Beyond the Standard Model

Lectures given at the DESY Summer Student Programme 2017, Hamburg, Germany

Felix Brümmer
Laboratoire Univers et Particules de Montpellier, Université de Montpellier, France
felix.bruemmer@umontpellier.fr

1 Introduction

Particle physics beyond the Standard Model is a vast subject, and two lectures are a pretty short time. I won't even try to give a comprehensive overview. I will hardly mention a lot of things that should be considered part of this field, such as flavour physics, baryogenesis, grand unification, string theory, the strong CP problem... Instead I will concentrate on what for many theorists is the most compelling argument for new physics beyond the Standard Model *not just anywhere but within the reach of present-day colliders*. Namely, the Higgs sector is said to suffer from the *electroweak hierarchy problem* when the Standard Model is embedded in a more complete theory. I will explain the hierarchy problem, and introduce the two currently most fashionable proposals for solving it, *supersymmetry* and *composite Higgs models*. Finally, I will sketch the present experimental situation regarding these two paradigms.

To convey the main ideas of even these limited topics, I'll have to gloss over quite a few subtleties, and on occasion I will even be forced to cheat about some technicalities for the benefit of, hopefully, getting the essential points across. If these lectures encourage some of you to learn more about these subjects properly and in detail, they will have served their purpose.

2 The electroweak hierarchy problem

Some common misunderstandings and half-truths

Here are a few somewhat misleading statements which you frequently hear in the introductory parts of talks on BSM phenomenology. These statements are usually preceded by "As everyone knows, ..." and go without being questioned, even by those who know better.

1. In the Standard Model, when calculating loop diagrams in perturbation theory, the Higgs boson mass is found to diverge quadratically with the ultraviolet momentum cutoff. Therefore you need to fine-tune a large bare Higgs mass against large counterterms in order to obtain a small renormalized Higgs mass. This is called the hierarchy problem.
2. You need to repeat this fine tuning at each order in perturbation theory, which makes the problem even worse.
3. The main culprit is the top quark because of its $\mathcal{O}(1)$ Yukawa coupling to the Higgs boson. Therefore new physics is needed first and foremost to cancel the (quadratically divergent pieces of the) top loop.

Let's not be afraid of being pedantic and try to disentangle what's right from what's wrong in these statements.

1. The quadratic divergence is indeed an indicator that something unnatural happens when we try to embed the Standard Model in some generic more fundamental theory (a "UV completion", or more precisely "UV continuation" since we're not necessarily asking it to be valid up to infinitely small distance scales). However, *the Standard Model on its own does not have a hierarchy problem.* It is a renormalisable quantum field theory. In renormalisable QFTs, while computations frequently involve formally divergent integrals, there exists a well-established procedure to extract perfectly finite physical answers to physical questions from these computations. This is done by *renormalising* the "bare" Lagrangian parameters by introducing counterterms, and arranging the divergences in the counterterms to cancel those in the bare parameters such that any physical observables are finite.

   Counterterms and bare masses themselves are not observable quantities, and the cancellation of their quadratic divergences is an artifact of regularising the theory with a momentum cutoff. If you insist on using a cutoff, you should take it to infinity at the end of your calculation, so the Higgs mass is not hierarchically small with respect to any physical scale of the theory (in fact, the Higgs mass is the *only* fundamental scale in the theory because the Higgs mass parameter is the *only* dimensionful parameter allowed by the symmetries of the Standard Model). There exist renormalisation schemes involving no explicit cutoff and no quadratic divergences.

2. Having to fine-tune the counterterms repeatedly at each order in perturbation theory is an artifact of the perturbation expansion. The 2-loop or n-loop counterterms are no more physical than the bare Higgs mass or the 1-loop counterterms. There is only one independent Higgs mass parameter in the full, all-orders theory.

3. The top quark is in no way special. From the point of view of a stubborn-minded particle theorist who decides to ignore astrophysics and cosmology, nothing is wrong with the Standard Model on its own,[1] and the top quark loop can happily diverge as long as we can absorb this divergence in the parameters of our theory. Which we can because the theory is renormalisable.

Neutrino masses I

Before moving on, let's qualify one of the statements I just made and, in doing so, introduce a few concepts that will come in handy momentarily. I stated that nothing is wrong with the Standard Model, even from a fundamentalist particle physicist's point of view who refuses to acknowledge that what astronomers are doing is also science and that gravity exists. This is not quite right since the Standard Model as a renormalisable QFT does not include neutrino masses, but experiment tells us that these are nonzero.

