

THE UNIVERSITY
of ADELAIDE

FACULTY OF SCIENCES
SCHOOL OF PHYSICAL SCIENCES

Bayesian and Frequentist Global Fits of
Minimal 4D Composite Higgs Models

Ethan Carragher

A thesis submitted for the degree of
MASTER OF PHILOSOPHY

Supervised by
Prof. Anthony G. Williams and A/Prof. Martin J. White

September 2020

Abstract

Since the famous discovery of the Higgs boson at the Large Hadron Collider (LHC) in 2012, the field of elementary particle physics has been at an impasse. Ever more precise measurements of the particle's properties have served only to verify the predictions of the current prevalent theory - the Standard Model (SM) of particle physics - even though this theory is known to be incomplete. However, the SM is only able to properly describe the observed boson with an incredibly unnatural fine-tuning of its parameters, which has led many to propose alternative models that can more naturally accommodate the particle in the hope of establishing a more complete theory. We concern ourselves in this work with one class of such models, in which the Higgs boson is not an elementary particle as is assumed by the SM, but rather a bound state of some as-yet undiscovered strong dynamics. We give a pedagogical introduction to the theory of such composite Higgs models (CHMs), and provide a complete description of three different versions of the Two-Site Minimal 4D CHM - the simplest calculable extensions of the SM in which the Higgs boson is composite, based on the $SO(5) \rightarrow SO(4)$ symmetry breaking pattern - that differ in their fermion sector embeddings. Convergent global fits are performed on these three models, under both frequentist and Bayesian frameworks, in order to find the regions of their parameter spaces that best fit a wide range of constraints, including recent Higgs measurements and exclusion bounds on heavy resonance production from Run II of the LHC. We use a novel technique to analyse the fine-tuning of the models, quantifying the tuning as the Kullback-Leibler divergence from the prior to the posterior probability on the parameter space. Each model is found to be able to satisfy all constraints at the 3σ level simultaneously, but the model that has fermions embedded in the fundamental representation of $SO(5)$, despite suffering from a “double tuning”, is clearly favoured by a Bayesian model comparison. As a by-product of the fits, we analyse the collider phenomenology of our models in these viable regions. We find clear predictions of the minimally-tuned models that the $gg \rightarrow H \rightarrow \gamma\gamma$ cross section is less than $\sim 90\%$ that predicted by the SM, which is already in slight tension with experiment and could potentially be ruled out in the future high-luminosity run of the LHC. In addition, the lightest fermions F arising from the new strong dynamics in these models are seen in general to lie between ~ 1.1 TeV and ~ 3.0 TeV, with the $F \rightarrow tW^+$ and $F \rightarrow \bar{b}W^+$ decays offering particularly promising channels for probing these models in future collider searches.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Ethan Nicholas Carragher

Acknowledgements

The fact that this thesis exists at all is a testament to the enormous support I have received throughout my years. To all who have helped, in ways big and small, noticed and unnoticed, I am thankful.

First and foremost, I am deeply indebted to my supervisors Anthony Williams and Martin White for their guidance throughout this project. I could not ask for a better pairing: Tony, whose sharp insights manage to cut to the heart of every problem, and Martin, who is an endless source of ideas and optimism.

I am also extremely grateful to Daniel Murnane for drawing me in to this project in the first place, and for spending the last months of his PhD helping me navigate the treacherous waters of composite Higgs theory. His collaboration was instrumental in getting the global fits up and running. I would like to extend my gratitude to Wei Su, who spent many nights collecting all of the recent LHC search analyses used in the fits, for his hard work on the project.

My sincere thanks to Peter Stangl for sharing his remarkable code with us that made our fits possible. His consultations were also very helpful and are much appreciated. Similarly, Will Handley has been invaluable in this work, with his great help in troubleshooting and interpreting our `PolyChord` results, and with his generous donation of computing resources to run some of our scans. Many thanks.

Thanks also to the Adelaide physics faculty, especially Silvana Santucci and Sharon Johnson, for keeping everything functioning smoothly, and to Fabien Voisin for his charitable (and tireless) assistance in getting our jobs to run on the Phoenix and Ronin computer clusters.

These years have been made more enjoyable than I could have hoped for by the companionship of Tajwar Tahabub, Albert Kong, Liam Hockley, and other friends at the University. I am not sure how I would have fared without them, and count myself lucky to have them in my life.

And finally, I cannot thank my family enough - Mum, Dad, and Daniel - for their endless support through all of my endeavours.

Contents

1	Introduction	1
I	Composite Higgs Models: Motivation and Theory	5
2	Preliminaries: Higgs Boson in the Standard Model	7
2.1	Content of the Standard Model	7
2.2	Symmetries	9
2.3	Higgs Mechanism	13
2.4	Electroweak Symmetry Breaking in the SM	14
2.5	Hierarchy Problem	18
3	Key Ideas of Composite Higgs Models	23
3.1	Spontaneously Broken Symmetry	25
3.2	Goldstone's Theorem	25
3.2.1	Example: the Linear σ -Model	27
3.2.2	Example: the Non-Linear σ -Model	31
3.3	Pseudo-Nambu-Goldstone Bosons	32
3.3.1	Example: Pions in QCD	33
3.4	CCWZ Construction	35
3.4.1	NGB Transformations Under Unbroken Symmetries	37
3.4.2	NGB-Gauge Interactions	38
3.4.3	NGB-Fermion Interactions	41
3.5	Vacuum Misalignment	42
3.5.1	EWSB in the Minimal Composite Higgs Model	46
3.6	Partial Fermion Compositeness	49
3.6.1	Composite Fermion Representations in the MCHM	53
3.6.2	$SO(4)$ Invariants	58
3.7	Coleman-Weinberg Potential	60
3.7.1	Higgs Potential in the MCHM	65
3.8	Multi-Site Models and Collective Breaking	69
3.8.1	Spontaneous Breaking to the Diagonal Subgroup	69
3.8.2	Hidden Local Symmetry	71
3.8.3	Multi-Site Formalism	75
3.8.4	Collective Breaking	82

4	The Two-Site Minimal 4D Composite Higgs Model	85
4.1	Symmetry Structure	85
4.2	Boson Sector	86
4.3	Fermion Sector	89
4.3.1	M4DCHM ⁵⁻⁵⁻⁵	90
4.3.2	M4DCHM ¹⁴⁻¹⁴⁻¹⁰	91
4.3.3	M4DCHM ¹⁴⁻¹⁻¹⁰	92
4.3.4	Elementary Fermions	93
4.4	Calculability	93
4.5	Integrating out Heavy Resonances	94
4.6	Particle Content	96
4.6.1	Boson Sector	96
4.6.2	Fermion sector	99
4.7	Status of Model Explorations	100
II	Global Fits of the Minimal Composite Higgs Models	103
5	Bayesian and Frequentist Global Fits	105
5.1	Bayesian Statistics	105
5.1.1	Nested Sampling	108
5.2	Frequentist Fits	110
5.2.1	Differential Evolution	111
6	Global Fit Procedures	115
6.1	Scan Parameters	115
6.2	Calculation of Observables	119
6.3	Constraints	121
7	Results	131
7.1	M4DCHM ⁵⁻⁵⁻⁵	132
7.1.1	Bayesian Fit	132
7.1.2	Frequentist Fit	141
7.2	M4DCHM ¹⁴⁻¹⁴⁻¹⁰	145
7.2.1	Bayesian Fit	145
7.2.2	Frequentist Fit	151
7.3	M4DCHM ¹⁴⁻¹⁻¹⁰	155
7.3.1	Bayesian Fit	155
7.3.2	Frequentist Fit	161
7.4	Discussion and Model Comparisons	164
7.5	Experimental Signatures	170
8	Conclusion	183

Appendices	
A Representation Theory	187
A.1 Groups and Lie Groups	187
A.2 Representations	188
B Charges and Generators of a Symmetry	195
C Quantum Effective Action	197
C.1 Interpretation of the Effective Potential	199
D Isomorphism Between $SO(4)$ and $SU(2) \times SU(2)$	203
E Custodial Symmetry	205
E.1 Custodial Symmetry in the SM	206
F Computation of the Nambu-Goldstone Boson Matrix	209
G $SO(5)$ Generators and Multiplets	211
G.1 Generators	211
G.2 Field Embeddings	211
H Mass Matrices in the M4DCHM	213
H.1 Boson Sector	213
H.2 Fermion Sector	215
H.2.1 M4DCHM ⁵⁻⁵⁻⁵	215
H.2.2 M4DCHM ¹⁴⁻¹⁴⁻¹⁰	217
H.2.3 M4DCHM ¹⁴⁻¹⁻¹⁰	219
H.2.4 Calculation of the Top and Bottom Masses	220
I Fermion Form Factors	223
I.1 M4DCHM ⁵⁻⁵⁻⁵	223
I.2 M4DCHM ¹⁴⁻¹⁴⁻¹⁰	224
I.3 M4DCHM ¹⁴⁻¹⁻¹⁰	225
J Collider Search Constraints	227
Bibliography	233

Introduction

It is a remarkable fact that all known material things appear to consist of a small variety of elementary particles. Perhaps more remarkable is that a single theory, the **Standard Model** (SM) of particle physics, successfully describes all known elementary particles and their electroweak (EW) and strong nuclear interactions, encompassing almost all known particle behaviour, to high precision. Formulated in the mid 1970s, the SM is a patchwork of the Weinberg-Salam theory of electroweak interactions, which by then had gained wide acceptance due to the discovery of weak neutral currents [1, 2], and quantum chromodynamics (QCD), the theory of strong interactions based on the quark model put forward to explain the patterns of observed hadrons [3, 4]. In the years since, SM predictions have been consistently verified, with early successes in the discoveries of the W^\pm and Z bosons that mediate the weak interaction. The SM posits that these particles have masses endowed through the Higgs mechanism, a by-product of which is the existence of a neutral scalar particle first explicitly noted by Higgs [5, 6], but also following the independent work of many others [7–9], dubbed the **Higgs boson**. The search for the Higgs boson was a major experimental undertaking that drove the construction of the Large Hadron Collider (LHC), culminating recently in the famous discovery of a 125 GeV Higgs-like particle [10, 11]. At present this particle is consistent with the SM Higgs boson, though more precise measurements of its couplings may reveal physics **Beyond the SM** (BSM).

The existence of the Higgs boson was by no means a foregone conclusion; its incorporation into the SM was simply the most economical explanation of the weak boson masses. Were it not to have been found, theoretical considerations practically guaranteed the LHC would have discovered some sort of BSM physics in its place. This is because the amplitude for the scattering of two longitudinally polarised W^\pm or Z bosons grows with energy, spoiling perturbative unitarity above ~ 3 TeV unless

- 1) some extra degrees of freedom, including a Higgs boson, are present to restore unitarity, and/or
- 2) perturbativity fails at these energies, and the theory becomes strongly coupled.

Nature seems to have chosen the first option. Early competitors to the SM such as Higgsless Technicolor, a theory that solved the unitarity problem through strongly coupled dynamics analogous to QCD, subsequently had to be abandoned. See Ref. [12] for a nice comparison of the SM and Technicolor paradigms.

This victory of the SM may, however, prove Pyrrhic. In a similar fashion to above, the existence of a 125 GeV Higgs boson can be used to argue for some type of new physics that supersedes the SM at a scale below around 10 TeV. In this case, the argument stems from naturalness: if the SM were accurate upwards of such energies then the Higgs mass would be expected to be many orders of magnitude larger than 125 GeV, and only with an overwhelming fine-tuning of the model parameters could such a light Higgs boson - and therefore any structure in the universe - exist. Many therefore view the SM as a deeply unsatisfying description of reality, calling this the **Hierarchy Problem**. Being metaphysical in nature, the Hierarchy Problem is one of the few theoretical issues present in the SM.

Of course the SM has its experimental limitations too: the theory does not include gravitational interactions, nor does it explain the matter-antimatter asymmetry, dark matter, or the oscillations of neutrino flavours. Understanding BSM physics is therefore a major focus of current research. But because the SM so accurately predicts particle behaviour at accessible energies, aside from neutrino mixings and some tensions with lepton flavour universality [13], there are no experimental leads as to how the SM is to be extended. For this reason, many prominent theories of BSM physics draw inspiration from the Hierarchy Problem and centre around mechanisms that would explain how the light Higgs boson arises naturally.

One particular class of such theories consists of models in which the Higgs is realised not as an elementary particle, but as a bound state of some currently undiscovered BSM dynamics. In such **Composite Higgs Models** (CHMs), the Higgs emerges as a pseudo-Nambu-Goldstone boson of some spontaneous symmetry breaking structure, making its expected mass significantly lighter than the BSM energy scale. For dynamics below around 10 TeV, this provides an adequate solution to the Hierarchy Problem. Many particular realisations of CHMs exist, and given the relatively low energy scales of their new physics, there is the exciting prospect that the LHC will soon be able to find evidence for, or rule out, certain models.

We concern ourselves in this thesis with the **Minimal CHM** (MCHM): the minimal extension to the SM in which the Higgs boson is composite. Our goal is to perform a convergent global fit of the MCHM, finding the regions of its parameter space that best fit a wide range of experimental constraints. Actually, the MCHM refers to a variety of models based on the $SO(5) \rightarrow SO(4)$ symmetry breaking pattern that differ in their BSM field content. We will specifically be fitting the Two-Site Minimal 4D CHM, which possesses the minimal structure necessary for calculability, using three different choices of $SO(5)$ representations for the BSM fields.

Although the subject of much theoretical work, these models have only been numerically explored to a small extent, and no such convergent global fits have yet been performed. This is primarily on account of the large parameter spaces of the models, making exhaustive parameter scans prohibitively computationally expensive, along with the fact that the parameters all have highly non-linear effects on both the SM and BSM sectors, so that even simplified analyses of the parameter spaces are dif-

difficult. Indeed for our fits we only consider limiting models in which out of the SM fermions only the third generation quarks couple to the BSM sector, so that we may have more manageable parameter spaces. It is common in the literature to use this same approximation, or even to have only the top quark couple to the BSM sector, as lighter fermions are not expected to have significant BSM interactions. Despite this simplification, our models still contain 15+ free parameters.

We will be fitting the models under both **Bayesian** and **frequentist** statistical frameworks. The frequentist fits aim to find those points that best fit the experimental constraints across the entire parameter space. To explore the spaces efficiently in these fits we employ a genetic optimisation algorithm known as **differential evolution**, which has proven particularly effective in optimising over high-dimensional spaces compared to a range of other algorithms [14], and which heuristically is quite suited to the particular case of CHMs. The Bayesian fits, on the other hand, optimise the fitness *weighted by the naturalness* of the regions in parameter space - a notion more in line with the philosophy of CHMs. For these fits we use a **nested sampling** algorithm that cleverly facilitates calculation of the Bayesian evidence, a quantity that measures the overall fitness-weighted-by-naturalness of each model, during exploration of the parameter spaces. As a by-product of the fits, the experimental signatures of the models we consider can be analysed to determine the prospects of probing the MCHM in future collider searches.

This thesis is structured in two parts. The first reviews the motivation and theory behind CHMs, with a special focus on the MCHM, in Chapters 2 to 4. Along with Appendices A to G, these chapters provide all the background necessary to understand and appreciate the structure of the models that we will be fitting, assuming a working knowledge of quantum field theory. The second part details the methodology of our scans and the results of the fits, in Chapters 5 to 8.

We begin with an overview of the SM Higgs boson in Chapter 2, not just because it is essential to understand the physics that CHMs intend to supersede, but also because the SM serves as a good introduction to many concepts that arise also in CHMs. Included here is a review of the particles of the SM and their interactions, and a discussion of the role symmetries play in physical theories - much of which should be familiar to the reader, but is covered anyway as such concepts are vital throughout the rest of this work. Afterwards we give a general explanation of the Higgs mechanism and show its enactment in the SM as an example. To conclude, there is a more in-depth discussion of the Hierarchy Problem and the various approaches taken to solve it, with a focus on CHMs.

We go on to present the main concepts utilised by general CHMs in Chapter 3, pedagogically working up from the barest consistency requirements of such theories. After describing (pseudo-)Nambu-Goldstone bosons, we illustrate how such particles may be included in a general theory using the CCWZ formalism. We then show how the weak bosons acquire their masses in CHMs through vacuum misalignment, and how fermions couple to the BSM sector in the partial compositeness paradigm. We also explain how the Higgs potential in a CHM arises from

the Coleman-Weinberg mechanism, and using this we analyse the naturalness of different $SO(5)$ representations in the MCHM. Finally, we discuss how the Higgs potential can be rendered finite and calculable with multi-site models and collective symmetry breaking. Throughout the chapter we illustrate most of these concepts with toy models, using examples that will prove helpful in understanding the MCHM where possible.

Chapter 4 is dedicated to providing a complete working description of the realistic versions of the MCHM that we will be fitting, using the theoretical tools of Chapter 3. We include some analysis of the particle content in each model, and discuss previous explorations of these models in the literature. The different statistical viewpoints - Bayesian and frequentist - are explained in Chapter 5 along with descriptions of the scanning algorithms we use, and our procedures for scanning and constraining the models' parameter spaces are detailed in Chapter 6. Results are presented in Chapter 7, with some discussion on the viability of each model in light of recent LHC data, as well as the prospects of probing the models in future collider searches. We conclude in Chapter 8, summarising the main points of the thesis and offering ideas for how this work may be extended in the future.

Part I

Composite Higgs Models: Motivation and Theory

Preliminaries: Higgs Boson in the Standard Model

2.1 Content of the Standard Model

The Standard Model is a quantum field theory whose Lagrangian (density) has the schematic form

$$\begin{aligned}
\mathcal{L}_{\text{SM}} = & -\frac{1}{4}\text{Tr} [G_{\mu\nu}G^{\mu\nu}] - \frac{1}{4}\text{Tr} [W_{\mu\nu}W^{\mu\nu}] - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} \\
& + i\bar{\psi}^f (\not{\partial} - ig\mathcal{W}^a T_L^a - ig' \not{B}Y - ig_s \not{G}^a T_c^a) \psi^f \\
& - y_{f_1 f_2} \bar{\psi}_L^{f_1} \Phi \psi_R^{f_2} - y_{f_1 f_2}^* \bar{\psi}_R^{f_1} \Phi^\dagger \psi_L^{f_2} \\
& + |(\partial_\mu - igW_\mu^a T_L^a - ig' B_\mu Y) \Phi|^2 + \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2. \quad (2.1)
\end{aligned}$$

The various parts of this Lagrangian will be explained in the following sections of this chapter. This section serves to catalogue the particles of the SM and outline their overall roles in the theory.

First, central to the SM is the symmetry group

$$SU(3)_c \times SU(2)_L \times U(1)_Y. \quad (2.2)$$

Each factor here is a group of transformations that can be performed on the SM fields such that \mathcal{L}_{SM} remains unchanged. Fields are categorised according to the *representations* of these groups under which they transform. A brief overview of the representation theory needed for this work is given in Appendix A.

Present in the SM are fermions, ψ^f , whose interactions are mediated by vector gauge bosons associated to each factor of the symmetry group. Among the gauge bosons are eight gluons G_μ of $SU(3)_c$, three W_μ fields of $SU(2)_L$, and one field, B_μ , of $U(1)_Y$. Each group gives rise to a type of “charge” for each particle - colour charge, weak isospin, and hypercharge, in the order of Equation (2.2) - that dictates their interactions with the associated gauge bosons. Particles in the trivial $\mathbf{1}$ representation of a group are “uncharged” and do not couple to the associated gauge bosons.

	Gen 1	Gen 2	Gen 3
Up-type quarks	u	c	t
Down-type quarks	d	s	b
Charged leptons	e	μ	τ
Neutrinos	ν_e	ν_μ	ν_τ

Table 2.1: Members of each fermion generation in the SM.

All fermions in the SM are Dirac fermions, with left- and right-handed components¹. Weak isospin only affects the left-handed fermions, organising them into doublets consisting of either a charged lepton and its associated neutrino, or of an up-type (electric charge $+2/3$) and a down-type (electric charge $-1/3$) quark. These doublets partition the quarks and leptons into three different “generations”, as listed in Table 2.1. Similarly, each quark flavour is organised into a triplet of its different colour charges.

The final member of the SM is the Higgs field, Φ . It possesses a non-zero vacuum expectation value so that interactions with the Higgs field yield mass terms in the Lagrangian. This so-called Higgs mechanism, explained in Section 2.3, mixes the $W_\mu^{1,2,3}$ and B_μ gauge fields into the massive W^\pm and Z bosons, which mediate the weak interaction, and into the photon, the massless mediator of electromagnetism. The SM particle content and transformation properties are summarised in Tables 2.2 and 2.3.

ψ^f	$SU(3)_c$	$SU(2)_L$	$U(1)_Y$
ℓ_R^i	1	1	-1
$l_L^i = (\nu_L^i, \ell_L^i)^\top$	1	2	$-1/2$
$q_L^i = (u_L^i, d_L^i)^\top$	3	2	$+1/6$
$u_R^i = (u_r^i, u_g^i, u_b^i)^\top$	3	1	$+2/3$
$d_R^i = (d_r^i, d_g^i, d_b^i)^\top$	3	1	$-1/3$

Table 2.2: Representations of SM fermions. Up-type and down-type quarks are labelled u and d , while ℓ refers to a charged lepton. Generations are indexed by i .

	$SU(3)_c$	$SU(2)_L$	$U(1)_Y$
G_μ	8	1	0
W_μ	1	3	0
B_μ	1	1	0
Φ	1	2	$+1/2$

Table 2.3: Representations of SM bosons.

¹Except the neutrinos, which only come left-handed.

2.2 Symmetries

In the development of the Standard Model, it became apparent that one of the most fruitful ways in which to view physics is through the lens of symmetries: transformations of physical systems that do not affect the outcome of any measurement. Nature has many symmetries, both exact and approximate, which severely constrain possible physical behaviours. Often, a physical symmetry refers to the stronger condition of a transformation S of fields ϕ that leaves the Lagrangian invariant, rather than just outcomes of measurements:

$$\phi \rightarrow S\phi \quad \text{such that} \quad \mathcal{L}(S\phi, \dots) = \mathcal{L}(\phi, \dots). \quad (2.3)$$

The SM Lagrangian is almost uniquely determined by specification of its symmetry properties and field content, showing the utility of the symmetry viewpoint.

The symmetries of a given Lagrangian form a mathematical group $G_{\mathcal{L}}$. All $G_{\mathcal{L}}$ transformations g of a quantum field ϕ can be imparted by linear operators S_g^ϕ chosen in such a way that the linear operators respect the group structure:

$$\phi \xrightarrow{g \in G_{\mathcal{L}}} S_g^\phi \phi \quad \text{such that} \quad S_g^\phi S_{g'}^\phi = S_{gg'}^\phi. \quad (2.4)$$

The collection $\{S_g^\phi \mid g \in G_{\mathcal{L}}\}$ forms a **representation** of the symmetry group. Different fields may transform under different representations. A brief overview of common representations is given in Appendix A. We shall only be interested in finite-dimensional representations of compact groups² $G_{\mathcal{L}}$, in which case the operators in the representation can always be taken to be unitary [16]. The distinction between group elements and operators is usually unimportant, so in what follows group elements will be regarded as their corresponding operators.

Perhaps the most useful aspect of symmetries is Nöther's theorem, guaranteeing that for each differentiable symmetry of the Lagrangian there is a current J^μ that is conserved whenever the equations of motion are satisfied, in the sense that $\partial_\mu J^\mu = 0$. From the divergence theorem, it then follows that the quantity, or “charge”,

$$Q(V, t) := \int_V d^3x J^0(\vec{x}, t) \quad (2.5)$$

is locally conserved. Examples of conserved quantities that arise in this way are energy, momentum, and electric charge.

A differentiable symmetry as in Nöther's theorem is formalised by the notion of a Lie group: a group with a manifold structure such that the group multiplication

²If $G_{\mathcal{L}}$ is not of the form $M \times \mathbb{R}^n$ for some compact Lie group M and some $n \geq 0$, kinetic terms for the gauge fields arising from $G_{\mathcal{L}}$ could not be properly defined [15]. We do not consider the case $n \neq 0$ because this would lead to long-range forces for which there is no experimental evidence (assuming no symmetry breaking), leaving only compact groups.

and inversion maps

$$(g_1, g_2) \mapsto g_1 g_2, \quad (2.6)$$

$$g \mapsto g^{-1}, \quad (2.7)$$

are smooth. Specifically, a differentiable symmetry is to be understood as a symmetry group of the Lagrangian containing all transformations in some neighbourhood of the identity of some Lie group G . For simplicity, we consider G itself to be this symmetry group, and remind that we take the case where G is compact. An infinitesimal transformation $S_\varepsilon \in G$ in a neighbourhood of the identity may be written in the form³

$$S_\varepsilon = \mathbb{1} + i\varepsilon^a T_a = e^{i\varepsilon^a T_a}$$

for some infinitesimal real parameters ε^a and some linear operators T_a , equal in number to the dimension of G . These operators, called “generators” of G , form a basis for the tangent space \mathfrak{g} of G at the identity. Note that since S_ε is taken to be unitary, each generator must be Hermitian. The tangent space \mathfrak{g} , when equipped with the commutator $[\cdot, \cdot]$, is known as the **Lie algebra** of G , and the structure constants $f_{ab}{}^c$ defined by

$$[T_a, T_b] =: i f_{ab}{}^c T_c \quad (2.8)$$

entirely specify the Lie algebra up to isomorphism. It is an important fact, demonstrated in Appendix B, that the generator T_a is nothing more than the conserved charge Q_a associated to the symmetry it generates.

Repeated application of infinitesimal symmetries will eventually produce finite symmetries of the Lagrangian that are connected to the identity, so for simplicity we also restrict to the case where G is connected. It is a basic result in differential geometry that the exponential map then extends from an infinitesimal neighbourhood of the identity to the whole group:

$$G = \{e^{i\theta^a T_a} \mid \vec{\theta} \in \mathbb{R}^{\dim G}\}. \quad (2.9)$$

Our attention is therefore restricted to connected, compact Lie groups under whose action the Lagrangian is invariant. The *modus operandi* of BSM physics research is to consider a Lagrangian with a (well-motivated) connected, compact Lie symmetry group, and then to determine if the theory corresponds to reality. Composite Higgs models are no different. For our purposes, we must also make the distinction between two different types of symmetries: global and local.

³We use the Einstein summation convention throughout this work.

A **global** symmetry is a symmetry transformation that does *not* depend on the spacetime point:

$$S_{global}(x) = e^{i\theta^a T_a}. \quad (2.10)$$

A **local** symmetry is a symmetry transformation that *does* depend on the spacetime point:

$$S_{local}(x) = e^{i\theta^a(x) T_a}. \quad (2.11)$$

Nöther's theorem applies for all differentiable global symmetries. In the case where the Lagrangian is invariant under all local transformations from some Lie group G , it is said to be **gauge invariant** and the Lie group is a **gauge group**.

At first glance, global symmetries may seem to be special cases of gauge symmetries, but in fact the two are entirely different. A global symmetry that does not extend to a gauge symmetry is a true symmetry of Nature: transforming the fields will produce a physically different configuration of fields that results in the same physical behaviour. With gauge symmetries, however, since physical behaviour is unchanged under *arbitrary local transformations* of the fields, any quantity affected by a gauge transformation must be unphysical, serving only as a trivial degree of freedom of the fields that has no effect whatsoever. It is often useful to set the fields' gauge configurations in some way when performing calculations in a theory with gauge symmetry, in a process called fixing the gauge. With that said, a gauge symmetry does indeed contain global symmetries as a subgroup, and so enjoys a conserved current from Nöther's theorem.

Gauging a Lagrangian

Often we are interested in Lagrangians that are gauge invariant under some group G . The easiest way to construct such a Lagrangian is to start with a Lagrangian with a global symmetry group G and altering it to accommodate local transformations, known as “gauging” the Lagrangian. This requires changing only the derivative terms, since other terms are insensitive to how the fields vary across spacetime. Quite generally, gauging a Lagrangian only modifies the existing terms by replacing partial derivatives with covariant derivatives:

$$\partial_\mu \rightarrow D_\mu := \partial_\mu - igA_\mu^a T_a. \quad (2.12)$$

Here we have introduced a coupling constant g and a vector gauge field A_μ^a for each generator T_a of G . Each simple group factor of G requires its own coupling constant [17], although here we have only shown one for simplicity. It is convenient to define $A_\mu := A_\mu^a T_a$, which by definition transforms under G as

$$A_\mu \xrightarrow{S \in G} SA_\mu S^{-1} - \frac{i}{g}(\partial_\mu S)S^{-1}, \quad (2.13)$$

so that under a gauge transformation $\phi \rightarrow S\phi$, covariant derivatives transform as

$$D_\mu \phi \rightarrow S D_\mu \phi, \text{ or formally } D_\mu \xrightarrow{S \in G} S D_\mu S^{-1}. \quad (2.14)$$

This is exactly the same transformation property as partial derivatives under a global transformation, so promoting partial derivatives to covariant derivatives makes a globally-invariant Lagrangian gauge invariant. Note if only a subgroup H of a global symmetry G is gauged, then not all elements of G may remain symmetries. In general the symmetry is explicitly broken to the **normaliser subgroup** of H in G , defined by $G_H := \{S \in G \mid ShS^{-1} \in H \text{ for all } h \in H\}$, for this is the necessary and sufficient property to preserve the covariant derivative transformation Equation (2.14) [18].

Kinetic terms for the gauge fields must also be introduced, and this is done through the **field strength tensor**

$$A_{\mu\nu} := \frac{i}{g} [D_\mu, D_\nu] = \partial_\mu A_\nu - \partial_\nu A_\mu + ig[A_\mu, A_\nu]. \quad (2.15)$$

The defining expression makes it clear that this transforms as $A_{\mu\nu} \xrightarrow{S \in G} S A_{\mu\nu} S^{-1}$, *i.e.* in the adjoint representation of G , and so $\text{Tr}[A_{\mu\nu} A^{\mu\nu}]$ is gauge invariant. It is often useful to choose generators such that $\text{Tr}[T_a T_b] = T \delta_{ab}$ for some number T , in which case the conventionally-normalised gauge-invariant kinetic term is

$$-\frac{1}{4T} \text{Tr}[A_{\mu\nu} A^{\mu\nu}]. \quad (2.16)$$

In this work we will always choose $T = 1$.

A similar gauge-invariant term in which gauge fields may appear is

$$\frac{1}{2} \varepsilon^{\mu\nu\rho\sigma} \text{Tr}[A_{\mu\nu} A_{\rho\sigma}] = \varepsilon^{\mu\nu\rho\sigma} \partial_\mu \text{Tr} \left[A_\nu A_{\rho\sigma} - \frac{2}{3} ig A_\nu A_\rho A_\sigma \right]. \quad (2.17)$$

However, being a four-divergence, this term does not affect the classical equations of motion. We will neglect such terms in this work, although they can produce non-perturbative physical effects.

This exhausts the list of possible ways a gauge field may enter a renormalisable Lagrangian. In total, then, gauging a Lagrangian is the process

$$\mathcal{L}_{\text{global}}(\phi, \partial\phi) \rightarrow \mathcal{L}_{\text{gauge}}(\phi, \partial\phi, A_\mu, \partial A_\mu) = \mathcal{L}_{\text{global}}(\phi, D\phi) - \frac{1}{4} \text{Tr}[A_{\mu\nu} A^{\mu\nu}]. \quad (2.18)$$

This whole business of considering gauge-invariant Lagrangians comes about because interacting massless vector particles such as the photon and gluons are necessarily gauge bosons [19]. Indeed, Maxwell's equations lead directly to the gauge invariance of electromagnetism. Gauge theories are also the only renormalisable theories of massive vector particles [20], and indeed the massive W^\pm and Z bosons of the weak interaction are understood to be gauge bosons with masses acquired through a technicality known as the Higgs mechanism.

2.3 Higgs Mechanism

Careful inspection of a gauged Lagrangian reveals that there are no gauge field mass terms $m_{ab}^2 A_\mu^{a*} A^{b\mu}$, except in the case where the covariant derivative acts on a field with a non-zero constant term, known as a Higgs field. To be clear, a Higgs field Φ is one that can be expressed as

$$\Phi(x) = \Phi_0 + H(x), \quad (2.19)$$

where Φ_0 is a constant non-zero operator, and H is a linear combination of normal-ordered products of creation and annihilation operators. Assuming the Higgs field is a multiplet of complex scalars (as it is in the SM), its kinetic term will contain the gauge boson mass terms

$$(D_\mu \Phi)^\dagger (D^\mu \Phi) \supset g_{(a)} g_{(b)} (T_a \Phi_0)^\dagger T_b \Phi_0 A_\mu^{a*} A^{b\mu}, \quad (2.20)$$

where the coupling $g_{(x)}$ is that of whatever simple group T_x is associated to. Rotating to the gauge field mass basis, *i.e.* diagonalising the mass matrix

$$m_{ab}^2 \equiv g_{(a)} g_{(b)} (T_a \Phi_0)^\dagger T_b \Phi_0, \quad (2.21)$$

reveals there are $\dim(K)$ -many massive gauge bosons (including antiparticles) and $\dim(G) - \dim(K)$ massless ones, where $K := \text{span}_{\mathbb{R}}\{T_1 \Phi_0, \dots, T_{\dim(G)} \Phi_0\}$. This method of mass generation is known as the Higgs mechanism.

One might now wonder how a Higgs field may come about in a theory. To start with, note that the constant Φ_0 is the expectation value of the Higgs field in the vacuum state $|\Omega\rangle$:

$$\langle \Omega | \Phi | \Omega \rangle = \langle \Omega | \Phi_0 | \Omega \rangle + \langle \Omega | \cancel{H} | \Omega \rangle \overset{0}{=} \Phi_0. \quad (2.22)$$

So a Higgs field is one that the Lagrangian dictates has a non-zero vacuum expectation value (**vev**). Incidentally, because the vacuum is Lorentz invariant, this observation implies that any Higgs field must be a scalar field. Excitations of Higgs fields are accordingly called Higgs bosons. Note there is no requirement that a Higgs boson be elementary.

The vev of a field is a configuration that locally minimises the **quantum effective potential** V_{eff} , a quantity discussed in detail in Appendix C. In short, writing the (free field) Lagrangian as the difference between kinetic and potential terms,

$$\mathcal{L}(\Phi, \partial\Phi) = \mathcal{T}(\Phi, \partial\Phi) - \mathcal{V}(\Phi), \quad (2.23)$$

the quantum effective potential is the classical potential \mathcal{V} with effects of loop diagrams taken into account:

$$V_{\text{eff}}(\Phi) = \mathcal{V}(\Phi) + [\text{terms next to leading order in } \hbar]. \quad (2.24)$$

More details about the radiative corrections will be given in Section 3.7. Including a Higgs field in a theory, then, simply requires setting up the classical potential and interactions in such a way as to give the quantum effective potential a minimum for a non-vanishing field configuration Φ_0 .

The effective potential acts in much the same way the classical potential acts in classical field theory. For example, expanding the effective potential around the vacuum Φ_0 yields the Higgs boson mass matrix

$$V_{\text{eff}}(\Phi) = V_{\text{eff}}(\Phi_0) + H^a \frac{\partial V_{\text{eff}}}{\partial \Phi^a} \Big|_{\Phi=\Phi_0} + \frac{1}{2} H^a H^b \frac{\partial^2 V_{\text{eff}}}{\partial \Phi^b \partial \Phi^a} \Big|_{\Phi=\Phi_0} + \mathcal{O}(H^3), \quad (2.25)$$

$$\implies m_{H_{ab}}^2 = \frac{\partial^2 V_{\text{eff}}}{\partial \Phi^b \partial \Phi^a} \Big|_{\Phi=\Phi_0}. \quad (2.26)$$

Diagonalising this gives the squared masses of any Higgs bosons, assuming the Higgs fields H are conventionally normalised. Note that since Φ_0 is a minimum of V_{eff} , the mass matrix is positive semidefinite, as should be physically expected.

2.4 Electroweak Symmetry Breaking in the SM

Here we detail a concrete realisation of the Higgs mechanism: the unified electroweak model devised by Glashow, Salam, Ward, and Weinberg [21–23]. This process, misleadingly referred to as **electroweak symmetry breaking** (EWSB), is how the weak gauge bosons acquire mass in the SM, and has proved extraordinarily successful as a description of reality.

In this model, the Lagrangian has an $SU(2)_L \times U(1)_Y$ gauge symmetry, giving $\dim SU(2) + \dim U(1) = 3 + 1$ gauge bosons. Those associated to $SU(2)_L$ are denoted W_μ^a , while the gauge field of $U(1)_Y$ is denoted B_μ . A Higgs field Φ exists that provides mass to three of the gauge bosons, identified with the W^\pm and Z of the SM, leaving a massless photon. In total, the Lagrangian has the form

$$\mathcal{L}_{\text{EW}} = -\frac{1}{4} \text{Tr} [W_{\mu\nu} W^{\mu\nu}] - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} + (D_\mu \Phi)^\dagger (D_\mu \Phi) - \mathcal{V}(\Phi). \quad (2.27)$$

For the correct phenomenology, the Higgs field is taken as a complex doublet⁴

$$\Phi = \begin{pmatrix} \phi_2 + i\phi_1 \\ \phi_4 - i\phi_3 \end{pmatrix}, \quad (2.28)$$

⁴The strange numbering and sign conventions are chosen so that $(\phi_1, \phi_2, \phi_3, \phi_4)$ transforms in the fundamental representation of $SO(4)$, using the local group isomorphisms presented in Appendix D.

that transforms as a doublet under $SU(2)_L$ and has hypercharge $Y = +1/2$. The covariant derivative is accordingly

$$D_\mu \Phi = \left(\partial_\mu - igW_\mu^a T_L^a - ig'B_\mu \frac{1}{2} \right) \Phi, \quad (2.29)$$

where the generators $T_L^a = \frac{1}{2}\sigma^a$ of $SU(2)_L$ are (half) the Pauli sigma matrices

$$\sigma^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma^3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (2.30)$$

Note the different coupling constants for each group factor.

As outlined in the previous section, the Higgs potential⁵ $\mathcal{V}(\Phi)$ must be gauge invariant and have a minimum at $\Phi_0 \neq 0$. This is realised by taking the most general renormalisable, gauge-invariant potential possible:

$$\mathcal{V}(\Phi) := -\mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2 \quad (2.31)$$

for some positive numbers μ^2 and λ . A lower-dimensional version of this potential is graphed in Figure 2.1. The potential has a sphere of minima given by

$$\Phi_0^\dagger \Phi_0 = \phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2 = \frac{\mu^2}{2\lambda}. \quad (2.32)$$

Which minimum is the correct one? The answer is immaterial. Any point on the sphere of minima is physically equivalent to any other, being connected by an $SU(2)_L \times U(1)_Y$ gauge transformation.




Figure 2.1: A plot of the Higgs potential in the $\phi_3 - \phi_4$ plane for $\phi_1 = \phi_2 = 0$ and some values of the parameters μ^2 and λ . Note the circle of minima, representing possible equivalent vacuum states. The rotational symmetry in this plane is a lower-dimensional manifestation of the $SU(2)_L \times U(1)_Y$ symmetry of the potential.

⁵For clarity in illustration, we leave aside the quantum contributions to the effective potential, although these radiative corrections will play a key role in composite Higgs theories.

At this stage it becomes convenient to spend the gauge symmetry to move to the **unitary gauge**, removing all the unphysical gauge degrees of freedom by fixing ϕ_4 to be the only non-zero component of Φ . In this gauge the vev is

$$\Phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ \sqrt{\frac{\mu^2}{\lambda}} \end{pmatrix} \equiv \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}, \quad (2.33)$$

where v is taken positive. Experimentally, $v \approx 246$ GeV.

With the given generators and vev, the gauge boson mass matrix Equation (2.21) is calculated to be

$$m_{\text{gauge}}^2 = \frac{1}{8} \left(\begin{array}{c|cccc} & W_\mu^1 & W_\mu^2 & W_\mu^3 & B_\mu \\ \hline W^{1\mu} & g^2 v^2 & -ig^2 v^2 & 0 & 0 \\ W^{2\mu} & ig^2 v^2 & g^2 v^2 & 0 & 0 \\ W^{3\mu} & 0 & 0 & g^2 v^2 & -g' g v^2 \\ B^\mu & 0 & 0 & -g' g v^2 & g'^2 v^2 \end{array} \right). \quad (2.34)$$

Diagonalising this yields

$$m_{\text{gauge}}^2 = \frac{1}{4} \left(\begin{array}{c|ccc} & W_\mu^+ & Z_\mu & A_\mu \\ \hline W^{-\mu} & g^2 v^2 & 0 & 0 \\ Z^\mu & 0 & (g^2 + g'^2) v^2 & 0 \\ A^\mu & 0 & 0 & 0 \end{array} \right), \quad (2.35)$$

in the basis

$$W_\mu^\pm = \frac{W_\mu^1 \mp iW_\mu^2}{\sqrt{2}}, \quad Z_\mu = \frac{gW_\mu^3 - g'B_\mu}{\sqrt{g^2 + g'^2}}, \quad A_\mu = \frac{g'W_\mu^3 + gB_\mu}{\sqrt{g^2 + g'^2}}. \quad (2.36)$$

As the notation suggests, these are the familiar W^\pm and Z bosons, and the massless photon. The weak bosons are seen to have masses

$$m_W = \frac{1}{2} g v, \quad m_Z = \frac{1}{2} \sqrt{g^2 + g'^2} v \quad \implies \quad \frac{m_W^2}{m_Z^2} = \frac{g^2}{g^2 + g'^2}. \quad (2.37)$$

This prediction has been verified to better than 1% accuracy, and would need to be reproduced in any viable model of EWSB. It is shown in Appendix E that Equation (2.37) is a consequence of a ‘‘custodial’’ $SU(2)$ symmetry of the W^a gauge bosons [24], a concept that will prove vital later on in this work.

Incidentally, if the covariant derivative is expressed in this mass basis, one sees the photon enters as

$$D_\mu \supset -i \frac{g'g}{\sqrt{g^2 + g'^2}} A_\mu (T_L^3 + Y). \quad (2.38)$$

The electromagnetic gauge coupling e and the electric charge operator Q are therefore identified as

$$e = \frac{g'g}{\sqrt{g^2 + g'^2}}, \quad Q = T_L^3 + Y, \quad (2.39)$$

as should be familiar from particle physics. Notice that $Q\Phi_0 = 0$. The unitary gauge is designed for this to conspire, ensuring the photon is massless. A different choice for the basis of $SU(2)_L$ generators would require a different form for the unitary gauge.

In the unitary gauge, Φ has only a single degree of freedom: the physical Higgs boson field, entering as

$$\Phi(x) = \Phi_0 + H(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}. \quad (2.40)$$

The other three degrees of freedom Φ started with manifest as the longitudinal polarisations of the massive gauge bosons. The Higgs boson is seen to be uncharged, with interactions

$$\begin{aligned} \mathcal{L}_h = & \frac{1}{2}(\partial_\mu h)(\partial^\mu h) + \left(m_W^2 W^{\mu+} W_\mu^- + \frac{1}{2} m_Z^2 Z^\mu Z_\mu \right) \left(1 + \frac{h}{v} \right)^2 \\ & - \mu^2 h^2 - \frac{\mu^2}{v} h^3 - \frac{\mu^2}{4v^2} h^4. \end{aligned} \quad (2.41)$$

Of the many interesting things to note here, we draw attention to the bare Higgs mass

$$m_h = \sqrt{2}\mu. \quad (2.42)$$

This mass is a free parameter and must be fixed by experiment.

Fermion Masses

In addition to giving the weak bosons mass, the SM Higgs field also provides masses to the fermions. It does this through so-called **Yukawa interactions** such as

$$\mathcal{L}_{\text{SM}} \supset -y_d^i (\bar{q}_L^i \Phi d_R^i + \bar{d}_R^i \Phi^\dagger q_L^i), \quad (2.43)$$

which by an incredible coincidence are invariant under the SM gauge group given the fermion representations listed in Table 2.2. Here i indexes the fermion generation, and the dimensionless Yukawa couplings y_d^i have been taken to be real and to only act generation-by-generation, by a redefinition of the quark fields. Demonstrating the gauge invariance explicitly, we perform an arbitrary $SU(2)_L$ transformation S and a $U(1)_Y$ phase transformation parameterised by α on the first term of Equation (2.43)

(the other term acts similarly):

$$\begin{aligned}\bar{q}_L^i \Phi d_R^i &= (q_L^i)^\dagger \gamma^0 \Phi d_R^i \rightarrow (e^{+\frac{\alpha}{6}} S q_L^i)^\dagger \gamma^0 (e^{+\frac{\alpha}{2}} S \Phi) (e^{-\frac{\alpha}{3}} d_R^i) \\ &= (e^{-\frac{\alpha}{6}} e^{+\frac{\alpha}{2}} e^{-\frac{\alpha}{3}}) \bar{q}_L^i (S^\dagger S) \Phi d_R^i \\ &= \bar{q}_L^i \Phi d_R^i,\end{aligned}\tag{2.44}$$

so Yukawa interactions are indeed unchanged under $SU(2)_L \times U(1)_Y$. Expanding these interactions in the unitary gauge produces

$$\mathcal{L}_{\text{SM}} \supset -y_d^i (\bar{d}_L^i d_R^i + \bar{d}_R^i d_L^i) \frac{(v+h)}{\sqrt{2}} = -\frac{y_d^i v}{\sqrt{2}} \left(\bar{d}^i d^i + \frac{1}{v} h \bar{d}^i d^i \right),\tag{2.45}$$

from which we see mass terms for the down-type quarks

$$m_{d^i} = \frac{y_d^i v}{\sqrt{2}}.\tag{2.46}$$

Replacing q_L^i and d_R^i with l_L^i and ℓ_R^i gives analogous mass terms for the charged leptons, with their own Yukawa couplings. The up-type quarks receive masses the same way, with a minor technical difference. Through a peculiarity of the $\mathbf{2}$ representation of $SU(2)$, the ‘‘conjugate’’ Higgs doublet defined by

$$\Phi^c := i\sigma^2 \Phi^* = \begin{pmatrix} \phi_4 + i\phi_3 \\ -\phi_2 + i\phi_1 \end{pmatrix}\tag{2.47}$$

is also a doublet of $SU(2)_L$, but has hypercharge $Y = -1/2$. This allows gauge-invariant Yukawa terms for up-type quarks to be written:

$$\mathcal{L}_{\text{SM}} \supset -y_u^i (\bar{q}_L^i \Phi^c u_R^i + \bar{u}_R^i \Phi^{\dagger} q_L^i),\tag{2.48}$$

yielding masses analogous to Equation (2.46).

Yukawa couplings are inputs to the theory, so none of the fermion masses are predictable. However, a clear consequence of this method of mass generation is that a fermion’s coupling to the Higgs boson is directly proportional to its mass. In practice, only the top and bottom quarks and the tau lepton are massive enough to have appreciable couplings to the Higgs.

2.5 Hierarchy Problem

‘‘Have you ever heard about the Higgs boson blues?’’

