Naturalness and New Approaches to the Hierarchy Problem

PiTP 2017

Nathaniel Craig

Department of Physics,
University of California,
Santa Barbara, CA 93106

No warranty expressed or implied. This will eventually grow into a more polished public document, so please don't disseminate beyond the PiTP community, but please do enjoy. Suggestions, clarifications, and comments are welcome.

Contents

1 Introduction

What are the natural sizes of parameters in a quantum field theory? The original notion is the result of an aggregation of different ideas, starting with Dirac's Large Numbers Hypothesis ("Any two of the very large dimensionless numbers occurring in Nature are connected by a simple mathematical relation, in which the coefficients are of the order of magnitude unity" [1]), which was not quantum in nature, to Gell-Mann's Totalitarian Principle ("Anything that is not compulsory is forbidden." [2]), to refinements by Wilson and 't Hooft in more modern language. In any event, for simplicity we will refer to this aggregate notion of naturalness as **Dirac naturalness**:

> In a theory with a fundamental scale Λ, given an operator O of the form
>
> $$\mathcal{L} \supset c_O O \tag{1}$$
>
> with scaling dimension Λ_O, the natural size of the coefficient c_O in natural units is
>
> $$c_O = \mathcal{O}(1) \times \Lambda^{4-\Delta_O} \tag{2}$$

This has the flavor of mere dimensional analysis, but it is reinforced by the nature of quantum corrections in QFT.

Of course, we have many examples of QFT which appear to violate this expectation. This leads to a refined notion of naturalness, due primarily to 't Hooft[1], which we will refer to as **technical naturalness**:

> Coefficients can be much smaller than their Dirac natural value if there is an enhanced symmetry of the theory when the coefficient is taken to zero. In this case, the natural size of the coefficient c_O is
>
> $$c_O = \mathcal{S} \times \mathcal{O}(1) \times \Lambda^{4-\Delta_O} \tag{3}$$
>
> where $\mathcal{S}$ is a parameter that violates the symmetry in question.

[1] Though certainly anticipated by Gell-Mann, who in the sentence after articulating the Totalitarian Principle notes "Use of this principle is somewhat dangerous, since it may be that while the laws proposed in this communication are correct, there are others, yet to be discussed, which forbid some of the processes that we suppose to be allowed."

The origin of this is fairly transparent: if the parameter $\mathcal{S}$ is zero, then there is an enhanced symmetry of the theory. Quantum corrections respect symmetries of the quantum action, and so radiative corrections will not regenerate c_O. If the symmetry is violated by nonzero $\mathcal{S}$, then there is a selection rule: radiative corrections must be proportional to the symmetry violation. We can formalize this at the level of spurion analyses, familiar from the chiral Lagrangianin QCD.

The two notions of naturalness are clearly on different footings. Ultimately, we expect Dirac naturalness to hold in all underlying field theories. But technical naturalness gives us the ability to understand how hierarchies observed in the infrared can be protected against radiative corrections that would otherwise spoil them.

These two notions of naturalness have been borne out countless times in nature, and provide a successful characterization of many of the parameters in the Standard Model. Two classic examples are the proton mass and flavor hierarchies.

1.0.1 The proton mass

This was the problem that originally motivated Dirac, and his own answer was wildly off the mark. Dirac understood that there was a mass scale associated with gravity, $M_{Pl} \sim 10^{18}$ GeV, as well as a mass scale associated with the proton, $m_p \sim 1$ GeV, and wished to understand why $m_p \ll M_{Pl}$.

Although the true explanation eluded Dirac, we now understand it to be a beautiful triumph of naturalness criteria. The answer is that the proton mass is dynamically generated by confinement, which in turn arises from the logarithmic evolution of a dimensionless coupling. This phenomenon, known as *dimensional transmutation*, explains the existence of exponentially different scales.

The essential idea is that the QCD coupling, like all couplings in the Standard Model, runs as a function of scale, giving rise to a renormalization group equation of the form

$$\frac{\partial \alpha_3}{\partial \ln \mu} = -7\frac{\alpha_3^2}{2\pi} + \dots \tag{4}$$

where $\alpha \equiv g^2/4\pi$. We can solve the RGE at one loop, starting from couplings defined at a fundamental scale (taken to be, e.g., M_{Pl}) down to some lower scale μ:

$$\frac{1}{\alpha_3(M_{Pl})} - \frac{1}{\alpha_3(\mu)} = \frac{7}{2\pi} \ln\left(\frac{M_{Pl}}{\mu}\right) \tag{5}$$

This tells us that, starting from a finite value of g, it will eventually diverge in the infrared. Although there is no rigorous proof, we understand this to be associated with confinement in QCD. At one loop, we can take the scale of confinement Λ_{QCD} to be the scale at which the coupling diverges, in which case

$$\frac{1}{\alpha_3(M_{Pl})} = \frac{7}{2\pi} \ln\left(\frac{M_{Pl}}{\Lambda_{QCD}}\right) \Rightarrow \Lambda_{QCD} = M_{Pl} e^{-\frac{2\pi}{7}\frac{1}{\alpha_3(M_{Pl})}} \tag{6}$$

Lo and behold, we observe a new scale that is exponentially far from the fundamental scale, where the exponential difference owes to the gentle, logarithmic evolution of a dimensionless coupling.

As the proton acquires most of its mass from confinement, $m_p \sim \Lambda_{QCD}$, we see there is a Dirac natural explanation for $m_p \ll M_{Pl}$: all parameters can take a Dirac-natural size at the scale M_{Pl}, and a new scale is generated dynamically.

1.0.2 Flavor hierarchies

A more subtle example is provided by the flavor hierarchies of the Standard Model. In the Standard Model, we see large hierarchies in fermion masses, e.g.

$$\frac{m_e}{m_t} \sim 10^{-5} \qquad \frac{m_\nu}{m_t} \sim 10^{-11} \tag{7}$$

Of course, in the Standard Model fermion masses are generated by electroweak symmetry breaking, so that these hierarchies emerge from hierarchies of Yukawa couplings.

These numerical hierarchies are not Dirac natural, but are technically natural. In the limit that the Yukawa couplings are taken to zero, there is an enhanced symmetry of the Standard Model, namely a $U(3)^5$ flavor symmetry. This corresponds to a $U(3)$ symmetry for each type of left-handed Weyl fermion, although it is often more conveniently decomposed into the following symmetries:

$$SU(3)_Q \times SU(3)_U \times SU(3)_D \times SU(3)_L \times SU(3)_E \tag{8}$$
$$\times U(1)_B \times U(1)_L \times U(1)_Y \times U(1)_{PQ} \times U(1)_E \tag{9}$$

We can then think of the yukawas as spurions for breaking the symmetry. For example, the up- and down-type Yukawa couplings break the $SU(3)^3_q \equiv SU(3)_Q \times SU(3)_U \times SU(3)_D$ global symmetry, while the lepton yukawas break the $SU(2)^2_\ell \equiv$

$SU(3)_L \times SU(3)_E$ symmetry. We can track the symmetry breaking by treating the yukawas as fields transforming in definite representations of the global flavor symmetry, whose vacuum expectation values spontaneously break the symmetry. In this sense the yukawas are "spurion fields" for the broken symmetry. Qua spurions, the various yukawas transform as

$$Y^u \sim (3, \bar{3}, 1)_{SU(3)_q^3} \tag{10}$$

$$Y^d \sim (3, 1, \bar{3})_{SU(3)_q^3} \tag{11}$$

$$Y^e \sim (3, \bar{3})_{SU(3)_\ell^2} \tag{12}$$

Consequently, radiative corrections to the yukawa couplings are proportional to these spurions. Any numerical hierachies in the spurions are therefore radiatively stable.

Of course, we would still like an explanation for the origin of the numerical hierarchies – why the yukawas might have hierarchical values to begin with – but this can be accomplished by model-building at some fundamental scale at which the yukawas are generated. Once the hierarchies are generated, they persist into the infrared.

2 The Electroweak Hierarchy Problem

Of course, there are places where our expectations of naturalness fail. The most notable examples are the cosmological constant problem, the electroweak hierarchy problem, and the strong CP problem. Each regards operators of different dimension – the parameters of interest have classical mass dimension four, two, and zero, respectively. The second of these, the electroweak hierarchy problem, will be our focus for the remainder of these lectures, and relates to the natural sizes of mass terms in a quantum field theory. There are various levels to the problem, but the essential issue is that the observed Higgs mass is some sixteen orders of magnitude smaller than the apparent cutoff of the Standard Model EFT, associated with the scale of quantum gravity. While this would not be a concern if the mass parameter were technically natural in the Standard Model, we are not so fortunate, and so we are faced with a striking violation of our notions of naturalness.

Of course, not all mass parameters need be problematic. Consider, for example, the mass of a Dirac fermion Ψ with a mass term of the form

$$m\bar{\Psi}\Psi \,. \tag{13}$$

This mass term is invariant under a vector-like $U(1)$ global symmetry under which $\Psi \to e^{i\alpha}\Psi$, but in the limit $m \to 0$ there is an additional symmetry, namely axial transformations of the form $\Psi \to e^{i\alpha\gamma_5}\Psi$. We could equivalently think of the symmetries in the massless limit as the two $U(1)$ symmetries of two free Weyl fermions.

Quantum corrections respect the symmetries of the quantum action, so provided that this axial symmetry is a good symmetry of the quantum theory (i.e., is not anomalous), when when $m = 0$ this implies that quantum corrections will not generate a mass term. Moreover, when the chiral symmetry is broken by $m \neq 0$, quantum corrections will be proportional to the symmetry-breaking term. Thus a large hierarchy between fermion masses is a curiosity, but not a deeply troubling one. If the fundamental theory of the universe generates fermions with very different masses, quantum corrections need not disturb the hierarchy.

The same does not in general hold for the mass terms for scalar fields. In particular, in the Standard Model the mass term

$$m^2 H^\dagger H \tag{14}$$

is in general a complete invariant under any gauge or global symmetry acting on H, and no symmetry is enhanced when the mass is zero. Thus we are without any argument to justify the stability of the Higgs mass parameter against radiative corrections. Indeed, we find in any theory with multiple mass scales that the Higgs accumulates radiative corrections from every scale with which it interacts, proportional to those scales. Unlike the case of spin-1/2 or spin-1, we do not have $\delta m^2 \propto m^2$, but rather $\delta m^2 \propto \Lambda^2$, where Λ stands for all other scales probed by the Higgs.

The hierarchy problem is often framed in the language of quadratic divergences. The idea is to consider the Standard Model as an effective field theory up to some cutoff Λ. One can infer the sensitivity of the Higgs mass parameter to the cutoff by computing one-loop radiative corrections up to the scale Λ, which give the famous quadratic divergence,

$$\delta m_H^2 = \frac{\Lambda^2}{16\pi^2}\left[6\lambda + \frac{9}{4}g_2^2 + \frac{3}{4}g_Y^2 - 6y_t^2 + \ldots\right] \tag{15}$$

On top of this, one should include a bare term, so that the expectation for the Higgs mass in the Standard Model EFT is

$$m_H^2 = c\Lambda^2 + \delta m_H^2 \tag{16}$$

There is a great deal of confusion about quadratic divergences and their significance, so it is worth parsing this result very carefully.

The first question is whether we need to treat the Standard Model as an EFT in the first place. In general, this is a sensible thing to do – even if it were not for the apparent cutoff imposed by strong gravity at the scale M_{Pl}, if the Standard Model were run up to arbitrarily high energies, it would hit a Landau pole in the hypercharge gauge coupling around 10^{41} GeV. More precisely, given the measured value of the hypercharge coupling at the Z pole, and the beta function

$$\frac{\partial \alpha_Y}{\partial \ln \mu} = \frac{41}{10}\frac{\alpha_Y^2}{2\pi} + \dots \tag{17}$$

the hypercharge coupling is fated to diverge around 10^{41} GeV. If this were to occur, then Standard Model fermions would form non-zero vacuum condensates in the UV, which is inconsistent with the long-range degrees of freedom in the IR. So the Standard Model is genuinely an effective field theory with cutoff Λ whether or not one is concerned about the implications of quantum gravity.

The second question is what to think of the quadratic divergence itself. We learn at an early age how to deal with divergent results in quantum field theory – we introduce counterterms and fix their coefficients according to some renormalization scheme, and then use this scheme to make finite predictions for observables at other scales. So at first glance, one might not be too troubled by the quadratic divergence. But even if one doesn't ascribe physical significance to the quadratic divergence alone, it signals the existence of sensitivity to UV physics.

