

Whence the Effectiveness of Effective Field Theories?

Alexander Franklin

2017

Abstract

Effective Quantum Field Theories (EFTs) are effective insofar as they apply within a prescribed range of length-scales, but within that range they predict and describe with extremely high accuracy and precision. The effectiveness of EFTs is explained by identifying the features – the scaling behaviour of the parameters – which lead to effectiveness. The explanation relies on distinguishing autonomy with respect to changes in microstates ($\text{autonomy}_{\text{ms}}$), from autonomy with respect to changes in microlaws ($\text{autonomy}_{\text{ml}}$), and relating these, respectively, to renormalisability and naturalness. It is claimed that the effectiveness of EFTs is a consequence of each theory's $\text{autonomy}_{\text{ms}}$ rather than its $\text{autonomy}_{\text{ml}}$.

Contents

1	Introduction	2
2	Renormalisability	5
2.1	Explaining renormalisability	11
3	Naturalness	14
3.1	An unnatural but renormalisable theory	17

4	Two Kinds of Autonomy	19
5	The Effectiveness of EFTs	24
6	Conclusion	26

1 Introduction

Effective Quantum Field Theories (EFTs) only apply within a prescribed range of length-scales, but within that range they may predict and describe with extremely high accuracy and precision: this is the effectiveness of EFTs which I seek to explain in this paper.

The title question is a special case of the more profound question: how is higher-level science possible at all? Answering the question in this more restricted context should help to provide clues regarding more general answers. In particular, while effective theories are ubiquitous throughout physics – any given theory will have limits to its applicability – a discussion of EFTs will allow for progress because the physics here is sufficiently mathematized that we may identify the particular theoretical property which accounts for the effectiveness of the theories. The title question is, moreover, interesting in its own right: one of our most fundamental theories – the standard model of particle physics – is an EFT; and there is already a small but live debate in the philosophy literature concerning various aspects of EFTs.¹

Two related questions, each of which has broader implications for the debates concerning inter-theoretic relations, are addressed below. The first question corresponds to the title: what about EFTs underlies their effectiveness? By ‘effective’, I mean to refer to the empirical success of EFTs within an explicitly delimited domain; the effectiveness of EFTs implies that when, say, energies are too high, or distances too small, the theory will generate inaccurate predictions. The question is how EFTs may be empirically successful despite leaving out details which are relevant at shorter distances and higher energies.² What accounts for the freedom to abstract from such

¹See e.g. Butterfield (2014), Castellani (2002), Crowther (2015), Hartmann (2001), and Williams (2015), and references therein.

²Note that this is distinct from the prior question concerning the empirical success of

details with negligible loss of predictive power? Answering this question is central to many aspects of the reduction-emergence debate, which consider how lower-energy abstractive descriptions arise.

In order to address the title question, a second related question must be answered: how do the low-energy and high-energy EFTs relate to each other? In what sense might it be said that the low-energy theory floats free from the high-energy theory? I've noted that the low-energy theory may leave out details relevant at higher energies, but in order to understand what allows such details to be discarded, it's necessary to establish which details are in fact left out.

To answer these questions, I appeal to two theoretical properties of EFTs: renormalisability and naturalness. Renormalisable theories are those where parameters at low energy may be redefined so as to take into account details relevant at high energies; while natural theories are those where the parameters at low energies are insensitive to putative changes in parameters at high energies. A principal aim of this paper is to unpick the differences between these properties. The question of how EFTs float free from one another is further motivated by considering the following puzzle. As just noted, EFTs may be renormalisable and/or natural: however, both renormalisability and naturalness are generally thought to make theories autonomous: how, then, is it possible that some theories are renormalisable and unnatural?

The puzzle is resolved by distinguishing two types of autonomy: autonomy from microstates ($\text{autonomy}_{\text{ms}}$) corresponds to invariance of the dynamics of a low-energy theory with respect to certain changes in the state at high energy – changes consistent with the high-energy dynamics; autonomy from microlaws ($\text{autonomy}_{\text{ml}}$) is invariance of the low-energy dynamics with respect to certain changes in the high-energy theory – changes in its laws or fixed parameters. I claim that renormalisability, together with a separation of mass scales, leads to $\text{autonomy}_{\text{ms}}$, while naturalness leads to $\text{autonomy}_{\text{ml}}$. Thus, unnatural renormalisable theories are those which are $\text{autonomous}_{\text{ms}}$ but not $\text{autonomous}_{\text{ml}}$. This distinction also allows the title question to be addressed: the effectiveness of EFTs – the fact that EFTs may be successful despite leaving out details relevant at higher-energy scales – is a consequence of their $\text{autonomy}_{\text{ms}}$, but is independent of their $\text{autonomy}_{\text{ml}}$. As such, renormalisability accounts for the effective-

scientific theories in general; in this paper I presume but do not discuss realism about empirically successful EFTs.

ness of EFTs. I go on to claim that renormalisability may be explained by the scaling behaviour of the theoretical parameters.

An intriguing consequence of this discussion is that autonomy_{ms} (and hence renormalisability) does not entail a total decoupling of physics at different length-scales. Unnatural EFTs exhibit extremely sensitive dependence between theoretical parameters defined at vastly different length-scales. I argue that renormalisability may lead to the effectiveness of even unnatural EFTs. Thus, the properties of the descriptions at different length-scales may turn out to be sensitively related even while the dynamics decouple. Thus, the properties of the descriptions at different length-scales may turn out to be sensitively related even while the dynamics decouple. Discussions of emergence and reduction in the philosophy of physics often emphasise that lower-energy descriptions exhibit autonomy with respect to higher-energy science. The discussion in this paper and, in particular, the distinction between two types of autonomy, clarifies the sense in which lower-energy science may exhibit such autonomy: autonomy sufficient for emergence is compatible with sensitive dependence between length-scales. Distinguishing between these senses of autonomy also contributes to the dialectic concerning the relations between EFTs and emergence; see e.g. Bain (2013b), Butterfield (2014), Castellani (2002), and Crowther (2015).

In §2 and §3, I discuss renormalisability and naturalness respectively; an example of an unnatural, renormalisable theory is given in §3.1. In §4, I distinguish ways in which a theory might be autonomous from an underlying theory, and thus answer the question ‘in what respect do low-energy EFTs float free of the corresponding high-energy theories?’ This sets the stage for §5’s claim (*pace* assumptions made in Williams (2015)) that it is renormalisability not naturalness which answers the question ‘what allows low-energy theories to abstract from details salient at higher energies?’; thus the effectiveness of EFTs is explained.

Two terminological comments: first, Hartmann (2001) discusses whether EFTs are best thought of as theories, models or *sui generis* scientific structures; I find that discussion illuminating, but for ease of reading I refer to EFTs as ‘theories’ throughout. Second, standard philosophical talk of higher and lower levels is inconsistent with standard physics usage. Thus, I talk of low-energy, large length-scale theories and high-energy, short length-scale theories, these correspond to what philosophers call ‘high-level’ and ‘low-level’ theories respectively.

2 Renormalisability

In §4 I argue that renormalisability leads to the autonomy_{ms} which accounts for the effectiveness of EFTs. In this section I define renormalisability and discuss what features in the world renormalisability tracks. The physics here is fairly schematic, the aim is not to teach effective field theory, for that see e.g. Petrov and Blechman (2016). I start by defining renormalisation and renormalisability, I go on to discuss effective renormalisability, and, in §2.1, I consider what explains renormalisability.

