3 Supersymmetry

A super-brief introduction

Having discussed the hierarchy problem, or more specifically its manifestation within our neutrino mass model, let us now turn to what has long been regarded as its leading candidate for a solution: supersymmetry. Supersymmetry is a symmetry relating fermions and bosons. Because fermions and bosons belong to different Poincaré representations, the algebra of generators for supersymmetry (of *supercharges*) is necessarily a nontrivial extension of the Poincaré algebra. These new charges transform as Lorentz spinors, which is a bit peculiar given that we are used to charges which are either Lorentz scalars, generating internal symmetries, or tensors, generating the Poincaré symmetry. But since spinor representations of the Lorentz algebra do exist, why not postulate some spinorial conserved charges and see what comes out?

The central equations of the simplest incarnation of the four-dimensional supersymmetry algebra are

$$\{Q_\alpha, \overline{Q}_{\dot\beta}\} = 2\sigma^\mu_{\alpha\dot\beta} P_\mu\,, \qquad [P_\mu, Q_\alpha] = [P_\mu, \overline{Q}_{\dot\beta}] = \{Q_\alpha, Q_\beta\} = \{\overline{Q}_{\dot\alpha}, \overline{Q}_{\dot\beta}\} = 0\,. \tag{18}$$

Here Q and $\overline{Q}$ are two supercharges transforming as a left-handed and right-handed spinor respectively. Because they are fermionic quantities, they satisfy anticommutation relations rather than commutation relations among themselves. The indices α and $\dot\beta$ are spinor indices and range from 1 to 2, and $\sigma^\mu = (\mathbb{1}_{2\times2}, \sigma^1, \sigma^2, \sigma^3)$ with σ^i the Pauli matrices. Finally, P_μ are the momentum generators of the Poincaré algebra. One could easily spend a few hours exploring Eqs. (18) and their consequences and generalisations in detail (and I encourage you to do so with the help of a good book on supersymmetry). Since we don't have a few hours I'll be very sketchy — let me just emphasize that the first of Eqs.(18) shows how the supercharges are nontrivially entangled with the Poincaré generators.

In a supersymmetric theory the concept of a *field* (which carries some representation of the Poincaré algebra) is extended to that of a *superfield* (which carries some representation of the supersymmetry algebra). In the simplest version of four-dimensional supersymmetry, the simplest superfields contain either a complex Lorentz scalar and a chiral (left-handed or right-handed) fermion,

$$\Phi = \begin{pmatrix} \phi \\ \psi_\alpha \end{pmatrix}, \qquad \text{where } [Q_\alpha, \phi] = \sqrt{2}\,\psi_\alpha\,, \tag{19}$$

or a Lorentz vector and a Majorana fermion,

$$V = \begin{pmatrix} A^\mu \\ \lambda^\alpha \end{pmatrix}, \qquad \text{where } [Q_\alpha, A^\mu] = \sigma^\mu_{\alpha\dot\beta}\bar\lambda^{\dot\beta} \quad \text{(up to a ``supergauge transformation'')}\,. \tag{20}$$

Gauge symmetries commute with supersymmetry.[10] So if ϕ carries some charge under some gauge symmetry, then ψ^α carries the same charge, and if A^μ is a gauge field, then λ^α transforms in the adjoint representation of the gauge group just like A^μ does. Furthermore, supersymmetry dictates that the mass of all states contained in a supermultiplet be the same. In fact this immediately follows from Eqs. (18) **(→Ex.)**. Supersymmetry also relates the quartic scalar $|\phi|^4$ couplings in the theory to the Yukawa couplings between ϕ and ψ^α, and to the gauge couplings relevant to ϕ; finally, it relates the gauge couplings for the gauge field A^μ to the Yukawa couplings involving λ^α. These relations between the dimensionless couplings are less easy to see explicitly.

[10] Unless you count rather exotic theories called "gauged supergravities" which almost certainly have very little to do with this our world.

There is now a very simple prescription for embedding the Standard Model into a supersymmetric theory: Take every Standard Model fermion ψ and complete it into a superfield Φ by adding a scalar *superpartner* ϕ with the same mass and gauge charges. Take every Standard Model gauge boson A_μ and complete it into a superfield V by adding a massless fermionic superpartner λ in the adjoint of the corresponding gauge factor. Take every Standard Model scalar (there is only one, namely the Higgs) and complete it into a superfield by adding its fermionic superpartner. Realize that this last step causes some problems[11] and repair these by adding a second Higgs superfield whose quantum numbers are conjugate to those of the Standard Model Higgs. This construction is known as the *Minimal Supersymmetric Standard Model* or MSSM. Here is its explicit particle content. We follow the usual naming convention of calling the fermionic superpartners of bosons "bos-inos" and the scalar superpartners of fermions "s-fermions", and write all chiral fermions as left-handed spinors.

