaccurate at some scale, we might still believe that these theories give approximately true descriptions of more fundamental entities and structures. For instance, we might believe that a low-energy continuum field theory provides an approximately accurate representation of the continuum field described by a more fundamental high-energy theory. The structure of effective theories prevents us from holding such beliefs, no matter what the new high-energy physics looks like.

6 Effective Field Theories and Selective Realism

We have seen that effective theories force us to adopt a differentiated attitude towards the entities and structures that we can reliably admit in the realist inventory. In particular, we cannot admit entities if their core properties are specified in regimes where the predictions of the effective theory of interest are likely to break down. Yet, these restrictions leave, in principle, ample space for making reliable and distinctively realist ontological commitments. In the Newtonian case, for instance, we can commit to the existence of sufficiently large massive bodies of center of mass $x_i(t)$ orbiting at sufficiently large distances from each another and moving at sufficiently low velocities, including black holes which, I take it, qualify as unobservables according to van Fraassen’s original distinction (e.g., 1980, pp. 13-9). I will now argue that, in some specific theoretical contexts including classical and quantum field theory, the restrictions imposed by the structure of effective theories still raise a challenge for more refined forms of scientific realism. To bring the point home, I will focus on Williams’s (2019b) and J. Fraser’s (2018; 2020) defense of selective realism in the context of QFT.¹⁸

The strategy of the selective realist is to defend the realist cause by conceding that our best theories do not get everything right and isolating their parts which both play an essential role in their explanatory and empirical success and are likely to be preserved under theory-change (see, e.g., Psillos, 1999; Chakravartty, 2007). Upon entering the realm of QFTs, the selective realist counts herself doubly fortunate, at least at first sight. First, she can use EFT methods to formulate and interpret our best current theories in a more epistemically reliable way. She has, in particular, efficient tools for evaluating the contributions of a theory in different regimes and eliminating, or “integrating out”, its theoretical constituents which are irrelevant in the regimes she is interested in. Second, she can also use the resources of renormalization theory and, in particular, the Wilsonian RG in order to analyze the scale-dependent structure of our best EFTs and increase her confidence in the robustness of

¹⁸I will leave aside Wallace’s account insofar as he is primarily concerned with defending the foundational and interpretative relevance of cut-off Lagrangian QFTs in (2006; 2011) and not scientific realism strictly speaking (or, more precisely, structural realism).

their low-energy theoretical descriptions. It is beyond the scope of this paper to give a detailed account of Wilsonian RG methods (for a recent review, see Williams, 2019a). Here, I will restrict myself to discussing the interpretative constraints that Williams and J. Fraser extract from EFT and RG methods and evaluating the success of their selective strategy.¹⁹

How, then, should we separate the theoretical descriptions of our best current EFTs if we want to implement the selective realist strategy? Since the structure of an EFT gives us good reasons to believe that its predictions break down at some high-energy scale, we should first restrict our attention to the parts of the theory which describe its low-energy content:

1. Isolate theoretical descriptions which are specified within the limited range of scales where the theory is likely to remain reliable (see, e.g., Williams, 2019b, p. 13).

As already discussed in section 4, constraint 1 purely follows from the structure of effective theories.

Some of these low-energy descriptions might still depend significantly on irrelevant parts of the theory or involve representational artifacts (e.g., the specific type of cut-off in Eq. 4). We need, therefore, to introduce further constraints if we want to isolate the parts of the theory which play an essential role in its explanatory and predictive success and which accomplish genuine representational work. Williams and J. Fraser remain somewhat ambiguous here. They highlight various ways in which EFT and Wilsonian RG methods allow us to gain confidence in the “robustness” of the low-energy content of EFTs. Yet, they also appear to put emphasis on two different robustness criteria. Williams seems to be more concerned with the relative insensitivity of the low-energy physics to the high-energy physics:

[...] it is one of the essential virtues of the RG that it provides a tool for determining how changes in the structure of the theory at the scale of the short-distance breakdown affect physics at longer distances where the theory is empirically reliable. What the RG shows is that the ‘fundamental’ short-distance structure with which standard interpreters are so concerned is largely irrelevant to the physical content of an EFT in the domain where we have any reason to consider it empirically reliable. (2019b, p. 16)

J. Fraser, by contrast, puts emphasis on a more general type of invariance, which includes the mathematical invariance of the low-energy descriptions of the theory under different parametrizations and other representational artifacts introduced when renormalizing the theory (e.g., J. Fraser, 2020, pp. 286-7; 2018, p. 1172; see also Ruetsche, 2018, pp. 11-2; 2020, pp. 305-6; Rosaler and Harlander, 2019, sec. 5.6).