- Nick Cave & The Bad Seeds

While the SM Higgs boson is an economical solution for the generation of masses for the gauge bosons and fermions, its couplings present theoretical difficulties. Consider, for example, the fermion loop diagram in Figure 2.2. This contributes to the




Figure 2.2: Loop diagram contributing to the renormalised Higgs mass.

renormalised Higgs mass

$$m_{\text{phys}}^2 = m_{\text{bare}}^2 + \Delta m^2 \quad (2.49)$$

an amount

$$\begin{aligned} \Delta m_\psi^2 &= iN_c \left(-i \frac{y_\psi}{\sqrt{2}} \right)^2 \int \frac{d^4 k}{(2\pi)^4} \frac{\text{Tr} [(\not{p} + \not{k} + m_\psi)(\not{k} + m_\psi)]}{((p+k)^2 - m_\psi^2 - i\varepsilon)(k^2 - m_\psi^2 - i\varepsilon)} \Big|_{p^2=m_H^2} \\ &= -\frac{3y_\psi^2}{8\pi^2} \Lambda^2 + \mathcal{O} \left(m_\psi^2 \log \frac{\Lambda}{m_\psi} \right), \end{aligned} \quad (2.50)$$

where Λ is the high-energy cutoff of the loop momentum - the energy scale at which the SM is no longer valid. We see there is a quadratic dependence on Λ . Similar loops with gauge and Higgs bosons also contribute, but the dominant contribution comes from the top quark due to its large Yukawa coupling. If Λ is much larger than, say, 10 TeV, the Higgs mass will then receive corrections that are enormous in comparison to its actual value $m_{\text{phys}} \approx 125$ GeV. In fact it gets worse, for the Higgs mass is UV sensitive: if there is *any* new physics that couples to the SM at a scale Λ_{UV} , then Δm^2 will scale as Λ_{UV}^2 [25–27]. Barring any remarkable cancellations, the Higgs should therefore be expected to have a mass on the order of Λ_{UV} . The fact that it has instead a relatively small mass is known as the **Hierarchy Problem**.

From an effective field theory point of view, the problem stems from scalar mass operators having a mass dimension of two, and so needing coupling constants scaling quadratically with the high-energy cutoff scale, say $c\Lambda^2$ for a dimensionless constant c . Typically c would be expected to be of order unity, raising the question of why the scalar Higgs particle has a mass so much smaller than Λ .

Another viewpoint is through the sensitivity of the Higgs mass to the parameters of the theory. The correction Δm^2 to the Higgs mass will be of the form $-c\Lambda_{\text{UV}}^2$ for some c of order unity and some extremely large energy scale Λ_{UV} . Supposing, for example, the UV energy is at the Planck scale ($\Lambda_{\text{UV}} \sim 10^{19}$ GeV), the physical Higgs mass is $m_{\text{phys}} \sim 10^{-17} \Lambda_{\text{UV}}$, making the bare mass

$$m_{\text{bare}}^2 \sim (c + 10^{-34}) \Lambda_{\text{UV}}^2. \quad (2.51)$$

Since the bare mass is a free parameter of the theory, such a value is perfectly allowed. However, many physicists are not satisfied by this, as the dimensionless parameter $m_{\text{bare}}^2/\Lambda_{\text{UV}}^2$ must be so precisely specified (to the $\sim 34^{\text{th}}$ decimal place!) to result in the observed Higgs mass, while there is no apparent symmetry that would make it so. It is clear that if the ratio varied even relatively slightly, the physical Higgs mass would be stupendously large. Such **fine-tuning** of the parameters goes against the concept of **naturalness**: the idea that the parameters of the theory describing our universe should have values that are more or less typical for such a theory.

If naturalness is to be upheld, it is clear there needs to be some new physics below around 10 TeV that protects the Higgs mass from large quantum corrections. Resolving the Hierarchy Problem is the driving force behind a large portion of current BSM physics research, with a diverse range of approaches having been put forward as solutions [28].

- Supersymmetry (SUSY) is perhaps the most well-known approach. It posits a spacetime symmetry between bosons and fermions, assuring a one-to-one matching between bosonic and fermionic degrees of freedom. The bosonic and fermionic contributions to the Higgs mass then mostly cancel each other out, resulting in only small Higgs mass corrections. While an attractive solution, no evidence for SUSY partner particles has yet been found. Bounds on the partner masses from the LHC are restricting the parameter space of the minimal SUSY model to quite unnatural regions, but this is not to say the minimal model can be ruled out by fine-tuning considerations just yet. A recent analysis has found some versions of the minimal SUSY models to be fine-tuned to $\mathcal{O}(1)\%$ precision in the model parameters [29], which is a slightly higher degree of tuning than the models we shall be considering in this work [30].
- Other approaches utilise higher dimensions, framing the Hierarchy Problem in terms of the apparent weakness of gravity (enormity of the Planck scale) in comparison to the gauge interactions (at the scale of the Higgs mass). For example, it may be that gravity is actually comparable in strength to the SM forces, and only appears extremely weak because there are other millimetre-length dimensions in which it can propagate [31]. A different solution is presented by Randall-Sundrum type models [32], which posit a finite interval extra dimension. The metric of usual 4D spacetime is warped according to the coordinate along the fifth dimension, leading to an exponential suppression of energy scales along the new dimension. Large hierarchies can then arise rather naturally depending on the locations in the fifth dimension where energy scales are embedded. In essence, in these models the Higgs *is* fundamentally at the Planck scale, but because of our position in the fifth dimension the Higgs mass is seen to be very suppressed.
- The final approach we discuss is the focus of this work: the idea that the Higgs is *not* an elementary scalar, but is instead a bound state arising from

some new strong dynamics at a scale m_* around a few TeV. In this case, the loop integrals contributing to the Higgs mass cannot sensibly be applied for momenta much above m_* ; above this scale, the constituent particles of the Higgs are able to be resolved and the Higgs as an entity becomes transparent to virtual particles, as illustrated in Figure 2.3. The Higgs mass should then be expected to be of order m_* , greatly reducing the hierarchy between the expected mass and physical mass. Of course m_* cannot be too large for fear of reinstating the Hierarchy Problem, nor can it be too small on account of the new strong dynamics having so far gone undiscovered.




Figure 2.3: Interaction between a composite Higgs boson and a virtual particle. The composite Higgs will have a finite diameter $\sim 1/m_*$, where m_* is the scale of the composite dynamics. For momenta greater than $\sim m_*$, the virtual particle will have a wavelength short enough to resolve the constituent particles of the Higgs boson.

Early variants of this idea, labelled under the umbrella of “Technicolor”, realised the Higgs boson as a bound state of two fermions from some strong dynamics that essentially mirrored QCD at a higher energy scale⁶ [25, 33–36]. However, such models were typically mired by issues with reproducing the fermion masses and flavour physics, and could not naturally account for the lack of particles with masses similar to the Higgs mass [28, 37].

Technicolor has since been supplanted by **composite Higgs models** (CHMs), which employ a number of more sophisticated theoretical tools for improved phenomenological viability. Having early roots in Refs. [38–40], CHMs did not garner much interest until after the discovery of the AdS/CFT correspondence [41], when it was realised that a composite Higgs boson appears in the dual 4D descriptions of the warped 5D models discussed above (corresponding to the fifth component of a 5D gauge field) [42–47]. This correspondence has spurred much development of CHMs, with natural features of the 5D models being used to guide the structure of 4D CHMs that will be detailed in Chapter 3.

⁶Very early variants of Technicolor did not even include a Higgs boson and instead relied on the strong dynamics to produce a condensate that breaks electroweak symmetry. Given the discovery of the Higgs particle, though, such Higgsless models are obviously ruled out.

Key Ideas of Composite Higgs Models

If the Higgs boson is to arise as a bound state of dynamics at an energy scale m_* , then generically there should be other composite resonances of typical mass m_* , constituting the so-called **composite sector**. As previously mentioned, the compositeness scale m_* must not be too large in order to avoid the Hierarchy Problem, while viable models must push the scale up to $\gtrsim 1$ TeV to explain the absence of composite resonances in high-energy collider experiments. The picture envisioned by composite Higgs models, then, is something as in Figure 3.1: there is a composite sector consisting of heavy resonances of mass $\sim m_* \gtrsim 1$ TeV along with the Higgs at a much lower mass scale, together with the **elementary sector** of the non-Higgs SM fields¹.

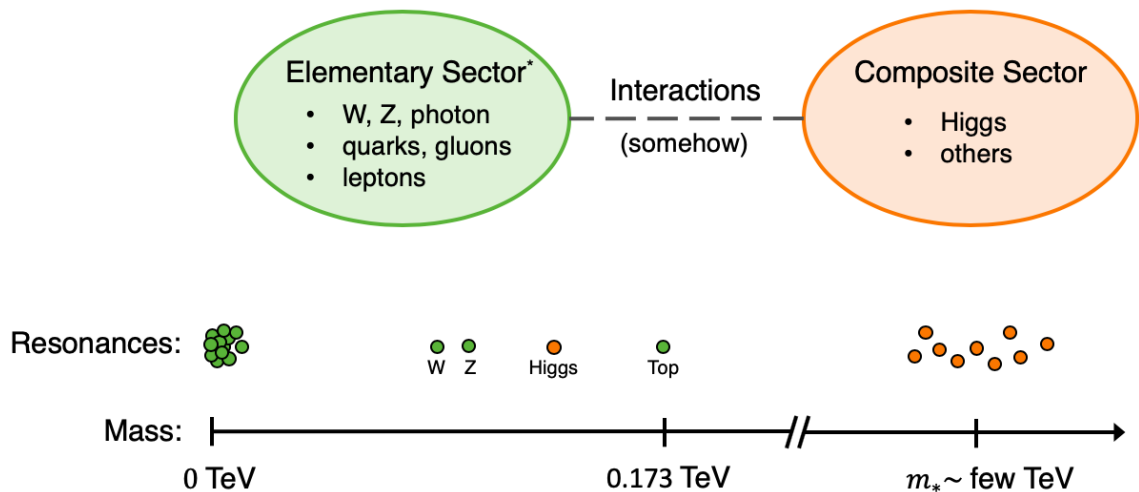


Figure 3.1: Structure of the particle content of a generic composite Higgs model. There is a “composite” sector consisting of resonances of the strong dynamics from which the Higgs arises, contrasted with an “elementary” sector resembling the SM with no Higgs. The composite resonances, aside from the Higgs, are clustered around a mass scale $m_* \gtrsim 1$ TeV. We shall see in Section 3.6 that SM fields are mixtures of elementary and composite fields. The top quark, shown as elementary here, will actually have a large composite component.

¹The SM fields are actually not exactly the elementary fields due to partial compositeness.

From this picture, obvious questions that must be answered in a CHM are:

1. What other resonances arise as bound states of the new strong dynamics?
2. Why is the Higgs so much lighter than the other composite resonances?
3. How do SM particles interact with the new composite sector?
4. How does the Higgs field, as part of the composite sector, break electroweak symmetry and impart mass to the SM fields?

This chapter will survey the theoretical tools that give us a good idea of how to answer these questions and show their application in simple (but relevant) examples. Armed with these tools, a full viable CHM will be put together in Chapter 4.

So that no section of this chapter feels unmotivated, I shall reveal the grand story of CHMs and spoil the surprise now.

3.0.1: Workings of a composite Higgs model

The Higgs boson is light compared to the other composite resonances because it arises from a **spontaneously broken symmetry** structure that would normally dictate it be a massless **Nambu-Goldstone boson**, but the symmetry is only approximate, so the Higgs is instead a naturally light **pseudo-Nambu-Goldstone boson**. The Goldstone symmetry is not exact because, among other things, it contains a gauged subgroup housing the electroweak gauge group, allowing a quantum effective potential for the Higgs to be generated at loop level through the **Coleman-Weinberg mechanism**. Elementary and composite fermions mix through linear couplings, making physical particles **partially composite** superpositions of elementary and composite fields and further spoiling the Goldstone symmetry, giving the Higgs effective potential a non-trivial minimum and triggering the Higgs mechanism on the electroweak gauge subgroup through **vacuum misalignment**. Masses are communicated to the SM particles via Higgs interactions with their composite components - heavier SM particles are typically more connected to the composite sector.

From this description it should be clear that CHMs have rather interconnected structures, perhaps making it difficult to appreciate the purpose of certain features and keep track of where each fits into the chain of cause and effect. It is best to keep the grand story in mind when going through this chapter to avoid losing the forest for the trees.

To pre-emptively clear up a potential source of confusion: throughout this work we will only be interested in low-energy descriptions of CHMs wherein the Higgs can be treated as an elementary field. Its compositeness will only be apparent through the existence of the composite sector. Microscopic realisations of CHMs may however be analysed on the lattice - see Ref. [48] for a review.

3.1 Spontaneously Broken Symmetry

Sometimes it is the case that a particular ground state, or vacuum state, of a theory does not share the same symmetry group as the Lagrangian underlying its dynamics. If the Lagrangian has a symmetry group G and the vacuum state $|\Omega\rangle$ is only invariant under a subgroup $H \subset G$, then it is said that G is **spontaneously broken** to H , written $G \rightarrow H$. The term is somewhat of a misnomer, as a spontaneously broken symmetry is still a symmetry of the Lagrangian. A more accurate name would be a non-linearly realised symmetry, as will be clear in later sections.

Equivalently, because the vacuum is specified by the vevs of all operators², a symmetry is spontaneously broken if it acts non-trivially on the vev of some operator Φ . Note that for this to be the case, the vev $\Phi_0 := \langle \Omega | \Phi | \Omega \rangle$ must be non-zero. That is, only Higgs fields can spontaneously break symmetry. The condition for a general spontaneously broken symmetry in terms of its generator T is

$$e^{i\varepsilon T} |\Omega\rangle \neq |\Omega\rangle, \quad (3.1)$$

from which it is seen that a symmetry is broken if and only if its generator does not annihilate the vacuum state.

Note that a gauge symmetry cannot be spontaneously broken since all physical states are invariant under a gauge transformation. The specific fact that the vacuum is invariant under a gauge transformation - even when physical states connected by gauge transformations are not identified - is known as Elitzur's theorem [49]. However, the notion of a “spontaneously broken gauge symmetry” is prevalent in the literature, and is used to describe when a field has a vev that is not annihilated by generators of the gauge symmetry. From Section 2.3, it should be clear that massive gauge bosons arise due to spontaneously broken gauge symmetries.

3.2 Goldstone's Theorem

The importance of spontaneously broken symmetries, at least for our purposes, lies in Goldstone's theorem, which assures the existence of a massless scalar particle for each differentiable symmetry of the Lagrangian that is spontaneously broken [50,51]. Stated formally:

3.2.1: Goldstone's Theorem

Suppose a Lagrangian \mathcal{L} has a global Lie symmetry group G that acts on an n -tuple of real scalar fields $\Phi = (\phi_1, \dots, \phi_n)$. If the vacuum state $|\Omega\rangle$ is invariant under only a (Lie) subgroup $H \subset G$, then there will be $\dim(G) - \dim(H)$ massless scalar particles in the theory, known as Nambu-Goldstone bosons (NGBs). There is no requirement that an NGB be elementary.

²It can be proven that the physical vacuum is not a superposition of vacua in which fields have different vevs [17].

There is a rather intuitive justification for this theorem. First, it can be shown that if G is a symmetry of the Lagrangian (and of the spacetime integration measure), then it will also be a symmetry of the quantum effective action [17], and therefore of the quantum effective potential too, for the (effective) kinetic terms can be made to vanish by setting the fields constant. The set of all possible vacuum states - that is, the minima of the effective potential - must then be invariant under the action of G . So there will be $\dim(G)$ -many independent directions around any particular vacuum state in which the fields Φ may fluctuate while keeping the effective potential constant, although $\dim(H)$ of these do not leave that vacuum state. Each independent excitation into a new state is a particle, which must be massless since its existence does not require any effective potential energy. The proof below, following the approach of Refs. [52, 53], captures this intuitive picture.

Proof. From the global symmetry G there arise $\dim(G)$ -many conserved Nöther currents J_a^μ ($a = 1, \dots, \dim G$), and hence conserved charges

$$Q_a = \int d^3x J_a^0(\vec{x}). \quad (3.2)$$

The operators J_a^0 of the spontaneously broken symmetries are what excite the Nambu-Goldstone bosons from the vacuum. Indeed, consider these states in the momentum representation:

$$|\Pi_a(\vec{p})\rangle := \int d^3x e^{i\vec{p}\cdot\vec{x}} J_a^0(\vec{x}) |\Omega\rangle. \quad (3.3)$$

First, notice such a state carries the same quantum numbers as $J_a^0 |\Omega\rangle$, and so describes a spin-0 mode since J_a^0 is constructed from scalar fields. Now recall a basic property of the momentum operator \mathbf{P} :

$$[P^i, J(\vec{y})] = i\partial_i J(\vec{y}). \quad (3.4)$$

Using the fact that the vacuum has zero momentum,

$$\begin{aligned} P^i |\Pi_a(\vec{p})\rangle &= \int d^3x e^{i\vec{p}\cdot\vec{x}} P^i J_a^0(\vec{x}) |\Omega\rangle \\ &= \int d^3x e^{i\vec{p}\cdot\vec{x}} [P^i, J_a^0(\vec{x})] |\Omega\rangle \\ &= \int d^3x e^{i\vec{p}\cdot\vec{x}} (i\partial_i J_a^0)(\vec{x}) |\Omega\rangle = p^i |\Pi_a(\vec{p})\rangle, \end{aligned} \quad (3.5)$$

having assumed the current falls off at infinity after integrating by parts. So $|\Pi_a(\vec{p})\rangle$ has a definite momentum \vec{p} . Taking $\vec{p} \rightarrow \vec{0}$ shows

$$|\Pi_a(\vec{0})\rangle = \int d^3x J_a^0(\vec{x}) |\Omega\rangle = Q_a |\Omega\rangle \quad (3.6)$$

is a state with zero momentum. It also has the same energy as the vacuum, for the charge, being conserved, commutes with the Hamiltonian. These are non-trivial statements only for a generator of a broken symmetry, which does not annihilate the vacuum. Putting everything together: for a generator Q_a of a broken symmetry, $Q_a |\Omega\rangle$ is a zero-momentum particle state with the same energy as the vacuum, and so must be a massless particle. \square

Actually, there is a slight cheat in this proof. The Fabri-Picasso theorem states that if $Q_a |\Omega\rangle \neq 0$, then $Q_a |\Omega\rangle$ does not exist in the Hilbert space of the theory: its norm is extensive with the volume of space, and so is (supposedly) infinite [54]. A more rigorous argument is that the energy of $|\Pi_a(\vec{p})\rangle$ becomes arbitrarily close to the vacuum energy as $\vec{p} \rightarrow \vec{0}$, which - if a mass gap exists - shows the particle is massless. Note that this is not a tree-level statement: even after renormalisation, NGBs are exactly massless. The Goldstone symmetry protects their masses from radiative corrections.

As an aside, there is a pervasive misrepresentation in the literature that NGBs are an integral part of the Higgs mechanism. Of course they cannot be, as Goldstone's theorem does not apply to the gauge symmetries spontaneously broken by the Higgs; the proof depends crucially on the non-degeneracy of vacuum states under the symmetry. However, it is true that in a certain gauge the degrees of freedom in the Higgs multiplet act like NGBs that are then "eaten" by the gauge bosons to become massive. Despite being unphysical, this viewpoint is useful for calculations on account of the NGB equivalence theorem: in the high-energy limit, the amplitude for emission or absorption of a longitudinally polarised gauge boson is equal to that of the would-be NGB eaten by that gauge boson [55, 56]. This can be used to show the perturbative unitarity of weak boson scattering at high energies.

3.2.1 Example: the Linear σ -Model

To showcase Goldstone's theorem in an explicit theory, we present the³ linear σ -model. While not the simplest possible example of Goldstone's theorem, this will pay dividends throughout the rest of this work being utilised in increasingly elaborate fashions. The linear σ -model is a theory of $n \geq 2$ real scalar fields $\Phi = (\phi_1, \dots, \phi_n)$ possessing a global $SO(n)$ symmetry that is spontaneously broken to $SO(n-1)$. $SO(n)$, the group of rotations of n -dimensional space, has dimension $n(n-1)/2$, so it is expected that $\dim SO(n) - \dim SO(n-1) = n-1$ NGBs will arise.

The $SO(n)$ -invariant Lagrangian is taken to be

$$\mathcal{L}_\sigma = \frac{1}{2} (\partial_\mu \Phi)^\top (\partial^\mu \Phi) - \frac{g_\sigma^2}{8} (\Phi^\top \Phi - f^2)^2, \quad (3.7)$$

³What we are calling the "linear σ -model" does not have a monopoly on the name. The term can generally refer to any theory of NGBs, no matter the symmetry breaking group. Such models were first introduced by Ref. [57].

where g_σ, f are some positive numbers and Φ is in the fundamental representation of $SO(n)$. If this looks familiar, it is because it is simply the real analogue of the non-gauged SM Higgs Lagrangian (*cf.* Section 2.4) with a slight change in notation.

The potential

$$\mathcal{V}(\Phi) = \frac{g_\sigma^2}{8} (\Phi^\top \Phi - f^2)^2 \quad (3.8)$$

is minimised on the sphere $|\Phi_0| = f$, giving the possible vacuum states⁴. The symmetry will be used to rotate the fields so the vacuum is in the n^{th} direction:

$$\Phi_0 = (0, 0, \dots, 0, f)^\top. \quad (3.9)$$

Clearly, the vacuum is not invariant under $SO(n)$ - a rotation will usually make Φ_0 point in a different direction. It is only the rotations that mix the first $n - 1$ field components that generally leave Φ_0 invariant, showing the symmetry has been broken to $SO(n - 1)$, as desired.

Now how are the resulting NGBs found? Drawing inspiration from the intuitive picture of Goldstone's theorem, the NGBs should be fields parameterising rotations of the vacuum vector Φ_0 into the other vacuum states, as depicted in Figure 3.2. Based on this, we reparameterise Φ as

$$\Phi(x) = e^{i\frac{\sqrt{2}}{f}\theta^a(x)\hat{T}_a} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ f + \sigma(x) \end{pmatrix}, \quad (3.10)$$

where \hat{T}_a are the $n - 1$ “broken” generators of $SO(n)$ that do not leave Φ_0 invariant. Here we have introduced the titular σ field of the model, as well as candidate NGB fields θ^a conventionally normalised by the factor $\sqrt{2}/f$. It is geometrically obvious from Figure 3.2 that such a reparameterisation is possible: the σ field governs the length of Φ , while the θ^a fields govern its direction so that Φ may take on any value.

The central disparity between CHMs and the SM is on display in Equation (3.10). In the SM, the Higgs boson is the analogue of the σ field, and the θ^a fields are the unphysical NGBs that can be removed by a transformation to the unitary gauge. CHMs, having no such gauge symmetry, instead realise the Higgs boson as part of the physical NGBs θ^a , while σ would be some other composite resonance. This exact comparison is not to be taken too seriously, however, for CHMs are *not* based on this kind of model, as the elementary scalar fields in Φ still lead to the Hierarchy Problem.

⁴As before, quantum corrections to the effective potential will be neglected for illustrative purposes.




Figure 3.2: Geometry of the linear σ -model for $n = 3$. The sphere of potential minima is invariant under $SO(3)$, but any particular minimum Φ_0 is only invariant under an $SO(2)$ subgroup. Expanding around Φ_0 , there are two NGB fields $\theta^{1,2}$ rotating Φ along the two dimensions of the sphere, and one field, σ , corresponding to radial perturbations. Since the potential is spherically symmetric, NGB excitations cost no energy and so are massless.

An explicit computation⁵ of Equation (3.10) results in

$$\begin{aligned}\Phi &= \begin{pmatrix} \mathbb{1} - \frac{\vec{\theta}\vec{\theta}^\top}{|\vec{\theta}|^2} \left(1 - \cos \frac{|\vec{\theta}|}{f}\right) & \frac{\vec{\theta}}{|\vec{\theta}|} \sin \frac{|\vec{\theta}|}{f} \\ -\frac{\vec{\theta}^\top}{|\vec{\theta}|} \sin \frac{|\vec{\theta}|}{f} & \cos \frac{|\vec{\theta}|}{f} \end{pmatrix} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ f + \sigma \end{pmatrix} \\ &= (f + \sigma) \begin{pmatrix} \frac{\vec{\theta}}{|\vec{\theta}|} \sin \frac{|\vec{\theta}|}{f} \\ \cos \frac{|\vec{\theta}|}{f} \end{pmatrix},\end{aligned}\quad (3.11)$$

where the phase fields have been collected into a vector $\vec{\theta} = (\theta^1, \dots, \theta^{n-1})^\top$. With this, the Lagrangian Equation (3.7) becomes

$$\begin{aligned}\mathcal{L}_\sigma &= \frac{1}{2}(\partial\sigma)^2 - \frac{1}{2}g_\sigma^2 f^2 \sigma^2 - \frac{1}{2}g_\sigma^2 f \sigma^3 - \frac{1}{8}g_\sigma^2 \sigma^4 \\ &\quad + \frac{1}{2} \left(1 + \frac{\sigma}{f}\right)^2 \frac{f^2}{|\vec{\theta}|^2} \left((\partial_\mu \vec{\theta})^\top (\partial^\mu \vec{\theta}) \sin^2 \frac{|\vec{\theta}|}{f} + \frac{1}{|\vec{\theta}|^2} \left(\frac{|\vec{\theta}|^2}{f^2} - \sin^2 \frac{|\vec{\theta}|}{f} \right) (\vec{\theta}^\top \partial \vec{\theta})^2 \right).\end{aligned}\quad (3.12)$$

A mass for the scalar σ resonance can be read off immediately:

$$m_\sigma = g_\sigma f. \quad (3.13)$$

Drawing the analogy to CHMs, m_σ can be thought of as the compositeness scale m_* . Interactions of the θ fields are found by expanding around $\vec{\theta} = \vec{0}$, yielding

$$\mathcal{L}_\sigma \supset \frac{1}{2} \left(1 + \frac{\sigma}{f}\right)^2 \left((\partial_\mu \vec{\theta})^\top (\partial^\mu \vec{\theta}) + \mathcal{O}\left(\frac{\theta^4}{f^2}\right) \right). \quad (3.14)$$

From this it is seen that the θ fields are indeed the massless NGBs, as was expected. Expansion to higher orders in $\vec{\theta}$ will reveal local interactions with an arbitrarily large number of NGBs. Despite this, the theory is renormalisable.

There are some important points to note from this example.

- The Lagrangian as expressed in Equation (3.12) displays an $SO(n-1)$ symmetry under which $\vec{\theta}$ transforms in the fundamental representation. That is, the NGBs transform linearly in a representation of the unbroken $SO(n-1)$ symmetry. As will be shown in Section 3.4, this holds in greater generality for other symmetry breaking patterns.
- It is no coincidence that mass terms for the phase fields are absent. Consider an arbitrary $SO(n)$ symmetry acting on Φ expressed as in Equation (3.10). This does not change the length of Φ , and so only amounts to changing the phase fields by some constant amount $\theta^a(x) \rightarrow \theta^a(x) + \delta\theta^a$. The fact that

⁵The computation is performed in Appendix F using a convenient basis for the broken generators.

NGBs display this **shift symmetry** forbids a wide range of interactions that are not invariant under the shift symmetry, including mass terms.

3.2.2 Example: the Non-Linear σ -Model

There is a massive scalar σ in the particle spectrum of the linear σ -model. Is there a theory of NGBs that does not include such extra resonances? The answer is yes: the non-linear σ -model. This is simply the linear σ -model in the limit where the coupling g_σ is sent to infinity, decoupling the σ resonance by making it infinitely heavy. In this limit, avoiding an infinite potential energy enforces the constraint

$$\Phi^\top \Phi - f^2 = (f + \sigma)^2 - f^2 = 0, \quad \implies \sigma = 0. \quad (3.15)$$

That is, the physical state Φ is constrained to lie on the sphere $|\Phi| = f$. The space of states being a non-linear manifold is where the model gets its name.

The theoretical description of the non-linear σ -model follows easily from setting $\sigma = 0$ in the linear σ -model equations in the previous section. Doing so is not particularly enlightening, however, and it will do well to write the equations in a notation more suggestive for generalisation. First, the Lagrangian Equation (3.7) reduces to

$$\mathcal{L}_\sigma = \frac{1}{2} (\partial_\mu \Phi)^\top (\partial^\mu \Phi). \quad (3.16)$$

The multiplet Φ can be written in terms of the vacuum vector $\Phi_0 = (0, 0, \dots, 0, f)^\top$ as

$$\Phi(x) = e^{i \frac{\sqrt{2}}{f} \theta^a(x) \hat{T}_a} \Phi_0 \equiv U[\vec{\theta}(x)] \Phi_0, \quad (3.17)$$

where the **NGB matrix**

$$U[\vec{\theta}(x)] := e^{i \frac{\sqrt{2}}{f} \theta^a(x) \hat{T}_a} \quad (3.18)$$

has been introduced. Note the constant f is not interpreted as the minimum of a potential in this model; instead it is usually referred to as the **NGB decay constant**⁶. With this, the Lagrangian takes the form

$$\mathcal{L}_\sigma = \frac{f^2}{2} [(\partial_\mu U)^\top (\partial^\mu U)]_{nn}. \quad (3.19)$$

In contrast to the linear σ -model, this theory is non-renormalisable with a high-energy cutoff⁷ of not much more than $\Lambda_f = 4\pi f$. The theory is best regarded as a low-energy description of the NGBs with some unknown high-energy completion,

⁶This comes from association with the successful description of pions as (pseudo-)NGBs of spontaneously broken chiral symmetry, where f is found to regulate pion decay. See Section 3.3.1.

⁷The argument, as given in Ref. [58], is that $\theta\theta \rightarrow \theta\theta$ scattering will be logarithmically divergent at one-loop level given Equation (3.19). Higher-order operators, weighted by inverse powers of

such as the linear σ -model or some strong dynamics from which the NGBs emerge as bound states. It will follow from Section 3.4 that Equation (3.19) is actually the general low-energy description of NGBs resulting from $SO(n) \rightarrow SO(n-1)$ symmetry breaking (at the two-derivative level), regardless of the high-energy dynamics. The fact that low-energy NGB dynamics are completely specified by the symmetry breaking pattern proves incredibly useful for constructing theories of NGBs.

3.3 Pseudo-Nambu-Goldstone Bosons

It is sometimes the case in nature that the Lagrangian of a system is not exactly invariant under a transformation, but only approximately invariant. How might Goldstone's theorem extend to such "approximate" symmetries⁸? The answer is somewhat intuitive: an exact symmetry results in an exactly massless scalar particle, while an approximate symmetry results in an approximately massless scalar particle, at least compared to the typical energy scale Λ_{break} of the dynamics that explicitly break the symmetry [59].

In terms of the geometrical picture of Goldstone's theorem, the small explicit breaking alters the equipotential surfaces so that oscillations of what were previously NGBs require a small amount of energy, imposing small masses onto the now pseudo-NGBs (pNGBs). Formalising this notion, we regard an approximately symmetric Lagrangian \mathcal{L} as an exactly symmetric Lagrangian \mathcal{L}_{sym} perturbed by terms $\mathcal{L}_{\text{break}}$ that explicitly break the symmetry:

$$\mathcal{L} = \mathcal{L}_{\text{sym}} + \varepsilon \mathcal{L}_{\text{break}},$$

where $\varepsilon \in \mathbb{R}$ parametrises the degree of symmetry breaking, of a size such that $\mathcal{L}_{\text{break}}$ has coupling constants of order unity. Retracing the steps in the proof of Goldstone's theorem with an obvious notation, there will be approximate Nöther currents

$$J^\mu = J_{\text{sym}}^\mu + \varepsilon J_{\text{break}}^\mu, \quad (3.20)$$

giving states

$$|\Pi(\vec{p})\rangle := \int d^3x e^{i\vec{p}\cdot\vec{x}} J^0(\vec{x}) |\Omega\rangle \quad (3.21)$$

that have a definite momentum \vec{p} . But now as $\vec{p} \rightarrow \vec{0}$ the state approaches

$$|\Pi(\vec{0})\rangle = Q_{\text{sym}} |\Omega\rangle + \varepsilon Q_{\text{break}} |\Omega\rangle. \quad (3.22)$$

the cutoff Λ , should be introduced to cancel the divergences, which only maintains perturbative unitarity for $\Lambda \lesssim 4\pi f$.

⁸An approximate symmetry is a symmetry inasmuch as particles can be modelled as transforming under representations of the approximate symmetry group, without having degenerate masses as required by an exact symmetry.

The first term on the right-hand side has the same energy as the vacuum, as before, while $Q_{\text{break}}|\Omega\rangle$ will have an expected energy of around Λ_{break} . Given the symmetry is approximate ($\varepsilon \ll 1$), the mass of our would-be NGB is seen to be non-zero but suppressed compared to Λ_{break} .

3.3.1 Example: Pions in QCD

These are not esoteric considerations; unlike NGBs, pNGBs are actually known to exist in Nature. Pions, for example, are understood to be the pNGBs of spontaneously broken approximate chiral symmetry. Here we give a brief overview of this understanding, based on the treatment in Ref. [53].

The QCD Lagrangian, including only massless u and d quarks for simplicity, is

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}\text{Tr}[G_{\mu\nu}G^{\mu\nu}] + i\bar{q}_L\not{D}q_L + i\bar{q}_R\not{D}q_R, \quad (3.23)$$

where $q_X = (u_X, d_X)^\top$. This Lagrangian is invariant under the transformations $q_X \rightarrow g_X q_X$ for $g_X \in SU(2)_X$. That is, the theory has a chiral $SU(2)_L \times SU(2)_R$ symmetry.

Now it has been experimentally found that the strong dynamics of QCD provides the quark bilinear $\bar{q}^i q^j$ a vev:

$$\langle\Omega|\bar{q}^i q^j|\Omega\rangle = \delta^{ij}\Lambda_{\text{break}}^3, \quad \Lambda_{\text{break}}^3 \approx -(250 \text{ MeV})^3. \quad (3.24)$$

Notice that the bilinear $\bar{q}q = \bar{q}_L q_R + \bar{q}_R q_L$ is only invariant under $SU(2)_L \times SU(2)_R$ transformations where $g_L = g_R$. In other words, the chiral symmetry is spontaneously broken to the diagonal subgroup $SU(2)_{L+R}$.

To describe the resulting NGBs, we follow the lead of the non-linear σ -model and consider the NGB matrix

$$U[\vec{\theta}] = e^{i\frac{\sqrt{2}}{f}\theta^a \hat{T}_a} = e^{i\frac{\sqrt{2}}{f}\theta_a \frac{\sigma_a}{2}} = \exp\left[\frac{i}{\sqrt{2}f}\begin{pmatrix} \theta_3 & \theta_1 - i\theta_2 \\ \theta_1 + i\theta_2 & -\theta_3 \end{pmatrix}\right], \quad (3.25)$$

once again using the (half) Pauli matrices as the generators of $SU(2)$. That the NGB matrix will indeed be of this form will follow from Section 3.8.1. There it will also be seen that the NGBs will transform under the chiral symmetry as a bidoublet,

$$SU(2)_L \times SU(2)_R : \quad U \xrightarrow{(g_L, g_R)} g_L U g_R^{-1}, \quad (3.26)$$

being described by a Lagrangian analogous to Equation (3.19) that contains no mass terms.

This changes, however, once quark masses are included. These *explicitly* break the chiral symmetry for the same reason the quark bilinear *spontaneously* breaks it,

for the mass terms are of the form

$$\mathcal{L}_m = -\bar{q} \begin{pmatrix} m_u & 0 \\ 0 & m_d \end{pmatrix} q \equiv -\bar{q} M q. \quad (3.27)$$

The non-zero mass matrix M , being the breaker of chiral symmetry, must convey masses to the (now pseudo-) NGBs. We can see how by using the trick of **spurion analysis**. Here, we imagine M is actually a dynamical field that transforms in a way that preserves chiral symmetry. Such a fictitious field is termed a spurion. This allows us to construct all chirally-invariant NGB interactions, which must be what impart mass to the NGBs once the spurion is fixed to its “background value” that explicitly breaks the symmetry. In this case the spurion must transform as

$$SU(2)_L \times SU(2)_R : \quad M \xrightarrow{(g_L, g_R)} g_R M g_L^{-1} \quad (3.28)$$

to leave the quark mass terms invariant⁹. The leading-order invariant NGB interaction is then

$$\mathcal{L} \supset \frac{\Lambda_{\text{break}}^3}{2} \text{Tr} [M U + M^\dagger U^\dagger]. \quad (3.29)$$

(The factor of Λ_{break}^3 is to match the vacuum energies of Equations (3.27) and (3.29)). Setting M to its background value $M = \text{diag}(m_u, m_d)$ results in

$$\mathcal{L} \supset \Lambda_{\text{break}}^3 (m_u + m_d) - \frac{\Lambda_{\text{break}}^3}{4f^2} (m_u + m_d) (\theta_1^2 + \theta_2^2 + \theta_3^2) + \mathcal{O}(\theta^3), \quad (3.30)$$

so the former NGBs now have masses

$$m_\theta^2 = \frac{\Lambda_{\text{break}}^3}{2f^2} (m_u + m_d). \quad (3.31)$$

This relation was found by Gell-Mann, Oakes, and Renner [60], who identified the θ fields with the three pions. Notice that here the pNGB masses are regulated by the quantity $m_u + m_d$, which serves the role of ε in Section 3.3. Experimentally, pions have masses of ~ 140 MeV. The next heaviest QCD bound states are the kaons (~ 500 MeV), which are also understood to be pNGBs of another symmetry, and the rho mesons (~ 770 MeV). Considering rho mesons have the same quark content as pions, the pions are indeed seen to be uncannily light.

Composite Higgs models contend that Nature has chosen to repeat this pattern with another symmetry group at a higher energy scale. The composite Higgs boson would be the analogue of the pions, and the other composite resonances analogues of the hadrons. The situation will be more complicated, however, as explicit mass terms for the composite Higgs will not appear in the Lagrangian as they do in

⁹Because M is Hermitian and $g_{L,R}$ are unitary, the transformation can equivalently be described as $M \rightarrow g_L M g_R^{-1}$, from which it follows the mass terms are invariant.

Equation (3.30). Instead, the Higgs mass must be found by computing the radiative corrections to the Higgs effective potential.

3.4 CCWZ Construction

Having seen some examples of theories with NGBs, we now take a detour to outline a general approach for incorporating NGBs into a theory. The approach, known as the Callan-Coleman-Wess-Zumino (CCWZ) construction [61, 62], provides a formalism for writing the general low-energy effective Lagrangian for any theory with a given spontaneously broken symmetry $G \rightarrow H$. CCWZ is an incredibly powerful construction, even lending itself to the description of spontaneously broken approximate symmetries, providing insight into the general properties of (p)NGBs.

The starting point of the construction is to note that since G is spontaneously broken to H , there must be some vacuum vector, conventionally denoted \vec{F} , invariant under H transformations but not generally invariant under G . As usual, G and H are assumed to be connected, compact Lie groups. Generators of G will be split into those that do not annihilate \vec{F} , denoted by \hat{T}_a (indexed by Roman letters), and those that do (*i.e.* those that generate H), denoted X_α (indexed by Greek letters). The construction rests on the fact that any group element $g \in G$ can be written uniquely as

$$g = U[\vec{\theta}]\tilde{h}[\vec{\omega}], \quad (3.32)$$

where

$$U[\vec{\theta}] := e^{i\frac{\sqrt{2}}{f}\theta^a\hat{T}_a}, \quad \text{and} \quad \tilde{h}[\vec{\omega}] := e^{i\frac{\sqrt{2}}{f}\omega^\alpha X_\alpha} \in H, \quad (3.33)$$

for some real numbers θ^a , ω^α restricted to suitable ranges.

To identify the NGBs, consider a multiplet Φ connected by local symmetry transformations to the vacuum vector \vec{F} :

$$\Phi(x) = g(x)\vec{F}. \quad (3.34)$$

If the local transformation $g(x)$ were instead global, then the state Φ would be equivalent to the vacuum \vec{F} , showing the degrees of freedom in the local transformation are massless. The point of the decomposition Equation (3.32) is to dispense with the unphysical degrees of freedom:

$$\Phi(x) = g(x)\vec{F} = U[\vec{\theta}(x)]\tilde{h}[\vec{\omega}(x)]\vec{F} = U[\vec{\theta}(x)]\vec{F}, \quad (3.35)$$

by the invariance of \vec{F} under H . The NGBs are therefore identified as the massless fields θ^a parameterising Φ . Of course this is nothing new; $U[\vec{\theta}(x)]$ is simply the NGB matrix seen in the non-linear σ -model.

Our goal is to understand how these NGB fields transform under G , so that G -invariant NGB interactions may be found and classified. Knowing¹⁰ Φ transforms under $g \in G$ as $\Phi \rightarrow g\Phi$, the transformed NGB fields $\vec{\theta}^{(g)}$ are defined such that

$$\begin{aligned} g\Phi &= U[\vec{\theta}^{(g)}]\vec{F} \\ \implies gU[\vec{\theta}]\vec{F} &= U[\vec{\theta}^{(g)}]\vec{F}. \end{aligned} \quad (3.36)$$

Since \vec{F} is invariant under H , we may conclude

$$gU[\vec{\theta}] = U[\vec{\theta}^{(g)}]h[g, \vec{\theta}], \quad \text{or} \quad U[\vec{\theta}] \xrightarrow{g \in G} U[\vec{\theta}^{(g)}] = gU[\vec{\theta}]h^{-1}[g, \vec{\theta}], \quad (3.37)$$

for some $h \in H$ that depends in some way on the transformation g and the NGB fields. Note that $h[g, \vec{\theta}]$ is unique by the uniqueness proposition in Equation (3.32). The H -invariance of the vacuum led directly to the non-linear transformation property Equation (3.37) of the NGBs, leading to the common terminology that a spontaneously broken symmetry is “non-linearly realised”. But is this transformation property of the NGB matrix consistent - as in, does it form a representation of G ? We explicitly verify this by performing two successive transformations g_1, g_2 :

$$\begin{aligned} U \left[\left(\vec{\theta}^{(g_2)} \right)^{(g_1)} \right] &= g_1 g_2 U[\vec{\theta}] \left(h \left[g_1, \vec{\theta}^{(g_2)} \right] h \left[g_2, \vec{\theta} \right] \right)^{-1} \\ &\equiv (g_1 g_2) U[\vec{\theta}] h^{-1} \left[g_1 g_2, \vec{\theta} \right] \\ &= U \left[\vec{\theta}^{(g_1 g_2)} \right]. \end{aligned} \quad (3.38)$$

This relation is precisely the rule Equation (2.4) that defines a representation, so the transformation Equation (3.37) is consistent. The second line that defines $h[g_1 g_2, \vec{\theta}]$ in Equation (3.38) is valid because of the uniqueness of the decomposition in Equation (3.32). Notice that this definition means $h[\cdot, \vec{\theta}]$ furnishes a representation of the unbroken subgroup H . In the general case this is all that can be said about the NGB transformation rule; without knowing more about the groups G and H , we cannot determine what $h[g, \vec{\theta}]$ actually is.

¹⁰It is not a given that Φ transforms in the fundamental representation of G , but we assume it here.

3.4.1 NGB Transformations Under Unbroken Symmetries

For $H \subset G$ transformations, the NGB transformation can be specified exactly. In this case let us call the symmetry transformation g_H . The transformation rule is

$$\begin{aligned} U \left[\vec{\theta}^{(g_H)} \right] h \left[g_H, \vec{\theta} \right] &\equiv g_H U \left[\vec{\theta} \right] \\ &= g_H \exp \left[i \frac{\sqrt{2}}{f} \theta^a \hat{T}_a \right] g_H^{-1} g_H \\ &= \exp \left[i \frac{\sqrt{2}}{f} \theta^a g_H \hat{T}_a g_H^{-1} \right] g_H. \end{aligned} \quad (3.39)$$

To evaluate this expression, we must note the Lie algebra commutation relations¹¹

$$[X_\alpha, X_\beta] = i f_{\alpha\beta}{}^\gamma X_\gamma + i f_{\alpha\beta}{}^c \hat{T}_c =: (t_\alpha^{\mathbf{Ad}_H})_\beta^\gamma X_\gamma, \quad (3.40)$$

$$[X_\alpha, \hat{T}_b] = i f_{\alpha b}{}^c \hat{T}_c + i f_{\alpha b}{}^\gamma X_\gamma =: (t_\alpha^{\mathbf{r}_{G/H}})_b^c \hat{T}_c. \quad (3.41)$$

The right-hand sides of these equations define the generators $t_\alpha^{\mathbf{Ad}_H}$ of the adjoint representation of H , as well as generators $t_\alpha^{\mathbf{r}_{G/H}}$ of a representation $\mathbf{r}_{G/H}$ of G such that the adjoint representation of G decomposes as

$$\mathbf{Ad}_G = \mathbf{Ad}_H \oplus \mathbf{r}_{G/H}. \quad (3.42)$$

By linearity, Equation (3.41) implies $[i\omega^\alpha X_\alpha, \hat{T}_b] = (i\omega^\alpha t_\alpha^{\mathbf{r}_{G/H}})_b^c \hat{T}_c$ for some numbers ω^α , which is simply the infinitesimal version of

$$e^{i\omega^\alpha X_\alpha} \hat{T}_b e^{-i\omega^\alpha X_\alpha} = \left(e^{i\omega^\alpha t_\alpha^{\mathbf{r}_{G/H}}} \right)_b^c \hat{T}_c. \quad (3.43)$$

Recognising $e^{i\omega^\alpha X_\alpha}$ as a general element of H , for example g_H , Equation (3.39) is calculated as

$$\exp \left[i \frac{\sqrt{2}}{f} \theta^a g_H \hat{T}_a g_H^{-1} \right] = \exp \left[i \frac{\sqrt{2}}{f} \theta^a \left(e^{i\omega^\alpha t_\alpha^{\mathbf{r}_{G/H}}} \right)_a^c \hat{T}_c \right] = U \left[e^{i\omega^\alpha t_\alpha^{\mathbf{r}_{G/H}}} \vec{\theta} \right], \quad (3.44)$$

so in total,

$$U \left[\vec{\theta}^{(g_H)} \right] h \left[g_H, \vec{\theta} \right] = U \left[e^{i\omega^\alpha t_\alpha^{\mathbf{r}_{G/H}}} \vec{\theta} \right] g_H. \quad (3.45)$$

From the uniqueness of the decomposition Equation (3.32), this gives $h[g_H, \vec{\theta}] = g_H$, and shows that the NGBs transform in the $\mathbf{r}_{G/H}$ representation of G under a H

¹¹The first relation comes from H being a Lie group, and the second follows from the first because the structure constants of G are completely antisymmetric, since G is compact. A good reference for many of the results used in this section is Ref. [16].

transformation:

$$\vec{\theta}^{(g_H)} = e^{i\omega^\alpha t_\alpha^{G/H}} \vec{\theta}. \quad (3.46)$$

For $G = SO(n)$ and $H = SO(n-1)$, $\mathbf{r}_{G/H}$ is simply the fundamental representation of H , as was observed in the linear σ -model and will be verified explicitly in Section 3.6.1.

3.4.2 NGB-Gauge Interactions

We have found that under a global transformation $g \in G$, the NGB matrix transforms as

$$U[\vec{\theta}(x)] \rightarrow gU[\vec{\theta}(x)]h^{-1}[g, \vec{\theta}(x)]. \quad (3.47)$$

The only ways the NGB matrix can enter the Lagrangian are through terms that are invariant under Equation (3.47). One such term is $\text{Tr}[U^\dagger U]$, except this is an inconsequential constant since U is unitary. Because of this, we need to use derivatives of the NGB matrix. But notice that even though the transformation in Equation (3.47) is global, the NGB matrix effectively undergoes a local transformation because of the spacetime dependence of the h^{-1} factor, necessitating an analogue of the covariant derivative to be used. It is therefore not much more work to consider the case where some subgroup $\mathcal{L} \subset G$ is gauged (\mathcal{L} for local), which we shall do for full generality.

To keep track of the various aspects of the model, it is helpful to represent the symmetry structure pictorially in what is known as a *moose*¹² diagram [63]. Many conventions for drawing moose diagrams exist, but in this work we use that of Refs. [64, 65]. For the present model the moose diagram is relatively simple:

$$\begin{array}{ccc} \text{Global:} & G & \\ \text{Site:} & \begin{array}{c} \text{---} \circ \text{---} \xrightarrow{U} \text{---} | \\ \text{---} \end{array} & H \cdot \\ \text{Gauged:} & \mathcal{L} & \end{array} \quad (3.48)$$

What this picture signifies is precisely that there is a global symmetry group G acting on a collection of fields (represented by the circle), spontaneously broken to H , giving an NGB matrix U transforming as in Equation (3.47). Of course the bottom row just shows that G has a gauged subgroup \mathcal{L} . More elements will be added to the diagrams as the models grow in complexity, but this diagram serves as a good starting point.