Gauge fields and gauginos:

spin-	1/2	1	SU(3)	SU(2)	U(1)
	λ_1	B_μ	**1**	**1**	0
	λ_2	W_μ	**1**	**3**	0
	λ_3	G_μ	**8**	**1**	0

(Anti-)quarks, (anti-)leptons, squarks and sleptons:

spin-	0	1/2	supermultiplet	SU(3)	SU(2)	U(1)
	$\tilde{q}_i$	q_{Li}	Q_i	**3**	**2**	1/6
	$\tilde{u}_i^*$	$u_{Ri}^\dagger$	U_i	$\overline{\mathbf{3}}$	**1**	−2/3
	$\tilde{d}_i^*$	$d_{Ri}^\dagger$	D_i	$\overline{\mathbf{3}}$	**1**	1/3
	$\tilde{\ell}_{Li}$	ℓ_{Li}	L_i	**1**	**2**	−1/2
	$\tilde{e}_i^*$	$e_{Ri}^\dagger$	E_i	**1**	**1**	1

Higgs fields and higgsinos:

spin-	0	1/2	supermultiplet	SU(3)	SU(2)	U(1)
	ϕ_u	$\tilde{\phi}_u$	Φ_u	**1**	**2**	1/2
	ϕ_d	$\tilde{\phi}_d$	Φ_d	**1**	**2**	−1/2

It turns out that, once you impose that the Standard Model Yukawa and gauge couplings should be reproduced, this requirement fixes all the dimensionless couplings of the MSSM. The dimensionful parameters, in particular the particle masses, are a different story though.

If supersymmetry were unbroken, it would predict that for each quark there exists a scalar "squark" superparter of the same mass, for each lepton there exists a mass-degenerate "slepton", and for each gauge boson there exists a "gaugino" (the photino and gluino being massless, the wino and zino being mass-degenerate with the W and Z bosons).[12] With unbroken supersymmetry and all these additional light states around our world would look nothing like the world we know. So supersymmetry, if it is realised in nature, should be spontaneously broken in much the same way that the $\mathrm{SU}(2)_L\times\mathrm{U}(1)_Y$ electroweak gauge symmetry is spontaneously broken. In other

[11] Among others, adding a single chiral fermion to the Standard Model induces a gauge anomaly which needs to be cancelled, for example by adding a second, conjugate chiral fermion.

[12] I'm cheating again: In fact with unbroken supersymmetry it's not so easy to break electroweak symmetry, so in fact the W and Z would be massless as would their superpartners.

words, supersymmetry should be a symmetry of the Lagrangian but not of the vacuum state. It turns out that you can parameterise the effects of spontaneous supersymmetry breaking on the MSSM by giving large masses $\sim \Lambda_{\rm SUSY}$ to precisely all the unobserved superpartner states.[13] This provides a nice explanation why supersymmetry hasn't been observed at colliders yet: the superpartners were simply too heavy to be produced, they decouple from the Standard Model at the scale $\Lambda_{\rm SUSY}$. The Standard Model again becomes an effective field theory at low energies.

Supersymmetry and the hierarchy problem

So what does supersymmetry have to do with the hierarchy problem?

It turns out that exact, unbroken supersymmetry imposes strong constraints on the structure of radiative effects. These constraints go under the name of *non-renormalisation theorems.* In particular, *there is no mass renormalisation* in a supersymmetric theory.

To stick with the neutrino example, in a supersymmetric theory we cannot just augment the theory by a right-handed neutrino ν_R, but we also need to include its scalar "sneutrino" superpartner $\tilde{\nu}_R$. This means that there are two sources of radiative corrections from the neutrino sector to the one-loop self energy of the Higgs:

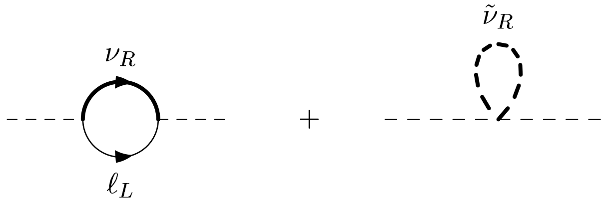

$$= \frac{iy^2}{16\pi^2}p^2\left(\frac{1}{\epsilon}+\frac{1}{2}-\log\frac{M^2}{\mu^2}\right) - \frac{iy^2}{8\pi^2}M^2\left(\frac{1}{\epsilon}+1-\log\frac{M^2}{\mu^2}\right) + \frac{iy^2}{8\pi^2}M^2\left(\frac{1}{\epsilon}+1-\log\frac{M^2}{\mu^2}\right) + \mathcal{O}(p^4) \tag{21}$$

We see that the two terms contributing to mass renormalisation cancel out precisely. This hinges on the quartic scalar coupling entering the second diagram being equal to the square of the Yukawa coupling y^2, and on the sneutrino mass being equal to the neutrino mass M, both of which is ensured by supersymmetry.

On general grounds, the absence of mass renormalisation in supersymmetry can be understood with the help of the symmetry argument which prevents fermions from picking up large radiative corrections to their masses. I argued above that fermion masses are protected by chiral symmetries: the symmetry of a theory is enhanced when the fermion mass is zero, so all corrections to the fermion mass have to be proportional to the symmetry breaking parameter, i.e. the fermion mass itself. Therefore small fermion masses are technically natural. But unbroken supersymmetry constrains the scalar masses to be the same as the masses of the corresponding fermionic superpartners. Thereby the combination of chiral symmetry and supersymmetry protectes also the scalar masses from large radiative corrections.

To the extent that supersymmetry is unbroken, there is therefore no electroweak hierarchy problem. Now I just argued that supersymmetry should be spontaneously broken in Nature. However, that doesn't present a problem as long as supersymmetry is broken at a scale $\Lambda_{\rm SUSY}$ close to the electroweak scale. Sure enough, parameterising the breaking by giving masses of, say, a few 100

[13] The Higgs superfields are an exception here, which is however irrelevant for the argument.