From the Wilsonian perspective, the quadratic divergence is really all there is. The underlying idea is that the fundamental theory is finite, and divergences in the EFT are physical (e.g. cutoff = lattice spacing, or mass scale of particles rendering the Higgs mass finite), and counterterms just manifest fine-tuning.

A less ambitious reading, but one that is much clearer to interpret than musings about cutoffs, is that the quadratic divergence is just a placeholder for physical thresholds. The detailed relationship between the cutoff and the mass of new physical particles is a bit subtle, but as an order of magnitude relationship, it typically holds true. And, indeed, when we know what those thresholds are, we can go ahead and compute explicitly to see what's going on. To see this, it helps to construct a toy model.

2.1 A toy model

Concretely, consider as a toy model a real scalar coupled to a Dirac fermion,

$$\mathcal{L} = \frac{1}{2}(\partial\phi)^2 - \frac{1}{2}m^2\phi^2 - \frac{\lambda}{4!}\phi^4 + \bar{\Psi}i\,\partial\!\!\!/\Psi - M\bar{\Psi}\Psi + y\phi\bar{\Psi}\Psi \tag{18}$$

The yukawa coupling of this particular toy model breaks the continuous chiral symmetry we discussed earlier, but retains a discrete chiral symmetry under which

$$\Psi \to \gamma_5\Psi \qquad \phi \to -\phi \tag{19}$$

Under this symmetry $\bar{\Psi}\Psi \to -\bar{\Psi}\Psi$, so the fermion mass M is rendered technically natural. But there is no additional symmetry that is manifest when $m \to 0$, so we expect to see a hierarchy problem.

We would like to imagine that we keep the scalar much lighter than the fermion, and to consider matching between the full theory and an effective theory in which the fermion has been integrated out. To avoid any confusion about quadratic divergences, we will work in terms of a mass-independent renormalization scheme, dimensional regularization with minimal subtraction ($\overline{MS}$). In this scheme, the mass parameters of the theory can be thought of as Lagrangian parameters that evolve as a function of scale. We deform the theory by non-integer dimension (e.g. $d = 4-\varepsilon$) to tame divergences, and the divergences are parameterized by $1/\varepsilon$ poles. The renormalization prescription is to choose our counterterms to cancel those poles plus some superfluous factors of 4π and γ.

We would like to carry out a matching procedure between the full theory and the effective field theory, matched at the scale M. To do so, we match the scalar two-point function in the EFT to the scalar two-point function in the full theory, at whatever order we care to compute. At one loop, the matching involves tree-level diagrams plus a one-loop diagram

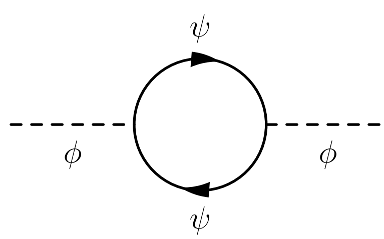

which evaluates to a contribution to the scalar self-energy of the form

$$\Sigma_2(p^2) = \frac{4y^2}{16\pi^2}\left[\left(\frac{3}{\bar{\epsilon}} + 1 + 3\log(\mu^2/M^2)\right)\left(M^2 - \frac{p^2}{6}\right) + \frac{p^2}{2} - \frac{p^2}{20M^2} + \dots\right] \quad (20)$$

where $\frac{1}{\bar{\epsilon}} = \frac{1}{\epsilon} - \gamma + \log(4\pi)$. Note that there are no logarithms involving m^2 or p^2, as these diagrams match on to an EFT that contains only a free scalar field at tree level, so there are no loop diagrams that could reproduce the logarithm.

Now we renormalize by adding counterterms to cancel the $1/\bar{\epsilon}$ pole and match at the scale $\mu = M$. The matched Lagrangian in the scalar theory is thus

$$\mathcal{L} = \left(1 - \frac{4}{3}\frac{y^2}{16\pi^2}\right)\cdot\frac{1}{2}(\partial\phi)^2 - \left(m^2 - \frac{4y^2}{16\pi^2}M^2\right)\cdot\frac{1}{2}\phi^2 + \dots \quad (21)$$

where $\dots$ includes higher-derivative terms and interactions.

It's clear that the mass in the effective field theory contains a threshold correction relative to the UV theory proportional to $\frac{4y^2}{16\pi^2}M^2$. We could have also calculated the above loop diagram with a hard momentum cutoff, and found a quadratically divergent contribution to the mass-squared

$$\delta m^2 \supset \frac{3\lambda^2}{4\pi^2}\Lambda^2 \quad (22)$$

In this sense, the quadratic divergence is just a stand-in for the finite threshold corrections. If we were infinitely powerful, we could compute everything explicitly and see the finite effects. But if we are not, and are only working from the bottom up, the quadratic divegences are a handy way to estimate the effects of new physics.

We can also see technical naturalness at play by reversing the setup, and considering a theory in which the fermion is light while the scalar is heavy. In this version, the threshold correction to the fermion mass is proportional to the fermion mass, rather than the scalar mass, a manifestation of the technical naturalness of the discrete chiral symmetry. Note that we could repeat the calculation with a hard cutoff, though here the result is subtle – we would naively see a quadratic divergence, but only because our regularization broke the chiral symmetry. Restoring the chiral symmetry with an appropriate choice of counterterms (as we should), the quadratic divergence for the fermion is absent – again consistent with technical naturalness of the fermion mass.

In any event, now we can extract the appropriate lesson from the naive quadratic divergence in the Standard Model. If physics enters to render the Higgs mass finite and calculable, then it will of course give contributions of this form. Indeed, this occurs for every theory in which the Higgs mass is rendered calculable, where the finite contributions are precisely from whatever new degrees of freedom render the Higgs mass finite. We will see such contributions in explicit examples.

But even if the physics in the far UV is mysterious and behaves differently from our expectations, it's also clear that there are finite contributions from other degrees of freedom entirely unrelated to the finiteness of the Higgs mass. For example,

Unification One of the first concrete settings in which the hierarchy problem became apparent was that of grand unification. In grand unified theories there are heavy gauge bosons associated with the scale of unification that interact with the Higgs boson.

Details depend on the precise model of unification, and the representation into which the Higgs is embedded. For example, in $SU(5)$ unification the SM gauge bosons are embedded into the 24 of $SU(5)$, which decomposes into the SM gauge bosons plus X gauge bosons transforming in the $(3,2)_{-5/6}$ + conjugate representation. Moreover, the Higgs is embedded in a $\bar{5}$ of $SU(5)$. In this case there are loops involving a triplet scalar Higgs and X boson of the form

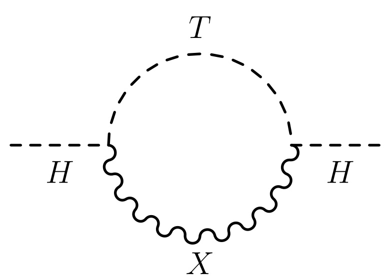

In general, these loops of heavy bosons give corrections of order

$$\delta m_H^2 \sim \frac{\alpha_{GUT}}{4\pi} M_{GUT}^2 \tag{23}$$

The original apparent scale of unification in nonsupersymmetric theories was $\mathcal{O}(10^{15})$ GeV, while bounds on proton decay now imply $M_{GUT} \gtrsim 10^{16}$ GeV. So grand unification implies a huge hierarchy problem.

Neutrino masses Now we can have a perfectly consistent universe without new electroweak fermions, but there are scenarios that favor the existence of new fermions. For example, the generation of neutrino masses may strictly be due to a dimension-five operator,

$$\mathcal{L} \supset \frac{(L^i H)(L^j H))}{M} + \text{h.c.} \tag{24}$$

without further ado. However, we expect that if the theory is genuinely renormalizable, this interaction arose from integrating out heavier states with mass $\sim M$. In particular, the Type-I seesaw entails right-handed neutrinos N with couplings

$$\mathcal{L} \supset -\frac{M_R^{ij}}{2} N_i N_j - y_{ij} L^i N^j H + \text{h.c.} \tag{25}$$

This provides a very concrete example of new fermions coupling to the Higgs. The leading one-loop correction to the Higgs mass is

$$\delta m_H^2 = -\frac{1}{4\pi^2} \sum_{ij} |y_{ij}|^2 M_j^2 \tag{26}$$

If all the RH neutrinos have a common mass M, the bound will be dominated by the combination of yukawas giving the heaviest SM neutrinos. In this case the naturalness bound is $M \lesssim 10^4$ TeV. This has amusing implications because thermal leptogenesis requires much higher values of M, on the order of $M \gtrsim 10^6$ TeV. So in this case naturalness would rule out thermal leptogenesis in a Type 1 see-saw.

Gravity Even giving up on these things, some UV completion is forced upon us. We have already encountered the physics of quantum gravity at a scale $M_P \sim 10^{19}$ GeV. Do not have a complete theory of quantum gravity, although it is likely that the answer lies in string theory. We are not yet able to compute the mass of the Higgs in a complete string theory, but the expectation is that string theory contains heavy states whose masses are close to the Planck scale that would give corrections to the Higgs mass.

It's clear that this is a problem, but we can make it even more apparent. Even new states coupling to the Higgs through loops of perturbative gravitons give a large threshold correction. For example, imagine there is some massive Dirac fermion Ψ with mass m_Ψ and it coupled to the Standard Model only gravitationally. Then as long as we are at energies $E \ll M_{Pl}$ we can compute loop diagrams including gravitons. The correction to the Higgs mass in this case arises at two loops,

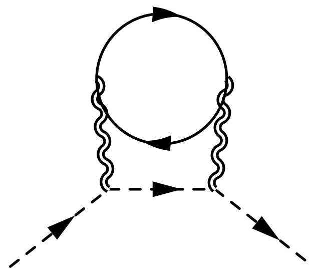

and gives a correction parametrically of order

$$\delta m_H^2 \sim \frac{m_H^2}{(16\pi^2)^2}\frac{m_\Psi^4}{M_{Pl}^4}$$

This correction is small because the graviton coupling to a massless, on-shell particle at zero momentum vanishes, and so the result is proportional to m_H.

However, we could also have a three-loop diagram where the graviton couples to a loop of top quarks,

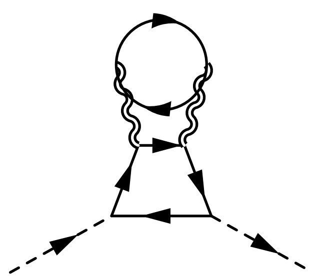

The correction from this diagram is parametrically of the form

$$\delta m_H^2 \sim \frac{6y_t^2}{(16\pi^2)^3}\frac{m_\Psi^6}{M_{Pl}^4}$$

and is much larger because now the gravitons are coupling to off-shell states.

If $m_\Psi \sim M_{Pl}$, correction is $\sim \frac{6y_t^2}{16\pi^2}\frac{M_{Pl}^2}{(16\pi^2)^2}$. Of course at this point we doubt the validity of our gravity EFT, but this parametrically validates our naive expectation from the cutoff argument, now with $\Lambda \sim M_{Pl}/16\pi^2$. So even gravitational physics is sufficient to feed through threshold corrections to the Higgs mass.

The conclusion is that if there are *any* other states out there, even ones that only couple to the Higgs gravitationally, they give a threshold correction to the Higgs mass that is proportional to the mass scale of the new states. We can see these corrections in $\overline{MS}$ or any other scheme; they are physical threshold corrections and have unambiguous value. The result using a hard cutoff was merely a placeholder for threshold corrections, which we could only see in $\overline{MS}$ if we had actual physical states in the theory.

2.2 The naturalness strategy

Now we can convert the UV sensitivity of the Higgs mass into a strategy for new physics. We imagine that the Higgs mass is natural because the theory changes not far from the weak scale. Even without specifying the details, we can first estimate where the change must occur. Considering the "quadratic divergence",

$$\delta m_H^2 = \frac{\Lambda^2}{16\pi^2}\left(-6y_t^2 + \frac{9}{4}g^2 + \frac{3}{4}g'^2 + 6\lambda\right) \tag{27}$$

we imagine that the scale Λ is such that contributions from the cutoff are of the natural size of the Higgs mass itself. If so, this implies $\Lambda \lesssim 500$ GeV. Higher cutoffs imply contributions larger than the observed mass, and a correspondingly increased tuning – for example, tuning at the percent level would correspond to a cutoff of 5 TeV.