Despite important overlaps, as we will see below, I think that it is crucial to distinguish between two main interpretative constraints to account for Williams’s and J. Fraser’s slightly different outlooks and for the variety of ways in which the low-energy content of an EFT amenable to RG methods is robust:

¹⁹See also Ruetsche (2018, 2020), Rosaler and Harlander (2019, sec. 5.6), and Rivat and Grinbaum (2020).

-
2. Isolate theoretical descriptions which are largely independent of high-energy physics;
 3. Isolate theoretical descriptions which are invariant under RG-transformations and independent of specific choices of renormalization methods.

Constraint 2 is mainly derived from the structure of effective theories, although RG methods often allow us to refine the analysis. As we saw above, part of what makes an effective theory distinctive is that its descriptions which are significant within a specific regime are largely independent of its descriptions which are significant within a different regime (e.g., lower- vs. higher-order interaction terms in Examples 1-2; light vs. heavy field dynamics in Eq. 6). In particular, it is usually possible to modify the high-energy content of an EFT without affecting much its low-energy content, including its low-energy predictions (e.g., by adding higher-order interaction terms in Examples 1-2). We can also usually show that different high-energy theories reduce to the same low-energy theory, or at least to similar ones (e.g., we can add a third heavy scalar field in Eq. 6 and obtain a similar low-energy theory after integrating out the two heavy fields and making appropriate approximations). In all these cases, the crucial point is that the low-energy content of the theory is robust under variations of its high-energy content. And, in general, the bulk of the low-energy content of the effective theory depends only on a finite number of free parameters (see Examples 1-2).

Constraint 3, by contrast, arises specifically from a RG analysis. In general, a theory can be renormalized in many different ways, and the specific renormalization method chosen usually requires us to introduce some arbitrary scale parameter (e.g., the parameter Λ in Example 2) and use some particular scheme to absorb the terms depending on this parameter (e.g., a mass-dependent renormalization scheme). Thus, constraint 3 requires us to isolate theoretical descriptions which are invariant under different renormalization methods and choices of scales (cf., Williams, 2019b, p. 12; J. Fraser, 2018, p. 1172; 2020, pp. 286-7).

We can, in fact, look at this constraint in two distinct ways. (i) If we consider some fixed high-energy theory, we can derive a series of low-energy theories by successively integrating out high-energy field configurations in their path integral formulation. In this case, constraint 3 is best understood as requiring us to isolate invariant theoretical descriptions in the series of low-energy theories. (ii) If we consider some low-energy theory with parameters fixed by means of experimental inputs, we can show that this theory and its parameters remain unaffected by changes in the high-energy theory from which it is originally derived, i.e., the so-called “bare” theory (cf. Wallace, 2006, p. 49; 2011, p. 6; Williams, 2019b, p. 12; J. Fraser, 2018, p. 1172; 2020, pp. 286-7). In this case, constraint 3 is best understood as requiring us to isolate theoretical descriptions which are not affected by changes in the value of the high-energy cut-off and in the parametrization of the high-energy theory.

Now, in addition to adopting constraints 1-3, the selective realist also needs to make sure that she is offering a genuine defense of the realist cause. First, in order to give a sufficiently informative and non-ambiguous explanation of the success of the theory, she needs to isolate a definite set of unobservable entities or structures with clear identity conditions—say, in the case of entities, with a

well-specified set of core properties that distinguish them from other entities, whether they are fundamental or not. For instance, in the Newtonian case, we might identify a system by means of its position, its velocity, its mass, and its dynamical behavior. If we simply give a functional characterization of the system by means of its mass, for instance, we are likely to pick out very different types of entities and leave the target of our commitments indeterminate. Likewise, in the QFT case, we might identify a system by means of the type and number of its degrees of freedom, its mass, its self-interacting parameters, and its dynamical behavior. If we simply specify the system by means of its dynamical behavior and its mass, for instance, there is still some ambiguity as to whether we pick out a lattice or a continuum field. Contrary to what Williams (2019b, p. 15) suggests, to simply “extract reliable ontological information” does not suffice (see also J. Fraser, 2020, pp. 286-7). The selective realist needs to give a sufficiently comprehensive account of a definite set of entities or structures in order to fulfill her explanatory duties.