The generators of \mathcal{L} will be denoted L_A (indexed with capital letters), possibly including any generator \hat{T}_a or X_α , having corresponding gauge fields $A_\mu = A_\mu^A L_A$

¹²So named because for more involved models the diagrams are allegedly reminiscent of moose antlers. They are also known as quiver diagrams, especially when utilised in string theory.

and gauge coupling $g_{\mathcal{L}}$. Note that some gauge bosons may become massive through the Higgs mechanism, making the corresponding NGBs unphysical. Also note that the gauging of \mathcal{L} will explicitly break G symmetry, leading to pNGBs, unless G itself is the normaliser group of \mathcal{L} in G (see Section 2.2). It is good to bear in mind that many common Lie groups (including $SO(n)$ for odd n , which we will be interested in later) are *simple*, meaning the exact symmetry group left over when a subgroup \mathcal{L} is gauged is \mathcal{L} itself.

The trick to constructing the covariant derivative for the NGBs lies in the quantity

$$\mathbf{a}_\mu[U, A] := iU^{-1}(\partial_\mu - ig_{\mathcal{L}}A_\mu)U, \quad (3.49)$$

which is a covariant version of the so-called Maurer-Cartan form $iU^{-1}\partial_\mu U$. This turns out to be in the Lie algebra of G , so it may be written as the linear combination

$$\mathbf{a}_\mu[U, A] = d_\mu^a[U, A]\hat{T}_a + e_\mu^\alpha[U, A]X_\alpha \equiv d_\mu[U, A] + e_\mu[U, A] \quad (3.50)$$

for some coefficients d_μ^a, e_μ^α . Why do we consider the strange object \mathbf{a}_μ ? Simply because it has nice transformation properties. Recalling the gauge field transformation rule Equation (2.13), \mathbf{a}_μ transforms under a local G transformation as (suppressing dependencies)

$$\begin{aligned} \mathbf{a}_\mu[U, A] &\rightarrow i(gUh^{-1})^{-1} \left(\partial_\mu - ig_{\mathcal{L}} \left(gA_\mu g^{-1} - \frac{i}{g_{\mathcal{L}}}(\partial_\mu g)g^{-1} \right) \right) (gUh^{-1}) \\ &= ihU^{-1}g^{-1}(\partial_\mu g)Uh^{-1} + ihU^{-1}(\partial_\mu U)h^{-1} + ih(\partial_\mu h^{-1}) \\ &\quad + ihU^{-1}(-ig_{\mathcal{L}}A_\mu)Uh^{-1} - ihU^{-1}g^{-1}(\partial_\mu g)Uh^{-1} \\ &= h\mathbf{a}_\mu[U, A]h^{-1} + ih(\partial_\mu h^{-1}). \end{aligned} \quad (3.51)$$

Notice the second term in Equation (3.51) is itself a Maurer-Cartan form of the H transformation h , so it is a linear combination of only the H generators X_α . In terms of the decomposition Equation (3.50), then, the transformation rules are

$$\begin{aligned} d_\mu[U, A] &\xrightarrow{g \in G} h[g, \vec{\theta}] d_\mu[U, A] h^{-1}[g, \vec{\theta}], \\ e_\mu[U, A] &\xrightarrow{g \in G} h[g, \vec{\theta}] e_\mu[U, A] h^{-1}[g, \vec{\theta}] + ih[g, \vec{\theta}](\partial_\mu h^{-1}[g, \vec{\theta}]). \end{aligned} \quad (3.52)$$

These transformation rules are very familiar, making it easy to construct invariant interactions. In fact, all invariant quantities are constructed out of combinations of these objects, apart from the Wess-Zumino-Witten term present when G has a global anomaly [66–68].

By the same reasoning as in Equation (3.44), the d_μ object is seen to transform in the $\mathbf{r}_{G/H}$ representation of G :

$$d_\mu^a[U, A] \rightarrow \left(e^{i\omega^\alpha t_\alpha^{G/H}} \right)_b^a d_\mu^b[U, A], \quad (3.53)$$

even for a transformation $g \notin H$, in contrast to the NGBs. The lowest-order invariant it forms is

$$\text{Tr} [d_\mu^\dagger d^\mu]. \quad (3.54)$$

The e_μ object, on the other hand, transforms just as a gauge field and thus enters the Lagrangian through the covariant derivative (with a coupling constant of unity). This allows couplings between NGBs and matter fields. The covariant derivative of the NGB matrix is

$$D_\mu U = \partial_\mu U - ig_{\mathcal{L}} A_\mu U + iU e_\mu, \quad (3.55)$$

transforming as

$$(D_\mu U) \rightarrow g(D_\mu U)h^{-1}. \quad (3.56)$$

It follows that $\text{Tr} [(D_\mu U)^\dagger (D^\mu U)]$ is also an invariant. But this turns out to be nothing but Equation (3.54); unravelling the definitions in Equation (3.55), it is found that $iU^\dagger D_\mu U = d_\mu$. To work out what this invariant actually contains, we expand Equation (3.49) up to first order in the NGB fields:

$$\alpha_\mu[U, A] = -\frac{\sqrt{2}}{f}(\partial_\mu \theta^a)\hat{T}_a + g_{\mathcal{L}} A_\mu^A L_A + \mathcal{O}(\theta^2). \quad (3.57)$$

By definition, then,

$$d_\mu[U, A] = -\frac{\sqrt{2}}{f}(\partial_\mu \theta^a)\hat{T}_a + g_{\mathcal{L}} A_\mu^{\tilde{A}} \hat{T}_{\tilde{A}} + \mathcal{O}(\theta^2), \quad (3.58)$$

where the index \tilde{A} runs over the gauged generators that also happen to be broken, *i.e.* the numbers such that $L_{\tilde{A}} = \hat{T}_{\tilde{A}}$. With this, recalling the generators are normalised such that $\text{Tr} [\hat{T}_a \hat{T}_b] = \delta_{ab}$, the invariant is calculated to be

$$\text{Tr} [d_\mu^\dagger d^\mu] = \frac{2}{f^2}(\partial_\mu \vec{\theta})^\top (\partial_\mu \vec{\theta}) - \frac{2\sqrt{2}}{f} g_{\mathcal{L}} (\partial_\mu \theta^{\tilde{A}}) A^{\mu \tilde{A}} + g_{\mathcal{L}}^2 A_\mu^{\tilde{A}} A^{\mu \tilde{A}} + \mathcal{O}(\theta^3) \quad (3.59)$$

It is shown in Refs. [61, 62] that this is the general invariant term that is leading-order in the number of derivatives. Much more discussion on possible invariants can be found in Ref. [69], but the rest of this work uses no more than is shown here.

In total, the conventionally-normalised low-energy Lagrangian for NGBs and possible gauge fields A_μ is

$$\mathcal{L}_{(2)} = \frac{f^2}{4} \text{Tr} [(D_\mu U)^\dagger (D^\mu U)] - \frac{1}{4} \text{Tr} [A_{\mu\nu} A^{\mu\nu}], \quad (3.60)$$

where the subscript serves to remind that this is at the two-derivative level. Note that these low-energy NGB-gauge interactions are entirely specified by the groups G and H , with strengths dictated by the NGB decay constant f .

3.4.3 NGB-Fermion Interactions

To finish off the construction, we consider how other resonances are to be included in the theory. The key point here is that since the G symmetry is spontaneously broken to H , at low energies the Lagrangian must be linearly invariant under only H , and not G . We therefore consider fields transforming under some linear representation ρ of H :

$$\psi \xrightarrow{h \in H} \rho(h)\psi. \quad (3.61)$$

The representation ρ may be reducible or irreducible, so it is without loss of generality that ψ can be considered as one large multiplet containing all resonances in the theory. Of course, if we were not ignorant of the high-energy dynamics from which \vec{F} arises then the G symmetry would be restored when \vec{F} is transformed appropriately. So actually the fields ψ must transform under G while only appearing to transform under H . A prime candidate for how the fields are to transform, then, is through the g -dependent H transformations of the NGB matrix,

$$\psi \xrightarrow{g \in G} \rho(h[g, \vec{\theta}])\psi. \quad (3.62)$$

That is, if the generators of representation ρ are denoted t_α^ρ , then

$$\psi \xrightarrow{g \in G} \exp \left[i\omega^\alpha [g, \vec{\theta}] t_\alpha^\rho \right] \psi \quad \text{for} \quad h[g, \vec{\theta}] = \exp \left[i\omega^\alpha [g, \vec{\theta}] X_\alpha \right]. \quad (3.63)$$

The interdependence of the transformations of ψ and the NGB matrix U make it possible to construct G -invariant terms. Consider, for example, the case where $\rho(h) = h$. Then the object $\Psi := U\psi$ will transform as

$$\Psi \xrightarrow{g \in G} (gU h^{-1}[g, \vec{\theta}]) (h[g, \vec{\theta}]\psi) = g\Psi, \quad (3.64)$$

making it easy to write down a G -invariant Lagrangian. This way of “dressing” the H multiplet ψ with the NGB matrix can be extended to other representations as a general method of turning H transformations into G transformations.

Similarly to the NGB matrix, ψ appears to undergo local transformations that must be compensated for by a covariant derivative. Fortunately, we already know

the gauge field analogue of the induced H transformations: e_μ . Including the \mathcal{L} gauge fields as well, the covariant derivative is simply

$$D_\mu\psi = (\partial_\mu - ig_{\mathcal{L}}A_\mu - ie_\mu^\alpha t_\alpha)\psi, \quad (3.65)$$

transforming as

$$D_\mu\psi \xrightarrow{g \in G} \rho(h[g, \vec{\theta}])D_\mu\psi. \quad (3.66)$$

This covariant derivative introduces interactions between the resonances and NGBs that are once again dictated by the decay constant f , as well as the representation ρ . The beauty of the CCWZ construction is that the most general Lagrangian with a $G \rightarrow H$ symmetry breaking pattern is built only from G -invariant pairings of the objects d_μ , ψ , and $D_\mu\psi$ that we have seen here [61, 62].

3.5 Vacuum Misalignment

Based on what has been presented so far, the basic ingredients of a composite Higgs model are:

- 1) a strongly interacting sector
- 2) acted upon by a global approximate symmetry group G
- 3) that is spontaneously broken to some subgroup H ,
- 4) delivering pNGBs consisting of (at least) the Higgs doublet, in order to
- 5) spontaneously break electroweak symmetry.

It will do well to consider how this structure might be enacted. Conditions (4) and (5) impose important constraints on the symmetry groups G and H , about which nothing is known at this point.

First though, it is wise to perform a sanity check as to whether a pNGB can actually act as a Higgs boson. After all, an NGB certainly cannot - by the shift symmetry noted in Section 3.2.1, an NGB vev amounts to a symmetry transformation and so is inconsequential. But this problem does not extend to pNGBs, for the associated symmetry is not exact in their case. The explicit symmetry breaking generates an effective potential for the pNGBs, giving them a definite vev. So this requirement is at least consistent.

Now let us see how electroweak symmetry may be broken. From the $G \rightarrow H$ spontaneous breaking, there must be a vacuum vector \vec{F} that is invariant under $H_F \cong H$ transformations but not generally invariant under G . It is not the concern of composite Higgs models how this breaking happens; it is simply taken for

granted that such an \vec{F} exists. Following the CCWZ construction, states will be parameterised by pNGB fields θ^a as

$$\Phi(x) = e^{i\frac{\sqrt{2}}{f}\theta^a(x)\hat{T}_a}\vec{F}, \quad (3.67)$$

where \hat{T}_a are the generators of G that do not annihilate \vec{F} . What happens if the pNGBs acquire a vev?

$$\begin{aligned} \langle\Phi\rangle &= e^{i\frac{\sqrt{2}}{f}\langle\theta^a\rangle\hat{T}_a}\vec{F} \\ &= \begin{cases} \vec{F} & \text{if } \langle\vec{\theta}\rangle = \vec{0}, \\ \text{not } \vec{F} & \text{if } \langle\vec{\theta}\rangle \neq \vec{0}. \end{cases} \end{aligned} \quad (3.68)$$

In other words, the true vacuum $\langle\Phi\rangle$ will not be the H_F -invariant vacuum \vec{F} if the pNGBs have a non-zero vev. This is an important point: if pNGBs acquire a non-zero vev, H_F symmetry will be spontaneously broken! The obvious course of action for composite Higgs models, then, is to choose the arbitrary reference vacuum \vec{F} such that the electroweak gauge group $G_{EW} = SU(2)_L \times U(1)_Y$ is embedded inside H_F . This type of mechanism for breaking electroweak symmetry, known as **vacuum misalignment**, was introduced by Ref. [38] and further explored in the context of composite Higgs models in Ref. [40].

The subtlety of vacuum misalignment is perhaps best recognised when contrasted with a theory such as Minimal Technicolor, which does not utilise the idea. Technicolor models are precursors to CHMs, also relying on a $G \rightarrow H$ symmetry breaking by strong dynamics, where the electroweak group G_{EW} is embedded inside G . However, in Minimal Technicolor, the unbroken group H is exactly the electromagnetic group $U(1)_Q$. Electroweak symmetry is therefore broken directly when G is broken to H , and so the EWSB scale is the same as the scale of $G \rightarrow H$ breaking. With vacuum misalignment, on the other hand, the $G \rightarrow H$ breaking does *not* break electroweak symmetry. Dynamics outside of the composite sector are instead required to misalign the vacuum to trigger EWSB, with the degree of misalignment controlling the scale of the EWSB. This allows EWSB to occur at a lower scale than $G \rightarrow H$ breaking, as is experimentally required.

Quite general results show that the explicit breaking introduced by the gauged electroweak group, or any vector boson fields, tends to keep the vacuum aligned in the EW-preserving direction in theories with pNGBs [70–72]. There must therefore be a different source of explicit breaking to misalign the vacuum and spontaneously break EW symmetry. In CHMs this is achieved through coupling elementary and composite fermions together, to be discussed in Section 3.6.

Vacuum misalignment lends itself nicely to geometric visualisation. The case for $G = SO(3)$ and $H = SO(2)$ is shown in Figure 3.3. Of course this is not a realistic model for EWSB because H is not large enough to contain the electroweak gauge group, but we are limited to only a three-dimensional visualisation. Fortunately,




Figure 3.3: Geometry of vacuum misalignment in the case $G = SO(3)$, $H = SO(2)$. The sphere of potential minima is invariant under G , but any particular minimum \vec{F} is only invariant under $H_F \cong SO(2)$. If the Goldstone fields $\vec{\theta}$ acquire a non-zero vev, the true vacuum $\langle \vec{\Phi} \rangle$ will not be the H_F -symmetric vacuum \vec{F} and will instead spontaneously break H_F symmetry.

the coset $SO(3)/SO(2)$ is indeed isomorphic to the 2-sphere¹³, so the geometric intuition will be correct in this case.

Let us imagine that H_F contains a gauged subgroup analogous to G_{EW} . To draw the connection to EWSB, the degree of the spontaneous breaking of H_F should be identified with the EWSB scale $v \approx 246$ GeV. This degree of breaking can be parameterised by the projection of the true vacuum $\langle \vec{\Phi} \rangle$ onto H_F , which from the geometry in Figure 3.3 is seen to be

$$v = f \sin \frac{|\langle \vec{\theta} \rangle|}{f}. \quad (3.69)$$

We shall see in Section 3.5.1 that this relation holds in a simple version of the minimal CHM, though it is not true for general symmetry breaking patterns.

It is instructive to consider Equation (3.69) in more detail. The meaning of Equation (3.69) is that the gauge boson arising from the gauged generator of H_F has a mass

$$m \equiv \frac{1}{2}gv = \frac{1}{2}gf \sin \frac{|\langle \vec{\theta} \rangle|}{f}, \quad (3.70)$$

where g is the gauge coupling. This makes geometric sense. First, if the unbroken generator X that rotates around the true vacuum $\langle \vec{\Phi} \rangle$ were gauged, then the corresponding gauge boson would be massless. On the other hand, if a generator \hat{T} orthogonal to X (rotating the true vacuum along a great circle) were gauged, then the resulting gauge boson would have a mass $m = \frac{1}{2}gf$, for f is the scale of the symmetry breaking¹⁴. But the generator that is actually gauged is T_3 - the one rotating about the vertical axis in Figure 3.3, leaving \vec{F} invariant. From the geometry, this is the linear combination

$$T_3 = \sin \frac{|\langle \vec{\theta} \rangle|}{f} \hat{T} + \cos \frac{|\langle \vec{\theta} \rangle|}{f} X, \quad (3.71)$$

where \hat{T} rotates about the axis perpendicular to $\langle \vec{\Phi} \rangle$ in the plane spanned by $\langle \vec{\Phi} \rangle$ and \vec{F} . It follows from Equation (2.21) that the mass of the boson associated to T_3 is $\sin |\langle \vec{\theta} \rangle|/f$ times that of the boson associated to \hat{T} , giving Equation (3.70).

It is common to express the symmetry breaking scale in a general composite Higgs model through the **vacuum misalignment parameter**

$$\xi := \frac{v^2}{f^2}, \quad (3.72)$$

¹³More generally, $SO(n)/SO(n-1)$ is isomorphic as a manifold to the $(n-1)$ -sphere by the orbit-stabiliser theorem.

¹⁴If this reasoning is unconvincing, these claims will be substantiated in Section 3.8.2.

which in the case of $SO(3) \rightarrow SO(2)$ symmetry breaking is evidently

$$\xi = \sin^2 \frac{|\langle \vec{\theta} \rangle|}{f}. \quad (3.73)$$

Given we have not yet discovered the compositeness scale f , a necessary constraint is $0 < \xi < 1$. Technicolor is reproduced by the limit $f \rightarrow v$, or $\xi \rightarrow 1$. What happens in the opposite limit, where $\xi \rightarrow 0$ or equivalently $f \rightarrow \infty$? Because the NGBs enter the Lagrangian through the combination θ/f , this limit will remove operators involving a sufficient number of NGBs. In fact, the only non-vanishing term in the NGB Lagrangian Equation (3.60) is

$$(D_\mu \vec{\theta})^\dagger (D^\mu \vec{\theta}). \quad (3.74)$$

For models that only deliver the Higgs doublet and no other NGBs, then, the $\xi \rightarrow 0$ limit reproduces the SM Higgs kinetic Lagrangian, including gauge interactions, exactly. We will see this explicitly for the minimal CHM in Section 3.5.1.

The experimental success of the SM indicates ξ must be rather small. In particular, precision electroweak tests typically require $\xi \lesssim 0.1$ [44]. But generically there is no reason to expect ξ to be particularly small. In fact if it were too small then we would have another Hierarchy Problem on our hands - the hierarchy between f and v - exactly what we are trying to avoid! To keep track of this, the fine-tuning of a model is commonly estimated as

$$\sim \frac{1}{\xi}, \quad (3.75)$$

so that $\xi \approx 0.1$ comes at a cost of $\sim 10\%$ precision of the parameters. This value is not too concerning and can be reasoned away by accidental cancellations in the NGB potential, but values $\xi \lesssim 0.01$ start becoming unacceptable. It is reasonable to conjecture some type of structure in the theory that naturally produces a small misalignment angle instead of relying on accidental cancellations. This is done, for example, in “little Higgs” variants of CHMs [73, 74], but these models tend to predict too large a Higgs mass.

3.5.1 EWSB in the Minimal Composite Higgs Model

We are now in a position to demonstrate electroweak symmetry breaking in the composite Higgs paradigm. We do so with the minimal group structure $G \rightarrow H$ that leads to a phenomenologically viable model, deemed the **Minimal CHM** (MCHM) [44]. Rather than simply stating the group structure of the MCHM, let us reason what it must be.

Group structure

First of all, from Section 3.5, we know that H must contain the electroweak group $SU(2)_L \times U(1)_Y$. But actually this is not all; since we would like the model to reproduce the experimentally verified prediction

$$\frac{m_W^2}{m_Z^2} = \frac{g^2}{g^2 + g'^2}, \quad (3.76)$$

we demand an enlarged $SU(2)_L \times SU(2)_R \cong SO(4)$ symmetry of the electroweak sector in order to guarantee Equation (3.76) through custodial $SU(2)_{L+R}$ symmetry (see Appendix E). In the minimal model, then, $H = SO(4)$. See Appendix D for more details on the isomorphism $SU(2)_L \times SU(2)_R \cong SO(4)$, which will be vital in the rest of this work.

We know from Appendix D that the Higgs doublet can be expressed as a vector in the fundamental representation of $SO(4)$. Our experiences in Sections 3.2.1 and 3.4.1 then establish $G = SO(5)$ as a viable symmetry group. In fact, since the breaking $SO(5) \rightarrow SO(4)$ delivers four NGBs, it delivers exactly the Higgs doublet and so must be precisely the minimal group structure! It is somewhat miraculous how well this symmetry structure fits with SM phenomenology. We shall see later that these groups must be extended to accommodate the composite matter sector, but the underlying symmetry breaking pattern remains the same.

The MCHM is by no means the only CHM, nor the only interesting CHM. The next-to-minimal model, for example, is based on $SO(6) \rightarrow SO(5)$ breaking and features an extra pNGB uncharged under the electroweak group that can serve as a dark matter candidate [18]. A wide range of other models is discussed in Ref. [75]. This work, however, focusses simply on the MCHM.

EWSB in the MCHM

As if by design, we have already figured out the low-energy NGB Lagrangian for the MCHM in Section 3.2.2:

$$\mathcal{L}_{\text{NGB}} = \frac{f^2}{2} (\partial_\mu \Phi)^\top (\partial^\mu \Phi), \quad \Phi = \begin{pmatrix} \frac{\vec{\theta}}{|\vec{\theta}|} \sin \frac{|\vec{\theta}|}{f} \\ \cos \frac{|\vec{\theta}|}{f} \end{pmatrix}. \quad (3.77)$$

Here $\vec{\theta} = (\theta^1, \dots, \theta^4)^\top$ is a fourplet of $SO(4)$. The usual Higgs doublet, following the isomorphisms of Appendix D, will be

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} \theta^2 + i\theta^1 \\ \theta^4 - i\theta^3 \end{pmatrix}. \quad (3.78)$$

So far this is just the non-linear σ -model. To make this a theory of EWSB, we need to gauge the electroweak generators of the unbroken $SO(4)$ subgroup as discussed in Section 3.5. It is convenient to express the $SO(4)$ generators in terms of the

generators $T_{L,R}^3$ of the subgroups locally isomorphic to $SU(2)_{L,R}$, which are given explicitly in Appendix G. The hypercharge operator is identified as T_R^3 , for this reproduces the correct hypercharge for the Higgs doublet, as shown in Appendix D. Gauging $SU(2)_L \times U(1)_Y$ is the process of introducing the covariant derivative

$$\partial_\mu \Phi \rightarrow D_\mu \Phi = (\partial_\mu - igW_\mu^a T_L^a - ig' B_\mu T_R^3) \Phi, \quad (3.79)$$

along with the gauge kinetic terms. Calculating the covariant derivative in full is rather formidable, so we take the easy way out and move to the unitary gauge where $\theta^1 = \theta^2 = \theta^3 = 0$ and $\theta^4 = h$. Here,

$$D_\mu \Phi = \frac{\partial_\mu h}{f} \begin{pmatrix} 0 \\ 0 \\ 0 \\ \cos \frac{h}{f} \\ -\sin \frac{h}{f} \end{pmatrix} + \frac{1}{2} \begin{pmatrix} gW_\mu^1 \\ gW_\mu^2 \\ gW_\mu^3 - g'B_\mu \\ 0 \\ 0 \end{pmatrix} \sin \frac{h}{f}, \quad (3.80)$$

giving the Lagrangian

$$\begin{aligned} \mathcal{L} = & -\frac{1}{4} \text{Tr} [W_{\mu\nu} W^{\mu\nu}] - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} + \frac{1}{2} (\partial h)^2 \\ & + \frac{1}{4} (g^2 W_\mu^1 W^{1,\mu} + g^2 W_\mu^2 W^{2,\mu} + (g' B_\mu - gW_\mu^3)^2) f^2 \sin^2 \frac{h}{f}. \end{aligned} \quad (3.81)$$

By now we are familiar with the pattern of EW boson interactions and perform the usual substitutions Equation (2.36) for the SM W^\pm and Z bosons:

$$\begin{aligned} \mathcal{L} = & -\frac{1}{4} \text{Tr} [W_{\mu\nu} W^{\mu\nu}] - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} + \frac{1}{2} (\partial h)^2 \\ & + \frac{1}{4} \left(g^2 W_\mu^- W^{\mu+} + \frac{1}{2} (g^2 + g'^2) Z_\mu Z^\mu \right) f^2 \sin^2 \frac{h}{f}. \end{aligned} \quad (3.82)$$

Now the gauged subgroup will explicitly break the $SO(5)$ symmetry, generating an effective potential for the Higgs doublet. We will not concern ourselves with the specifics at this point, and just declare the Higgs field will receive some vev $\langle h \rangle$ and redefine $h \rightarrow \langle h \rangle + h$. Expanding around $h = 0$ produces mass terms in the second line of Equation (3.82):

$$m_W = \frac{1}{2} g f \sin \frac{\langle h \rangle}{f}, \quad m_Z = \frac{1}{2} \sqrt{g^2 + g'^2} f \sin \frac{\langle h \rangle}{f}, \quad (3.83)$$

explicitly confirming the custodial relation Equation (3.76), as we knew must have been the case. Matching to the SM mass term $m_W = \frac{1}{2} g v$ also confirms the vacuum

misalignment

$$\xi \equiv \frac{v^2}{f^2} = \sin^2 \frac{\langle h \rangle}{f} \quad (3.84)$$

predicted in Equation (3.69). Further expansion of the Lagrangian reveals the infinite series of interaction terms

$$\begin{aligned} \mathcal{L} = & -\frac{1}{4} \text{Tr} [W_{\mu\nu} W^{\mu\nu}] - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} + \frac{1}{2} (\partial h)^2 \\ & + \left(m_W^2 W^{\mu+} W_{\mu}^- + \frac{1}{2} m_Z^2 Z^\mu Z_\mu \right) \left(1 + 2\sqrt{1-\xi} \frac{h}{v} + (1-2\xi) \frac{h^2}{v^2} + \mathcal{O} \left(\xi \frac{h^3}{v^3} \right) \right). \end{aligned} \quad (3.85)$$

Notice that in the limit $\xi \rightarrow 0$, this reduces to the SM gauge-Higgs Lagrangian Equation (2.41), as was reasoned previously.

We have found a precise prediction of the MCHM: the gauge-Higgs interactions are given by exactly Equation (3.85), needing only one free parameter ξ to be fully specified. It is then in principle very easy to falsify the MCHM. If these couplings are measured and found to not follow the pattern of predicted strengths, the model can be discarded. As of yet, though, these couplings have not been measured precisely enough to provide much evidence for or against the model. Beware, however, that this is only a simple incarnation of the MCHM, and the models we ultimately consider in this work have an extended gauge sector with couplings that are not so simple.

In all, EWSB in composite Higgs models is very similar to EWSB in the SM: the Higgs field interacts with the EW gauge fields through the covariant derivative, and once it acquires a vev the gauge fields receive masses. It is only in the details of how the Higgs enters the Lagrangian and how it develops a vev that the two approaches differ.

3.6 Partial Fermion Compositeness

We come now to the question of how the SM fermions are to fit in to all of this. They must couple to the composite sector somehow, firstly in order to explicitly break the global composite symmetry and trigger EWSB (*cf.* Section 3.5), and secondly to obtain mass through interactions with the composite Higgs. Following the approach of Refs. [12, 69], we first show an approach that is *not* viable in order to contrast the modern viewpoint of CHMs.

A Non-Viable Type of Fermion Coupling

Perhaps the first approach one might try would be to generalise the SM Yukawa interactions into terms such as

$$\mathcal{L} \supset \frac{\Delta_t(\Lambda)}{\Lambda^{[\mathcal{O}]-1}} \bar{q}_L \mathcal{O} t_R + \text{h.c.} \quad (3.86)$$

where $[\mathcal{O}]$ denotes the mass (scaling) dimension of some composite operator \mathcal{O} that has the same quantum numbers as the Higgs field. Here Λ is the UV cutoff of the theory, and $\Delta_t(\Lambda)$ is some dimensionless coupling of order unity at that scale. Indeed, this is what was done historically in early Technicolor models [25, 33–35] and even early CHMs [39, 40]. This approach, however, invariably leads to tensions with either the Hierarchy Problem, with the fermion masses, or with experimentally disfavoured flavour-changing neutral currents (FCNCs).

In the SM, $[\mathcal{O}] \approx 1$. In this case, the operator \mathcal{O}^2 will have dimension $[\mathcal{O}^2] \approx 2$ [76], which was the cause of the Hierarchy Problem from an effective field theory viewpoint.

Technicolor avoids this pitfall by using a bilinear composite fermion operator $\mathcal{O} = \bar{\psi}\psi$ with dimension $[\mathcal{O}] \approx 3$ that acts as a Higgs field. The problem in this case is that the Yukawa couplings in Equation (3.86) naïvely evolve to the energy scale $\mu < \Lambda$ as

$$\Delta_t(\mu) \approx \Delta_t(\Lambda) \left(\frac{\mu}{\Lambda}\right)^{[\mathcal{O}]-1} \approx \Delta_t(\Lambda) \left(\frac{\mu}{\Lambda}\right)^2, \quad (3.87)$$

becoming highly suppressed at low scales $\mu \ll \Lambda$ and tending to produce extremely light fermions. The suppression could be alleviated if Λ were lower, but this turns out to result in more abundant FCNCs resulting from the four-fermion interactions. The fact that all couplings evolve by the same factor also does not satisfyingly resolve the observed hierarchy of fermion masses. Models such as Walking Technicolor [77], with more sophisticated proposals regarding the evolution of the couplings, or Extended Technicolor [33, 34], utilising higher symmetry scales, offer solutions to these problems, but we will not discuss them further.

Fermion Couplings in CHMs

The modern approach of CHMs was put forward by Kaplan [78], wherein all elementary fermions couple *linearly* to their own composite operators:

$$\mathcal{L} \supset \frac{\Delta_{t_L}(\Lambda)}{\Lambda^{[\mathcal{O}_{t_L}]-\frac{5}{2}}} \bar{t}_L \mathcal{O}_{t_L} + \frac{\Delta_{t_R}(\Lambda)}{\Lambda^{[\mathcal{O}_{t_R}]-\frac{5}{2}}} \bar{t}_R \mathcal{O}_{t_R} + \text{h.c.} \quad (3.88)$$

with analogous terms for the other fermions. Although it may seem unusual to postulate such couplings, this coupling paradigm is a natural consequence in the four-dimensional models that are dual to five-dimensional models of EWSB [79]. With this coupling structure the composite operators must be fermionic to respect

Poincaré symmetry. By simple dimensional analysis, the couplings will evolve as

$$\Delta_{t_{L,R}}(\mu) \approx \Delta_{t_{L,R}}(\Lambda) \left(\frac{\mu}{\Lambda}\right)^{[\mathcal{O}_{t_{L,R}}]-\frac{5}{2}}. \quad (3.89)$$

Notice that all of the problems with the Yukawa-like interactions in Equation (3.86) can be avoided with this scaling behaviour. FCNCs can be rendered negligible by having a large cutoff Λ , and this would not significantly suppress the couplings as long as the operators have dimensions $[\mathcal{O}_{t_{L,R}}] \approx 5/2$. The Hierarchy Problem is then not a concern, as the operators $\mathcal{O}_{t_{L,R}}^2$ will be irrelevant. Additionally, couplings for each fermion may evolve differently to others depending on the dimensions of their associated composite operators, providing a reasonable explanation for the observed hierarchy of fermion masses.

Clearly, the fields in Equation (3.88) are not expressed in the mass basis. The mass eigenstates - the particles that are physically observed - must instead be superpositions of elementary and composite fields. For this reason, the coupling paradigm of Equation (3.88) is known as **partial compositeness**. SM fields are therefore mixtures of elementary and composite fields, and we assign to the elementary sector one field for each (non-Higgs) field of the SM, carrying the same quantum numbers as its associated SM field. Note that although here we are focussing on the fermion sector, the gauge bosons are also subject to partial compositeness.

The partial compositeness hypothesis shines light on not just the elementary sector, but the composite sector also. At this point, all that is known about the composite sector is that it has a $G \rightarrow H$ symmetry breaking pattern, where H contains the electroweak gauge group, and the groups G and H are such that the NGB spectrum contains the Higgs doublet. With partial compositeness, though, for mixing between elementary and composite states to occur, there must be at least one spin-1/2 composite field sharing the same $SU(3)_c \times SU(2)_L \times U(1)_Y \equiv G_{\text{SM}}$ quantum numbers as each SM fermion. This implies the composite symmetry group G must also contain a colour group $SU(3)$ under which the composite operators coupling to the elementary quarks transform as triplets. In general, composite fields must transform in representations of G that decompose under G_{SM} to give representations of SM fermions. Suddenly there is not so much freedom in the composite sector!

Let us consider Equation (3.88) in more detail, focussing at first on only the left-handed elementary top quark t_L^0 . (We know the elementary fields cannot be the SM fields so we distinguish them with a superscript 0). From above, there should be a composite Dirac fermion T partnered to t_L^0 . We shall assume for simplicity that the composite operator \mathcal{O}_{t_L} is exactly this field. Then the mass-mixing terms of the elementary left-handed top quark will be

$$\begin{aligned} \mathcal{L}_{t_L} &= -m_L \bar{T}T + \Delta_{t_L} \bar{t}_L^0 T + \Delta_{t_L}^\dagger \bar{T} t_L^0 \\ &= -(\bar{t}_L^0 \quad \bar{T}_L \quad \bar{T}_R) \begin{pmatrix} 0 & 0 & -\Delta_{t_L} \\ 0 & 0 & m_L \\ -\Delta_{t_L}^\dagger & m_L & 0 \end{pmatrix} \begin{pmatrix} t_L^0 \\ T_L \\ T_R \end{pmatrix}. \end{aligned} \quad (3.90)$$

The composite partner is given a Dirac mass m_L consistent with the symmetries. A field with such a mass term is known as a “vector-like” fermion. In this simple example a field redefinition may be performed to make Δ_{t_L} real. The mass basis is then found to be

$$t_L = \frac{m_L t_L^0 + \Delta_{t_L} T_L}{\sqrt{m_L^2 + \Delta_{t_L}^2}}, \quad T'_L = \frac{m_L T_L - \Delta_{t_L} t_L^0}{\sqrt{m_L^2 + \Delta_{t_L}^2}}, \quad T'_R = T_R, \quad (3.91)$$

giving the Lagrangian

$$\mathcal{L}_{t_L} = -(\bar{t}_L \quad \bar{T}') \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{m_L^2 + \Delta_{t_L}^2} \end{pmatrix} \begin{pmatrix} t_L \\ T' \end{pmatrix}. \quad (3.92)$$

The same analysis can be done for the right-handed elementary top quark, having its own composite partner \tilde{T} with Dirac mass m_R and coupling Δ_{t_R} . We see there is a massless field $t = t_L + t_R$, and two composite partners T' and \tilde{T}' with respective masses $M_{L,R} = \sqrt{m_{L,R}^2 + \Delta_{t_{L,R}}^2}$. We identify t as the SM top quark, expecting its mass to arise after EWSB.

It is seen from Equation (3.91) that the physical particle $t_{L,R}$ is a superposition of the elementary and composite fields, being mostly elementary if $M_{L,R}$ is much larger than the mixing $\Delta_{t_{L,R}}$, and mostly composite if the mixing is much larger. The dimensionless number $\Delta_{t_{L,R}}/M_{L,R}$ is often used to measure the *compositeness* of $t_{L,R}$. Of course these considerations also apply to the rest of the SM fermions, which have their own partners and mixings. But note that in realistic models, the elementary fermions do not couple to just one single composite field each, making the mixings significantly more complicated.

Masses for the SM particles arise from their composite components' interactions with the composite Higgs. To illustrate the general idea, imagine the composite sector has a Yukawa-like coupling to the Higgs such as

$$\mathcal{L} \supset -Y_t \bar{T}_L h \tilde{T}_R. \quad (3.93)$$

Again, things are not so simple in realistic models (not least because a pNGB Higgs cannot interact in this way), but the details are not important. By inverting Equation (3.91) to express the composite fields in terms of the mass basis, we get the top-Higgs coupling

$$\mathcal{L} \supset -Y_t \frac{\Delta_{t_L}}{M_L} \frac{\Delta_{t_R}}{M_R} \bar{t}_L h t_R, \quad (3.94)$$

yielding the top quark Yukawa coupling

$$y_t = Y_t \frac{\Delta_{t_L}}{M_L} \frac{\Delta_{t_R}}{M_R}, \quad (3.95)$$




Figure 3.4: Generation of Yukawa couplings for the top quark. The large circle in the middle represents unknown composite dynamics. Once the Higgs develops a vev, the top quark will gain a mass through this diagram. The compositeness of the left- and right-handed components obviously dictate the amplitude of this process.

and similarly for the other fermions. We see that Y_t acts as a sort of “proto-Yukawa” coupling for the top quark. Also note that an SM fermion mass is dependent on the compositeness of its left- and right-handed components. More composite particles are more massive, and less composite particles are less massive. This qualitative relation should be expected in general, as Figure 3.4 makes clear. Lighter fermions are then expected to be mostly elementary (at least as long as their different chiral components share a similar degree of compositeness), while the top quark must be significantly composite because of its large Yukawa coupling. The top quark, and to a lesser degree the bottom quark and tau lepton, offer the most promising portals into the composite sector.

3.6.1 Composite Fermion Representations in the MCHM

To move into more concrete territory, let us discuss the possible composite fermionic operators $\mathcal{O}_{L,R}$ in the Minimal CHM. Recall that the MCHM is based on the symmetry breaking pattern $SO(5) \rightarrow SO(4)$. The elementary-composite couplings are to exist at high energies where the $SO(5)$ symmetry is unbroken, so the composite operators must furnish representations of $SO(5)$ to reflect this symmetry in the Lagrangian. Our goal, then, is to explore the representations of $SO(5)$ and find how each decomposes under $SU(2)_L \times U(1)_Y = G_{EW}$ in order to match with the representations of SM fermions. We will only focus on partnering the third generation quarks to the composite sector, for the other fermions, being light, are expected to mix only weakly.

For background on different types of representations, see Appendix A. We will be interested in the irreducible representations of $SO(5)$, for these are the ones that are expected to have small mass dimensions and so give the least suppressed operators at the energy scales we are capable of reaching. For the same reason, we will only be concerned with the lowest-dimensional irreducible representations, among which

are the trivial **1**, the spinorial **4**, the fundamental **5**, the antisymmetric **10**, and the symmetric traceless **14**.

The Fundamental 5

We start with the most familiar representation: the fundamental **5**. To make the notation more familiar and explicit, let us write a fermionic operator transforming in this representation as $\Psi_{\mathbf{5}}$. This is most easily written as a fiveplet

$$\Psi_{\mathbf{5}} = (\Psi_1, \dots, \Psi_5)^\top \quad (3.96)$$

that transforms under $SO(5)$ by regular matrix multiplication:

$$\Psi_{\mathbf{5}} \xrightarrow{g_5 \in SO(5)} g_5 \Psi_{\mathbf{5}}. \quad (3.97)$$

How does this decompose under G_{EW} ? We break the answer down into multiple steps, overexplaining in this simple case so that the procedure is clear for less straightforward representations.

First, we decompose the **5** under $SO(4)$, since G_{EW} is embedded into this subgroup. With the convention for $SO(5)$ generators we use throughout this work (given in Appendix G), $SO(4)$ is embedded into the top left 4×4 block of $SO(5)$ transformations. That is, a transformation $g_4 \in SO(4)$ is enacted via the transformation $g_5 \in SO(5)$ given by

$$g_5 = \begin{pmatrix} g_4 & \vec{0} \\ \vec{0}^\top & 1 \end{pmatrix}. \quad (3.98)$$

Under such an $SO(4)$ transformation, the fiveplet transforms as

$$\Psi_{\mathbf{5}} \xrightarrow{g_4 \in SO(4)} \begin{pmatrix} g_4 & \vec{0} \\ \vec{0}^\top & 1 \end{pmatrix} \Psi_{\mathbf{5}} = \begin{pmatrix} g_4 \Psi_{\mathbf{4}} \\ \Psi_5 \end{pmatrix}, \quad (3.99)$$

where $\Psi_{\mathbf{4}} = (\Psi_1, \dots, \Psi_4)$ are the first four components of $\Psi_{\mathbf{5}}$. Evidently, $\Psi_{\mathbf{4}}$ is in the fundamental **4** of $SO(4)$, while Ψ_5 is a singlet (does not transform). This shows that the **5** decomposes under $SO(4)$ as

$$\mathbf{5} \rightarrow \mathbf{4} \oplus \mathbf{1}. \quad (3.100)$$

The second step is to decompose these $SO(4)$ representations under $G_{EW} \subset SO(4)$. As an intermediate step we write the $SO(4)$ representations in the language of $SU(2)_L \times SU(2)_R$. The singlet representation is trivial: it remains a singlet. And from Appendix D, we know that the **4** is nothing but the bidoublet $(\mathbf{2}, \mathbf{2})$ of $SU(2)_L \times SU(2)_R$. Tracing the isomorphisms of Appendix D backwards, $\Psi_{\mathbf{4}}$ can be written in

terms of fields Ψ^{n_L, n_R} having $SU(2)_{L,R}$ quantum numbers $n_{L,R}$ as

$$\Psi_4 = \frac{1}{\sqrt{2}} \begin{pmatrix} i\Psi^{-,-} - i\Psi^{+,+} \\ \Psi^{-,-} + \Psi^{+,+} \\ i\Psi^{+,-} - \Psi^{-,+} \\ i\Psi^{-,+} - \Psi^{+,-} \end{pmatrix}, \quad (3.101)$$

where \pm superscripts denote $\pm 1/2$. Essentially, Ψ_4 is here being expressed as a linear combination of eigenvectors of $T_{L,R}^3$ with eigenvalues $n_{L,R}$, and Ψ^{n_L, n_R} are the coefficients¹⁵. These fields are a useful basis, for their G_{EW} quantum numbers can easily be read off from their superscripts. Recall from Section 3.5.1 that the hypercharge is embedded such that $Y = T_R^3$, so a field Ψ^{n_L, n_R} has hypercharge n_R .

Evidently the **4** decomposes into $SU(2)_L$ doublets with hypercharges $\pm 1/2$. In total, then, the decomposition is

$$\begin{array}{ccccccc} SO(5) & & SO(4) & & SU(2)_L \times SU(2)_R & & G_{EW} \\ \mathbf{5} & \rightarrow & \mathbf{4} \oplus \mathbf{1} & \cong & (\mathbf{2}, \mathbf{2}) \oplus (\mathbf{1}, \mathbf{1}) & \rightarrow & \mathbf{2}_{-\frac{1}{2}} \oplus \mathbf{2}_{+\frac{1}{2}} \oplus \mathbf{1}_0. \end{array} \quad (3.102)$$

The representation of $U(1)_Y$ is typically written as a subscript.

Here we hit a problem: none of these G_{EW} representations match the representation of any SM quark (see Table 2.2). Elementary quarks therefore cannot couple to operators in the **5** of $SO(5)$, lest they explicitly break electroweak symmetry. It is a similar story for other representations of $SO(5)$. Fortunately, this can be salvaged by extending the global $SO(5)$ symmetry to $SO(5) \times U(1)_X$ for some new quantum number X , and redefining the hypercharge so that

$$Y = T_R^3 + X. \quad (3.103)$$

This merely has the effect of adding X to all the hypercharges in Equation (3.102). All of the representations of the various quarks can then be fit into two fiveplets, one with an X charge of $+2/3$ and the other with $X = -1/3$:

$$\begin{array}{ccc} SO(5) \times U(1)_X & & G_{EW} \\ \mathbf{5}_{+\frac{2}{3}} & \rightarrow & \underbrace{\mathbf{2}_{+\frac{1}{6}} \oplus \mathbf{2}_{+\frac{7}{6}}}_{q_L} \oplus \underbrace{\mathbf{1}_{+\frac{2}{3}}}_{u_R} \\ \mathbf{5}_{-\frac{1}{3}} & \rightarrow & \mathbf{2}_{-\frac{5}{6}} \oplus \underbrace{\mathbf{2}_{+\frac{1}{6}}}_{q_L} \oplus \underbrace{\mathbf{1}_{-\frac{1}{3}}}_{d_R}. \end{array} \quad (3.104)$$

Notice that since the electric charge is determined by Equation (2.39)

$$Q = T_L^3 + Y = n_L + n_R + X, \quad (3.105)$$

¹⁵The bidoublet representation uses generators in the fundamental of $SU(2)$, which have eigenvalues $\pm 1/2$. These eigenvalues are better recognised as the possible spin projection values of a spin-1/2 particle.

coupling the elementary quarks to a composite multiplet in the **5** representation necessitates composite resonances with exotic charges $Q = 4/3$ and $Q = 5/3$, along with a few extra resonances with the same charges as the quarks.

This shows how the particle content of the composite sector may be found through analysis of the $SO(5)$ representations. Before moving on to other representations, let us demonstrate how partial compositeness works in practice using operators in the **5** as an example. For the sake of simplicity, we focus on only partnering the elementary top quark with the composite sector.

Based on Equation (3.104), we could partner the left-handed third generation quark doublet $q_L^0 = (t_L^0, b_L^0)^\top$ to a composite multiplet $\Psi^q \sim \mathbf{5}_{+2/3}$, and partner the right-handed top quark t_R^0 to another multiplet $\tilde{\Psi}^t \sim \mathbf{5}_{+2/3}$. The composite multiplets consist of all the fields of Equations (3.96) and (3.101). The elementary fields will be embedded into respective spurions (*cf.* Section 3.3.1) ψ_L^q and ψ_R^t in the same representations as their composite partners for ease in writing electroweak symmetric interactions. The elementary-composite quark mixing terms will be

$$\mathcal{L}_{\text{mix}} = \Delta_q \bar{\psi}_L^q \Psi^q + \Delta_t \bar{\psi}_R^t \tilde{\Psi}^t + \text{h.c.}, \quad (3.106)$$

for these are invariant under $SO(5)$. But once the spurions take on their background values

$$\psi_L^q = \frac{1}{\sqrt{2}} \begin{pmatrix} ib_L^0 \\ b_L^0 \\ it_L^0 \\ -t_L^0 \\ 0 \end{pmatrix}, \quad \psi_R^t = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ t_R^0 \end{pmatrix}, \quad (3.107)$$

(matching the quarks' quantum numbers to the fields in Equation (3.101)), the $SO(5)$ symmetry will be explicitly broken. This demonstration should not be taken too literally; it will turn out that a workable theory will need two different $SO(5)$ symmetries acting separately on the elementary and composite sector, and NGBs that transform under both of those symmetries are needed to “link” the elementary and composite multiplets together into invariant terms. Equation (3.106) nevertheless shows the basic mechanics of enacting partial compositeness.

The Traceless Symmetric **14** and Antisymmetric **10**

Next we consider the traceless symmetric **14** and antisymmetric **10** representations of $SO(5)$, but we will not go in to the same level of detail as for the **5**. Multiplets in the **14** or **10** can be embedded into traceless 5×5 matrices that are respectively symmetric (Ψ_+) and antisymmetric (Ψ_-). We consider these representations together because they both transform adjointly:

$$\Psi_{\pm} \xrightarrow{g_5 \in SO(5)} g_5 \Psi_{\pm} g_5^{-1}. \quad (3.108)$$

Indeed, the **10** is the adjoint representation of $SO(5)$.