GeV to squarks, sleptons, gauginos and higgsinos will have a significant effect on the resulting electroweak scale. But the contributions from all these states are proportional to their masses, $\mathcal{O}(\Lambda_{\rm SUSY})$ ~(few 100 GeV), and most of them are even suppressed by loop factors. Therefore we would expect them to be not much larger than the electroweak scale itself, and hence to lead to no significant fine-tuning.

The latter argument actually remains true even if some of the superpartners are quite heavy. For example, you can give rather large masses to the squarks and sleptons of the first two generations with impunity, because they are coupled to the Higgs sector only through small Yukawa couplings (at the tree level). Their contributions to the electroweak scale are correspondingly suppressed. Keeping only those states light which have to be light to avoid fine-tuning is sometimes called "natural supersymmetry", although this term is not very sharply defined. Roughly, it means that at least the Higgsinos, gauginos, and third-generation squarks shouldn't be much heavier than a TeV.

The suggestion for a supersymmetric solution to the Hierarchy Problem can be summarised as follows:

- The fundamental Lagrangian is supersymmetric.
- In a supersymmetric theory scalar masses are tied to fermion masses by supersymmetry, and fermion masses are protected from large radiative corrections by a chiral symmetry which becomes exact at zero fermion mass.
- Therefore, in a supersymmetric theory there is no mass renormalisation, and in particular no sensitivity of the electroweak scale to UV-scale physics. To the extent that supersymmetry is exact, we can embed our theory into a see-saw model for generating neutrino masses, a model of grand unification, a superstring theory... at very high scales without generating a hierarchy problem.
- Observation tells us that supersymmetry must, in fact, be spontaneously broken. The breaking is parameterised by supersymmetry breaking masses $\mathcal{O}(\Lambda_{\rm SUSY})$ for the superpartners. However, if $\Lambda_{\rm SUSY}$ is not too far above the Fermi scale $\Lambda_{\rm Fermi}$, this breaking does not entail a hierarchy problem. Roughly speaking, there is no large hierarchy between $\Lambda_{\rm SUSY}$ and $\Lambda_{\rm Fermi}$, and above the scale $\Lambda_{\rm SUSY}$ supersymmetry is restored anyway.
- If the model is supposed not to suffer from some residual fine-tuning, at least the higgsinos, gauginos, stop and sbottom squarks should have masses close to the electroweak scale, and hence be detectable at the LHC or at a next-generation e^+e^- collider; these are the states whose masses have the largest effect on the prediction of the electroweak scale.

At least this was the picture before LHC started looking for superpartners and didn't find any. The absence of any hints for supersymmetry at the LHC implies that $\Lambda_{\rm SUSY}$ is probably larger than many people are comfortable with: there seems to exist at least a *little hierarchy* between $\Lambda_{\rm Fermi}$ and $\Lambda_{\rm SUSY}$, and there is an ongoing debate as to in how far this weakens the case for supersymmetry. We'll get to the experimental situation in more detail later.

Exercises

- Show that the supersymmetry algebra Eq. (18) implies that all states within the same supermultiplet have the same mass.

- Supersymmetry allows for Yukawa terms (scalar-fermion-fermion interactions) of the type $\phi\psi_L\psi'_L + \phi^*\psi_L^\dagger\psi'^\dagger_L$, where ψ_L and ψ'_L are left-handed spinors and ϕ is *the scalar superpartner of a left-handed spinor*. It does *not* allow for Yukawa terms of the type $\phi^*\psi_L\psi'_L$ + h.c.

  Find all Yukawa interactions between the MSSM Higgs and matter fields and their superpartners that are allowed by supersymmetry and by gauge invariance. Among them, identify the subset which contains an even number of superpartners — and which are therefore allowed by the discrete *R-parity* symmetry which acts as (SM field) $\to$ (SM field) and (superpartner) $\to$ ($-$ superpartner). Give an argument why the minimal supersymmetric extension of the Standard Model needs at least two Higgs superfields, independently of anomaly cancellation.

4 Composite Higgs models

Why is there no strong hierarchy problem?

To understand the leading contestor of supersymmetry for a solution of the electroweak hierarchy problem, it helps to get some intuition from QCD. Why is there no "strong hierarchy problem"? After all, QCD at low energies contains a host of massive hadrons with masses around the scale $\Lambda_{\rm QCD} \sim$ a few 100 MeV. Why aren't their masses sensitive to new physics at high energy scales $\Lambda_{\rm UV} \gg \Lambda_{\rm QCD}$?

The answer is quite simple if you look at QCD at sufficiently high energies, where it is just a gauge theory of quarks and gluons. This theory *does not contain any dimensionful parameters*. At high energies there are no scalars in the spectrum, just fermions and gauge bosons; this already is a strong hint that there should be no hierarchy problem, as fermion masses are protected from picking up masses by chiral symmetry and gauge bosons by gauge symmetry. In fact, the only way by which QCD knows about scales, i.e. the only way to distinguish a QCD process at an energy scale Λ_1 from a QCD process at a different energy scale Λ_2, is through the *running gauge coupling*. That is to say that, rather than being a "coupling constant", the strong coupling α_s depends on the energy scale, $\alpha_s = \alpha_s(\Lambda)$. QCD is what's called *asymptotically free*: As $\Lambda \to \infty$, α_s tends to zero. As Λ decreases, α_s becomes larger and ultimately diverges. The scale where α_s becomes non-perturbatively large, and the description of QCD as a theory of weakly coupled quarks and gluons breaks down, is by definition precisely $\Lambda_{\rm QCD}$.