Second, the selective realist needs to give a literal interpretation of some privileged parts of the theory, as it is often assumed in the literature (e.g., Psillos, 1999; Chakravartty, 2007). In the Newtonian case, for instance, the selective realist can take the theory to literally describe a black hole with a center of mass specified by the position $x(t)$ and which interacts gravitationally with other bodies. The gravitational force can be interpreted as a concrete structure, i.e., as a variable relation with a specific strength depending on the relative position and the masses of the bodies. Although Williams and J. Fraser do not give much details about their preferred version of selective realism, they both seem to endorse this semantic constraint, i.e., that the privileged set of descriptions that we take be trustworthy should be understood in their standard sense and taken to be approximately true or false *simpliciter*.²⁰ In the same vein, the selective realist should avoid modifying too much the original mathematical structure of the theory or engaging into any other form of *post hoc* interpretative practice. Otherwise, she will fail to take the original theory at face value and explain its explanatory and predictive success in its own terms. This is well illustrated, for instance, by attempts to draw conclusions about the ontological content of our best current QFTs based on their putative algebraic reformulation, despite the fact that they have not yet been successfully formulated in algebraic terms.²¹

The difficulty now is that it is not clear what the selective realist should commit to if she endorses these constraints in the case of our best current EFTs, as it has been acknowledged by J. Fraser (2018, p. 1172; 2020, p. 289). I will expand on Ruetsche’s recent discussion in (2018; 2020) by looking at the most obvious candidates—correlations, particles, and lattices—and argue that they do not allow us to meet constraints 1-3 or make distinctively realist ontological commitments.

²⁰At the very least, this seems to be implicit in the central question underlying Williams’s and J. Fraser’s interpretative stance—“given that this theory provides an approximately true description of our world, what is our world approximately like?” (Williams, 2019b, p. 2). Reference to particular physical scales seems to be included in the properties of the target system (see, e.g., J. Fraser’s reference to the “bulk properties” of a fluid when he illustrates the idea of large distance features of the world in 2018, p. 1173).

²¹See, e.g., D. Fraser (2008) for such an attempt and Williams (2019b) for a criticism, emphasizing the importance of paying attention to how QFTs are successfully implemented in practice.

Correlations: J. Fraser proposes to focus on low-energy correlation functions:

[...] a preliminary strategy is to point to correlation functions over distances much longer than the cutoff scale as appropriate targets for realist commitment. These quantities are preserved by the renormalization group coarse-graining transformation and encode the long distance structure of a QFT model. They are also directly connected to its successful predictions—you cannot vary the long distance correlation functions of a theory without drastically affecting its low energy scattering cross sections. (2018, p. 1172)

We face several issues here. First, it is not clear how we should interpret correlation functions. In the standard QFT framework, they correspond to vacuum expectation values of time-ordered products of field operators at different space-time points. The simplest textbook interpretation in the simple case of two field operators $\hat{\phi}(x)$ and $\hat{\phi}(y)$ is to take the expectation value $\langle 0|\mathcal{T}\{\hat{\phi}(x)\hat{\phi}(y)\}|0\rangle$ to measure the probability (once squared) that a particle is created at some earlier point x , propagates, and is annihilated at some later point y (assuming $x_0 < y_0$). This interpretation is controversial, in large part because of the difficulties associated with the interpretation of quantum fields and particles in interacting QFTs. The crucial point here is that however we interpret these entities (I discuss the two cases below), we need to commit to something more than correlations if we follow this standard textbook interpretation. Likewise, if we interpret correlation functions more generally as standing for the degrees of co-variation or coordination between two variables at two distinct points, we need to commit to something more than degrees of co-variation (I discuss the case of physical degrees of freedom below).

