Naturalness, the Autonomy of Scales, and the 125 GeV Higgs

Porter Williams

Columbia University

Abstract

The recent discovery of the Higgs at 125 GeV by the ATLAS and CMS experiments at the LHC has put significant pressure on a principle which has guided much theorizing in high energy physics over the last 40 years, the principle of naturalness. In this paper, I provide an explication of the conceptual foundations and physical significance of the naturalness principle. I argue that the naturalness principle is well-grounded both empirically and in the theoretical structure of effective field theories, and that it was reasonable for physicists to endorse it. Its possible failure to be realized in nature, as suggested by recent LHC data, thus represents an empirical challenge to certain foundational aspects of our understanding of QFT. In particular, I argue that its failure would undermine one class of recent proposals which claim that QFT provides us with a picture of the world as being structured into quasi-autonomous physical domains.

1 Introduction

This paper is an analysis of a conceptual problem that has been at the center of high-energy physics research for approximately 40 years: the problem of "naturalness". I introduce the effective field theory framework in which the problem arises, survey the many faces of naturalness in the physics literature with an eye toward clarifying its (oftentimes misunderstood) physical significance, and discuss the implications that a failure to solve the naturalness problem would have for the ontology of quantum field theory. This latter issue is particularly pressing after the first run of the LHC: the discovery of the Higgs boson with a mass of 125 GeV and no additional particles not predicted by the Standard Model has put significant pressure on proposed solutions to the main problem of naturalness in the Standard Model, the Hierarchy Problem (for details see "Supersymmetry: Theory," "Supersymmetry: Experiment," and "Dynamical Symmetry Breaking: Implications of the H" in [38])[1]. The motivation for and significance of naturalness in quantum

[1] This is not the place for a detailed discussion of Beyond the Standard Model phenomenology at the LHC, so I refer the reader to [17], [21], [12], and [28] for further discussions of the status of solutions to the Hierarchy problem after the first run of the LHC. For a pedagogical introduction to the phenomenology of such supersymmetric solutions, see [36].

field theory is a hotly contested topic, with some claiming it to be ill-motivated or a mere theorists' prejudice[2]; I argue that, to the contrary, naturalness is essentially a prohibition of sensitive correlations between widely separated physical scales. I further argue that this is an expectation which is well-motivated within the effective field theory framework, justified on both theoretical and empirical grounds. The fact that the prospects for naturalness appear dim in light of the discovery of the 125 GeV Higgs, then, constitutes an empirical challenge to our current understanding of certain foundational features of quantum field theory (QFT).

The structure of the paper is as follows. I begin by introducing the notion of an effective field theory (EFT), since it is in this context that the naturalness problem arises. In Section 3, I use 2 simple models – a massive Yukawa theory and a scalar field theory– to illustrate the essential features of the naturalness problem. This then leads to a survey of the physics literature on naturalness, where there is a remarkable discordance of opinion. In Section 4, I draw on general features of the structure of effective field theories as well as particular examples of past successful realizations of naturalness in particle physics to present a physically transparent characterization, arguing that naturalness is best understood as a prohibition of correlations between widely separated physical scales. This understanding of naturalness has existed in the physics literature for quite some time; indeed, it is clearly stated in some of the earliest discussions of the concept in the late 1970's[3]. However, over time a variety of technical conditions for ensuring naturalness have developed and the understanding of naturalness has shifted in a direction that obscures its physical content (for a historical overview, see [31]). I argue that understanding naturalness in the proposed way illustrates that these superficially discordant technical understandings of naturalness in the physics literature are simply diverse attempts to formalize a shared unease about correlations between widely separated physical scales. This understanding of naturalness then forms the basis of an argument that naturalness is a well-motivated expectation in particle physics whose apparent failure requires a significant revision of our understanding of the effective field theoretic description of nature. Section 5 discusses one such revision: I examine how such a failure affects recent proposals that quantum field theory supports an ontological picture on which our world consists of a hierarchy of quasi-autonomous physical domains.

2 Effective Field Theory

Problems of naturalness arise in effective field theories. An effective field theory (EFT) is a quantum field theory which is known to become inapplicable above some energy scale Λ. This energy scale is called the

[2] See, for example, [51], [42], or from a slightly different angle [5].
[3] See [46], [47], or [39].

ultraviolet cutoff[4] of the EFT. There are both conceptual and pragmatic reasons for treating QFTs as EFTs. I postpone the technical details of naturalness to Section 3; in this section, my aim is to introduce core conceptual and technical features of EFT. In particular, I claim that the severity of naturalness problems is amplified by the fact all empirically applicable QFTs are best understood as EFTs, and provide a short argument to that effect.

The central obstacle to formulating a QFT that is a plausible candidate for describing our world up to arbitrarily high energies is gravity. All matter fields couple to the gravitational field. At low energies the effects due to this coupling are negligibly weak but gravity becomes strong when considering physical processes occurring at energies on the order of the Planck scale Λ_{Pl}. This means that while one can ignore gravity when doing particle physics at relatively low energies – such as those energies currently being probed at the LHC – gravitational effects must be taken into account at higher energies. Now, it is well known that a quantum field theoretic description of gravity is not perturbatively renormalizable. New divergences arise at each order of perturbation theory and in order to describe gravitational processes occurring up to arbitrarily high energies, eliminating these divergences would require considering an infinite number of parameters in the Lagrangian. The upshot of these two features of gravity is that while gravitational effects must be incorporated as we experimentally probe higher energies, it appears that quantum field theory is not the appropriate framework for doing so[5]. The fact that no quantum field theoretic description of gravitation is available in the regime where it is needed most, then, leads one to conclude that any QFT which aims to describe the world should not be trusted at energies above, at the highest, the Planck scale.

For the general class of QFTs that possess non-abelian gauge symmetries[6] gravity is the only clear obstacle to treating them as fundamental. In particular, there is no clear mathematical obstacle to treating such theories as fundamental. These theories are called *asymptotically free*; the interactions get weaker at higher energies, ultimately going to zero as the ultraviolet cutoff is taken arbitrarily large. This is quite different from QFTs which are *not* non-abelian gauge theories, such as quantum electrodynamics (QED) or the ϕ^4-theory. When formulated on a background spacetime of dimension greater than or equal to 4, the strength of interactions in these theories gets stronger at higher energies, eventually diverging at some finite energy scale. These singularities are known as Landau poles, and indicate that these theories are not mathematically consistent to arbitrarily high energies – at some finite energy scale, they break down on

[4]In what follows, unless otherwise noted all usages of Λ will stand for a generic ultraviolet cutoff whose particular value is not of central importance.

[5]It is worth noting that there is an active research program predicated on the possibility that gravity may be a consistent quantum field theory to arbitrarily high energies after all, in the sense that the renormalization group equation describing the behavior of the gravitational coupling at different energy scales hits a fixed point as we consider the coupling's behavior at higher and higher energies. This is called "asymptotic safety"; see [37] for a review.

[6]To be precise, they must also not contain too many matter fields since large numbers of these fields can spoil the asymptotic freedom of the theory. In QCD, for example, the theory is only asymptotically free if it includes 16 or fewer quark flavors.

purely mathematical grounds[7].

These two cases cover theories that describe physics at low energies but which are incapable of being extended above some very high energy scale. However, the use of EFTs also has a strong pragmatic motivation. Very commonly one wants to study physics at a relatively *low* energy scale but possesses a theory that describes physics up to some much higher scale. Typically, the high energy theory includes degrees of freedom distinct from those that are dominant at the low scale; in quantum chromodynamics (QCD), for example, the dominant degrees of freedom at low energy are hadrons that do not even appear in the Lagrangian of the high energy theory describing quarks and gluons. Similarly, given a theory of atomic physics and a desire to study some low-energy process like ocean wave propagation, the dominant degrees of freedom are not those found in the high energy theory; the atomic constituents of the sea can be neglected in providing an accurate description of the propagation of ocean waves. In that case, one employs an effective hydrodynamical theory describing ocean waves as disturbances in a continuous fluid, ignoring the underlying atomic structure of the ocean. In these cases, the main benefit of EFTs is essentially a pragmatic one. One could in principle describe the low-energy phenomena using the full theory (or at least many think we could; I doubt whether anyone has actually tried to study ocean wave propagation using the Standard Model). The central issue is rather that not only is doing calculations in the full theory more complicated, it also generally yields a *less* informative qualitative description of the phenomena in question. Degrees of freedom appropriate for vastly different scales become mixed up and a tractable understanding of the physics becomes much more difficult[8]. Effective field theories provide a better understanding of which degrees of freedom are dominant at different energy scales and of the relationship between the dominant physical processes that at these different scales.

The upshot of this discussion is that, due mathematical inconsistency or the inability to incorporate gravitational degrees of freedom (or both), all quantum field theories that purport to describe the world come with an ultraviolet cutoff beyond which they become inapplicable. For physics below this scale, however, these theories provide the most accurate agreement between theory and experiment in scientific history. I turn now to the technical details of effective field theories.

2.1 Effective Field Theories: How do they work?

It is well known that interacting QFTs experience divergences in perturbation theory that come from including contributions from field modes of arbitrarily high momenta in calculations[9]. Eliminating these divergences is a two step process. First, one chooses a method for *regularizing* the divergent integrals. There are a number of regularization methods but the central idea is the following. Upon encountering a divergent

[7] For discussion in both perturbative and nonperturbative contexts, see [1], [14], [35], or [43].

[8] For a philosophical discussion of the difficulties of modeling systems over many scales of length, see [11].

[9] For the purposes of this paper I will ignore infrared divergences arising from field modes of arbitrarily low momenta.

integral I in a perturbative expansion, one introduces a new parameter ϑ so that the integral now becomes $I(\vartheta)$, a function of ϑ. The dependence of $I(\vartheta)$ on ϑ is such that as the regulator is removed by taking $\vartheta \to \infty$, the result is the original divergent integral I. However, for finite ϑ the regularized integral $I(\vartheta)$ is a finite, well-defined mathematical object on which one can perform meaningful mathematical operations. For example, a popular regularization involves replacing Minkowski spacetime with a hypercubic lattice with lattice spacing a. This renders divergent integrals finite by restricting them to include only momenta k such that $-\frac{\pi}{a} \leq k \leq \frac{\pi}{a}$. Sending the lattice spacing to zero then returns the originally divergent integral.

Step two is renormalization. The renormalization procedure identifies the terms in $I(\vartheta)$ which diverge as $\vartheta \to \infty$ and then modifies the original Lagrangian by subtracting off the divergent pieces of $I(\vartheta)$, leaving a quantity which is convergent and independent of the regulator after ϑ is eliminated. This is done by modifying the couplings of terms in the original Lagrangian by adding terms designed to subtract off the divergent parts of the regularized integrals; the original Lagrangian is called the bare Lagrangian, the added terms are called counterterms, and the sum of the bare and counterterm parts of the Lagrangian is called the renormalized Lagrangian. The values of the couplings in the renormalized Lagrangian are the measured, physical values of those couplings and must be taken from experiment. When one can eliminate all divergences occurring in perturbation theory with finitely many renormalizations of the bare couplings, a theory is *perturbatively renormalizable*. In theories which are *not* perturbatively renormalizable, like gravity, one must include infinitely many renormalized parameters to cancel all divergences. Since values of renormalized parameters must be taken from experiment, these theories were long thought useless for making predictions: one would need to perform infinitely many measurements before calculating anything at all.

This procedure of regularizing divergent integrals, modifying the Lagrangian to subtract off divergent contributions to integrals, and removing the regulator to reintroduce arbitrarily high momenta into the theory has produced much discomfort since it was introduced in 1940s. It was thought to be merely clever trickery for hiding inconsistencies that remained in the theory after the regulator was removed. An important feature of EFTs is that one *need not* eliminate the regulator in order to make sensible predictions; indeed, it turns out that low-energy physics is remarkably insensitive to the details of the theory at much higher energies, including details about how it is regulated. The rare cases in which a high degree of sensitivity is present are generally cases in which problems of naturalness exist.