This is a strategy for new physics, not a necessity. Nothing fails in the field theory if the expectation is violated; it is simply difficult to understand from the perspective of naturalness. Of course, one might wonder whether this strategy is justified – after all, nature does not care much about our level of puzzlement.

It turns out that there are many instances of naturalness in nature, including naturalness at the level of mass parameters. One of my favorite is the mass splitting between the charged and neutral pions, which differ by about 5 MeV. These states are all goldstones of the spontaneously broken chiral symmetries of QCD, and these symmetry arguments lead one to expect the pions to be nearly degenerate. The answer is that we have radiative corrections from the explicit breaking of chiral symmetries by QED. The charge matrix for three flavors is

$$\begin{pmatrix} +2/3 & 0 & 0 \\ 0 & -1/3 & 0 \\ 0 & 0 & -1/3 \end{pmatrix}$$

This matrix breaks the chiral symmetry associated with the generators of the charged pions and kaons (i.e., it only commutes with the generators associated with the neutral pions – i.e., $[Q, T^a\pi^a] \sim f(\pi^\pm, K^\pm)$). So the charged pions and kaons can get a mass contribution from electromagnetic loops. If we compute the photon loop that would give a mass correction, using a hard cutoff to estimate the threshold correction we get

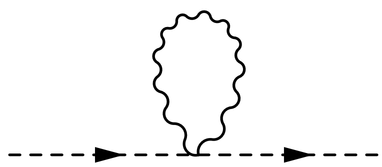

$$\sim \frac{3e^2}{16\pi^2}\Lambda^2 = \delta m^2 \qquad (28)$$

Given the size of the charged-neutral meson splittings, $m^2_{\pi^\pm} - m^2_{\pi_0} \sim (35.5\,\text{MeV})^2$, we expect the loop should be cut off around 850 MeV if electromagnetic loops explain the mass difference. In fact, the ρ meson enters at 770 MeV, which provides a cutoff for the effective theory. Here the ρ meson is a proxy for compositeness, as it is the first QCD bound state outside of the chiral lagrangian. Thus there is perfect agreement between the size of the mass correction based on cutoff-based arguments and the scale at which new physics enters.

Another beautiful example is the mass difference between the K^0_L and K^0_S states. Computed in the effective theory at the scale of the kaons, the splitting is

$$\frac{M_{K^0_L} - M_{K^0_S}}{M_{K^0_L}} = \frac{G_F^2 f_K^2}{6\pi^2}\sin^2\theta_c \Lambda^2 \qquad (29)$$

where $f_K = 114$ MeV is the kaon decay constant and $\sin\theta_C = 0.22$ is the Cabibbo angle. Requiring this correction to be smaller than the measured value $(M_{K^0_L} - M_{K^0_S})/M_{K^0_L} = 7 \times 10^{-15}$ gives $\Lambda < 2$ GeV. And lo, the charm quark enters with mass

$m_c \sim 1.2$ GeV to modify the short-distance behavior of the theory by implementing the GIM mechanism. Moreover, this is not merely rationalization; this was the actual argument used by Gaillard and Lee to compute the mass of the charm quark before its discovery.

On the other hand, the cosmological constant is a tremendous failure of naturalness. The observed value of the c.c. is on the order of $(2.4 \times 10^{-3}\ \text{eV})^4$. The prediction obtained by simply computing vacuum loops up to a cutoff Λ is proportional to Λ^4 itself. There is no apparent new physics at the eV scale related to cutting off contributions to the vacuum, and even if these loops were cutoff not far above the weak scale, there would be many orders of magnitude of unexplained hierarchy.

At this point, it's helpful to summarize our understanding of the hierarchy problem diagrammatically:

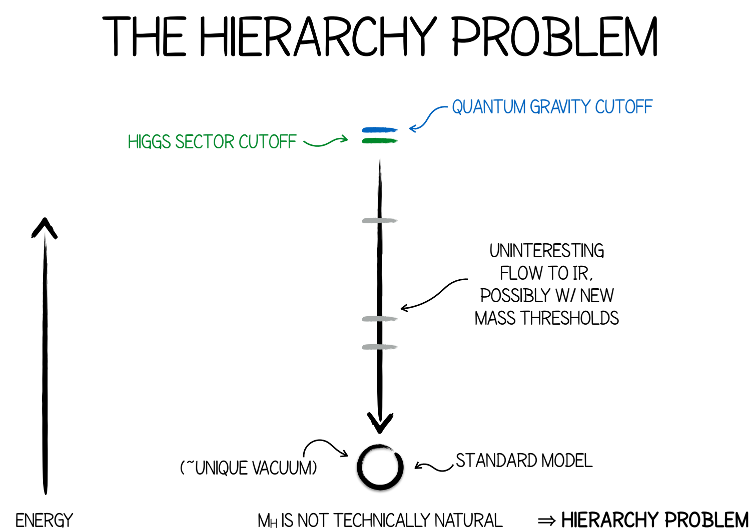

Figure 1: Cartoon of the hierarchy problem

Having this cartoon in hand will help us to evaluate possible solutions.

3 Old Hierarchy Solutions

Inspecting the cartoon of the hierarchy problem, there are more or less three obvious things to try, which we can illustrate with their own cartoons:

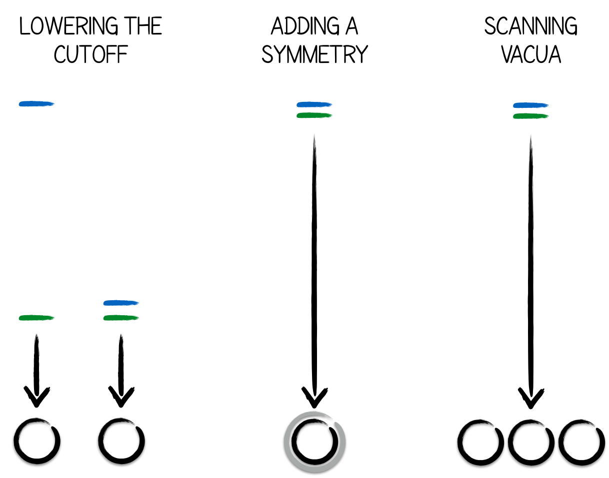

Figure 2: Ways to solve the hierarchy problem

3.1 Lowered cutoff

The first thing one is tempted to do when confronted by the hierarchy problem is to erase the apparent hierarchy itself, bringing down the cutoff of the Higgs sector or the entire Standard Model. Indeed, this was the nature of the first attempted solution to the hierarchy problem, *technicolor*, which attempted to replicate the success of the proton mass prediction by imagining that electroweak symmetry was broken by the vacuum condensate of a strongly coupled group. The five-dimensional holographic duals of technicolor are Randall-Sundrum models, specifically ones on a finite interval with branes at either end. In these cases, the Higgs is not an elementary degree of freedom, and the cutoff is provided by compositeness of the Higgs itself.

Alternately, we could imagine leaving the Higgs alone and lowering the scale of quantum gravity, so that all field theoretic physics reaches an end at the cutoff. This is the nature of solutions such as large extra dimensions.

The problem with pure lowered-cutoff solutions is that they generically do not allow a small bare mass term for the Higgs. That is to say, the natural expectation of the Higgs mass is of order

$$m_H^2 = c\Lambda^2 + \delta m_H^2 \tag{30}$$

As such, the cutoff of the theory must be close to the Higgs mass, rather than parametrically separated. Such theories then predict a host of particles near in mass to the Higgs, as well as a host of higher-dimensional operators suppressed by a low cutoff. The nonobservation of new particles close in mass to the Higgs, as well as strong bounds on dimension-6 operators, suggests that this mechanism is not operative on its own. This brings us to...

3.2 Symmetries

The idea behind symmetry solutions is to enlarge the Standard Model so that the Higgs mass becomes a technically natural parameter. What possible symmetries can we use? Coleman-Mandula theorem constrains options in four dimensions:

The Coleman-Mandula theorem (1967): *in a theory with non-trivial interactions (scattering) in more than 1+1 dimensions, the only possible conserved quantities that transform as tensors under the Lorentz group are the energy-momentum vector P_μ, the generators of Lorentz transformations $M_{\mu\nu}$, and possible scalar symmetry charges Z_i corresponding to internal symmetries, which commute with both P_μ and $M_{\mu\nu}$.* For theories with only massless particles, this can be extended to include generators of conformal transformations.

The Coleman-Mandula theorem can be generalized to include spinorial symmetry charges, giving rise to supersymmetry. First identified by Golfand and Likhtman, the full set of possible generalizations were identified by Haag, Sohnius, and Lopuszanski.

So possible options seem to be: Spinorial internal symmetry (supersymmetry); scalar internal symmetry (global symmetry); and potentially conformal symmetry.

3.2.1 Supersymmetry

Here I will assume you have some familiarity with SUSY, and focus on the essential aspects for the hierarchy problem. The idea is to extend Poincare algebra to include conserved supercharges minimally four supercharges $Q_\alpha, \tilde{Q}_{\dot{\alpha}}$ in four dimensions. As a Weyl spinor, the transformation properties of Q_α with respect to the Poincare group are known, namely

$$[P_\mu, Q_\alpha] = [P_\mu, \tilde{Q}^{\dot{\alpha}}] = 0$$
$$[M^{\mu\nu}, Q_\alpha] = i(\sigma^{\mu\nu})_\alpha^\beta Q_\beta$$
$$[M^{\mu\nu}, Q^{\dot{\alpha}}] = i(\bar{\sigma}^{\mu\nu})^{\dot{\alpha}}_{\dot{\beta}} \tilde{Q}^{\dot{\beta}}$$

We also need anticommutators $\{Q, \tilde{Q}\}$ and $\{Q, Q\}$ to close the algebra. The only option is for $\{Q, \tilde{Q}\}$ to be proportional to $P_{\alpha\dot{\beta}}$, since this is the only conserved operator with the appropriate index structure. The choice of normalization gives us

$$\{Q_\alpha, \tilde{Q}_{\dot{\beta}}\} = 2P_\mu(\sigma^\mu)_{\alpha\dot{\beta}}$$

Finally, the consistent choice for $\{Q_\alpha, Q_\beta\}$ is

$$\{Q_\alpha, Q_\beta\} = \{\tilde{Q}_{\dot{\alpha}}, \tilde{Q}_{\dot{\beta}}\} = 0$$

though of course nonzero values of the anti-commutator are possible given a larger number of supercharges.

Fields will be arranged into supermultiplets, transforming as irreducible representations of super-Poincare. For example, the chiral multiplet contains scalar and fermion related by infinitesimal SUSY rotation,

$$\phi \to \phi + \delta\phi \qquad \psi \to \psi + \delta\psi$$

where

$$\delta\phi = \epsilon^\alpha \psi_\alpha \tag{31}$$
$$\delta\psi_\alpha = -i(\sigma^\nu \epsilon^\dagger)_\alpha \partial_\nu \phi \tag{32}$$

where ϵ_α is a Grassmann variable that you can think of as an infinitestimal parameter multiplying a SUSY generator; it has mass dimension $[\epsilon] = -1/2$.

We see clearly that supersymmetry relates a scalar to a fermion, and so relates a scalar mass to a fermion mass protected by chiral symmetry. This already suggests the sense in which supersymmetry will solve the hierarchy problem: by making the mass of a scalar (the Higgs) proportional to that of a fermion (the appropriately defined superpartner thereof).

The salient properties of supermultiplets are straightforward to work out:

1. Computing the expectation value of the Witten index within a supermultiplet, we have $\mathrm{tr}\left[(-1)^{N_f}\right] = 0 \to n_F = n_B$, i.e., supermultiplets contain the same number of bosonic and fermionic degrees of freedom.

2. From $[P^2, Q_\alpha] = [P^2, \tilde{Q}_{\dot\alpha}] = 0$ we see that the components of a supermultiplet all have the same mass.

3. There is at most one $U(1)$ global symmetry that does not commute with the supercharges,
$$[R, Q_\alpha] = -Q_\alpha \qquad [R, \tilde{Q}_{\dot\alpha}] = \tilde{Q}_{\dot\alpha} \tag{33}$$
which implies that components of a supermultiplet have the same gauge and global quantum numbers apart from their $U(1)_R$ charges.

The supersymmetric extension of the Standard Model is fairly straightforward, entailing the incorporation of all Standard Model fields into corresponding supermultiplets, with the addition of a second Higgs multiplet. This is necessary on account of both anomalies and holomorphy.