To back up all these words by equations, at one loop in perturbation theory the gauge coupling of an SU(N_c) gauge theory with N_f flavours of massless fundamental fermions evolves according to its renormalisation group equation:

$$\frac{d}{d\log\Lambda}\alpha = \beta^{\text{1-loop}}(\alpha) \tag{22}$$

where

$$\beta^{\text{1-loop}}(\alpha) = \frac{\alpha^2}{2\pi}b, \qquad\qquad b = -\frac{11}{3}N_c + \frac{2}{3}N_f\,. \tag{23}$$

In real-life QCD $N_c = 3$ and, at energies where all quark masses are negligible, $N_f = 6$, so $b = -7$ and $\beta^{\text{1-loop}}(\alpha_s) = -7\alpha_s^2/(2\pi)$. The ODE (22) is solved by **(→Ex.)**

$$\alpha_s^{\text{1-loop}}(\Lambda) = \left(-\frac{b}{2\pi}\log\frac{\Lambda}{\Lambda_{\rm QCD}}\right)^{-1}. \tag{24}$$

Here $\Lambda_{\rm QCD}$ appears as a constant of integration. For asymptotically free theories such as QCD, b is negative, so $\alpha_s(\Lambda)$ is well defined at energies $\Lambda > \Lambda_{\rm QCD}$ where the log is positive. As Λ

decreases and approaches $\Lambda_{\rm QCD}$, the log approaches zero, perturbation theory becomes unreliable, and eventually the gauge coupling becomes formally singular at $\Lambda = \Lambda_{\rm QCD}$. One says that the scale $\Lambda_{\rm QCD}$ is *dynamically generated* or generated through *dimensional transmutation* (exchanging a dimensionless quantity α_s for a dimensionful quantity $\Lambda_{\rm QCD}$). The only way in which physics at high energy scales affects $\Lambda_{\rm QCD}$ is through the beta function Eq. (23): N_f changes according to the number of quarks (or possibly other hypothetical fields charged under QCD) which are effectively massless at the energy scale one is considering. Clearly $\Lambda_{\rm QCD}$ can depend only logarithmically on the scale of new physics where any hypothetical new degrees of freedom decouple. Therefore a dynamically generated scale such as $\Lambda_{\rm QCD}$ is naturally exponentially different from any fundamental scale associated to some UV completion, the running of α_s being only logarithmic. If some grand-unified theory exists at a very large scale, it is sufficient for this theory to induce a moderately small effective α_s in order to give rise to a huge hierarchy of scales. For example, a moderately small value $\alpha_s|_{\Lambda_{\rm UV}} \approx$ (few percent) at a UV scale $\Lambda_{\rm UV} \sim 10^{16}$ GeV results in a strong-coupling scale $\Lambda_{\rm QCD}$ of the right order, some sixteen orders of magnitude below $\Lambda_{\rm UV}$.

In summary, QCD has no hierarchy problem because its scale $\Lambda_{\rm QCD}$ is dynamically generated.

The Higgs as a composite pseudo Nambu-Goldstone boson

Could the electroweak scale $\Lambda_{\rm Fermi}$ also be dynamically generated?

This is exactly the question which composite Higgs models are trying to address. They have some older cousins called *technicolor* theories, trying to do the same without even having a Higgs particle in the spectrum, but especially after the discovery of the Higgs boson, technicolor has become somewhat unfashionable, and I'll not discuss it here. Instead I'll sketch a way to obtain the Standard Model, including the Higgs, as an effective field theory from a more complete theory which, at high energies, possesses no elementary scalars. Much as in QCD, the UV completion is thought to be an asymptotically free gauge theory with some fermions ("technifermions"), and the Higgs boson is thought to correspond to some scalar meson of this theory — a bound state of techniquarks held together by new short-distance interactions. The strong-coupling scale $\Lambda_{\rm CH}$, the equivalent of $\Lambda_{\rm QCD}$, should not to be too far from the electroweak scale $\Lambda_{\rm Fermi}$.

How far is "not too far"? Experimentally, we have seen the Higgs boson, as well as all the Standard Model particles, but we have not seen any hint of Higgs compositeness. QCD comes with all sorts of bound states with masses of the order of $\Lambda_{\rm QCD}$ — where are their equivalents in composite Higgs models? Or is there a reason why some composite state, in this case the Higgs, should be parametrically lighter than the actual compositeness scale $\Lambda_{\rm CH}$? It turns out that there is, and for this we can again draw upon the QCD analogy.