In all cases I restrict discussion to renormalisation and renormalisability in the framework of perturbation theory. See e.g. Berges, Tetradis, and Wetterich (2002) for details of the physics of non-perturbative renormalisability; while a philosophical discussion of this topic would be interesting, it's not clear that the claims made here would straightforwardly carry over. Thus, for the purposes of this paper, I focus on the perturbative domain, which has also been the primary focus of the philosophical literature on this topic. Further work on non-perturbative renormalisability should build on the conceptual progress made here.

Renormalisability is crucial to the effectiveness of EFTs because it limits a given theory's dependence on higher-energy physics: a renormalisable high-energy theory affects the corresponding low-energy theory only by way of finitely many modifications to the theory's parameters. Thus the dynamics decouple and an effective theory may be developed without requiring the input of detailed knowledge concerning high-energy goings-on.

Equation (1) gives the Green's function for a generic quantum field theory involving fields ϕ with momenta p_1, \dots, p_N , masses m, M , and couplings g and g_H , μ stands for the energy at which the system is probed. The Green's function is central to extracting the probabilities for outcomes in scattering experiments.

$$\begin{aligned} &G_N(p_1, \dots, p_N; g, g_H, m, M, \mu) \\ &= \langle 0 | \phi(p_1) \dots \phi(p_N) | 0 \rangle_{\text{full theory}} \end{aligned} \tag{1}$$

This theory is well defined only if it is regularised at some high-energy cutoff scale Λ_H , which serves as the upper limit on integration and eliminates certain problematic divergences. Renormalisation is the procedure

whereby various terms in the theory may be redefined such that they take account of this cutoff. This allows for the further step of removing the cutoff without bringing back divergences. One may understand the processes of regularisation and renormalisation as a procedure for demarcating the expected applicability of the theory: it will only apply at energies significantly below the cutoff scale ($E \ll \Lambda_H$).

If one then considers a lower energy regime where $E < M$ and M is the mass of the heavier particle, we may transition to a theory corresponding to equation (2) below. In this equation, terms referring to M and the heavy coupling g_H are integrated out and the couplings and mass terms for other particles represented by g and m are redefined. This redefinition is crucial, and corresponds to an additional renormalisation procedure.

$$\begin{aligned} & G_N^*(p_1, \dots, p_N; g^*, m^*, \mu) [1 + \mathcal{O}(1/\Lambda_L)] \\ &= \langle 0 | \phi^*(p_1) \dots \phi^*(p_N) | 0 \rangle [1 + \mathcal{O}(1/\Lambda_L)] \end{aligned} \quad (2)$$

In the low-energy theory $\Lambda_L \sim M$ is now the cutoff scale and the theory is thus applicable at energies significantly below that scale. Once the theory has been renormalised with the new cutoff, and a low-energy effective theory has been defined, the cutoff may once again be removed; this does not return us to the old theory because the intermediate integrating-out of high-energy terms and subsequent renormalisation define a new theory.³

One may understand a quantum field theory by considering the full Lagrangian as a sum: $\mathcal{L} = \sum_n \mathcal{L}_n$. Each \mathcal{L}_n corresponds to an order of the perturbative expansion of the Lagrangian. Insofar as perturbation theory is valid (which is all that's of concern in this paper), each higher order is of decreasing relevance. The perturbative expansion is in general infinite, as such any evaluation will involve a truncation at some order and, consequently, a degree of approximation. However the contribution of neglected higher orders often leads to discrepancies far smaller than those which may be experimentally measured.

Renormalisation involves appeal to renormalisation group (RG) equations; in principle these probe the relationship between theoretical param-

³This is somewhat simplified: in general one applies Gell-Mann's 'totalitarian principle' which demands that all operators not ruled out by symmetry principles be included for each new theory; see Grinbaum (2008) for a discussion.

eters as some control parameter is varied.⁴ In the context of quantum field theory, that control parameter is the length-scale, or, equivalently, the energy.⁵ RG equations track the changes in parameters as the energy or distance cutoff is changed; alternatively, where the cutoff is modelled as placing the system on a lattice, renormalisation may be understood to re-model the system on lattices with different minimal spacings.

Renormalisation is, in fact, rather messier than the abstract description just given. A theory's Lagrangian may be expressed as a sum of operators, with the theory's parameters included as coefficients of those operators. In order to recover appropriate parameter values as the cutoff is lowered, counterterms are introduced which are constrained to match renormalisation conditions. The counterterms may be subtracted from the original terms in such a way that the theory is re-expressed as a sum of operators with only the values of the parameters having varied.⁶ The introduction of counterterms and subsequent reparameterisation allows the construction of a theory which putatively describes the world in a different energy domain from that of the original theory. However, the complex mathematical process just described does not work for all theories; the theories for which it works are termed 'renormalisable'. Renormalisable theories are defined as those which allow the cancellation of divergences by a finite number of counterterms.

An important feature of renormalisable theories is that, at low energy, the high-energy terms don't show up explicitly; rather all effects of high-energy scales are absorbed into the redefinition of the mass and coupling parameters. For a given theory, at a given scale, such parameters are fixed; thus, if a theory is renormalisable, changes at high energies consistent with

⁴In practice one must often appeal to experimentally sourced information in order to set various theoretical parameters.

⁵Note that RG equations are interpreted differently in quantum field theoretic (QFT) contexts from those used in condensed matter physics: for some condensed matter systems, the control parameter may be the temperature – this means that we may empirically track how particular systems' evolution is described by an RG flow. In QFT contexts we can't change the scale but we can probe the system at different energies. Thus we can track the RG flows which describe parameters' change with a set of different experiments at various energies.

⁶The introduction of counterterms will differ on different renormalisation schemes. Bain (2013a) argues that the choice of such a scheme is consequential, however, while certain schemes allow for greater mathematical tractability, whether or not a theory counts as renormalisable is scheme-independent, thus Bain's concerns are orthogonal to issues at stake here.

the high-energy dynamics will not affect the low-energy description.⁷ It is this aspect of renormalisability which allows for autonomy_{ms}, see §4: while the theory depends on and, indeed, is generally derived from a higher-energy theory, the dynamics may decouple.

Compare equations (1) and (2): the effective theory is renormalisable insofar as all effects of M and g_H can be absorbed into redefinitions of the parameters $g \rightarrow g^*$ and $m \rightarrow m^*$; thus the effective theory is renormalisable in the limit as $\Lambda_L \rightarrow \infty$. Sending the cutoff to infinity allows for mathematical tractability and, as will be discussed shortly, can be used to remove non-renormalisable contributions.

A variety of interpretative issues accompany discussions of renormalisation in QFT. While I won't discuss these in detail here, it's worth flagging some relevant issues. The modern conception of renormalisation owes a great deal to the work of Kenneth Wilson and collaborators (see e.g. Wilson and Kogut (1974)). David Wallace (2011) builds on this to argue that the best way to understand renormalisation is that all QFTs are EFTs (with limited domains of applicability), and that the otherwise odd process of subtracting infinities is better understood as a procedure undertaken with the aim of ensuring mathematical tractability. James Fraser (2017) elaborates on this point and argues that the fact that cutoffs are generally taken to their infinite limit does not warrant viewing these models as continuum models – rather we should view the infinite limit as allowing for approximations to the intractable finite theories. Effective theories thus generally display two approximations, one due to the truncation of their perturbative expansion, as discussed above, and the other due to sending the cutoff to infinity and consequently eliminating terms which depend on inverse powers of the cutoff; both approximations are often sufficiently small that the theories lead to extremely accurate predictions.