It is an important feature of EFTs that their Lagrangians generally contain infinitely many terms, both renormalizable and nonrenormalizable[10]. For the sake of concreteness in the following discussion, consider a generic[11] effective scalar field theory regulated by an ultraviolet cutoff Λ, which excludes field modes of

[10] This is unproblematic because the regulator Λ is never removed and problematic divergences do not arise. For details, see [20].

[11] The symmetry $\varphi \to -\varphi$ has also been imposed for the sake of simplifying expressions.

momenta $|k| \geq \Lambda$ from the theory. Integrating the Lagrangian over all spacetime to obtain an action and introducing couplings a_n for each term in the Lagrangian, the result is the Wilsonian effective action

$$S_W = \int d^4x \mathcal{L}_\Lambda(\varphi, \partial_\mu\varphi, \ldots) = \int d^4x \frac{1}{2}(\partial_\mu\varphi_\Lambda)^2 + \sum_{n\geq 0} a_n \varphi_\Lambda^{2+n} + \sum_{n\geq 0} a'_n (\partial_\mu\varphi_\Lambda)^2 \varphi_\Lambda^n + \cdots$$

where the field $\varphi_\Lambda(x)$ is defined to be

$$\varphi_\Lambda(x) = \int_{|k|<\Lambda} \frac{d^4k}{(2\pi)^4} \widetilde{\varphi}(k) e^{-ik\cdot x}$$

For the $a_0(\varphi_\Lambda)^2$ and $a_2(\varphi_\Lambda)^4$ terms, the couplings are the familiar mass term and the dimensionless quartic coupling λ, respectively. However, there are an infinite series of additional interactions described by operators with mass dimension greater than or equal to 5. Since the action S_W is a dimensionless quantity, these higher-dimensional operators must be multiplied by negative powers of some dimensionful quantity. Indeed, dimensional analysis shows that for S_W to be dimensionless, the couplings a_n must have mass dimension $2-n$ and the a'_n must have dimension $(-n)$. The only dimensionful quantity available is the UV cutoff Λ; accordingly, one defines the couplings as some dimensionless quantity g_n or g'_n multiplied by powers of the cutoff, $a_n \equiv g_n\Lambda^{2-n}(n = 0, 2, 4, \ldots)$ and $a'_n \equiv g'_n\Lambda^{-n}(n = 2, 4, 6, \ldots)$.

Dimensional analysis can accurately estimate how the various operators will contribute to scattering amplitudes for processes occurring at a given energy scale $E \ll \Lambda$. The operators will make contributions proportional to some power of the energy scale E; specifically, operators containing $2+n$ powers of the field and no derivatives will contribute $\sim E^{n-2}$, while operators containing m derivatives and n powers of the field contribute $\sim E^{n+m}$. Including the contributions of the coupling constants, the contributions of every term in the action S_W to a process at energy E can be estimated as $g_n\left(\frac{E}{\Lambda}\right)^{n-2}$ or $g'_n\left(\frac{E}{\Lambda}\right)^{n+m}$. Of great importance for the issue of naturalness are operators of dimension *less than four*: in the present theory, these are operators containing no derivatives and for which $n = 0$ or $n = 1$, meaning their contributions are proportional to the ratio $\left(\frac{\Lambda}{E}\right)$ to some positive power. This means that their contribution grows for processes occurring at lower and lower energies. Such operators are called *relevant*[12]. To see that there is already a *prima facie* problem with relevant operators, note that the coefficient $g_n\Lambda^2$ of the relevant operator φ^2 represents the *mass* of the φ field. The problem is that this effective theory is intended to describe physics *below* the ultraviolet cutoff Λ; how could it possibly contain a field of mass Λ^2? This issue will become central when turning to discuss naturalness in Section 3.

Even setting this question aside for the moment, one issue remains. Although this theory no longer

[12]One can distinguish three classes operators depending on whether they scale independently of Λ, as a positive power of Λ, or a negative power of Λ; these correspond respectively to marginal, relevant, and irrelevant operators.

includes arbitrarily high momenta, it will still include field modes over a large range of momenta in calculations. For example, consider a scattering process involving φ taking place at a center of mass energy $E \approx 1$ GeV. This theory will incorporate effects from processes ranging from arbitrarily low momenta up to the cutoff scale, a range which may be very large. Not only is this calculationally unpleasant, it also makes a qualitative understanding of the scattering process more difficult. Luckily, effective field theories can improve this. Beginning with a theory that includes momentum modes up to the cutoff Λ, one can recursively eliminate additional momentum scales from the theory, resulting in a new theory that breaks down at a new momentum scale $\mu \ll \Lambda$. Briefly, this is done as follows. Starting with the original effective theory with cutoff Λ, one can integrate out modes of the field φ_Λ with momenta $\mu \leq |k| \leq \Lambda$ from the functional integral, yielding a new effective theory applicable only up to the scale μ. In this new theory, the field becomes φ_μ and the Lagrangian now contains polynomials of this new field φ_μ and its derivatives. This procedure can be iterated in the obvious way, and by doing so one can recursively reduce the range of momenta included in calculations in order to isolate the contributions coming from below given energy scale to the physical process being examined, occurring at some scale much lower than the ultraviolet cutoff.

Of course, although the low energy predictions of EFTs are generally insensitive to the high energy physics integrated out of the theory as the cutoff is reduced, the high energy physics does contribute to some degree to physical processes at the low energy scale. How is this sensitivity incorporated as we repeatedly integrate out high energy physics?

Remarkably, all of the sensitivity of the low energy theory to the high energy physics is carried by the *couplings* a_n and a'_n. After integrating out momentum modes with $\mu \leq |k| \leq \Lambda$, the result is a new effective theory described by the Lagrangian $\mathcal{L}_\mu$ describing the new lower-energy field φ_μ. This theory will *also* have modified couplings $a_n(\mu)$ and $a'_n(\mu)$ that now depend on the new cutoff scale μ, replacing an original dependence on the cutoff Λ. Of course, iteratively lowering the cutoff through the procedure described above causes the couplings $a_n(\mu)$ and $a'_n(\mu)$ to change with each new cutoff scale. Imposing as boundary conditions at the initial cutoff scale Λ that $a'_0(\Lambda) = 1$, $a_n(\Lambda) = a_n = g_n \Lambda^{2-n}$ and $a'_n(\Lambda) = a'_n = g'_n \Lambda^{-n}$, the change in the couplings as the ultraviolet cutoff is changed can be described by a set of differential equations called the *renormalization group equations*. As I discuss in Section 3.1, precisely this process of integrating out high energy physics and modifying the couplings of operators according to the renormalization group equations provides a particularly physically transparent way of understanding problems of naturalness.

2.2 The Decoupling Theorem

Before turning to discuss naturalness, it is necessary to introduce a theorem which will be important to the discussion going forward: a theorem of Appelquist & Carazzone [4], often called the Decoupling Theorem, which describes aspects of the relationship between two widely separated energy scales in a quantum field theory. Appelquist and Carazzone prove that if one starts with a perturbatively renormalizable theory – the "full theory" – containing a set of fields all of which are much heavier than the remaining fields, the correlation functions describing physics occurring at energies much lower than the mass of the heavy fields can be obtained from an effective theory which contains *only* the light fields. The only remnant of the heavy fields is that the couplings of the light fields in the effective theory may be different than the couplings of the light fields in the full theory. Furthermore, for quantum field theories containing only fermions or gauge bosons, no relevant operators appear[13] and the contribution of the heavy fields to the light field couplings is merely logarithmic, a small correction to the original light field couplings. Virtually all quantum field theories employed in elementary particle physics are of this form; until the recent discovery of the Higgs particle, no elementary particle described by a scalar field had been discovered in our world. This is of central importance for naturalness since the presence of an elementary scalar field in a theory introduces a relevant operator whose coupling, representing the mass of the scalar field, receives very large corrections from high-energy physics. Setting aside for the moment the conceptual complication introduced by relevant operators, the essence of the Decoupling Theorem is that low-energy physics can be accurately described by a theory including only low-energy degrees of freedom. Past experience with physical theories leads one to expect that low-energy physics like billiard ball trajectories will not depend sensitively on the behavior of high-energy atomic physics; in QFT this observed fact that physical processes at widely separated physical scales are largely independent of one another becomes a theorem[14].

3 Naturalness

> Now we play the effective field theory game, writing down all terms allowed by symmetry and seeing how they scale. If we find a relevant term we lose: the theory is unnatural [41] (p.14).

[13] This may be because the particles the theory describes are massless, or because relevant operators violate certain symmetries of the Lagrangian; we return to this second possibility in Section 3.2.2.

[14] It may hardly seem worthy of proving a theorem to demonstrate something apparently obvious. The fact that such a theorem can be proven in QFT is surprising in part because perturbative corrections to correlation functions – so-called "loop" corrections – automatically include contributions from field modes of all momentum scales included in the theory's domain of applicability. The fact that these contributions can be confined entirely to corrections to the couplings is a noteworthy result.

3.1 We Have a Problem

I begin by outlining two calculations which put on the table all the essential features of the naturalness problem. The second calculation, done in a scalar field theory, illustrates how failures of naturalness manifest in the renormalization group flow for unnatural couplings. The first is done in a theory describing a massive scalar field ϕ with (bare) mass m and a massive fermion field Ψ with (bare) mass M, interacting via a Yukawa interaction with coupling g. Note that there are *two* relevant operators in this theory[15]: the ϕ^2 term, which has dimension 2, and the $\overline{\Psi}\Psi$ term, which has dimension 3. Based on our above discussion, we might expect these terms – in particular, the couplings of these terms – to give us trouble; recall that they were the terms that became multiplied by some positive power of the cutoff in the EFT in Section 2, apparently entailing that the effective theory intended for describing physics below the cutoff Λ included a particle with a mass many orders of magnitude larger than that cutoff.

First, assume the fermion is much lighter than the scalar so that $m \gg M$. The Lagrangian describing this theory[16] is

$$\mathcal{L} = \frac{1}{2}\partial_\nu\phi\partial^\nu\phi - \frac{1}{2}m^2\phi^2 - \frac{\lambda}{4!}\phi^4 + i\overline{\Psi}\gamma^\nu\partial_\nu\Psi - M\overline{\Psi}\Psi + g\phi\overline{\Psi}\Psi$$

Consider studying physics taking place at some low energy scale $E \ll M$. Then the Decoupling Theorem states that the appropriate thing to do is use perturbation theory to calculate the effects of the heavy scalar field in the full theory, integrate this scalar field out of the path integral as described above, and incorporate the effects of the heavy scalar field that we calculated using the full theory into the couplings of the light fermion field in the low energy effective theory. This will produce an effective theory involving *only* the light fermion field Ψ, with modified effective couplings M^* and g^*.

Up to terms of $\mathcal{O}(1/m^4)$, the resulting effective Lagrangian is

$$\mathcal{L} = i\overline{\Psi}\gamma^\nu\partial_\nu\Psi - M^*\overline{\Psi}\Psi + \frac{g^*}{2m^2}(\overline{\Psi}\Psi)^2$$

Consider the mass term $M^*\overline{\Psi}\Psi$. The coupling M^* is the effective mass of the fermion in the low energy theory; in general, it is equal to the bare mass M plus corrections ΔM coming from the high energy physics. What do these corrections look like? Including the first perturbative correction coming from the high energy physics (i.e. the one-loop correction), the effective fermion mass becomes

[15] This is another way to classify relevant operators. In d spacetime dimensions, any operator with mass dimension less than d is relevant, and will have to be multiplied by positive powers of the UV cutoff – the only dimensionful scale in the theory – in order to ensure the action is dimensionless.