Names		spin 0	spin 1/2	$SU(3)_C,\ SU(2)_L,\ U(1)_Y$
squarks, quarks	Q	$(\widetilde{u}_L\ \ \widetilde{d}_L)$	$(u_L\ \ d_L)$	$(\mathbf{3},\ \mathbf{2}\ ,\ \frac{1}{6})$
(×3 families)	$\overline{u}$	$\widetilde{u}^*_R$	$u^\dagger_R$	$(\ \overline{\mathbf{3}},\ \mathbf{1},\ -\frac{2}{3})$
	$\overline{d}$	$\widetilde{d}^*_R$	$d^\dagger_R$	$(\ \overline{\mathbf{3}},\ \mathbf{1},\ \frac{1}{3})$
sleptons, leptons	L	$(\widetilde{\nu}\ \widetilde{e}_L)$	$(\nu\ \ e_L)$	$(\mathbf{1},\ \mathbf{2}\ ,\ -\frac{1}{2})$
(×3 families)	$\overline{e}$	$\widetilde{e}^*_R$	$e^\dagger_R$	$(\mathbf{1},\ \mathbf{1},\ 1)$
Higgs, higgsinos	H_u	$(H_u^+\ \ H_u^0)$	$(\widetilde{H}_u^+\ \ \widetilde{H}_u^0)$	$(\mathbf{1},\ \mathbf{2}\ ,\ +\frac{1}{2})$
	H_d	$(H_d^0\ \ H_d^-)$	$(\widetilde{H}_d^0\ \ \widetilde{H}_d^-)$	$(\mathbf{1},\ \mathbf{2}\ ,\ -\frac{1}{2})$

Names	spin 1/2	spin 1	$SU(3)_C,\ SU(2)_L,\ U(1)_Y$
gluino, gluon	$\widetilde{g}$	g	$(\mathbf{8},\ \mathbf{1}\ ,\ 0)$
winos, W bosons	$\widetilde{W}^\pm\ \ \widetilde{W}^0$	$W^\pm\ \ W^0$	$(\mathbf{1},\ \mathbf{3}\ ,\ 0)$
bino, B boson	$\widetilde{B}^0$	B^0	$(\mathbf{1},\ \mathbf{1}\ ,\ 0)$

Of course, supersymmetry cannot be an exact symmetry of nature, otherwise we would have seen selectrons degenerate with electrons. So in general we must include soft terms, which can be worked out using the appropriate generalization of spurion techniques to superfields; in the case of the MSSM these take the form

$$\begin{aligned}
\mathcal{L}_{\rm soft}^{\rm MSSM} \;=\;& -\frac{1}{2}\left(M_3\widetilde{g}\widetilde{g} + M_2\widetilde{W}\widetilde{W} + M_1\widetilde{B}\widetilde{B} + \text{h.c.}\right) \\
& - \left(\widetilde{\overline{u}}\,\mathbf{a_u}\,\widetilde{Q}H_u - \widetilde{\overline{d}}\,\mathbf{a_d}\,\widetilde{Q}H_d - \widetilde{\overline{e}}\,\mathbf{a_e}\,\widetilde{L}H_d + \text{c.c.}\right) \\
& - \widetilde{Q}^\dagger\,\mathbf{m^2_Q}\,\widetilde{Q} - \widetilde{L}^\dagger\,\mathbf{m^2_L}\,\widetilde{L} - \widetilde{\overline{u}}\,\mathbf{m^2_{\overline{u}}}\,\widetilde{\overline{u}}^\dagger - \widetilde{\overline{d}}\,\mathbf{m^2_{\overline{d}}}\,\widetilde{\overline{d}}^\dagger - \widetilde{\overline{e}}\,\mathbf{m^2_{\overline{e}}}\,\widetilde{\overline{e}}^\dagger \\
& - m^2_{H_u}H_u^*H_u - m^2_{H_d}H_d^*H_d - \left(bH_uH_d + \text{c.c.}\right). \qquad (34)
\end{aligned}$$

It is straightforward to check that when supersymmetry is broken by these dimensionful soft terms, corrections due to breaking are proportional to these terms.

Now the Higgs mass is calculable by the introduction of supersymmetry. There are lots of ways to see it, but perhaps the simplest is constructive: we know particle content of MSSM, and can again work from an effective field theory perspective where we allow unknown new physics at a cutoff scale Λ. We assume new physics at the cutoff respects the symmetry, and then can compute loops up to cutoff as way of parameterizing our ignorance.

Relative to SM, there are now cancellations between loops of opposite statistics, e.g. top-stop loop

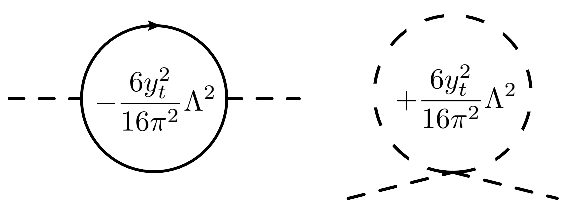

Figure 3: Quadratic divergence cancellation in the top sector of the MSSM.

Carrying out the calculation, we find

$$\delta m^2_{H_u} = -\frac{6y_t^2}{16\pi^2}\Lambda^2 + \frac{6y_t^2}{16\pi^2}\Lambda^2 - \frac{3y_t^2}{4\pi^2}m_{\tilde{t}}^2\ln\left(\Lambda/m_{\tilde{t}}\right) + \dots \qquad (35)$$

The quadratic pieces cancel. There is no longer UV sensitivity! The key assumption is that Λ is same for both loops, true for UV physics respecting supersymmetry. Obviously if supersymmetry were broken by a large amount in another sector, this would spoil the cancellation. In addition to the elimination of UV sensitivity, we are left only with physical threshold corrections (which we can compute in any scheme) from hew heavy states. At most there is logarithmic sensitivity to the cutoff Λ, and even this can be fixed by writing down explicit theory to break SUSY.

Now that mass is finite, can use naturalness argument to determine where the new particles should enter. Now we see the hierarchy problem very explicitly. We have rendered the Higgs mass calculable; now depends on masses of new partner particles, which cannot be too large without increasing fine-tuning.

There are two direct sources of concern, corresponding to tree-level contributions and loop-level contributions. Both play a role primarily through the relation between the weak scale and soft parameters, viz.

$$m_h^2 = -2(m_{H_u}^2 + |\mu|^2) + \ldots \tag{36}$$

Then corrections to Higgs mass come from three places:

- The first is the tree-level potential, which involves certain combinations of soft masses that set the weak scale vev. At tree-level the naturalness of the weak scale implies something about the soft parameters $m_{H_u}^2$ and μ, which itself controls the higgsino masses. Higgsinos should be light! Naturalness suggests $\mu \lesssim 200$ GeV and correspondingly light Higgsinos.
- The second is immediate loop-level corrections. The soft mass parameter $m_{H_u}^2$ accumulates one-loop corrections from other soft parameters. By far the largest is due to the stops, since the top chiral superfields couple most strongly to H_u, with correction of order

$$\delta m_{H_u}^2 = -\frac{3y_t^2}{4\pi^2} m_{\tilde{t}}^2 \ln\left(\Lambda/m_{\tilde{t}}\right) \tag{37}$$

  Naturalness requires stops ~ 400 GeV with a cutoff $\Lambda \sim 10$ TeV. Other particles are also tied to naturalness, though less directly. After the SM top loop, the gauge and Higgs loops drive the mass corrections, so unsurprisingly the wino and higgsino corrections play a role, with

$$\delta m_{H_u}^2 = -\frac{3g^2}{8\pi^2}(m_{\tilde{W}}^2 + m_{\tilde{h}}^2) \ln\left(\Lambda/m_{\tilde{W}}\right) \tag{38}$$

Having already bounded Higgsinos, for winos this translates to $m_{\tilde{W}} \lesssim$ TeV. Note that sbottoms need not be directly connected to naturalness, but since the left-handed sbottom gauge eigenstate transforms in the same $SU(2)$ multiplet as the left-handed stop gauge eigenstate, at least one sbottom is typically found in the same mass range as the stops.

- The third is two-loop corrections, due to the naturalness of other sparticles. The stop mass is corrected by the gluino mass due to the size of g_3, so it is hard to separate the gluino substantially from the stops, with

$$\delta m_{\tilde{t}}^2 = \frac{2g_s^2}{3\pi^2} m_{\tilde{g}}^2 \ln\left(\Lambda/m_{\tilde{g}}\right) \tag{39}$$

  which ties $m_{\tilde{g}} \lesssim 2m_{\tilde{t}}$. Indeed, these corrections typically tie the masses of the gluino and all squark flavors quite tightly given even a modest amount of running.

The problem, of course, is that we have yet to observe any evidence for supersymmetry, with bounds exceeding the TeV scale. The implication is that tuning is approaching the percent level in supersymmetric scenarios.

3.2.2 Global symmetry

Let's now turn to the alternate symmetry possibility. The general idea is that the Higgs will be a (pseudo-) goldstone boson of a spontaneously broken global symmetry, which will render the Higgs mass technically natural.

While there are many possible global symmetry structures that lead to the Standard Model at low energies, for simplicity we will focus on some simple examples. To begin with, consider an $SU(N)$ global symmetry, spontaneously broken by the vacuum expectation value of a fundamental scalar ϕ. This breaks $SU(N) \to SU(N-1)$, and the usual counting of goldstones gives us

$$[N^2 - 1] - [(N-1)^2 - 1] = 2N - 1 \tag{40}$$

goldstones. It's convenient to organize these $2N-1$ real degrees of freedom into $N-1$ complex scalars $\pi_1, \ldots, \pi_{N-1}$ plus one real scalar π_0. To study the low energy theory of the goldstones well below the scale of spontaneous symmetry breaking, we can expand ϕ in terms of the goldstones via

$$\phi = \exp\left[\frac{i}{f}\left(\begin{array}{ccc|c} & & & \pi_1 \\ & & & \vdots \\ & & & \pi_{N-1} \\ \hline \pi_1^* & \cdots & \pi_{N-1}^* & \pi_0/\sqrt{2} \end{array}\right)\right]\left(\begin{array}{c} 0 \\ \vdots \\ 0 \\ f \end{array}\right) \equiv e^{i\pi/f}\phi_0 \tag{41}$$

In this basis, the $SU(N)$ generators corresponding to the unbroken $SU(N-1)$ take the form

$$U_{N-1} = \left(\begin{array}{cc} \hat{U}_{N-1} & 0 \\ 0 & 1 \end{array}\right) \tag{42}$$

and it is relatively straightforward to check explicitly the transformation properties of the goldstones under the unbroken $SU(N-1)$. In particular, a transformation involving the unbroken generators acting on ϕ amounts to

$$\phi \to U_{N-1}\phi = (U_{N-1}e^{i\pi/f}U_{N-1}^\dagger)U_{N-1}\phi_0 = e^{\frac{i}{f}(U_{N-1}\pi U_{N-1}^\dagger)}\phi_0 \tag{43}$$

from which we see the goldstones transform as

$$\left(\begin{array}{c|c} & \vec{\pi} \\ \hline \vec{\pi}^\dagger & \pi_0/\sqrt{2} \end{array}\right) \to U_{N-1}\left(\begin{array}{c|c} & \vec{\pi} \\ \hline \vec{\pi}^\dagger & \pi_0/\sqrt{2} \end{array}\right)U_{N-1}^\dagger = \left(\begin{array}{c|c} & \hat{U}_{N-1}\vec{\pi} \\ \hline \vec{\pi}^\dagger\hat{U}_{N-1}^\dagger & \pi_0/\sqrt{2} \end{array}\right) \tag{44}$$

This is just the familiar result that the goldstones transform in the fundamental representation of the unbroken symmetry.

The transformation of the goldstones under the broken generators is far more complicated, but for infinitesimal transformations has the leading form

$$\vec{\pi} \to \vec{\pi} - \vec{\alpha} + \ldots \tag{45}$$

which manifests the shift symmetry of the goldstones under the broken generators. This is just the familiar shift symmetry of goldstone bosons, which forbids non-derivatively-coupled interactions.

So much for generalities. If we would like to see how a spontaneously broken global symmetry might protect the Higgs, it suffices to consider the simple toy model of

$SU(3) \to SU(2)$. In this case the goldstones arrange themselves naturally into the form

$$\pi = \left(\begin{array}{cc|c} -\eta/2 & 0 & H_1 \\ 0 & -\eta/2 & H_2 \\ \hline H_1^* & H_2^* & \eta \end{array}\right) \tag{46}$$

where the H_i arrange themselves naturally to form a doublet of the unbroken $SU(2)$.