The accounts due to Wallace and Fraser apply straightforwardly to renormalisable theories. Below, I consider EFTs which are not traditionally renormalisable. In such cases, while we take the infinite limit for certain terms in certain calculations, we may also maintain finite-valued cutoffs in order to take account of the contribution of non-renormalisable terms. Fraser's claims are further evidenced in these contexts, for such EFTs are explicitly not to be regarded as continuum theories.

⁷Changes in high-energy parameters would affect the low-energy description, but such changes are dynamically impermissible.

How may one establish whether a theory is renormalisable? The renormalisation group is crucial here: it provides a schema for probing parameter variation over a range of length-scales without requiring evaluation of the theory at each scale involved. RG equations generate a flow which allows for the classification of operators into classes termed ‘relevant’ and ‘irrelevant’ according to how their contributions vary as the flow approaches a fixed point.⁸ Theories with non-trivial fixed points are renormalisable if there are scales at which the irrelevant operators vanish. The RG is thus a framework for determining renormalisability since the RG can provide information about higher-order terms without explicit calculation.⁹

Renormalisable theories are such that one can change the length-scale, adjust finitely many parameters and the theory still purports to describe the world. Thus, a renormalisable theory does not describe the world at all scales, but its renormalisability means that it can, in principle, slot into a level-based description at any level.¹⁰ On the other hand, EFTs which are not traditionally renormalisable are only usable within a certain range of length-scales; beyond their so-called ‘breakdown scale’ they may fail to be perturbatively tractable or lead to violations of unitarity.

In fact, the connection between a cutoff and non-renormalisability is not accidental. Anthony Duncan (2012, p. 653) notes that: “any physically sensible theory should include a cutoff at some high-energy / momentum scale, and that the necessary result of such a cutoff was the appearance of an *infinite* number of operators including the baleful non-renormalisable ones”. While the cutoff remains finite, the non-renormalisable operators will not go to zero. Thus effective theories with finite cutoffs will be traditionally non-renormalisable.

However, there is a weaker condition which non-renormalisable EFTs may satisfy. This I call ‘effective renormalisability’. As noted above, a theory’s Lagrangian may be expressed as a sum of Lagrangians. Each Lagrangian will contain operators which may be assigned numbers known as the degrees of divergence: if the degree of divergence of its operators is too high (usually greater than 4), then the Lagrangian, in the traditional sense,

⁸Operators may also be marginal – it is more difficult to assess the divergence behaviour in this case.

⁹See Polchinski (1984) for a proof of the efficacy of this approach for perturbative renormalisability.

¹⁰For my purposes, I define theory identity as follows: a single theory may have parameters with different values at different scales, the theory would only count as a new theory if we eliminate certain terms, such as the mass and coupling terms for high-energy particles.

counts as non-renormalisable.¹¹ The procedure which assigns degrees of divergence to operators is known as ‘power counting’.

If we are able to organise the theory’s Lagrangians in order of increasing degree of divergence, so $\mathcal{L} = \mathcal{L}_d + \mathcal{L}_{d+1} + \mathcal{L}_{d+2} + \dots$, then the theory may be effectively renormalisable. This is because, for some high-energy cutoff Λ , we may be assured that the Lagrangian will scale as $\langle \mathcal{L}_{d+n} \rangle \sim (E/\Lambda)^n$ where E is the energy scale at which the system is measured. Such terms will not diverge so long as $E \ll \Lambda$ for $n > 0$.

If a theory is traditionally renormalisable then, to any order, only finitely many corrections are required to cancel the theory’s divergences. Theories which are not traditionally renormalisable and, thus, require infinitely many corrections to cancel all of their divergences, may nonetheless be effectively renormalisable. Within their domain of applicability, to a given order in the energy-momentum expansion of the theory, effectively renormalisable theories only require finitely many corrections to cancel their divergences. That is, counterterms or equivalent can be constructed such that, order by order, they result in a theory with non-renormalisable terms which are proportional to inverse powers of the cutoff. Effectively renormalisable theories will thus have finite predictions for low energies, but uncontrolled contributions beyond the breakdown scale, at which point we do not expect the theory to generate physically meaningful predictions. While this distinction has not been much discussed in the philosophical literature, it is well known in the physics literature; Matthew Schwartz (2014, p. 392) observes that:¹²

[a]s long as we are only interested in physics at low energy, only a finite number of terms in this [infinite] series will be important. Thus, we can fit those terms with a few measurements and then predict the complete momentum dependence. In this way, the [traditionally] non-renormalizable theory is predictive even at tree-level.

While one can often discover fixed points and thus establish the renormalisability of a given theory, rigorous demonstration of effective renor-

¹¹Non-renormalisable operators will be classed as irrelevant if the theory in question lies in the basin of attraction of a fixed point.

¹²See also Stewart (2014, p. 9).

malisability is much harder to come by.¹³ Effective renormalisability is assessed by appeal to power counting schemes tailored to each theory. These, to some extent, have to be justified empirically – van Kolck (2002) details how a particular power counting technique for chiral EFT is justified on the basis of both calculations and empirical comparison. Note that many questions over the status of effectively renormalisable theories have not yet been settled. For example, there are ongoing discussions concerning the validity of particular power counting schemes within the chiral EFT community; see e.g. Epelbaum and Meißner (2013) and Nogga, Timmermans, and van Kolck (2005).

Such complexities are no doubt interesting, but can be put to one side for our present purposes. While it's difficult to work out if a given EFT is effectively renormalisable or not, any theory which is not, at least, effectively renormalisable, is inadequate for the EFT project and generally calculationally intractable. In §4 I relate renormalisability and effective renormalisability to autonomy_{ms}.

2.1 Explaining renormalisability

Given that I account for the effectiveness of EFTs in terms of their (effective) renormalisability, I should say what explains renormalisability – otherwise I might fail fully to answer the title question. In this section, I consider an explanation of renormalisability suggested by Jeremy Butterfield (2014), and offer my own competing response.

Butterfield appeals to the fixed point structure of the renormalisation group (RG) in order to explain theories' renormalisability; were his claims accepted, the RG might be seen to provide an alternative answer to the title question – it is, thus, additionally important that I address his claims.

[T]he modern approach to renormalization ... explains, indeed deduces, a striking feature (namely, renormalizability) of a whole class of theories. It does this by making precise mathematical sense of the ideas of a space of theories and a flow on the space, called the renormalization group (RG).

¹³As such, some theories may be termed 'effectively renormalisable' in the literature even while it is not established that their terms can be arranged in order of increasing degree of divergence.

...
it is surely uncontroversial that one very satisfying way to explain the good fortune reported at the end of section III.1 [that our best QFTs are renormalizable] would be to show, not merely that some given theory is renormalizable, but that any theory, or more modestly, *any* of a large and/or generic class of theories, is renormalizable.

[Butterfield (2014, pp. 14, 25-26), original emphasis]

Butterfield asks why our best theories are renormalisable, and responds that almost all theories at low energies are renormalisable. The demonstration relies on deriving that, as theories are evaluated at lower energies, the contribution of non-renormalisable terms dwindles, and all that's left are the renormalisable terms; many theories which are otherwise distinct away from a fixed point will converge on the same renormalisable theory at the fixed point. Butterfield argues that this result follows directly from the nature of RG flows, and he acknowledges that it depends on the non-trivial fact that "interactions that are strong at short distances should be weak at long distances" (ibid., p. 27).