[16] I have suppressed the spacetime dependence of the fields, the spinor indices on the fermions and the γ matrices, and the dependence of the fields on the chosen UV cutoff Λ

$$M^* = \left[M + M\frac{g}{16\pi^2}\ln\left(\frac{\Lambda}{M}\right)\right] = (M + \Delta M)$$

Note that that the correction ΔM to the bare mass is proportional to the bare mass itself. Since the bare mass of the fermion was small in the full theory, this means that it will remain small in the low-energy effective theory; the corrections it receives are just the bare mass itself multiplied by a quantity roughly of order 1.

Now consider an alternative scenario in which the fermion Ψ is much heavier than the scalar ϕ, so that $M \gg m$. Again, one can integrate out Ψ to obtain an EFT, now with effective couplings $(m^*)^2$ and g'^* that each receive corrections from the integrated-out high energy physics. The effective scalar mass is given by $(m^*)^2 = (m^2 + \Delta m^2)$ where Δm^2 is calculated using the full theory. Unlike the effective fermion mass, the effective scalar mass is proportional to the cutoff squared:

$$(m^*)^2 = m^2 + \frac{g}{16\pi^2}\left[\Lambda^2 + M^2 + m^2\ln\left(\frac{\Lambda}{M}\right) + \mathcal{O}\left(\frac{M^4}{\Lambda^4}\right)\right]$$

This is radically different than the previous case; here, starting with a small scalar mass perturbative corrections generated a scalar mass far larger than the scale Λ at which *the full theory itself* is expected to break down. Something has gone terribly wrong.

Indeed, this seems to produce a straightforward contradiction[17]. We began by claiming that we would integrate out all heavy fields – the fermions, in this case – leaving a theory involving only a light scalar field ϕ. However, any attempt to do this instead produces a theory containing a ϕ with a mass term proportional to the cutoff squared!

This is one face of the naturalness problem. In the Standard Model, this is called the Hierarchy Problem: the scalar field is the Higgs field and the heavy fermion could be the top quark. The scale at which the Standard Model breaks down, Λ, is unknown but is commonly thought to be, at the very highest, the Planck scale; at the lowest, evidence from the first run of the LHC now suggests that the Standard Model is applicable up to approximately 1 TeV. Either way, the fact that the mass of the recently discovered Higgs boson is 125 GeV presents a problem: even a correction of $(1 \text{ TeV})^2$ produces a mass for the Higgs boson 6 orders of magnitude larger than its measured value.

For a second, particularly physically transparent way of presenting the issue, consider the problem of naturalness through the lens of the renormalization group. Recall the statement in Section 2.1 that the appropriate tool for examining the relationship between different scales in an effective field theory was the

[17] Actually, this apparent contradiction was present when we wrote down the effective field theory for the single scalar field φ_Λ above. Recall that we found then that mass a_0 was proportional to Λ^2, meaning that it shouldn't have appeared in the effective theory at all.

renormalization group. In that spirit, I include a calculation using the renormalization group which makes salient the extreme sensitivity of relevant parameters to high energy physics, as well as the stark contrast between the sensitivity of relevant and irrelevant operators to such physics. Specifically, the calculation demonstrates how the respective values of relevant and irrelevant parameters at low energy depend on their initial, high energy values at the cutoff scale; in the example of the scalar EFT in Section 2, these initial values would be the couplings $a_n(\Lambda) = g_n \Lambda^{2-n}$ and $a'_n(\Lambda) = g'_n \Lambda^{-n}$. What the RG provides is another perspective on the sensitivity of such couplings to their initial values at the UV cutoff scale Λ_H when they are evaluated for some other, much lower UV cutoff $\Lambda_L \ll \Lambda_H$.

Consider the renormalization group equations (RGEs) for a toy scalar theory including only two couplings: one marginal, dimensionless coupling g_4 (corresponding to an operator with mass dimension 4) and one irrelevant coupling g_6 with mass dimension (-2) (corresponding to a dimension 6 operator)[18]. Recall that the RGEs determine how the values of the couplings change as we iteratively lower the chosen value Λ_H for the UV cutoff. In this case these equations take the form

$$\Lambda \frac{d}{d\Lambda} g_4 = \beta_4(g_4, \Lambda^2 g_6)$$

$$\Lambda \frac{d}{d\Lambda} g_6 = \frac{1}{\Lambda^2} \beta_6(g_4, \Lambda^2 g_6)$$

where the β_4 and β_6 are some unspecified, generally complicated functions and the factors of Λ^2 are required by dimensional analysis. For simplicity, define $\lambda_4(\Lambda) = g_4$ and $\lambda_6(\Lambda) = g_6\Lambda^2$, so that the RGEs become

$$\Lambda \frac{d}{d\Lambda} \lambda_4 = \beta_4(\lambda_4, \lambda_6)$$

$$\Lambda \frac{d}{d\Lambda} \lambda_6 - 2\lambda_6 = \beta_6(\lambda_4, \lambda_6)$$

One can solve these RGEs by perturbatively expanding the β-functions in λ_4 and λ_6. There are several ways to justify this expansion – one can expand around some exact solution, or assume that β_4 and β_6 are small – but here simply assume such an expansion is allowed. To first order in the couplings, this gives

$$\Lambda \frac{d}{d\Lambda} \lambda_4 = a\lambda_4 + b\lambda_6$$

[18] For the details of the calculation of the behavior of the irrelevant coupling see [44] (§ 23.6.1). My calculation below for the relevant coupling follows an essentially identical outline.

$$\Lambda \frac{d}{d\Lambda}\lambda_6 = c\lambda_4 + (d+2)\lambda_6$$

where a, b, c, d are small coefficients. One can solve this system of RGEs to obtain general solutions for $\lambda_4(\Lambda)$ and $\lambda_6(\Lambda)$ which depend on the initial values of these functions at the cutoff scale[19]. To see how the irrelevant coupling λ_6 at low energies $\Lambda_L \ll \Lambda_H$ depends on its initial value at the cutoff scale Λ_H, choose some arbitrary value of $\lambda_6(\Lambda_H)$ – for simplicity, let it be zero. Combining the solutions to the RGEs to express $\lambda_6(\Lambda)$ as a function of $\lambda_6(\Lambda_H)$ and $\lambda_4(\Lambda)$ yields

$$\lambda_6(\Lambda) = \lambda_4(\Lambda) \frac{2c\left[\left(\frac{\Lambda}{\Lambda_H}\right)^{\Delta} - 1\right]}{(2+d-a+\Delta) - (2+d-a-\Delta)\left(\frac{\Lambda}{\Lambda_H}\right)^{\Delta}}$$

where $\Delta = \sqrt{4bc + (d-a+2)^2}$.

Since $a, b, c, d \ll 1$ to ensure the validity of perturbation theory, to a good approximation $\Delta \approx 2$. To examine the dependence of λ_6 at low-energy on its chosen value of zero at the cutoff scale set $\Lambda = \Lambda_L \ll \Lambda_H$ to get

$$\lambda_6(\Lambda_L) = \frac{c}{2}\left[\left(\frac{\Lambda_L}{\Lambda_H}\right)^2 - 1\right]\lambda_4(\Lambda_L)$$

The limit $\Lambda_H \to \infty$ exists, and taking that limit shows that the initial value of λ_6 at Λ_H is totally irrelevant to its low energy value; at low energies, it is a function of only the low energy value of the marginal coupling $\lambda_4(\Lambda_L)$.

Repeating this process with a *relevant* coupling, the result is quite different. Consider the same marginal coupling g_4 and a new relevant coupling g_2 with mass dimension 2. One can similarly rewrite these couplings as $\lambda_4 = g_4$ and $\lambda_2 = \frac{1}{\Lambda^2} g_2$ and the RGEs are then

$$\Lambda \frac{d}{d\Lambda}\lambda_4 = a\lambda_4 + b\lambda_2$$

$$\Lambda \frac{d}{d\Lambda}\lambda_6 = c\lambda_4 + (d-2)\lambda_2$$

Solving these gives us general solutions for $\lambda_4(\Lambda)$ and $\lambda_2(\Lambda)$. Now combine these into an equation for the value of the relevant coupling $\lambda_2(\Lambda)$ as a function of $\lambda_4(\Lambda)$ and the initial, cutoff-scale value $\lambda_2(\Lambda_H)$ for the relevant coupling. Unlike for the irrelevant coupling, let $\lambda_2(\Lambda_H)$ have some nonzero value. The resulting

[19] See [44] (§ 23.6.1) for details.

equation is rather ugly, but simplifying it by again noting that $\Delta \approx 2$ gives the following expression for $\lambda_2(\Lambda)$:

$$\lambda_2(\Lambda) = \lambda_4(\Lambda)\frac{c}{2}\left[\left(\frac{\Lambda}{\Lambda_H}\right)^2 - 1\right] + \lambda_2(\Lambda_H)\frac{cb}{2}\left[\left(\frac{\Lambda}{\Lambda_H}\right)^{-2} - 1\right] + \lambda_2(\Lambda_H)\left(\frac{\Lambda}{\Lambda_H}\right)^{-2}$$

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 $\left(\frac{\Lambda_L}{\Lambda_H}\right)^2 = 10^{-28}$. Now alter the 20th decimal place of $\lambda_2(\Lambda_H)$, sending $\lambda_2(\Lambda_H) \rightarrow \left(\lambda_2(\Lambda_H) + 10^{-20}\right)$. 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! This is in sharp contrast to the low-energy behavior of irrelevant operators, which were entirely independent of their value at the cutoff scale. It is precisely this remarkable sensitivity of relevant couplings at low energies to their initial, cutoff-scale values that makes relevant operators unnatural; that a small change in the value of a coupling at the Planck scale could produce an enormous jump in the value of that coupling at a scale 14 orders of magnitude lower flies directly in the face of a "central dogma" of effective field theory discussed in the following section.

3.2 ...What Kind of Problem, Exactly?

Something has gone awry in the construction of this effective field theory. We began by integrating out all of the heavy fields – a single heavy fermion, in this case – in order to study low-energy physics using only the light scalar field. However, in the process the scalar mass was made extraordinarily heavy. Since it was claimed at the outset that the end result would be an EFT describing only light fields, this seems to be a contradiction[20].

What went wrong? In what follows, I outline five common ways of stating the naturalness problem in the physics literature. I then argue that while many of these capture certain features of the naturalness problem, all of them fail to adequately capture the heart of the problem. Ultimately, I claim, the reason that failures of naturalness are problematic is that they violate a "central dogma" of the effective field theory approach: that phenomena at widely separated scales should decouple. This central dogma is well-supported both theoretically and empirically. Theoretical support is to be found most clearly in the Decoupling Theorem and in the successful applications of the renormalization group; it also enjoys some (weak) inductive support from the fact that prior to the discovery of the Higgs, we had never encountered a quantum field theoretic description of elementary particles that required the inclusion of any problematic relevant operators. Em-

[20]In Section 3.2.3 I will show that there is a way out of this apparent contradiction – one can set the value of the bare mass m of the scalar to a very particular value in order to cancel the enormous corrections it receives, but this does not turn out to be a particularly satisfying way of dealing with the problem.

pirical support comes from the fact that (i) there are several examples from the history of particle physics in which enforcing this central dogma led to predictions of novel phenomena that were later confirmed by experiments, and (ii) the precise agreement between Standard Model predictions and experimental data entail that any corrections coming from currently unknown high-energy physics must have a small effect on the comparatively low-energy phenomena described by the Standard Model. To see why failures of naturalness violate this central dogma, I now turn to an examination of the many characterizations of naturalness in the physics literature.