Following the previous example, we first work out the decompositions under $SO(4)$. To do this, let us write Ψ_{\pm} in block form for convenience:

$$\Psi_{\pm} = \left(\begin{array}{c|c} A & \vec{b} \\ \hline \pm \vec{b}^{\top} & c \end{array} \right), \quad (3.109)$$

where A is a 4×4 matrix. Under an $SO(4)$ transformation g_4 , this goes to

$$\left(\begin{array}{c|c} A & \vec{b} \\ \hline \pm \vec{b}^{\top} & c \end{array} \right) \rightarrow \left(\begin{array}{c|c} g_4 A g_4^{-1} & g_4 \vec{b} \\ \hline \pm (g_4 \vec{b})^{\top} & c \end{array} \right). \quad (3.110)$$

So A transforms adjointly under $SO(4)$, \vec{b} transforms in the fundamental **4**, and c is a singlet. Given A is symmetric or antisymmetric, it will furnish the symmetric **9** or antisymmetric¹⁶ **6** of $SO(4)$. We therefore have the decompositions

$$\begin{array}{ccc} SO(5) & & SO(4) \\ \mathbf{10} & \rightarrow & \mathbf{6} \oplus \mathbf{4}, \\ \mathbf{14} & \rightarrow & \mathbf{9} \oplus \mathbf{4} \oplus \mathbf{1}, \end{array} \quad (3.111)$$

where **10** does not contain a singlet because its singlet component vanishes. We see that both representations contain the **4** of $SO(4)$ that can house the left-handed quark doublets as before. As for the unfamiliar representations, since $SO(4)$ is locally isomorphic to $SU(2)_L \times SU(2)_R$, its adjoint is the direct sum of the adjoints of each group factor: $\mathbf{6} \cong (\mathbf{3}, \mathbf{1}) \oplus (\mathbf{1}, \mathbf{3})$. We shall see this explicitly in Chapter 4. Similarly, the **9** of $SO(4)$ is the symmetric two-index tensor and so translates as $\mathbf{9} \cong (\mathbf{3}, \mathbf{3})$. Eigenvalues of the $SU(2)$ generators in the adjoint representation are $-1, 0$, and $+1$ (the possible spin projections of a spin-1 particle), so these decompose to

$$\begin{array}{ccc} SU(2)_L \times SU(2)_R & & SU(2)_L \times U(1)_Y \\ (\mathbf{3}, \mathbf{1}) & \rightarrow & \mathbf{3}_0 \\ (\mathbf{1}, \mathbf{3}) & \rightarrow & \mathbf{1}_{-1} \oplus \mathbf{1}_0 \oplus \mathbf{1}_{+1} \\ (\mathbf{3}, \mathbf{3}) & \rightarrow & \mathbf{3}_{-1} \oplus \mathbf{3}_0 \oplus \mathbf{3}_{+1}. \end{array} \quad (3.112)$$

Note that if Ψ_{\pm} is given an X charge of $+2/3$ then all of the SM quark representations are present, except for that of the right-handed down-types in the **14**. With this X charge, the **14** and **10** both deliver (electric) charge $8/3$ particles in addition to the exotic particles of the **5**.

¹⁶These labels just come from the dimensions of (anti)symmetric 4×4 matrices.

Expressing the matrices Ψ_{\pm} in terms of fields Ψ^{n_L, n_R} with definite $SU(2)_{L,R}$ quantum numbers is rather unwieldy, so these are relegated to Appendix G.2.

The Spinorial 4

The last low-dimensional representation of $SO(5)$ we discuss is the spinorial **4**. But we will not give many details, for coupling quarks to this representation turns out to be phenomenologically unviable. In short, the problem stems from the $Z\bar{b}_L b_L$ coupling, which has been measured in agreement with the SM prediction to decent precision, receiving inadmissably large corrections from operators in this representation. The Z coupling to a fermion ψ is in general given by

$$\left[\sqrt{g^2 + g'^2} Q_L^3 - g' Q \right] Z_{\mu} \bar{\psi} \gamma^{\mu} \psi, \quad (3.113)$$

where Q is the electric charge and Q_L^3 is some other operator. In the SM this operator is T_L^3 , but it may be different in CHMs due to partial compositeness, though it so happens that there are no corrections to Q_L^3 if the fermion couples to a representation that satisfies $T_L^3 = T_R^3$ [80]. Since **4** decomposes as

$$\begin{aligned} SO(5) & \quad SU(2)_L \times SU(2)_R \\ \mathbf{4} & \quad \rightarrow \quad (\mathbf{2}, \mathbf{1}) \oplus (\mathbf{1}, \mathbf{2}), \end{aligned} \quad (3.114)$$

there is no way to embed q_L in a **4** in a way that protects the $Z\bar{b}_L b_L$ coupling. Because of this, we do not consider the spinorial representation any further. Note the representations we previously discussed *do* protect this coupling, although they do not guarantee protection for the $Z\bar{t}_L t_L$ coupling, which has not yet been precisely measured.

3.6.2 $SO(4)$ Invariants

Having gone through Section 3.6.1, one might be under the impression that the purpose of the composite $SO(5)$ multiplets is to construct an $SO(5)$ -invariant Lagrangian. This is not so. Remember the moral of the CCWZ construction: at the low energies we are interested in, our Lagrangian need only be linearly invariant under $SO(4)$, not $SO(5)$. It is only by realising the linear $SO(4)$ symmetry as a special case of a non-linear representation according to the CCWZ prescription that the $SO(5)$ invariance of the Lagrangian becomes manifest. The $SO(5)$ multiplets simply provide us with the operators with which we may build $SO(4)$ -invariant terms for the Lagrangian.

To illustrate this, we need to reverse the arguments used in the CCWZ construction. Consider, for example, an $SO(5)$ multiplet $\Psi_{\mathbf{5}} \sim \mathbf{5}$. From what we have seen in Section 3.4.3, this must be nothing but a multiplet $\Psi_{SO(4)}$ that transforms under some representation of $SO(4)$, dressed by the NGB matrix U . The fields

transforming under $SO(4)$ that we desire, then, are¹⁷

$$\Psi_{SO(4)} = U^{-1}\Psi_5. \quad (3.115)$$

In this case $\Psi_{SO(4)}$ consists of a fourplet Ψ_4 in the **4** and a singlet Ψ_1 in the **1** of $SO(4)$:

$$\Psi_{SO(4)} = (\Psi_4, \Psi_1)^\top. \quad (3.116)$$

It is these $SO(4)$ multiplets $\Psi_{1,4}$ that belong in the Lagrangian. Leading-order terms will be contractions of two fields in the same representation such as

$$\bar{\Psi}_4\Psi_4, \quad \bar{\Psi}_1\Psi_1, \quad \bar{\Psi}_4\tilde{\Psi}_4, \quad \bar{\Psi}_1\tilde{\Psi}_1, \quad \dots, \quad (3.117)$$

having introduced another set of fields distinguished by tildes. In principle each of these operators may enter the Lagrangian with independent coupling constants, unless there is an extra symmetry preventing it.

Of course it is not too enlightening to just write Equation (3.117); we would like to know what these invariants contain. In particular we are most interested in the invariants showing how the SM fields enter the Lagrangian. Let us consider the case where Ψ_5 and $\tilde{\Psi}_5$ are the spurions of q_L and t_R defined by Equation (3.107). The mixing terms between q_L and t_R are then the last two terms in Equation (3.117). Because of the orthogonality of the spurions

$$0 = \bar{\Psi}_5\tilde{\Psi}_5 = \bar{\Psi}_4\tilde{\Psi}_4 + \bar{\Psi}_1\tilde{\Psi}_1, \quad (3.118)$$

there is actually only one invariant mixing term up to a factor. An explicit calculation of the invariant using the NGB matrix Equation (3.11) in the unitary gauge yields

$$\begin{aligned} \bar{\Psi}_1\tilde{\Psi}_1 &= (\bar{\Psi}_5 U)_5 (U^{-1}\tilde{\Psi}_5)_5 \\ &= \left(-\sin\left(\frac{h}{f}\right) \frac{\bar{t}_L}{\sqrt{2}} \right) \left(\cos\left(\frac{h}{f}\right) t_R \right). \end{aligned} \quad (3.119)$$

We see Yukawa-like interactions with the top quark involving an arbitrarily large number of Higgs insertions. This operator will typically be accompanied by a coefficient with a factor of f , so the limit $f \rightarrow \infty$ will reproduce the SM top Yukawa coupling. Of course if the composite partners of other fermions are also included, then there will be other sorts of mixings that can lead to interesting effects.

It should be clear from this analysis how to find the $SO(4)$ invariants for other representations. All it requires is writing the $SO(5)$ multiplets as $SO(4)$ multiplets dressed by the NGB matrix, extracting the (sub)multiplets that transform under irreducible representations of $SO(4)$, and contracting them together. See Ref. [69]

¹⁷Recalling the NGB matrix transformation rule Equation (3.47), $\Psi_{SO(4)}$ will transform as $\Psi_{SO(4)} \rightarrow (gUh^{-1})^{-1}(g\Psi_5) = h\Psi_{SO(4)}$, where $g \in SO(5)$ and $h \in SO(4)$.

for a comprehensive analysis of invariants resulting from the other $SO(5)$ representations we have discussed.

Despite drawing attention to $SO(4)$ invariants here, the point will turn out to be moot for the models we will consider. In our models, we will have extra symmetries that force the $SO(4)$ invariants to come from $SO(5)$ invariants, so we will only need to construct Lagrangians with manifest $SO(5)$ symmetry anyway. The reason we bring attention to the issue is because the possible invariants arising from a given representation determine the structure of the Higgs effective potential, and some representations will turn out to be more theoretically favourable as a result.

3.7 Coleman-Weinberg Potential

In the SM, the Higgs boson breaks electroweak symmetry because it is given a classical potential

$$\mathcal{V}(\Phi) = -\mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \quad (3.120)$$

that is minimised for $\Phi_0 \neq 0$. CHMs cannot take the same approach of course, for the Higgs boson, being a pNGB, has no classical potential. Instead, the quantum corrections for the effective potential must be relied upon to give the Higgs doublet a non-zero vev. Here we show how these quantum corrections may be calculated to leading order, resulting in the so-called Coleman-Weinberg potential [81] that is an essential tool in the analysis of CHMs.

Sufficient background (for our purposes) on the quantum effective potential is provided in Appendix C. The upshot of it all is that the effective potential is a function of classical scalar fields $\Phi_c = (\phi_c^1, \phi_c^2, \dots)$, among which will be the Higgs field, given by

$$V_{\text{eff}}(\Phi_c) = - \sum_{N=1}^{\infty} \frac{1}{N!} \Gamma^{(N)}(0, \dots, 0) (\Phi_c(x))^N. \quad (3.121)$$

Here, $\Gamma^{(N)}(0, \dots, 0)$ is the sum of all one-particle-irreducible Feynman diagrams (those that cannot be disconnected by removing a propagator) with N external lines of scalar fields that have vanishing momenta. It should be understood that $(\Phi_c(x))^N$ is shorthand for the product of any N scalar fields, with all such products being included. The tree-level diagrams just reproduce the classical potential $\mathcal{V}(\Phi_c)$. Diagrammatically, the calculation of V_{eff} is represented in Figure 3.5.

The diagrams contributing to V_{eff} in Figure 3.5 have been ordered suggestively, showing two possible methods for approximating the effective potential. The first is an expansion in the number of loops, and the second is an expansion in the order of coupling constants. Coleman and Weinberg argue in their classic paper [81] that expansion in terms of the number of loops is at least as good as expansion in the coupling constants, with the added benefit of being unaffected by field shifts

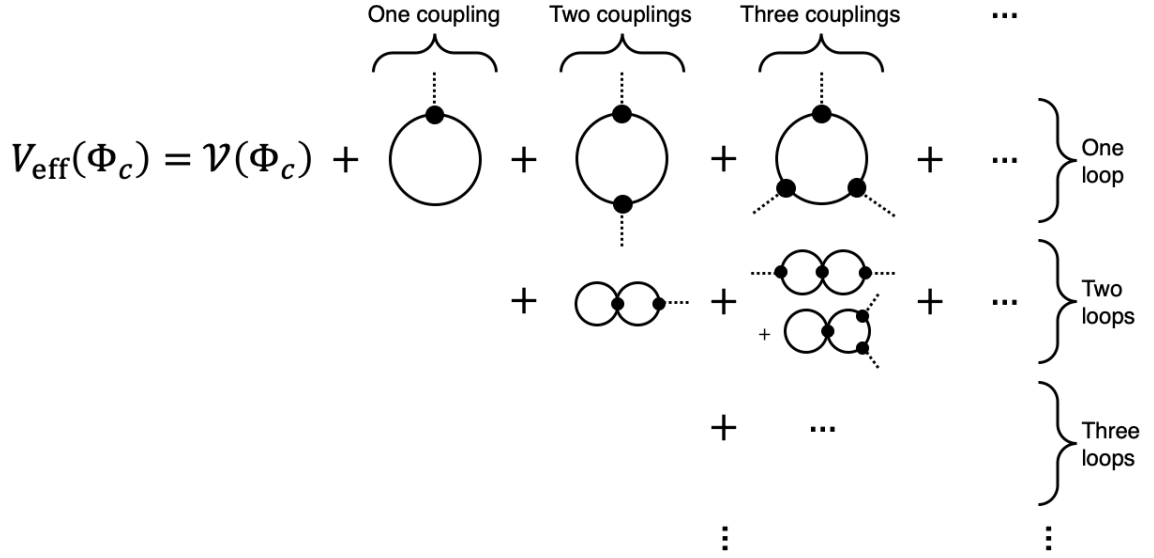


Figure 3.5: Expansion of the effective potential for scalar fields Φ_c in terms of one-particle-irreducible Feynman diagrams. Here the dashed lines are the classical scalar fields, and the internal solid lines can be any scalar, fermionic, or gauge particle that couples to Φ_c . All possible such diagrams are to be included. The external lines have zero momentum and may represent one or more scalars depending on the various interaction terms.

$\Phi(x) \rightarrow h(x) = \Phi(x) - \Phi_0$. The loop expansion is therefore particularly well-suited for studying Higgs phenomena.

Functional methods can be used to calculate V_{eff} to any number of loops [82], but this technique is too difficult to apply in realistic interacting theories. Instead, we make do with calculating the effective potential to one-loop level - that is, using only the first row of diagrams in Figure 3.5. It can be shown that the one-loop diagrams provide contributions to V_{eff} that are next-to-leading-order in \hbar compared to \mathcal{V} , with higher-loop diagrams contributing at higher powers of \hbar [83]. We therefore assume the one-loop truncation is sufficiently accurate for our purposes.

Fortunately, the one-loop effective potential has been calculated explicitly by Coleman and Weinberg [81] to have the simple form

$$V_{\text{eff}}^{1\text{-loop}}(\Phi_c) = \mathcal{V}(\Phi_c) + \sum_{\text{particles } i} \frac{c_i}{64\pi^2} \left(2m_i^2(\Phi_c)\Lambda^2 + m_i^4(\Phi_c) \log \left(\frac{m_i^2(\Phi_c)}{\Lambda^2} \right) \right). \quad (3.122)$$

Here, Λ is the high-energy cutoff of the loop momentum, $m_i(\Phi_c)$ is the Φ_c -dependent tree-level mass of particle i in the mass basis, and c_i is a combinatorial factor count-




Figure 3.6: Possible vertices entering the one-loop diagrams contributing to the effective potential in Figure 3.5: (a) for internal scalar lines (b) for internal fermion lines (c) for internal gauge boson lines. The vertex factors \tilde{M} are different for the different cases.

ing the degrees of freedom of particle i :

$$c_i = \begin{cases} 3 & \text{for neutral gauge bosons,} \\ 6 & \text{for charged gauge bosons,} \\ -12 & \text{for coloured Dirac fermions.} \end{cases} \quad (3.123)$$

The negative sign for fermions comes from the usual negative sign when evaluating fermion loops. In the models we are considering, we shall see that all dependence on $\Lambda \sim m_*$ drops out of the potential.

Since Equation (3.122) is central to this work, we reproduce here the original calculation in Ref. [81]. The starting point is to notice that all the one-loop diagrams in Figure 3.5 contain scalar couplings to exactly two internal lines; any more could not produce a one-loop diagram. There are then three types of interactions, shown in Figure 3.6, to consider:

- Interactions with scalar fields ϕ^a, ϕ^b . Let us collect all the interactions between these two fields and other scalar fields Φ_c into the term $-\tilde{M}_{ab}^2(\Phi_c)\phi^a\phi^b$. Each vertex will then carry the numerical factor $\tilde{M}_{ab}^2(\Phi_c)$. Note this is distinct from the field-independent mass term $-M_{ab}^2\phi^a\phi^b$.




Figure 3.7: Example of a one-loop diagram with $N = 5$ interaction points. Each interaction is as in Figure 3.6.

- Interactions with fermions ψ^a, ψ^b . Their interactions with scalar fields will be written in the form $-\tilde{M}_{ab}(\Phi_c)\bar{\psi}^a\psi^b$, so the vertex in Figure 3.6 has the numerical factor $\tilde{M}_{ab}(\Phi_c)$. (The mass-like matrices \tilde{M} are of course different for the different types of interactions, but we use the same letter for convenience).
- Interactions with gauge fields A_μ^a, A_μ^b . These will have vertex factors $\tilde{M}_{ab}^2(\Phi)$ (just as for the scalars) if the gauge-scalar interactions are written in the form $+\frac{1}{2}\tilde{M}_{ab}^2(\Phi_c)A_\mu^a A^{b\mu}$.

The contribution to the effective potential can be calculated in more or less the same way for each type of interaction. We proceed with the calculation of the scalar contribution, whose logic carries over to the gauge boson contribution also¹⁸.

Consider a one-scalar-loop diagram with $N \geq 1$ interaction points, as in Figure 3.7. This will contribute to the effective potential an amount

$$i \frac{1}{2N} \int \frac{d^4k}{(2\pi)^4} \text{Tr} \left[\frac{(\tilde{M}^2(\Phi_c))^N}{(k^2 - M^2 + i\varepsilon)^N} \right], \quad (3.124)$$

having taken into account the $2N$ symmetries (rotations and reflections) of the diagram¹⁹. The overall factor of i ultimately comes from the factor of i in the definition of the Schwinger functional (see Appendix C). The $1/N_E!$ factor of Equation (C.13), where N_E is the number of external lines, has been cancelled by the $N_E!$ ways the external lines can be arranged.

¹⁸For gauge boson loops the vertices and propagators also have Lorentz indices that must be contracted, but these make little difference since they project onto the transverse components, giving integrals just like the those for the scalars.

¹⁹The diagram is a regular N -gon, whose symmetries form the dihedral group D_N of order $2N$.

Recognising the Taylor series $\ln(1-x) = -\sum_{N=1}^{\infty} \frac{x^N}{N}$, summing all such diagrams gives

$$\begin{aligned} V_{\text{scalar}}^{1\text{-loop}}(\Phi_c) &= -\frac{i}{2} \int \frac{d^4k}{(2\pi)^4} \text{Tr} \left[\ln \left(1 - \frac{\tilde{M}^2(\Phi_c)}{k^2 - M^2 + i\varepsilon} \right) \right] \\ &= -\frac{i}{2} \int \frac{d^4k}{(2\pi)^4} \text{Tr} \left[\ln \left(1 - \frac{M^2 + \tilde{M}^2(\Phi_c)}{k^2 + i\varepsilon} \right) - \ln \left(1 - \frac{M^2}{k^2 + i\varepsilon} \right) \right]. \end{aligned} \quad (3.125)$$

The second term in the trace is independent of Φ_c , so it has no effect and is safe to ignore. The first term contains the quantity $M^2 + \tilde{M}^2(\Phi_c)$, which we recognise as the tree-level mass (squared) matrix of the scalars. Its trace (the sum of its eigenvalues) is the sum of the squared tree-level scalar masses, making the effective potential

$$V_{\text{scalar}}^{1\text{-loop}}(\Phi_c) = -\frac{i}{2} \sum_{\text{scalars } i} \int \frac{d^4k}{(2\pi)^4} \ln \left(1 - \frac{m_i^2(\Phi_c)}{k^2 + i\varepsilon} \right). \quad (3.126)$$

This can be calculated by Wick rotating the momentum to Euclidean space, taking $k^2 \rightarrow -k_E^2$ and $d^4k \rightarrow id^4k_E$. The resulting integral is spherically symmetric, picking up a 4D solid angle of $2\pi^2$:

$$\begin{aligned} V_{\text{scalar}}^{1\text{-loop}}(\Phi_c) &= \frac{1}{2} \sum_{\text{scalars } i} \int_0^{\Lambda^2} \frac{dk_E^2}{16\pi^2} k_E^2 \ln \left(1 + \frac{m_i^2(\Phi_c)}{k_E^2} \right) \\ &= \frac{1}{64\pi^2} \sum_{\text{scalars } i} \left[k_E^4 \ln \left(1 + \frac{m_i^2}{k_E^2} \right) + k_E^2 m_i^2 - m_i^4 \ln (m_i^2 + k_E^2) \right]_{k_E^2=0}^{\Lambda^2} \\ &= \frac{1}{64\pi^2} \sum_{\text{scalars } i} \left[\Lambda^4 \ln \left(1 + \frac{m_i^2}{\Lambda^2} \right) + \Lambda^2 m_i^2 + m_i^4 \ln \left(1 + \frac{\Lambda^2}{m_i^2} \right) \right]. \end{aligned} \quad (3.127)$$

Since the integral is ultraviolet divergent, a high-energy cutoff Λ has been introduced. The limit $\Lambda \gg m_i(\Phi_c)$ reproduces the form of Equation (3.122), as desired.

Fermionic contributions are calculated largely the same way. The slight difference comes from the propagators, giving integrands of the form

$$-\text{Tr} \left[\dots \tilde{M} \frac{\gamma^\mu k_\mu}{k^2} \tilde{M} \frac{\gamma^\nu k_\nu}{k^2} \dots \right] = -\text{Tr} \left[\dots \frac{1}{k^2} \tilde{M} \tilde{M}^\dagger \dots \right]. \quad (3.128)$$

Only diagrams with even numbers of fermions contribute because the trace of the product of an odd number of gamma matrices vanishes, so all mass-like matrices in the integrand can be paired in this way. There is no reflection symmetry in the diagrams anymore because fermion lines are directed, but the factor of 1/2 is compensated for by only summing the graphs with even numbers of fermions. The logic from the scalar calculation then follows through, identifying \tilde{M}^2 in that

case with $\tilde{M}\tilde{M}^\dagger$ in this case, because the eigenvalues of $\tilde{M}\tilde{M}^\dagger$ are indeed the field-dependent contributions to the squared fermion masses. This demonstrates the validity of Equation (3.122).

We should warn that in the literature of CHMs, the effective potential is usually not expressed in the form of Equation (3.122) but rather is expressed in terms of gauge boson and fermion form factors. We avoid the form factor viewpoint here because it adds a layer of complication that is unnecessary at present, although we do introduce it in Section 4.5, as we use this method in some of our numerical scans later. But no matter the approach of calculating it, the effective potential in a given model is a horrendously complicated function of the Lagrangian parameters that has no hope of being expressed analytically. This fact makes it difficult to study the parameter spaces of CHMs, and is a major reason why global fits of such models have not been attempted in detail before this work.

3.7.1 Higgs Potential in the MCHM

Despite the uncooperative nature of the Higgs potential, symmetry considerations can be used to reason about its form in a given model. To demonstrate, we focus once again on the MCHM, where Φ_c is simply the Higgs field h and all Higgs dependence is through the quantity

$$s_h := \sin \frac{h}{f}. \quad (3.129)$$

The potential can therefore generically be expanded as²⁰

$$V_{\text{eff}}(h) = -\gamma s_h^2 + \beta s_h^4 + \mathcal{O}(s_h^6) \quad (3.130)$$

for some model-dependent coefficients γ and β . Being interested in only the low-energy region $s_h \ll 1$, we truncate the potential at order s_h^4 . Reproducing the EWSB scale enforces the constraint

$$\frac{\partial V_{\text{eff}}(h)}{\partial h} \Big|_{h=\langle h \rangle} = 0 \quad \implies \quad \frac{\gamma}{2\beta} = s_{\langle h \rangle}^2 \equiv \frac{v^2}{f^2} =: \xi, \quad (3.131)$$

and similarly for the Higgs mass,

$$m_h^2 = \frac{\partial^2 V_{\text{eff}}(h)}{\partial h^2} \Big|_{h=\langle h \rangle} = \frac{8\xi(1-\xi)\beta}{f^2}. \quad (3.132)$$

With these relations, the potential takes the equivalent form

$$V_{\text{eff}}(h) = \frac{m_h^2 v^2}{8\xi^2(1-\xi)} (s_h^2 - \xi)^2. \quad (3.133)$$

²⁰The spinorial $\mathbf{4}$ representation introduces $\sin(h/2f)$ terms into the potential, but we do not consider this case here.

As with the Higgs interactions we have seen previously, the potential is completely specified up to some unknown parameter ξ . But this simplification of the potential is not the main benefit of Equations (3.131) and (3.132); these relations instead can tell us how finely-tuned the potential must be to result in the observed EWSB scale and Higgs mass. The whole point of considering CHMs, recall, is to obtain a theory that is less finely-tuned than the SM. If CHMs turn out to have the same fine-tuning problems, then they are not well motivated and would offer little theoretical benefit over the SM.

Fine-Tuning in the MCHM

To quantify the fine-tuning of a given model we use the Barbieri-Giudice (BG) measure [84], defined as the maximum relative sensitivity of a given observable \mathcal{O} across all input parameters x_i

$$\Delta_{\text{BG}}(\mathcal{O}) := \max_i \left| \frac{x_i}{\mathcal{O}} \frac{\partial \mathcal{O}}{\partial x_i} \right|_{\mathcal{O}=\mathcal{O}_{\text{exp}}} . \quad (3.134)$$

Roughly, $\Delta_{\text{BG}}(\mathcal{O})^{-1}$ measures the per cent precision in the parameters needed to obtain the experimental value $\mathcal{O} = \mathcal{O}_{\text{exp}}$. In the SM, the sensitivity of the Higgs mass to the input parameters leads to an incredibly large fine-tuning. The main fine-tuning concern of CHMs, on the other hand, is in the separation between the symmetry breaking scales v and f , so let us calculate the fine-tuning of v^2 . By Equation (3.131),

$$\Delta_{\text{BG}}(v^2) = \max_{x=\gamma,\beta} \left| \frac{x}{v^2} \frac{\partial}{\partial x} \left(\frac{\gamma f^2}{2\beta} \right) \right|_{v=246 \text{ GeV}} = \frac{\gamma}{2\beta} \left(\frac{f}{246 \text{ GeV}} \right)^2 = \frac{\gamma}{2\beta} \frac{1}{\xi}, \quad (3.135)$$

yielding a fine-tuning similar to the estimate Equation (3.75). But notice the naïve estimate $\Delta_{\text{BG}}(v^2) \sim 1/\xi$ may be parametrically enhanced or suppressed by the factor $\gamma/2\beta$. Here we shall estimate this factor in some explicit realisations of the MCHM, ultimately finding that some fermion representations are more favourable than others on the basis of fine-tuning.

As a warm up, let us first consider the gauge boson contributions to the Higgs potential. Since the potential vanishes when the $SO(5)$ symmetry is exact, the trick to analysing the gauge boson contributions lies once again in spurion analysis. In this case, the spurions are imagined to be $SO(5)$ gauge fields

$$\mathcal{G}_\alpha W_\mu^\alpha = \mathcal{G}_{\alpha,A} T^A W_\mu^\alpha, \quad \mathcal{G}' B_\mu = \mathcal{G}'_A T^A B_\mu, \quad (3.136)$$

transforming in the adjoint of $SO(5)$, where A indexes all generators of $SO(5)$ and α indexes a triplet under $SU(2)_L$. These explicitly break $SO(5)$ once they take on their background values

$$\mathcal{G}_\alpha = g T_L^\alpha, \quad \mathcal{G}' = g' T_R^3, \quad (3.137)$$

that transform only under G_{EW} , thereby generating a Higgs potential. At leading order in the elementary couplings g and g' , there are only two $SO(5)$ -invariant operators [85]:

$$f^4 \Phi^\dagger \mathcal{G}_\alpha \mathcal{G}^\alpha \Phi = \frac{3}{4} f^4 g^2 s_h^2, \quad \text{and} \quad f^4 \Phi^\dagger \mathcal{G}' \mathcal{G}' \Phi = \frac{1}{4} f^4 g'^2 s_h^2, \quad (3.138)$$

where $\Phi = U\Phi_0$, as in Section 3.5.1. These are the only operators that may enter the Higgs potential, accompanied by dimensionless coefficients. We see the gauge bosons only contribute to the s_h^2 term in V_{eff} , giving contributions of order $f^4 g^{(\prime)2}$ to the γ coefficient. It follows that gauge bosons by themselves cannot provide a satisfactory Higgs vev (as we have stated before), so we next look to the fermion contributions.

Double Tuning

We begin with the contributions to the Higgs potential from fermions in the $\mathbf{5}$ of $SO(5)$. Again, we only consider the third generation quarks, since these are the most composite. As for the gauge bosons, we find the contributions through spurion analysis. In a slight deviation from the example in Section 3.6.2, the spurions here will not be the multiplets of elementary fermions, but rather the elementary-composite coupling constants themselves, following Ref. [86]. The mixings will be written as

$$\mathcal{L}_{\text{mix}} = \bar{q}_L \tilde{\Delta}_q \Psi_R^q + \bar{t}_R \tilde{\Delta}_t \tilde{\Psi}_L^t + \text{h.c.}, \quad (3.139)$$

where the spurions $\tilde{\Delta}_{q,t}$ are in the antifundamental representation $\bar{\mathbf{5}}$ and take on the background values

$$\tilde{\Delta}_q = \frac{\Delta_q}{\sqrt{2}} \begin{pmatrix} 0 & 0 & -i & -1 & 0 \\ -i & 1 & 0 & 0 & 0 \end{pmatrix}, \quad \tilde{\Delta}_t = \Delta_t \begin{pmatrix} 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (3.140)$$

This is equivalent to the formulation in Section 3.6.2. From here finding the contributions to the potential is simply an exercise in writing all $SO(5)$ invariants using $\tilde{\Delta}_{q,t}$ and Φ . At quadratic order in the spurions, there are two invariant operators (in the unitary gauge)

$$\Phi^\dagger \tilde{\Delta}_q^\dagger \tilde{\Delta}_q \Phi = \frac{\Delta_q^2}{2} s_h^2, \quad \text{and} \quad \Phi^\dagger \tilde{\Delta}_t^\dagger \tilde{\Delta}_t \Phi = \Delta_t^2 (1 - s_h^2). \quad (3.141)$$

We see that, once again, only the s_h^2 term in the potential receives any contributions. It is only once we go to quartic order in the spurions that s_h^4 terms appear. By naïve dimensional analysis, the coefficients in the potential must therefore be of the form

$$\gamma = a \frac{N_c}{16\pi^2} \Delta^2 m_*^2, \quad \beta = b \frac{N_c}{16\pi^2} \Delta^4, \quad (3.142)$$

where a and b are dimensionless parameters of order unity, $N_c = 3$ is the number of colours, and Δ is shorthand for Δ_q or Δ_t . Barring any significant cancellations in a and b , the fine-tuning of the model Equation (3.135) is then expected to be²¹

$$\Delta_{\text{BG}} \approx \frac{1}{2} \left(\frac{m_*}{\Delta} \right)^2 \frac{1}{\xi}. \quad (3.143)$$

So in the **5** model, the fine-tuning is **enhanced** by a factor $(m_*/\Delta)^2$ from the naïve estimate. This phenomenon was noticed in Ref. [85] and discussed more comprehensively in Refs. [86, 87]. It is understood as a **double tuning**: first there needs to be a tuning of order $(m_*/\Delta)^2$ to make the s_h^4 term comparable in size to the s_h^2 term, and on top of that the overall scale of the potential needs to be tuned to result in a small misalignment ξ .

It should be clear that the double tuning will be present in any model in which γ arises at quadratic order in the spurions while β only arises at quartic order. From the possible invariants in the different representations of $SO(5)$ listed in Ref. [87], it is seen that the **10** representation also suffers from a double tuning, and only the **14** representation avoids it with an invariant in s_h^4 at quadratic order. As long as the left-handed quark doublet is partnered to a **14**, the coefficients in the potential will be of order

$$\gamma = a \frac{N_c}{16\pi^2} \Delta^2 m_*^2, \quad \beta = b \frac{N_c}{16\pi^2} \Delta^2 m_*^2, \quad (3.144)$$

giving the minimal fine-tuning

$$\Delta_{\text{BG}} \approx \frac{1}{2\xi}. \quad (3.145)$$

By this criterion, models with $q_L \sim \mathbf{14}$ are preferable to others. It would appear at first that minimally tuned models have the disadvantage of tending to predict too large a Higgs mass because of a lack of parametric suppression of β , but a further analysis reveals the Higgs mass is equally tuned across the different models [69].

As a final note, we mention that the BG fine-tuning measure is not particularly suited for quantifying tuning in a doubly-tuned model. With double tuning, two separate parameters need to be tuned to produce the correct value of an observable, but the BG measure only takes the maximum tuning of all parameters. A more sophisticated tuning measure that takes into account the dependencies of double- and higher-order tunings has recently been developed for use in CHMs [30, 88], though we will not go into the details here.

²¹The gauge contribution is neglected, since it only provides small corrections.

3.8 Multi-Site Models and Collective Breaking

If we are to analyse composite Higgs models, the Coleman-Weinberg potential poses a clear problem: it is quadratically and logarithmically divergent in the high-energy cutoff Λ of the theory, generically sabotaging its calculability. How can we construct predictable models in light of this? Well, when the Goldstone symmetry G is exact, the effective potential vanishes and everything is perfectly predictable. It is once G is explicitly broken that operators are introduced into the Lagrangian that allow loop diagrams that give rise to a divergent effective potential. The key to constructing a predictable model, then, lies in restricting the allowable operators so that none produces a quadratic or logarithmic divergence in the effective potential. This is done by extending the symmetry structure of the theory, leading to the idea of multi-site models and collective symmetry breaking [42, 85]. To describe such models, though, we must acquaint ourselves once again with the CCWZ formalism for special cases of symmetry breaking structures.

3.8.1 Spontaneous Breaking to the Diagonal Subgroup

Ultimately, the effective potential will be rendered calculable by extending the symmetry group G to a product group $G_L \times G_R$, which then spontaneously breaks to G_{L+R} . Here the groups are all isomorphic to G , with the L and R subscripts denoting whether the group acts on left- or right-handed fields. G_{L+R} is the diagonal subgroup of $G_L \times G_R$, containing all elements of the product group that have equal left- and right-handed transformations. Fortunately this is an exceptional type of breaking that allows the NGB transformation rule Equation (3.47) to be given explicitly in terms of known group elements. To show this, we summarise the detailed exposition in Ref. [65].

The exception comes from the fact that $(G_L \times G_R)/G_{L+R}$ is a *symmetric space*, meaning that in addition to the Lie algebra commutation relations in Equations (3.40) and (3.41), also the following is satisfied:

$$[\hat{T}_a, \hat{T}_b] = if_{ab}{}^\gamma X_\gamma + if_{ab}{}^c \hat{T}_c, \quad (3.146)$$

where the same convention for generators is being used as in Section 3.4, *i.e.* X_α are the generators for the diagonal subgroup and \hat{T}_a are the remaining generators of the product group. By inspection, swapping the sign of all \hat{T}_a generators preserves all of the commutation relations and is therefore an automorphism of the Lie algebra. We shall denote this operation τ . It follows that any valid equation will remain valid under the action of τ . In particular, since the NGB matrix is given by

$$U[\vec{\theta}] = e^{i\frac{\sqrt{2}}{f}\theta^a \hat{T}_a}, \quad (3.147)$$

we have

$$\tau(U[\vec{\theta}]) = U^{-1}[\vec{\theta}], \quad (3.148)$$

and from the definition of the Maurer-Cartan form Equation (3.49), this implies

$$\begin{aligned} \mathbf{a}_\mu[U^{-1}] &= \tau(\mathbf{a}_\mu[U]) = -d_\mu^a[U]\hat{T}_a + e_\mu^\alpha[U]X_\alpha, \\ \implies d_\mu[U] &= \frac{1}{2}(\mathbf{a}_\mu[U] - \mathbf{a}_\mu[U^{-1}]), \end{aligned} \quad (3.149)$$

where we are not considering the gauge field A_μ for now. Direct substitution of Equation (3.149) using Equation (3.49) into the NGB Lagrangian yields

$$\mathcal{L}_{(2)} = \frac{f^2}{4} \text{Tr} [d_\mu^\dagger d^\mu] = \frac{f^2}{16} \text{Tr} [(\partial_\mu \Sigma^{-1})(\partial^\mu \Sigma)], \quad (3.150)$$

where $\Sigma(x) = U^2(x)$. The automorphism can be used to find an equivalent transformation rule for U :

$$\begin{aligned} U^{-1} &= \tau(U) \xrightarrow{g \in G} \tau(gUh^{-1}) = \tau(g)\tau(U)\tau(h^{-1}) = \tau(g)U^{-1}h^{-1} \\ \implies U &\xrightarrow{g \in G} hU\tau(g)^{-1}, \end{aligned} \quad (3.151)$$

resulting in a simple transformation rule for Σ :

$$\Sigma \xrightarrow{g \in G} (gUh^{-1})(hU\tau(g)^{-1}) = g\Sigma\tau(g)^{-1}. \quad (3.152)$$

So far this analysis applies to any symmetric space. Now to see how this applies in the case of diagonal breaking. To make the notation more convenient, group elements will be written as matrices: for $g \in G_L \times G_R$, we can write

$$g = \begin{pmatrix} g_L & 0 \\ 0 & g_R \end{pmatrix} \quad (3.153)$$

for some $g_{L,R} \in G_{L,R}$. Then the generators will in general be given by

$$X_a = \frac{1}{\sqrt{2}} \begin{pmatrix} T_a & 0 \\ 0 & T_a \end{pmatrix}, \quad \hat{T}_a = \frac{1}{\sqrt{2}} \begin{pmatrix} T_a & 0 \\ 0 & -T_a \end{pmatrix}, \quad (3.154)$$

where T_a are the generators of the original group G . Notice that the automorphism τ , which changes the sign of \hat{T}_a , simply interchanges the blocks in \hat{T}_a - the only distinguishing feature between left- and right- group elements. That is, the automorphism swaps left- and right- group transformations.

To arrive at the transformation rule for Σ , we first calculate

$$U[\vec{\theta}] = e^{i\frac{\sqrt{2}}{f}\theta^a\hat{T}_a} = \begin{pmatrix} e^{\frac{i}{f}\theta^a T_a} & 0 \\ 0 & e^{-\frac{i}{f}\theta^a T_a} \end{pmatrix}, \quad (3.155)$$

so that

$$\Sigma(x) = U^2(x) = \begin{pmatrix} e^{i\frac{2}{f}\theta^a(x)T_a} & 0 \\ 0 & e^{-i\frac{2}{f}\theta^a(x)T_a} \end{pmatrix} =: \begin{pmatrix} \Omega(x) & 0 \\ 0 & \Omega^{-1}(x) \end{pmatrix}. \quad (3.156)$$

And under a $G_L \times G_R$ transformation, this will transform as

$$\Sigma \rightarrow g\Sigma\tau(g)^{-1} = \begin{pmatrix} g_L & 0 \\ 0 & g_R \end{pmatrix} \begin{pmatrix} \Omega(x) & 0 \\ 0 & \Omega^{-1}(x) \end{pmatrix} \begin{pmatrix} g_R^{-1} & 0 \\ 0 & g_L^{-1} \end{pmatrix}, \quad (3.157)$$

or more simply,

$$\Omega(x) \rightarrow g_L\Omega(x)g_R^{-1}. \quad (3.158)$$

Heavy use will be made of this transformation property throughout the rest of this work. As a final note, it is conventional to rescale $f \rightarrow \sqrt{2}f$ for models of this form so that the NGB fields in Ω have the same normalisation factor as in the usual NGB matrix U . Doing this results in the NGB Lagrangian

$$\mathcal{L}_{(2)} = \frac{f^2}{4} \text{Tr} [(\partial_\mu\Omega^{-1})(\partial^\mu\Omega)], \quad \text{for } \Omega(x) := e^{i\frac{\sqrt{2}}{f}\theta^a(x)T_a}. \quad (3.159)$$

If gauge fields are to be included, the partial derivatives should be converted to covariant derivatives in accordance with the NGB transformation rule Equation (3.158).

In terms of moose diagrams, this model is represented precisely by

$$\begin{array}{ccc} \text{Global:} & G_L & G_R \\ \text{Site:} & \begin{array}{c} \bigcirc \xrightarrow{\Omega} \bigcirc \end{array} & \\ \text{Gauged:} & \emptyset & \emptyset \end{array} \quad (3.160)$$

Here a new ‘‘site’’ of fields has been introduced, whose placement shows it transforms under G_R and not G_L . The fact that Ω lies above an arrow directed from G_R to G_L means it transforms under the product group as in Equation (3.158). It is to be understood by the arrangement of the groups that the global groups on each site are spontaneously broken to their diagonal subgroup.

3.8.2 Hidden Local Symmetry

There is a remarkable connection between models of $G_L \times G_R \rightarrow G_{L+R}$ symmetry breaking and the non-linear σ -model of spontaneously broken $G \rightarrow H$ symmetry, known as Hidden Local Symmetry (HLS) [89–91]. Specifically, the two models are gauge-equivalent (under certain conditions) if G_R has a gauged subgroup $H_R \cong H$.

Pictorially,

$$\begin{array}{l}
\text{Global:} \\
\text{Site:} \\
\text{Gauged:}
\end{array}
\begin{array}{c}
G \\
\begin{array}{c} \circ \xrightarrow{U} \text{---} | H \end{array} \\
\emptyset
\end{array}
\cong
\begin{array}{c}
G_L \qquad G_R \\
\begin{array}{c} \circ \xrightarrow{\Omega} \text{---} \circ \end{array} \\
\emptyset \qquad H_R
\end{array}
. \quad (3.161)$$

Many models of composite Higgs, and indeed the models considered in this work, utilise the HLS framework not just to extend the symmetry structure of the theory, but also as a method of introducing heavy composite vector resonances into theories based on the simple non-linear $G \rightarrow H$ model.

The equivalence is seen by first considering the diagonal breaking $G_L \times G_R \rightarrow G_{L+R}$. As in the previous section, the NGBs will be described by a matrix

$$\Omega(x) = e^{i\frac{\sqrt{2}}{f}\theta_\Omega^a(x)T_a}, \text{ transforming as } \Omega(x) \xrightarrow{(g_L, g_R) \in G_L \times G_R} g_L \Omega(x) g_R^{-1}, \quad (3.162)$$

where T_a are the generators of $G \cong G_L \cong G_R$. Note there will be $\dim(G)$ -many NGBs, which Equation (3.162) shows transform in the adjoint representation of G_{L+R} . But actually not all of these are physical; the gauge bosons of H_R will eat those NGBs in the adjoint representation of H_R and become massive, leaving $\dim(G) - \dim(H)$ many NGBs transforming under H_R in the leftover representation $\mathbf{r}_{G/H}$ defined by Equation (3.42). Already parallels with the non-linear $G \rightarrow H$ model are apparent. This gauge structure will be fleshed out in greater generality shortly.

Now the gauged subgroup H_R *explicitly* breaks the $G_L \times G_R$ symmetry to only $G_L \times H_R$, assuming H_R is its own normaliser subgroup in G_R . Under this exact group, the NGB matrix transforms as

$$\Omega(x) \xrightarrow{(g_L, h_R) \in G_L \times H_R} g_L \Omega(x) h_R^{-1}(x). \quad (3.163)$$

As it stands, Ω still contains the unphysical degrees of freedom of the NGBs that are eaten by the gauge bosons. These can be removed by a similar trick as before in Equation (3.35), except now by using a gauge transformation instead of a vacuum vector. Specifically, we express Ω as the product of two group elements

$$\Omega(x) = U_\Omega[\vec{\theta}(x)] \cdot \Xi[\vec{\omega}(x)], \quad (3.164)$$

where

$$U_\Omega[\vec{\theta}(x)] = e^{i\frac{\sqrt{2}}{f}\theta^a(x)\hat{T}_a}, \quad \Xi[\vec{\omega}(x)] = e^{i\frac{\sqrt{2}}{f\Xi}\omega^\alpha(x)X_\alpha}. \quad (3.165)$$

Again, X_α are the generators of $H \cong H_R$ and \hat{T}_a are the generators of G that do not generate H . Observe, the factor Ξ can be removed by a gauge transformation

using $h_R = \Xi$:

$$\Omega(x) \rightarrow \Omega(x)\Xi[\vec{\omega}(x)]^{-1} = U_\Omega[\vec{\theta}(x)]. \quad (3.166)$$

This is the unitary gauge, having no unphysical NGBs in its description. As the notation suggests, the equivalence between this model and the one of spontaneous $G \rightarrow H$ symmetry breaking comes from identifying U_Ω with the usual NGB matrix U of the non-linear model. However, it is a little premature to declare equivalence just yet - the transformation properties of U_Ω under G_L need to be compared to the transformation Equation (3.47) of U under G . Under G_L , U_Ω will generically not be expressible as an exponential of broken generators and must instead be decomposed similarly to Ω in Equation (3.164):

$$U_\Omega[\vec{\theta}(x)] \xrightarrow{g_L \in G_L} g_L U_\Omega[\vec{\theta}(x)] \equiv U_\Omega[\vec{\theta}^{(g_L)}(x)] \cdot \Xi[\vec{\omega}^{(g_L)}(x)]. \quad (3.167)$$

But with another gauge transformation, the second factor can again be removed. As long as the H_R gauge symmetry is continually spent in this way to remain in the unitary gauge, the U_Ω transformation will then take the form

$$U_\Omega[\vec{\theta}(x)] \xrightarrow{g_L \in G_L} g_L U_\Omega[\vec{\theta}(x)] h_R^{-1}[g_L, \vec{\theta}(x)], \quad (3.168)$$

where $h_R^{-1}[g_L, \vec{\theta}(x)]$ is the appropriate gauge transformation cancelling out $\Xi[\vec{\omega}^{(g_L)}(x)]$. But this is exactly the transformation rule Equation (3.47)! Save for the massive gauge bosons, there is only a notational difference between the two theories. The equivalence can be made exact by decoupling the gauge bosons, sending their gauge coupling to infinity to make them infinitely heavy (see Equation (2.21)). All trace of the gauge symmetry disappears in this equivalence: there are no gauge bosons, and we are not free to perform a gauge transformation lest we leave the unitary gauge. For this reason, this type of equivalence is known as a Hidden Local Symmetry.

For greater generality, HLS also applies when a subgroup $\mathcal{L}_L \subset G_L$ is gauged, if the corresponding $\mathcal{L} \subset G$ is also gauged in the non-linear model. In this case, the considerations above remain the same except for a slight change in the NGB and gauge boson spectra. It will be convenient for this analysis to split the generators L_a and X_α of \mathcal{L}_L and H_R into those that generate the intersection $\mathcal{L} \cap H$ (which will be indexed by letters with tildes), and those that do not (indexed by letters without tildes). This is because those generating the intersection can be linearly combined into unbroken generators of G_{L+R} , therefore coupling to massless gauge bosons, while others will couple to massive gauge bosons.