As the nature of this explanation is not entirely clear in Butterfield's article, I think it worth examining further. The RG *qua* mathematical framework seems ill-suited to provide a sufficient explanation of renormalisability. While the RG allows for the mathematical expression of facts about renormalisability, the mathematical derivation of renormalisability is better described as a codification rather than an explanation. However, this codification does allow us to identify the features responsible for renormalisability. For example, the demonstration of the convergence of RG flows to an infrared fixed point tells us that a class of theories, all of which are renormalisable have a particular commonality: they have parameters which have diminishing value as we move to lower-energy descriptions.

If a theory is renormalisable, its parameters change with scale in just such a way that they can be absorbed when constructing a lower-energy theory. The renormalisability of successful theories is explained by the fact that structures in the world have particular features: that the mass values and coupling strengths scale appropriately. The convergence of the RG flows allows us to recognise this explanation. Thus we may view the RG as a mathematical framework which allows us to establish whether or not a given theory is renormalisable, and what that renormalisability depends on.

There are two further reasons to consider the RG as codifying rather than explaining renormalisability. The first is that the fixed point structure of the RG is not generic: there are many theories, not least those which are effectively renormalisable, for which the fixed points correspond to unrealistic non-interacting theories. As such, the properties of renormalisability and effective renormalisability apply more generally than the RG fixed point analysis.

Although fixed point arguments do not always apply, the explanation of renormalisability and effective renormalisability may be unified as suggested above: theories which describe the world are (effectively) renormalisable if their theoretical parameters scale appropriately. This allows a theory to include the effects from other energy scales in reparameterisations.

My second concern with RG arguments relates to a particular formulation of the RG explanation. The existence of an RG infrared fixed point tells us that whatever the value of the couplings in the high-energy theory, many theories will converge on the same low-energy description. Thus, one might argue, the existence of RG flows guarantees that many different theories are renormalisable. As noted above, this argument form allows us to identify what's responsible for renormalisability, but the convergence of flows does not by itself explain renormalisability. This is because there is only one system which in fact flows to a given fixed point.

Unlike in the context of continuous phase transitions and the universality of critical phenomena, where common behaviour is exhibited by multiple distinct physical systems, there is not, in general, common behaviour for distinct physical systems described by EFTs; the only universality in this case is that between the actual system, and a number of hypothetical systems with varying behaviour away from the fixed point. So, the convergence of flows is a mathematical rather than a physical phenomenon. For phase transitions, the convergence of RG flows plays an explanatory role insofar as each convergent trajectory can be shown to correspond to a different system tending towards the same fixed point, in the EFT case different physical systems tend towards different fixed points and exhibit different scaling behaviours.¹⁴ These observations demonstrate that the convergence of RG flows relates differently to the renormalisability of EFTs than to the universality of critical phenomena.¹⁵

¹⁴See e.g. Batterman (2017) and Franklin (2018) for contrary analyses of similar appeals to the RG in explanations of universality.

¹⁵This is not to suggest that EFT-described systems never exhibit universality, rather that

My disagreement with Butterfield's explanatory claim may be associated with a distinction between two kinds of explanation. Laudan (1990, p. 55) considers the following question "why are Olympic runners more successful than the runners in the local high school?" and observes that this could either be answered by comparing the anatomies and physiologies of the respective groups, or by listing the procedures used by the Olympic selection panels. That each answer might be relevant in different contexts may provide a way of deflating the debate over the two ways of explaining renormalisability. The RG flows are akin to selection procedures for dismissing non-renormalisable theories, while the anatomy comparison is like discussion of the parameter scaling. Just as pointing to the selection procedure provides a good answer for why the cohort of Olympic runners does better than the cohort of school runners, discussion of RG flows provides a good explanation for why renormalisability is fairly generic amongst EFTs. The scaling properties, like the account of physiology, provide a better answer to the question about particular theories or runners. So, to explain why a given theory is renormalisable one ought to discuss its parameter scaling, just as 'why is Usain Bolt such a good runner?' is answered by considering his training and physiology.

Renormalisability is about how the parameters tracked by our theories vary relative to each other at different scales. The RG may be used, as described above, to bring out features of a theory, including the theory's renormalisability, and systematically to express such theoretical properties; although the RG may also play a role in explaining the genericity of renormalisability among EFTs, claims that this explains particular theories' renormalisability, or accounts for their effectiveness, are ill-motivated. Instead, it is the parameters, including the mass values and coupling strengths, and the variation of such parameters with changing length-scale, which better explain a theory's renormalisability, and, thus, help to explain the effectiveness of EFTs.

3 Naturalness

Having discussed renormalisability, I turn to naturalness, a term that's played a significant heuristic role in the development of various Beyond

universality *qua* commonality between multiple, actual systems is insufficiently generic to be relevant to explanations of renormalisability.

the Standard Model theories in high-energy physics. Naturalness is relevant to the concerns of this paper for two reasons: firstly, Porter Williams (2015) has suggested that naturalness is closely related to the effectiveness of EFTs, and, secondly, consideration of the type of autonomy which corresponds to naturalness will allow for the clarification of the sense in which EFTs at different scales inter-depend.

Naturalness is a tricky concept to define, and seems to be used in a variety of different ways within the physics community. Williams (2015) sorts through a number of competing definitions and settles on the view that naturalness is a form of theoretical autonomy: it corresponds to the degree of sensitivity of parameters defined at low energies to their high-energy counterparts; thus unnatural theories are those which exhibit incredibly sensitive dependence of low-energy parameters on high-energy parameters. In this paper, I follow Williams' definition and will not rehearse the arguments he considers in its favour.

To get a better handle on the nature of naturalness, consider a pair of low and high-energy EFTs. As described above, the low-energy theory will be constructed such that its dependencies on details of the high-energy theory will only appear in modifications to its coupling parameters. The values of such parameters are determined by the renormalisation group equations which take into account the effects due to details at different length-scales.¹⁶ It is standard for such parameters to exhibit logarithmic sensitivity to the values of parameters at high energies. However, in cases of unnaturalness, parameters may exhibit quadratic sensitivity.

To see the sensitivity of $\lambda_2(\Lambda_L)$ to the cutoff-scale value $\lambda_2(\Lambda_H)$, set $\Lambda = \Lambda_L = 10^5$ GeV and the UV cutoff at the Planck scale $\Lambda_H = 10^{19}$ GeV, so that the ratio of scales appearing in the equation becomes $(\frac{\Lambda_L}{\Lambda_H})^2 = 10^{-28}$. Now alter the 20th decimal place of $\lambda_2(\Lambda_H)$, sending $\lambda_2(\Lambda_H) \rightarrow (\lambda_2(\Lambda_H) + 10^{-20})$. Plugging this value for $\lambda_2(\Lambda_H)$ into the above equation shows that this tiny change in $\lambda_2(\Lambda_H)$ causes the low-energy value $\lambda_2(\Lambda_L)$ to jump by a factor of 10^8 !

[Williams (2015, p. 87)]

So-called natural theories do not exhibit sensitive dependence of low-energy parameters on values at high energies. For natural theories the cou-

¹⁶In practice, experimental information is often also required.

plings at very high energies could take a range of values with negligible impact on experimental results at low-energy scales.