3.2.1 Quadratic Divergences

Perhaps the most common way of stating the naturalness problem is as the "problem of quadratic divergences." The corrections to the effective scalar mass from high energy physics are quadratic in the high energy cutoff Λ, so removing the cutoff by taking $\Lambda \to \infty$ at the end of the calculation causes the scalar mass in the effective theory to diverge quadratically. This way of stating the problem is extremely widespread. For example, [34] points out that "the naturalness problem for scalar mass parameters is often said to be a consequence of the fact that scalar mass parameters are quadratically divergent in the UV" (p.8). For a larger (but far from exhaustive) sample, see [40] (p.788), [22] (p.314), or [8] (p.190).

This formulation is misleading in at least two ways. The first is relatively benign: because we have adopted an effective field theory perspective, there is no reason to let $\Lambda \to \infty$ and so nothing diverges. A better way of describing the problem would be as a quadratic *sensitivity* to the high energy scale Λ where the theory breaks down. Stated this way, the problem concerns the high degree of sensitivity of a low-energy parameter – a light scalar mass – to physics at some very high energy scale Λ.

The second problem is more serious. Thus far, a sharp momentum cutoff Λ has been employed in order to exclude field modes of arbitrarily high momenta from the theory, regularizing integrals that would otherwise be divergent. The problem is that one could choose a different way of regulating the theory, in which case calculations do not yield any quadratic divergences whatsoever. To demonstrate this, consider repeating the same calculation using *dimensional regularization* in order to render divergent integrals finite.

In dimensional regularization field modes of arbitrarily high momenta remain in the theory, but calculations are carried out in a *non-integer* number of dimensions $(d-\epsilon)$ for some appropriate number of dimensions d. Using this regulator, *all* potential divergences will now appear as corrections multiplied by $\frac{1}{\epsilon}$; when one analytically continues back to 4 dimensions by taking $\epsilon \to 0$, these terms will diverge. The divergences are removed by adding counterterms to the Lagrangian, as outlined in Section 2.2.1. For present purposes, the most significant feature of calculations performed using dimensional regularization is that terms quadratic in

Λ simply do not appear[21]! For example, using dimensional regularization to calculate the correction to the effective scalar mass in the above Yukawa theory gives $\Delta m^2 = \frac{g^2}{16\pi^2}M^2$ where M is the mass of the heavy fermion. Note, however, that a quadratic dependence on the heavy fermion mass M remains. Thus even in the absence of quadratic divergences, the light scalar field is sensitive to the heavy fermion in a way that again to produces a contradiction: the "light" scalar has a mass on the order of the heavy particles that were supposed to have been integrated out of the theory to begin with.

Treating the naturalness problem as a problem of sensitivity of low-energy parameters to high energy physics, we can see that the problem is in no way dependent on the regulator we choose. However, considered as a problem of quadratic divergences, misconceptions are forthcoming: [2] (p.236) warn that "as [dimensional regularization] regularizes the quadratic divergences to zero it seems that the whole hierarchy problem results from using a clumsy regulator, and that by using [dimensional regularization] we could shield the Higgs mass from the scale of new physics." In a similar spirit, [48] (p.56) describes such a conclusion as a "knowing-just-enough-to-be-dangerous naive way to look at the Standard Model" arising precisely from equating the appearance of quadratic divergences with the naturalness problem. Eliminating the quadratic divergences by performing calculations with a different regulator really does not eliminate the problem at all; thus, quadratic divergences cannot actually have been the central issue in the first place.

3.2.2 Symmetry Principle

Recall the first calculation done in the Yukawa theory: the scalar field of mass m was much heavier than the fermion field of mass M, and an effective theory suitable for studying physics at energies $E \ll m$ was constructed, describing only the fermion field Ψ. Both the mass term $\overline{\Psi}\Psi$ for the fermion and the mass term ϕ^2 for the scalar field are relevant operators, so one might expect both of these operators to present naturalness problems. However, the corrections to the effective fermion mass from high energy physics (represented by the heavy scalar) were small – proportional to the fermion mass M times a logarithm. However, repeating the calculation for a heavy fermion and a light scalar, the scalar received enormous corrections. Why is the fermion so much less sensitive than the scalar to high energy physics?

One answer has to do with symmetry properties of the Lagrangian. When writing down the original Lagrangian we actually broke the rules of the "effective field theory game" described above by Polchinski: a term disallowed by the symmetries of the theory was included. Setting the fermion mass to zero increases the symmetry of the theory: with no fermion mass term, the left and right-handed components of the Dirac spinor Ψ do not mix and the theory has a *chiral symmetry.* Adding a fermion mass term breaks the chiral

[21] One can, in fact, find a kind of "quadratic divergence" in dimensional regularization, but this arises in the form of additional divergences for $d < 4$, not as divergences quadratic in ϵ. See [2] Ch.12.

symmetry, but that symmetry should be restored in the $M \to 0$ limit; this is only the case if the perturbative corrections to M are proportional to M itself. In that case, letting $M \to 0$ sends the bare mass appearing in the Lagrangian as well as all corrections to zero and the chiral symmetry is restored. If the fermion mass received additive corrections proportional to Λ^2 or some power of the scalar mass m, even if we started with a chirally symmetric theory perturbative corrections would break the chiral symmetry. It is a remarkable fact about perturbation theory that perturbative corrections do not generate symmetry breaking terms: if setting some parameter to zero enhances the symmetry of the Langrangian, perturbative corrections to that parameter will be proportional to the parameter itself[22]. This is why the initially light fermion was able to remain light; the small fermion mass was "protected" by the chiral symmetry.

On the other hand, no symmetry is added to the theory by sending the scalar mass $m \to 0$. The scalar mass term does not break any symmetry, so there is no reason why perturbative corrections should not generate a scalar mass term even if we start with an initially massless scalar field in the Lagrangian. No symmetry "protects" the small scalar mass and perturbative corrections need not be proportional to m; instead, they are quadratic in some very heavy scale and thus cause the scalar to become very heavy.

This has led to another common way of stating what it means for a parameter to be natural, due to [47]:

> at any energy scale μ, a physical parameter or set of physical parameters $\alpha_i(\mu)$ is allowed to be very small only if the replacement $\alpha_i(\mu) = 0$ would increase the symmetry of the system...this is what we mean by naturalness (p.353).

't Hooft's connection of naturalness with symmetry has led some to suggest that the naturalness criterion is (or could have been) justified as a type of symmetry principle. [31], for example, claims that "[b]ased upon 't Hooft's definition, it could have received a...conceptual foundation similar to that of symmetry," where that conceptual foundation is taken to be partially aesthetic and partially pragmatic (p.616)[23].

Defining the principle of naturalness as equivalent to 't Hooft's symmetry principle offers an impoverished characterization of naturalness for several reasons. First, 't Hooft has only presented a sufficient technical condition for naturalness: a small parameter is natural *if* setting it to zero increases the symmetry of the theory[24]. This simply tells us how to mathematically ensure that a theory is natural, without giving much insight into the physical content of the principle of naturalness itself. Much more troublingly, it turns out that naturalness has no necessary relationship with symmetry after all. Perhaps the paradigmatic "naturally" small physical parameter is the the QCD scale Λ_{QCD}. The question of why this scale is small compared to

[22] The exception to this rule are so-called "anomalies". The chiral symmetry in the Yukawa theory is not anomalous, so these are of no consequence here.

[23] See also [30] (§ 5.3).

[24] One class of theories with naturally small parameters which do not satisfy 't Hooft's criterion – setting them to zero does not restore any symmetry in the Lagrangian – are those which ensure this naturalness via some dynamical mechanism. The smallness of the QCD scale relative to the Planck scale is perhaps the most well-known realization of such a mechanism.

the Planck scale dates back to Dirac (see [17] (p.7)), and the explanation of the smallness of this parameter has nothing essential to do with any symmetry principle. Rather, its value is explained by dimensional transmutation, which is a feature of the dynamics of QCD (see [8] (§ 9.12)). Similarly, one proposed natural solution to the Hierarchy problem is called Technicolor: in these models the Hierarchy problem is solved by extending the Standard Model, appending a QCD-like sector at higher energies that removes the sensitivity of the Higgs boson mass to high-energy physics through the same dynamical mechanism that stabilizes Λ_{QCD}. Defining naturalness as a symmetry principle leaves one in the awkward position of either requiring a separate account of naturalness describing cases in which it is achieved through a dynamical mechanism like dimensional transmutation, or describing a paradigmatically natural scenario as "unnatural" due to the inapplicability of the symmetry principle.

3.2.3 Fine-Tuning

In almost any discussion of naturalness, one encounters the claim that unnatural theories require some form of "fine-tuning." This has led many to classify problems of naturalness as fine-tuning problems. The fine-tuning enters in several places, and always concerns relevant operators not protected by any symmetry. Consider the generic EFT constructed for a single real scalar field in section 2. Recall the scale of all dimensionful parameters in the effective theory had to be set by the only dimensional scale present in the theory, the UV cutoff Λ, with the result that the coefficient a_0 of the relevant operator ϕ^2 was written $a_0 \equiv g_0\Lambda^2$. Since the coefficient of ϕ^2 represents the mass of the ϕ field, this suggested that the ϕ field didn't belong in the low-energy effective theory at all.

This is the point where fine-tuning first enters. Imagine that the ϕ field is the Higgs field, with physical mass 125 GeV. In order to set the mass parameter $a_0 = m_\phi^2$ in the Lagrangian equal to the physical mass, one must assign a very small value to the dimensionless coupling g_0, on the order of $\frac{m_\phi^2}{\Lambda^2} \sim 10^{-34}$ GeV for a UV cutoff at the Planck scale. Many physicists find such tiny dimensionless ratios unsettling; as [52] (p.419) puts it, they "naturally expect that dimensionless ratios of parameters in our theories should be of order unity...say anywhere from 10^{-2} or 10^{-3} to 10^2 or 10^3."[25] The idea behind this discomfort is essentially that the dimensional analysis arguments that were employed in Section 2.1 in order to estimate the size of the contribution of a given operator at a given scale E should always work: the contribution of any operator in an effective theory to a physical process occurring at some energy scale E should be determined more or less entirely by the physical scales involved in the problem, which are just E and the cutoff Λ. Setting a dimensionless parameter g_n in such a way as to vitiate this sort of reasoning runs counter to the sorts of dimensional arguments that are ubiquitous, and almost always successful, in effective field theories.

[25] See also [20] (p.578).

Essentially the only instances in which they break down are when relevant operators are involved.

However, imagine that one somehow come to grips with this fine tuning, and sets $g_0 \sim \frac{m_\phi^2}{\Lambda^2}$. The need for fine-tuning does not end there. Recall the earlier calculation of the effective scalar mass in the Yukawa theory: regardless of the value of the bare scalar mass parameter appearing in the Lagrangian, it will receive quadratic corrections from high-energy physics, driving the value of the effective mass far above its physical value. Thus, in order to keep the effective mass near the value of the physical mass, one will again have to set the bare mass m^2 in such a way as to cancel very delicately the corrections Δm^2 coming from high-energy physics.

One sometimes hears the following objection at this point (cf. [13] (§1.4)): why can't one simply absorb these large corrections into a large (but finite) renormalization of the bare mass m? It is, after all, standard practice to add counterterms to bare quantities in order to subtract off divergences in perturbation theory – that is precisely the renormalization procedure outlined in Section 2.1. As [9] (p.3) puts the objection, "aren't we supposed to talk only of physical renormalized quantities, with all divergences suitably reabsorbed?" Why not simply do that here?