If H is indeed a goldstone, it naturally inherits a series of irrelevant interactions as an expansion in f, namely

$$f^2|\partial_\mu \phi|^2 = |\partial_\mu H|^2 + \frac{H^\dagger H |\partial_\mu H|^2}{f^2} + \dots \tag{47}$$

Among other things, this immediately implies that this goldstone Higgs will have coupling deviations relative to the Standard Model, encoded by the higher-dimensional operators and unavoidable if the Higgs is a goldstone. These terms also point to the cutoff of our goldstone EFT. The loop expansion parameter in this theory is of the form $\frac{1}{f^2}\frac{\Lambda^2}{16\pi^2}$, so that a well-defined loop expansion implies a cutoff $\Lambda \lesssim 4\pi f$. Physics at Λ could be strongly coupled, as in Composite Higgs models, or weakly coupled, as in SUSY UV completions of a linear sigma model.

Of course, our discussion thus far has been a little trivial, since it neglects interactions. Consider the consequences of turning on the top yukawa coupling, which takes the form

$$\mathcal{L} \supset -\lambda_t t_R^\dagger \tilde{H} Q_3 \tag{48}$$

where $\tilde{H} = (i\sigma_2 H)^\dagger$ and $Q_3 = (t_L, b_L)$. This represents an explicit breaking of the global symmetry, and correspondingly gives rise to the usual quadratic divergence in the mass of the H. This is not surprising: the yukawas and gauge couplings of the Standard Model all violate the $SU(3)$ symmetry, and so the global symmetry offers no protection to UV physics entering through SM couplings.

Of course, this does not mean that all is lost. The global symmetry does explain why the Higgs mass is not of order $m_H^2 \sim \Lambda^2$. Assuming that Standard Model couplings are the only things that explicitly violate the global symmetry, then our notion of technical naturalness dictates that contributions to the Higgs mass coming through other Standard Model fields arise at loop level, rather than tree level. And, to a certain extent, some radiative contributions are unavoidable. After all, if the goldstone Higgs is to break electroweak symmetry, it must accumulate some quartic and quadratic terms.

However, given the lightness of the Higgs mass this is somewhat unsatisfying, as we have seen no evidence for the existence of new states beneath a TeV consistent with the requisite low cutoff. As such, it's compelling to extend the Standard Model to make the $SU(3)$ a good global symmetry, at least of the largest couplings. At the level of the top yukawa, this can be accomplished by extending the $SU(2)$ doublet quark Q_3 into a triplet of a global $SU(3)$ via $Q_3 \to \hat{Q}_3 = (\sigma_2 Q_3, T_L)$ (see, e.g. [3]). In order to marry up all degrees of freedom appropriately, we can also extend the $SU(2)$-singlet quark t_R via $t_R \to \hat{t}_R + \hat{T}_R$. Now the top Yukawa can originate from an $SU(3)$ symmetric coupling in the UV of the form

$$\mathcal{L} \supset -(\lambda_1 \hat{t}_R^\dagger + \lambda_2 \hat{T}_R^\dagger)\phi^\dagger \hat{Q}_3 + \text{h.c.} \tag{49}$$

After ϕ acquires a vacuum expectation value, this leads to goldstone couplings of the form

$$\mathcal{L} = -f(\lambda_1 t_R^\dagger + \lambda_2 T_R^\dagger)T_L - \lambda_1 t_R^\dagger \tilde{H} Q_3 + \frac{\lambda_1}{2f}(H^\dagger H) t_R^\dagger T_L + \text{h.c.} + \dots \tag{50}$$

where the first term represents an explicit soft breaking of the global symmetry consistent with our choice of vacuum alignment. This means the Higgs can acquire radiative corrections proportional to this soft breaking, much as in supersymmetric theories.

The approximate mass eigenstates are T_L, t_L, and the linear combinations

$$t_R = \frac{\lambda_2 \hat{t}_R - \lambda_1 \hat{T}_R}{\sqrt{\lambda_1^2 + \lambda_2^2}} \tag{51}$$

$$T_R = \frac{\lambda_1 \hat{t}_R + \lambda_2 \hat{T}_R}{\sqrt{\lambda_1^2 + \lambda_2^2}} \tag{52}$$

In terms of the approximate mass eigenstates, we see the emergence of what we recognize as the usual top yukawa coupling, as well as an additional yukawa involving the new singlet fermions and irrelevant operators dictated by the vacuum manifold. These take the form

$$\mathcal{L} = -\lambda_t t_R^\dagger \tilde{H} Q_3 - \lambda_T T_R^\dagger \tilde{H} Q_3 + \frac{\lambda_1^2}{m_T}(H^\dagger H) T_R^\dagger T_L + \text{h.c.} + \ldots \tag{53}$$

where $m_T = \sqrt{\lambda_1^2 + \lambda_2^2} f$ and the yukawa couplings are related to the original interactions by

$$\lambda_t = \frac{\lambda_1 \lambda_2}{\sqrt{\lambda_1^2 + \lambda_2^2}} \tag{54}$$

$$\lambda_T = \frac{\lambda_1^2}{\sqrt{\lambda_1^2 + \lambda_2^2}} \tag{55}$$

which satisfy $\lambda_1^2 = \lambda_t^2 + \lambda_T^2$. Having promoted the top yukawa to an $SU(3)$-symmetric form, we know that the goldstone Higgs will be protected from radiative corrections through the top yukawa. But the explicit cancellation mechanism is a bit amusing: it amounts to a cancellation between the normal corrections coming from the yukawas, and the irrelevant interaction enforced by the goldstone nature of the Higgs:

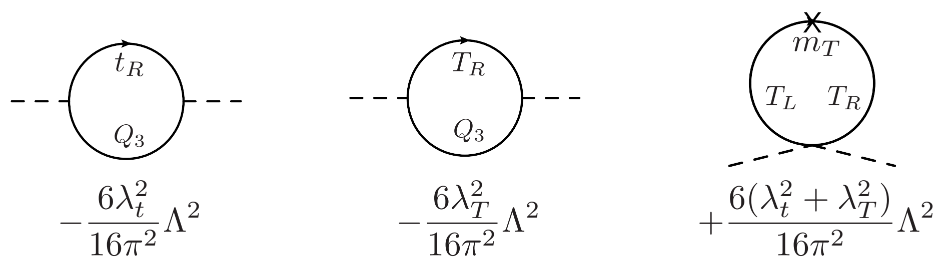

Figure 4: Quadratic divergence cancellation in the global symmetry case with light top partners.

Again the UV sensitivity is removed, and replaced with finite corrections proportional to the masses of the new states that restore the $SU(3)$ symmetry, exactly in analogy with SUSY. Of course, the non-observation of new physics then puts these models on the same footing as supersymmetry in terms of fine-tuning.

3.3 Vacuum selection

A final possibility is that nothing protects the Higgs mass, but rather there are many vacua of the Standard Model over which the Higgs mass varies according to some statistical distribution. If there is then a mechanism for selecting from the tail of the distribution with smaller Higgs masses, one has an explanation for the observed Higgs mass that does not rely on symmetries or a low cutoff.

For example, you can imagine an anthropic pressure fixing the weak scale in a universe where the dimensionful parameters of the Standard Model (i.e., the Higgs mass, or equivalently the vacuum expectation value) vary, but the dimensionless quantities are held fixed. In this case, v is bounded from above to be near its observed value by an argument known as the Atomic Principle [4].

Recall that for $v = v_{SM}$ the lightest baryons are the proton and neutron, of which the proton is lighter because the splitting due to quark masses exceeds the electromagnetic energy splitting:

$$m_n - m_p = (3v/v_{SM} - 1.7)\,\text{MeV}$$

So free neutrons decay into protons, with a reaction energy

$$Q = m_n - m_p - m_e = (2.5v/v_{SM} - 1.7)\,\text{MeV}$$

But in nuclei there is a binding energy that stabilizes the nuclei. The details are a bit complicated. The long-range part of the nucleon-nucleon potential is due to single pion exchange, with a range of $\sim 1/m_\pi$. For small u, d masses $m_\pi \propto ((m_u + m_d) f_\pi)^{1/2}$, so (neglecting the weak dependence of Λ_{QCD} on v) we have $m_\pi \sim v^{1/2}$.

We can mock up the binding energy in deuterons, the most weakly bound system, as a square well with a hard core to mimick short-range repulsion, which accounting for the dependence of the potential on v via m_π gives

$$B_d \simeq \left[2.2 - 5.5\left(\frac{v - v_{SM}}{v_{SM}}\right)\right]\,\text{MeV}$$

for small $v - v_{SM}$.

Now we see that as we increase v, we will eventually reach the point where $B_d < Q$ and the neutron is no longer stabilized by nuclear binding energy. This occurs for

$$v/v_{SM} \gtrsim 1.2$$

which is a tight bound, indeed! The deuteron is fairly important, since all primordial and stellar nucleosynthesis begins with deuterium. But this is not an airtight bound, as nuclei could form in violent astrophysical processes. The binding energies for heavier nuclei are larger, but for

$$v/v_{SM} \gtrsim 5$$

typical nuclei no longer stabilize the neutron against decay.

Assuming that stable protons and complex atoms are required for observers to form, this provides an anthropic pressure that favors $v \lesssim v_{SM}$. But it is clear that a robust constraint only exists if dimensionless couplings are held fixed; variation of the yukawas allows these constraints to be naturally evaded.

Indeed, it is possible to imagine a "weak-less" universe where the gauge group of the Standard Model is $SU(3)_c \times U(1)_{EM}$, and fermions appear in vector-like representations. It has been argued that such a universe undergoes big-bang nucleosynthesis, matter domination, structure formation, and star formation – i.e., sufficient stages of development to produce some form of observers. Of course, truly demonstrating that such a theory is capable of reproducing the physics necessary for forming observers is beyond the scope of a handful of theorists, but suffices to indicate that anthropic reasoning applied to the weak scale is sufficiently permeable.

4 New Hierarchy Solutions

4.1 Twin Higgs / Neutral naturalness

One interesting direction is to retain the symmetry-based approach but expand the scope of possible symmetries. The natural possibility is to work with discrete symmetries, rather than continuous ones. The idea is that the new particles required by a discrete symmetry need not carry the same Standard Model quantum numbers, and so are less strongly constrained by data from the LHC.

There are by now many different examples of neutral naturalness, but the simplest is the original: the Twin Higgs [5]. The idea is to introduce a mirror copy of the Standard Model along with a $\mathbb{Z}_2$ symmetry exchanging each field with its mirror counterpart. On top of this, one needs to assume an approximate global symmetry in the Higgs sector, which may be $U(4)$ or $O(8)$ depending on one's level of ambition. This global symmetry need not be exact, and is violated by all SM yukawa and gauge couplings, but should be an approximate symmetry of the Higgs potential.

For simplicity, we will consider the perturbative case where it suffices to work in terms of a $U(4) \simeq SU(4) \times U(1)$ approximate global symmetry, gauged by the Standard Model and twin electroweak interactions (i.e., gauging the $SU(2)\times SU(2)\times U(1)$ subgroup of $SU(4)$ and the additional $U(1)$). We can assemble the Higgs doublets H_A and H_B into a fundamental of $SU(4)$,

$$H = \left(\begin{array}{c} H_A \\ H_B \end{array} \right) \tag{56}$$

and, under the assumption that the Higgs sector potential is approximately $SU(4)$ symmetric, write down a potential of the form

$$V(H) = m^2 H^\dagger H + \lambda (H^\dagger H)^2 \tag{57}$$

For $m^2 < 0$, H acquires a vev and the $SU(4) \times U(1)$ is spontaneously broken to $SU(3) \times U(1)$, yielding seven goldstones. Depending on the vacuum alignment, all goldstones will be eaten, but it's also possible to align the vev entirely in the A sector or B sector by judicious adjustment of the potential. This adjustment is accomplished by terms that softly break the $U(4)$, and so ultimately will induce finite corrections to the goldstone mass through radiative corrections.