In detail, SM gauge invariance forbids a neutrino mass term $m_{ij}\nu_{Li}\nu_{Lj}$. (I remind you that in the SM the neutrino forms an $\mathrm{SU}(2)_L$ doublet ℓ_{Li} together with the left-handed charged lepton, $\ell_{Li}=\binom{\nu_{Li}}{l^-_{Li}}$; $i = 1, 2, 3$ is a generation index.) But there is a fairly simple possible remedy for this: Just include three new fermions ν_{Ri} which are uncharged under all SM gauge interactions. Call them right-handed neutrinos, and write down a Yukawa interaction between them, their

[1] The hypercharge sector of the Standard Model does have some UV problems, which manifest themselves at very large energies (it is not well defined as a quantum field theory at the microscopic level). But these are usually ignored since the Standard Model does not include gravity anyway, and coupling it to gravity in a consistent manner raises a lot more problems much earlier. So our stubborn-minded particle theorist would have to ignore gravity before starting to worry about these issues.

left-handed cousins, and the SM Higgs field ϕ:

$$\mathcal{L} \supset y_{ij}\ \phi \bar{\ell}_{Li} \nu_{Rj} + \text{h.c.} \tag{1}$$

Upon replacing ϕ by its vacuum expectation value $v = 246$ GeV after electroweak symmetry breaking, this becomes an effective Dirac neutrino mass $M_{ij} = y_{ij}\, v$. The downside of this idea is that, to reproduce the famously tiny neutrino masses which sum up to less than about an eV, the Yukawa couplings y_{ij} would need to be minuscule, of the order of 10^{-12}. This is therefore a cheap but ugly solution to account for neutrino masses — one remains in a renormalisable setting, but one postulates couplings which are suppressed by twelve orders of magnitude without any explanation. But you might argue that the electron Yukawa coupling is already about 10^{-6}, so who cares about another factor two in the exponent. And if you buy these small couplings, you get yourself a renormalisable model which can explain every particle physics experiment there is. Moreover, I re-emphasize that in this model there is no hierarchy of scales (just a hierarchy of dimensionless Yukawa couplings).[2] This minimal extension of the Standard Model indeed has no hierarchy problem, just like the original version with massless neutrinos.

Now what precisely is the hierarchy problem? I won't be able to answer this question in just one sentence, so first for a bit of context.

Effective field theory

The hierarchy problem arises when you think of the Standard Model as an *effective field theory.* An effective field theory is valid only at energies below a certain energy scale, or equivalently above a certain distance scale. Contributions to physical processes from effects at higher energies, or shorter distances, are absorbed in the couplings of the effective field theory. The canonical example is Fermi's theory of weak interactions. Suppose we were interested in lepton physics at sub-GeV energies. The low-energy interactions of muons and electrons can be described by an effective Lagrangian:

$$\begin{aligned}\mathcal{L} = &-\frac{1}{4} F_{\mu\nu} F^{\mu\nu} + \bar{e}(i\not{D} - m_e)e + \bar{\nu}_e\, i\not{\partial}\, \nu_e + \bar{\mu}(i\not{D} - m_\mu)\mu + \bar{\nu}_\mu\, i\not{\partial}\, \nu_\mu \\ &+ \frac{G_{\text{F}}}{\sqrt{2}} \left(\bar{\nu}_e \gamma^\alpha (1-\gamma^5) e\right) \left(\bar{\nu}_\mu \gamma_\alpha (1-\gamma^5)\mu\right)\end{aligned} \tag{2}$$

The first line is just QED, with the charged states coupled to the photon and the neutrinos hanging around. The second line is a four-fermion interaction which allows the decay $\mu \to e\bar{\nu}_e\nu_\mu$. This operator has mass dimension 6 (recall that a fermion has dimension 3/2 as can be read off from the Dirac term above), therefore the coupling constant G_{F} can be written as $G_{\text{F}} = 1/\Lambda^2_{\text{Fermi}}$, where Λ_{Fermi} is some mass scale. By measuring the muon lifetime, Λ_{Fermi} is found to be around 300 GeV.