In fact one can, in principle. However, this provides temporary relief at best, and ultimately doesn't allow one to avoid fine-tuning. The reason for this is straightforward[26]. As discussed in §2, the process of integrating out high-energy physics to obtain a low-energy Wilsonian action S_W is recursive; having integrated out the high energy degrees of freedom between Λ and some lower scale μ, one can then repeat the process and produce a new S'_W appropriate for studying physics below some new scale μ'. In general, this will involve integrating out more and more heavy particles from the initial Wilsonian action S_W and each these will make large contributions to the value of the relevant coupling – in this case, the effective mass of the scalar. However, to absorb the problematic corrections coming from high energy physics, one must pick some single scale at which to renormalize the bare mass m; having done so, one cannot do it again. One can't simply re-renormalize the mass with every move to a new S'_W appropriate for a new cutoff scale μ'. Thus, even if one absorbs the large corrections into a renormalization of m at some scale μ, equally problematic corrections will appear for every subsequent move to lower energies; even after renormalization, the RG flow of the *physical*, renormalized mass remains sensitive to high-energy physics. Ultimately, there is no option but to simply fine-tune m in the high energy theory in order to cancel very delicately Δm^2, the sum of *all* corrections from successively integrating out heavy particles when moving lower energies, in such a way as to ensure that the effective mass in the low energy theory will reproduce the physical value of the mass of the scalar particle.

It is undeniable that there is fine-tuning taking place: certain free parameters are required to take very

[26]For an excellent discussion of this in the context of the Cosmological Constant, see [13].

specific numerical values for no reason other than that this produces correct predictions. However, to lump naturalness problems in with other fine-tuning problems is to ignore central and qualitatively distinctive features of problems of naturalness. Furthermore, simply describing naturalness as a prohibition on fine-tuning is to make a fairly weak argument for taking naturalness seriously, while a much stronger one is available, as I argue in Section 4.

There are a number of puzzles in physics that get classified as fine-tuning problems – the flatness and horizon problems in cosmology, the apparent low entropy of the early universe, and so on – where the fine-tuning in question is qualitatively quite different than that present in naturalness problems. In particular, the degree of fine-tuning in naturalness problems, such as the Hierarchy problem or the problem of the cosmological constant, is determined by the size of the ratio of scales over which parameters need to be correlated. Indeed, what one could call the informal measure for the "unnaturalness" of some parameter is simply the ratio of the physical scales across which the parameter's low and high energy values must be correlated; for example, physicists often speak of the Hierarchy problem as requiring a tuning of the Higgs mass to ~ 34 decimal places, which is simply the ratio of the physical Higgs mass m_H^2 to the Planck scale Λ_{Pl}^2. In contrast, the other listed instances of fine-tuning involve no such correlations between widely separated physical scales. The charge that a low-entropy initial state of the universe necessitates fine-tuning, for example, stems from the small Lebesgue measure of such low-entropy regions in the phase space of the early universe; it has nothing to do with sensitivity between widely separated scales.

On the other hand, if naturalness problems really are just fine-tuning problems, there is a good argument that one popular way of understanding effective field theories suggests that we shouldn't care about them. This popular way of understanding effective field theory goes roughly as follows (see [26] and [32] for discussions). Accepting that the theory one is working with is not valid up to arbitrarily high energy scales, one simply takes the measured, low-energy scalar mass as an empirical input parameter that must be taken from experiment. This is sometimes accompanied by a hope that the UV completion of the low-energy EFT, whatever it may be, will eventually explain the values of the low energy parameters. However, stuck with an effective theory, one is free to arrange the values of free parameters at high energy – those appearing in the original Lagrangian at the original cutoff scale Λ – however is needed to make accurate predictions for empirically accessible low energy physics, which is all the EFT can reasonably purport to describe anyway. The only concern in doing this is that one isn't fooled into thinking that by being forced to make specific choices for the values of the high-energy parameters, they have thereby learned something meaningful about physics near the cutoff scale. Christof Wetterich[27] to the captures this view succinctly: "we do not need to

[27]Note that in context, Wetterich's point is not that naturalness problems are not genuine problems. Rather, it is that they are best understood as sensitivity to mass thresholds in the RG flow of renormalized, physical parameters, as I discussed above in this section and again in Section 4.1.

know the exact formal relation between physical and bare parameters (which furthermore depends on the regularization scheme), and *it is not important if some particular expansion method needs fine-tuning in the bare parameters or not.* The relevant parameters are the physical parameters, since any predictions of a model must finally be expressed in terms of these" [50] (p.217; emphasis mine).

Now, one is free to resist this understanding of effective field theory, and I do not intend my discussion of it to be an endorsement. However, it is indeed common in the physics community. Thus, on the charitable assumption that broad swaths of the physics community do not hold manifestly inconsistent views about problems of naturalness, the fact that this understanding of effective field theory (i) strongly suggests that we should not worry about fine-tuning at high-energies and (ii) is fairly widely accepted by the same physics community that takes naturalness problems quite seriously is at least a *prima facie* worry for simply equating naturalness problems with numerical fine-tuning. As the discussion of formal measures of fine-tuning in the following section will show, there is even greater reason to resist defining problems of naturalness to be merely problems of fine-tuning.

3.2.4 An Aesthetic Criterion

Often in conjunction with talk of fine-tuning, one finds naturalness problems described as problems of aesthetics. In many such cases, declaring naturalness problems to be aesthetic problems is just a way of declaring them not to be genuine *physical* problems at all. For example, [31] concludes that the force of naturalness problems stems largely from "down-to-earth sociological factors" in the particle physics community, involving a unique interpretation of formal measures of fine-tuning and heavily influenced by factors like "the choice at the leading universities of professors with a particular taste in physics" and the historically contingent fact that certain physicists who were "the first to fix the arbitrary convention of what is natural and what is not" maintain much influence as a result. Importantly, Grinbaum does *not* think that naturalness problems have the force of "normal scientific arguments"; they are not, in short, properly *physical* problems. Similarly, [45] (p.4) states that naturalness is a philosophical constraint, which one should feel free to ignore if they don't like it. On the other hand, even those who take naturalness quite seriously often speak in this way: [3], [6], and [19] each refer to a preference for natural theories as some variety of an "aesthetic preference," all while arguing that this preference should be taken seriously and developing quantitative measures of fine-tuning to facilitate comparisons of naturalness between models.

In order to understand these claims, I think it is important to recognize that all of these authors consider naturalness problems to be fine-tuning problems *and* take quite seriously formal measures of fine-tuning. Conflating these formal measures and their use in the physics community with naturalness itself, I claim, is plausibly what leads such authors to conclude that naturalness is an aesthetic or sociological restriction.

First of all, fine-tuning problems in general are always susceptible to being dismissed as pseudo-problems. In large part, this is because they involve dissatisfaction with certain features of a theory that present neither strict logical nor empirical problems (and generally don't run counter to deeply embedded features of the theory's structure); rather, the problem is alleged to lie elsewhere, often diagnosed as a failure of the theory to explain the finely-tuned feature. One way to deny this is a problem is simply to adopt an especially pragmatic attitude toward physical theories, such as the understanding of effective field theory discussed near the end of Section 3.2.3. Any inclination to dismiss such problems in general is then compounded in discussions of naturalness, I claim, by focusing on the results given by formal measures of fine-tuning.

While there are a variety of these measures in use in physicists' discussions of naturalness, all of them determine the degree of fine-tuning present in a selected model by something like the following procedure (see [21] for a very clear discussion): First one chooses a specific measure, which will be a function of some set of parameters of the model. Then one selects some set of parameters from the high energy theory and another set in the low energy theory, where the low energy parameters are functions of the high-energy ones. Common choices of the low energy parameters in discussions of the Hierarchy problem, for example, are the mass of the W and Z bosons, which are functions of the Higgs vacuum expectation value, which is in turn a function of the Higgs mass. The degree of fine tuning in this model is then the maximum "sensitivity" of the low energy parameters to changes in the high energy parameters. This is typically captured by taking the maximum value of all the derivatives of the low-energy parameters with respect to the the high-energy ones – the exact form of this prescription will vary from measure to measure. In the example of the Hierarchy problem, this would aim to capture the sensitivity of the physical W and Z boson masses to changes in the value of the bare mass of the Higgs appearing in the Lagrangian at the scale of the UV cutoff. There is much subjectivity in making these choices; which parameters one chooses to consider at low and high energy can drastically change the amount of fine-tuning assigned to a model by the chosen measure. Different measures will, of course, also assign different degrees of fine-tuning. Given the large number of subjective choices that go into quantitatively determining the fine-tuning of a given model, one can already see why one might conclude that "naturalness can at best be understood as a sociological heuristic" [31] (p.627).

Examining the history of quantitative naturalness prescriptions (as is done in [31] and [17]) seems to lend further support to this conclusion. This history reveals that "the acceptable value of [fine-tuning] has drifted up over the years, starting around ~ 10 and floating towards 100. Many would now call 1000 a reasonable value" [17] (pp.6-7). Not only is the choice of measure a somewhat arbitrary and sociologically influenced choice, even the acceptable value of the sensitivity parameter has changed within the physics community over time. As experimental constraints have increasingly put pressure on proposed natural solutions to the Hierarchy problem, what was once considered a degree of sensitivity far too high to be considered "natural" has come to be considered acceptable. This also makes naturalness look more like a prejudice of theorists than a "normal scientific argument" or even an expectation well-motivated by the effective field theory framework. As the notion of naturalness seems amenable to gradual tailoring to accommodate disconfirming empirical data, this seemingly supports claims about its predominantly sociological or aesthetic nature.

There are several things to say about this understanding. First, the idea that naturalness is an aesthetic problem in the sense that particle theorists are unwilling to tolerate "ugly" or "inelegant" theories is not consistent with popular opinions on low-energy supersymmetry (SUSY), the most popular natural solution to the Hierarchy problem. Indeed, while *unbroken* SUSY is commonly asserted to produce beautiful, elegant theories, many particle physicists openly acknowledge that when broken at low energies (as it must be in order to describe our world), SUSY is a rather unattractive theory. For a representative sample, consider that [21] (p.353) acknowledges "weak-scale supersymmetry is neither ravishingly beautiful (and hasn't been for decades) nor excluded," while [45] (p.4) states that "[a]lthough theoretically supersymmetry is a beautiful concept, the corresponding phenomenology was and still is less than elegant," and [42] complains that even though the invention of low-energy supersymmetry offered a natural solution to the Hierarchy problem, "[t]he price of this invention is 124 new constants, which I always thought was too high a price to pay... In this case a conceptual nicety was accompanied by an explosion in arbitrary parameters." Thus it appears that the concepts of naturalness and aesthetic appeal can and do come apart, even in the most popular natural solution to the Hierarchy problem[28]. Given the untenability of understanding the demand for naturalness as a demand for elegant or beautiful theories, then, it seems those quoted above describing naturalness as an aesthetic requirement are probably best understood as using "aesthetic requirement" as shorthand for something like Grinbaum's claim that naturalness problems do not constitute normal scientific arguments, presenting neither empirical nor logical problems for unnatural theories.

The second two things to say involve clarifying the usefulness of quantitative measures of naturalness. First, one should be sure not to lose sight of the original motivations for these measures. Importantly, when one looks at the physical principles that are invoked to justify these measures, a picture of naturalness emerges that is much closer to the requirement that widely separated scales not be sensitively correlated. [3] offer this physical motivation, describing the Hierarchy problem as a problem in which "[t]he appearance of [a] heavy mass scale in turn requires demonstrably large, unexplained cancellations among heavy masses in order to maintain a light weak-scale," arguing that the "fine-tuning" in this scenario manifests itself as an unexplained

[28]An anonymous referee raised the question of whether one is right even to look for beauty in applicable physical models, as opposed to the mathematical methods used to construct such theories. I will restrict myself to stating that the cited authors, at minimum, apparently think it reasonable to expect beauty in applicable models, as do those who lament the "inelegance" of the Higgs mechanism (see [23] and the many references therein). I have considerable sympathy with such views, and while I think the referee's question is an interesting one, it would require a far longer discussion than can be given here to be addressed responsibly.

sensitivity of the values of parameters at low energies to certain high-energy parameters. Similarly, in their presentation of the first quantitative measure of naturalness, [10] (p.73) motivate their naturalness measure by pointing out that "[t]here is no known example of cancellation between a quadratic divergence in the low energy theory and contributions from shorter distances," arguing that the sensitive correlations between the two scales required to ensure such a cancellation are unacceptable. The measures, then, are motivated by underlying physical ideas about the prohibition of correlations between scales; however, throwing away the ladder and focusing only on the numbers that the quantitative measures produce can easily cause one to forget about the original physical motivation and conclude that naturalness is simply a aesthetic or sociological (and flexible) constraint imposed on effective field theories by the particle theory community.