This view of naturalness might face some of the interpretative puzzles considered in the last section: historically it was believed that the bare mass and bare couplings defined at high energies represent true properties of the system and that the renormalisation and introduction of counterterms correspond merely to radiative corrections. This old view was problematic because these quantities were infinite in theories where the cutoff is taken to infinity, and one thus had to engage in the opaque process of subtracting infinities. The rise of the EFT view is sometimes seen to have resolved these problems; the EFT view tends to interpret the high-energy couplings as mathematical devices and the infinities as arising only if the theory is incorrectly viewed as applying outside its proper energy regime.¹⁷

The claim that unnaturalness involves sensitive coupling between scales may restore the old view that bare parameters are more than mere mathematical devices – that they are, in fact, the parameters of the corresponding high-energy theory. However, on the EFT view, all these theories are to be regarded as effective theories; as such, even the high-energy theories are defined with their own cutoffs and with validity below that cutoff. While the bare parameters, which are defined at the cutoff, may still be regarded as mere mathematics, this view allows that there are physically meaningful parameters defined at high energies below the high-energy cutoff, which sensitively relate to low-energy parameters. Thus, the high-energy theory ought not to be conceived of as a continuum model with an infinite cutoff; instead, following Fraser (2017), we should view the high-energy theory as having a finite cutoff, and the infinite limit as a way to approximate certain properties of the finite theory. This step avoids the conceptual problems which previously arose when renormalisation was viewed as a process for correcting infinite bare parameters, but allows there to be connections between parameter values defined at different length-scales.

Before assessing whether or not naturalness accounts for the effectiveness of EFTs, it's worth understanding the relation between naturalness and renormalisability. The following case study makes this clear.

¹⁷See Castellani (2002) for a discussion.

3.1 An unnatural but renormalisable theory

Natural theories may be renormalisable or effectively renormalisable: renormalised theories absorb details from high-energy scales into parameters at low energies. One way this is done is via the specification of high-energy terms and counterterms. In natural theories, one can vary the terms defined at high energies while leaving the low-energy terms unchanged. As such, particular high-energy parameters may not be thought of as physically meaningful. In unnatural theories, the values of the parameters at high energies are tightly coupled to the values of the renormalised parameters at low energies. Note that in many unnatural theories one may produce counterterms to deal with the unnatural divergences, and, thus, retain renormalisability.

Consider the following Lagrangian:

$$\mathcal{L} = \frac{1}{2}\phi(\square + m^2)\phi + \lambda\phi\bar{\psi}\psi + \bar{\psi}(i\not{\partial} - M)\psi \quad (3)$$

This is a toy theory which describes the evolution and interaction of a scalar boson (ϕ) with bare mass m and a fermion (ψ) with bare mass M ; the particles interact with coupling strength λ . Following Schwartz (2014, p. 410) we may consider the \overline{MS} one-loop renormalisation of the mass term in this theory which leads to equation (4); there the measured pole mass m_P is derived from a correction to the renormalised mass $m_{\overline{MS}}$. In other words, one may renormalise the bare mass m to obtain $m_{\overline{MS}}$ but further corrections are required to obtain the experimentally measured value. These further corrections expose the unnatural scale dependence.

$$m_P^2 - m_{\overline{MS}}^2(\mu) = \frac{\lambda^2}{24\pi^2}(6M^2 - m_P^2) - \frac{3\lambda^2}{4\pi^2} \int_0^1 dx [M^2 - m_P^2 x(1-x)] \ln \frac{M^2 - m_P^2 x(1-x)}{\mu^2} \quad (4)$$

In equation (4) the pole mass is determined by experiment and used as a renormalisation condition. Thus, this value does not get corrections at any order in perturbation theory. But the difference between the pole mass and the \overline{MS} mass is determined theoretically as a correction: the difference is proportional to the square of the mass of the fermions which couple to

the scalar. The scalar mass gets quadratically divergent corrections, though these may be removed with counterterms.

This theory may be shown to be renormalisable: we may iterate renormalisation and, at each iteration, only finitely many corrections are needed. However, as the table below shows, this theory is unnatural: the bigger the separation between the fermion mass M and the experimentally measured boson mass m_P , the greater the correction. That is, the theory exhibits a very sensitive parameter dependence on high-energy values. So we end up with a correction of 140% rather than 1% for a change in the boson mass m_P : $125\text{GeV} \rightarrow 30\text{GeV}$.

M	λ	m_P	$m_{\overline{MS}}$	$m_P^2 - m_{\overline{MS}}^2(\mu)$	%
163GeV	0.93	125GeV	123.6GeV	$(18.6\text{GeV})^2$	1%
163GeV	0.93	30GeV	72GeV	$(66\text{GeV})^2$	140%

Schwartz further analyses what would happen if we took the theory to have a finite completion at the Planck scale. We would then take $\Lambda \sim 10^{19}\text{GeV}$. For a boson at 125GeV , this would lead to an extremely sensitive renormalisation. The renormalised physical mass for the boson is approximately equal to the difference between the high-energy mass and the cut-off: $m_P^2 \approx m_{\text{Planck}}^2 - \Lambda^2$. That gives the high-energy mass as $m_{\text{Planck}}^2 \approx (1 + 10^{-34})\Lambda^2$.

Given that the standard model is unnatural, if no new particles were to be discovered beyond the standard model and below the Planck scale, we have what's called the 'naturalness problem', which asks 'how could there be such a sensitive dependence between values for couplings at the Planck scale and empirically accessible scales?' While further philosophical analysis of this problem is worthwhile that's not my focus here. Instead I aim to clarify two different senses of 'sensitive dependence between length-scales'. One is apparent in cases of unnatural theories, and the other is found in all (effectively) renormalisable EFTs. Distinguishing between these two senses is interesting in its own right, and helps identify what accounts for the effectiveness of EFTs. The upshot of this subsection is that not all (effectively) renormalisable theories are natural.

4 Two Kinds of Autonomy

In this paper, I consider autonomy with a view to characterising the ways in which EFTs float free of EFTs defined at different energy scales. In this section, I explicate two senses of autonomy, which I term ‘autonomy from microstates’ ($\text{autonomy}_{\text{ms}}$) and ‘autonomy from microlaws’ ($\text{autonomy}_{\text{ml}}$). Having established that renormalisability and naturalness are conceptually distinct in previous sections, here I pair them with $\text{autonomy}_{\text{ms}}$ and $\text{autonomy}_{\text{ml}}$ respectively; towards the end of this section I relate autonomy to emergence and multiple realisability. In §5, I go on to argue that the effectiveness of EFTs should be understood as a consequence of renormalisability rather than naturalness.

$\text{Autonomy}_{\text{ms}}$ licences the development of theories in ignorance of the goings-on at much higher-energy scales. Where a theory is $\text{autonomous}_{\text{ms}}$, we may abstract from dynamics in empirically inaccessible regimes when testing that theory and constructing a local description of the world.

If one can absorb salient details from high-energy scales into the low-energy description, then the low-energy description is $\text{autonomous}_{\text{ms}}$. For example, the trajectory of a spherical bouncy ball, when dropped, may be predicted at macroscopic scales on the basis of just a few parameters.¹⁸ The internal cohesion and spherical symmetry of bouncy balls makes the bouncy ball theory $\text{autonomous}_{\text{ms}}$ from details of the particular microstate of the molecules which constitute the ball. This $\text{autonomy}_{\text{ms}}$ allows the macroscopic predictions of the bouncy ball theory to be accurate, while the precise movements of the molecules may be unknown.