Fermi theory is a good description of weak interaction processes at energies $\ll \Lambda_{\text{Fermi}}$ (such as muon decays). At energies of the order of Λ_{Fermi}, it breaks down, because new states can be excited: the W and Z gauge bosons. Diagramatically, at leading order in perturbation theory, in the effective field theory the muon decays because of a local interaction

[2] I should mention that the strange hierarchical patterns of SM Yukawa couplings may themselves serve as a motivation for regarding the SM as incomplete. In fact there are several proposed extensions of the SM where all fundamental parameters are $\mathcal{O}(1)$ and where the Yukawa hierarchies emerge from some dynamical mechanism.

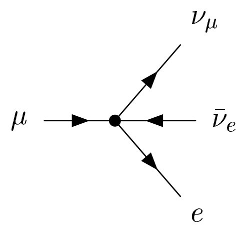

while in the full electroweak theory this decay is mediated by a virtual W boson which is far off shell:

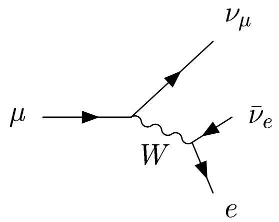

For a process at very low momenta $p \ll \Lambda_{\text{Fermi}}$ the effects of the dimension-6 operator become very small (that's why the weak interactions are called weak), since their matrix elements are suppressed by powers of $p^2/\Lambda^2_{\text{Fermi}}$. For example, for muon decay the characteristic energy scale is $p = m_\mu$, so one finds that the muon decay width is roughly $m_\mu^5/\Lambda^4_{\text{Fermi}}$ up to some numerical constant, with the extra factor of m_μ coming from kinematics.

In fact, even for processes such as $e^+e^- \to \mu^+\mu^-$ at low energies, there is a contribution from a four-fermion term with coefficient G_{F}, coming from Z exchange in the full theory. However, in this case there is another contribution which exists both in the full and in the effective theory, namely, photon exchange. At low energies, it's a pretty good approximation to forget about the four-fermion interaction and to just consider QED:

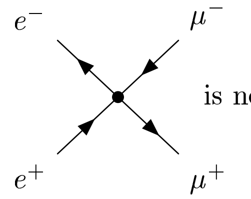

is negligible with respect to

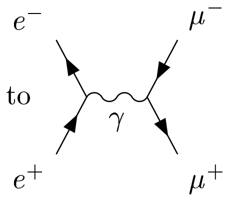

Thus lepton physics at low energies is described by a renormalisable gauge theory, QED, supplemented by some operators of dimension > 4 which become ever more irrelevant in the IR (but may still be important for processes such as muon decay where there is no QED contribution): an effective field theory. For processes whose energies approach or surpass the suppression scale Λ_{Fermi} of these operators, the effective field theory description breaks down and you need to take the W and Z bosons into account. The appropriate UV embedding is nothing but the Standard Model. The couplings of the full SM and of the effective theory are related since both should give the same predictions at low energies: by *matching* the SM to Fermi theory one can determine, say, G_F as a function of the SM parameters.

Neutrino masses II

Now back to the Standard Model and the hierarchy problem. I told you that in its original version neutrinos are massless, which contradicts experiment. Instead of introducing right-handed neutrinos and tiny Yukawa couplings as we did above, let's consider an alternative, arguably more elegant solution of the neutrino mass problem. We add a dimension-5 operator to the Standard Model:

$$\mathcal{L} = \frac{1}{2}\frac{\kappa_{ij}}{\Lambda_\nu}\,\ell_{Li}\ell_{Lj}\phi^\dagger\phi^\dagger \;+\; \text{h.c.} \tag{3}$$

When you replace the Higgs fields by their vacuum expectation value v, you obtain a Majorana-type neutrino mass matrix

$$M_{ij} = \kappa_{ij} \frac{v^2}{\Lambda_\nu} \,. \tag{4}$$

If the couplings in the matrix κ_{ij} are $\mathcal{O}(1)$, then we know from the tiny neutrino masses that $\Lambda_\nu \gtrsim 10^{13}$ GeV **(→Ex.)**. This is a hint that there should be new states with masses of the order of Λ_ν around; their effects at low energies are described by the higher-dimensional operator of Eq. (3). In this picture, there is an explanation for why neutrinos are light: they are light because new physics is heavy. The SM becomes a low-energy effective field theory of a more complete theory which includes the degrees of freedom of that new physics.[3]