We might opt for a more minimal interpretation of correlation functions as encoding structural physical information independently of the physical objects or variables they relate. In the case of EFTs, we can interpret correlation functions as encoding the correlations of the target system at sufficiently large distances, where ‘correlation’ refers to a set of numbers characterizing the degree of correlation between two space-time points or regions. If we take this route, however, the empiricist might raise doubts about the distinctively realist character of these commitments and, instead of rejecting them altogether as she usually does, simply re-appropriate them as her own as Ruetsche (2020, pp. 306-7) rightly notes. It turns out that the framework of QFT even gives her good reasons to do so. Typically, in high energy physics, we summarize empirical information about the correlations between the initial and final states of some scattering process in a mathematical object called the S-matrix, and the S-matrix can be derived by taking the appropriate asymptotic limit of a sum over all the possible correlations between initial and final states by means of the LSZ reduction formula (see, e.g., Schwartz, 2013, sec. 6.1). If we take the effects of a field disturbance to be in principle detectable in any sufficiently large region of space-time, nothing seems to prevent the empiricist from understanding the numbers specified by correlation functions as simply summarizing the empirical information that *would* be gathered about the correlations between two states of the system if we were to make measurements in this space-time region.

Even if the structural realist finds a way of avoiding this empiricist re-appropriation, she still faces one important issue. Strictly speaking, correlation functions in QFT are not RG-invariant in the general case contrary to what J. Fraser claims. If we implement a coarse-graining procedure by integrating out high-energy field configurations, for instance, the different correlation functions obtained at low energies are multiplied by “wave function normalization” factors. In general, these multiplicative factors depend on other variables, such as the couplings of the theory. And so it does not appear that there is an invariant and therefore unambiguous characterization of the degree of correlation between two distinct space-time points since it depends on the way we parametrize the low-energy theory. By contrast, S-matrix elements are invariant under these different parametrizations. Similarly, the path integral used to generate the set of correlation functions is also invariant under different coarse-graining procedures. Yet, it seems to be even more difficult to interpret the S-matrix and the path integral in distinctively realist terms compared to correlation functions. And, again, the empiricist might simply re-interpret the S-matrix and the path integral as bookkeeping devices for all the possible empirical information that we could gather about the correlations between initial and final states of the system in sufficiently large space-time regions.

Particles: Another option, perhaps more likely to enable us to make distinctively realist ontological commitments, is to focus on particles, such as protons, neutrons, gluons, and photons (see, e.g., Williams, 2019b, p. 20, p. 22). The concept of particle in interacting QFTs which involve an infinite number of degrees of freedom is controversial (see, e.g., Teller, 1995; Bain, 2000; D. Fraser, 2008; Ruetsche, 2011). In the modern understanding of QFT, it is common to understand particles in terms of patterns of excitations in the fields (as it is rightly noted by Wallace, 2006; 2019, sec. 4, for instance). This understanding is robust whether we deal with the perturbative or exact, non-interacting or interacting formulation of a QFT with an infinite or finite number of degrees of freedom (ignoring the mathematical issues inherent in realistic continuum QFTs). And, to be more precise, we can interpret particles in terms of sufficiently well-behaved and localized patterns in the field configurations in regimes where the interactions described by the theory are sufficiently weak.

Again, the main issue here is that neither field configurations nor energy-momentum states are RG-invariant. In general, RG-transformations mix both field operators and the states of different kinds of particles with one another. The only notion of “particle” that does not suffer from these issues is the one specified by the asymptotic states in the non-interacting version of the theory. But insofar as we seek a realist interpretation of *interacting* QFTs, we cannot simply restrict our commitments to the free particles that we prepare and detect in experiments. And even if we were to take this extreme route and leave aside potential empiricist re-appropriations, we would still not be able to commit to the existence of particles such as quarks and gluons insofar as the quark and gluon fields do not have asymptotic elementary particle states.

Lattice fields: A third option is to focus on low-energy degrees of freedom (e.g., as represented by the field operators associated with the variables $\tilde{\phi}(k)$ for $k \ll \Lambda$ in Example 2). Agreed, many of the properties associated with these degrees of freedom do vary under RG-transformations, including cou-

pling parameters and the specific form of the variables used to specify these degrees of freedom (which depends, in particular, on how we separate low- and high-energy degrees of freedom). Yet, whether we integrate out a large or a small range of high-energy field configurations, the number of degrees of freedom at sufficiently low energies remains exactly invariant. We could, therefore, consider them to be an appropriate target for the selective realist, as Williams sometimes seems to suggest (2019b, p. 13; pp. 14-5). The main issue here is that this might not be enough for the realist. We can interpret a degree of freedom as a determinable dynamical property of some system. However, without a specification of the low-energy system, any appeal to low-energy degrees of freedom will remain too indeterminate for the realist and therefore undermine her attempt to provide a sufficiently informative and unambiguous explanation of the success of the theory. After all, these degrees of freedom could perfectly stand for the properties of radically different low-energy systems. They could be, for instance, the degrees of freedom of low-energy lattice fields with different types of spatial structures.