Before we demonstrate this algebraically, however, there is a subtlety in this situation that must be addressed. The subtlety is that the Lagrangian

$$\mathcal{L}_{(2)} = \frac{f^2}{4} \text{Tr} [d_\mu^\dagger d^\mu] \quad (3.169)$$

is no longer the most general possible at the two-derivative level. The difference comes from there now being gauge fields $A_{H,\mu} = A_{H,\mu}^\alpha X_\alpha + A_{H,\mu}^{\tilde{\alpha}} X_{\tilde{\alpha}}$ transforming under H_R as²²

$$A_{H,\mu} \rightarrow h(x)A_{H,\mu}h^{-1}(x) + \frac{i}{g_H}h(x)\partial_\mu h^{-1}(x), \quad (3.170)$$

so that by Equation (3.52), the quantity $(g_H A_{H,\mu} - e_\mu)$ transforms as

$$(g_H A_{H,\mu} - e_\mu) \rightarrow h(x)(g_H A_{H,\mu} - e_\mu)h^{-1}(x), \quad (3.171)$$

leading to a new invariant term for the Lagrangian

$$\text{Tr} [(g_H A_{H,\mu} - e_\mu)(g_H A_H^\mu - e^\mu)]. \quad (3.172)$$

It can be seen from Equation (3.57) that to first order, $e_\mu = g_{\mathcal{L}} A_{\mathcal{L},\mu}^{\tilde{a}} L_{\tilde{a}}$. Then recalling the generators are normalised such that $\text{Tr}[X_a X_b] = \delta_{ab}$, the general Lagrangian is seen to have the form

$$\mathcal{L}_{\text{NGB}} = \frac{f^2}{4} \sum_a g_{\mathcal{L}}^2 A_{\mathcal{L},\mu}^a A_{\mathcal{L}}^{a,\mu} + \frac{f_\Xi^2}{4} \left(\sum_\alpha g_H^2 A_{H,\mu}^\alpha A_H^{\alpha,\mu} + \sum_{\tilde{a}} (g_H A_{H,\mu}^{\tilde{a}} - g_{\mathcal{L}} A_{\mathcal{L},\mu}^{\tilde{a}})^2 \right), \quad (3.173)$$

having introduced a new NGB scale f_Ξ for the new invariant. This Lagrangian yields the gauge field mass-mixing matrix

$$m_{\text{gauge}}^2 = \frac{1}{4} \begin{pmatrix} & A_{\mathcal{L},\mu}^a & A_{H,\mu}^\alpha & A_{\mathcal{L},\mu}^{\tilde{a}} & A_{H,\mu}^{\tilde{\alpha}} \\ \hline A_{\mathcal{L}}^{a,\mu} & f^2 g_{\mathcal{L}}^2 & 0 & 0 & 0 \\ A_H^{\alpha,\mu} & 0 & f_\Xi^2 g_H^2 & 0 & 0 \\ A_{\mathcal{L}}^{\tilde{a},\mu} & 0 & 0 & f_\Xi^2 g_{\mathcal{L}}^2 & -f_\Xi^2 g_{\mathcal{L}} g_H \\ A_H^{\tilde{\alpha},\mu} & 0 & 0 & -f_\Xi^2 g_{\mathcal{L}} g_H & f_\Xi^2 g_H^2 \end{pmatrix}. \quad (3.174)$$

Noting this is just the SM electroweak mass matrix Equation (2.34) with slightly different notation, the mass basis is found to have the same form as the EW sector, with fields

$$A_{\mathcal{L}}^{a,\mu}, \quad A_H^{\alpha,\mu}, \quad A_{\mathcal{L}-H}^{\tilde{a},\mu} := \frac{g_H A_H^{\tilde{a},\mu} - g_{\mathcal{L}} A_{\mathcal{L}}^{\tilde{a},\mu}}{\sqrt{g_{\mathcal{L}}^2 + g_H^2}}, \quad A_{\mathcal{L}+H}^{\tilde{\alpha},\mu} := \frac{g_{\mathcal{L}} A_H^{\tilde{\alpha},\mu} + g_H A_{\mathcal{L}}^{\tilde{\alpha},\mu}}{\sqrt{g_{\mathcal{L}}^2 + g_H^2}}, \quad (3.175)$$

having respective masses

$$m_{\mathcal{L}}^2 = \frac{1}{2} f^2 g_{\mathcal{L}}^2, \quad m_H^2 = \frac{1}{2} f_\Xi^2 g_H^2, \quad m_{\mathcal{L}-H}^2 = \frac{1}{2} f_\Xi^2 (g_{\mathcal{L}}^2 + g_H^2), \quad m_{\mathcal{L}+H}^2 = 0. \quad (3.176)$$

²²This is simply using Equation (2.13) along with the fact that $(\partial_\mu h)h^{-1} = -h(\partial_\mu h^{-1})$.

In total, then, there are $\dim(\mathcal{L} \cap H)$ -many massless gauge fields, as was reasoned intuitively, and $\dim(\mathcal{L} \cup H)$ -many massive ones. If $\mathcal{L} \cup H = G$, and in particular if $\mathcal{L} = G$, then all NGBs will be rendered unphysical, each being eaten by a massive gauge boson and being able to be gauged away. What happens in the limit $g_H \rightarrow \infty$ that previously reproduced the non-linear model? In this case, all of the $A_H^{\alpha,\mu}$ and $A_{\mathcal{L}-H}^{\tilde{\alpha},\mu}$ fields will be decoupled, becoming infinitely massive, leaving the massless fields $A_{\mathcal{L}+H}^{\tilde{\alpha},\mu} = A_{\mathcal{L}}^{\tilde{\alpha},\mu}$ and the finite-mass fields $A_{\mathcal{L}}^{\alpha,\mu}$. But of course this is exactly the spectrum of the non-linear model when $\mathcal{L} \subset G$ is gauged, demonstrating the equivalence between the two models once again. The general picture for HLS is then

$$\begin{array}{l}
 \text{Global:} \quad G \\
 \text{Site:} \quad \begin{array}{c} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \end{array} \\
 \text{Gauged:} \quad \mathcal{L}
 \end{array}
 \xrightarrow{g_H \rightarrow \infty}
 \begin{array}{l}
 G_L \qquad G_R \\
 \begin{array}{c} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \end{array} \\
 \mathcal{L} \qquad H_R
 \end{array}
 \quad (3.177)$$

This equivalence can be used in either direction, hiding a local symmetry to produce a non-linear model, or extending a non-linear model to a theory with extra gauge symmetry.

3.8.3 Multi-Site Formalism

HLS shows that the one-site model of $G \rightarrow H$ symmetry breaking has an almost equivalent description as a two-site model with the extended symmetry structure $G_L \times G_R \rightarrow G_{L+R}$, where $H_R \subset G_R$ is gauged. Combining the ideas of the previous two sections can generalise this even further, resulting in multi-site models that have sufficiently elaborate symmetry structures to admit a finite Higgs potential. It is these sorts of models that we will be exploring in later sections.

Adding a Level of Heavy Vector Resonances

Let us consider the HLS extension of $G_L \times G_R \rightarrow G_{L+R}$ symmetry breaking:

$$\begin{array}{l}
 \text{Global:} \quad G_L \times G_R \\
 \text{Site:} \quad \begin{array}{c} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \end{array} \\
 \text{Gauged:} \quad \emptyset
 \end{array}
 \rightarrow
 \begin{array}{l}
 G_L \times G_R \qquad G'_L \times G'_R \\
 \begin{array}{c} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \end{array} \\
 \emptyset \qquad G_{L+R}
 \end{array}
 \quad (3.178)$$

Here the primes on the groups, not the L and R subscripts, distinguish the sites they act on. This symmetry structure is not particularly relevant as it stands, but it will let us uncover some surprising connections similar to HLS between different types of models. Once again we proceed by summarising the excellent exposition in Ref. [65].

In Equation (3.178) the NGB matrix is denoted as $\tilde{\Omega}$ in order to distinguish it from the usual NGB matrix Ω that does not transform under such product groups. From Equation (3.162), $\tilde{\Omega}$ transforms as

$$\tilde{\Omega} \xrightarrow{(\tilde{g}, \tilde{g}') \in \tilde{G} \times \tilde{G}'} \tilde{g} \tilde{\Omega} \tilde{g}'^{-1}, \quad (3.179)$$

where for convenience, the product groups are denoted by $\tilde{G}^{(\prime)} := G_L^{(\prime)} \times G_R^{(\prime)}$. Expressing the product group elements and NGBs in matrix notation (compare with Equation (3.153)),

$$\tilde{g}^{(\prime)} = \begin{pmatrix} g_L^{(\prime)} & 0 \\ 0 & g_R^{(\prime)} \end{pmatrix}, \quad \tilde{\Omega} = \begin{pmatrix} \Omega_1 & 0 \\ 0 & \Omega_2^{-1} \end{pmatrix}, \quad (3.180)$$

for some fields $\Omega_{1,2}$, the transformation Equation (3.179) manifests as

$$\begin{aligned} \Omega_1 &\xrightarrow{(\tilde{g}, \tilde{g}') \in \tilde{G} \times \tilde{G}'} g_L \Omega_1 g_L'^{-1} \\ \Omega_2 &\xrightarrow{(\tilde{g}, \tilde{g}') \in \tilde{G} \times \tilde{G}'} g_R' \Omega_2 g_R^{-1}. \end{aligned} \quad (3.181)$$

We recognise that Ω_1 and Ω_2 respectively transform as NGB matrices of $G_L \times G_L'$ and $G_R' \times G_R$ symmetries broken to their diagonal subgroups. But note that $G_L' \times G_R'$ is *explicitly* broken to the gauged diagonal subgroup G_{L+R} . Under the exact symmetry group, where $g_L' = g_R' \equiv g'$, the matrices transform as

$$\begin{aligned} \Omega_1(x) &\xrightarrow{(\tilde{g}, \tilde{g}') \in \tilde{G} \times G_{L+R}} g_L \Omega_1(x) g'^{-1}(x) \\ \Omega_2(x) &\xrightarrow{(\tilde{g}, \tilde{g}') \in \tilde{G} \times G_{L+R}} g'(x) \Omega_2(x) g_R^{-1}, \end{aligned} \quad (3.182)$$

showing the model Equation (3.178) has an equivalent description as

$$\begin{array}{ccccc} G_L & & G_{L+R} & & G_R \\ \circ & \xrightarrow{\Omega_1} & \circ & \xrightarrow{\Omega_2} & \circ \\ \emptyset & & G_{L+R} & & \emptyset \end{array}. \quad (3.183)$$

We can go even further by noting that the product $\Omega \equiv \Omega_1 \Omega_2$ transforms only under $G_L \times G_R$:

$$\Omega(x) \xrightarrow{(\tilde{g}, \tilde{g}') \in \tilde{G} \times G_{L+R}} g_L \Omega(x) g_R^{-1}, \quad (3.184)$$

the exact transformation found in Section 3.8.1! The model therefore has yet another equivalent description as

$$\begin{array}{lcl}
 \text{Global:} & G_L & G_R \\
 \text{Site:} & \text{---} \circ \xrightarrow{\Omega} \circ \text{---} & \\
 \text{Gauged:} & \emptyset & \emptyset
 \end{array} \quad (3.185)$$

where also a set of heavy G_{L+R} gauge fields is present. It is not difficult to believe, based on what was found previously, that these relations between models extend in greater generality to

$$\begin{array}{lcl}
 \text{Global:} & G & G_L \quad G_R & G_L \quad G_{L+R} \quad G_R \\
 \text{Site:} & \text{---} \circ \xrightarrow{U} \text{---} | H & \xrightarrow[H]{\text{gauge}} \text{---} \circ \xrightarrow{\Omega} \circ \text{---} & \xrightarrow[G_{L+R}]{\text{gauge}} \text{---} \circ \xrightarrow{\Omega_1} \circ \xrightarrow{\Omega_2} \circ \text{---} \\
 \text{Gauged:} & \mathcal{L} & \mathcal{L} \quad H & \mathcal{L} \quad G_{L+R} \quad H
 \end{array} \quad (3.186)$$

This method of introducing vector resonances is equivalent to a wide range of other approaches [92]. Adding a level of heavy vector resonances, then, only requires factoring the NGB matrix Ω of $G_L \times G_R \rightarrow G_{L+R}$ symmetry breaking into a product of matrices $\Omega_1 \Omega_2$ and introducing another G symmetry. Without a relation between the groups associated to Ω_1 and Ω_2 , these NGB matrices may have independent NGB decay constants f_1, f_2 , making the NGB Lagrangian Equation (3.159) into

$$\mathcal{L}_{(2)} = \sum_{i=1}^2 \frac{f_i^2}{4} \text{Tr} [(\partial_\mu \Omega_i^{-1})(\partial^\mu \Omega_i)]. \quad (3.187)$$

Note there is still a G_{L+R} gauge freedom that may be spent on Ω_1 and Ω_2 , but this will not affect the product matrix Ω . Since either Ω_1 or Ω_2 may be set to the identity with a suitable gauge transformation, the physical NGBs of $G \rightarrow H$ breaking cannot be parametrised by either matrix individually, but only by the product matrix Ω .

The N-Site Model

Of course, there is no reason to stop at adding only one level of heavy vector resonances. We could just as well add another level of heavy G_{L+R} resonances by splitting Ω_2 itself into a product of NGB matrices, and so on *ad infinitum*. Doing so results a model of the form

$$\begin{array}{lcl}
 \text{Global:} & G_1 & G_2 \quad G_3 & \dots & G_N \quad G_{N+1} \\
 \text{Site:} & \text{---} \circ \xrightarrow{\Omega_1} \circ \xrightarrow{\Omega_2} \circ \text{---} & \dots & \text{---} \circ \xrightarrow{\Omega_N} \circ \text{---} & \\
 \text{Gauged:} & \mathcal{L} & G_2 \quad G_3 & & G_N \quad H
 \end{array} \quad (3.188)$$

having relabelled the subscripts on the groups for clarity. Introduced in Ref. [85], Equation (3.188) is named the **(N+1)-site Discrete CHM** (DCHM) on account of its moose structure. Actually the name runs deeper, for the sites of fields in this construction are viewed as lattice points in a discretisation of a fifth dimension. This is the sense in which CHMs are connected to the 5D Randall-Sundrum models mentioned in Section 2.5 - DCHMs are interpreted as deconstructed versions of the extra-dimensional models. See Ref. [93] for the origins of dimensional deconstruction.

The full 5D theory is recovered in the limit $N \rightarrow \infty$, with the gauge couplings of the sites becoming a continuously varying warp factor of the metric along the fifth dimension. A finite number of sites is simply regarded as a truncation of the infinite Kaluza-Klein tower of resonances present in the 5D models, valid in the low-energy limit [94]. Interestingly, this type of construction can be used to describe QCD, and there it is found that the geometry of the fifth dimension is not of much consequence for the 4D theory, and truncating at around three levels of resonances to fit experiment is about just as good as using the full 5D model [95].

This extra-dimensional viewpoint motivates much of the structure of DCHMs, to be detailed below.

Boson Sector

First of all, let us unpack what Equation (3.188) represents. The diagram tells us that the (N+1)-site DCHM features $N + 1$ sites of fields (obviously), with Site k acted upon by a global symmetry group G_k . Symmetries of neighbouring sites spontaneously break to their diagonal subgroups:

$$G_k \times G_{k+1} \rightarrow G_{k+(k+1)}, \quad k = 1, \dots, N, \quad (3.189)$$

and N matrices Ω_k parameterise the resulting NGBs. The NGB matrices transform as

$$G_k \times G_{k+1} : \Omega_k \rightarrow g_k \Omega_k g_{k+1}^{-1}, \quad (3.190)$$

with covariant derivatives

$$D_\mu \Omega_k = \partial_\mu \Omega_k - i g_k A_{k,\mu} \Omega_k + i g_{k+1} \Omega_k A_{(k+1),\mu}, \quad (3.191)$$

where in a slight abuse of notation g_k stands for an element of G_k in Equation (3.190) but is a gauge coupling in Equation (3.191), while $A_{k,\mu}$ is the gauge field of G_k (or of \mathcal{L} for $k = 1$, or H for $k = N + 1$). The generalisation of the NGB Lagrangian

Equation (3.187), including gauge field kinetic terms, is²³

$$\mathcal{L}_{\text{NGB}} = \sum_{k=1}^N \frac{f_k^2}{4} \text{Tr} [(D_\mu \Omega_k^{-1})(D^\mu \Omega_k)] - \frac{1}{4} \sum_{k=1}^{N+1} \text{Tr} [A_{k,\mu\nu} A_k^{\mu\nu}]. \quad (3.192)$$

Many of these NGBs are unphysical, being eaten by the $N - 1$ levels of heavy vector resonances in the adjoint of G . This leaves $\dim(G)$ -many NGBs which will be eaten by the gauge bosons of \mathcal{L} and H according to the procedure specified by Section 3.8.2. The specifics are different in this case because the gauge bosons of the different sites will mix into massless diagonal combinations and massive axial combinations, but the counting remains the same.

The product

$$\Omega := \prod_{k=1}^N \Omega_k \quad (3.193)$$

transforms by Equation (3.190) as the NGB matrix of $G_1 \times G_{N+1} \rightarrow G_{1+(N+1)}$ symmetry breaking. Ultimately we will only be interested in models with NGBs of $G \rightarrow H$ breaking, so hereafter we focus on the HLS limit where the H gauge fields are decoupled²⁴:

$$\begin{array}{ccccccc} \text{Global:} & G_1 & G_2 & G_3 & \dots & G_N & \\ \text{Site:} & \text{---} \circ \xrightarrow{\Omega_1} \text{---} \circ \xrightarrow{\Omega_2} \text{---} \circ \xrightarrow{\dots} \text{---} \circ \xrightarrow{\Omega_N} \text{---} | H. & & & & & \\ \text{Gauged:} & \mathcal{L} & G_2 & G_3 & & G_N & \end{array} \quad (3.194)$$

We will refer to Equation (3.194) as the **N-site model**, not to be confused with the N-site DCHM. In this limit Ω will act as the NGB matrix for $G_1 \rightarrow H$ symmetry breaking, transforming as

$$G_1 \times H : \Omega(x) \rightarrow g_1 \Omega(x) h^{-1}(x). \quad (3.195)$$

Generally, Ω can be written in the form

$$\Omega(x) = e^{i \frac{\sqrt{2}}{f} \theta_\Omega^a(x) T_a}, \quad (3.196)$$

where T_a are the generators of G . The individual matrices Ω_k can be put similarly, with their own normalisations for the phase fields:

$$\Omega_k(x) = e^{i \frac{\sqrt{2}}{f_k} \frac{f}{f_k} \theta_\Omega^a(x) T_a}. \quad (3.197)$$

²³This Lagrangian assumes that the decay constants for each of the factors of the decomposition $\Omega_k(x) = U_{\Omega_k}(x) \Xi_k(x)$ are equal: $f_k = f_{\Xi_k}$.

²⁴We did not consider this limit from the start because the boson content is only understood by first considering the general DCHM.

Notice that to match with Equation (3.193), or to canonically normalise the θ^a kinetic terms in Equation (3.192), these decay constants must be related by

$$\frac{1}{f^2} = \sum_{k=1}^N \frac{1}{f_k^2}. \quad (3.198)$$

Much gauge freedom exists with which to manipulate the forms of Ω_k . One simple possibility is the **Site N holographic gauge**, defined by

$$\begin{aligned} \Omega_k(x) &= \mathbb{1} \quad \text{for } k = 1, \dots, N-1, \\ \Omega_N(x) &= U(x) = e^{i\frac{\sqrt{2}}{f}\theta^\alpha(x)X_\alpha}, \end{aligned} \quad (3.199)$$

where X_α are the generators of H . This removes the unphysical NGBs from the description and shifts all the physical NGB dependence onto Ω_N . It is reached from the original gauge Equation (3.197) by the gauge transformations

$$\begin{aligned} g_k(x) &= \prod_{n=1}^{k-1} \Omega_n(x) \quad \text{for } k = 2, \dots, N, \\ h(x) &= U(x)^{-1}\Omega(x) \in H. \end{aligned} \quad (3.200)$$

Another is the **Site 1 holographic gauge**, defined similarly but with $\Omega_1 = U(x)$ and $\Omega_N = \mathbb{1}$. It is reached from the Site N holographic gauge with the gauge transformations $g_k = U(x)^{-1}$ for $k = 2, \dots, N$. We will be using this gauge throughout the rest of this work as it produces relatively uncomplicated couplings between fields.

Fermion Sector

The sites of the N-site model must of course each be populated by fields transforming under that site's symmetry group. To respect all of the symmetries, fields of two different sites can only couple together through NGB matrices Ω_k . For the sake of illustration, let us assume the fermions of each site transform in the fundamental representation:

$$\Psi_k \xrightarrow{g_k \in G_k} g_k \Psi_k. \quad (3.201)$$

Then an appropriate invariant interaction term will be

$$\mathcal{L}_{\text{int}} \supset \bar{\Psi}_k \Omega_k \Psi_{k+1}. \quad (3.202)$$

We could also have, for example,

$$\mathcal{L}_{\text{int}} \supset \bar{\Psi}_k \Omega_k \Omega_{k+1} \Psi_{k+2}, \quad (3.203)$$

along with all of the other obvious generalisations. But this unwieldy set of fermion interactions may be curtailed with guidance from the 5D models. Specifically, we will not be considering couplings between non-neighbouring sites as in Equation (3.203) because these translate, in the extra-dimensional models, to a violation of spacetime locality [96]. This “locality in theory space” is not necessary in a general CHM, but does provide the minimal structure necessary for calculability, passed down from the calculability of the 5D models. Fermion interactions, including partial compositeness, are therefore enacted entirely through terms of the form Equation (3.202) in the models we consider.

The fermion content is refined by following the example of the 5D models further. A fermion in the 5D bulk gives rise to a Dirac fermion²⁵ at each site in the 4D theory, and with the appropriate boundary conditions yields a massless chiral fermion on the leftmost site [12, 44]. We therefore identify this as the elementary site, so each chirality ψ_L or ψ_R of an elementary fermion will be coupled to an associated tower of composite Dirac fermions Ψ^k or $\tilde{\Psi}^k$ throughout the remaining sites. These towers cannot interact at any site $k < N$, for otherwise ψ_L could couple to ψ_R through a chain of interactions that does not include the Higgs field and ψ would possess a mass before EWSB (to see this, take the Site N holographic gauge). Fermion couplings are then constrained to have the very particular structure

$$\begin{array}{ccccccc}
 \underbrace{\left(\begin{array}{c} \psi_L \\ \psi_R \end{array} \right)}_{\text{Elem.}} & \xrightarrow{\Delta_L^1} & \underbrace{\left(\begin{array}{c} \Psi^2 \\ \tilde{\Psi}^2 \end{array} \right)}_{\text{Site 2}} & \xrightarrow{\Delta_L^2} & \underbrace{\left(\begin{array}{c} \Psi^3 \\ \tilde{\Psi}^3 \end{array} \right)}_{\text{Site 3}} & \cdots & \underbrace{\left(\begin{array}{c} \Psi^N \\ \tilde{\Psi}^N \end{array} \right)}_{\text{Site N}} \\
 & & & & & & \uparrow m_Y, Y \\
 & & & & & & \downarrow m_Y, Y
 \end{array} \tag{3.204}$$

where $\Delta_{L,R}^k$ are the coupling strengths between sites, m_Y is a mass-like coupling, and Y is a Yukawa-like coupling. Each chiral component of each composite fermion will only be coupled to either the preceding or the following site, and not both at once, simply because this is the minimal working structure. The chiral components at each site will be coupled together with a Dirac mass term. In total, the fermion Lagrangian has the form

$$\begin{aligned}
 \mathcal{L}_{\text{fermion}} = & \bar{\psi} i \not{D} \psi & \left. \vphantom{\bar{\psi} i \not{D} \psi} \right\} & \text{elementary kinetic} \\
 & + \sum_{k=2}^N \left(\bar{\Psi}^k (i \not{D} - m_L^k) \Psi^k + \bar{\tilde{\Psi}}^k (i \not{D} - m_R^k) \tilde{\Psi}^k \right) & \left. \vphantom{\sum_{k=2}^N} \right\} & \text{composite kinetic} \\
 & + \sum_{k=1}^N \left(\Delta_L^k \bar{\Psi}_L^k \Omega_k \Psi_R^{k+1} + \Delta_R^k \bar{\tilde{\Psi}}_R^k \Omega_k \tilde{\Psi}_R^{k+1} \right) + \text{h.c.} & \left. \vphantom{\sum_{k=1}^N} \right\} & \text{link} \\
 & - m_Y \bar{\Psi}_L^N \tilde{\Psi}_R^N - Y \bar{\Psi}_L^N \Phi \Phi^\dagger \tilde{\Psi}_R^N + \text{h.c.} & \left. \vphantom{- m_Y \bar{\Psi}_L^N \tilde{\Psi}_R^N} \right\} & \text{Yukawa-like.}
 \end{aligned} \tag{3.205}$$

²⁵All fermions in 5D are Dirac fermions [97].

Here, the fields Ψ^1 and $\tilde{\Psi}^1$ on Site 1 are the respective elementary fields ψ_L and ψ_R , and $\Phi := \Omega_N \Phi_0$ is defined in terms of the vacuum vector Φ_0 invariant under H , so that Φ transforms only under G_N . Again, these fields are assumed to be in the fundamental representation. Other representations will be considered in Chapter 4.

3.8.4 Collective Breaking

Having now seen some models with extended symmetry sectors, let us investigate whether these are sufficiently restrictive for the Higgs potential to be finite and calculable. To start with, let us consider the simple extension of the symmetry structure $G \rightarrow H$ provided by HLS (see Section 3.8.2),

$$\begin{array}{l}
 \text{Global:} \\
 \text{Site:} \\
 \text{Gauged:}
 \end{array}
 \begin{array}{c}
 G \\
 \begin{array}{c} \circ \xrightarrow{U} \text{---} | H \end{array} \\
 \mathcal{L}
 \end{array}
 \rightarrow
 \begin{array}{c}
 G_L \qquad G_R \\
 \begin{array}{c} \circ \xrightarrow{\Omega} \text{---} \circ \end{array} \\
 \mathcal{L}_L \qquad H_R
 \end{array}
 . \quad (3.206)$$

Now there are two independent symmetries G_L and G_R , under which the NGB matrix transforms as

$$\Omega(x) \rightarrow g_L \Omega(x) g_R^{-1}. \quad (3.207)$$

In the same way the original NGB symmetry G had to be broken to generate an effective potential, now *both* G_L and G_R need to be broken for an effective potential to be generated. Indeed, any NGB vev Ω_0 , could be removed by a symmetry transformation $g_L = \Omega_0^{-1}$ or a transformation $g_R = \Omega_0$. Only with a ‘‘collective breaking’’ of both symmetries can the NGB vevs become physical, signalling they are actually pNGBs with a non-vanishing potential [42].

Does this extra symmetry actually protect the Higgs potential from divergences though? Naïve dimensional analysis shows that the one-loop contributions to the effective potential, using the leading order NGB Lagrangian, are of the form [85]

$$\sim \frac{\Lambda^4}{16\pi^2} \left(\frac{\theta}{f}\right)^{N_E} \left(\frac{gf}{\Lambda}\right)^{2\eta} \left(\frac{\mu}{\Lambda}\right)^\chi \propto \Lambda^{4-2\eta-\chi}, \quad (3.208)$$

where N_E is the number of external NGB legs, 2η is the number of gauge field insertions, χ is the number of fermion mass-like insertions, and θ , g , μ are the corresponding NGB fields, gauge couplings, and mass terms, respectively. We see the number of gauge and fermion insertions in a given loop dictate its degree of divergence.

Spurion analysis can be used to determine the number of these insertions. The gauge couplings and fermion mass-like couplings are imagined to be dynamical fields, transforming under the symmetries as in Section 3.7.1, which explicitly break the symmetries when they assume their actual constant values. This allows operators

that contain the explicit symmetry breaking terms that contribute to the pNGB potential to be constructed. Focussing on the gauge couplings, for example, the spurions $\mathcal{G}_{\mathcal{L}}^a$ and \mathcal{G}_H^a will transform as

$$\begin{aligned} G_L : \quad \mathcal{G}_{\mathcal{L}}^a &\rightarrow g_L \mathcal{G}_{\mathcal{L}}^a g_L^{-1}, \\ G_R : \quad \mathcal{G}_H^a &\rightarrow g_R \mathcal{G}_H^a g_R^{-1}. \end{aligned} \quad (3.209)$$

When they are set to their physical values proportional to the group generators (using the same notation as previous sections)

$$\begin{aligned} \mathcal{G}_{\mathcal{L}}^a &\rightarrow g_{\mathcal{L}} L^a, \\ \mathcal{G}_H^a &\rightarrow g_H X^a, \end{aligned} \quad (3.210)$$

they *explicitly* break $G_L \times G_R$ to only $\mathcal{L}_L \times H_R$ because they no longer transform under the full groups as in Equation (3.209). With the transformation properties Equations (3.207) and (3.209), it is seen that the leading one-loop gauge contribution to the effective potential comes from operators such as

$$\text{Tr} [(\mathcal{G}_{\mathcal{L}} \Omega \mathcal{G}_H)(\mathcal{G}_H \Omega^\dagger \mathcal{G}_{\mathcal{L}})]. \quad (3.211)$$

This contains $2\eta = 4$ gauge field insertions, corresponding in Equation (3.208) to a logarithmic divergence in the effective potential. A similar analysis can be done for the fermion contributions, yielding the same result (though we shall not introduce the machinery here). Success! The degree of divergence has been reduced from quadratic to logarithmic. In fact, it can be shown that *any* extension of the naïve non-linear $G \rightarrow H$ model without a quadratically divergent potential must have a two-site structure [69].

Of course, we still would like the logarithmic divergence to go away. Motivated by the success of extending the symmetry by HLS, we take the bold step of extending the symmetry once again and consider the three-site DCHM:

$$\begin{array}{l} \text{Global:} \quad G_L \qquad \qquad G' \qquad \qquad G_R \\ \text{Site:} \quad \begin{array}{c} \text{---} \text{---} \text{---} \\ \circ \text{---} \xrightarrow{\Omega_1} \text{---} \circ \text{---} \xrightarrow{\Omega_2} \text{---} \circ \text{---} \end{array} \\ \text{Gauged:} \quad \mathcal{L}_L \qquad \qquad G' \qquad \qquad H_R \end{array} \quad (3.212)$$

Now to explicitly break the NGB symmetry, an operator must include *both* NGB matrices $\Omega_{1,2}$ because the G' gauge symmetry can be used to shift the NGB dependence of Ω_1 and Ω_2 at will. Leading-order operators (in the gauge sector) for the effective potential are then of the form

$$\text{Tr} [(\mathcal{G}_{\mathcal{L}} \Omega_1 \mathcal{G}')(\mathcal{G}' \Omega_1^\dagger \mathcal{G}_{\mathcal{L}})] \cdot \text{Tr} [(\mathcal{G}' \Omega_2 \mathcal{G}_H)(\mathcal{G}_H \Omega_2^\dagger \mathcal{G}')] \quad (3.213)$$

This contains $2\eta = 8$ gauge field insertions, ensuring a finite contribution to the Higgs potential! Again, the same is true for the fermionic contributions too [85]. But we warn that while removing the quadratic divergence from the potential requires a two-site structure, a removal of the logarithmic divergence does not necessitate a three-site DCHM. In fact, the slightly more minimal two-site model, where the H gauge bosons are decoupled through HLS, is sufficient to avoid the logarithmic divergence in the MCHM.

Because the quadratic and logarithmic divergences in the scalar potential Equation (3.122) are avoided in models with at least three sites, it must be true that the **Weinberg Sum Rules**

$$(1) \quad \sum_i m_i^2(\Phi_c) \text{ is independent of } \Phi_c, \quad (3.214)$$

$$(2) \quad \sum_i m_i^4(\Phi_c) \text{ is independent of } \Phi_c, \quad (3.215)$$

are automatically satisfied in such models for sums over particles of each species, so that the potential is equivalent to²⁶

$$V_{\text{eff}}(\Phi_c) = \sum_{\text{particles } i} \frac{c_i}{64\pi^2} m_i^4(\Phi_c) \log(m_i^2(\Phi_c)). \quad (3.216)$$

This would not be expected to hold in general; different coupling constants could be assigned to composite fields in different H multiplets, for example, which would not satisfy the Weinberg Sum Rules. With the symmetry structure imposed, however, coupling constants must apply to whole composite G multiplet, which ensures the rules are satisfied in the MCHM [98].

²⁶Any units of mass can be used in Equation (3.216) - changing units only adds a constant, by the Weinberg Sum Rules.

The Two-Site Minimal 4D Composite Higgs Model

This chapter is dedicated to providing a full description of the minimal viable 4D composite Higgs model (M4DCHM), using the ideas of Chapter 3. First put forward in Ref. [99], variations of this model have been studied extensively in the literature - see in particular Refs. [30, 87, 98–101]. Here we give the complete Lagrangians for three different realisations of the M4DCHM that vary in their fermion representations in Sections 4.1 to 4.3, and analyse their particle content in Section 4.6. In preparation for our global fits of these models in Part II of this thesis, we also review previous studies of the models in Section 4.7.

4.1 Symmetry Structure

First and foremost, the M4DCHM assumes the minimal symmetry content that yields the Higgs doublet in the NGB spectrum and provides a finite and calculable Higgs potential, while facilitating composite fermion partners. We have already reasoned that the minimal model is based on an overall $SO(5) \rightarrow SO(4)$ symmetry breaking pattern so that the Higgs doublet is produced. We also know from partial compositeness that $SU(3)_c$ group factors must be present to respect the SM colour symmetry, and a final group factor of $U(1)_X$ was shown to be necessary in Section 3.6.1 for quark partners to be assigned the correct hypercharge under the definition


$$Y = T_R^3 + X, \quad (4.1)$$

where the isospin operator T_R^3 is the third generator¹ of $SU(2)_R \subset SO(5)$. So in total, the breaking pattern is

$$\underbrace{SU(3)_c \times SO(5) \times U(1)_X}_G \rightarrow \underbrace{SU(3)_c \times SO(4) \times U(1)_X}_H. \quad (4.2)$$

¹See Appendix G for our convention with $SO(5)$ generators.

The breaking is enacted through HLS, using the two-site structure that is sufficient for a finite Higgs potential:



Site 0 is the elementary sector², consisting of massless fields ψ with the same quantum numbers as the non-Higgs SM fields. It mixes with the composite fields $\Psi, \tilde{\Psi}$ of Site 1, so that physical fields are superpositions of elementary and composite states. A global symmetry group G^k , isomorphic to the G of Equation (4.2), acts on Site k . By Section 3.8.2, the unbroken gauge group of the theory (before EWSB) will be the vector subgroup of $\mathcal{L}^0 \times G^1$, which is isomorphic to \mathcal{L}^0 . Since we wish to reproduce the gauge symmetry of the SM, the elementary gauge group is taken to be

$$\mathcal{L}^0 := SU(3)_c^0 \times SU(2)_L^0 \times U(1)_Y^0. \quad (4.4)$$

The SM gauge bosons will be mixtures of those of \mathcal{L}^0 and G^1 .

4.2 Boson Sector

The spectrum of bosons can be reasoned from the discussion in Section 3.8.3. To spell it out, there will be a set of heavy vector resonances in the adjoint of G , including ten gauge bosons of $SO(5)$, an abelian resonance of $U(1)_X$, and eight massive $SU(3)_c$ resonances (befittingly called “heavy gluons”). In addition, there will be the massless gauge fields $G_\mu, W_\mu,$ and B_μ of the SM. This leaves four uneaten NGBs that constitute the Higgs doublet, parameterised by the matrix

$$\Omega = \Omega_1 \Omega_2. \quad (4.5)$$

The physical W^\pm and Z bosons are obtained after EWSB.

These mass eigenstates are not what we explicitly put into the Lagrangian. The gauge sector will instead consist of $G_\mu^0, W_\mu^0,$ and B_μ^0 - the elementary counterparts of the SM gauge fields - and $\rho_{G_\mu}, \rho_\mu,$ and ρ_{X_μ} - the $SU(3), SO(5),$ and $U(1)_X$ gauge fields on the composite site. The sixteen $SU(3)$ gauge bosons G_μ^0 and ρ_{G_μ} will mix into vector and axial combinations, giving respectively the SM gluons and the heavy gluons. Similarly, the remaining fields mix into the heavy $SO(5)$ and $U(1)_X$ resonances, and the electroweak gauge bosons, as above.

²In contrast to the previous chapter, here we start counting at Site 0 to fit with the convention of labelling elementary quantities with superscript zeros.

It is straightforward to write the bosonic part of the Lagrangian for the M4DCHM; indeed, it is simply that of Equation (3.192). To be explicit, we separate the Lagrangian into gauge and NGB contributions:

$$\mathcal{L}_{\text{boson}} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{NGB}}. \quad (4.6)$$

The gauge kinetic terms are simply

$$\begin{aligned} \mathcal{L}_{\text{gauge}} = & \left. \begin{aligned} & -\frac{1}{4}\text{Tr} [G_{\mu\nu}^0 G^{0\mu\nu}] - \frac{1}{4}\text{Tr} [W_{\mu\nu}^0 W^{0\mu\nu}] - \frac{1}{4}B_{\mu\nu}^0 B^{0\mu\nu} \end{aligned} \right\} \text{ elementary} \\ & \left. \begin{aligned} & -\frac{1}{4}\text{Tr} [\rho_{G_{\mu\nu}} \rho_G^{\mu\nu}] - \frac{1}{4}\text{Tr} [\rho_{\mu\nu} \rho^{\mu\nu}] - \frac{1}{4}\rho_{X_{\mu\nu}} \rho_X^{\mu\nu}, \end{aligned} \right\} \text{ composite} \end{aligned} \quad (4.7)$$

where the field strength tensors ($G_{\mu\nu}^0, \dots$) have the form

$$X_{\mu\nu} = \partial_\mu X_\nu - \partial_\nu X_\mu + ig_{(X)}[X_\mu, X_\nu] \quad (4.8)$$

from Equation (2.15).

As for the NGBs, recall the matrices Ω_k parameterise the NGBs of the breaking of $G^{k-1} \times G^k$ to its vector subgroup. To write the NGB Lagrangian, it will therefore be useful to first recast Ω_k into independent matrices for each simple group factor in G . In this set will be $\Omega_{1,X,G}$ respectively associated to the $SO(5)$, $U(1)_X$, and $SU(3)_c$ vector subgroup breakings between the elementary and composite sites, and also Ω_2 for the breaking $SO(5)^1 \rightarrow SO(4)$. Each will carry its own NGB decay constant $f_{1,2,X,G}$ that determines the scale of the associated symmetry breaking. By Section 3.8.1, the NGB matrices will transform under the symmetries at each site as

$$\begin{aligned} SO(5)^0 \times SO(5)^1 : \quad \Omega_1 &\rightarrow g_0 \Omega_1 g_1^{-1}, & U(1)_X^0 \times U(1)_X^1 : \quad \Omega_X &\rightarrow g_0 \Omega_X g_1^{-1}, \\ SO(5)^1 \times SO(4) : \quad \Omega_2 &\rightarrow g_1 \Omega_2 h^{-1}, & SU(3)_c^0 \times SU(3)_c^1 : \quad \Omega_G &\rightarrow g_0 \Omega_G g_1^{-1}, \end{aligned} \quad (4.9)$$

where transformations g_k come from Site k , and $h \in SO(4)$. Accordingly, their covariant derivatives are given by

$$\begin{aligned} D_\mu \Omega_1 &= \partial_\mu \Omega_1 - i(g_0 W_\mu^{0a} T_L^a + g'_0 B_\mu^0 T_R^3) \Omega_1 + ig_\rho \Omega_1 \rho_\mu, \\ D_\mu \Omega_2 &= \partial_\mu \Omega_2 - ig_\rho \rho_\mu \Omega_2, \\ D_\mu \Omega_X &= \partial_\mu \Omega_X - ig'_0 B_\mu^0 \Omega_X + ig_X \Omega_X \rho_{X_\mu}, \\ D_\mu \Omega_G &= \partial_\mu \Omega_G - ig_s^0 G_\mu^0 \Omega_G + ig_G \Omega_G \rho_{G_\mu}. \end{aligned} \quad (4.10)$$

We have introduced elementary gauge couplings g_0 , g'_0 , and g_s^0 mirroring the SM couplings of Equation (2.1), as well as new composite couplings g_ρ , g_X , and g_G of $SO(5)^1$, $U(1)_X^1$, and $SU(3)_c^1$. Notice there are no H gauge fields, since these have been decoupled through HLS. The generators $T^A = \{T_L^a, T_R^a, \hat{T}^a\}$ of $SO(5)$ have been

grouped into $SU(2)_L$, $SU(2)_R$, and coset components (see Appendix G), allowing the $SO(5)$ resonances to be expressed as

$$\rho_\mu = \rho_\mu^A T^A = \rho_{L\mu}^a T_L^a + \rho_{R\mu}^a T_R^a + \mathbf{a}_\mu^a \hat{T}^a. \quad (4.11)$$

While the $SO(5)$ generators have been written explicitly in Equation (4.10), the generators of $SU(3)_c$ are implicitly taken to be the usual (half) Gell-Mann matrices $\lambda^a/2$.

Since $SO(5)$ is spontaneously broken to $SO(4)$, there is a vacuum vector Φ_0 left invariant under $SO(4)$, so the product $\Omega_2 \Phi_0$ transforms as a fundamental of $SO(5)^1$. We will use a basis where $\Phi_0 = (0, 0, 0, 0, 1)^\top$, as before. With these, the leading-order NGB Lagrangian is

$$\mathcal{L}_{\text{NGB}} = \sum_{i=1, X, G} \frac{f_i^2}{4} \text{Tr} [(D_\mu \Omega_i)^\dagger (D^\mu \Omega_i)] + \frac{f_2^2}{2} (D_\mu \Omega_2 \Phi_0)^\top (D^\mu \Omega_2 \Phi_0), \quad (4.12)$$

using the normalisation found in Section 3.2.2 for the non-linear σ -model matrix Ω_2 . A clever choice of gauge allows this Lagrangian to be expressed more transparently. We will spend the gauge symmetry to move to the Site 0 holographic gauge, removing the unphysical NGBs by setting

$$\Omega_1 = e^{i \frac{\sqrt{2}}{f_1} h^a \hat{T}^a}, \quad \Omega_2 = \Omega_X = \Omega_G = \mathbb{1}, \quad (4.13)$$

where h^a are the Higgs doublet components. From here, the electroweak gauge symmetry can be used to go to the SM unitary gauge in which $h^a = (0, 0, 0, h)^\top$. Hereafter, all equations should be understood to be in this gauge unless otherwise stated. Now gauge boson mixing terms in Equation (4.12) become readily apparent:

$$\begin{aligned} \mathcal{L}_{\text{NGB}} = & \frac{f_1^2}{4} \text{Tr} [(D_\mu \Omega_1)^\dagger (D^\mu \Omega_1)] - \frac{f_2^2 g_\rho^2}{2} (\rho_\mu \Phi_0)^\top (\rho^\mu \Phi_0) \\ & + \frac{f_X^2}{4} (g'_0 B_\mu^0 - g_X \rho_{X\mu})^2 + \frac{f_G^2}{4} (g_s^0 G_\mu^0 - g_G \rho_{G\mu})^2. \end{aligned} \quad (4.14)$$

We see from Appendix F that in the unitary gauge the NGB matrix comes out to

$$\Omega \equiv \Omega_1 \Omega_2 = \Omega_1 = \begin{pmatrix} 1 & & & & \\ & 1 & & & \\ & & 1 & & \\ & & & \cos \frac{h}{f_1} & \sin \frac{h}{f_1} \\ & & & -\sin \frac{h}{f_1} & \cos \frac{h}{f_1} \end{pmatrix}, \quad (4.15)$$

so the first line of Equation (4.14) hides a mixing between \mathbf{a}_μ^4 and the Higgs:

$$\mathcal{L}_{\text{NGB}} \supset \frac{1}{2} (\partial_\mu h) (\partial^\mu h) + \frac{f_1}{\sqrt{2}} g_\rho \mathbf{a}_\mu^4 \partial^\mu h + \frac{g_\rho^2}{4} (f_1^2 + f_2^2) \mathbf{a}_\mu^4 \mathbf{a}^{4\mu}, \quad (4.16)$$

which also follows from Equation (3.59). This mixing is removed, keeping the Higgs kinetic term conventionally normalised, by redefining the fields

$$\mathbf{a}_\mu^4 \rightarrow \mathbf{a}_\mu^4 - \frac{\sqrt{2}}{g_\rho} \frac{f}{f_2^2} \partial_\mu h, \quad h \rightarrow \frac{f_1}{f} h, \quad \text{where} \quad \frac{1}{f^2} \equiv \frac{1}{f_1^2} + \frac{1}{f_2^2}. \quad (4.17)$$

Note that the Higgs interactions now depend only on the quantity

$$s_h := \sin\left(\frac{h}{f}\right). \quad (4.18)$$

The rest of the bosonic interactions may be gleaned from the presented information, though we will not explicitly expand the Lagrangian any further here.

4.3 Fermion Sector




Figure 4.1: Structure of couplings between elementary and composite fermions.

At this point, the remaining freedom in specifying a particular model lies in establishing the fermionic content of the composite site and the representations of $SO(5) \times U(1)_X$ under which it transforms. As mentioned in previous chapters, we restrict to the case where the composite sector mixes only with the third generation elementary quarks in the interests of reducing the parameter space of the theory. This is a reasonable limiting case, for the lighter fermions must mix with the composite sector only weakly if their left- and right-handed chiralities mix with roughly equal strengths.

We will be considering three different models in this work: the M4DCHM⁵⁻⁵⁻⁵, the M4DCHM¹⁴⁻¹⁴⁻¹⁰, and the M4DCHM¹⁴⁻¹⁻¹⁰, all of which have been outlined in Ref. [100]. The labels M4DCHM^{q-t-b} specify the representations $(\mathbf{q}, \mathbf{t}, \mathbf{b})$ of $SO(5)$ in which the composite partners $(\Psi^q, \tilde{\Psi}^t, \tilde{\Psi}^b)$ of the elementary $(q_L^0 = (t_L^0, b_L^0)^\top, t_R^0, b_R^0)$ transform. To help keep track, the structure of the elementary-composite interactions is depicted graphically in Figure 4.1.

Our reasoning for choosing these models stems largely from the considerations in Sections 3.6.1 and 3.7.1. Namely, because they all provide custodial protection for the $Zb_L\bar{b}_L$ coupling from tree-level corrections, and because it will be interesting to compare models with qualitatively different fine-tuning behaviours. The M4DCHM⁵⁻⁵⁻⁵, recall, suffers from a double tuning in the Higgs potential, while

the other models need not have such extra tuning. The M4DCHM¹⁴⁻¹⁴⁻¹⁰ and M4DCHM¹⁴⁻¹⁻¹⁰ each also offer interesting prospects, with the former having symmetries that allow for two independent proto-Yukawa couplings for the top quark, and the latter being able to feature an entirely composite right-handed top quark, since the interactions of the singlet t_R^0 would not violate the composite symmetry (though we do not take the right-handed top quark as composite in this work). These are not the only possible models that fit our criteria, but it has been found in Ref. [100] that the various other representation combinations either tend to give experimentally unfavourable predictions, or are qualitatively similar to the three we consider here.

In specifying the Lagrangian for the fermion sector of each model, we split up the contributions as

$$\mathcal{L}_{\text{fermion}}^{q-t-b} = \mathcal{L}_{\text{comp. quark}}^{q-t-b} + \mathcal{L}_{\text{elem. quark}} + \mathcal{L}_{\text{lepton}}. \quad (4.19)$$

The high-energy Lagrangian for the (partially) composite quarks, consisting of the third generation elementary quarks and their composite partners Ψ and $\tilde{\Psi}$, will be given below for each model. Elementary fields will be embedded into incomplete multiplets ψ in the same representation as their partners for convenience in writing gauge-invariant interactions. In all cases the covariant derivatives of the fields are given by

$$\begin{aligned} D_\mu \psi_R &= \left(\partial_\mu - ig'_0 B_\mu^0 Y - ig_s^0 G_\mu^{0a} \frac{\lambda^a}{2} \right) \psi_R, \\ D_\mu \psi_L &= \left(\partial_\mu - ig_0 W_\mu^{0a} T_L^a - ig'_0 B_\mu^0 Y - ig_s^0 G_\mu^{0a} \frac{\lambda^a}{2} \right) \psi_L, \\ D_\mu \Psi &= \left(\partial_\mu - ig_\rho \rho_\mu^A T^A - ig_X \rho_X X - ig_G \rho_G^a \frac{\lambda^a}{2} \right) \Psi, \end{aligned} \quad (4.20)$$


where the hypercharges Y are those from Table 2.2, and λ^a are the Gell-Mann matrices. It should be kept in mind that a multiplet Ψ in the symmetric **14** or antisymmetric **10** of $SO(5)$, when expressed as a matrix, is acted upon by the generators as

$$T^A \Psi = [T^A, \Psi]. \quad (4.21)$$

Refer to Appendix G for the explicit embeddings of fields into representations of $SO(5)$.