Consider, by contrast, dropping a bag containing a collection of bouncy balls and tacks – this is a collection which may sometimes bounce and sometimes not bounce depending on the internal alignment of the constituent objects. This combined object will, thus, be chaotic in certain ways: it would require not only detailed knowledge of fine-grained initial conditions but also knowledge of the microstate and its variations (the relative movements of each object within the bag) in order to be able to predict how or if it will bounce. $\text{Autonomy}_{\text{ms}}$ is essential to the abstractions involved in describing a system at a lower-energy scale. Where it’s violated, as in the bag of balls and tacks case, the macroscopic description allows

¹⁸Cross (1999) includes the mass, diameter, coefficient of restitution and spring constant of the ball.

for probabilistic predictions at best. This illustrates the role of $\text{autonomy}_{\text{ms}}$ throughout physics: for example, there may be changes to the high-energy state consistent with dynamics of beyond the standard model theories, and it is only our theories' $\text{autonomy}_{\text{ms}}$ which allows such changes to be disregarded at currently accessible scales.

$\text{Autonomy}_{\text{ms}}$ is commonplace across science. A theory is $\text{autonomous}_{\text{ms}}$ if it is invariant with respect to variations at higher energies such as those which may differ on different experimental runs but leave the laws unchanged.¹⁹ The $\text{autonomy}_{\text{ms}}$ of the relevant sciences explains how, for example, expert cell biologists may be ignorant of high-energy physics. $\text{Autonomy}_{\text{ms}}$ would be violated if we had a strong coupling between energy scales such that abstraction was impossible (at least without introducing probabilities). $\text{Autonomy}_{\text{ms}}$ is often violated within a theory, in particular, if aspects of the theoretical description exhibit inter-scale dependencies.²⁰

Below, I argue that renormalisability leads to $\text{autonomy}_{\text{ms}}$, and that effective renormalisability leads to a more constrained form of $\text{autonomy}_{\text{ms}}$ —in both cases high-energy effects are absorbed into redefinitions of low-energy parameters.

A regularity which is $\text{autonomous}_{\text{ml}}$ is invariant with respect to changes in the laws or fixed parameters. That is, $\text{autonomy}_{\text{ml}}$ is autonomy with respect to (counterlegal) variations which do not occur; such variations may be imagined to take us to possible worlds with different laws. Where $\text{autonomy}_{\text{ms}}$ is invariance with respect to changes consistent with the laws, but which differ on different experimental runs, $\text{autonomy}_{\text{ml}}$ is invariance with respect to changes which distinguish, at most, different possible worlds.

If a theory is $\text{autonomous}_{\text{ml}}$ then it underdetermines the corresponding high-energy theory: the high-energy theory may have any of a range of laws or parameter values, and the world described by the low-energy theory would be unchanged. A striking feature of $\text{autonomous}_{\text{ml}}$ theories is that the same low-energy theory will apply to multiple worlds which have different high-energy laws.

¹⁹This will include variations in initial conditions and variation due to irreducible stochasticity.

²⁰Chaotic systems may also violate $\text{autonomy}_{\text{ms}}$: their dynamics do not allow for a low-high energy decoupling, nor is it possible determinately to predict the future condition of some such systems without knowledge of the fine-grained initial conditions.

If a theory is autonomous_{ms}, but it's compatible with very few high-energy theories, then it is not autonomous_{ml}. Although autonomy_{ms} and autonomy_{ml} are conceptually distinct, one might wonder whether any realistic autonomous_{ms} theory would violate autonomy_{ml}. A consequence of discussion in the rest of this section is that the case study considered in §3.1, which is renormalisable but unnatural, is a putative example of an autonomous_{ms} but not autonomous_{ml} theory. As such, the present discussion not only informs us about the nature of EFTs but reveals something interesting about inter-theoretic relations more generally: not all theories which exhibit extremely sensitive parameter dependence violate autonomy_{ms}; it turns out that, for putatively realistic theories, autonomy_{ms} and autonomy_{ml} may come apart. In any context where inter-theoretic relations are non-trivial (i.e. excluding cases of strong emergence) there will be some degree of inter-scale parameter dependence, and consequent autonomy_{ml} violation; however, even those theories which violate autonomy_{ml} in a particularly stark manner may still be autonomous_{ms}.

The connection between autonomy_{ms} and renormalisability should by now be fairly straightforward. Assuming a sufficient separation of scales, renormalisable theories allow for the absorption of details of higher-energy theories including couplings and mass terms into parameters at low energies. Thus the dynamics which are relevant to theoretical descriptions at high energies, are, if the theory is renormalisable, irrelevant at low energies. Insofar as a theory is renormalisable we may discard details of high energies and include their effects in modifications at low energies. Thus renormalisability implies autonomy_{ms}:

the dominant effect of integrating out non-renormalizable operators between a high UV cutoff scale and a low-energy scale (up to small corrections involving inverse powers of the large ultraviolet scale) was to produce modifications, potentially of order unity, in the couplings associated with marginal and relevant operators.

[Duncan (2012, p. 653)]

Once such couplings have been modified, any potential changes due to the non-renormalisable high-energy operators have been absorbed into the low-energy theory. Changes at high energies consistent with the high-energy dynamics are then no longer relevant to an accurate description of the world at the low-energy scale.

Importantly, however, a high-energy theory's renormalisability does not mean that just any low-energy theory constructed will be able to abstract from high-energy details. In particular, non-renormalisable couplings may only go to zero once the cutoff is taken to infinity. $\text{Autonomy}_{\text{ms}}$ is specified with respect to a particular set of variations: if a low-energy theory still includes non-renormalisable terms, and, as such, not all details from the high-energy theory can be absorbed, then the theory is not $\text{autonomous}_{\text{ms}}$ with respect to variation of such details. Theories which are only effectively renormalisable will be $\text{autonomous}_{\text{ms}}$ in a wide variety of contexts – however, insofar as we may not abstract away from all high-energy details there will be variations with respect to which the theory in question will not be $\text{autonomous}_{\text{ms}}$.

Given that EFTs are renormalisable or effectively renormalisable (otherwise they're unusable), we can declare that all EFTs are $\text{autonomous}_{\text{ms}}$ to some extent. The next step is to think about the sense in which EFT parameters depend on parameters of high-energy theories. There are two ways such dependence could manifest. Firstly, it could be that the high-energy parameters vary for the same low-energy system on different experimental runs – it could be that these parameters were sensitive to the goings-on in the high-energy system. But this would be a failure of $\text{autonomy}_{\text{ms}}$ – it would mean that the low-energy dynamics were failing to capture some pertinent aspect of the high-energy dynamics.

The other possible mode of dependence corresponds to the sensitivity of low-energy parameters to the values of the parameters at high energies. Such values are, according to our theories, fixed. It is a theoretical assumption that at any given energy scale, the value of the coupling is a particular number and no known dynamics change this.²¹

If an EFT is unnatural, then, as argued above, there is an extremely high degree of sensitivity between far removed energy scales; this sensitivity is certainly troubling to the conception of science whereby theories are viewed as autonomous in all respects from theories at different length-scales. However, as demonstrated in §3.1, we may formulate unnatural renormalisable theories which are, thus, $\text{autonomous}_{\text{ms}}$ but not $\text{autonomous}_{\text{ml}}$. Therefore, the only type of autonomy incompatible with

²¹There are proposals for theories of early universe cosmology which include mechanisms for theory selection – on such theories the laws and parameters may well change; such theories are, however, speculative, and we have good reason to think that, on time-scales relevant to EFTs, the laws and parameters are fixed.

the unnaturalness of theories is $\text{autonomy}_{\text{ml}}$.