In order to avoid the issue of underdetermination at low energies, we can perhaps isolate a privileged set of low-energy lattice fields for our best current EFTs. If we put a given EFT on a lattice of finite size and spacing, we can indeed integrate out high-energy degrees of freedom, obtain low-energy lattices, and eventually derive empirically equivalent low-energy *predictions* which do not significantly depend on the details of the short-distance physics and on the way we eliminate high-energy degrees of freedom (cf. Wallace, 2006, pp. 48-50). In addition, these low-energy lattices are well-specified within the limited range of energy scales where the EFT of interest is likely to remain reliable, and they do appear to enable us to make distinctively realist ontological commitments.

Yet, we still face a severe issue of underdetermination both at low and high energies. If we formulate an EFT on a lattice and interpret its low-energy descriptions in their literal sense, the RG coarse-graining transformations appear to force us to commit to the existence of different lattice fields at different low-energy scales. We might solve this issue by claiming that these lattice fields are more or less coarse-grained partial instantiations of the same high-energy lattice field. If we fix any of the low-energy lattice representations, however, RG methods allow us to change the high-energy lattice representation without affecting the low-energy lattice one. And this introduces some pernicious form of underdetermination about what the low-energy lattice representations are supposed to stand for.

There are two additional points that make the matter even worse. First, if we start with a given lattice field, we can implement a specific type of coarse-graining procedure that defines a lattice field with a different number of degrees of freedom but with the same lattice spacing. We simply need to rescale the original lattice spacing and adjust the parameters of the theory after having integrated out high-energy degrees of freedom. And the two lattice field representations are, of course, empirically equivalent (see, e.g., Hollowood, 2013, sec. 1.2, for a simple explanation of this specific way of implementing RG-transformations). Second, the specific form of the low-energy lattices depends on the type of coarse-graining procedure we implement in the first place. We might separate low- and high-energy degrees of freedom in very different ways, or define new low-energy degrees of freedom by averaging over

high-energy ones in a particular way. In each case, the procedure yields a different set of low-energy lattices. And overall, then, it appears that low-energy lattices do not allow us to satisfy constraint 3.

Now, if we are to make distinctively realist ontological commitments about entities or structures in the case of our best current EFTs and maintain Williams's and J. Fraser's robustness constraints, continuum quantum fields appear to be ideal candidates. Assuming that we do not latticeize the theory, we may either take a smooth cut-off or a sharp cut-off (in which case we eliminate high-energy states of the field), and keep higher-order interaction terms or eliminate them (depending on the desired accuracy). Either way, the theory describes a RG-invariant continuous system with an infinite number of degrees of freedom, at least for a finite range of scales. If we keep all the degrees of freedom in the theory, we do not face the issues encountered with lattices. And if we do not focus on the specific values of the properties of the continuum field, such as the value of its mass, the strength of its interactions, or the value of its field configurations on space-time, we also avoid the issues encountered with correlation functions and particles.

The main issue here comes from the domain restriction claim.²² On the face of it, we are committing to entities with core properties specified in regimes where the predictions of the EFT of interest are likely to break down, and this should be a good enough reason not to make such commitments (as Williams and J. Fraser would probably agree). In response, we might insist that we are committing to the existence of continuum quantum fields insofar as they are approximately similar to large distance scale features of the world. If we wish to endorse the literalness constraint, however, we cannot make such a claim. As we saw in section 5, if we take the descriptions of a continuum quantum field itself at face value, i.e., as being either (approximately) true or false, we are forced to attribute degrees of freedom to some putative entity in arbitrarily small regions of space-time, and the structure of effective theories gives us reasonable grounds not to commit to the existence of such entities.

We might also try to escape the difficulty by taking the representation of the putative continuum field to contain a finite part that does the appropriate representational work at large distances, say, a finite representation of a lattice field. The issue here is that any specification of such finite representation involves a particular specification of an arbitrary lattice spacing, or at least of a finite number of degrees of freedom, and therefore brings us back to the issues discussed above. The best RG-invariant representations of putative entities in our best current EFTs appear to be the representations of continuous systems with an infinite number of degrees of freedom. And we cannot simply embed these representations in finite ones without losing their representational value altogether.