In the SM, the quarks obtain masses from their Yukawa couplings to the Higgs. For two of the quarks, the u and the d, the Yukawa couplings are so small that the resulting masses are $\ll \Lambda_{\rm QCD}$. For the sake of the argument let's pretend that they are actually massless and that the four other quarks are infinitely massive. Let's also forget about the electroweak interactions for the moment. Setting the quark mass to zero enhances the symmetries of the theory. In the case at hand, for QCD with two massless quarks, there is now an $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$ global non-abelian flavour symmetry. Under the $\mathrm{SU}(2)_L$, the left-handed u and d form a doublet, and under $\mathrm{SU}(2)_R$, the right-handed u and d form a doublet. This non-abelian flavour symmetry is spontaneously broken to the diagonal subgroup $\mathrm{SU}(2)_V$ of simultaneous rotation in the spaces of left-handed and right-handed fields. This happens due to the gauge interactions becoming nonperturbatively strong, and the relevant scale is $\sim \Lambda_{\rm QCD}$: QCD exhibits *chiral symmetry breaking*. The surviving subgroup $\mathrm{SU}(2)_V$ is nothing but *isospin* symmetry which you're probably familiar with from nuclear physics. A spontaneously broken global symmetry famously gives rise to one massless Nambu-Goldstone boson (NGB) per broken generator. For $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R \to \mathrm{SU}(2)_V$ breaking the three NGBs can be interpreted as bound states of u and d quarks, and identified with the pions.

Now we know that in real-life QCD u and d quarks aren't actually massless, the other quarks aren't actually infinitely massive, and with electromagnetism there exists at least a remnant of the electroweak interactions which is relevant even at the QCD scale. The pions, therefore, are not exactly massless and degenerate NGBs. But they are *pseudo-Nambu-Goldstone bosons*: The explicit breaking of the $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$ symmetry is a very small effect, as it is only due to the tiny u and d masses of the order of a few MeV and due to the small electromagnetic gauge coupling $\alpha \approx 1/137$. One can estimate

$$m_\pi^2 \approx (m_u + m_d)\frac{\Lambda_{\mathrm{QCD}}^3}{f_\pi^2} \tag{25}$$

in the limit of negligible α (here f_π is the pion decay constant). Electromagnetic effects will further induce a small mass difference between the neutral and the charged pions. The pions are light — indeed they are much lighter than all other mesons — because their masses are proportional to the small $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$ breaking parameters m_u and m_d. If m_u and m_d were zero, the pions would be exactly massless NGBs.

Now back to the Higgs: If we postulate that the *Higgs boson is a pseudo-NGB* of an approximate global symmetry G of our theory of new short-range interactions, this can naturally explain why $m_\phi < \Lambda_{\mathrm{CH}}$ and thus $\Lambda_{\mathrm{Fermi}} < \Lambda_{\mathrm{CH}}$. The Standard Model, once more, becomes an effective theory which is valid at scales below Λ_{CH}. Around the scale Λ_{CH}, usually a few TeV, we expect to start seeing other bound states of the theory which, unlike the Higgs, have no special reason to be light. Once again the small electroweak scale can be understood in terms of a symmetry: as m_ϕ, the quartic Higgs self-coupling λ, the electroweak gauge couplings and the Yukawa couplings all approach zero, the Higgs becomes an exact NGB with a flat potential, and acquires a shift symmetry $\phi \to \phi + c$ which is simply the non-linearly realized global symmetry G. The nonzero gauge and Yukawa couplings break G explicitly, and as a consequence a nonzero Higgs potential will be generated by loop effects.

Pseudo-NGBs are conveniently described by non-renormalisable effective Lagrangians whose structure is however very well understood. Indeed they are the only tool needed to efficiently capture the LHC phenomenology of a pseudo-NGB Higgs, since a compositeness scale Λ_{CH} of a few TeV is quite far away by LHC standards, so an effective theory approach works very well — for practical purposes it's not necessary to construct a full UV model comprising techniquarks charged under some new gauge group, and mimicking QCD quark condensation and confinement. To obtain a viable composite Higgs model, all you need is a pattern of global symmetry breaking, and perhaps an effective description of the lightest few accessible other composite states and their interactions with the pseudo-NGBs (as well as a good deal of confidence that a full theory exists which follows this pattern). Examples for UV completions are known,[14] which is reassuring, but they are not necessarily the best tool for studying the resulting Higgs physics. This, again, is just as in QCD, where talking about quarks and gluons doesn't get you all that far for modelling pion physics.

[14] In fact, recent interest in this field was started by people building models in warped five-dimensional space-times, and then realizing that their weakly coupled 5D theories should have a strongly coupled 4D dual description where the Higgs is a pseudo-NGB, via the AdS/CFT correspondence. These 5D models can be considered as calculable UV realizations of the composite Higggs idea.

A lot of the phenomenology, naturally, depends on what the approximate global symmetry G is taken to be. In the simplest *minimal composite Higgs models* $G \simeq \mathrm{SO}(5)$ which is broken spontaneously through strong dynamics at the scale Λ_{CH} to $H \simeq \mathrm{SO}(4) \simeq \mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$. This construction yields the right number of pseudo-NGBs (**→Ex.**) and has the great advantage of preserving an approximate *custodial symmetry*: the tree-level Higgs potential in the Standard Model is accidentally SO(4)-invariant, and any significant deviation from this would quickly lead to conflict with electroweak precision tests.