Neutrino physics isn't the only place where the SM, despite being self-consistent and all, seems to be lacking. In fact there are all these things which I promised not to go into at the beginning of the lecture: the seemingly arbitrary structure of fermion masses and mixings, resulting from the Yukawa coupling hierarchies; the want for a more common origin of the SM gauge forces; the non-observation of CP violation in the strong interactions... all this cries for an explanation within some more fundamental theory. Some physicists would argue that these issues can be "discussed away" as being mere aesthetic faults, in the same sense that Yukawa couplings of the order of 10^{-12} might seem odd but are perfectly self-consistent. But even ignoring these possible hints for new physics that are built into the structure of the SM itself, there remains an evident conflict with astrophysics and cosmology: Dark matter exists, there is overwhelming evidence for inflation, and the matter-antimatter asymmetry of the universe cannot be accounted for by Standard Model physics. And finally, there is empirical evidence that things fall down when dropped, so gravity should be part of the ultimate picture as well.

For the rest of these lectures, I will therefore assume that new physics at high energies will somehow resolve the above issues. The Standard Model is not to be viewed in isolation but is an effective field theory itself. It ceases to be valid at some scale $\Lambda \gg 100$ GeV. It is this assumption which leads to the hierarchy problem.

Let's investigate how the hierarchy problem arises in a concrete case, using again our neutrino example which should be a bit more familiar by now. For clarity I'll pretend that there is just one fermion generation, and thus only one left-handed neutrino, so that we can drop the flavour indices. All arguments can be easily generalised to the physical case of three generations. We would like to UV-complete our effective theory, defined by the Standard Model plus the dimension-5 operator of Eq. (3), by embedding it into a renormalisable field theory. To generate the operator of Eq. (3) we introduce, again, a new neutral fermion ν_R and couple it to the left-handed neutrino with the Yukawa coupling (1). However, this time we don't tune the Yukawa coupling to be minuscule but assume it is some $\mathcal{O}(1)$ number. Moreover, we also include a Majorana mass term for ν_R — there is no reason not to, if we don't insist on lepton number[4] being a fundamentally conserved quantity. The Lagrangian for ν_R is

$$\mathcal{L} = \bar{\nu}_R i \partial\!\!\!/ \nu_R + \frac{1}{2} M \nu_R \nu_R + \text{h.c.} + y\, \phi \bar{\ell}_L \nu_R + \text{h.c.} \tag{5}$$

Consider the case where M is much larger than the Higgs vacuum expectation value,

$$M \gg v \,. \tag{6}$$

[3]There are several models on the market which flesh out this idea, including new heavy states capable of inducing the effective operator of Eq. (3) at low energies. The simplest model of this kind will be discussed momentarily.

[4]More precisely, baryon minus lepton number in the context of the full Standard Model.

In that case the ν_R will be an approximate mass eigenstate. It won't be an exact mass eigenstate, because the Yukawa coupling y causes ν_R to mix with $\bar{\nu}_L$ after electroweak symmetry breaking. That is, in the $(\bar{\nu}_L,\ \nu_R)$ space the neutrino mass matrix is

$$\mathcal{M}_\nu = \begin{pmatrix} 0 & yv \\ yv & M \end{pmatrix} \tag{7}$$

after replacing $\phi \to v$, and the exact mass eigenstates are the eigenvectors of this matrix. But you can easily convince yourself **(→Ex.)** that ν_R is a mass eigenstate up to small corrections which are suppressed by yv/M and which we will neglect from now on.

At momenta $p \ll M$ the ν_R will decouple from the theory: it will continue to contribute to physical processes as a virtual particle, but there is not enough energy to excite it as a real particle, so it can no longer propagate. Interactions mediated by ν_R will look like local interactions at low energies, much like interactions mediated by W and Z bosons look like local interactions in Fermi theory. Our goal is now to construct an effective field theory for the light states of the theory $(\nu_L, \phi$, and all the other Standard Model fields) which reproduces the complete theory (including ν_R and the interactions of Eq. (5)) at energies far below the ν_R mass scale M. More precisely, it should reproduce all the correlation functions of the light fields in the complete theory, at least to zeroth order in an expansion in p/Λ_ν where p is any external momentum. This is easy to achieve at leading order in perturbation theory. The ν_L and Higgs 2-point functions don't care about ν_R, so they remain as they were with ν_R included. The same is true for the Higgs self-interaction and for the SM Yukawa couplings. However, the full theory had a ℓ_L-ℓ_L-Higgs-Higgs interaction mediated by ν_R:[5]