7 Conclusion

I will briefly conclude with a more radical suggestion to defend the realist cause in the case of our best current EFTs. To summarize the main points of the

²²Another set of issues that I will not discuss here is related to the existence of empirically equivalent field representations (for a discussion about Borchers classes, for instance, see Haag, 1996, sec. II.5.5; Wallace, 2006, sec. 2.2, 3.3).

paper first, we have seen that the structure of effective theories across physics is best characterized by the fact that they contain intrinsic empirical limitations. In a slogan: effective theories “predict” their own predictive failure at some scale. We have also seen that the most straightforward realist explanation of this predictive pattern is to take effective theories to accurately represent limited parts of the world, which provides one central constraint for the sort of entities and structures that a realist might reliably include in his inventory if he takes effective theories seriously. I gave one concrete example of the sort of entities that the traditional scientific realist cannot commit to if he interprets the core descriptions of effective theories in literal terms: namely, he cannot commit to the existence of continuum fields since their individuating properties are specified in regimes where the predictions of the theory are likely to break down. Yet, the domain of empirical validity of an effective theory leaves, at least in principle, enough space for the realist to commit to the existence of unobservable entities or structures (as we have seen in the Newtonian case). As I have argued in the last section, this is not always straightforward. In particular, the structure of our best current EFTs is such that it is not clear what we should commit to if we want to make distinctively realist ontological commitments and avoid making these commitments depend on irrelevant or artifactual features.

I suspect that many of us still entertain the hope of a robust form of scientific realism that does not totally fail to adhere to its original letter and which is concerned with explaining the success of our best theories in their own terms. In the case of our best current EFTs, a potential candidate for making distinctively realist ontological commitments appears to be continuum quantum fields. And if we want to commit to the existence of such entities at low energies, one potential solution is to modify the traditional semantic tenet of scientific realism (but keep its ontological and epistemological tenets as summarized in, e.g., Psillos, 1999, p. xvii). Instead of taking the descriptions of a continuum field at face value, that is, as being either (approximately) true or false, we need to take them to be (approximately) true or false relative to a specific range of physical scales. That is, when we speak about a continuum field with properties assigned at every point of space-time, we are not literally making the claim that the field has properties at arbitrarily short distances *simpliciter*. We are making a claim about the structure of matter *at large distances*. And the descriptions of an effective theory are approximately true or false relative to these scales up until we discover that the theory breaks down at some limiting scale, in which case we need to work with a new theory. If the new theory is effective, we will be again making claims relative to a specific range of physical scales. This strategy requires us to modify one of the central tenets of scientific realism usually endorsed by selective realists. But it might enable us to explain the success of our best theories in their own terms.

Acknowledgments: I am thankful to David Albert, Jonathan Bain, Jeremy Butterfield, Jenann Ismael, Michael Miller, Achille Varzi, David Wallace, and Porter Williams, as well as audiences at Oxford, Columbia (SC), Gif-sur-Yvette, and Louvain-la-Neuve, for helpful discussions and comments. I would also like to thank the referees for their constructive criticisms and useful suggestions.