This theory accumulates radiative corrections from the usual couplings, for example top yukawas of the form

$$\lambda_A H_A Q_A t_A + \lambda_B H_B Q_B t_B \tag{58}$$

This gives the usual quadratic divergence,

$$\delta m^2 = -\frac{6}{16\pi^2}\Lambda^2 \left(\lambda_A^2 |H_A|^2 + \lambda_B^2 |H_B|^2 \right) \tag{59}$$

but the $\mathbb{Z}_2$ symmetry enforces $\lambda_A = \lambda_B = \lambda$, so that

$$\delta m^2 = -\frac{6\lambda^2}{16\pi^2}\Lambda^2 \left(|H_A|^2 + |H_B|^2 \right) \tag{60}$$

Now we can see the magic of the discrete symmetry. At the level of mass terms, the quadratic divergences respect the $U(4)$ symmetry. Thus the goldstones of the spontaneous breaking of the $U(4)$ symmetry will be protected against UV contributions.

We could continue to study the linear model (see, e.g., [6]), but it's convenient to focus on the low-energy theory of the goldstones [7]. In the limit where the vacuum expectation value lies entirely in the B sector, in B-sector unitary gauge we have

$$H = \begin{pmatrix} H_A \\ H_B \end{pmatrix} = \exp\left[\frac{i}{f}\left(\begin{array}{ccc|c} & & & h_1 \\ & & & h_2 \\ & & & 0 \\ \hline h_1^* & h_2^* & 0 & 0 \end{array}\right)\right]\begin{pmatrix} 0 \\ 0 \\ 0 \\ f \end{pmatrix} \equiv e^{i\pi/f} H_0 \tag{61}$$

Expanding out the exponential, we then get (up to a phase on h)

$$H = \begin{pmatrix} h \frac{if}{\sqrt{h^\dagger h}} \sin\left(\frac{\sqrt{h^\dagger h}}{f}\right) \\ 0 \\ f \cos\left(\frac{\sqrt{h^\dagger h}}{f}\right) \end{pmatrix} \tag{62}$$

where $h = (h_1, h_2)^T$ Then we can immediately expand out H_A and H_B in terms of the goldstone modes, obtaining

$$H_A = h \frac{f}{\sqrt{h^\dagger h}} \sin\left(\frac{\sqrt{h^\dagger h}}{f}\right) = h + \ldots \tag{63}$$

$$H_B = \begin{pmatrix} 0 \\ f \cos\left(\frac{\sqrt{h^\dagger h}}{f}\right) \end{pmatrix} = \begin{pmatrix} 0 \\ f - \frac{1}{2}\frac{h^\dagger h}{f} + \ldots \end{pmatrix} \tag{64}$$

The goldstones inherit Yukawa couplings (which break the $U(4)$)

$$\lambda_A H_A Q_A t_A + \lambda_B H_B Q_B t_B \to \lambda_A h Q_A t_A + \lambda_B \left(f - \frac{h^\dagger h}{2f}\right) Q_B t_B + \ldots \tag{65}$$

and now we can see in detail the cancellation of quadratic divergences, in exact analogy with the case of a continuous global symmetry: As in our other symmetry examples, the UV sensitivity is replaced by finite corrections coming from the mass of the SM-enutral top partners, which violate the accidental $U(4)$ through the soft breaking terms in the potential.

But in contrast to older continuous symmetry approaches, now there are no direct constraints on the partner particles. There is still v/f tuning and Higgs coupling deviations, so the scenario is bounded, but limits are at the 10% level and unlikely to improve significantly during the remainder of the LHC era.

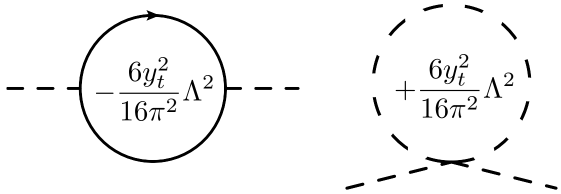

Figure 5: Quadratic divergence cancellation in the discrete global symmetry case.

There are various generalizations of this idea. One can trivially construct $\mathbb{Z}_N$ models, which generalize naturally to multiple sectors. Alternately, one can construct "fraternal" models where the $\mathbb{Z}_2$ symmetry is only a good symmetry for the states most relevant to the Higgs potential [6]. One can also construct more elaborate symmetry structures using orbifold projections [8]. The signatures are rich and interesting and worth looking for enthusiastically in the remaining lifetime of the LHC.

4.2 Relaxion

In some sense, neutral naturalness is the most conservative "new" idea, retaining an old mechanism (symmetry protection) and pushing the specific realization in a new direction. But an even more exciting thing about the modern era is that we are now beginning to see radically new ideas that don't fit into traditional paradigms. The most exciting recent ideas involve dynamics to select the Higgs mass from a range of values consistent with a cutoff well above the weak scale.

4.2.1 QCD/QCD′ Relaxion

The simplest and original incarnation [9] is that of a QCD axion-like particle ϕ coupled to the Standard Model, with an additional inflationary sector whose properties will turn out to be somewhat special. Here we emphasize *axion-like* because the axion-like field will not be manifestly compact, but rather possess only a shift symmetry. This shift symmetry will be broken by a small, dimensionful coupling to the Higgs. We will circle back to these features, and their potential relation to technical naturalness arguments, towards the end.

We envision enlarging the Standard Model with the following terms:

$$\delta\mathcal{L} = (-M^2 + g\phi)|H|^2 + V(g\phi) + \frac{1}{32\pi^2}\frac{\phi}{f}\tilde{G}^{\mu\nu}G_{\mu\nu} \tag{66}$$

where M is of the order of the cutoff of the SM Higgs sector, H is the Higgs doublet, g is the dimensionful coupling that breaks the shift symmetry, and

$$V(g\phi) \sim gM^2\phi + g^2\phi^2 + \ldots$$

parameterizes the non-derivative terms solely involving ϕ. We will be interested in field values of ϕ that greatly exceed f, so we should understand it as a non-compact field. Now clearly when $g/M \to 0$ the Lagrangian has a shift symmetry $\phi \to \phi + 2\pi f$, and g can be treated as a spurion for breaking of the shift symmetry.

Below the QCD confinement scale, the coupling between ϕ and the gluon field strength gives rise to the familiar periodic axion potential

$$\frac{1}{32\pi^2}\frac{\phi}{f}\tilde{G}^{\mu\nu}G_{\mu\nu} \to \Lambda^4 \cos(\phi/f) \tag{67}$$

For values of the Higgs vev near the Standard Model value, the height of the cosine potential is

$$\Lambda^4 \sim f_\pi^2 m_\pi^2 \sim yvf_\pi^3 \tag{68}$$

where m_π^2 changes linearly with the quark masses, and so the barrier height is linearly proportional to the Higgs vev (at least roughly speaking; there are of course logarithmic corrections from the contributions to QCD running).

Now the idea is clear: starting at values of ϕ such that the total Higgs mass is large and positive, and assuming the slope of the ϕ potential causes it to evolve in a direction that lowers the Higgs mass, the ϕ potential will initially be completely dominated by the $g\phi$ potential terms, until the point at which the total Higgs mass-squared goes from positive to negative and the Higgs acquires a vacuum expectation value. At this point the wiggles due to the quark masses grow linearly in the Higgs vev, and generically ϕ will stop when the slope of the QCD-induced wiggles matches the slope of $V(\phi)$. This classical stopping point occurs for

$$g \sim \frac{yvf_\pi^3}{M^2 f} \tag{69}$$

This allows for a light Higgs (i.e., a small total Higgs mass-squared and small electroweak scale) relative to a cutoff M provided $g/M \ll 1$. For example, with a QCD axion decay constant $f = 10^9$ GeV and $M \sim 10^7$ GeV we have $g/M \sim 10^{-30}$.

So far we have only accounted for the parametrics of the potential, neglecting the actual dynamical process. In the minimal realization of the relaxion mechanism, ϕ is made to roll slowly by imagining that its evolution occurs during a period of inflation, such that Hubble friction provides efficient dissipation of kinetic energy in ϕ.

Now there are various considerations that must be taken into account. They are:

1. In order to sensibly yield a Higgs mass much smaller than the cutoff, ϕ must scan over a sufficiently large range such that m_H^2 varies from $\mathcal{O}(M^2)$ to $\mathcal{O}(0)$. Thus the field range of interest is $\Delta\phi \sim M^2/g$. Inflation must endure for the entirety of this scanning. During N e-folds of inflation, the field rolls by an amount

$$\Delta\phi \sim N\dot{\phi}/H \sim NV_\phi'/H^2 \sim NgM^2/H^2 \tag{70}$$

   where the first expression just relates the displacement to the velocity of the slow-rolling field and the duration of inflation, the second uses slow-roll conditions for ϕ, and the third uses the leading form of V_ϕ. Requiring that this cover a change of order M^2/g implies the number of e-folds of inflation is at least

$$N \gtrsim \frac{H^2}{g^2} \tag{71}$$

2. The scanning of ϕ results in a change in vacuum energy of order M^4. We require the vacuum energy during inflation to exceed this change so that the dynamics is dominated by inflation throughout the evolution of ϕ. This amounts to requiring

$$H > \frac{M^2}{M_{Pl}} \tag{72}$$

3. During the inflationary epoch, the evolution of ϕ involves both classical rolling and quantum fluctuations. If this were not the case, different patches of the universe could end up in different electroweak vacua. Classical rolling beats quantum fluctuations.

$$H < \frac{V_\phi'}{H^2} \Rightarrow H < (gM^2)^{1/3} \tag{73}$$

4. Finally, it should be the case that the barriers from QCD are higher than the Hubble scale during inflation, so that the barriers are sufficient to stop scanning. This amounts to

$$H < \Lambda_{QCD} \tag{74}$$

which in general is superseded by the previous requirement.

Putting everything together, we can see that the cutoff of the theory is at most

$$M \lesssim \left(\frac{\Lambda^4 M_{Pl}^3}{f}\right)^{1/6} \sim 10^7 \text{ GeV} \times \left(\frac{10^9 \text{ GeV}}{f}\right)^{1/6} \tag{75}$$

It is worth pausing to work out the numerical consequences. Maximizing the cutoff, we have $g \simeq 10^{-23}$ GeV, so $H < 1$ MeV, $N = 10^{40}$, and the field range is $\Delta\phi = 10^{47}$ GeV. While the first two problems are aesthetic in nature, the third is more severe. It requires the relaxion potential to be valid over field ranges vastly in excess of the Planck scale. In general it is difficult to protect a potential over trans-Planckian field ranges, and – as we will discuss more shortly – particularly so in this case.

Unfortunately, even if all of these criteria are satisfied, there is an observational problem with this simplest scenario. The field ϕ stops not at the minimum of the QCD cosine potential (for which the effective θ angle is zero), but is rather displaced by an amount proportional to the slope of ϕ. This amounts to $\theta \sim 1$, which is excluded by bounds on the neutron EDM that constrain $\theta \lesssim 10^{-11}$. So the mechanism is ruled out by a natural prediction, though it is certainly no fault of the mechanism.

A simple solution is to repeat all of the same ingredients, but make the relaxion a non-compact axion of another gauge group for which constraints on the θ parameter are weaker or nonexistent. This scenario should involve quarks of a new gauge group that are also charged under the electroweak gauge group. For example, consider adding vector-like lepton doublets L, L^c, N, N^c with charges

Field	$SU(3)_N$	$SU(3)_C$	$SU(2)_L$	$U(1)_Y$
L	$\square$	–	$\square$	$-1/2$
L^c	$\overline{\square}$	–	$\square$	$+1/2$
N	$\square$	–	–	$+1$
N^c	$\overline{\square}$	–	–	$+1$

This model is now subject to a variety of additional constraints, namely

1. The quarks of the new gauge group must get most of their mass from the Higgs:

$$\mathcal{L} \supset m_L L L^c + m_N N N^c + y H L N^c + y' H^\dagger L^c N \tag{76}$$

2. The new gauge group must confine with light flavor,

$$\Lambda^4 \simeq 4\pi f_{\pi'}^3 m_N \tag{77}$$

3. The natural size of the smallest mass from the see-saw (assuming a heavy L) is

$$m_N \geq y y' v^2 / m_L \tag{78}$$

4. The see-saw mass is at least as large as the radiative Dirac mass

$$m_N \geq \frac{y y'}{16\pi^2} m_L \log(M/m_L) \tag{79}$$

5. The wiggles in the potential due to EWSB exceed the wiggles due to confinement alone

$$m_N \geq y y' f_{\pi'}^2 / m_L \tag{80}$$

Taken together, these bounds imply $f_{\pi'} < v$ and

$$m_L < \frac{4\pi v}{\sqrt{\log(M/m_L)}} \tag{81}$$

That is to say, although the mechanism lives in a sector distinct from the Standard Model, the scale of new physics still lies near the weak scale.