4.3.1 M4DCHM⁵⁻⁵⁻⁵

Unlike the other models we will be exploring, the M4DCHM⁵⁻⁵⁻⁵ requires *two* composite multiplets $\Psi^{t,b}$ to couple to q_L^0 , simply because a single partner multiplet does not provide sufficient couplings for both the top and the bottom quark to have mass. So instead of Figure 4.1, the structure of interactions is as in Figure 4.2. The partner


 Figure 4.2: Elementary and composite quark couplings in the M4DCHM⁵⁻⁵⁻⁵.

content of this model consists of Ψ^t , $\tilde{\Psi}^t$ in the $\mathbf{5}_{+\frac{2}{3}}$ representation of $SO(5) \times U(1)_X$, and Ψ^b , $\tilde{\Psi}^b \sim \mathbf{5}_{-\frac{1}{3}}$. Their interactions are specified by the quark Lagrangian

$$\begin{aligned}
 \mathcal{L}_{\text{comp. quark}}^{5-5-5} = & \bar{q}_L^0 i \not{D} q_L^0 + \bar{t}_R^0 i \not{D} t_R^0 + \bar{b}_R^0 i \not{D} b_R^0 & \left. \vphantom{\mathcal{L}_{\text{comp. quark}}^{5-5-5}} \right\} \text{ elementary} \\
 & + \bar{\Psi}^t (i \not{D} - m_t) \Psi^t + \bar{\tilde{\Psi}}^t (i \not{D} - m_{\tilde{t}}) \tilde{\Psi}^t & \left. \vphantom{\mathcal{L}_{\text{comp. quark}}^{5-5-5}} \right\} \text{ composite} \\
 & + \Delta_{tL} \bar{\psi}_L^t \Omega_1 \Psi^t_R + \Delta_{tR} \bar{\psi}_R^t \Omega_1 \tilde{\Psi}^t_L & \left. \vphantom{\mathcal{L}_{\text{comp. quark}}^{5-5-5}} \right\} \text{ link} \\
 & - m_{Y_t} \bar{\Psi}^t_L \tilde{\Psi}^t_R - Y_t \bar{\Psi}^t_L \Phi \Phi^\dagger \tilde{\Psi}^t_R & \left. \vphantom{\mathcal{L}_{\text{comp. quark}}^{5-5-5}} \right\} \text{ Yukawa} \\
 & + (t \rightarrow b) + \text{h.c.} & (4.22)
 \end{aligned}$$

To save space, we have defined $\Phi := \Omega_2 \Phi_0$. Here the third generation elementary quarks are embedded into the incomplete $SO(5)$ fundamentals

$$\psi_L^t = \frac{1}{\sqrt{2}} \begin{pmatrix} b_L^0 \\ -ib_L^0 \\ t_L^0 \\ it_L^0 \\ 0 \end{pmatrix}, \quad \psi_R^t = \begin{pmatrix} \vec{0} \\ t_R^0 \end{pmatrix}, \quad \psi_L^b = \frac{1}{\sqrt{2}} \begin{pmatrix} t_L^0 \\ it_L^0 \\ -b_L^0 \\ ib_L^0 \\ 0 \end{pmatrix}, \quad \psi_R^b = \begin{pmatrix} \vec{0} \\ b_R^0 \end{pmatrix}. \quad (4.23)$$

Note that this is not the most general Lagrangian obeying the symmetries of the theory. Terms such as $\psi_L \Omega_1 \tilde{\Psi}_R$ or $\tilde{\Psi}_L \Psi_R$, if included, would reduce the model to the two-site DCHM. In that case the Higgs potential is logarithmically divergent, so we avoid such couplings. Similar considerations exist for the other representations.

4.3.2 M4DCHM¹⁴⁻¹⁴⁻¹⁰

Here we have two composite multiplets, Ψ^a and $\tilde{\Psi}^t$, in the traceless symmetric $\mathbf{14}$ of $SO(5)$, and one multiplet $\tilde{\Psi}^b$ in the antisymmetric $\mathbf{10}$. Each composite multiplet carries an X charge of $+\frac{2}{3}$. All of them transform adjointly,

$$\Psi \xrightarrow{g \in SO(5)} g \Psi g^{-1}, \quad (4.24)$$

4.3.4 Elementary Fermions

In a truly consistent CHM, all of the lighter fermions would be incorporated through partial compositeness in a similar manner to the above. However in our work, we treat them as entirely elementary for simplicity. We reproduce the SM lepton sector with bilinear couplings to the low-energy Higgs pNGB:

$$\mathcal{L}_{\text{lepton}} = \sum_{\text{generations}} \left(i\bar{l}_L \not{D} l_L + i\bar{\ell}_R \not{D} \ell_R - \frac{m_{\text{SM}}}{v} \bar{l}_L \begin{pmatrix} 0 \\ h \end{pmatrix} \ell_R + \text{h.c.} \right). \quad (4.28)$$

The covariant derivatives here couple the leptons to the elementary gauge fields, with SM quantum numbers as in Equation (4.20). The elementary quarks, on the other hand, are given simple Dirac mass terms

$$\mathcal{L}_{\text{elem. quark}} = \sum_{q=u,d,c,s} \bar{q}(i\not{D} - m_q)q \quad (4.29)$$

with no Higgs couplings, and again coupling to the elementary gauge fields. Admittedly there is no reason to treat the mass terms of the elementary quarks and leptons differently. However, the difference is largely inconsequential due to the small couplings to the Higgs involved, and because there is little distinction between a mass term and a Higgs coupling below the EW scale.

4.4 Calculability

The information above completely specifies the models we are considering in this work. All of the interactions and observables of the theory may be calculated, with sufficient perseverance, by substituting the definitions of the multiplets and generators into the Lagrangians given here. Though we use all of these couplings for our work, we do not explicitly provide the coupling matrices because of the sheer number present. We do, however, give the mass-mixing matrices for each model in Appendix H, on account of their importance. With these matrices, the Weinberg Sum Rules Equations (3.214) and (3.215) can be directly verified in each model by calculating

$$\sum_{\text{particles } i} m_i^2(h) = \text{Tr} [M^\dagger M] \quad \text{and} \quad \sum_{\text{particles } i} m_i^4(h) = \text{Tr} [(M^\dagger M)^2], \quad (4.30)$$

and observing the quantities are independent of the Higgs field h for each individual mass matrix. It follows that the Higgs potential is calculable and given by Equation (3.216),

$$V_{\text{eff}}(h) = \sum_{\text{particles } i} \frac{c_i}{64\pi^2} m_i^4(h) \log(m_i^2(h)), \quad (4.31)$$

which can be minimised to find the Higgs vev $\langle h \rangle$. Once the vev is known, the Higgs mass is calculated as

$$m_h^2 = \partial_h^2 V_{\text{eff}}(h)|_{h=\langle h \rangle}, \quad (4.32)$$

and the masses of other particles can be calculated by diagonalising the mass matrices at the point $h = \langle h \rangle$. However, not every set of masses is sufficiently balanced to give the potential a non-trivial minimum, and indeed a large portion of parameter space is not phenomenologically viable because of this.

4.5 Integrating out Heavy Resonances

Up until this point, this thesis has been structured around using mass matrices to calculate observables. This method has its benefits - namely that it is accurate at all energies and is conceptually clear - but is somewhat slow for numerical calculations. Indeed, we initially used this method in our frequentist explorations of these models, but decided it would be more beneficial to switch to a quicker, approximate, approach for the Bayesian scans. To this end, we give here a brief summary of this alternative *form factor* based calculation method.

In the low-energy range that we are capable of probing, it is often useful to work with *effective* theories that retain only the low-energy degrees of freedom (*i.e.* the elementary fields), having the heavy composite fields “integrated out”. This is an approximation wherein the composite fields are set as satisfying their classical equations of motion, valid for energies well below the masses of the composite fields. Once this procedure is carried out (more details can be found in Ref. [100]), the effective fermionic Lagrangian, in momentum space, will be of the form

$$\begin{aligned} \mathcal{L}_{\text{comp. quark}}^{\text{eff}} = \sum_{\psi=t,b} & [\bar{\psi}_L^0 \not{p} (1 + \Pi_{\psi_L}(p^2)) \psi_L^0 + \bar{\psi}_R^0 \not{p} (1 + \Pi_{\psi_R}(p^2)) \psi_R^0 \\ & + \bar{\psi}_L^0 M_\psi(p^2) \psi_R^0 + \text{h.c.}] \end{aligned} \quad (4.33)$$

for some model-dependent functions Π_ψ and M_ψ , which we call form factors, that for our models are provided in Appendix I.

It is possible to calculate the effective potential in terms of these functions with largely the same method as in Section 3.7, but with the form factors in place of the mass matrices used there³. The result is that the fermionic contribution to the potential, with only partially composite third generation quarks, is

$$V_{\text{eff}}^{\text{fermion}}(h) = 2iN_c \sum_{\psi=t,b} \int \frac{d^4p}{(2\pi)^4} \ln [p^2 (1 + \Pi_{\psi_L}(p^2)) (1 + \Pi_{\psi_R}(p^2)) + |M_\psi(p^2)|^2], \quad (4.34)$$

³The steps are performed explicitly in Ref. [12].

where $N_c = 3$ is the number of colours. This expression is in Minkowski space, but the integral may be simplified with a Wick rotation to Euclidean space, $p^2 \rightarrow -p_E^2$, $d^4p \rightarrow id^4p_E$. The resulting integrand is spherically symmetric, so the trivial angular integrals give a factor of $2\pi^2$ and we are left with the one-dimensional integral

$$V_{\text{eff}}^{\text{fermion}}(h) = -2N_c \sum_{\psi=t,b} \int \frac{dp_E^2}{16\pi^2} \ln [(-p_E^2) (1 + \Pi_{\psi_L}(-p_E^2)) (1 + \Pi_{\psi_R}(-p_E^2)) + |M_\psi(-p_E^2)|^2]. \quad (4.35)$$

The benefit of this formalism is that from here it is quite simple to find the minimum of the potential numerically, since the integrand can be expanded in powers of s_h to obtain the coefficients γ , β defined previously:

$$V_{\text{eff}}(h) =: -\gamma s_h^2 + \beta s_h^4, \quad (4.36)$$

that dictate the minimum occurs at

$$s_{\langle h \rangle} = \frac{\gamma}{2\beta}. \quad (4.37)$$

Generally, a single one-dimensional integral is required to calculate each of γ and β (or at least the fermionic contributions to them), which is several orders of magnitude faster than minimising the potential numerically with the mass matrix approach.

We neglect detailing the incorporation of the gauge boson contributions to the potential in this framework, since they are of secondary importance. Suffice to say that the gauge boson contribution to γ can be calculated analytically, with the result that [30]

$$\gamma_{\text{gauge}} = -\frac{9m_\rho^4 (m_a^2 - m_\rho^2) t_\theta}{64\pi^2 (m_a^2 - (1 + t_\theta) m_\rho^2)} \ln \left[\frac{m_a^2}{(1 + t_\theta) m_\rho^2} \right] \quad (4.38)$$

to first order in

$$t_\theta := \frac{g_0}{g_\rho}. \quad (4.39)$$

Here we have the approximate masses of the lightest composite gauge bosons (see Section 4.6)

$$\begin{aligned} m_\rho^2 &:= \frac{1}{2} g_\rho^2 f_1^2, \\ m_a^2 &:= \frac{1}{2} g_\rho^2 (f_1^2 + f_2^2). \end{aligned} \quad (4.40)$$

We use this formula in our scans, and ignore the negligible contributions to β from the gauge bosons.

4.6 Particle Content

One interesting signature of CHMs is the existence of new resonances in the few-TeV mass range. In this section we give an overview of the particle content in each of the models we are considering, including some analysis of the particles' masses.

4.6.1 Boson Sector

As stated above, there will be eight heavy gluons, a heavy neutral abelian resonance, and ten $SO(5)$ gauge bosons in each realisation of the M4DCHM, along with the SM gauge bosons. Since the first two species are familiar from our understanding of the SM gauge sector, let us consider the $SO(5)$ resonances in more detail, following Ref. [102].

$SO(5)$ Gauge Bosons

Recall from Equation (4.11) that we have grouped the $SO(5)$ gauge fields as

$$\rho_\mu = \rho_\mu^A T^A = \rho_{L_\mu}^a T_L^a + \rho_{R_\mu}^a T_R^a + \mathbf{a}_\mu^a \hat{T}^a. \quad (4.41)$$

Notice that by the normalisation of the generators, we have

$$\rho_{L,R_\mu}^a = \text{Tr} [\rho_\mu T_{L,R}^a], \quad \mathbf{a}_\mu^a = \text{Tr} [\rho_\mu \hat{T}^a]. \quad (4.42)$$

It is straightforward to see how these transform. Under $h \in SO(4)$, for example,

$$\begin{aligned} \rho_{L,R_\mu}^a &\rightarrow \text{Tr} \left[\left(h \rho_\mu h^{-1} - \frac{i}{g_\rho} (\partial_\mu h) h^{-1} \right) h T_{L,R}^a h^{-1} \right] \\ &= h \rho_{L,R_\mu}^a h^{-1} - \frac{i}{g_\rho} \text{Tr} [h^{-1} (\partial_\mu h) T_{L,R}^a], \end{aligned} \quad (4.43)$$

recalling that a transformation is a simple change of basis for the generators. In the second term, $h^{-1}(\partial_\mu h)$ is in the Lie algebra of $SO(4)$, and so is a linear combination of the form

$$h^{-1}(\partial_\mu h) = [h^{-1}(\partial_\mu h)]_{L,a} T_L^a + [h^{-1}(\partial_\mu h)]_{R,a} T_R^a. \quad (4.44)$$

This reduces the transformation to

$$\rho_{L,R_\mu}^a \rightarrow h \rho_{L,R_\mu}^a h^{-1} - \frac{i}{g_\rho} [(\partial_\mu h) h^{-1}]_{L,R}^a. \quad (4.45)$$

We see ρ_{L,R_μ} transforms as a gauge field in the adjoint of $SU(2)_{L,R}$. That is, $\rho_{L_\mu} \sim (\mathbf{3}, \mathbf{1})$ and $\rho_{R_\mu} \sim (\mathbf{1}, \mathbf{3})$ under $SU(2)_L \times SU(2)_R$. With our knowledge of the electroweak sector, we recognise that for each group factor there will be a linear

combination

$$\rho_{L,R_\mu}^\pm := \frac{\rho_{L,R_\mu}^1 \mp i\rho_{L,R_\mu}^2}{\sqrt{2}} \quad (4.46)$$

of unit electric charge, along with an uncharged boson ρ_{L,R_μ}^3 . The axial fields, on the other hand, can be seen to transform as

$$\mathbf{a}_\mu \rightarrow h\mathbf{a}_\mu h^{-1}, \quad (4.47)$$

without the inhomogeneous term in Equation (4.45) because their generators do not overlap with $T_{L,R}^a$. This is the bidoublet $(\mathbf{2}, \mathbf{2})$ - the same representation as the Higgs doublet. We therefore recognise another unit-charged boson

$$\mathbf{a}_\mu^\pm := \frac{\mathbf{a}_\mu^1 \mp i\mathbf{a}_\mu^2}{\sqrt{2}}, \quad (4.48)$$

and two more neutral bosons $\mathbf{a}_\mu^3, \mathbf{a}_\mu^4$. Incidentally, this analysis confirms the decomposition $\mathbf{10} \rightarrow (\mathbf{3}, \mathbf{1}) \oplus (\mathbf{1}, \mathbf{3}) \oplus (\mathbf{2}, \mathbf{2})$ of Section 3.6.1.

Boson Masses

In total, we see that in addition to the SM gauge bosons there are the heavy gluons, five neutral bosons, and three bosons of unit electric charge. Their mass-mixing matrices are given as functions of the Lagrangian parameters in Appendix H. Singular values of these matrices at the point $h = \langle h \rangle$ are the squared tree-level masses of the resonances.

Unfortunately, the mass matrices are not simple enough to result in useful analytic expressions for all of the gauge boson masses. The only bosons with easily calculable masses are the heavy gluons and the neutral \mathbf{a}^4 resonance, which have squared masses of

$$m_G^2 = \frac{1}{2}f_G^2((g_s^0)^2 + g_G^2), \quad m_{\mathbf{a}^4}^2 = \frac{1}{2}\frac{f_1^2}{f_1^2 - f^2}f_1^2g_\rho^2, \quad (4.49)$$

and the charged bosons, whose mass spectrum is given by

$$m_{\text{charged}}^2 = \left\{ \frac{1}{2}f_1^2g_\rho^2, m_1^2, m_2^2, m_3^2 \right\}, \quad (4.50)$$

where the last three entries are the solutions to

$$m^2 \left(m^2 - \frac{1}{2}(g_0^2 + g_\rho^2)f_1^2 \right) \left(m^2 - \frac{g_\rho^2f_1^2}{2}\frac{f_1^2}{f_1^2 - f^2} \right) = \frac{1}{16}g_0^2g_\rho^4\frac{f_1^6f^2}{f_1^2 - f^2}s_{\langle h \rangle}^2. \quad (4.51)$$

These solutions do not have particularly understandable analytical forms, but useful formulae may be obtained by expanding the masses as power series in $s_{\langle h \rangle} :=$

$\sin(\langle h \rangle / f)$. The W mass, for example, distinguished as being $\mathcal{O}(s_{\langle h \rangle})$, is given by

$$m_W = \frac{1}{2} \tilde{g} f s_{\langle h \rangle} + \frac{1}{4f_1^2} \left(\frac{1}{g_0^2 + g_\rho^2} + \frac{f_1^2 - f^2}{g_\rho^2 f_1^2} \right) \tilde{g}^3 f^3 s_{\langle h \rangle}^3 + \mathcal{O}(s_{\langle h \rangle}^5), \quad (4.52)$$

where

$$\frac{1}{\tilde{g}^2} := \frac{1}{g_0^2} + \frac{1}{g_\rho^2}. \quad (4.53)$$

Since EW precision constraints typically require $s_{\langle h \rangle} \lesssim 0.3$, only the leading-order term in Equation (4.52) is of any relevance [44]. Neglecting the higher-order terms reproduces the W mass in the simplified model Equation (3.83),

$$m_W = \frac{1}{2} g f s_{\langle h \rangle}, \quad (4.54)$$

provided we match to the SM coupling⁴

$$g \equiv \tilde{g} = \frac{g_0 g_\rho}{\sqrt{g_0^2 + g_\rho^2}}. \quad (4.56)$$

This leads once again to the relation $s_{\langle h \rangle} = v/f$, but we stress that this is only accurate when the W mass is close to its experimental value $\frac{1}{2} g v$. The same sort of analysis can be done for the other charged resonances, but the contributions to their masses from the Higgs field are only relatively minor. We simply give their masses in the $s_{\langle h \rangle} \rightarrow 0$ limit, which will be approximately satisfied by realistic points and can be used as a zeroth order estimate:

$$\lim_{s_{\langle h \rangle} \rightarrow 0} m_{\text{charged}}^2 = \left\{ 0, \frac{1}{2} f_1^2 g_\rho^2, \frac{1}{2} f_1^2 (g_0^2 + g_\rho^2), \frac{1}{2} \frac{f_1^2}{f_1^2 - f^2} f_1^2 g_\rho^2 \right\}. \quad (4.57)$$

This same limit can be used to understand the neutral boson masses (including the photon and Z boson, and \mathbf{a}^4):

$$\lim_{s_{\langle h \rangle} \rightarrow 0} m_{\text{neutral}}^2 = \left\{ 0, 0, \frac{1}{2} f_1^2 (g_0^2 + g_\rho^2), \frac{1}{2} \frac{f_1^2}{f_1^2 - f^2} f_1^2 g_\rho^2, \right. \\ \left. \frac{1}{2} \frac{f_1^2}{f_1^2 - f^2} f_1^2 g_\rho^2, \frac{1}{2} M_\pm^2 (f_1 g_\rho, f_X g_X, f_X g_\rho, f_1 g_0') \right\}. \quad (4.58)$$

⁴Similar matchings can also be done to yield relations for the remaining SM gauge couplings in terms of the elementary and composite couplings:

$$\frac{1}{g'^2} = \frac{1}{g_0'^2} + \frac{1}{g_\rho^2} + \frac{1}{g_X^2}, \quad \text{and} \quad \frac{1}{g_s^2} = \frac{1}{(g_s^0)^2} + \frac{1}{g_G^2}, \quad (4.55)$$

which we use throughout this work.

Here we have defined the functions⁵

$$M_{\pm}^2(x_1, x_2, x_3, x_4) := \frac{|\vec{x}|^2}{2} \pm \sqrt{\frac{|\vec{x}|^4}{4} - (x_1^2 x_2^2 + x_2^2 x_4^2 + x_3^2 x_4^2)}, \quad (4.59)$$

which will also be useful in expressing the fermion masses.

Note there is some degeneracy among the masses of the charged and neutral bosons in the $s_{\langle h \rangle} \rightarrow 0$ limit. In fact, numerical calculation of the masses for realistic parameter points with small non-zero $s_{\langle h \rangle}$ suggests that two charged bosons will always separately be very close in mass to two neutral bosons, and the last charged boson will be very close in mass to two other neutral bosons (“very close” here meaning within $\sim 0.1\%$).

4.6.2 Fermion sector

The fermion sector of each model can be analysed using the explicit embeddings of fields in the different representations of $SO(5)$ given in Appendix G. We group the particles according to their electric charge, which is determined through the relation

$$Q = T_L^3 + Y = T_L^3 + T_R^3 + X. \quad (4.60)$$

It was seen in Section 4.3 that these models deliver some up-type (U) and down-type (D) resonances, along with particles of exotic charge $Q_{4/3}$, $Q_{5/3}$, and $Q_{8/3}$. The number of new particles in each model is listed in Table 4.1.

Model	U	D	$Q_{4/3}$	$Q_{5/3}$	$Q_{8/3}$
M4DCHM ⁵⁻⁵⁻⁵	8	8	2	2	0
M4DCHM ¹⁴⁻¹⁴⁻¹⁰	16	9	2	9	2
M4DCHM ¹⁴⁻¹⁻¹⁰	11	6	1	6	1

Table 4.1: Number of BSM up-type (U) and down-type (D) particles, and particles Q_x that have exotic electric charge x in each model.

The (tree-level) masses of the fermions are found as the singular values of the fermion mass matrices in Appendix H. For convenience, the non-SM fermion masses in each model are provided in Figure 4.3. Similarly to the bosons, many fermion masses are not easily expressible analytically, and we resort to giving some masses in the $s_{\langle h \rangle} \rightarrow 0$ limit. Note the many degeneracies and approximate degeneracies among the particle masses in each model.

⁵For some intuition about the masses $M_{\pm}(m_1, m_2, m_Y, \Delta)$: if $\Delta = 0$ then one mass will be less than $(\min(m_1^2, m_2^2) + m_Y^2/2)^{1/2}$ and the other greater than $(\max(m_1^2, m_2^2) + m_Y^2/2)^{1/2}$. If m_Y is also small, then the two masses come out nearly equal to m_1 and m_2 . In general both masses increase as Δ increases in magnitude.

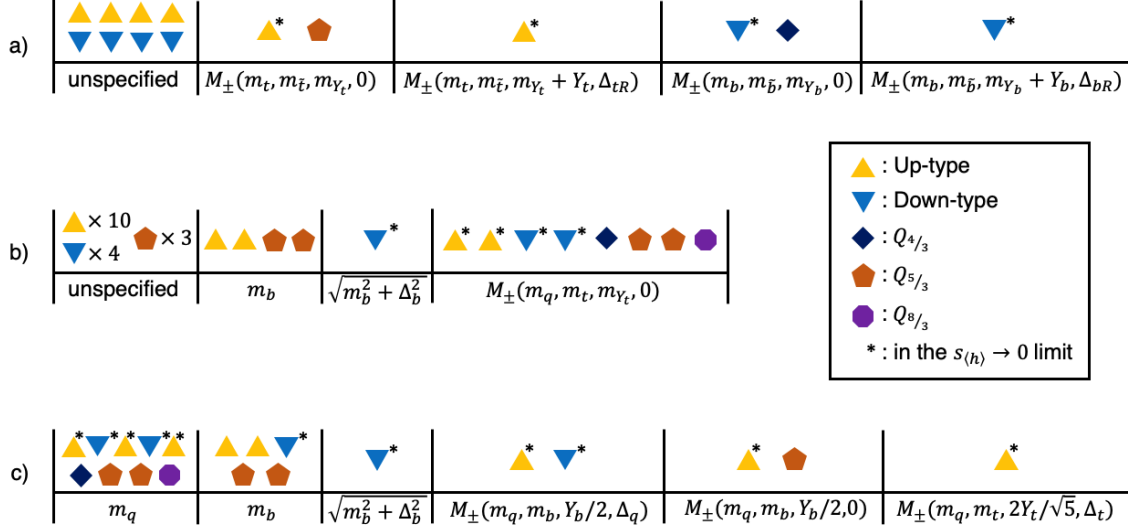


Figure 4.3: Tree-level masses of the non-SM fermions (a) in the M4DCHM⁵⁻⁵⁻⁵, (b) in the M4DCHM¹⁴⁻¹⁴⁻¹⁰, and (c) in the M4DCHM¹⁴⁻¹⁻¹⁰. Each icon corresponds to one particle, or to a pair of particles for masses given in terms of the functions M_{\pm} defined by Equation (4.59). Some masses cannot be expressed in a reasonably understandable form, in which case they are left unspecified.

Of particular interest are the top and bottom quark masses, which we present to first order in $s_{\langle h \rangle}$ for each model in Equations (H.23) to (H.25). With the approximation $s_{\langle h \rangle} = v/f$, these expressions can be used to estimate the SM quark masses without needing to minimise the Higgs potential, providing a computationally inexpensive way to constrain the various parameter spaces.

4.7 Status of Model Explorations

Being the simplest experimentally viable composite Higgs models, M4DCHMs have been the subject of considerable theoretical research, and we provide here a brief summary of this previous work to put our fits of these models into greater context.

When this class of models was first proposed in Ref. [99], the Higgs boson had not been officially discovered. Working with only a partially composite top quark for simplicity, the main focus of Ref. [99] was on the behaviour of the Higgs potential and the relation of the Higgs mass to the mass m_F of the lightest fermionic composite partners in the M4DCHM⁵⁻⁵. It was estimated that in this model, the Higgs mass is roughly proportional to m_F/f , so that for lower tunings (smaller f), lighter composite partners are preferable. This was numerically verified with a simple scan over the fermionic parameters for fixed values of f and other gauge parameters. It was also noted that for EWSB to take place in this model, the two chiralities of the top quark must share a similar compositeness, $\Delta_{tL}/m_t \approx \Delta_{tR}/m_{\bar{t}}$, which constrains the parameter space somewhat.

After the discovery of the Higgs boson, Ref. [87] went further and considered our three models above, though again with only a partially composite top quark. Here the tunings of the models were examined theoretically and verified with some numerical scans that found points with realistic Higgs masses, again with simplifying assumptions on some of the parameters. It was shown that all three models generically predict too large a Higgs mass unless the typical mass scale m_* of the composite fermions is relatively low, $m_* \sim f$. It was also pointed out that under this condition, and as can be seen from Equation (3.143), if all the fermionic parameters are also of order f , then the double tuning in the M4DCHM⁵⁻⁵ disappears! But the condition $m_* \sim f$ is not necessary in all cases: a realistic Higgs mass is also possible in the M4DCHM⁵⁻⁵ when the composite mass spectrum is non-generic and there is an anomalously light resonance, but this requires some fine-tuning. Similarly, the Higgs mass can be accommodated in the M4DCHM¹⁴⁻¹ without any light partners if there is some extra tuning in the Higgs potential of around the same size as in double tuning. It was concluded that for a moderate tuning, there must be composite partners of roughly 1 TeV to reproduce the observed Higgs mass in all three models.

From there, more substantial explorations of the various M4DCHMs began. Ref. [100] performed parameter scans for many different models, including the exact ones we focus on (though without the composite $SU(3)$ sector), with a focus on the Higgs phenomenology and tuning in each model. As with the previous work, these scans used simplifying assumptions on the parameters, setting all mass parameters equal and imposing relations between the gauge parameters. There were several interesting findings. First, it was found that the production cross section for the Higgs boson from gluon-gluon fusion was always suppressed compared to the SM prediction for realistic points, even though this is not a requirement of the theory. Second, the Higgs branching ratios may be suppressed or enhanced compared to the SM predictions. The combined effect is that the total cross section for the gluon-gluon fusion production of a Higgs boson and its subsequent decay into some final state (quantities that we take as constraints in our fits, but that had not been measured when Ref. [100] was published) is typically lower in M4DCHMs than in the SM, unless the branching ratio of the Higgs boson into that final state is greatly enhanced. And lastly, the M4DCHM⁵⁻⁵⁻⁵ was found to have considerably lower tuning than the M4DCHM¹⁴⁻¹⁴⁻¹⁰ and M4DCHM¹⁴⁻¹⁻¹⁰, primarily because the latter two tend to predict too heavy a Higgs. This is in agreement with Ref. [87], showing that perhaps the double tuning is not an important factor.

In all of these studies, the parameter spaces were explored with random scans that did not attempt to optimise for points that better fit with experiment. Such a fitting procedure was first undertaken in Ref. [101], which employed an incredibly large variety of observables as constraints to find those points that are consistent with EW precision tests, Higgs physics, flavour physics, and (Run I) LHC resonance searches. But the models explored in Ref. [101] were rather unwieldy, for they took into account the partial compositeness of *all* quarks. Ref. [30] is similar, fitting

points to the observed particle masses in models with all third generation fermions (including leptons) partially composite. Neither of these works purported to have performed rigorous, convergent statistical fits of these models, which would have simply been impractical given their large (40+ dimensional) parameter spaces.

The purpose of this thesis is to extend this research with the first full convergent fits of the M4DCHMs presented in this chapter. This is made possible through the use of some sophisticated sampling algorithms, detailed in Chapter 5, and because our models have been chosen to be sufficiently simplified so as to have somewhat manageable parameter spaces, while still containing enough physics to be interesting.

Part II

Global Fits of the Minimal Composite Higgs Models

Bayesian and Frequentist Global Fits

It is important to be clear about what we hope to achieve by fitting the minimal composite Higgs models specified in Chapter 4. Ostensibly, we aim to see whether such models can successfully describe our world - that is, whether there are any values of the model parameters that lead to accurate predictions of a wide range of experimental observations - and to map the regions of parameter space that best fit those observations. Such a task is well-suited to a **frequentist** global fit, which assigns to each parameter point a likelihood based on how well it reproduces experiment, and finds those parameter points that maximise the likelihood. But this is not all we are interested in. After all, the Standard Model gives incredibly accurate predictions, but is still quite unfavourable because it is unnatural: its accuracy is present over only a small region of its parameter space. Just as we judge the SM by its naturalness, so too must we do the same for the composite Higgs models. If one of the models could reproduce observations to an acceptable level of accuracy over a broad region of its parameter space, it would be more favourable than the SM under the metric of naturalness, even if its predictions are not quite as accurate as those of the SM. This notion of weighing the accuracy of a model's predictions against the volume of parameter space in which those predictions are accurate is formalised in the framework of **Bayesian statistics** by a quantity called the **evidence** of a model. It is this Bayesian, rather than frequentist, viewpoint, that interests us most.

We give a brief overview of Bayesian statistics in Section 5.1, and the algorithm we use to perform Bayesian global fits of the M4DCHMs in Section 5.1.1. We also fit the models under a frequentist framework, whose theory is described in Section 5.2, using a differential evolution algorithm detailed in Section 5.2.1.

5.1 Bayesian Statistics

Bayesian statistics, at its core, is a framework to quantify how one's *degree of belief* in some statement should be updated when one is presented with new information. It is based on the simple observation that if $P(A)$ is the probability (degree of belief) of a proposition A being true, and $P(A|B)$ is the probability that A is true given

that B is true, then

$$P(A \text{ and } B) \equiv P(A|B)P(B). \quad (5.1)$$

The symmetry between A and B on the left-hand side implies **Bayes' theorem**:

$$P(A|B) = \frac{P(B|A)}{P(B)}P(A). \quad (5.2)$$

In other words, if we initially assign A a probability $P(A)$ of being true prior to knowing that B is true, and then we learn that B is true, then our belief in A must update to $P(A|B)$ by this formula. $P(A)$ is suitably called the **prior**, and $P(A|B)$ the **posterior** probability.

In the context of our global fits, this framework allows us to quantify our degree of belief that the universe is described by a particular parameter point \mathbf{p} in one of our models \mathcal{M} by combining our (suggestively labelled)

- (A) initial degree of belief that the universe would be described by \mathbf{p} , simply based on the probability of picking \mathbf{p} out of the parameter space of the model \mathcal{M} , and
- (B) experimental data about the universe (conditioned on \mathcal{M}).

To make the notation more convenient, let us denote the prior $P(A)$ - the probability density of picking \mathbf{p} out of the parameter space of \mathcal{M} - as $\pi(\mathbf{p}|\mathcal{M})$, and let us denote our experimental data (conditioned on \mathcal{M}) by $(\mathcal{D}|\mathcal{M})$. The probability of observing the data given the specific parameter point in the model, $P(\mathcal{D}|\mathbf{p}, \mathcal{M})$, is simply the likelihood $\mathcal{L}(\mathbf{p})$ that would be assigned to \mathbf{p} in a frequentist global fit - see Section 5.2. With this notation, the posterior probability that \mathbf{p} describes our universe given the data is then

$$P(\mathbf{p}|\mathcal{D}, \mathcal{M}) = \frac{\mathcal{L}(\mathbf{p})\pi(\mathbf{p}|\mathcal{M})}{P(\mathcal{D}|\mathcal{M})}. \quad (5.3)$$

Ultimately, the definitions of the functions on the right-hand side are subjective choices. If little is known about the parameter space, one might choose to assign a *flat prior* such that every point is given an equal probability of being picked out. Or if the even the *scale* of a parameter is unknown, it might be advantageous to use a *logarithmic prior* to assign equal weight across each order of magnitude of the parameter (thereby also rewarding smaller absolute values), and so on. It is hoped that by going through the process of using Bayes' theorem, the prior and the likelihood function work together so that the posterior distribution will be largely independent of the prior. How we define the likelihood function is detailed in Section 6.3.

There is one more factor in Equation (5.3) that we have so far not discussed: the total probability $P(\mathcal{D}|\mathcal{M})$ of observing the experimental data if the universe is described by the model \mathcal{M} . The problem of assigning this value to a model is usually pushed aside because it would require knowledge of the entire parameter

space, and because it is simply a normalisation factor for the posterior. But it can be calculated in principle if the posterior is conventionally normalised so that its integral over the (n -dimensional) parameter space is unity, in effect fixing

$$P(\mathcal{D}|\mathcal{M}) = \int d^n p \mathcal{L}(\mathbf{p})\pi(\mathbf{p}|\mathcal{M}). \quad (5.4)$$

Notice that $P(\mathcal{D}|\mathcal{M})$ is calculated by weighing the fitness of each point against its prior likelihood in parameter space, exactly matching how the naturalness of a model is quantified! This is the **Bayesian evidence** mentioned above, typically denoted as \mathcal{Z} .

Beware: it is commonly thought that a model with a greater evidence is more likely (in a Bayesian sense) to describe our universe, but this not necessarily correct. The evidence is the probability of observing the *data* given the *model*, but to compare models one needs instead the probability of the *model* being true given the *data*. The two are related by Bayes' theorem:

$$P(\mathcal{M}|\mathcal{D}) = \frac{P(\mathcal{D}|\mathcal{M})}{P(\mathcal{D})}P(\mathcal{M}). \quad (5.5)$$

Typically, out of ignorance, one would assign the same prior likelihood $P(\mathcal{M})$ to each model being considered, so a larger evidence *would* imply a more likely model, but this does not need to be the case. It should also be kept in mind that the evidence is not a uniquely defined number for a given model because it is dependent on the prior, so comparing models based on their evidences is an inexact science. The traditional interpretation, known as the **Jeffreys scale**, is that a model comparison is inconclusive if the evidence of one model is within a factor of three of the evidence of another, while outside a factor of 100 the comparison is decisive, with varying degrees of confidence in between [103].

Let us consider the evidence Equation (5.4) further, assuming a one-dimensional parameter space for simplicity. If it is the case, as is quite common, that the posterior is peaked around a parameter value p^{\max} , with width Δp , then the evidence will be approximately

$$\mathcal{Z} \approx \mathcal{L}(p^{\max})\pi(p^{\max}|\mathcal{M})\Delta p. \quad (5.6)$$

The so-called **Occam factor** $\pi(p^{\max}|\mathcal{M})\Delta p$ is less than unity because of the normalisation of the prior, so the evidence is less than $\mathcal{L}(p^{\max})$. This generalises if the posterior peaks across multiple dimensions. We can see from this that the Bayesian framework penalises models whose posteriors are sharply peaked in parameter space, and also penalises the addition of parameters to a model if the additional parameters do not sufficiently improve the fit to the data [104].

Finally, a helpful quantity that may be defined is the **Kullback-Leibler** (KL) **divergence** between the posterior and the prior,

$$D_{\text{KL}} := \int d^n p P(\mathbf{p}|\mathcal{D}, \mathcal{M}) \ln \left(\frac{P(\mathbf{p}|\mathcal{D}, \mathcal{M})}{\pi(\mathbf{p}|\mathcal{M})} \right), \quad (5.7)$$

which measures, in a sense, the extra information gained when going from the prior to the posterior [105]. Using Bayes' theorem and the normalisation of the posterior, this may be rewritten as

$$D_{\text{KL}} = \int d^n p P(\mathbf{p}|\mathcal{D}, \mathcal{M}) \ln(\mathcal{L}(\mathbf{p})) - \ln(\mathcal{Z}), \quad (5.8)$$

or rearranging for the average log-likelihood weighted by the posterior [106, 107],

$$\langle \ln(\mathcal{L}) \rangle_P = D_{\text{KL}} + \ln(\mathcal{Z}). \quad (5.9)$$

Once the evidence and KL divergence are known, this is a useful way to calculate the posterior-averaged log-likelihood as a quick check of how well the model fits the data.

In performing a Bayesian global fit, the main quantities of interest are the posterior likelihood of the parameter space, and the Bayesian evidence. Since our models are high-dimensional, we will actually describe results in terms of the **marginalised posterior** of each parameter: one-dimensional functions obtained by integrating the posterior over all other parameters, giving the relative probability distribution (in a Bayesian sense) of the leftover parameter. And the evidence, as we have seen, quantifies how well (in a Bayesian sense) the model describes our universe.

5.1.1 Nested Sampling

While the theory of Bayesian global fits is straightforward, the fact that the evidence is given by an integral of the likelihood over the parameter space means that in many cases it is too computationally expensive to calculate. There is an algorithm for scanning the space, however, that simplifies the calculation immensely, known as **nested sampling** [108]. It makes use of the clever observation that if we define a variable X as a function of the likelihood by

$$X(L) := \int_{\mathcal{L}(\mathbf{p}) > L} d^n p \pi(\mathbf{p}|\mathcal{M}), \quad (5.10)$$

then the evidence Equation (5.4) can be written as the *one-dimensional* integral

$$\mathcal{Z} = \int_0^1 dX \mathcal{L}(X), \quad (5.11)$$

given the prior is normalised such that

$$X(0) = \int d^n p \pi(\mathbf{p}|\mathcal{M}) = 1. \quad (5.12)$$

Geometrically, $X(L)$ is the prior-weighted volume of the region of parameter space in which the likelihood is greater than L . As long as the prior-weighted volumes X_i are known for a number of increasing likelihood cutoff values L_i , the evidence Equation (5.11) can be approximated as the simple Riemann sum

$$\mathcal{Z} \approx \sum_i (X_{i-1} - X_i) L_i. \quad (5.13)$$

The problem of actually finding these values X_i is ordinarily highly non-trivial, but is elegantly solved by the nested sampling algorithm.

In nested sampling, a certain number n_{live} of “live” points are generated, sampled according to the prior distribution $\pi(\mathbf{p}|\mathcal{M})$. An iterative process is then used to find regions of higher likelihood. Once the likelihoods of all the live points have been calculated at iteration i , the lowest-likelihood point is removed (and classified as “dead”) and replaced with another live point drawn according to the prior, under the condition that the new point has a higher likelihood. The likelihood of the newly dead point is taken as L_i , so X_i will approximately be the prior-weighted volume of the remaining live points, assuming sufficient coverage of the parameter space.

Now here is the key point: since all points are sampled from the prior, the prior-weighted volume decreases on average by a factor of $n_{\text{live}}/(n_{\text{live}} + 1)$ each time a new live point is generated¹. So at iteration i , the average² volume is simply

$$\langle X_i \rangle = \left(\frac{n_{\text{live}}}{n_{\text{live}} + 1} \right)^i. \quad (5.14)$$

Nested sampling algorithms use these average values in place of the exact X_i . The error this introduces to the calculated evidence is derived in, *e.g.*, Ref. [109].

As the process continues, higher likelihoods will be found, with the live points squeezing into ever-decreasing volumes. Eventually the volumes X_i will become small enough for the remaining contributions to the evidence to become negligible, at which point the sampling will stop and the algorithm is said to have **converged**. For our scans, we set the algorithm to terminate when the evidence carried by the remaining live points

$$\mathcal{Z}_{\text{live}} \approx X_i \langle \mathcal{L} \rangle_{\text{live}} \quad (5.15)$$

falls below 0.001 times that contributed by the dead points.

¹To see this, one only needs to consider, what amounts to the same thing, n_{live} evenly distributed points on a unit line segment.

²Average taken over repeated runs with different initial live points.

There are many particular realisations of nested sampling algorithms, differing in the methods they use to generate new live points that satisfy the likelihood criterion $\mathcal{L} > L_i$. For this work we use the recently developed nested sampling package `PolyChord`, which uses a highly sophisticated version of *slice sampling* to generate points efficiently. Explaining the particularities of `PolyChord` would take us too far afield from the focus of this work; we simply use the default settings recommended in the original papers [109, 110].

5.2 Frequentist Fits

In contrast to Bayesian fits, frequentist fits seek only to maximise the likelihood function \mathcal{L} over the parameter space, with no regard to any prior prejudices about the space. This approach is useful for finding those points that best fit experimental constraints, even if those points exist in a finely-tuned region.

The object of interest here is solely the likelihood function, dictating how well each point agrees with experimental data. Details about the likelihood function we use for our fits are given later in Section 6.3, but for our current purposes it suffices to know that we take the likelihood as being Gaussian in the data:

$$\mathcal{L}(\mathbf{p}) = e^{-\frac{1}{2}\chi^2(\mathbf{p})}, \quad \chi^2(\mathbf{p}) = (\mathcal{O}^{\text{theo}}(\mathbf{p}) - \mathcal{O}^{\text{exp}})^\top C^{-1}(\mathcal{O}^{\text{theo}}(\mathbf{p}) - \mathcal{O}^{\text{exp}}), \quad (5.16)$$

where $\mathcal{O}^{\text{theo}}(\mathbf{p})$ and \mathcal{O}^{exp} are the predicted and experimental values of the considered observables, and C is the covariance matrix.

In showing the results of our frequentist fits, we will not marginalise over the likelihood as we do for the posterior in the Bayesian fits, but instead give the **profile likelihood ratio** of each parameter. The profile likelihood ratio Λ of a parameter p_i at a value \tilde{p}_i is simply the maximum value taken by the likelihood when $p_i = \tilde{p}_i$, normalised by the global maximum likelihood:

$$\Lambda(\tilde{p}_i) := \frac{\max_{\mathbf{p}|p_i=\tilde{p}_i} \mathcal{L}(\mathbf{p})}{\max_{\mathbf{p}} \mathcal{L}(\mathbf{p})}. \quad (5.17)$$

A virtue of the profile likelihood ratio is that it provides a straightforward method of estimating confidence intervals for each parameter. Specifically, it follows from Wilk's theorem that if the true global maximum of the likelihood has been found in the sampling space, the $1 - \alpha$ confidence intervals for the parameter p_i can be approximated as those values \tilde{p}_i for which the test statistic

$$\left(\min_{\mathbf{p}|p_i=\tilde{p}_i} \chi^2(\mathbf{p}) \right) - \left(\min_{\mathbf{p}} \chi^2(\mathbf{p}) \right) \quad (5.18)$$

fails to be statistically significant under the null hypothesis that $p_i = \tilde{p}_i$, where α is the significance level [111]. From the definition of Λ , this simply translates to those values \tilde{p}_i for which $\Lambda(\tilde{p}_i)$ lies above some α -dependent value.

We warn that technically this procedure is only correct for finding confidence intervals for the *data*, since the likelihood function is Gaussian in the data, while it prescribes confidence intervals to the *input parameters* that are only approximate, due to the non-linear transformations relating the data to the input parameters. We do not perform coverage tests to determine the accuracy of the confidence intervals we find for the input parameters in this way on account of the computing expense required, but it has been shown to be a fairly good approximation even in quite pathological cases [112, 113].

Of course when assessing the validity of a model it is not much help to know the *relative* likelihoods of parameter points; we would also like to know whether any of the points are likely at all. To this end we place particular focus on experimentally viable “valid” points, which we define as those that satisfy all individual experimental constraints at the 3σ level. That is, a point is considered valid if $\chi_i^2 < 9$ for each individual constraint³.

5.2.1 Differential Evolution

For our frequentist fits we opt to use a genetic optimisation algorithm known as **Differential Evolution** (DE), which has proven particularly successful at quickly exploring large parameter spaces and optimising difficult likelihood functions [14]. This type of algorithm takes a specified number of parameter points and from them “breeds” successive “generations” of points that better satisfy the experimental constraints in a manner analogous to natural selection. Aggressively optimising variants of DE with self-adaptive parameters have been put forward [115], and we use one such version known as λ jDE that is provided by the `Diver` package [116].

Like all optimisation algorithms, the objective of λ jDE is to maximise a given function - in our case, our likelihood function \mathcal{L} - over a specified parameter space. To do this, the algorithm begins with an initial “population” of N points in parameter space, generated randomly with a given probability distribution, that constitutes generation $G = 1$. Given a population of points \mathbf{p}_i^G ($i = 1, \dots, N$) in generation G , the next generation is calculated using the following three-step process, represented schematically in Figure 5.1.

The first step, *mutation*, generates a population of mutant vectors \mathbf{m}_i^G . In λ jDE, the mutant vectors are constructed as

$$\mathbf{m}_i^G = \lambda_i^G \mathbf{p}_{\text{best}}^G + (1 - \lambda_i^G) \mathbf{p}_{a_i}^G + F_i^G (\mathbf{p}_{b_i}^G - \mathbf{p}_{c_i}^G). \quad (5.19)$$

Here, a_i , b_i , and c_i are random integers distinct from each other and from i , and $\mathbf{p}_{\text{best}}^G$ is the vector with the highest likelihood of those in generation G . The constants $\lambda_i^G \in [0, 1]$ and $F_i^G \in [0.1, 0.9]$ have randomly assigned values for $G = 1$, which self-adapt through subsequent generations. In each successive generation each constant

³Some constraints will have non-zero minimum possible χ^2 values, which are subtracted from the χ^2 s when calculating their 3σ bounds. Also, the oblique S and T constraints, being correlated, have instead a 3σ bound of $\chi^2 = 11.83$ [114]. See Section 6.3.