References

- Adams, A., N. Arkani-Hamed, S. Dubovsky, A. Nicolis, and R. Rattazzi (2006). Causality, Analyticity and an IR Obstruction to UV Completion. *Journal of High Energy Physics* 2006(10), 014.
- Aydemir, U., M. Anber, and J. Donoghue (2012). Self-healing of Unitarity in Effective Field Theories and the Onset of New Physics. *Physical Review D* 86(1), 014025.
- Bain, J. (2000). Against Particle/Field Duality: Asymptotic Particle States and Interpolating Fields in Interacting QFT (Or: Who's Afraid of Haag's Theorem?). *Erkenntnis* 53(3), 375–406.
- Bain, J. (2013). Effective Field Theories. In R. Batterman (Ed.), *The Oxford Handbook of Philosophy of Physics*, pp. 224–254. Oxford: Oxford University Press.
- Baumann, D. and D. Green (2012). Supergravity for Effective Theories. *Journal of High Energy Physics* 2012(3), 1.
- Baumann, D. and L. McAllister (2015). *Inflation and String Theory*. Cambridge: Cambridge University Press.
- Blanchet, L. (2014). Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Reviews in Relativity* 17(1), 2.
- Bueno, O. and S. French (2011). How Theories Represent. *The British Journal for the Philosophy of Science* 62(4), 857–894.
- Burgess, C. (2004). Quantum Gravity in Everyday Life: General Relativity as an Effective Field Theory. *Living Reviews in Relativity* 7(1), 5.
- Burgess, C. (2017). Introduction to Effective Field Theories and Inflation. arXiv:1711.10592: <https://arxiv.org/abs/1711.10592>.
- Burgess, C. and G. Moore (2006). *The Standard Model: A Primer*. Cambridge: Cambridge University Press.
- Butterfield, J. (2014). Reduction, Emergence and Renormalization. *Journal of Philosophy* 111(1), 5–49.
- Butterfield, J. and N. Bouatta (2014). On Emergence in Gauge Theories at the 't Hooft Limit. *European Journal for Philosophy of Science* 5(1), 55–87.
- Butterfield, J. and N. Bouatta (2015). Renormalization for Philosophers. In T. Bigaj and C. Wuthrich (Eds.), *Metaphysics in Contemporary Physics*, pp. 437–485. Leiden: Brill Rodopi.
- Calmet, X. and R. Casadio (2014). Self-healing of Unitarity in Higgs Inflation. *Physics Letters B* 734, 17–20.
- Cao, T. Y. and S. S. Schweber (1993). The Conceptual Foundations and the Philosophical Aspects of Renormalization Theory. *Synthese* 97(1), 33–108.

-
- Castellani, E. (2002). Reductionism, Emergence, and Effective Field Theories. *Studies in History and Philosophy of Science Part B* 33(2), 251–267.
- Chakravartty, A. (2007). *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge: Cambridge University Press.
- da Costa, N. C. A. and S. French (2003). *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*. Oxford: Oxford University Press.
- Donoghue, J. (1995). Introduction to the Effective Field Theory Description of Gravity. arXiv:gr-qc/9512024: <https://arxiv.org/abs/gr-qc/9512024>.
- Dubovsky, S., L. Hui, A. Nicolis, and D. T. Son (2012). Effective Field Theory for Hydrodynamics: Thermodynamics, and the Derivative Expansion. *Physical Review D* 85(8).
- Egg, M., V. Lam, and A. Oldofredi (2017). Particles, Cutoffs and Inequivalent Representations. *Foundations of Physics* 47(3), 453–466.
- Endlich, S., A. Nicolis, R. Rattazzi, and J. Wang (2011). The Quantum Mechanics of Perfect Fluids. *Journal of High Energy Physics* 2011(4), 102.
- Fraser, D. (2008). The Fate of 'Particles' in Quantum Field Theories with Interactions. *Studies in History and Philosophy of Science Part B* 39(4), 841–859.
- Fraser, D. (2009). Quantum Field Theory: Underdetermination, Inconsistency, and Idealization. *Philosophy of Science* 76(4), 536–567.
- Fraser, D. (2011). How to Take Particle Physics Seriously: A Further Defence of Axiomatic Quantum Field Theory. *Studies in History and Philosophy of Science Part B* 42(2), 126–135.
- Fraser, J. D. (2018). Renormalization and the Formulation of Scientific Realism. *Philosophy of Science* 85(5), 1164–1175.
- Fraser, J. D. (2020). Towards a Realist View of Quantum Field Theory. In S. French and J. Saatsi (Eds.), *Scientific Realism and the Quantum*, pp. 276–292. Oxford: Oxford University Press.
- Goldberger, W. and I. Rothstein (2006). An Effective Field Theory of Gravity for Extended Objects. *Physical Review D* 73(10), 104029.
- Gripaios, B. (2015). Lectures on Effective Field Theory. arXiv:1506.05039: <https://arxiv.org/abs/1506.05039>.
- Gripaios, B. and D. Sutherland (2015). Quantum Field Theory of Fluids. *Physical Review Letters* 114(7), 021602.
- Haag, R. (1996). *Local Quantum Physics: Fields, Particles, Algebras* (2nd ed.). Springer.
- Hartmann, S. (2001). Effective Field Theories, Reductionism and Scientific Explanation. *Studies in History and Philosophy of Science Part B* 32(2), 267–304.