Natural theories are $\text{autonomous}_{\text{ml}}$ because they are not particularly sensitive to putative variations in high-energy theoretical parameters. As such, were, *per impossibile*, the high-energy theoretical parameters to change, a natural theory's parameters would be unaffected. All theories, whether natural or unnatural, will exhibit some dependence on high-energy parameter values due to renormalisation, but that dependence will be vanishingly small for natural theories. $\text{Autonomy}_{\text{ml}}$ is invariance with respect to changes in laws or parameters at high energies, and naturalness, as defined above, is insensitivity of parameters to parameters at high energies. Thus natural theories are more $\text{autonomous}_{\text{ml}}$ than unnatural theories.

Having distinguished two kinds of autonomy, it's worth discussing how this relates to emergence. This issue is particularly pressing because emergence is generally presumed to involve autonomy. Indeed, Karen Crowther (2015, 2016) defends the view that emergence is "novelty and autonomy of the low-energy level compared to the high-energy level" (Crowther (2015, p. 421)).²² Crowther claims that a lower-energy theory is autonomous from the relevant high-energy theory if "it is impervious to changes in the high-energy system ... [i]n the physical examples discussed here, autonomy stems from the high-energy theory being *severely underdetermined* by the low-energy physics" (ibid., p. 433 , original emphasis).

Thus, it seems that Crowther does not distinguish between $\text{autonomy}_{\text{ms}}$ and $\text{autonomy}_{\text{ml}}$; as such her view may have the (unfortunate) consequence that only natural EFTs are to be regarded as emergent. While I endorse much that she says about EFTs, I propose a friendly amendment to her view whereby EFTs are emergent if they are novel and $\text{autonomous}_{\text{ms}}$ with respect to higher-energy theories. To complete the definition, an account of novelty is needed: the view on offer in Franklin and Knox (unpublished) suggests that novelty implies that new explanations are available which are not expressible in terms of the variables of the higher-energy theory. $\text{Autonomous}_{\text{ms}}$ EFTs may thus be regarded as emergent if they are novel in this sense.

²²Bain (2013a,b) also argues that low-energy EFTs are emergent and autonomous from their high-energy counterparts, but he links this to in-principle underderivability. However, Bain's arguments merely establish that a certain class of EFTs (those which appeal to mass-independent renormalisation schemes) are in-practice underivable. As in-practice underivability is independent of $\text{autonomy}_{\text{ms}}$, Bain's autonomy does not allow one to answer the question of the effectiveness of EFTs, or to comment on the inter-dependence of energy scales in the same generality as allowed by considering $\text{autonomy}_{\text{ms}}$ and $\text{autonomy}_{\text{ml}}$.

Taking the argument of this section together with that of §2.1, I claim that the scaling properties of parameters explain renormalisability and that renormalisability leads to autonomy_{ms}. Thus, the scaling properties of parameters explain the autonomy which allows for emergence.

This discussion also bears relations to distinctions sometimes made concerning multiple realisability (MR); see e.g. Bickle (2016) and Hüttemann, Kühn, and Terzidis (2015). MR may be defined such that phenomena are multiply realised only if the same phenomenon occurs in at least two different systems where no dynamical process can transform one system into the other. This may be distinguished from robustness, defined as invariance with respect to dynamically allowed changes. Autonomy_{ms} leads to robustness, and autonomy_{ml} makes MR possible. However, explanations which advert to multiply realised phenomena are more desirable than those which appeal to merely autonomous_{ml} phenomena; MR is of interest insofar as there are multiple different systems which exhibit the same higher-level phenomenon – explanations and theories which refer to that phenomenon may apply to a broader variety of real-world systems. Insofar as unification is an explanatory desideratum, such explanations are valuable. On the other hand, theories which are autonomous_{ml}, but only apply to one kind of system in the actual world, do not have these explanatory advantages.

5 The Effectiveness of EFTs

The last section gave us the distinction between the two kinds of autonomy and, thus, a clear answer to the question concerning how low-energy and high-energy EFTs inter-depend. Now it's time to return to the question with which I started: whence the effectiveness of EFTs? By which I mean, what property explains our capacity to construct theories in ignorance of the details at higher energies, or to start with some high-energy theory and discard aspects of the high-energy description in constructing an empirically successful low-energy theory?

Williams (2015) suggests that the success of the EFT programme provides grounds to assert the 'EFT dogma':

Ultimately, I claim, the reason that failures of naturalness are problematic is that they violate a "central dogma" of the effec-

tive field theory approach: that phenomena at widely separated scales should decouple. This central dogma is well supported both theoretically and empirically.

[Williams (2015, p. 87)]

I claim, *pace* Williams, that the EFT dogma, insofar as it is supported both theoretically and empirically by the success of the EFT programme, warrants the requirement that our theories are (effectively) renormalisable rather than natural. In other words, the aim of constructing effective theories with limited empirical domains warrants the stipulation that all EFTs be (effectively) renormalisable. Importantly, a sufficient separation of mass scales is also required in order for EFTs to be effective; without such a separation, renormalisability is not sufficient to allow construction of empirically adequate low-energy theories.

EFTs conceptually require decoupling between length-scales. I have argued that what's required for decoupling sufficient to explain the effectiveness of EFTs is autonomy_{ms}, and, thus, (effective) renormalisability.²³ An EFT can violate naturalness and autonomy_{ml} while still exhibiting sufficient decoupling to satisfy the EFT dogma. That is, insofar as an EFT is autonomous_{ms}, it may be constructed in ignorance of high-energy physics, or discard details of such high-energy physics. The autonomy_{ml} of a theory has no bearing on those capacities.

Does naturalness failure imply failure of autonomy_{ms}? If the answer to this question were 'yes', then we would need to know very precise details about the high-energy scales in order to generate predictions from low-energy unnatural theories. But this is not how unnatural EFTs work. The high-energy, short-distance facts on which unnatural EFTs depend are the same in every experiment and there is no theoretical mechanism for modifying such values. While unnatural theories do violate autonomy_{ml}, they may nonetheless be autonomous_{ms}.

I agree with Williams' identification of the EFT dogma: "the expectation that widely separated scales should largely decouple in EFTs". But this dogma, so stated, is ambiguous. The effectiveness of EFTs which is warranted by the success of the EFT programme, corresponds to autonomy_{ms},

²³For EFTs with sufficient separation of scales, renormalisability and autonomy_{ms} are co-extensional, however they are not generally equivalent: not all autonomous_{ms} theories need be renormalisable – for instance it's not clear what it would mean to claim that cell biology is renormalisable.

and is due to (effective) renormalisability and the separation of mass scales. The decoupling consequent upon autonomy_{ms} is sufficient for EFTs only to apply within a limited range of length-scales. While autonomy_{ml} would correspond to further decoupling of scales, that's not required for EFTs' effectiveness.

For renormalisable (and effectively renormalisable) theories, renormalisation allows us to absorb the dependencies on high-energy facts into corrections to parameters. Therefore it is the (effective) renormalisability, and the consequent autonomy_{ms} which accounts for the effectiveness of EFTs. As argued in §2.1, renormalisability is explained by facts about the scaling properties of the couplings, these properties thus help to explain the effectiveness of EFTs.