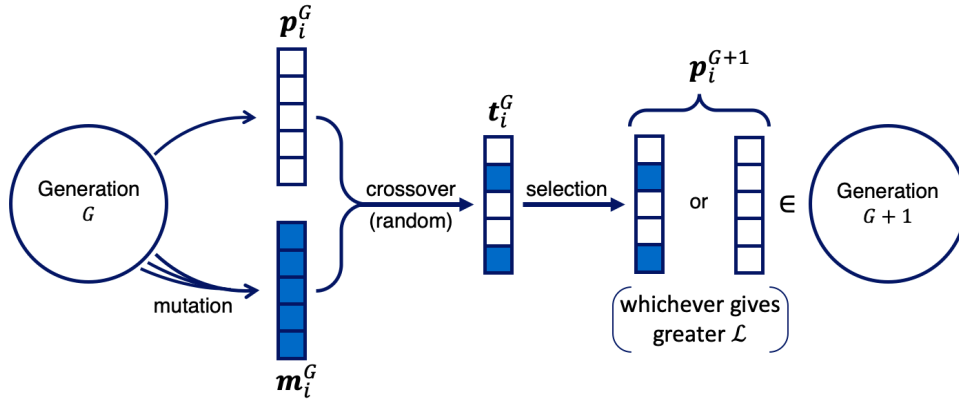


Figure 5.1: Basic form of the differential evolution algorithm.

has a 10% chance of taking on a new random value, and otherwise adopts the value of the corresponding constant from an earlier generation that was used in the construction of \mathbf{p}_i^G .

Mutant vectors do not form the next generation directly. Instead, a trial vector \mathbf{t}_i^G for the next generation is formed from each pair $(\mathbf{p}_i^G, \mathbf{m}_i^G)$ in the *crossover* stage. The trial vectors are formed as follows.

1. A random component of \mathbf{t}_i^G is set equal to the corresponding component of \mathbf{m}_i^G .
2. Each other component of \mathbf{t}_i^G is set to the corresponding component of either \mathbf{m}_i^G or \mathbf{p}_i^G , having a fixed probability $P_i^G \in [0, 1]$ of using \mathbf{m}_i^G .

The probability P_i^G is calculated using the same procedure as for λ_i^G . If a trial vector lies outside of the parameter space, we reflect the offending components back into the space.

The final step, *selection*, simply sets the i^{th} vector in the next generation, \mathbf{p}_i^{G+1} , equal to whichever point out of $(\mathbf{p}_i^G, \mathbf{t}_i^G)$ has the highest likelihood. Iterating these three steps, a sequence of generations is produced with populations that successively migrate towards regions that maximise the likelihood function. The benefit provided by calculating generation G is measured by the fractional improvement

$$1 - \frac{\sum_i \ln(\mathcal{L}(\mathbf{p}_i^{G-1}))}{\sum_i \ln(\mathcal{L}(\mathbf{p}_i^G))}. \quad (5.20)$$

The process stops and is said to have **converged** when the average fractional improvement provided by the last few generations falls below some given amount, called the **convergence threshold**. We choose to use the last ten generations in taking this average, recommended by Ref. [116] as a reasonable number to avoid premature convergence. As for the remaining settings, for each of our scans we use a large population size of $N = 50,000$ points, along with a rather strong convergence threshold of 10^{-5} .

Note that by no means does convergence mean a global optimum has been found, only that the explored regions are sufficiently likely that the algorithm struggles to find significantly more likely regions. With that said, DE is often successful at finding global optima, even for badly behaved likelihood functions. We choose to use DE for this reason, and also because it may be particularly well suited for exploring CHMs: since the algorithm uses the *differences* of population vectors when constructing mutant vectors (see Equation (5.19)), it tends to generate points along various hyperplanes on which the populations lie, which may provide the necessary cancellations for EWSB take place. This is only heuristic reasoning, however, and we have not tested whether DE is more efficient at generating points that trigger EWSB than other algorithms.

Global Fit Procedures

In this chapter we detail exactly how we perform our global fits of the three models specified in Chapter 4. We discuss the model parameters and the bounds we impose on them for the scans in Section 6.1. The procedure for calculating observables from the parameters is outlined in Section 6.2, and the experimental constraints used to determine the validity of a given parameter point are discussed in Section 6.3.

6.1 Scan Parameters

For convenience, the Lagrangian parameters for each model, defined in Sections 4.2 and 4.3, have been collected in Table 6.1.

M4DCHM	5 – 5 – 5	14 – 14 – 10	14 – 1 – 10
Decay constants	f, f_1, f_X, f_G	f, f_1, f_X, f_G	f, f_1, f_X, f_G
Gauge couplings	g_ρ, g_X, g_G	g_ρ, g_X, g_G	g_ρ, g_X, g_G
Link couplings	$\Delta_{t_L}, \Delta_{t_R}, \Delta_{b_L}, \Delta_{b_R}$	$\Delta_q, \Delta_t, \Delta_b$	$\Delta_q, \Delta_t, \Delta_b$
On-diagonal masses	$m_t, m_{\bar{t}}, m_b, m_{\bar{b}}$	m_q, m_t, m_b	m_q, m_t, m_b
Off-diagonal masses	m_{Y_t}, m_{Y_b}	m_{Y_t}	
Proto-Yukawa couplings	Y_t, Y_b	Y_t, Y_b, \tilde{Y}_t	Y_t, Y_b
Dimensionality	19	17	15

Table 6.1: Parameters present in each model.

At the onset of this project, we were performing our fits by simply scanning over the parameters as they appear in Table 6.1 in units of GeV, subject to some conditions specified below. While a valid approach, we were finding that our scans (using only `Diver` during this time) were having trouble getting sufficient coverage around the best-fit regions to provide consistent convergent results. The main complication was that small changes in the Higgs decay constant, f , need to be compensated by very particular changes in the other parameters so as not to drastically affect the Higgs potential. For points that reproduce the correct EWSB scale v , varying the f parameter by even 5 GeV while keeping the other parameters fixed leads to wildly

unrealistic EWSB scales. This fact makes satisfactory exploration of the parameter spaces particularly difficult, as has been noted in previous works [101, 117].

Fortunately, there is a far more clever parameterisation of the models avoiding this problem that has been utilised in previous explorations of CHMs [30, 88, 100]. It rests on the fact that the Higgs potential Equation (3.216) is a function of the particle masses, *in any units of mass*. So if all of the mass-dimension parameters are expressed in units of f , the Higgs potential will not depend on the actual value of f at all! The potential could be minimised to find the misalignment $s_{\langle h \rangle}$ as usual, with the mass scale only being set later by *defining*

$$f \equiv \frac{v}{s_{\langle h \rangle}} = \frac{246}{s_{\langle h \rangle}} \text{ GeV}, \quad (6.1)$$

in accordance with Equation (3.84). Notice that not only does this remove f as a scan parameter, this automatically reproduces the correct EWSB scale, and by extension the experimental W and Z masses. Regrettably, we did not consider this approach until after having performed the frequentist scans, so this scanning method was only used with the Bayesian fits. The Bayesian results are therefore considered to be the main results of this thesis, and the frequentist results only secondary.

Our treatment of the parameters in each scanning approach is summarised in Tables 6.2 and 6.3, and explained in greater detail below.

NGB Decay Constants

The Higgs decay constant f is one of the more consequential parameters in the theory, defining the approximate mass scale $m_* \sim g_\rho f$ as well as the naïve energy cutoff $\Lambda_f = 4\pi f$. For our frequentist fits, where we take f as an input, we choose to scan over the wide range

$$f \in [0.5 \text{ TeV}, 5.0 \text{ TeV}]. \quad (6.2)$$

The Bayesian fits instead take f as a derived quantity, and so do not scan over f directly.

The ranges of other parameters are limited by the value of f . For example, f_1 must be greater than f by virtue of Equation (4.17), and less than $\sqrt{3}f$ to maintain partial unitarisation of NGB scattering [98]. The other decay constants are constrained by

$$\frac{f_1}{2} \leq f_{X,G} \leq 2f_1 \quad (6.3)$$

to avoid decoupling any resonances. We use these constraints to define the bounds for these parameters, with the added condition that all are greater than 0.5 TeV. Points that are generated within the bounds but that do not satisfy these consistency conditions are immediately discarded - a process that simply amounts to a modification of the prior such that only consistent points are admitted.

	Parameters	Scan Range	Prior
All Models	f	[0.5 TeV, 5.0 TeV]	Uniform
	f_1	[0.5 TeV, $\sqrt{3}f_{\max}$]	
	f_X, f_G	[0.5 TeV, $2\sqrt{3}f_{\max}$]	
	g_ρ, g_X, g_G	[1.0, 4π]	
M4DCHM ⁵⁻⁵⁻⁵	Δ_{bL}, Δ_{bR}	[$e^{2.0}$ GeV, $\Lambda_{f_{\max}}$]	Logarithmic
	Δ_{tL}, Δ_{tR}	[$e^{5.0}$ GeV, $\Lambda_{f_{\max}}$]	
	$m_t, m_{\tilde{t}}, m_b, m_{\tilde{b}}$	[0.5 TeV, $\Lambda_{f_{\max}}$]	
	m_{Y_t}, m_{Y_b}	[1.0 GeV, $\Lambda_{f_{\max}}$]	
	$m_{Y_t} + Y_t, m_{Y_b} + Y_b$	[1.0 GeV, $2\Lambda_{f_{\max}}$]	
M4DCHM ¹⁴⁻¹⁴⁻¹⁰	Δ_q	[$e^{6.0}$ GeV, $\Lambda_{f_{\max}}$]	Logarithmic
	Δ_t	[$e^{5.5}$ GeV, $\Lambda_{f_{\max}}$]	
	Δ_b	[$e^{3.5}$ GeV, $\Lambda_{f_{\max}}$]	
	m_q, m_t, m_b	[0.5 TeV, $\Lambda_{f_{\max}}$]	
	m_{Y_t}	[1.0 GeV, e^9 GeV]	
	$m_{Y_t} + \frac{1}{2}Y_t$	[1.0 GeV, $1.5\Lambda_{f_{\max}}$]	
	$m_{Y_t} + \frac{4}{5}(Y_t + \tilde{Y}_t)$	[$e^{6.0}$ GeV, $2.6\Lambda_{f_{\max}}$]	
Y_b	[$e^{3.5}$ GeV, $\Lambda_{f_{\max}}$]		
M4DCHM ¹⁴⁻¹⁻¹⁰	Δ_q, Δ_t	[$e^{7.0}$ GeV, $\Lambda_{f_{\max}}$]	Logarithmic
	Δ_b	[$e^{5.0}$ GeV, $\Lambda_{f_{\max}}$]	
	m_q, m_b	[0.5 TeV, $\Lambda_{f_{\max}}$]	
	m_t	[1.0 GeV, $\Lambda_{f_{\max}}$]	
	Y_t	[$e^{6.0}$ GeV, $\Lambda_{f_{\max}}$]	
	Y_b	[$e^{3.0}$ GeV, $\Lambda_{f_{\max}}$]	

Table 6.2: Ranges and priors used for the parameters in our frequentist fits. Here, $f_{\max} = 5.0$ TeV and $\Lambda_{f_{\max}} = 4\pi f_{\max}$. Some parameters are subject to further constraints specified in the main text. Note that the priors in these frequentist scans dictate only the sampling density of the parameters, and should have no effect on the profile likelihood distributions the scans find as long as the sampling density provides sufficient coverage of the space.

	Parameters	Scan Range	Prior
All Models	$m_\rho/f, m_a/f$	$[1/\sqrt{2}, 4\pi]$	Uniform
	$f_X/f, f_G/f$	$[0.5, 2\sqrt{3}]$	
	g_ρ, g_X, g_G	$[1.0, 4\pi]$	
M4DCHM ⁵⁻⁵⁻⁵	$\Delta_{t_L}/f, m_{Y_b}/f$	$[e^{-0.25}, 4\pi]$	Logarithmic
	Δ_{t_R}/f	$[e^{-0.75}, 4\pi]$	
	$\Delta_{b_L}/f, m_{Y_t}/f$	$[e^{-8.50}, 4\pi]$	
	Δ_{b_R}/f	$[e^{-1.25}, 4\pi]$	
	$m_t/f, m_b/f, m_{\tilde{b}}/f$	$[e^{-0.50}, 4\pi]$	
	$m_{\tilde{t}}/f$	$[e^{-1.00}, 4\pi]$	
	$(m_{Y_t} + Y_t)/f$	$[e^{-0.75}, 8\pi]$	
$(m_{Y_b} + Y_b)/f$	$[e^{-8.50}, 8\pi]$		
M4DCHM ¹⁴⁻¹⁴⁻¹⁰	$\Delta_q/f, \Delta_t/f, m_q/f, m_b/f$	$[e^{-2.5}, 4\pi]$	Logarithmic
	$\Delta_b/f, m_t/f, m_{Y_t}/f, Y_b/f$	$[e^{-8.5}, 4\pi]$	
	$(m_{Y_t} + \frac{1}{2}Y_t)/f$	$[e^{-5.2}, 1.5 \times 4\pi]$	
	$(m_{Y_t} + \frac{4}{5}(Y_t + \tilde{Y}_t))/f$	$[e^{-8.5}, 2.6 \times 4\pi]$	
M4DCHM ¹⁴⁻¹⁻¹⁰	$\Delta_q/f, \Delta_t/f, Y_t/f$	$[e^{-5.0}, 4\pi]$	Logarithmic
	Δ_b/f	$[e^{-7.0}, 4\pi]$	
	$m_q/f, m_b/f$	$[e^{-3.0}, 4\pi]$	
	m_t/f	$[e^{-9.0}, 4\pi]$	
	Y_b/f	$[e^{-6.0}, 4\pi]$	

Table 6.3: Ranges and priors used for the parameters in our Bayesian fits. The masses m_ρ and m_a , defined in Equation (4.40), reparameterise the decay constants into a form more suitable for calculating the gauge boson contribution to the Higgs potential. The normalisation factor f is determined after the potential is minimised. Some parameters are subject to further constraints specified in the main text.

Gauge Couplings

Given that the composite sector is strongly coupled and we can only perform calculations in the semi-perturbative regime, we must take the gauge couplings in Table 6.1 to be between 1 and 4π . In addition to these couplings, we also vary the SM gauge couplings within their experimental limits (at the scale of the top mass). The elementary gauge couplings are then determined by the relations in Equation (4.55). For g_s^0 to be real we require that $g_G > g_s$. We also impose the restrictions

$$\frac{1}{\sqrt{2}}f_1g_\rho < \Lambda_f, \quad \frac{1}{\sqrt{2}}f_Xg_X < \Lambda_f, \quad \frac{1}{\sqrt{2}}f_Gg_G < \Lambda_f, \quad (6.4)$$

to avoid vector resonance masses above the cutoff Λ_f .

Link, Mass, and Proto-Yukawa Couplings

The remaining parameters in Table 6.1 have dimensions of mass and so are given upper bounds of Λ_f . Instead of directly scanning over the off-diagonal masses and proto-Yukawa couplings in Table 6.1, however, it turns out to be more convenient to scan over particular linear combinations that appear in the mass matrices and form factors. Specifically, we scan over

$$\begin{aligned} \text{M4DCHM}^{5-5-5} & : m_{Y_u}, \quad m_{Y_d}, \quad m_{Y_u} + Y_u, \quad m_{Y_d} + Y_d, \\ \text{M4DCHM}^{14-14-10} & : m_{Y_u}, \quad Y_d, \quad m_{Y_u} + \frac{1}{2}Y_u, \quad m_{Y_u} + \frac{4}{5}(Y_u + \tilde{Y}_u), \\ \text{M4DCHM}^{14-1-10} & : Y_u, \quad Y_d. \end{aligned} \quad (6.5)$$

All of these parameters that we scan over are taken to be positive through field redefinitions, and in our initial test scans they were given arbitrarily chosen lower bounds of 1 GeV. These ranges were further restricted into those listed in Tables 6.2 and 6.3 based on the results of these tests. It was also apparent that the experimental constraints were drawing many of these parameters towards lower values, so scanning them with logarithmic priors was found to be advantageous.

We further constrain these parameters by discarding all points for which any fermion mass listed in Figure 4.3 is below 500 GeV, since the existence of such light resonances has essentially been ruled out. This is done without problem in PolyChord, but for Diver we needed to modify the source code so that the algorithm will keep generating trial vectors until all pass this check. This modification has made its way into the official release of Diver, in version 1.0.5.

6.2 Calculation of Observables

In setting up our models, we chose to use a convenient basis of fields in which it was straightforward to write Lagrangians obeying the symmetries of the theory. But since *physical* fields are mixtures of the fields in this *gauge basis* (gb), the Lagrangians as written are not conducive to the calculation of physical observables.

For this, the fields must be rotated to the physical *mass basis* (mb), in which the mass matrices are diagonal (*after* EWSB), so that the Feynman rules of the theory may become known. To illustrate, a generic coupling between two fermions Ψ^a, Ψ^b and a vector field A_μ^c can be written in terms of the gauge basis coupling matrix $g^{(gb)}$ as

$$\mathcal{L} \supset g_{abc}^{(gb)} \bar{\Psi}^a \gamma^\mu \Psi^b A_\mu^c. \quad (6.6)$$

In general, rotating to the mass basis requires unitary transformations

$$\begin{aligned} \Psi_{L,R}^{(mb)} &= U_{L,R}^\Psi \Psi_{L,R}^{(gb)}, \\ A_\mu^{(mb)} &= U^A A_\mu^{(gb)}, \end{aligned} \quad (6.7)$$

so that when expressed in the mass basis, the coupling matrix (for each handedness) is given by

$$g_{abc}^{(mb)} = (U^\Psi)_{ai}^\dagger g_{ijk}^{(gb)} U_{jb}^\Psi U_{kc}^A. \quad (6.8)$$

Similar considerations also exist for the other types of couplings, of course. It is these mass basis couplings from which the Feynman rules may be read.

The steps for calculating observables are therefore as follows. First, the Higgs potential must be minimised to find the Higgs vev for a given parameter point. In practice this can only be done numerically, through either the mass matrix or the form factor method, since there is no tractable exact formula for the Higgs potential as a function of the model parameters. The Higgs vev is then substituted into the mass matrices, which are then diagonalised (again, numerically) to obtain the particle masses after EWSB. In the process, the unitary matrices U^Ψ and U^A that diagonalise the mass matrices become known, and these are used to calculate the couplings in the mass basis given all of the gauge basis couplings. The resulting Feynman rules may then be applied to calculate any number of observables of interest.

For this whole procedure we use the Python package `pypngb` developed for Refs. [101, 117], and discussed in greater detail in Refs. [65, 118]. Given the coupling matrices of the models as inputs to the code, `pypngb` will minimise the Higgs potential, diagonalise the mass matrices, and calculate a large number of observables for a given parameter point. If the potential is already known (by using the form factor method, for example), the Higgs mass and vev may be passed as inputs to `pypngb` to skip the minimisation step.

Most of the (non-mass) observables of interest in this work make use of the partial width $\Gamma(R \rightarrow ij)$ of a resonance R decaying to a final state of two particles i, j . For particles of mass m_R, m_i, m_j , the partial width is calculated as [114, 119]

$$\Gamma(R \rightarrow ij) = \frac{\sqrt{m_R^4 + m_i^4 + m_j^4 - 2(m_R^2 m_i^2 + m_R^2 m_j^2 + m_i^2 m_j^2)}}{16\pi m_R^3} |\mathcal{M}|^2, \quad (6.9)$$

where $|\overline{\mathcal{M}}|^2$ is the squared amplitude for the decay, averaged over the initial states while summed over the final states. For explicit expressions for the amplitudes $|\overline{\mathcal{M}}|^2$ in particular processes, refer to Ref. [65]. From the partial widths, we may calculate the total width Γ_R of the resonance (assuming only two-body decays are relevant)

$$\Gamma_R = \sum_{i,j} \Gamma(R \rightarrow ij), \quad (6.10)$$

as well as its branching ratio into i and j ,

$$\text{BR}(R \rightarrow ij) = \frac{\Gamma(R \rightarrow ij)}{\Gamma_R}. \quad (6.11)$$

We note that in calculating the partial widths, `pypngb` takes into account only tree-level processes. Theoretical uncertainties are assigned to some observables to account for this simplification.

6.3 Constraints

We constrain our models to reproduce a wide range of observables in accordance with experiment. Namely, we employ SM masses, the EWSB energy scale, EW precision observables, Z boson decays, Higgs signal strengths, and also bounds on the production of new heavy resonances from direct collider searches as constraints. Our treatment of the constraints builds on that of Ref. [101] (and the subsequent modifications of Ref. [117]), so we only give a brief overview of each constraint below and refer the reader to Section 3.1 of Ref. [101] for further details and discussion. The experimental values have been updated for this work and are given in Table 6.4.

Finding the parameter values that best fit these observations is equivalent to finding those points \mathbf{p} that maximise some likelihood function \mathcal{L} that compares the predicted values of the observables $\mathcal{O}^{\text{theo}}(\mathbf{p})$ to the experimental values \mathcal{O}^{exp} . We model the likelihood as a multivariate Gaussian function in the observables:

$$\mathcal{L}(\mathbf{p}) = e^{-\frac{1}{2}\chi^2(\mathbf{p})}, \quad \chi^2(\mathbf{p}) = (\mathcal{O}^{\text{theo}}(\mathbf{p}) - \mathcal{O}^{\text{exp}})^\top C^{-1} (\mathcal{O}^{\text{theo}}(\mathbf{p}) - \mathcal{O}^{\text{exp}}), \quad (6.12)$$

where C is the covariance matrix taking into account uncertainties, both theoretical and experimental, and also correlations between observables. Most observables \mathcal{O}_i are not correlated with any other, having the simple (additive) contribution

$$\chi_i^2(\mathbf{p}) = \frac{(\mathcal{O}_i^{\text{theo}}(\mathbf{p}) - \mathcal{O}_i^{\text{exp}})^2}{\sigma_i^2}, \quad (6.13)$$

while for two observables that share a correlation ρ , the relevant submatrix is

$$C^{-1} = \frac{1}{\sigma_i^2 \sigma_j^2 (1 - \rho^2)} \begin{pmatrix} \sigma_j^2 & -\rho \sigma_i \sigma_j \\ -\rho \sigma_i \sigma_j & \sigma_i^2 \end{pmatrix}. \quad (6.14)$$

Observable	Value(s)	Ref.
m_b	4.18(4) GeV	[114]
m_t	$(173.0(4) \pm 1.0^{\text{theory}})$ GeV	[114, 120]
m_W	80.379(12) GeV	[114]
m_Z	91.1876(21) GeV	[114]
m_H	125.18(16) GeV	[114]
G_μ	$1.1663787(6) \times 10^{-5}$ GeV ⁻²	[114]
S	0.02(10)	[114]
T	0.07(12)	[114]
R_e	20.804(50)	[121]
R_μ	20.785(33)	[121]
R_τ	20.764(45)	[121]
R_b	0.21629(66)	[121]
$\mu_{\tau\tau}^{gg}$	1.0(6), 1.05(50), 0.97(56)	[122–124]
μ_{WW}^{gg}	0.84(17), 1.35(20), 1.20(20)	[122–124]
μ_{ZZ}^{gg}	1.13(33), 1.22(22), 1.03(16)	[122–124]
$\mu_{\gamma\gamma}^{gg}$	1.10(23), 1.15(15), 0.97(15)	[122, 124, 125]

Table 6.4: Experimental values used for constraints in our scans. The observables are grouped (in order) into SM masses, the EWSB scale (Fermi constant), oblique parameters, Z decay ratios, and Higgs signal strengths. The latter employ multiple independent measurements.

Here, σ_i is the total uncertainty of \mathcal{O}_i , obtained by summing its experimental and theoretical uncertainties in quadrature. Note that because of the multiple independent measurements of the Higgs signal strengths, the minimum possible χ^2 is 5.405. This simply acts as an inconsequential normalisation factor for the likelihood.

SM Masses and the EWSB Scale

All of the masses of SM particles that are predictions of the theory are used as constraints. This includes the top and bottom masses, the W and Z masses, and the Higgs mass. Once the Higgs mass is found by the method outlined at the end of Section 4.3, the other masses are found at tree level by diagonalising the mass matrices in Appendix H at the point $h = \langle h \rangle$. The top and bottom quarks are identified as the third lightest up- and down-type particles in the theory, and similarly the W is identified as the lightest charged boson and the Z as the lightest massive neutral boson.

We use the formulae Equations (H.23) to (H.25) to target the regions of parameter space that give realistic top and bottom masses, using the estimate $s_{\langle h \rangle} = v/f$.

Because this estimate is only accurate for points that give the correct W mass, it is unfortunately not effective to use these formulae to eliminate two free parameters in favour of the known quark masses - doing so makes the scans converge to regions in which none of the SM masses are reasonable. Instead, we constrain the parameter space by only considering those points for which we estimate

$$30 \text{ GeV} \leq m_t \leq 350 \text{ GeV}, \quad (6.15)$$

$$0.1 \text{ GeV} \leq m_b \leq 8 \text{ GeV}. \quad (6.16)$$

These bounds were chosen so as to cover a vast majority of those points that provide reasonable quark masses while excluding a large portion that do not, based on some tests of randomly sampled points across every model.

We interpret the calculated masses as $\overline{\text{MS}}$ running masses at the scale m_t . The bottom mass is then run to the scale m_b using `RunDec` [126] in order to compare it to experiment. We assign theoretical uncertainties of 5% to the top, W , Z , and Higgs masses, and of 1% to the bottom mass to conservatively account for these assumptions.

In addition to the correct SM masses, we also constrain points to reproduce the correct EWSB scale v . This is accomplished by using the Fermi constant for muon decays as a constraint, which at tree level is given by the relation

$$G_\mu^{\text{tree}} \equiv \frac{1}{\sqrt{2}v^2} = \frac{1}{\sqrt{2}f^2 s_{\langle h \rangle}^2}, \quad (6.17)$$

having used Equation (3.84). A theoretical uncertainty of 1% is included in calculating this observable to account for only using the tree-level value. Note that in the more sophisticated scanning method we use for our Bayesian fits, the W and Z mass constraints, and the EWSB scale constraint, are automatically satisfied.

Oblique Parameters

Important constraints for CHMs come from EW precision observables, which restrict the non-linear dynamics of the pNGB Higgs and place lower bounds on composite resonance masses. Such observables are conveniently parameterised by the Peskin-Takeuchi S and T parameters [127, 128], which we take as constraints. These “oblique” parameters measure the vacuum polarisation of the electroweak bosons, and as such quantify the effects of any new physics that affects all fermion generations equally.

The T parameter is defined in terms of the transverse parts of the W and Z vacuum polarisations as

$$T = \frac{1}{\alpha_{\text{em}}} \left(\frac{\Pi_{WW}^T(p^2 = 0)}{m_W^2} - \frac{\Pi_{ZZ}^T(p^2 = 0)}{m_Z^2} \right), \quad (6.18)$$

where α_{em} is the electromagnetic fine-structure constant. This parameter receives no tree-level contributions in our models as a result of the custodial symmetry from the $SO(5)/SO(4)$ coset [69]. We use the tree-level masses and calculate the form factors Π_{VV}^T at one-loop level, only including the dominant fermion contributions. An absolute theoretical uncertainty of 0.10 is assigned to this calculation of T to accommodate the typical size of the neglected gauge and NGB loop contributions.

The S parameter, on the other hand, is calculated at tree-level by the formula

$$\alpha_{\text{em}}S = \frac{1}{4} \left(1 - \frac{m_W^2}{m_Z^2} - \frac{1}{2} \frac{g_{Zee}^R}{g_{Zee}^R - g_{Zee}^L} \right), \quad (6.19)$$

where $g_{Zee}^{L,R}$ are the couplings of the Z boson to left- and right-handed electrons. This relation for S is only valid in models where T receives no tree-level contributions. We give a theoretical uncertainty of 0.05 to S to account for neglected loop effects.

Note the values for S and T in Table 6.4 have a correlation coefficient of +0.92 that is taken into account in the χ^2 calculation. The theoretical uncertainties are assumed to be uncorrelated.

Z Decays

While the oblique parameters are sufficient to constrain corrections to observables relating to the elementary fermions, they are not sensitive to observables relating to the significantly composite third generation quarks, for these receive non-universal corrections. The compositeness of the third generation quarks is instead tightly constrained by the $Zb_L\bar{b}_L$ coupling, which has been measured precisely at the LEP [121]. To this end we employ the observables

$$R_b := \frac{\Gamma(Z \rightarrow b\bar{b})}{\Gamma_{\text{had}}}, \quad R_\ell := \frac{\Gamma_{\text{had}}}{\Gamma(Z \rightarrow \ell\bar{\ell})}, \quad (6.20)$$

(for $\ell = e, \mu, \tau$), where

$$\Gamma_{\text{had}} = \sum_{q=u,d,c,s,b} \Gamma(Z \rightarrow q\bar{q}) \quad (6.21)$$

is the total Z hadronic width. Since all fermions here are taken to be entirely elementary except for the bottom quark, all of these observables are effectively constraints on the $Zb_L\bar{b}_L$ coupling.

The partial widths $\Gamma(Z \rightarrow x\bar{x})$ are calculated at leading order, and are numerically modified with NNLO contributions from the SM [129] so that the SM predictions [130]

$$R_b^{\text{SM}} = 0.215835, \quad R_e^{\text{SM}} = 20.7389, \quad R_\mu^{\text{SM}} = 20.7391, \quad R_\tau^{\text{SM}} = 20.7860 \quad (6.22)$$

are accurately reproduced when the composite sector is decoupled.

Higgs Signal Strengths

We use Higgs signal strengths to further constrain the non-linear dynamics of the pNGB Higgs. Specifically, we use the signal strengths from gluon fusion production, which is the dominant Higgs production mechanism and is well measured at the LHC. The signal strength μ_X^{gg} for the decay into a final state X is defined as the ratio of the measured cross section to the predicted SM value:

$$\mu_X^{gg} := \frac{[\sigma(gg \rightarrow h)\text{BR}(h \rightarrow X)]_{\text{exp}}}{[\sigma(gg \rightarrow h)\text{BR}(h \rightarrow X)]_{\text{SM}}}. \quad (6.23)$$

We include as observables μ_X^{gg} for $X = \tau\tau$, WW , ZZ , and $\gamma\gamma$. The partial widths for decays into massive particles are calculated at tree level, and those for decays into massless particles are calculated at one-loop order including all fermionic and bosonic contributions. The production cross section σ is calculated with the formula Equation (6.25) given later.

For each of these observables we use three independent measurements - one using combined ATLAS and CMS data from Run I of the LHC, and the others separately from ATLAS and CMS using Run II data.

Collider Searches

Searches for heavy resonances in colliders place useful upper bounds on the production cross section times branching ratio of composite resonances for various decay modes. For each composite resonance R , this value is calculated for all relevant SM final states (i, j) to compare with the upper bounds. Note that since these searches do not constrain an observable to any particular value but rather only provide upper bounds, there is a slight departure from Equation (6.13) in calculating the χ^2 contribution of each decay. In this case, the “experimental value” $\mathcal{O}_i^{\text{exp}}$ is taken to be the Gaussian central value of the expected upper bound and the observed upper bound on the cross section at the mass of the appropriate resonance. If the observed bound is weaker than expected, *i.e.* there is an excess of detected events, the χ^2 contribution is indeed given by Equation (6.13) so as to favour those resonances that could explain the surplus events. However, if the observed bound is stronger than expected, *i.e.* there are less events detected than expected, the contribution is instead

$$\chi_i^2 = \frac{(\mathcal{O}_i^{\text{theo}}(\mathbf{p}) - \mathcal{O}_i^{\text{exp}})^2 - (\mathcal{O}_i^{\text{exp}})^2}{\sigma_i^2}, \quad (6.24)$$

so that a vanishing cross section gives the highest likelihood. Here, σ_i is the Gaussian standard deviation of the bound¹.

¹Assuming the collider analyses give a 95% confidence level bound, this is simply 1/1.96 times the expected bound.

All of the experimental analyses that provide constraints we use are featured in Tables J.1 to J.5 in Appendix J. This builds on the list of analyses used in Ref. [117] with an additional 40 LHC searches at $\sqrt{s} = 13$ TeV. We have digitised the bounds provided by these analyses using the Python package `collim` developed by Peter Stangl. For reference, the upper bounds of various processes provided by the 13 TeV analyses are shown in Figures 6.1 and 6.2.

The χ^2 associated to each decay can be calculated with `collim`, but unfortunately it is rather slow to perform these calculations for all of the decays in our models. Including these direct search constraints in our scans increases the time spent evaluating each point from roughly 0.02 seconds to 0.2 seconds. We included these constraints in our frequentist fits, but due to time constraints opted not to do so for the Bayesian fits, instead only analysing them for the dead points after the scans converged.

For the quark partners we only employ constraints from analyses of pair-produced fermionic resonances. This is because the pair-production cross sections can be found quickly by interpolating the model-independent NNLO values provided by the `Hathor` package [131]. When comparing a production cross section times branching ratio to the upper bound from one of these analyses, the calculated branching ratio is appropriately modified depending on whether the analysis required one, both, or at least one of the resonances in the pair to decay.

Vector resonance production cross sections are calculated with the narrow width approximation (NWA) [114, 132]

$$\sigma(pp \rightarrow R) = \frac{16\pi^2 S_R c_R}{m_R} \sum_{i,j} \frac{1 + \delta_{ij}}{S_i S_j c_i c_j} \Gamma(R \rightarrow ij) \frac{\mathcal{L}_{ij}(s, m_R)}{s}, \quad (6.25)$$

where S_i and c_i are respectively the number of polarisations and colours of particle i , m_R is the resonance mass, \sqrt{s} is the centre of mass energy of the collider, and $\mathcal{L}_{ij}(s, m_R)$ is the parton luminosity of partons i, j with centre of mass energy m_R in a proton-proton collision with collider energy \sqrt{s} . Refer to Ref. [65] for more information on this calculation. The contribution to the χ^2 from the bosonic resonance R is suppressed by the factor

$$\frac{4}{\pi^2} \arctan^2 \left(\frac{m_R}{10\Gamma_R} \right), \quad (6.26)$$

so that the contribution is negligible when the NWA is not reasonably valid.

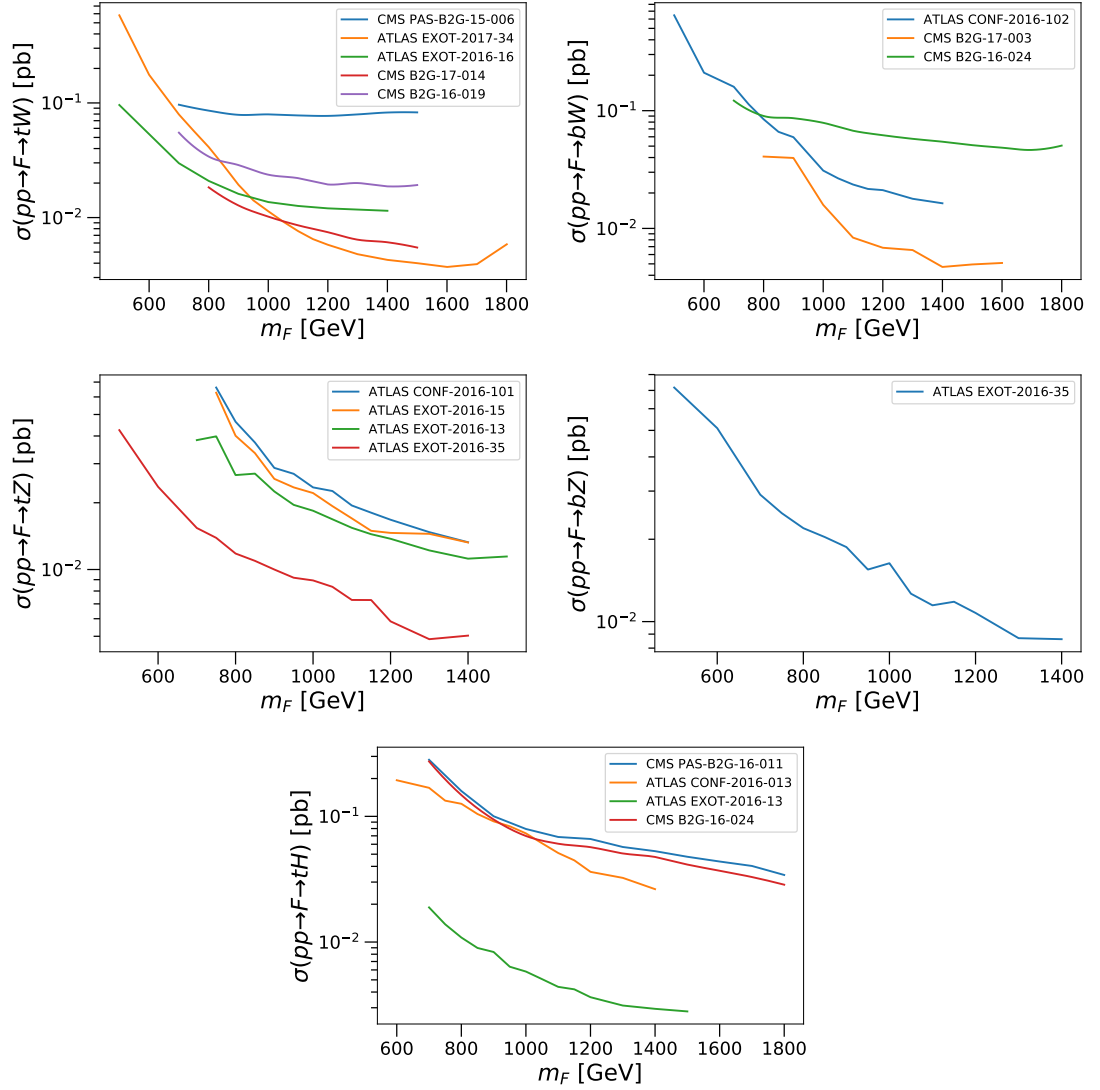


Figure 6.1: Observed 95% CL upper bounds for cross sections of pair-produced fermions F decaying to SM final states in proton-proton collisions at $\sqrt{s} = 13$ TeV that we use as constraints.

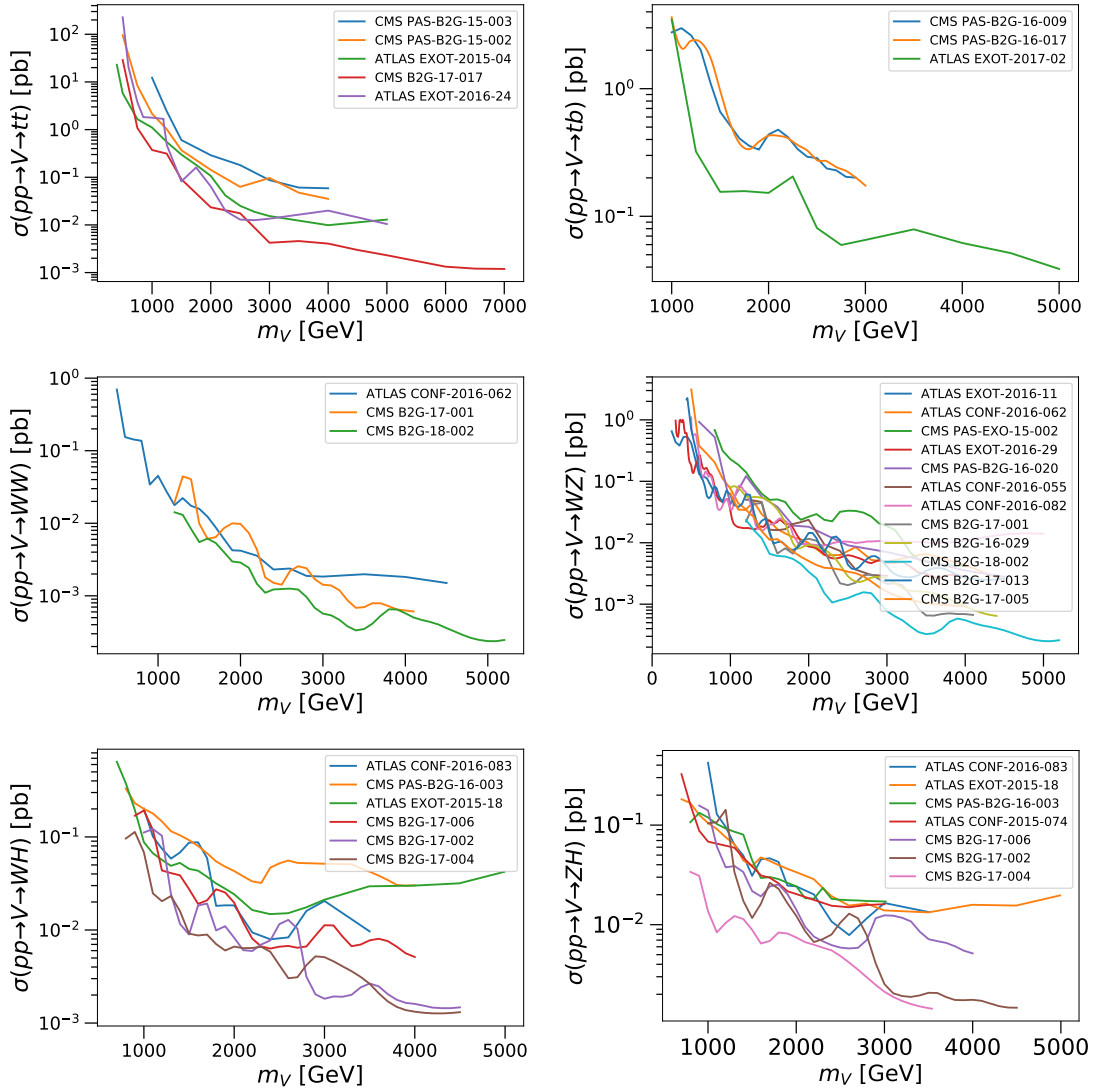


Figure 6.2: Observed 95% CL upper bounds for cross sections of a vector boson V decaying to SM final states in proton-proton collisions at $\sqrt{s} = 13$ TeV that we use as constraints.

In calculating the χ^2 for each decay, we only use the maximum χ^2 contribution from all the analyses restricting the decay. While a tighter bound would be given by combining the contributions from independent analyses, we find that the improvement in the excluded cross section at the 3σ level when doing this is often smaller than the error introduced by the NWA, and so implementing this would not be worthwhile, especially given the difficulty in combining analyses due to the lack of public information regarding the covariances between the analyses in question.


Indeed, from a naïve separation of the analyses into groups that use independent data sets based on the channel each looks at (which gives a charitable interpretation of “independent”), we find that the decays whose 3σ exclusion bounds would be most improved by combining the independent contributions are

- $V \rightarrow t\bar{t}$, by $\sim 10\text{--}20\%$ over the probed mass range,
- $V \rightarrow WH$, by $\sim 20\text{--}30\%$, and
- $V \rightarrow WZ$, by $\sim 30\text{--}50\%$, where V denotes a vector boson.


Now the relative error in the cross section introduced by the NWA is typically of order Γ_R/m_R , although it can be significantly larger depending on the kinematics of the decay [133]. Figure 6.3 displays this ratio for randomly distributed parameter points in the M4DCHM¹⁴⁻¹⁻¹⁰, showing that only a small proportion of points produce predictions precise enough to distinguish the improvements in the exclusion bounds. We are therefore justified in using only the maximum χ^2 contributions.

Mass Cuts

As a final measure, we impose harsh penalties for any new fermionic resonances below 500 GeV because such a resonance would likely have already been discovered. Ruling out points using hard mass cuts below 500 GeV makes it prohibitively time consuming to find an initial population of viable points, so instead we assign a steep one-sided Gaussian likelihood to the fermionic masses. This makes initial populations of points easier to find, and they will evolve towards points that give more reasonable masses as the scans proceed.



(a)



(b)

Figure 6.3: Plots of the estimated relative error Γ/m in the cross sections calculated with the NWA approximation for (a) the lightest charged boson W_2 and (b) the lightest uncharged boson Z_3 for randomly chosen points in the M4DCHM¹⁴⁻¹⁻¹⁰. The heavier bosons typically give larger relative errors.

Results

The results of our global fits of the M4DCHM⁵⁻⁵⁻⁵, the M4DCHM¹⁴⁻¹⁴⁻¹⁰, and the M4DCHM¹⁴⁻¹⁻¹⁰ are presented below in Sections 7.1 to 7.3, and are collectively discussed in Section 7.4. As a by-product of the fits, we analyse the experimental signatures of the three models in Section 7.5 in the hope of guiding the search for evidence of these models at the LHC.

But first, it should be stressed how difficult it was to obtain these final results with consistency between repeated scans. Dozens of test scans were performed with various tweaks to the models and scan settings in the process. Initially, we treated all of the quarks as partially composite in our models, but it soon became apparent that fitting these was infeasible. There was still difficulty after simplifying to only partially composite third generation quarks, and much experimenting was done regarding the population sizes, convergence criteria, model features, and constraints included in the scans. Removing the composite bottom sector led to high reproducibility in the results, which we took as a promising sign for obtaining consistent results with more complicated models. We then took steps increasing the model complexity by adding the partially composite bottom quark back in, but removing the composite $SU(3)$ sector. This again gave consistent fit results, after which we moved on to scanning the final versions of our models described in Chapter 4. In the interest of brevity, we forego presenting the results of these tests, although they were necessary to settle on the procedures for the final scans, and are broadly consistent with the final results. This was an incredibly time consuming process; many of these test scans would take a week or longer to converge running on a high-performance computational cluster, which severely hindered progress. The final `Diver` scans, using populations of 50,000 points and convergence thresholds of 10^{-5} , each had a runtime of at least one week on 256 cores, while the `PolyChord` scans each took several days to a week to converge.

As a final note, recall from Chapter 5 that the Bayesian fit results will be given in terms of marginalised posteriors, and the frequentist fit results in terms of profile likelihood ratios. We use `anesthetic` to analyse and plot the Bayesian results [134], and `pippi` to plot the profile likelihood ratios [135]. The Bayesian scan parameters are rescaled in the plots to match those used in the frequentist scans in order to facilitate easier comparisons between the fits.