-
- Hellerman, S., S. Maeda, J. Maltz, and I. Swanson (2014). Effective String Theory Simplified. *Journal of High Energy Physics 2014* (9), 183.
- Hollowood, T. (2013). *Renormalization Group and Fixed Points*. Springer-Briefs in Physics. Heidelberg: Springer.
- Hossenfelder, S. (2013). Minimal Length Scale Scenarios for Quantum Gravity. *Living Reviews in Relativity 16*(1), 2.
- Huggett, N. and R. Weingard (1995). The Renormalisation Group and Effective Field Theories. *Synthese 102*(1), 171–194.
- Kaplan, J. (2016). Lectures on AdS/CFT from the Bottom Up. Available at: <http://sites.krieger.jhu.edu/jared-kaplan/files/2016/05/AdSCFTCourseNotesCurrentPublic.pdf>.
- Kuhlmann, M. (2010). Why Conceptual Rigour Matters to Philosophy: On the Ontological Significance of Algebraic Quantum Field Theory. *Foundations of Physics 40*(9-10), 1625–1637.
- Laudan, L. (1981). A Confutation of Convergent Realism. *Philosophy of Science 48*(1), 19–49.
- Petrov, A. A. and A. E. Blechman (2016). *Effective Field Theories*. Singapore: World Scientific.
- Poisson, E. and C. Will (2014). *Gravity: Newtonian, Post-Newtonian, Relativistic*. Cambridge: Cambridge University Press.
- Polchinski, J. and A. Strominger (1991). Effective String Theory. *Physical Review Letters 67*(13), 1681–1684.
- Porto, R. (2016). The Effective Field Theorist’s Approach to Gravitational Dynamics. *Physics Reports 633*, 1–104.
- Psillos, S. (1999). *Scientific Realism : How Science Tracks Truth*. London: Routledge.
- Rivat, S. (2019). Renormalization Scrutinized. *Studies in History and Philosophy of Science Part B 68*, 23–39.
- Rivat, S. and A. Grinbaum (2020). Philosophical Foundations of Effective Field Theories. *The European Physical Journal A 56*(3), 90.
- Robinson, D. (1992). Renormalization and the Effective Field Theory Programme. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association 1992*, 393–403.
- Rosaler, J. and R. Harlander (2019). Naturalness, Wilsonian Renormalization, and ”Fundamental Parameters” in Quantum Field Theory. *Studies in History and Philosophy of Science Part B 66*, 118–134.
- Ruetsche, L. (2011). *Interpreting Quantum Theories: The Art of the Possible*. Oxford: Oxford University Press.

-
- Ruetsche, L. (2018). Renormalization Group Realism: The Ascent of Pessimism. *Philosophy of Science* 85(5), 1176–1189.
- Ruetsche, L. (2020). Perturbing Realism. In S. French and J. Saatsi (Eds.), *Scientific Realism and the Quantum*, pp. 293–314. Oxford: Oxford University Press.
- Schwartz, M. (2013). *Quantum Field Theory and the Standard Model*. Cambridge: Cambridge University Press.
- Schweber, S. S. (2015). Hacking the Quantum Revolution: 1925-1975. *The European Physical Journal H* 40(1), 53–149.
- Stanford, P. K. (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.
- Teller, P. (1995). *An Interpretive Introduction to Quantum Field Theory*. Princeton: Princeton University Press.
- van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Clarendon Press.
- Wallace, D. (2006). In Defence of Naivete: The Conceptual Status of Lagrangian Quantum Field Theory. *Synthese* 151(1), 33–80.
- Wallace, D. (2011). Taking Particle Physics Seriously: A Critique of the Algebraic Approach to Quantum Field Theory. *Studies in History and Philosophy of Science Part B* 42(2), 116–125.
- Wallace, D. (2019). The Quantum Theory of Fields. In E. Knox and A. Wilson (Eds.), *The Routledge Companion to the Philosophy of Physics*. Routledge. Forthcoming. Philsci preprint: <http://philsci-archive.pitt.edu/15296/>.
- Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.
- Wells, J. D. (2012). *Effective Theories in Physics: From Planetary Orbits to Elementary Particle Masses*. Heidelberg: Springer Science & Business Media.
- Williams, P. (2019a). Renormalization Group Methods. In E. Knox and A. Wilson (Eds.), *The Routledge Companion to the Philosophy of Physics*. Routledge. Forthcoming. Philsci preprint: <http://philsci-archive.pitt.edu/15346/>.
- Williams, P. (2019b). Scientific Realism Made Effective. *The British Journal for the Philosophy of Science* 70(1), 209–237.