Partially composite fermions

There is however something that doesn't easily fit into this picture, and that's the top Yukawa coupling. Experiment tells us that the top is heavy, so its Yukawa coupling must be large. If the Higgs boson is secretly a composite state of technifermions, say $\phi \sim \langle \bar{\psi}' \psi \rangle$, how does one generate the operator

$$\mathcal{L} \supset y_t\, \phi \bar{q}_L t_R \sim \frac{\tilde{y}_t}{\Lambda_t^2}\, \bar{\psi}' \psi \bar{q}_L t_R\ ? \tag{26}$$

This is a four-fermion operator similar to the one we discussed when we were talking about Fermi theory. Thanks to the new strong dynamics we expect the techniquarks to condense. Naively we can replace $\psi' \psi = \Lambda_{\mathrm{CH}}^2 \phi$ in the effective theory, up to some uncalculable $\mathcal{O}(1)$ factors, hence

$$y_t = \tilde{y}_t \left(\frac{\Lambda_{\mathrm{CH}}}{\Lambda_t} \right)^2 . \tag{27}$$

I remind you that Λ_{CH} is the strong-coupling scale of the new strongly interacting gauge group. By contrast Λ_t is the scale at which the four-fermion operator of Eq. (26) is generated. We expect $\Lambda_t > \Lambda_{\mathrm{CH}}$ if the technifermion theory is to make sense as an effective theory of perturbatively coupled technifermions between the scales Λ_{CH} and Λ_t. On the other hand, we know that the numerical value of the top Yukawa coupling at the weak scale is $y_t \approx 1$. Suppose that $\tilde{y}_t \approx 4\pi$, so that the loop expansion parameter $\frac{\tilde{y}_t}{4\pi}$ is just on the brink of becoming too large to trust the perturbation series (in fact we're being very generous here, as the actual loop expansion parameter also contains some group-theoretical factors which tend to work against us). Even this allows only for a separation of the scales Λ_{CH} and Λ_t by a modest factor of $\sqrt{4\pi} \approx 3.5$: there is definitely no parametrically large separation of scales.

But we are talking about strongly interacting theories, while Eq. (27) relies on dimensional analysis with canonical field dimensions, which in turn is valid if the fields are approximately free. Perhaps we should replace Eq. (27) by

$$y_t = \tilde{y}_t \left(\frac{\Lambda_{\mathrm{CH}}}{\Lambda_t} \right)^{2-\gamma} . \tag{28}$$

where the parameter γ reflects the fact that the scaling dimension of an operator in a strongly interacting theory can significantly differ from its classical value? If γ were negative, this would allow for a larger separation of the two scales, and better theoretical control. There are, however, strong arguments that $\gamma > -1$ in general field theories, so the top Yukawa coupling is always suppressed by at least one power of $\frac{\Lambda_{\mathrm{CH}}}{\Lambda_t}$. Moreover, at the same scale Λ_t where we're generating the four-fermion operator of Eq. (26), we'll also generate a four-fermion operator with four *Standard Model* fermions, for example

$$\mathcal{L} \supset \frac{1}{\Lambda_t^2} \bar{q}_i q_j \bar{q}_k q_n . \tag{29}$$

These operators will give rise to *flavour-changing neutral currents*, which are extremely tightly constrained by experiment. For $\mathcal{O}(1)$ couplings and a suppression scale $\Lambda_t \lesssim 10$ TeV, the experimental bounds on flavour-changing neutral currents are generically exceeded by many orders of magnitude in this model. A lot of additional model-building gymnastics is required to evade this conclusion.

The more promising avenue to generate a large top Yukawa coupling is therefore not to use the operator of Eq. (26) at all, and to instead rely on a mechanism called *partial compositeness*. I remind you that our new strongly interacting gauge theory is supposed to contain an approximate global symmetry, $G = \mathrm{SO}(5)$ say, of which the Higgs bosons are the pseudo-NGBs. The proposed "proto-Yukawa coupling" of Eq. (26) explicitly breaks G. What if the third-generation quarks were to instead approximately respect G? For concreteness, let's introduce fermionic composite operators O_L and O_R in some representation of G. They could secretly be made, for instance, of three techniquarks. Let us also write down a G-breaking mixing term between O_L and some elementary fermions which carry the quantum numbers of the ordinary Standard Model quarks:

$$\mathcal{L} = \lambda_L \, \bar{q}_L O_L + \lambda_R \bar{u}_R O_R \,. \tag{30}$$

We now postulate that the physical, observed quarks are mixtures between the composite fermions contained in O_L and the elementary fermions q_L. The third generation is mostly composite (and thus feels the strongest coupling to the Higgs) whereas the first two generations are mostly elementary (and therefore their couplings to the Higgs are weaker, hence their masses are smaller). Concerning in particular the top, Eq. (28) is replaced by

$$y_t \sim \left(\frac{\Lambda_{\mathrm{CH}}}{\Lambda_t}\right)^{\gamma_L+\gamma_R} \tag{31}$$

with $\gamma_{L,R}$ the anomalous dimensions of the O_L and O_R operators. Now the scale Λ_t can be very large without suppressing the top Yukawa coupling, and flavour-changing neutral currents can be accordingly suppressed.

Why am I telling you all these details about partial compositeness? Not just because they are interesting from the model-builder's point of view (they are!) but also because they lead to observable consequences at the LHC. Since we have introduced complete G multiplets, in particular for the third generation, this means that our theory contains additional "top partners" beyond the degrees of freedom of the Standard Model. Given that the Higgs boson is relatively light at 125 GeV, it has been argued that generically at least some of these fermionic top partners should be relatively light as well (sub-TeV), otherwise the model would need to be fine-tuned. I won't go into the whys and hows of this argument here, but it is interesting to see that, just as supersymmetry wants its scalar top partners to be light for naturalness reasons, composite Higgs models want their fermionic top partners to be light for the same reason. Now "light" means in particular "kinematically accessible at the LHC", so this is one of the ways in which a composite Higgs model can be tested.