The details of the inflationary scenario are similar, though now the axion is not a QCD axion so the constraints on the PQ scale are not as stringent. Taking it to be of the same order a the cutoff, the cutoff in this case is pushed to

$$M < 2 \times 10^8\,\mathrm{GeV} \left(\frac{f_{\pi'}}{30\,\mathrm{GeV}}\right)^{4/7} \left(\frac{M}{f}\right)^{1/7} \tag{82}$$

Caveats At this point it is worth mentioning several caveats to these scenarios that may compromise or spoil the mechanism. To be clear, the mechanism is brilliant, and the problems are modest compared to the originality of the mechanism. In any event, the first caveat relates to the cosmological constant. In symmetry solutions to the hierarchy problem, one can effectively factorize the solution to the CC problem from the solution to the hierarchy problem – because there is one value for the weak scale, one tuning (or other mechanism) can then set the CC to the observed value. In the relaxion scenario, the cosmological constant changes by large amounts from minimum to minimum over which the changes to the Higgs mass are negligible. From one minimum to another, we have $\Delta\phi \sim f$ and thus $\Delta V \sim gfM^2 \sim \Lambda^4$, while the change in the Higgs mass-squared is infinitesimally small, $\Delta m_H^2 \sim gf \sim \Lambda^4/M^2$. So while there are many vacua with Higgs masses-squared at the electroweak scale, the changes in the cosmological constant from vacuum to vacuum are all vastly larger than the observed cosmological constant.

From this it is tempting to argue that one needs enough vacua to scan the full range of the CC for each viable electroweak minimum. This would require an even larger tuning than one would require to tune the CC in a theory with a unique electroweak vacuum. On the other hand, these arguments are not necessarily well-defined.

The second issue relates to the technical naturalness of the scenario, or the lack thereof. As constructed, ϕ possesses a non-compact shift symmetry. While the parameter g breaks the shift symmetry, when $g \to 0$ no compact global symmetry is restored. Thus the theory does not satisfy the typical considerations of technical naturalness.

Let's pull this apart a bit more. If we have a theory with an exact global symmetry that is spontaneously broken, then the effective action of the theory has a continuous symmetry under which

$$\phi \to \phi + \alpha f \tag{83}$$

for any real α. But a subgroup of this is gauged, in the sense that

$$\phi \to \phi + 2\pi k f \tag{84}$$

is a gauge symmetry for $k \in \mathbb{Z}$ because ϕ is really an angle, and no local operator can break the angular periodicity. In practice, in QCD this means that quark masses and anomaly couplings break the continuous shift symmetry but preserve the discrete one.

Concretely, one can think of a $U(1)$ global symmetry spontaneously broken by the vev of a complex scalar Φ, expressed via a non-linear mapping $\Phi \to \rho e^{i\phi/f}$. This parameterization has a clear invariance under $\phi \to \phi + 2\pi k f$ in the sense that it maps Φ back to itself. Explicit breaking of the symmetry in terms of operators involving Φ will still have this invariance purely from the mapping between Φ and ϕ.

As far as the relaxion goes, the coupling g breaks both the global symmetry and the gauge symmetry. If the relaxion were to be a genuine goldstone (or, more restrictively, expressly the QCD axion), then the potential – and the Higgs mass – would need to be a periodic in $2\pi f$ if it's to come from a local QFT. Since the theory requires a non-periodic field excursion of order $\phi \sim M^2/g$, this would imply $f > M^2/g$. This ultimately forces the cutoff of the theory to live down at the weak scale, giving no parametric improvement [10].

Thus we are forced to conclude that the relaxion is not an axion, and the shift symmetry does not arise from a compact global symmetry. So what if the relaxion is *not* an axion? In this case, there is no compact global symmetry, and no mechanism to protect the shift symmetry over field excursions beyond the Planck scale. One expects quantum gravity effects to alter the picture significantly, preventing the large field excursions required for the mechanism to operate. There have been attempts to model-build a relaxion from multiple compact fields arranged to give larger effective periods, but it is not obvious that these attempts are successful.

4.2.2 Interactive Relaxion

Given the challenges facing the original relaxion mechanism, it is worth asking if there are other mechanisms that might work along similar lines. Indeed, there are several, of which one is worth briefly sketching here. Whereas the initial realization of the relaxion has an omnipresent source of dissipation and a potential that turns on near $m_H^2 = 0$, this alternative has an omnipresent potential and a source of dissipation that turns on near $m_H^2 = 0$ [11].

The basic idea is to start with a relaxion of the familiar form, for an Abelian Higgs toy model

$$\delta\mathcal{L} = (-M^2 + g\phi)|H|^2 + V(g\phi) + \Lambda^4 \cos\frac{\phi}{f'} + \frac{\phi}{4f} F\tilde{F} \tag{85}$$

where the cosine potential is an axion potential generated from the confinement of some non-SM gauge group – so that Λ is not related to the Higgs vev – and $F\tilde{F}$ is some abelian gauge group (we'll get to the SM version momentarily). The coupling between the relaxion and the $U(1)$ gauge field will be a source of particle production, which will provide dissipation.

The essential idea is for ϕ to start at some large field value, $\phi \sim M^2/g$ with some nonzero velocity, and from the direction in which the Higgs mass-squared is large and negative. In this case, the abelian gauge group is Higgsed, with correspondingly large masses. For $\dot{\phi} \gtrsim \Lambda^2$, ϕ then rolls down its potential without slowing on the cosine bumps, such that the Higgs mass-squared decreases in magnitude. Eventually, the vev becomes small enough that ϕ can dissipate kinetic energy through production of gauge bosons.

For simplicity, we will consider the process at zero temperature. The equations of motion for the transverse modes of the gauge field A – call them $A_\pm$ – in unitary gauge ($\partial_\mu A^\mu = 0$) are

$$\ddot{A}_\pm + \left(k^2 + m_A^2 \pm \frac{k\dot{\phi}}{f} \right) A_\pm = 0 \tag{86}$$

Neglecting backreaction on ϕ, and treating $\dot{\phi}$ as constant, the solutions are

$$A_\pm(k) \propto e^{i\omega_\pm t} \tag{87}$$

$$\omega_\pm^2 = k^2 + m_A^2 \pm \frac{k\dot{\phi}}{f} \tag{88}$$

There is a tachyonic growing mode for imaginary frequencies, corresponding to

$$\omega_\pm^2 = k^2 + m_A^2 \pm \frac{k\dot{\phi}}{f} < 0 \Rightarrow |\dot{\phi}| \gtrsim 2fm_A \tag{89}$$

The tachyonically growing mode drains the kinetic energy of ϕ exponentially quickly, as the growing mode backreacts on ϕ.

Of course, for a fully accurate picture the analysis must be repeated at finite temperature. While the qualitative picture persists, some subtleties arise, including the fact that exponential growth only occurs for abelian gauge fields at finite temperature.

To implement this mechanism in the Standard Model, ϕ must couple to a linear combination of electroweak gauge bosons, but cannot couple to pairs of photons. If it coupled to pairs of photons, it could dissipate energy into photon pair production irrespective of the value of the Higgs vev. Rather, we require it to dissipate energy only to gauge fields acquiring mass through the Higgs mechanism. The natural candidate is thus

$$\mathcal{L} \supset \frac{\phi}{f}\left(\alpha_Y B\tilde{B} - \alpha_2 W\tilde{W}\right) \tag{90}$$

where this linear combination contains all appropriate pairs of electroweak gauge bosons except $\gamma\gamma$ Such a coupling might look like a fine-tuning, but can be protected in a UV model for the axion where the SM electroweak group is embedded in an $SU(2)_L \times SU(2)_R$ gauge theory. In such a theory there is a PQ symmetry under which

$$\phi \to \phi + \alpha \tag{91}$$

$$\theta_L \to \theta_L - \alpha \tag{92}$$

$$\theta_R \to \theta_R + \alpha \tag{93}$$

where $\theta_{L,R}$ are the θ angles of $SU(2)_{L,R}$ respectively. This forces ϕ to couple to $W_L\tilde{W}_L - W_R\tilde{W}_R$. The combination $\gamma\tilde{\gamma}$ is invariant under the PQ symmetry, and so can only appear in the combination $\propto (\theta_L + \theta_R)\gamma\tilde{\gamma}$, i.e., cannot couple to ϕ. In this way we can forbid the $\gamma\tilde{\gamma}$ coupling with symmetries.

There are various other subtleties in this scenario, too many to enumerate here, but hopefully we have articulated the sense in which there are multiple possible realizations of the essential relaxion mechanism.

4.3 NNaturalness

An alternative that proceeds from similar inspiration is to put many copies of the Standard Model in the same universe, but explain why one copy acquires the dominant energy density [12].

The idea is to envision N sectors which are mutually decoupled. For simplicity, we could take it to be N copies of the Standard Model, though this is not an important restriction. From copy to copy, we imagine the Higgs mass parameters are distributed in some range from $-\Lambda_H^2$ to Λ_H^2 according to some probability distribution. For a wide range of distributions, the generic expectation is that some sectors have accidentally small Higgs masses, $m_H^2 \sim \Lambda_H^2/N$. For large enough N, this implies that there is a sector whose electroweak scale is well below the cutoff, which we might identify with "our" Standard Model.

Reversing the argument, this implies that the cutoff of the theory should be

$$\Lambda_H \sim \sqrt{N}|m_H|$$

E.g. a cutoff of 10 TeV corresponds to $N = 10^4$, whereas a cutoff of 10^{10} GeV requires $N = 10^{16}$.

There is another factor in play when N is large. While the naive scale of quantum gravity is M_{Pl}, in the presence of a large number of species the scale at which gravity becomes strongly coupled is lowered,

$$\Lambda_G^2 \sim M_{Pl}^2/N$$

You can think of this as just coming from wavefunction renormalization of the graviton by N fields whose contributions are dominated by the scale N. This implies the effective Planck scale should be at least $M_{Pl}^2 \sim N\Lambda G^2$. Solving the entire hierarchy problem this way would entail $N = 10^{32}$. However, this lowers the cutoff of quantum gravity to the weak scale, and gives us the usual problems associated with a low cutoff.

But we would naturally have one sector with the observed value of the weak scale and a Higgs cutoff associated with the cutoff of quantum gravity for $N = 10^{16}$, for which $\Lambda_H = \Lambda_G = 10^{10}$ GeV. Alternately, we could preserve a notion of grand unification for $N = 10^4$, for which quantum gravity grows strong at 10^{16} GeV, and something like supersymmetry enters at $\Lambda_H = 10$ TeV to cut off the Higgs sector.

The question, then, is to explain why this sector with "our" Standard Model is populated, while all of the other sectors are not. As with the relaxion, this is accomplished through cosmology. In a universe with many sectors, the universe is populated by whatever sectors are abundant. If all sectors had a thermal abundance, there would be an enormous contribution to the energy density of the universe, and we would not have any ability to understand why we are the sector with the smallest scales. Thus we can imagine a cosmological mechanism that preferentially reheats sectors with smaller scales.

The simplest way to accomplish this is to imagine an inflationary epoch, followed by reheating due to the decay of some reheaton. To avoid tuning, this reheaton should couple universally to all sectors. The Standard Model can be preferentially reheated (i.e., absorb most of the energy from the reheaton decays) if the branching ratio of the reheaton to each sector scales like an inverse power of the (absolute value of the) Higgs mass-squared in each sector.

The simplest example is of a scalar ϕ with couplings

$$\mathcal{L} \supset -a\phi \sum_i |H_i|^2 - \frac{1}{2}m^2\phi^2 \tag{94}$$

The branching ratios of ϕ to each sector depend on its mass and whether or not electroweak symmetry is broken in each sector (in general, it will be broken in half and unbroken in the other half). If we imagine that $m_\phi \ll |m_H|$ in all the sectors, then we can work out the branching ratios by integrating out the Higgses and gauge bosons (when massive) in each sector.