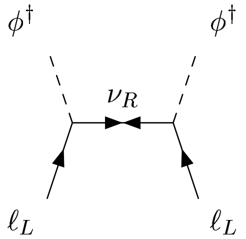

$$= \quad (-iy)\frac{iM}{p^2 - M^2}(-iy) = \frac{iy^2}{M} + \mathcal{O}\left(\frac{p^2}{M^3}\right)$$

So to match this interaction to leading order in p/M we need to introduce a vertex

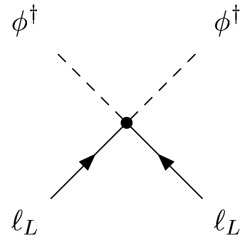

$$= \quad \frac{iy^2}{M} \tag{8}$$

which corresponds precisely to the dimension-5 neutrino mass operator of Eq. (3) with the identifications $\Lambda_\nu = M$ and $\kappa = y^2$,

$$\mathcal{L} = \frac{1}{2}\frac{y^2}{M}\,\ell_L \ell_L \phi^\dagger \phi^\dagger\,. \tag{9}$$

This mechanism of generating small neutrino masses from a heavy right-handed singlet neutrino is called *type-I see-saw mechanism.*

[5] We are using 2-component spinor notation here, which is the most appropriate for a theory of chiral fermions such as the SM. If you have never seen it before, rest assured that this formalism is perfectly equivalent to the 4-component Dirac spinor formalism which you may be more familiar with.

The hierarchy problem, finally

The hierarchy problem becomes apparent when we try to match our low-energy effective field theory to the full theory not just at the tree level but at one loop. Let's take a look at the Higgs two-point function in the complete theory with the right-handed neutrino included. It is given by

$$\int d^4x\, e^{ipx} \langle 0|\mathrm{T}\, \phi^\dagger(x)\phi(0)|0\rangle = \frac{i}{p^2 - m_\phi^2 - \Pi(p^2)} \tag{10}$$

where the *self-energy* $\Pi(p^2)$ of the Higgs field is some function which can be computed in perturbation theory. Expanding it as

$$\Pi(p^2) = \Pi(0) + p^2\Pi'(0) + \ldots \tag{11}$$

we can roughly identify the first term on the RHS with a higher-order correction to the Higgs mass.

At one loop, the self-energy receives a contribution from a neutrino:

$$i\Pi(p^2) =$$

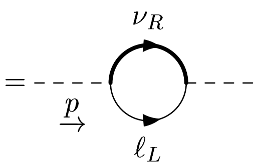

$$\begin{aligned} &= -\,\mathrm{tr}\int \frac{d^4q}{(2\pi)^4}(-iy)\frac{i\,(\not{p}+\not{q})}{(p+q)^2 - M^2}(-iy)\frac{i\,\bar{\not{q}}}{q^2} \\ &= \frac{iy^2}{16\pi^2}p^2\left(\frac{1}{\bar\epsilon} + \frac{1}{2} - \log\frac{M^2}{\mu^2}\right) - \frac{iy^2}{8\pi^2}M^2\left(\frac{1}{\bar\epsilon} + 1 - \log\frac{M^2}{\mu^2}\right) + \mathcal{O}(p^4/M^2)\,. \end{aligned} \tag{12}$$

Here $\not{q} \equiv q_\mu \sigma^\mu_{\alpha\dot\beta}$ and $\bar{\not{q}} \equiv q_\mu \bar\sigma^{\mu\dot\beta\alpha}$, and the σ^μ and $\bar\sigma^\mu$ are 2×2 matrices playing the same role as Dirac's γ matrices for 4-spinors. I've used dimensional regularisation with the shorthand $\frac{1}{\bar\epsilon} = \frac{1}{\epsilon} - \gamma_E + \log 4\pi$. The number of space-time dimensions is $d = 4 - 2\epsilon$, and the renormalisation scale is μ. If you have never calculated a loop diagram before, you might feel a bit lost at this point, but don't worry — the technical details are not important for understanding the main issue.