In summary, according to composite Higgs models:

- There exists a new asymptotically free gauge interaction which becomes strongly coupled at a scale $\Lambda_{\mathrm{CH}} \approx$ few TeV.
- Above the scale Λ_{CH} there are no elementary scalars, hence no hierarchy problem. The scale Λ_{CH} is dynamically generated by dimensional transmutation.
- Below that scale, the Higgs field emerges as a light composite state. It is light because it is the pseudo-NGB of an approximate global symmetry G, broken to H by the new strong gauge interactions, in analogy to chiral symmetry breaking in QCD.

- The Higgs potential is generated by radiative effects.
- The top Yukawa coupling can be generated through the mechanism of partial compositeness, where the third-generation quarks are identified with some composite fermionic operator, up to a small admixture by an elementary fermion field. This leads to fermionic partners for the third-generation quarks, and the LHC is looking for them.

Exercises

- Verify that Eq. (24) solves the renormalisation group equation Eq. (22) for the strong gauge coupling.
- Show that the number of Goldstone bosons obtained from SO(5) $\to$ SO(4) breaking corresponds to the number of real degrees of freedom of the SM Higgs field.

5 LHC phenomenology and constraints

There are essentially two ways to look for physics beyond the Standard Model at high-energy particle colliders. Either you try to produce the new particles directly, and detect their Standard Model decay products (or, noticing that the energy/momentum budget of the final state doesn't add up, you conclude that you have produced a stable particle which has escaped the detector as "missing energy"). Alternatively you can measure some Standard Model process very precisely and infer on the presence of new physics from its virtual corrections to that process. I'll talk about the LHC here, which does both. On the one hand there are the two multi-purpose detectors ATLAS and CMS which are well suited for directly detecting possible new states produced in the beam collisions. On the other hand there is LHCb, which is dedicated to measuring the properties of B mesons in detail — these particles are copiously produced, and since their properties are quite sharply predicted by the Standard Model, any deviation from these predictions would point towards new physics.

Let me focus on direct searches. The LHC as a proton-proton collider is of course very good at producing *coloured particles*. In supersymmetry the particles with the largest production cross sections are therefore typically the scalar superpartners of the first-generation quarks and the fermionic superpartner of the gluino. Furthermore, in most supersymmetric models there is an additional (exact or very weakly broken) discrete $\mathbb{Z}_2$ symmetry known as *R-parity*, under which the superpartners are odd while Standard Model particles are even. This has important experimental consequences: The superpartners can only be produced in pairs, since we start from an even purely Standard Model initial state, so we have to end up with an even final state which necessarily contains an even number of superpartners. R-parity conservation also implies that the lightest supersymmetric particle (LSP) is stable, since there is no R-parity odd state which it could decay into. A typical signature for supersymmetry is therefore generated by the pair-production of squarks and/or gluinos, which decay in cascades into other supersymmetric particles plus Standard Model states, until they end up in the LSP which escapes the detector unseen.[15] The experimental signature is jets (from hadronizing high-energy quarks and gluons, as the originally produced squarks and gluinos need to shed their colour charges somehow) and "missing transverse energy" (from the escaping LSP).

[15] If the LSP is stable, it should be neutral with respect to colour and electromagnetism — otherwise it would completely spoil cosmology. It can however be undergo weak interactions, and might be a good candidate for dark matter.

Here is a sketch of a typical process:

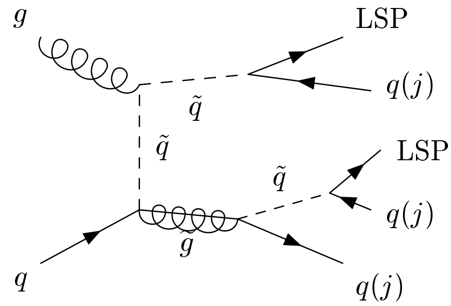

In this example you end up with three jets from the hadronizing quarks, as well as two LSPs, which will only be detected when adding up the total jet momentum transverse to the beam and noticing that some contribution is missing.

Searches for multiple jets and missing transverse energy, possibly with additional leptons, are among the most powerful tools to put constraints on supersymmetric models. If the detector sees more of these processes than you would expect in the Standard Model, in some kinematic domain, then they must be due to new physics. By contrast, if all you see is consistent with the Standard Model, you can put constraints on the production cross sections of the new particles, which translate into constraints on their masses. As the expected signatures are strongly model dependent, so are the constraints you obtain. So a statement along the lines of "gluinos with masses below 2 TeV are ruled out by LHC" must always be interpreted very carefully: Gluinos in the MSSM? With the assumption that R-parity is conserved? Using what further assumptions on the superparticle mass spectrum?

Having said that, I invite you to study the exclusion plots from ATLAS and from CMS and the ensuing mass limits on gluinos and first-generation squarks. I won't go through them here, but those of you who are interested in the details should now have the necessary background to make sense of them. For those of you who aren't, suffice it to say that the present LHC limits at 13 TeV on squark and gluino masses are typically around 2 TeV in optimistic scenarios, depending on the details, but can be weaker in special situations. This is starting to become a serious problem for supersymmetry. The gluino has a large influence on the prediction of the electroweak scale, because it couples strongly to third-generation squarks which in turn couple strongly to the Higgs sector. A difference of more than an order of magnitude between the Z boson mass and the gluino mass is possible to accommodate in the MSSM, but it requires some fine-tuning of independent parameters, and of course our initial motivation for introducing supersymmetry was to obtain a completely natural model. It will be exciting to see if the ongoing run of the LHC eventually brings a gluino to supersymmetry's rescue, but at the moment things don't look very bright.