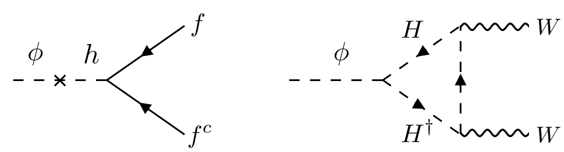

Figure 6: Dominant decays when $\langle H \rangle \neq 0$ (left) and $\langle H \rangle = 0$ (right)

For sectors where electroweak symmetry is broken, the dominant decay is into fermions, via

$$\mathcal{L} \supset ay\frac{v}{m_h^2}\phi q q^c \tag{95}$$

whereas when electroweak symmetry is unbroken the dominant decay is into gauge bosons, via

$$\mathcal{L} \supset a\frac{g^2}{16\pi^2}\frac{1}{m_H^2}\phi W_{\mu\nu}W^{\mu\nu} \tag{96}$$

Thus the decay rate into broken-phase sectors scales as $1/m_h^2$, while the decay into unbroken-phase sectors scales as $1/m_H^4$. Reheaton decays prefer a sector with broken electroweak symmetry and the smallest possible value of m_h.

The resulting energy density of each sector is proportional to the decay width,

$$\frac{\rho_i}{\rho_{us}} \simeq \frac{\Gamma_i}{\Gamma_{us}} \tag{97}$$

This leads to some energy density in the sectors nearest to ours in mass, with attendant predictions for dark radiation within the reach of future CMB experiments.

5 Rampant Speculation

5.1 UV/IR mixing

Let's end with an excursion into radically different territory, which marks a sharp departure from the types of solutions considered thus far. One way to frame the hierarchy problem is as a separation of UV physics from IR physics in effective field theory: the theory in the far UV knows nothing about the theory in the far IR.

From this perspective, one might hope to work around the hierarchy problem by linking the far UV and the far IR. This would represent a sharp departure from effective field theory, and the challenge is to make the departure well posed.

Thankfully, we have examples of UV/IR mixing. Perhaps the most famous is in quantum gravity. We can imagine accelerating two protons to Planckian energies and smashing them together to create Planck-length-sized black holes. You might then hope to probe distances shorter than the Planck length by increasing the energy of the two protons above the Planck energy. But when you do so, you create larger and larger black holes. More energetic protons mean more massive black holes, which have larger radii. Instead of probing shorter distances, you produce large black holes which resolve only longer distances – exciting the theory in the UV really probes the physics of the IR.

A more precise version of the same thing happens with T duality, which relates string theories propagating on some circle of radius R and $1/R$.

Of course, absent a complete theory of quantum gravity, it is difficult to understand what bearing this might have on the hierarchy problem. So it is fruitful to look into quantum field theories in which similar UV/IR mixing arises. Thankfully, we have a well-posed example in the guise of quantum field theory on noncommutative backgrounds [13].

The starting point is to imagine a nonvanishing commutator between coordinates on $\mathbb{R}^4$,

$$[x^\mu, x^\nu] = i\Theta^{\mu\nu}$$

where Θ is a constant, real, antisymmetric noncommutativity matrix. The algebra of functions on this noncommutative space can be viewed as an algebra of ordinary functions on the usual $\mathbb{R}^4$ with the product deformed to the noncommutative, associative star product,

$$(\phi_1 \star \phi_2)(x) = e^{\frac{i}{2}\Theta^{\mu\nu}\partial_\mu^y \partial_\nu^z} \phi_1(y)\phi_2(z)\Big|_{y=z=x} \tag{98}$$

So we are studying theories whose fields are functions on ordinary $\mathbb{R}^4$ with ordinary actions, except that products of fields are replaced by the star product.

To see evidence for UV/IR mixing, it suffices to consider the appropriate generalization of ϕ^4 theory. This is a theory with a mass gap and quadratic divergences in the commutative version. The non-commutative version is simply

$$S = \int d^4x \left(\frac{1}{2}(\partial\phi)^2 + \frac{1}{2}m^2\phi^2 + \frac{1}{4!}\lambda\phi \star \phi \star \phi \star \phi \right) \tag{99}$$

where the star product in the quadratic pieces of the action reduces to the normal commutative products up to total derivatives; only the interactions are modified. This amounts to modifying the Feynman rules so that the interaction vertex has an additional phase factor of the form

$$e^{-\frac{i}{2}\sum_{i<j} k_i \times k_j}$$

where k_i is the momentum flowing into the vertex through the ith field and the "cross product" is

$$k_i \times k_j \equiv k_{i\mu}\Theta^{\mu\nu}k_{j\nu}$$

This phase factor is invariant under cyclic permutations, but not arbitrary permutations. In a Feynman diagram with fixed external legs, there are then "planar" graphs, where propagators don't cross on their way to external states, and "non-planar" graphs where propagators cross.

At one loop, the two-point function receives corrections from one planar graph and one non-planar graph:

The two diagrams give

$$\text{Planar} \sim \frac{\lambda}{3}\int \frac{d^4k}{(2\pi)^4}\frac{1}{k^2+m^2} \tag{100}$$

$$\text{Non}-\text{planar} \sim \frac{\lambda}{6}\int \frac{d^4k}{(2\pi)^4}\frac{e^{ik\times p}}{k^2+m^2} \tag{101}$$

$$\tag{102}$$

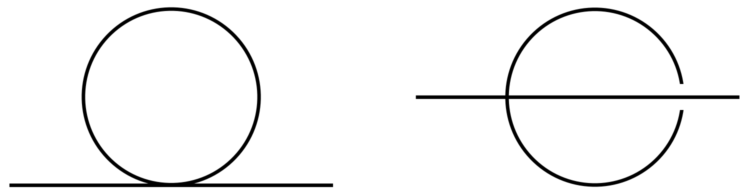

where the planar one is just the usual quadratically divergent graph, and the non-planar one picks up a phase factor from the crossing of an internal line. To see the effect of the phase factor, we can re-write the propagators in terms of Schwinger parameters

$$\frac{1}{k^2+m^2} = \int_0^\infty d\alpha e^{-\alpha(k^2+m^2)} \tag{103}$$

to get gaussian integrals

$$\mathrm{Planar} \sim \frac{\lambda}{48\pi^2}\int \frac{d\alpha}{\alpha^2} e^{-\alpha m^2} \tag{104}$$

$$\mathrm{Non-planar} \sim \frac{\lambda}{96\pi^2}\int \frac{d\alpha}{\alpha^2} e^{-\alpha m^2 - p\circ p/\alpha} \tag{105}$$

$$\tag{106}$$

where $p \circ q = -p_\mu \Theta^2_{\mu\nu} q_\nu$ has dimensions of $1/\mathrm{mass}^2$. These integrals are divergent, which we regulate by multiplying the integrand by a smooth cutoff $e^{-1/(\Lambda^2\alpha)}$. Then we find the graphs give the following contributions

$$\mathrm{Planar} \sim \frac{\lambda}{48\pi^2}\left(\Lambda^2 - m^2\log(\Lambda^2/m^2) + \dots\right) \tag{107}$$

$$\mathrm{Non-planar} \sim \frac{\lambda}{96\pi^2}\left(\Lambda^2_{eff} - m^2\log(\Lambda^2_{eff}/m^2) + \dots\right) \tag{108}$$

$$\tag{109}$$

where

$$\Lambda^2_{eff} = \frac{1}{1/\Lambda^2 + p\circ p}$$

In this latter case, taking $\Lambda \to \infty$ gives $\Lambda_{eff} = \frac{1}{p\circ p}$. Taking $p \to 0$ then gives $\Lambda_{eff} \to \infty$. That is to say, the non-planar diagram generates an IR divergence from what we normally think of as a UV divergence.

Needless to say, this represents a striking breakdown of Wilsonian EFT, and Wilsonian renormalization fails. But it provides a suggestive hint. A Wilsonian effective field theorist would see the existence of an IR divergence and interpret it as a new light particle. In particular, the above IR divergence could be mimicked by adding to the theory a new light field coupled to ϕ,

$$\delta S = \int d^4x \left(\frac{1}{2}\partial\chi \circ \partial\chi + \frac{1}{2}\Lambda^2(\partial \circ \partial\chi)^2 + \frac{i\sqrt{\lambda}}{\sqrt{96\pi^2}}\chi\phi \right) \tag{110}$$

The existence and lightness of the field χ is inexplicable from the perspective of Wilsonian EFT, but can be understood merely as an interpretation of the IR divergences resulting from UV/IR mixing in the non-commutative theory.

Of course, this is a long way from solving the hierarchy problem. The field χ doesn't look anything like a standard propagating degree of freedom in Lorentzian signature. But it points to a qualitatively interest direction in which to probe the hierarchy problem, one which is unlike any we have encountered before. If the hierarchy problem is solved by radically new ideas in quantum field theory, I am willing to bet that it will proceed somewhere along these lines of UV/IR mixing.

6 Conclusion

Thus we come to an end. Hopefully I have illustrated to you what the hierarchy problem is, and what it is not. There are old solutions which are compelling but in tension with data, and new solutions which are born of necessity and take us in wildly new directions. Time considerations have prevented us from exploring the full set of new directions, and some of my favorites which you may wish to investigate further include approaches using conformal symmetry [14] and ones using disorder [15].

If nothing else, hopefully is far from clear to you that we have systematically studied all such solutions. Many new directions remain, some of which proceed along avenues sketched here, and some which have yet to be imagined. Null results in conventional channels free us to break new ground. This is where you come in!

References

[1] P. A. M. Dirac, "New basis for cosmology," *Proc. Roy. Soc. Lond.* **A165** (1938) 199–208.

[2] M. Gell-Mann, "The interpretation of the new particles as displaced charge multiplets," *Nuovo Cim.* **4** (1956) no. S2, 848–866.

[3] M. Perelstein, M. E. Peskin, and A. Pierce, "Top quarks and electroweak symmetry breaking in little Higgs models," *Phys. Rev.* **D69** (2004) 075002, `arXiv:hep-ph/0310039 [hep-ph]`.

[4] V. Agrawal, S. M. Barr, J. F. Donoghue, and D. Seckel, "The Anthropic principle and the mass scale of the standard model," *Phys. Rev.* **D57** (1998) 5480–5492, `arXiv:hep-ph/9707380 [hep-ph]`.

[5] Z. Chacko, H.-S. Goh, and R. Harnik, "The Twin Higgs: Natural electroweak breaking from mirror symmetry," *Phys. Rev. Lett.* **96** (2006) 231802, `arXiv:hep-ph/0506256 [hep-ph]`.

[6] N. Craig, A. Katz, M. Strassler, and R. Sundrum, "Naturalness in the Dark at the LHC," *JHEP* **07** (2015) 105, `arXiv:1501.05310 [hep-ph]`.

[7] G. Burdman, Z. Chacko, R. Harnik, L. de Lima, and C. B. Verhaaren, "Colorless Top Partners, a 125 GeV Higgs, and the Limits on Naturalness," *Phys. Rev.* **D91** (2015) no. 5, 055007, `arXiv:1411.3310 [hep-ph]`.

[8] N. Craig, S. Knapen, and P. Longhi, "Neutral Naturalness from Orbifold Higgs Models," *Phys. Rev. Lett.* **114** (2015) no. 6, 061803, `arXiv:1410.6808 [hep-ph]`.

[9] P. W. Graham, D. E. Kaplan, and S. Rajendran, "Cosmological Relaxation of the Electroweak Scale," *Phys. Rev. Lett.* **115** (2015) no. 22, 221801, `arXiv:1504.07551 [hep-ph]`.

[10] R. S. Gupta, Z. Komargodski, G. Perez, and L. Ubaldi, "Is the Relaxion an Axion?," *JHEP* **02** (2016) 166, `arXiv:1509.00047 [hep-ph]`.

[11] A. Hook and G. Marques-Tavares, "Relaxation from particle production," *JHEP* **12** (2016) 101, `arXiv:1607.01786 [hep-ph]`.

[12] N. Arkani-Hamed, T. Cohen, R. T. D'Agnolo, A. Hook, H. D. Kim, and D. Pinner, "Solving the Hierarchy Problem at Reheating with a Large Number of Degrees of Freedom," *Phys. Rev. Lett.* **117** (2016) no. 25, 251801, `arXiv:1607.06821 [hep-ph]`.

[13] S. Minwalla, M. Van Raamsdonk, and N. Seiberg, "Noncommutative perturbative dynamics," *JHEP* **02** (2000) 020, `arXiv:hep-th/9912072 [hep-th]`.

[14] P. H. Frampton and C. Vafa, "Conformal approach to particle phenomenology," `arXiv:hep-th/9903226 [hep-th]`.

[15] I. Z. Rothstein, "Gravitational Anderson Localization," *Phys. Rev. Lett.* **110** (2013) no. 1, 011601, `arXiv:1211.7149 [hep-th]`.