This result diverges as $d \to 4$, but the divergence doesn't affect any observable quantities. In the low-energy effective theory, there is no corresponding one-loop renormalisation of the Higgs two-point function (the vertex (8) doesn't actually contribute — you might think that you can just tie its fermion legs together to get a one-loop contribution, but the direction of the arrows don't match). There are some additional diagrams involving Standard Model fermions, gauge bosons, and the Higgs itself, but they exist in both the full and in the effective theory, so they cancel out in comparing the two.

We now require that the two-point functions in the full and in the effective theories match at the scale $\mu = M$:

$$\begin{aligned} &Z_{\text{eff}} D_\mu\phi^\dagger D^\mu\phi - \left(m^2_{\phi,\text{eff}} + (\delta m_\phi^2)_{\text{eff}}\right)\phi^\dagger\phi + \ldots \\ \overset{!}{=}\; &Z_{\text{full}} D_\mu\phi^\dagger D^\mu\phi - \left(m_\phi^2 + (\delta m_\phi^2)_{\text{full}}\right)\phi^\dagger\phi + \ldots \end{aligned} \tag{13}$$

To get rid of the infinities we use the $\overline{\text{MS}}$ subtraction scheme, which amounts to dropping the $\frac{1}{\bar\epsilon}$ terms; renormalisation theory tells us that, as long as we do so consistently, this is as good a choice of renormalisation scheme as any other. Comparing with the p^2 term in Eq. (12) we find

$$Z_{\text{full}}^{\overline{\text{MS}},1-\text{loop}}\Big|_{\mu=M} = 1 + \frac{y^2}{32\pi^2} + (\text{some corrections not involving the neutrino sector})\,. \tag{14}$$

Since there is no contribution from the additional vertex in the effective theory, we have

$$Z_{\text{eff}}^{\overline{\text{MS}},1-\text{loop}}\Big|_{\mu=M} = 1 + (\text{the same neutrino-independent corrections as in the full theory})\,. \quad (15)$$

This tells us that the field ϕ in the full theory is related to the field ϕ in the effective theory by a rescaling by $(1+\frac{y^2}{32\pi^2})$. This "wave function renormalisation" translates into a multiplicative correction to the Higgs mass parameter and is not important for illustrating the origin of the hierarchy problem.

What is important instead is the mass renormalisation, or the loop contribution to $\Pi(0)$, for which we find

$$(\delta m_\phi^2)_{\text{full}}^{\overline{\text{MS}},1-\text{loop}}\Big|_{\mu=M} = -\frac{y^2}{8\pi^2}\,M^2 + (\text{neutrino-independent terms}) \quad (16)$$

and

$$(\delta m_\phi^2)_{\text{eff}}^{\overline{\text{MS}},1-\text{loop}}\Big|_{\mu=M} = (\text{the same neutrino-independent terms})\,. \quad (17)$$

To obtain an effective Higgs mass-squared parameter $m_{\phi,\text{eff}}^2$ of the order of $(100\ \text{GeV})^2$, one needs to arrange that the one-loop contribution to (δm_ϕ^2) coming from M^2 cancels out the m_ϕ^2 term to great precision. Note that M is not some mathematical device that we use to regularise the theory; it's a physical parameter of the theory which governs the right-handed neutrino masses. If $M = 10^{13}$ GeV, and $y \sim \mathcal{O}(1)$, then you need a cancellation of about one part in 10^{20}. Clearly the independent parameters m_ϕ and M *can* be fine-tuned such that their difference, with all the prefactors and higher loop corrections accounted for, comes out this tiny. But just as clearly this cries for an explanation. Why *does* Nature choose this peculiar relation between M and m_ϕ? Just to confuse us?[6]

Now, with the help of dimensional regularisation and the above procedure of matching the effective theory to the full theory, we have been able to cleanly identify the source of the hierarchy problem. *The hierarchy problem is due to UV degrees of freedom.* In the case at hand, it manifests itself in the fact that heavy right-handed neutrinos with a mass M tend to induce an effective Higgs mass squared of the order $M^2/(8\pi^2)$, which needs to be cancelled by a similarly large m_ϕ^2 if you want $m_{\phi,\text{eff}}^2$ to be small. The fact that, when using a momentum cutoff Λ, the Higgs mass diverges quadratically as Λ^2 is an indicator that the theory *will* suffer from a hierarchy problem once it's coupled to UV degrees of freedom. But there is actually no need to mention quadratic divergences at all once you specify the UV completion.[7]