That said, despite their underlying physical motivations, I am reluctant to assign too much significance to quantitative measures of naturalness in general. We have already seen that there is a large amount of arbitrariness involved in choosing any particular measure, in choosing which parameters to consider at each scale after choosing a measure, and so on. This is only the beginning of the arbitrariness; both [17] and [21] offer excellent lists of similar difficulties with formal measures of fine-tuning. However, much more alarmingly, they also point out that these measures typically assign a high degree of fine-tuning to scenarios that we take to be *paradigmatically* natural, such as the dynamical preservation of the smallness of the QCD scale Λ_{QCD} with respect to the Planck scale! This strongly suggests that while the physical motivation underlying these quantitative measures – that low-energy physics should not depend sensitively on physics at much higher energies – is eminently reasonable within the EFT framework, the measures are often unsuccessful at capturing it; [17] (p.7) underscores this by pointing out that "[o]ne frequently comes across models that are constructed using a legalistic interpretation of naturalness that fails an intuitive sniff test." In light of the fact that these quantitative measures involve large amounts of arbitrariness and render incorrect judgments of paradigmatically natural scenarios, I endorse the suggestion of [17] that rather than placing undue weight on quantitative measures, "we exercise our physical judgment when weighing naturalness." Since, as we have seen, the physical motivation for these quantitative measures amounts to a prohibition on low energy physics depending sensitively on physics at much higher energy scales, it seems that once we take Craig's suggestion seriously, this is the most appropriate understanding of naturalness to adopt.

4 Naturalness and Interscale Sensitivity

In section 3 we saw that naturalness is subject to many different, sometimes seemingly unrelated descriptions in the physics literature. However, lurking in each of these descriptions was an underlying idea that in the effective field theory framework, sensitive correlations between high and low energy physics are prohibited.

Here I present a positive argument that this is indeed the appropriate way to understand the naturalness criterion: we should understand naturalness as the requirement that theories should be able to describe physics at low energies in ways that do not invoke a sensitive dependence on those theories' descriptions of physics at much higher energies. I present three (largely independent) claims in support of this understanding of naturalness: (i) it provides a uniform notion which can undergird the myriad descriptions of naturalness seen above, (ii) it enables us to make a much more compelling argument for naturalness from within the effective field theory framework than any of those offered above, and (iii) it offers an accurate description of the reasoning on display in commonly cited historical examples of successes (and failures) of the naturalness principle in quantum field theory.

4.1 A Uniform Notion of Naturalness

I will begin by briefly recapitulating and synthesizing the conclusions of Section 3 that understanding naturalness as a prohibition on correlations between widely separated scales unifies the apparently discordant notions of naturalness in the physics literature. In the case of quadratic divergences (or more accurately, quadratic sensitivity to physics at the cut-off scale), simply choosing a different regularization method eliminated the problem, making it clear that naturalness has nothing to do with the appearance of perturbative corrections quadratic in the cutoff. However, a correction quadratic in the mass of the heavy fermion remained. This suggested that understood as a requirement that low energy physics – in this case, the scalar mass – not be sensitively dependent on high energy physics, the notion of naturalness is independent of the choice of regulator. This conception remains able to justify the widespread discomfort with quadratic sensitivity to cutoff scale physics – this is just the special case in which the high energy scale to which the low energy physics is sensitive is the scale of the UV cutoff Λ. Likewise, understanding naturalness as a symmetry principle excluded paradigmatically natural scenarios in which the independence of low and high energy physics is ensured by a *dynamical* mechanism, such as the preservation of the smallness of the QCD scale Λ_{QCD} through dimensional transmutation. I argued that the relationship between naturalness and symmetry was much weaker than is sometimes claimed; the presence of a symmetry simply amounts to a sufficient technical condition for eliminating sensitivity of certain low-energy parameters to high-energy physics.

Attempts to identify naturalness with some kind of aesthetic preference of theoretical physicists also encountered difficulty: it sits in tension with a number of theorists' acknowledgment that the most popular natural solution to the Hierarchy problem is not at all elegant, and the identification seemed to rest heavily on problematic formal measures of fine-tuning[29]. Furthermore, the underlying motivation for these formal

[29] Ironically, identifying naturalness as an aesthetic preference of theorists also obscured the physical content of naturalness.

measures was the belief that low energy physics should not depend sensitively on physics at high energy; the measures are (problematic) attempts to formalize this physical requirement. Simply taking this last observation as the criterion for naturalness allowed the maintenance of a physically motivated standard of naturalness while jettisoning talk of "aesthetics" or "sociological factors" along with its concomitant problems. Finally, acknowledging that there is certainly fine-tuning that must occur in unnatural theories, I argued that simply equating naturalness problems with fine-tuning problems was to give short shrift to naturalness. In particular, by focusing on the "unlikely" numerical precision required when setting certain free parameters or on the output of specific formal measures of fine-tuning, it ignored the ways in which naturalness problems were qualitatively different from other fine-tuning problems in physics by obscuring the important role that interscale sensitivity plays in problems of naturalness. Furthermore, identifying problems of naturalness as fine-tuning problems opens an avenue for denying that naturalness problems are anything to worry about at all. However, treating naturalness problems as fundamentally concerning sensitive dependencies between scales, we found that this underwrote widely used "informal" measures of fine-tuning, preserved the physical motivation underlying the formal fine-tuning measures, and made it much more difficult to deny the seriousness of naturalness problems from within the effective field theory framework. The upshot of the review of the myriad naturalness notions in the physics literature in Section 3, recapitulated here, was that they could all be unified by adopting the physically motivated understanding of naturalness that I am advocating.

Consider, now, the "central dogma" of effective field theory mentioned in Section 3.2: the expectation that widely separated scales should largely decouple in EFTs. What justifies this expectation? Of course, that such decoupling occurs in other areas of physics is a convenient and ubiquitous fact of life; the trajectory of a billiard ball is insensitive to small changes in its atomic constituents. However, in QFT one can do better than simply pointing to historical precedent. Indeed, as discussed in Section 2.2 and 3.2, the expectation of decoupling is based on a number of theoretical and empirical results in effective field theory which can ensure that, in various senses, low energy observables are not sensitively dependent on high energy physics. Perhaps the most significant theoretical result that justifies this expectation is the Decoupling Theorem. As outlined in Section 2.2, the theorem tells us that in a QFT satisfying certain conditions, one can construct an EFT suitable for describing low energy physics by integrating out all fields with masses above some high energy threshold. The only consequence of this integration is that the couplings in the low energy effective theory will be different than those in the full theory, where this change of the couplings incorporates all of the contributions coming from the high energy physics.

The irony is that it is precisely that physical content which motivates the formal measures of fine-tuning, which then turn around and play an important role in motivating an identification of naturalness as an aesthetic criterion which ignores the original physical content.

There are limitations, however, on when we can invoke the Decoupling Theorem to license our expectations about interscale insensitivity. In particular, we cannot point to the Decoupling Theorem to justify our decoupling expectation unless (i) our starting point for constructing a low energy EFT is a *perturbatively renormalizable* high energy theory, and (ii) we renormalize the theory using a *mass-dependent* renormalization scheme. In fact, in many cases one or both of these requirements are not satisfied. It is typically most convenient to regularize and renormalize using dimensional regularization and a mass-*independent* scheme like $\overline{\text{MS}}$, and it is part and parcel of the effective field theory approach that when writing down a QFT, we include all terms – renormalizable and nonrenormalizable alike – consistent with cluster decomposition, Lorentz invariance, and the symmetries of the theory. Nonrenormalizable operators will be suppressed by powers of the high-energy cutoff and can be ignored for most purposes, but strictly speaking virtually any EFT formulated according to the prescription in Section 2 will be nonrenormalizable. However, this alone should not undermine our belief in the decoupling of scales in QFT. The reason is that we also have excellent empirical evidence that scales do, in fact, largely decouple. On the one hand, calculations using renormalization schemes in which the decoupling of scales is not manifest and must be inserted by hand yield excellent agreement with experiment; this empirical fact is ultimately what justifies performing decoupling by hand in mass-independent renormalization schemes (see [26] or [7]). On the other, the spectacular empirical successes of the Standard Model are evidence for the fact that whatever new physics may await at higher energy scales, it has little influence on the low energy physics currently being probed at the LHC and described with remarkable precision by the Standard Model. Understood, then, as the principle that low energy physics should not depend sensitively on physics at high energies, naturalness is a physically transparent and well-motivated expectation within the effective field theory framework.

In fact, amidst the historical shift toward talk of quadratic divergences, fine-tuning, and aesthetic preferences, this early understanding of naturalness has been maintained by various authors. Evaluating how the results of the first run of the LHC bear on prospects for naturalness, [9] says that the core difficulty of the Hierarchy problem is that "it is against the notion that the physics at the Fermi scale should not depend on details of what happens at shorter distances...The quadratic divergence of the Higgs mass is not the problem per se, but the sign of the sensitivity of the Higgs mass to any threshold encountered at higher energy scales." Describing how the naturalness criterion should manifest itself at the LHC, [27] tells us that "The naturalness criterion, as stated above, excludes the possibility that the parameters that encode the information of physics at very short distances are correlated with dynamics of the effective theory occurring at large distances. Such a correlation would signal a breakdown of the philosophy underlying the effective-theory approach." Writing after the first run of experiments at the LHC, [28] again reiterates that "[i]t is the consequence of a reasonable criterion that assumes the lack of special conspiracies between phenomena occurring at very different length scales. It is deeply rooted in our description of the physical world in terms of effective theories." Similarly, in his excellent review of the naturalness problem associated with the cosmological constant [13] writes that ' "cosmology appears to provide one of the few loopholes to the general observation that Nature decouples; that is, that the details of short-distance physics do not matter much for describing physics on much larger distance scales...Cosmology is special because the things that seem needed to describe observations – things like a vacuum energy or a *very light scalar field* [my emphasis] – are among the few kinds of things that *are* sensitive to the details of small-distance physics." Finally, at the risk of belaboring the point, we include a (too-strong) statement from [18]: "The Higgs boson mass diverges quadratically! The Higgs boson thus does not obey the decoupling theorem...This is known as the 'hierarchy problem'." Dawson's claim is actually incorrect – while the sensitivity of the Higgs mass certainly violates the spirit of the Decoupling Theorem, this sensitivity manifests itself only in a modification of the coupling of the mass term of the scalar, which is entirely allowed by the theorem. However, I include her remark because it captures rather dramatically the core physical idea that I am advocating here: that the central issue in naturalness problems is that low energy physics is unduly sensitive to physics taking place at high energies in a way that violates the "central dogma" of Section 3.2.

An important question may remain[30]: even granting that naturalness problems violate a "central dogma" of effective field theory, given that these correlations are unobservable, should one really worry about them? I think that one should, for several reasons. First, even if unobservable, naturalness problems do affect the RG flow of renormalized, *physical* parameters (see [9] (pp.2-3) or [13], §1.4). Recall from Section 3.2.3 that even if one renormalized the scalar mass at some chosen renormalization scale μ, the value of this coupling still sustains quadratic corrections each time it crosses a mass threshold, i.e. each time a heavy particle is integrated out of the theory as the ultraviolet cutoff is iteratively lowered. The renormalized, *physical* value of the coupling "jumps" dramatically with each mass threshold it crosses unless the initial renormalization of the mass is chosen to cancel every correction it sustains as the RG flows to lower energies. Although unobservable, this sensitivity seems to be a clearly *physical* problem. Second, from a methodological point of view, even if naturalness problems themselves do not have observable consequences, proposed solutions to such problems certainly do, and as I argue below attempting to solve naturalness problems has in fact proven a good guide to nature in the 20$^{\text{th}}$ century (though not infallible – see the discussion of the cosmological constant below). While neither of these logically necessitates that one take naturalness problems seriously, I do think that it makes a compelling case for viewing them as genuine physical problems worth taking seriously and attempting to solve.