6 Conclusion

In this paper, I distinguished two theoretical properties – renormalisability and naturalness – with a view to evaluating which is responsible for the effectiveness of EFTs. While (effective) renormalisability allows for an (effective) decoupling of dynamics, a *sine qua non* for effectiveness, naturalness corresponds to a decoupling of parameters, which is not in itself essential to effectiveness.

I related these properties to two types of autonomy. Autonomy_{ms} is an autonomy of low-energy dynamics from changes a system may actually undergo. Failure of autonomy_{ms} implies that the dynamics do not decouple. Autonomy_{ml} is rather about a decoupling of laws and parameter values; such decoupling is not required for the development of predictively successful low-energy theories. Parameters are fixed for our world, putative changes are not changes the system may actually undergo. Autonomy_{ms} is thus required for the effectiveness of EFTs, while autonomy_{ml} is not.

As we are discussing inter-theoretic relations there will always be some violation of autonomy_{ml}, for there will be some dependence relations between parameters defined at different length-scales. Unnatural theories are those theories which violate autonomy_{ml} in a particularly stark or extreme way. I have demonstrated that such violation is nonetheless consistent with renormalisability. As renormalisability is sufficient for autonomy_{ms}, it is thus consistent for a theory starkly to violate autonomy_{ml} but to be

autonomous_{ms}. All renormalisable theories violate autonomy_{ml} somewhat, but only a small subclass of renormalisable theories are unnatural and, as such, severely violate autonomy_{ml}. This distinction allowed me to explain how EFTs might be renormalisable yet unnatural: such theories will be autonomous_{ms} but not autonomous_{ml}.

That theories may be autonomous_{ms} and not autonomous_{ml} is interesting and non-trivial: it tells us something about the relations between different levels of description. The world may be such that one can construct successful and accurate low-energy theories in ignorance of high-energy goings-on, and yet such low-energy theories may strongly constrain details of the as-yet-undiscovered high-energy theories.

EFTs are effective insofar as they are autonomous_{ms}. As renormalisability, together with a separation of mass scales is sufficient for autonomy_{ms}, renormalisability accounts for the effectiveness of EFTs. Renormalisability is not explained by general features of the renormalisation group but, in each theory, by the scaling behaviour of particular theoretical parameters such as the coupling and mass terms. Therefore, the structure of the theoretical parameters helps to explain the effectiveness of EFTs.

Acknowledgements

I am grateful to Mike Birse, Eleanor Knox, Bryan Roberts and Porter Williams for comments on various versions of this paper. Thanks also to audiences at the BPS annual conference 2017, the Commissariat à l'Énergie Atomique Saclay workshop on Effective Field Theories, and the Bristol Philosophy of Physics Group for helpful questions and discussion. This work was supported by a London Arts and Humanities Partnership Research Council Studentship.

References

- Bain, Jonathan (2013a). "Effective Field Theories". In: *The Oxford Handbook of Philosophy of Physics*. Ed. by Robert W. Batterman. Oxford University Press, pp. 224–254.
- (2013b). "Emergence in effective field theories". In: *European Journal for Philosophy of Science* 3.3, pp. 257–273.

- Batterman, Robert W. (2017). "Autonomy of Theories: An Explanatory Problem". In: *Nous*. DOI: 10.1111/nous.12191.
- Berges, Jürgen, Nikolaos Tetradis, and Christof Wetterich (2002). "Non-perturbative renormalization flow in quantum field theory and statistical physics". In: *Physics Reports* 363.4, pp. 223–386.
- Bickle, John (2016). "Multiple Realizability". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2016. Metaphysics Research Lab, Stanford University.
- Butterfield, Jeremy (2014). "Reduction, Emergence, and Renormalization". In: *The Journal of Philosophy* 111.1, pp. 5–49.
- Castellani, Elena (2002). "Reductionism, emergence, and effective field theories". In: *Studies in History and Philosophy of Modern Physics* 33.2, pp. 251–267.
- Cross, Rod (1999). "The bounce of a ball". In: *American Journal of Physics* 67.3, pp. 222–227.
- Crowther, Karen (2015). "Decoupling emergence and reduction in physics". In: *European Journal for Philosophy of Science* 5.3, pp. 419–445.
- (2016). *Effective Spacetime: Understanding Emergence in Effective Field Theory and Quantum Gravity*. Springer.
- Duncan, Anthony (2012). *The Conceptual Framework of Quantum Field Theory*. Oxford University Press, USA.
- Epelbaum, E. and Ulf-G. Meißner (2013). "On the Renormalization of the One-Pion Exchange Potential and the Consistency of Weinberg's Power Counting". In: *Few-Body Systems* 54.12, pp. 2175–2190. DOI: 10.1007/s00601-012-0492-1.
- Franklin, Alexander (2018). "On the Renormalization Group Explanation of Universality". In: *Philosophy of Science* 85.2. DOI: 10.1086/696812.
- Franklin, Alexander and Eleanor Knox (unpublished). *Emergence Without Limits: The Case of Phonons*. URL: <http://philsci-archive.pitt.edu/13397/>.
- Fraser, James D. (2017). "The Real Problem with Perturbative Quantum Field Theory". In: *The British Journal for the Philosophy of Science*. Forthcoming.
- Grinbaum, Alexei (2008). "On the eve of the LHC: conceptual questions in high-energy physics". In: *arXiv preprint arXiv:0806.4268*.
- Hartmann, Stephan (2001). "Effective field theories, reductionism and scientific explanation". In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 32.2, pp. 267–304.
- Hüttemann, Andreas, Reimer Kühn, and Orestis Terzidis (2015). "Stability, Emergence and Part-Whole Reduction". In: *Why More is Different: Philo-*

- sophical Issues in Condensed Matter Physics and Complex Systems*. Ed. by Margaret Morrison and Brigitte Falkenburg. Springer. Chap. 10.
- Laudan, Larry (1990). "Normative Naturalism". In: *Philosophy of Science* 57.1, pp. 44–59.
- Nogga, A., R. G. E. Timmermans, and U. van Kolck (2005). "Renormalization of one-pion exchange and power counting". In: *Phys. Rev. C* 72 (5), p. 054006. DOI: 10.1103/PhysRevC.72.054006.
- Petrov, Alexey A. and Andrew E. Blechman (2016). *Effective Field Theories*. World Scientific Publishing Co.
- Polchinski, Joseph (1984). "Renormalization and Effective Lagrangians". In: *Nuclear Physics B* 231.2, pp. 269–295.
- Schwartz, Matthew D. (2014). *Quantum Field Theory and the Standard Model*. Vol. 1. Cambridge University Press.
- Stewart, Iain W. (2014). "Effective Field Theory". EFT Course 8.851, Lecture Notes, Massachusetts Institute of Technology.
- van Kolck, U. (2002). "Recent developments in nuclear effective field theory". In: *Nuclear Physics A* 699.1-2, pp. 33–40.
- Wallace, David (2011). "Taking particle physics seriously: A critique of the algebraic approach to quantum field theory". In: *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics* 42.2, pp. 116–125.
- Williams, Porter (2015). "Naturalness, the autonomy of scales, and the 125 GeV Higgs". In: *Studies in History and Philosophy of Modern Physics* 51, pp. 82–96.
- Wilson, Kenneth G. and John Kogut (1974). "The Renormalization Group and The ϵ Expansion". In: *Physics Reports* 12.2, pp. 75–199.