The hierarchy problem, which we have discovered with the help of our example model, can be phrased more generally as follows: *When the Standard Model is regarded as an effective field theory, and UV-completed at a scale $\Lambda > \Lambda_{\text{Fermi}}$ into a more fundamental theory including states of mass $\sim \Lambda$ with sizeable couplings to the SM, then the resulting effective electroweak scale is generically $\mathcal{O}(\Lambda)$. A large hierarchy $\Lambda \gg \Lambda_{\text{Fermi}}$ can only be obtained by fine-tuning independent parameters in the UV theory.*

Before introducing the currently most popular attempts for a solution to the problem, supersymmetry and compositeness, let me stress two important facts.

First, the problem is *not* specific to the new physics model being a type-I see-saw model for neutrino masses. It applies to any generic extension of the Standard Model at high scales with

[6]Suggestion to creationists: how about a cult centered around a Mischievous Designer?

[7]Every now and then the claim surfaces that the hierarchy problem is absent in dimensional regularisation, because by construction you never come across any quadratic divergences in dimensional regularisation. From our discussion it should be evident that this is nonsense.

additional heavy fermions, scalars or gauge fields, provided that they are significantly coupled to the Higgs (even if just indirectly through loops).

Second, the problem *is* specific to the Higgs being an elementary scalar field. Indeed you might ask at this point, what about the electron mass in Fermi theory? Doesn't it also receive large radiative corrections from the high-energy degrees of freedom, in this case, from the W and Z bosons and from the Higgs field? Doesn't, therefore, having $m_e = 511$ keV represent a huge fine-tuning when considering that these states have masses around Λ_{Fermi}? The answer to all these questions is no. You can check explicitly that loops do not induce any corrections to m_e proportional to m_W or m_Z or m_ϕ, to any order in perturbation theory that you can be bothered to calculate. Or you can save yourself all that work and instead look for a simple argument why any such corrections must be absent to all orders. One argument goes like this: When setting m_e to zero, the *symmetry of the theory is enhanced* (because in that limit the left-handed and right-handed electrons decouple from each other, which allows one to rotate their phases independently). Since m_e is the only parameter which breaks this enhanced symmetry, all radiative corrections which contribute to m_e have to be proportional to m_e itself. In other words, once a tiny value for m_e is chosen, it will remain tiny under renormalisation. This argument still doesn't provide a reason why the electron is light, but at least it allows you to relegate the question to some UV completion (there are several known mechanisms for generating large fermion mass hierarchies dynamically); afterwards there is no danger of spoiling the smallness of m_e by radiative corrections from heavy mass thresholds. Small fermion masses are said to be *technically natural.*[8] Small scalar masses are not, since $m_\phi = 0$ does not enhance the symmetries of the theory.[9]

Exercises

Suppose that neutrino masses are generated by the type-I see-saw mechanism.

- Calculate the mass eigenvalues and eigenstates of the neutrino mass matrix Eq. (7) to leading order in an expansion in yv/M.
- Convince yourself that, in order to obtain masses $\lesssim \mathcal{O}(\text{eV})$ for the three light neutrinos with $\sim \mathcal{O}(1)$ neutrino Yukawa couplings, the three heavy neutrinos should have masses $\gtrsim 10^{13}$ GeV.

[8] Or "natural in the sense of 't Hooft". The general definition of technical naturalness is just that: A small parameter is technically natural if, in the limit of that parameter going to zero, the symmetry of the theory is enhanced. In that case the smallness of this parameter will be stable under renormalisation.

[9] Or does it? Having no dimensionful parameter in the theory does enhance the *space-time* symmetries of the theory; for instance, the classical action is now invariant under a suitable simultaneous rescaling of all fields and of the space-time metric. Closer inspection reveals that it is invariant under a full set of *conformal transformations*, enhancing the Poincaré symmetry group to the *conformal group*. However, this symmetry is present only at the classical level and broken by quantum anomalies. While there are attempts to make sense of classical conformal symmetry as a symmetry principle to protect the Higgs mass, they are somewhat controversial and we won't discuss them here.