Finally, I turn now to considering the major successes of the naturalness principle in the 20$^{\text{th}}$ century.

[30]I thank an anonymous referee for raising this question.

These are instructive for two reasons: (i) they further my argument that one has good reason to take naturalness problems seriously by demonstrating that this has been an accurate guiding principle for predicting new physics in the past and (ii) the specific reasoning employed in these past successes further supports my proposed understanding of naturalness. There are three major successes that are often cited as motivation for the naturalness principle, lending further inductive reason to expect effective field theories to be natural: the divergent self-energy of the classical electron leading to the prediction of the positron, the difference between the charged and neutral pion masses leading to the prediction of the ρ meson, and the mass difference between the "long" and "short" neutral kaon states K_L^0 and K_S^0 leading to the prediction of the mass of the charm quark. Here I focus only on the last of these three because while the first two are examples of situations in which naturalness arguments *would* have led to the eventual solution[31], in the case of the kaon mass difference a naturalness argument of precisely the sort currently being used to motivate extensions of the Standard Model actually *was* invoked to predict the mass of the charm quark. Instructive descriptions of the relationship of naturalness to the classical electron self-energy and the charged/neutral pion mass difference can be found respectively in [36] and [27].

Before addressing successes, however, I should first note its most infamous failure: the cosmological constant[32]. In the Standard Model the cosmological constant is simply the identity operator, which has mass dimension zero, and so must be multiplied by *four* powers of the ultraviolet cutoff in order to make the Standard Model action dimensionless. Without any fine-tuning, this results in a theoretically predicted value which differs from its measured value by anywhere from 50 to 122 orders of magnitude, depending on what one considers as "the" cosmological constant problem (see [13] §1); in any case, reproducing its measured value requires fine-tuning even more delicate than in the Hierarchy problem. Furthermore, no natural solution to the cosmological constant problem seems likely: as Burgess puts it, "[p]robably the most well-established attitude to the cosmological constant problem in particle physics is one of pragmatic despair" (p.17). As it stands, it appears that one may just have to accept the presence of such sensitive correlations between the cutoff-scale and low-energy values of the cosmological constant. This is just to remind the reader that when considering past successes of naturalness-based reasoning in what follows, one should not lose sight of the fact that expectations of naturalness can be and indeed have been mistaken in the past.

I begin with a very brief summary of some background to the prediction of the charm quark mass. In 1970, Glashow, Iliopouos, and Maiani (GIM) [29] introduced a feature of the theory of weak interactions – the GIM mechanism – to account for certain previously unexplained empirical properties of weak decays[33].

[31] For another, more speculative bit of counterfactual history based on naturalness arguments, see [49].

[32] For an excellent review, see [13].

[33] Of particular relevance here, the GIM mechanism banishes so-called flavor-changing neutral currents (FCNC) from the theory. The presence of these currents implied that certain kaon decay processes, such as $k \to \mu\overline{\mu}$, should occur at rates much higher than experimentally observed.

A necessary condition for the implementation of this mechanism was the existence of a fourth flavor of quark, called *charm*. There was no experimental indication at the time that such a quark flavor existed, but its existence was a prediction of the GIM mechanism. While the GIM mechanism itself made no firm predictions for the mass of the charm quark, the presence of the charm quark would have significant effects on a number of observable properties of kaon physics.

In [24], the authors considered one of these observable properties – the mass difference between the two neutral kaon states K_L^0 and K_S^0. Computed in an EFT appropriate for describing physics at the kaon scale, the mass difference is

$$\frac{M_L - M_S}{M_K} \approx \frac{G_F f_K^2}{4\pi} \sin^2\theta_c \cos^2\theta_c \, \epsilon_0 \approx \epsilon_0 \times 1.4 \times 10^{-3}$$

where f_K is the kaon decay constant and θ_c is the Caibbibo angle. What matters here is the ϵ_0 term, which contains a dependence on the then-unknown mass of the charm quark. Because the charm quark is not included in this effective theory of kaon physics, its mass functions as the scale at which the effective theory breaks down so we can equivalently view ϵ_0 as containing a dependence on the cutoff of the EFT. Experimentally, it was known at the time that $\frac{M_L - M_S}{M_K} \approx 7 \times 10^{-15}$. Now, Gaillard and Lee pointed out that there were two possible scenarios consistent with this experimental result: either (i) both the strange quark and charm quark masses are approximately equal and much larger than 1 GeV, but these high energy parameters cancel in such a way as to ensure the measured low-energy result for the K_L^0–K_S^0 mass difference, or (ii) the charm quark mass alone appears in ϵ_0, which would require it to have a mass of roughly 1.5 GeV. Gaillard and Lee argue for the second option, stating that "in the absence of a symmetry argument, the cancellation [of option (i)] appears purely fortuitous" [25] (p.279) – in more modern language, one might say that the cancellation between high energy parameters in order to reproduce the known experimental result would be unnatural. Famously, they were correct – the charm quark (discovered in 1974) has a mass of approximately 1.3 GeV [38]. Adhering to the requirement that the effective kaon theory be natural in the sense that it required no sensitive correlations between high and low energy physics, even though those correlations were had no directly observable effects, was in this case a very good guide to the physical world[34].

[34] An anonymous referee asked whether the symmetry-based understanding of naturalness could not also account for this prediction. My response to this is twofold. First, it is true that one could predict the mass difference between the charm quark and it's SU(2) partnered strange quark should be small, since setting the mass difference to zero would enhance the SU(2) symmetry. This gives only guidance regarding the mass *splitting* between the charm and strange quarks, but scant guidance as to the mass of the charm itself; perhaps both the charm and strange quark are approximately equal and both very heavy. Indeed, Gaillard and Lee consider and dismiss precisely this possibility, as I discuss above. Second, even if the symmetry-criterion could account for this, it would still be limited in cases of dynamically preserved naturalness in the ways I have discussed above.

5 Naturalness and the Ontology of QFT

I believe that this understanding of naturalness has consequences for recent attempts to evaluate the ontological picture presented by effective field theories. In particular, I believe that it teaches that one must be more careful when discussing whether effective field theories underwrite a picture of the world as consisting of "quasi-autonomous" physical domains, and that the apparent unnaturalness of the Standard Model vitiates recent attempts to describe the ontology of our world as a hierarchy of such domains. I will be intentionally imprecise about the exact meaning of "quasi-autonomous" in what follows; the reason for this is that I believe the discussion will make clear that any precisification of "quasi-autonomous" that aims to be strong enough to underwrite the ontological hierarchy advocated by such proposals will fail to be realized in EFTs which are unnatural.

Perhaps the most well-known (and certainly the most ambitious) attempt to argue that EFTs present a picture of the world as a hierarchy of quasi-autonomous domains is [15]. They claim that "[i]n this picture, [nature] can be considered as layered into quasi-autonomous domains, each layer having its own ontology and associated 'fundamental' laws" (p.72). They justify this claim by leaning heavily on the Decoupling Theorem, stating that

> with the decoupling theorem and the concept of EFT emerges a hierarchical picture of nature offered by QFT, one that explains why the description at any one level is so stable and is not disturbed by whatever happens at higher energies, and thus justifies the use of such descriptions (p.64).

Throughout [15], the authors make much of the stability and independence of EFT descriptions of physics at different energy scales, arguing that such stability and independence follows directly from the Decoupling Theorem[35]. What makes the proposal of Cao and Schweber particularly ambitious is that in addition to the hierarchal picture of nature just described, they also conclude that EFTs "support a pluralism in theoretical ontology, an antifoundationalism in epistemology and an antireductionism in methodology" (p.69). I will return briefly to this more extreme claim when discussing the criticism of [32], but for the moment I simply endorse the conclusion of [33] that this constitutes a significant overreach and is not at all entailed by the success of the effective field theory approach. However, one need not subscribe wholesale to Cao and Schweber's conclusions in order to hold onto the belief that an ontology of quasi-autonomous domains is licensed by the Decoupling Theorem.

That said, even the more modest ontological claim of Cao and Schweber has been criticized by [32] on

[35] See, for example, their claim about the implications of realistic interpretations of the UV cutoff Λ on (p.53) and the discussion of the decoupling theorem in § 3.4.

the grounds that the conditions required for the proof of the Decoupling Theorem may not always hold. In particular, the proof requires that we start with a renormalizable theory as the starting point for the construction of an EFT; as Hartmann rightly points out, such a starting point may not always be available to us. Indeed, insofar as the claim of Cao and Schweber that EFTs entail "antireductionism in methodology" means that physics is to be described by a never-ending tower of (in general, nonrenormalizable) effective theories, it would in fact be wrongheaded to expect such a renormalizable high energy theory to be lying at the base of the tower. If this is the understanding of "antireductionism in methodology" that Cao and Schweber have in mind, then, it seems that it straightforwardly undermines their heavy reliance on the Decoupling Theorem. Furthermore, [7] has emphasized that the proof of the Decoupling Theorem depends on the use of a mass-dependent renormalization scheme; if one uses a mass-independent scheme like dimensional regularization and minimal subtraction, one can no longer prove the theorem and the decoupling of scales is no longer manifest. So the criticism of Cao and Schweber's use of the Decoupling Theorem to underwrite an ontology of quasi-autonomous domains is threefold; their position seems internally inconsistent, its validity depends on a particular choice of renormalization scheme, and there is no reason to expect in general that the starting point will be an underlying renormalizable theory. These all entail that one cannot use the Decoupling Theorem to justify the claim that, in general, EFTs entail a picture of the world as a hierarchy of quasi-autonomous domains.

However, the criticisms of Hartmann and Bain undermine the very strong claim that the Decoupling Theorem, applied to EFTs *in general*, does not support such a picture; one can still reasonably ask whether *our* world seems amenable to such a description. In fact, it seems *prima facie* as though it might. Ignoring for the moment the reasons stated in Section 2 for treating all QFTs as effective theories, the Standard Model is a perturbatively renormalizable theory containing widely separated mass scales, and thus satisfies the conditions of the Decoupling Theorem. Furthermore, it describes the (non-gravitational) physics of our world extraordinarily well. Additionally, even if one prefers a renormalization scheme in which the Decoupling Theorem cannot be proved, one can still ensure decoupling by simply putting it in by hand, where the reader will recall that this is justified by the empirical fact that calculations done in this way give answers that are in extraordinary agreement with experiment[36]. Thus it seems that starting from the Lagrangian of the Standard Model, one can construct low-energy effective theories and invoke either the Decoupling Theorem or the observed empirical fact that high-energy physics does in general decouple from low energy observables in experiments to conclude that the appropriate ontological picture of *our* world is one in which nature separates into a hierarchy of quasi-autonomous physical domains after all. In fact, both [7] and [16] do tentatively endorse this "quasi-autonomous domains" ontology; I will return to their claims

[36]Again, see [7] or, for the technical details of this process, [26].

in a moment, after discussing the problems that failures of naturalness pose for this ontological picture.