What about third-generation squarks? We argued before that they should also be relatively light in order to avoid fine-tuning, because their masses have an even stronger influence on the electroweak scale. Unfortunately LHC is less sensitive to stops and sbottoms, because a proton doesn't contain any significant amount of them. You can still produce them with the help of the gluinos, if these are not too heavy,

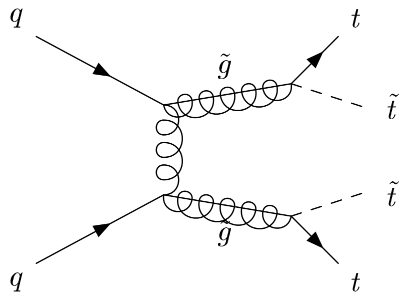

or even directly but at a lower cross section,

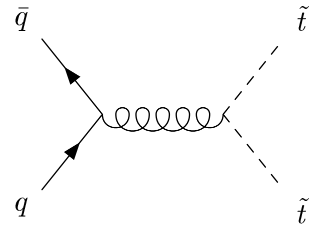

Looking for stops and sbottoms at the LHC is another main line of attack in the experimental programme of supersymmetry searches. A key signature are hard jets and missing energy with bottom quarks in the final state. Now jets originating from a b quark can be "b-tagged" by the detector because they have some characteristic features; most importantly, B mesons are long-lived enough to decay away from the interaction point at a resolvable distance. All in all, even though the LHC is less sensitive to third-generation squarks because of their lower cross section, limits on stop masses can already be established. They vary wildly depending on the assumptions, but range up nearly a TeV by now in the most optimistic settings for direct stop production.

Incidentally, these searches can also lead to interesting limits on models of the second kind we discussed, namely composite Higgs models. I told you that in composite Higgs models with partially composite top quarks, the tops come with light fermionic top partners. These, too, will eventually decay into b-quark jets, so mass limits on stops can be translated into mass limits on composite top partners.

Unfortunately no-one has seen any hints of physics beyond the Standard Model yet (as of July 2017). Every now and then there are vague indications for an anomaly in the LHC data, but so far none of these is significant enough to be taken seriously. Eventually the LHC will have the potential to rule out the "natural" parameter space of both supersymmetry and of composite Higgs models. In fact even today it seems that any concrete such model which is compatible with the present exclusion bounds suffers from some more or less severe fine-tuning of parameters, at least of the order of a percent. Still better than one part in 10^{20}, but not too encouraging given that their main motivation was to avoid fine-tuning altogether. If the ongoing LHC run doesn't see anything beyond the Standard Model, we'll have to rethink our notions of naturalness, and I'll have to rewrite most of these notes…

Conclusions and further reading

Hopefully these lectures got you interested in getting a little deeper into the subject of particle physics beyond the Standard Model. Just to summarise: We have argued that the Standard Model should be thought of as an effective field theory, and that coupling it to UV degrees of freedom will cause the effective electroweak scale to become extremely sensitive to the precise choice of UV parameters. This is called the electroweak hierarchy problem. The hierarchy problem has led us to consider two interesting scenarios, either of which could be realised in Nature. The world could be supersymmetric, or the recently discovered Higgs boson could be a composite state. Both these scenarios have consequences which are in principle testable with present-day experiments, since they predict additional states not too far from the electroweak scale.

What if we don't find any evidence for eiter scenario? Perhaps we are fundamentally misunderstanding something about quantum field theory in general, and when using the right formulation the hierarchy problem ceases to be a problem. Or perhaps the small electroweak scale is environmentally selected (there are attempts to sharpen arguments along the lines of "if there are five gazillion universes around, and the electroweak scale is basically random but needs to be small for a universe to support intelligent observers, then it's no wonder that we happen to find ourselves in a universe where it is indeed small"). Perhaps it is cosmologically selected, a very recent idea whose concrete implementations are however still somewhat unconvincing. Who knows? Unfortunately we may never know — experiment doesn't actually tell us at what scale new physics can be found, if not at the electroweak scale, but the vague hints we do have (neutrino masses...) seem to indicate that this scale may be far beyond our experimental means.

Of all the subjects I have touched I have really just exposed the tip of the iceberg. And of course there is much more literature than I could do justice to, so I won't even try to complile an extensive bibliography.[16] Two exceptions though. If you want to know more about phenomenological supersymmetry, a good starting point (besides various textbooks) is Steve Martin's excellent review `http://arxiv.org/abs/hep-ph/9709356`. If you want to know more about composite Higgs models, try Roberto Contino's lecture notes `http://arxiv.org/abs/1005.4269`, to which chapter 4 of the present text owes a great deal. Finally, the easiest way to get an overview of the present experimental constraints on beyond-the-SM models from direct searches is to visit the `TWiki` pages of the ATLAS and CMS experiments (just google "`cms twiki susy`" for example).

[16] OK, maybe I'm just lazy.