In contrast to this picture of our world decoupling into quasi-autonomous domains, the calculations in Section 3 made it especially salient that certain parameters in an "unnatural" low-energy theory – the relevant couplings – depend sensitively on high-energy physics. Consider the renormalization group calculation: the low-energy value of the relevant coupling was extremely sensitive to variations in its high-energy value. There seems to be something very wrong with describing low-energy physics as "stable and not disturbed by whatever happens at higher energies" when changing the value of a relevant coupling at the Planck scale by a factor of one part in 10^{20} results in the value of that parameter evaluated at, say, 10^5 GeV – a scale 14 orders of magnitude lower – jumping by a factor of 10^8. Whatever precise meaning is given to the notion of a "quasi-autonomous domain," it should certainly prohibit this sort of extreme sensitivity between the physics in domains which are supposed to be largely independent of one another. However, in theories which are "unnatural" this kind of sensitivity is ineliminable; indeed, it is precisely what earns these theories the moniker "unnatural." In particular, the Standard Model is unnatural, containing two relevant parameters: the Higgs mass parameter and the cosmological constant. My claim, then, is that despite the fact that the Standard Model satisfies the conditions of the Decoupling Theorem, it still cannot underwrite even the more restricted picture of our actual world as a hierarchy of quasi-autonomous domains: the problem is not that the Decoupling Theorem fails to hold, but rather that the Standard Model is unnatural, entailing interscale sensitivity that vitiates the quasi-autonomy that was *prima facie* underwritten by the Decoupling Theorem.

Failures of naturalness, then, have ontological import. Furthermore, they raise an important methodological point for philosophers of quantum field theory: the Decoupling Theorem, *even in those situations in which it holds*, is too weak to underwrite an ontology of quasi-autonomous domains. At minimum, an additional necessary condition is that the EFT in question be *natural*.

What, then, of the conclusions of [7] and [16] that EFTs do support an ontology of quasi-autonomous domains? For Bain – and, as Bain plausibly argues, for Castellani as well – their conclusions are justified by the fact that in EFTs constructed using *mass-independent* renormalization schemes[37], decoupling can simply be performed by hand. Doing this, to reiterate, is justified by the fact that putting in this decoupling by hand gives remarkably accurate agreement between calculations and the results of experiments. For Bain, then, the question of whether EFTs support an ontology of quasi-autonomous physical domains depends on the renormalization scheme, albeit in a perhaps unexpected way: it is in EFTs renormalized according to precisely those schemes in which the Decoupling Theorem *cannot* be proven that we find support for such an ontology. Any attempt to underwrite such a picture in a mass-dependent scheme by appealing to the Decoupling Theorem, says Bain, falls victim to Hartmann's critique of Cao and Schweber, which was

[37] EFTs constructed in this way are often called continuum EFTs, following [26].

discussed above.

Failures of naturalness, however, are independent of the choice of renormalization scheme: as seen in Section 3.2.1, the sensitivity of relevant parameters to high-energy physics does not vanish when switching from a mass-dependent renormalization scheme to a mass-independent one. The upshot of this is that just as failures of naturalness invalidate the claim of Cao and Schweber that the Decoupling Theorem justifies an ontology of quasi-autonomous domains, so too do they invalidate Bain's claim that "continuum EFTs are, by themselves, capable of supporting an ontology of quasi-autonomous domains" [7] (p.241). The reason for this, of course, is that on any reasonable precisification of "quasi-autonomy," the extremely delicate interscale sensitivity of relevant parameters renders Bain's claim ineluctably false.

6 Conclusion

The aim of this paper has been twofold. First, I sought to provide a physically transparent and uniform notion of naturalness that would be both well-motivated within the framework of effective field theory and could provide a physical grounding for the seemingly discordant notions in use in the physics literature. My conclusion was that the understanding of naturalness which satisfied these desiderata was one which had been all too often left on the sidelines of the more recent discussions of naturalness, according to which the requirement of naturalness amounts to a prohibition on sensitive correlations between widely separated physical scales in effective field theories. Second, I argued that the unnaturalness of the Standard Model, increasingly pressing in light of recent LHC data, had ontological consequences. In particular, it undermines recent attempts to argue that effective field theory entails an ontology of quasi-autonomous physical domains, largely independent from one another and each with its own ontology and laws. I concluded that the extraordinarily sensitive dependence of relevant parameters at low energies – the mass of the Higgs particle in the Standard Model, for example – on their initial high-energy values at the scale of the UV cutoff vitiated this ontological picture of our world. Phrased slightly differently, I concluded that any attempt to justify such an ontology of our world by pointing to features of the effective field theory framework is doomed to fail unless the theory in question is *natural*.

Acknowledgments

I would like to thank David Albert, Sidney Felder, Jon Lawhead, Michael Miller, and David Wallace for helpful comments and discussions. I would also like to thank two anonymous referees for insightful comments and criticism.

References

[1] Michael Aizenman. Geometric analysis of ϕ^4 fields and ising models i and ii. *Communications in Mathematical Physics*, 86(1):1–48, 1982.

[2] Luis Alvarez-Gaumé and Miguel A Vázquez-Mozo. *An invitation to quantum field theory*, volume 839. Springer Science & Business Media, 2011.

[3] G.W. Anderson, D.J Castano, and A. Riotto. Naturalness lowers the upper bound on the lightest higgs boson mass in supersymmetry. *Phys. Rev. D*, 55(5):2950–2954, 1997.

[4] T. Appelquist and J. Carazzone. Infrared singularities and massive fields. *Phys. Rev. D*, 11(10):2856–2861, 1975.

[5] N. Arkani-Hamed and S. Dimopolous. Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the lhc. *Journal of High Energy Physics*, 2005(06):073, 2005.

[6] P. Athron and D.J. Miller. New measure of fine-tuning. *Phys. Rev. D*, 76(7):07510, 2007.

[7] J. Bain. Effective field theories. In *Oxford Handbook of Philosophy of Physics*. Oxford University Press, Oxford, 2013.

[8] T. Banks. *Modern Quantum Field Theory*. Cambridge University Press, Cambridge, 2008.

[9] R. Barbieri. Electroweak theory after the first lhc phase. http://arxiv.org/abs/1309.3473, 2013.

[10] R. Barbieri and G. F. Guidice. Upper bounds on supersymmetric particle masses. *Nuclear Physics B*, 306(1):63–76, 1988.

[11] R. Batterman. The tyranny of scales. In *Oxford Handbook of Philosophy of Physics*. Oxford University Press, Oxford, 2013.

[12] E. Bertuzzo. Susy after lhc8: a brief overview. http://arxiv.org/abs/1307.0318, 2013.

[13] C. Burgess. The cosmological constant problem: why it's hard to get dark energy from microphysics. http://arxiv.org/abs/1309.4133, 2013.

[14] D.J.E. Callaway. Triviality pursuit: can elementary scalar particles exist? *Phys. Rep.*, 167(5):241–320, 1988.

[15] T.Y. Cao and S. Schweber. The conceptual foundations and the philosophical aspects of renormalization theory. *Synthese*, 97(1):33–108, 1993.

[16] E. Castellani. Reductionism, emergence, and effective field theories. *Studies in History and Philosophy of Modern Physics*, 33(2):251–267, 2002.

[17] N. Craig. The state of supersymmetry after run i of the lhc. arXiv:1309.0528, 2013.

[18] S. Dawson. Introduction to electroweak symmetry breaking. In A. Masiero, G. Senjanovic, and A. Smirnov, editors, *ICTP High Energy Physics and Cosmology Summer School*, 1998.

[19] J. Donoghue. *Universe or Multiverse?*, chapter The fine-tuning problems of particle physics and anthropic mechanisms. Cambridge University Press, Cambridge, 2007.

[20] A. Duncan. *The conceptual framework of quantum field theory.* Oxford University Press, Oxford, 2012.

[21] J.L. Feng. Naturalness and the status of supersymmetry. *Annual Review of Nuclear and Particle Science*, 63(1):351–382, 2013.

[22] P. Frampton. *Gauge Field Theories.* Wiley-VCH, 2000.

[23] Simon Friederich, Robert Harlander, and Koray Karaca. Philosophical perspectives on ad hoc hypotheses and the higgs mechanism. *Synthese*, 191(16):3897–3917, 2014.

[24] M.K. Gaillard and B.W. Lee. Rare decay modes of the k mesons in gauge theories. *Phys. Rev. D*, 10(3):897–916, 1974.

[25] M.K. Gaillard and B.W. Lee. Search for charm. *Reviews of Modern Physics*, 47(2):277–310, 1975.

[26] H. Georgi. Effective field theories. *Annual Review of Nuclear and Particle Science*, 43(1):209–252, 1993.

[27] G.F. Giudice. Naturally speaking: the naturalness criterion and physics at the lhc. *Perspectives on LHC physics*, pages 155–178, 2008.

[28] G.F. Giudice. Naturalness after lhc8. http://arxiv.org/abs/1307.7879, 2013.

[29] S. Glashow, J. Iliopoulos, and L. Maiani. Weak interactions with lepton-hadron symmetry. *Phys. Rev. D*, 2(7):1285–1292, 1970.

[30] A. Grinbaum. On the eve of the lhc: conceptual questions in high-energy physics. http://arxiv.org/abs/0806.4268, 2008.

[31] Alexei Grinbaum. Which fine-tuning arguments are fine? *Foundations of Physics*, 42(5):615–631, 2012.

[32] S. Hartmann. Effective field theories, reductionism, and scientific explanation. *Studies in History and Philosophy of Modern Physics*, 32(2):267–304, 2001.

[33] N. Huggett and R. Weingard. The renormalisation group and effective field theories. *Synthese*, 102(1):171–194, 1995.

[34] M. Luty. 2004 tasi lectures on supersymmetry breaking. http://arxiv.org/abs/hep-th/0509029, 2004.

[35] I. Montvay and G. Munster. *Quantum fields on a lattice*. Cambridge University Press, Cambridge, 1997.

[36] H. Murayama. Supersymmetry phenomenology. http://arxiv.org/abs/hep-ph/0002232, 2000.

[37] M. Niedermaier and M. Reuter. The asymptotic safety scenario in quantum gravity. *Living Reviews in Relativity*, 9(5), 2006.

[38] K.A. Olive et al. Review of Particle Physics. *Chin.Phys.*, C38:090001, 2014.

[39] Burt Ovrut and Howard J. Schnitzer. New approach to effective field theories. *Phys. Rev. D*, 21:3369–3387, Jun 1980.

[40] M. E. Peskin and D. V. Schroeder. *An introduction to quantum field theory*. Westview Press, 1995.

[41] J. Polchinski. Effective field theory and the fermi surface. http://arxiv.org/abs/hep-th/9210046, 1992.

[42] B. Richter. Theory in particle physics: theological speculation versus practical knowledge. *Physics Today*, 59(10):8–9, 2006.

[43] V. Rivasseau. *From constructive to perturbative renormalization*. Princeton University Press, Princeton, 1991.

[44] M. Schwartz. *Quantum field theory and the standard model*. Cambridge University Press, Cambridge, 2014.

[45] M. Shifman. Reflections and impressionistic portrait at the conference "frontiers beyond the standard model," ftpi, oct. 2012. http://arxiv.org/abs/1211.0004, 2012.

[46] Leonard Susskind. Dynamics of spontaneous symmetry breaking in the weinberg-salam theory. *Phys. Rev. D*, 20:2619–2625, Nov 1979.

[47] Gerard t Hooft. Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking. *NATO Adv. Study Inst. Ser. B Phys.*, 59(PRINT-80-0083 (UTRECHT)):135, 1979.

[48] J. D. Wells. *Effective theories in physics: from planetary orbits to elementary particle masses*. Springer, 2012.

[49] J. D. Wells. The utility of naturalness, and how its application to quantum electrodynamics envisages the standard model and the higgs boson. http://arxiv.org/abs/1305.3434, 2013.

[50] C. Wetterich. Fine-tuning problem and the renormalization group. *Physics Letters B*, 140(3–4):215–222, 1984.

[51] K. G. Wilson. The origins of lattice gauge theory. *Nuclear Physics B – Proceedings Supplements*, 140:3–19, 2005.

[52] A. Zee. *Quantum Field Theory in a Nutshell*. Princeton University Press, Princeton, 2010.