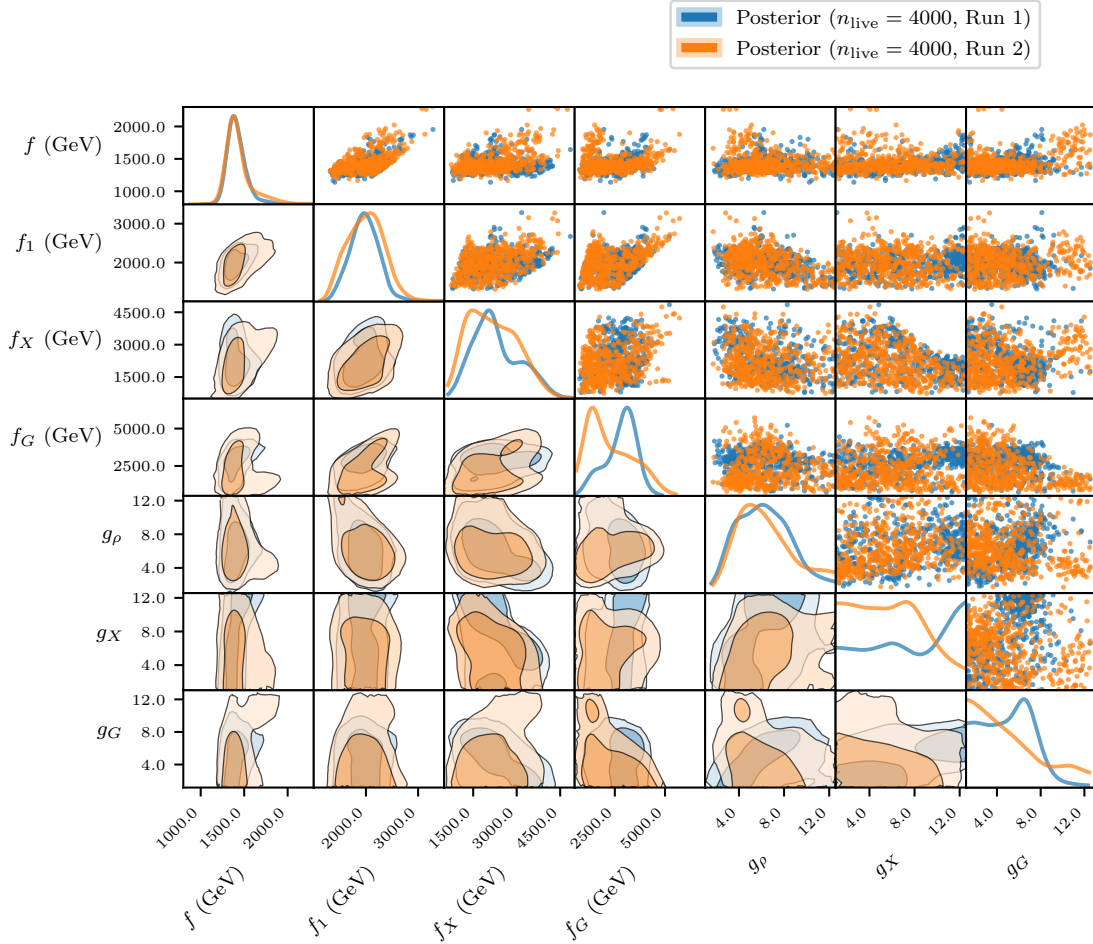


Figure 7.1: 1D and 2D marginalised posteriors for the gauge sector parameters in the M4DCHM⁵⁻⁵⁻⁵ found in two different runs with 4000 live points.

7.1 M4DCHM⁵⁻⁵⁻⁵

7.1.1 Bayesian Fit

Our results for this model come from four different PolyChord scans, performed to verify reproducibility. The first two scans, which each used 2000 live points, did not display very good agreement with each other, but more agreement was seen after increasing the sampling density to 4000 live points. The posteriors found by these 4000 point scans are shown in Figures 7.1 to 7.3. In the type of figure shown here, the diagonal plots show the 1D marginalised posteriors of the parameters, with the 2D marginalised posteriors of the various parameter pairings filling the off-diagonal plots. The upper-right plots show samples drawn directly from the posterior, while the lower-left plots are kernel density estimates of this distribution, with iso-likelihood contours containing 66% and 95% of the probability mass.

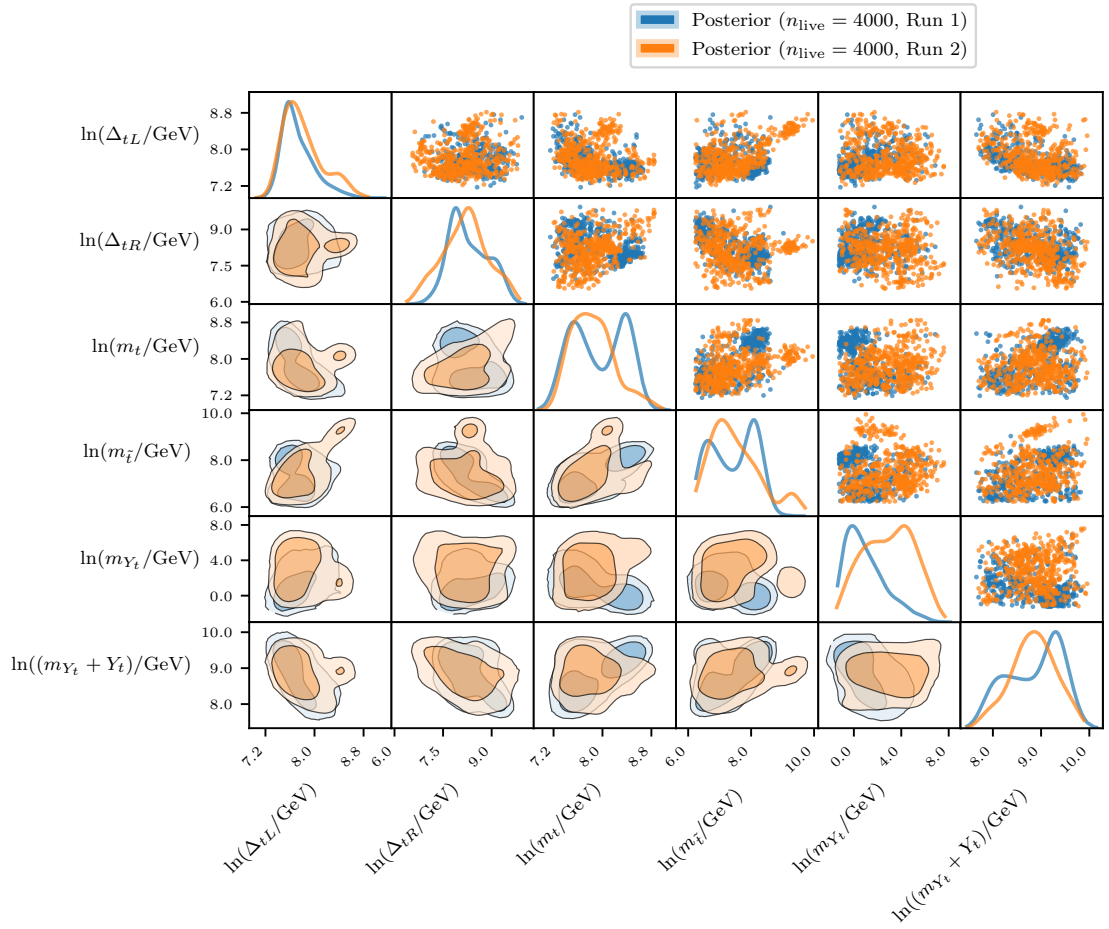


Figure 7.2: 1D and 2D marginalised posteriors for the top partner parameters in the $M4DCHM^{5-5-5}$ found in two different runs with 4000 live points.

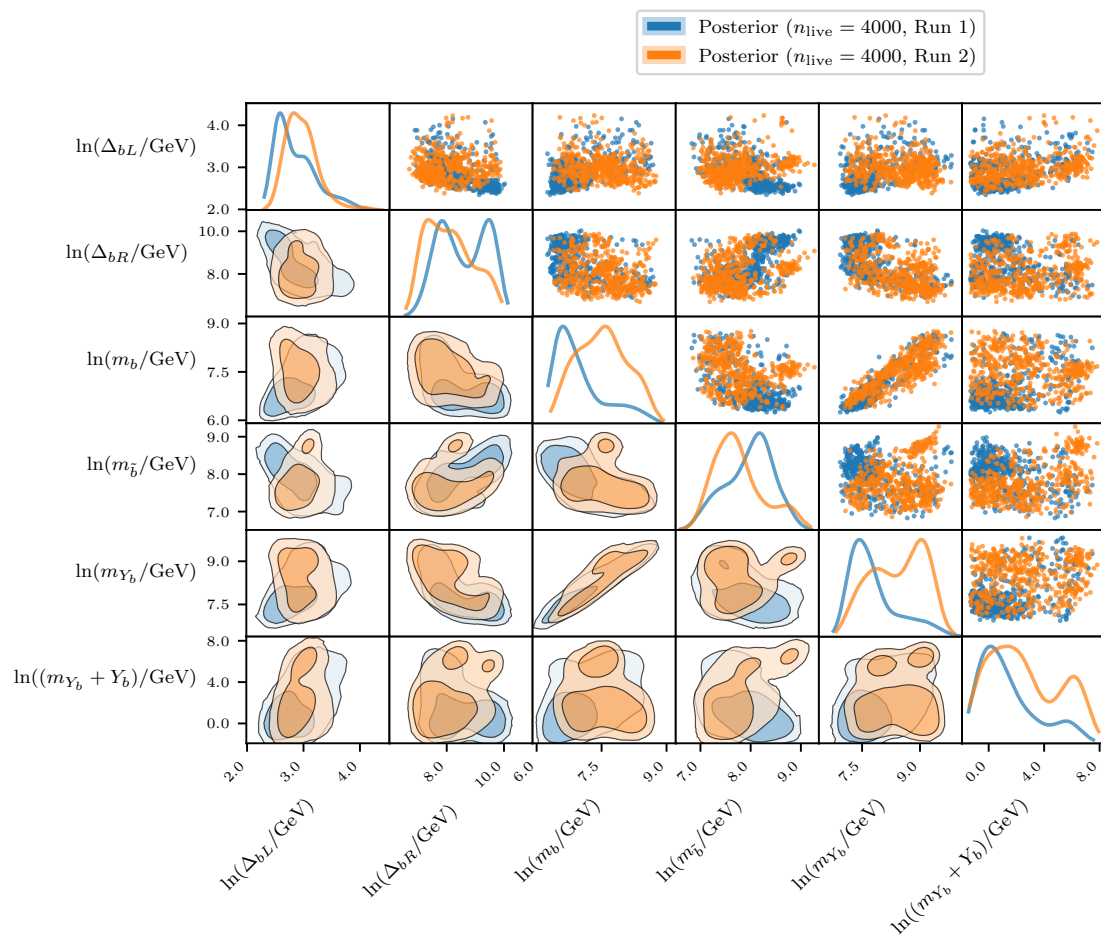


Figure 7.3: 1D and 2D marginalised posteriors for the bottom partner parameters in the M4DCHM⁵⁻⁵⁻⁵ found in two different runs with 4000 live points.

While there is not exact agreement in the posteriors in Figures 7.1 to 7.3, there are no major differences in the general regions that each scan found (apart from a minor mode found in the second run that is apparent in the plots of the fermion parameters), and specific features such as the correlation between m_b and m_{Y_b} are present in both scans. Given how difficult sampling the space has proven to be, we regard this as an acceptable level of agreement. It is also encouraging that the Bayesian evidences found in the 4000 point scans were in agreement:

$$\begin{aligned}\ln(\mathcal{Z})_{\text{Run 1}} &= -27.85 \pm 0.06, \\ \ln(\mathcal{Z})_{\text{Run 2}} &= -27.87 \pm 0.06,\end{aligned}\tag{7.1}$$

which we take as further indication that the parameter space has been reliably explored.

For our final analysis, we combine the samples from all four scans¹ in order to maximise the robustness of our conclusions. Our main results, the priors and posteriors of the parameters from the combined runs, are shown in Figures 7.4, 7.6 and 7.7. Recall that the prior distribution is that of the points drawn randomly within the imposed bounds subject only to consistency conditions such as the production of EWSB, while the posterior is a balance between the prior and the likelihood of each point. A posterior distribution similar to the prior distribution indicates the likelihood does not considerably constrain the space, while a posterior significantly different from the prior indicates the likelihood constrains the space to quite unnatural regions.

We look first at the results for the gauge parameters, in Figure 7.4. There is a clear pattern for the NGB decay constants: their priors all favour lower values of the constants and drop off rapidly for higher values, while their posteriors occupy the higher values. We remind the reader that in these Bayesian scans, these priors were not imposed directly, but are the result of fixing f to be whatever is needed to give the correct EWSB scale. Notice that these priors successfully encode the notion that smaller decay constants are more natural (see Equation (3.75)), and that larger decay constants require more fine-tuning. The fact that the posteriors differ significantly from these priors shows that the experimental constraints strongly disfavour smaller decay constants. We already know this to be the case from precision electroweak tests [44], but it will do well to examine other factors present. The ideal way to do this would be to run many scans, including only one experimental constraint at a time to see the effect each has on the posterior, but doing so would take more resources than are available to us. Instead, we simply examine the contribution to the χ^2 from each constraint for all of the dead points in the scans, and if any is small (large) over a specific parameter range, then we know that that constraint contributes to the enhancement (suppression) of the posterior in that range.

¹Nested sampling is trivially parallelisable, meaning that any collection of nested sampling runs can be combined by simply merging their samples and ordering them according to their likelihoods, and the result is equivalent to a single run that uses a number of live points equal to the sum of the live points from each individual run. This merging functionality is provided by `anesthetic`.

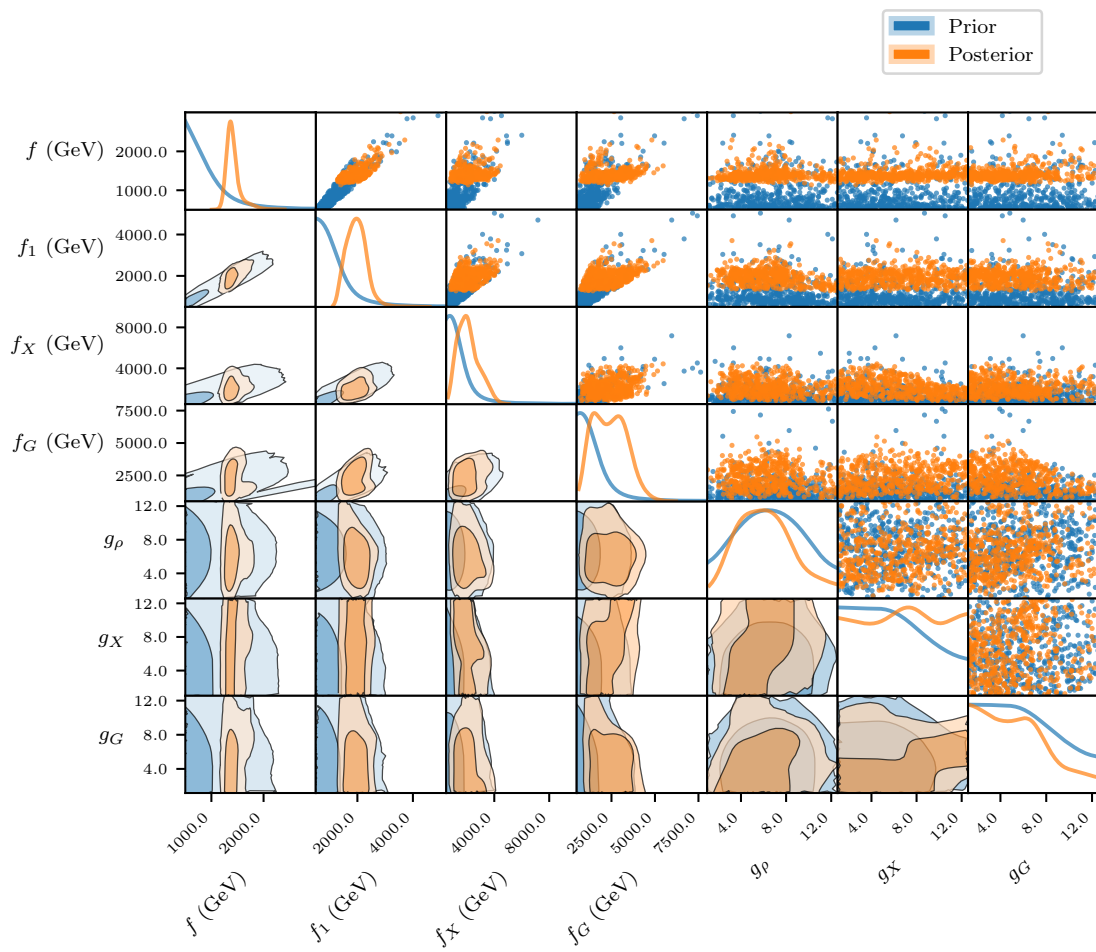


Figure 7.4: 1D and 2D marginalised priors and posteriors for the gauge sector parameters in the M4DCHM⁵⁻⁵⁻⁵.




Figure 7.5: Contributions to the total χ^2 from the oblique (S and T) and Higgs-to-diphoton signal strength constraints plotted against the Higgs decay constant f . Dead points from all Bayesian scans of the $M4DCHM^{5-5-5}$ are featured.

Let us go through the analysis process in explicit detail for the Higgs decay constant f , to illustrate. From Figure 7.4, the posterior for f lies mostly between 1.2 TeV and 1.9 TeV, with a maximum at ~ 1.4 TeV. The upper bound here can be explained easily, as the prior greatly suppresses the posterior beyond this region, but the same is not true for the lower bound. It turns out that the lower bound is due to a variety of constraints; namely, the oblique constraints, the Higgs signal strength constraints, and the Z decay constraints. All of these are understandable: the oblique constraints, which parameterise the electroweak precision observables, are sensitive to low-mass vector resonances and thus low f ; Higgs signal strengths are highly dependent on the non-linearities of the pNGB Higgs and hence on f , since the Higgs field enters the Lagrangian as h/f ; and the mixings of the vector bosons, dependent on f , would affect the Z decays. The dependence of the χ^2 contributions from some of these constraints on f is shown in Figure 7.5. Notice from Figure 7.5 that it is not *impossible* for such constraints to be individually well-satisfied at low f ; it is simply that many points in this range do not satisfy the constraints well, and this large low-likelihood region would suppress the posterior over this range because of the marginalisation process². This same sort of analysis will be done for the other parameters also, though we will not provide all of the χ^2 plots here.

Moving on to the other decay constants in Figure 7.4, there does not seem to be any structure in their posteriors beyond the consistency conditions outlined in Section 6.1. That is, the $SO(5)$ decay constant f_1 has no particular preference for any values between its theoretical bounds of f and $\sqrt{3}f$, having a posterior peaking around 2 TeV, and likewise for the $U(1)_X$ and $SU(3)_c$ decay constants f_X and f_G , which have similar posterior distributions between ~ 1 TeV and ~ 4 TeV. These parameters are constrained by the same factors as f mentioned above.

²It is also possible that satisfying one constraint well necessitates a violation in another constraint, which would lead to a large *total* χ^2 and a direct suppression of the posterior from the likelihood, rather than from a volume effect.

As for the gauge couplings, the posterior for the $SO(5)^1$ coupling g_ρ extends across the entire range $[1, 4\pi]$ with a preference for values between ~ 2 and ~ 8 , while that for the $U(1)_X^1$ coupling g_X is almost uniform, and that for the $SU(3)_3^1$ coupling g_G favours the lower end of the range. These posteriors are largely driven by the prior, though there are some additional experimental factors at play. All constraints can be satisfied well for all values of the coupling strengths, so it is difficult to discern exactly which experiments are behind this behaviour. The Higgs signal strength, SM mass, and oblique constraints all seem to slightly disfavour values for g_ρ at the extreme ends of the range, while the Z decay constraint R_b disfavour values below ~ 2 , in opposition to R_e and R_μ , which favour those values. The Z decays act similarly for g_X , and here the Higgs signal strength and oblique constraints promote stronger couplings. It is unsurprising that there is little structure in the posteriors of f_G and g_G since we do not include any constraints on heavy gluon decays or flavour physics in these Bayesian scans.

Moving on to the top sector parameters in Figure 7.6, we see the posteriors are significantly localised in comparison to the priors, demonstrating the effects of the constraints on the space. However, the posteriors all peak around the same regions as the priors, so it seems that satisfying the constraints does not require excessive fine-tuning. Indeed, the elementary-composite mixing and composite mass parameters all tend to be constrained around $\Delta_{tL} \sim \Delta_{tR} \sim m_t \sim m_{\bar{t}} \sim 1.8$ TeV, which, from the findings of Ref. [87] mentioned in Section 4.7, indicates double tuning is not so much of an issue. All of these parameters are primarily constrained by the SM masses, except for the lower bound of 0.5 TeV on $m_{\bar{t}}$ that comes from the resonance mass cutoff³ mentioned in Section 6.3. Not particularly well constrained is the off-diagonal mass $m_{Y_{\bar{t}}}$, which may take any value below ~ 2 TeV. The last top sector parameter, $m_{Y_t} + Y_t$, clearly prefers larger values from around 2 TeV to 20 TeV, and this is due largely to the SM mass and oblique constraints.

Finally, we discuss the bottom sector parameters in Figure 7.7. There is not as much uniformity among these as for the top parameters, presumably because these play a subdominant role in the Higgs potential. Interestingly, Δ_{bL} is very narrowly constrained, between around 10 GeV and 30 GeV. Much of this is due to the prior, which suppresses values above ~ 150 GeV, but again there are other factors present. The lower bound on its posterior is easily identified as coming from the bottom quark mass constraint, but the upper bound is more subtle, and there are two contributing factors at play:

- Similarly to the case with f , there is a large region of points with $\Delta_{bL} \gtrsim 50$ GeV that significantly violate the Z decay constraints, suppressing the marginalised posterior over this range. This is on display in Figure 7.8. It agrees with intuition that the mixing parameter Δ_{bL} of the left-handed quark doublet is sensitive to the tight constraints on the $Z\bar{b}_L b_L$ coupling.

³We could not impose this bound at the prior level because we were imposing the prior on the dimensionless variable $m_{\bar{t}}/f$ for an unspecified scale f .

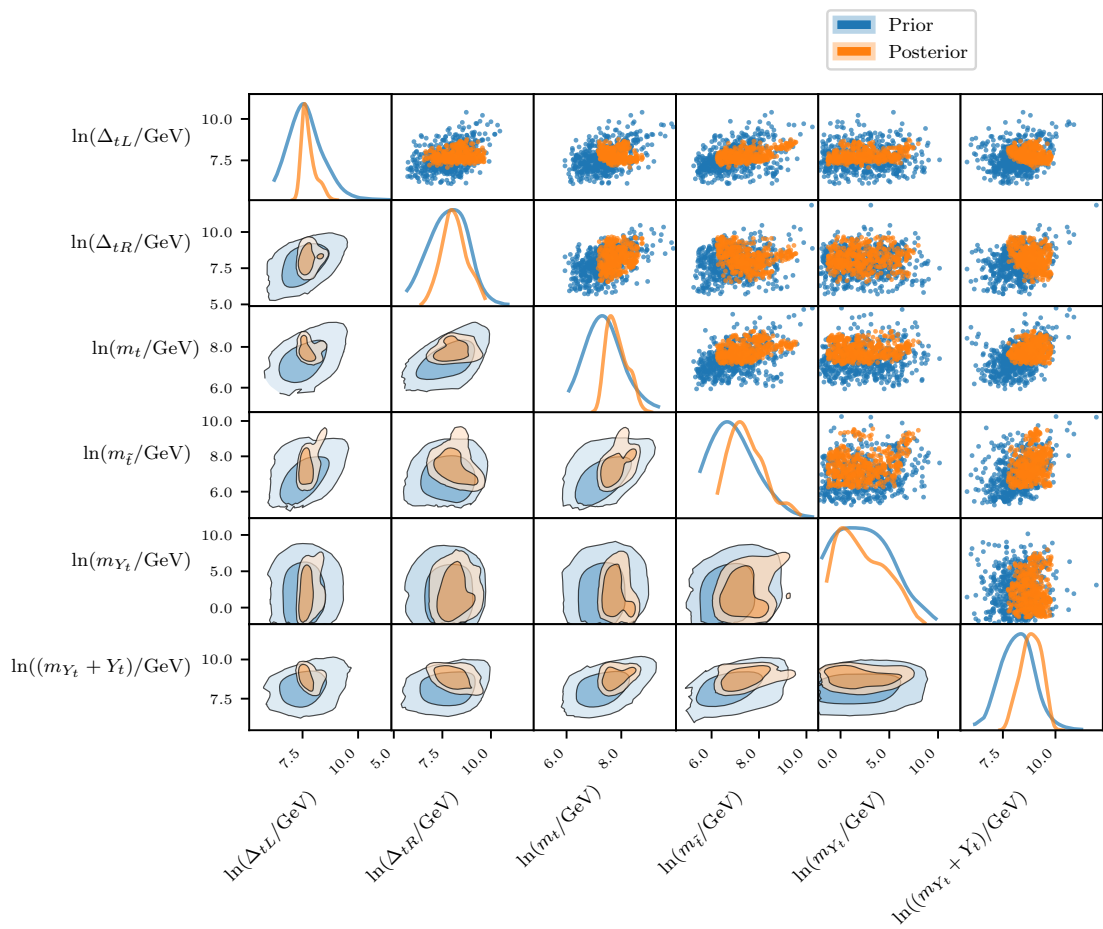


Figure 7.6: 1D and 2D marginalised priors and posteriors for the top partner parameters in the $M4DCHM^{5-5-5}$.

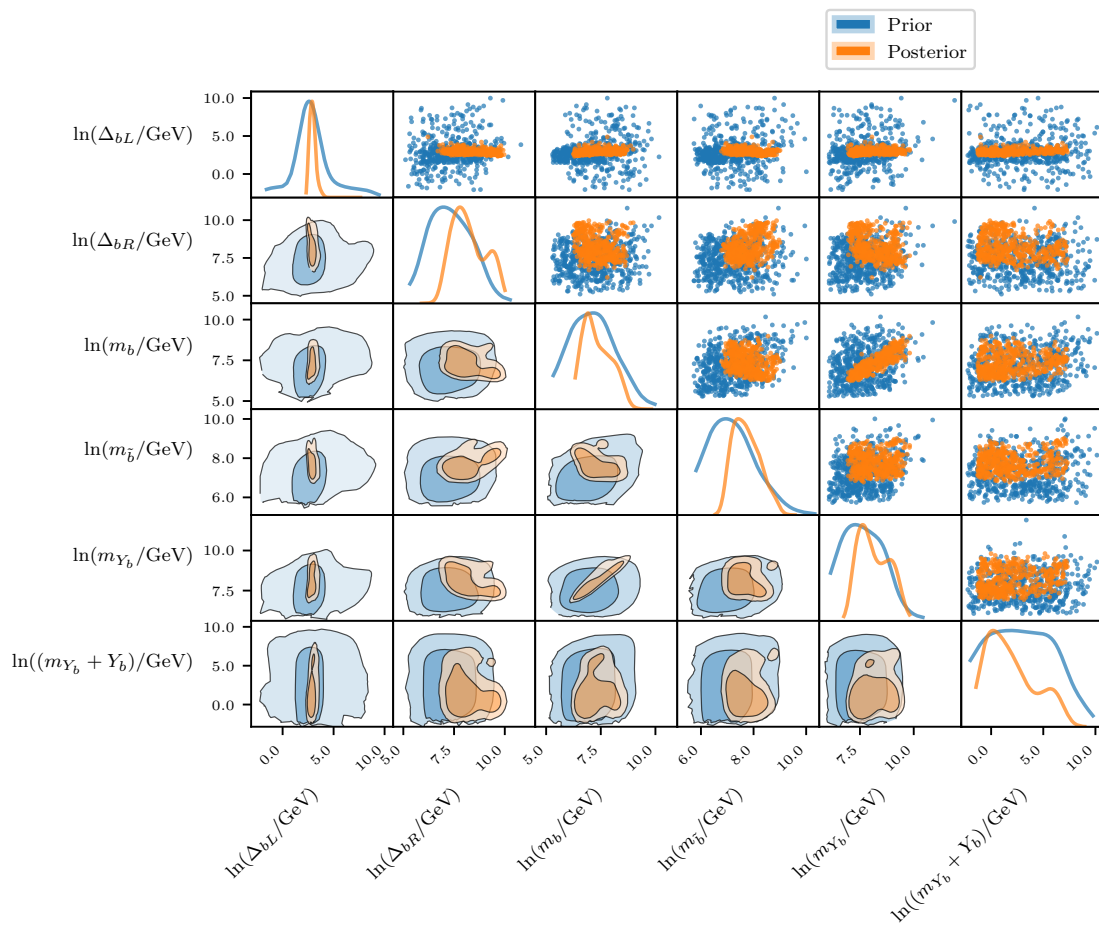


Figure 7.7: 1D and 2D marginalised priors and posteriors for the bottom partner parameters in the M4DCHM⁵⁻⁵⁻⁵.




Figure 7.8: Plots helping to explain the localisation of the posterior of Δ_{bL} . *Left*: contributions to the χ^2 from the Z decay constraint R_b . Notice the large region of points that significantly violate this constraint for $\Delta_{bL} \gtrsim 50$ GeV. *Right*: correlations displayed by the posterior of Δ_{bL} with those of other parameters.

- Correlations exist among the posteriors of Δ_{bL} and Δ_{tL} , and of Δ_{bL} and m_t , also shown in Figure 7.8. That is, one or more experiments must constrain some non-trivial combination of these parameters. Other constraints on Δ_{tL} and m_t thereby inadvertently act to constrain Δ_{bL} also. In this case, the upper bound on Δ_{tL} and lower bound on m_t both contribute, to some extent, to the upper bound on Δ_{bL} .

Easier to analyse are the parameters Δ_{bR} and m_{Y_b} , which share similar posterior distributions ranging between roughly 1 TeV and 10 TeV, with lower bounds coming from the bottom quark and Higgs mass constraints. The resonance mass cutoff sets a lower bound of 0.5 TeV on the posterior of m_b , which ranges up to ~ 4.5 TeV. The bottom mass constrains $m_{\bar{b}}$ to lie above ~ 1 TeV, and the prior below ~ 5 TeV. And finally, $m_{Y_b} + Y_b$ lies mostly below 2 TeV, with a preference towards the lower values stemming from the SM mass and Z decay constraints.

One interesting feature to note here is that the correlation between the posteriors of m_b and m_{Y_b} mentioned previously is *not* present in the prior and so must be enforced by one or more of the experimental constraints, though it is difficult to pinpoint the cause of this correlation. Other structures to note in Figures 7.6 and 7.7 are slight correlations in the priors and posteriors of the pairs $(m_t, m_{Y_t} + Y_t)$ and $(m_t, m_{\bar{t}})$, a mild correlation in the posteriors of $(\Delta_{bR}, m_{\bar{b}})$, and posteriors that seem to avoid small values for both parameters in the pairs (Δ_{tL}, m_t) , $(\Delta_{tL}, m_{Y_t} + Y_t)$, $(m_b, m_{\bar{b}})$, and (Δ_{bR}, m_{Y_b}) . More will be said about these results when comparing models in Section 7.4.

7.1.2 Frequentist Fit

In contrast to our Bayesian fit of this model, and despite our best efforts, we have unfortunately not been able to arrive at reproducible results for our frequentist fit of the M4DCHM⁵⁻⁵⁻⁵. The best-fit points found in several `Diver` scans with pop-

ulation sizes of 20,000+ points were always sparsely scattered and never in similar regions, neither of which are to be expected if the space has been well-sampled. This may be due in part to our scanning method for these fits, which recall makes it difficult to explore along the f dimension, but we cannot repeat the scans using the more sophisticated parameterisation on account of time constraints. Nevertheless, out of interest, we include the combined results in Figures 7.9 to 7.13. For reference, the best-fit point found in this model carries a likelihood of $\ln(\mathcal{L}) = -15.3$.




Figure 7.9: Profile likelihoods for the NGB decay constants in the M4DCHM⁵⁻⁵⁻⁵. The 1σ and 2σ confidence level bounds are overlaid. Stars mark the best-fit point.




Figure 7.10: Profile likelihoods for the gauge couplings in the M4DCHM⁵⁻⁵⁻⁵.




Figure 7.11: Profile likelihoods for the couplings linking the elementary and composite sectors in the $M4DCHM^{5-5-5}$.




Figure 7.12: Profile likelihoods for the on-diagonal mass parameters in the $M4DCHM^{5-5-5}$.




Figure 7.13: Profile likelihoods for the off-diagonal mass and proto-Yukawa parameters in the $M4DCHM^{5-5-5}$.

It is immediately apparent in these plots that the profile likelihood ratio is not smoothly varying across many of the parameters. This is not to be expected (after all, the likelihood function is continuous!) and is an indication of the poor sampling of the space. As such, these results are not trustworthy, and we shall not discuss them in detail or in particularly definite terms, leaving the majority of the frequentist analysis to the other models.

We first note that the likelihood distributions for the NGB decay constants in Figure 7.9 differ significantly from the posteriors from the Bayesian fit. Here, larger values are preferred, especially for f . If the confidence interval bounds in Figure 7.9 are to be believed, f lies above ~ 2.1 TeV at the 2σ confidence level, and maximises the likelihood at ~ 4.5 TeV. The discrepancy between this interval and the posterior, which lies below 2 TeV, highlights the distinction between the fitness of a region and its naturalness. Evidently, this model must be quite finely-tuned to best fit the data. This range on f corresponds to a vacuum misalignment $s_{\langle h \rangle} \lesssim 0.11$, or a fine-tuning of $\gtrsim 70$ (see Equation (3.75)). As for the gauge couplings, the likelihood for g_ρ in Figure 7.10 roughly matches its posterior, preferring values between 2 and 8, as is to be expected by the similarity of its prior and posterior. Contrastingly, the g_X and g_G likelihood distributions do not match with their posteriors, tending towards smaller and larger values respectively.

The likelihood distributions for the fermion sector parameters are more in line with the Bayesian posterior results. A notable exception is Δ_{bL} , whose posterior was narrowly constrained around $\mathcal{O}(10)$ GeV, here having a 2σ confidence interval of roughly 0.1 TeV to 3 TeV. Recall that the posterior for Δ_{bL} was heavily suppressed

both by the prior and by a large low-likelihood region over this range, so this is at least not necessarily inconsistent. It is also important to keep in mind that these scans included the LHC resonance bounds as constraints while the Bayesian scans did not, which may contribute to some of the differences between the results. This matter will be discussed further for the other models, where more satisfactory results were obtained.

7.2 $M4DCHM^{14-14-10}$

7.2.1 Bayesian Fit

Unfortunately, in this model we only had the resources to perform two Bayesian scans, each with 2000 live points, and their results are not entirely in agreement, as can be seen in Figures 7.14 and 7.15. The most obvious difference between these scan results is the many minor modes of the posterior among f and the fermion parameters that were found in the second scan, but not the first. Their evidences,

$$\begin{aligned}\ln(\mathcal{Z})_{\text{Run 1}} &= -40.1 \pm 0.2, \\ \ln(\mathcal{Z})_{\text{Run 2}} &= -42.0 \pm 0.3,\end{aligned}\tag{7.2}$$

are also not in good agreement. We aim to improve these results in later work by scanning with higher sampling densities, but presently we are resigned to giving only the results of these scans for this model. Still, we see there is loose agreement in the dominant modes of the posteriors found in each scan, so we can draw some preliminary conclusions about the space here.

The combined results of our two scans are given in Figures 7.16 and 7.17. The priors and posteriors for the gauge sector parameters, in Figure 7.16, look much the same as in the $M4DCHM^{5-5-5}$, with the NGB decay constant posteriors concentrated well above the small values preferred by the prior. The posterior for f enjoys a slightly higher range than before, from around 1.5 TeV to 2.2 TeV, again due largely to the Higgs signal strength, Z decay, and oblique constraints. And in contrast to the previous model, here we *do* see some interesting structure between the remaining decay constants on top of that imposed by consistency conditions. Specifically, f_1 and f_G tend towards the higher ends of their possible ranges, respectively spanning from ~ 1.5 TeV to ~ 3.8 TeV and from ~ 0.8 TeV to ~ 6.5 TeV, while f_X has a modest preference for the lower end of its range, spanning between ~ 0.5 TeV and ~ 5.7 TeV. Unfortunately, these preferences are too slight for their causes to be discerned. Moving on, the gauge coupling g_ρ has a similar distribution to before, preferring mid-range values but perhaps slightly skewed in favour of weaker couplings, while g_X is more constrained in this model, having a definite tendency towards lower values. Both seem to be driven once again by the priors in combination with the Higgs signal strength and Z decay constraints. Stronger g_G couplings are disfavoured, as they appear to entail undesirably small values of f_G .

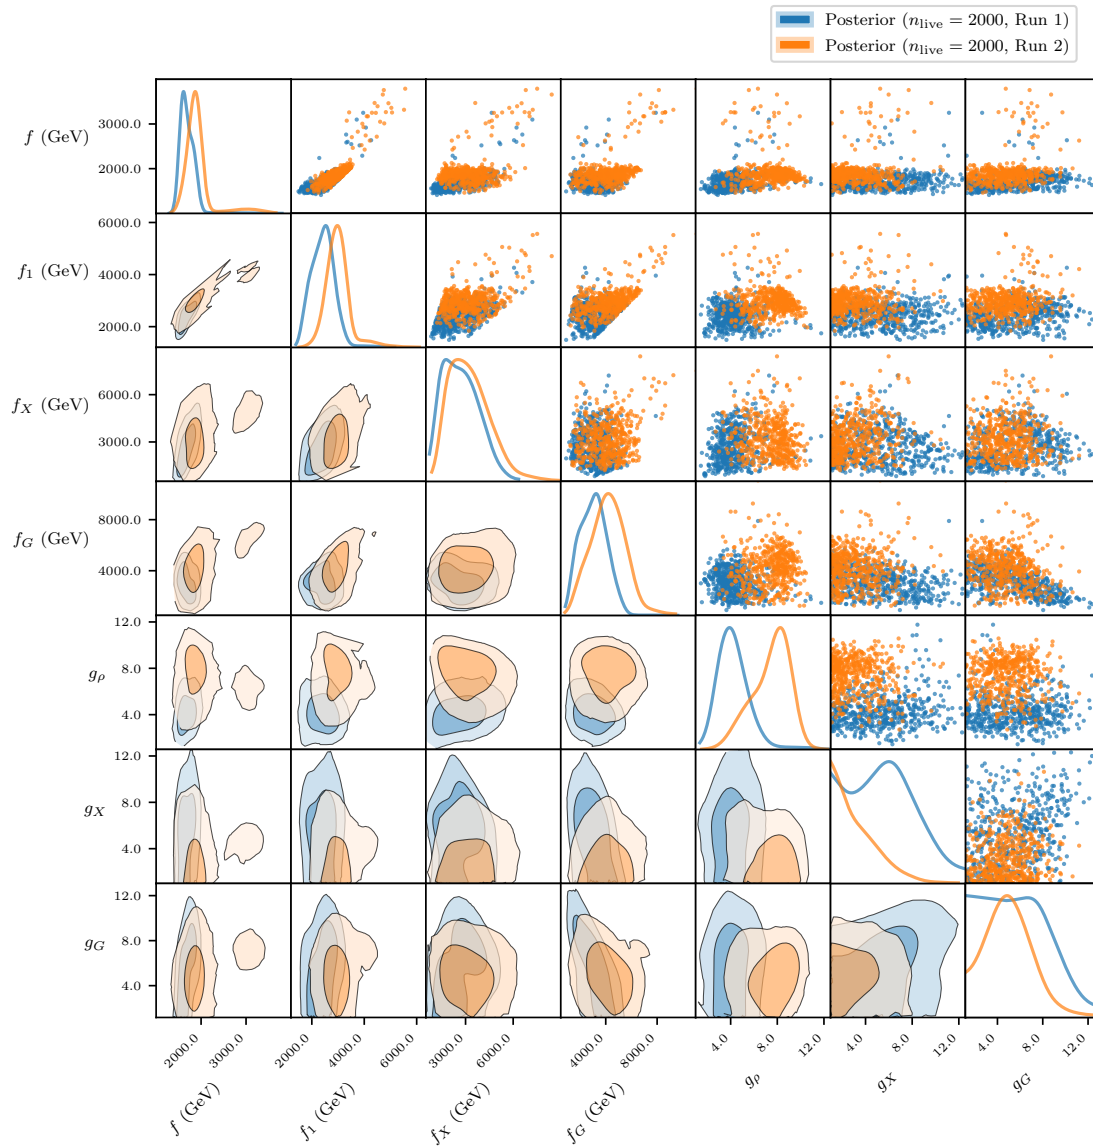


Figure 7.14: 1D and 2D marginalised posteriors for the gauge sector parameters in the M4DCHM¹⁴⁻¹⁴⁻¹⁰ found in two different runs with 2000 live points.

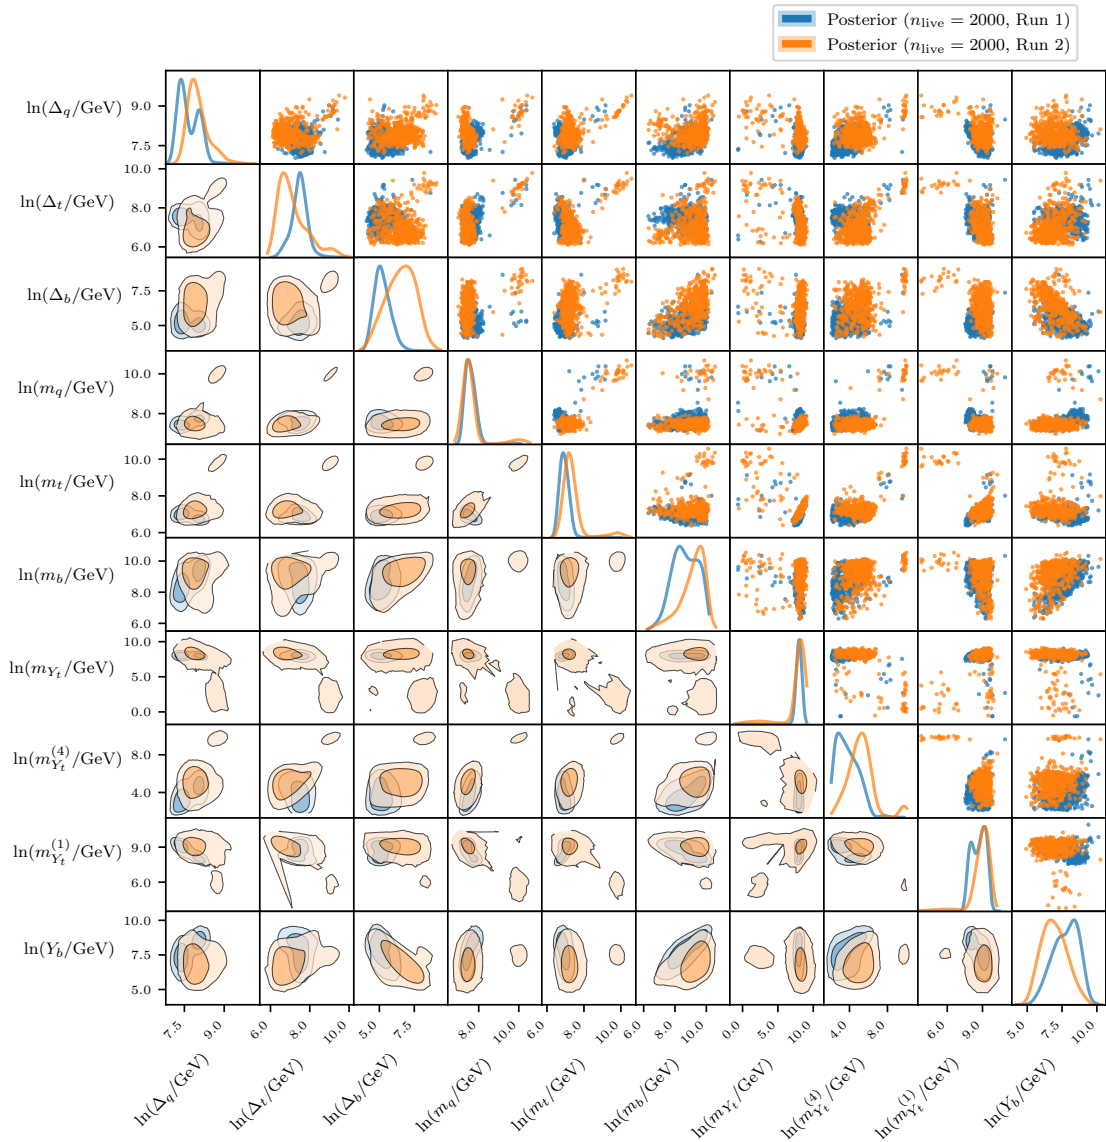


Figure 7.15: 1D and 2D marginalised posteriors for the fermion sector parameters in the $M4DCHM^{14-14-10}$ found in two different runs with 2000 live points.

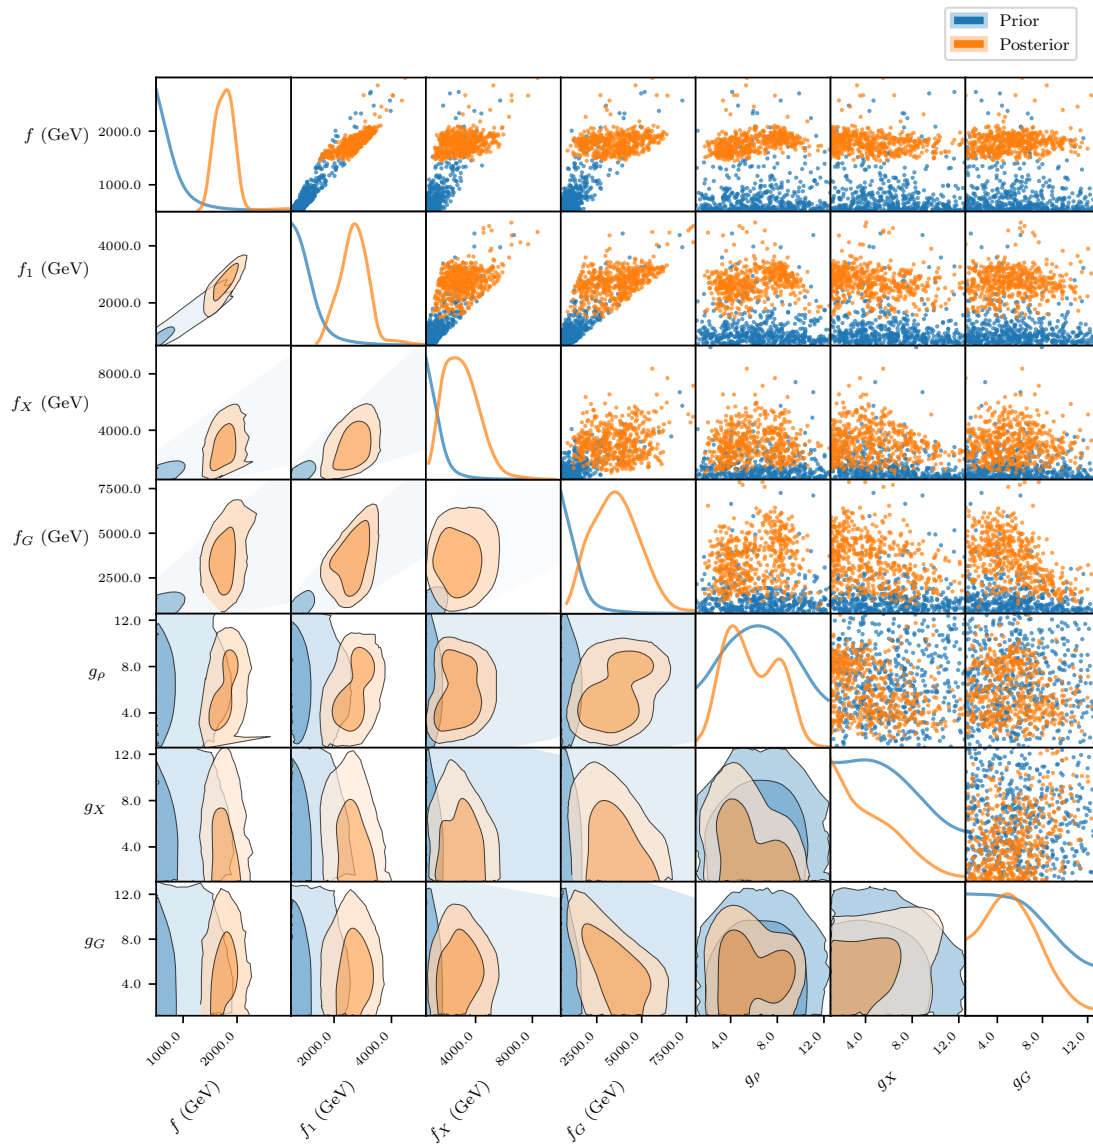


Figure 7.16: 1D and 2D marginalised priors and posteriors for the gauge sector parameters in the M4DCHM^{14–14–10}.

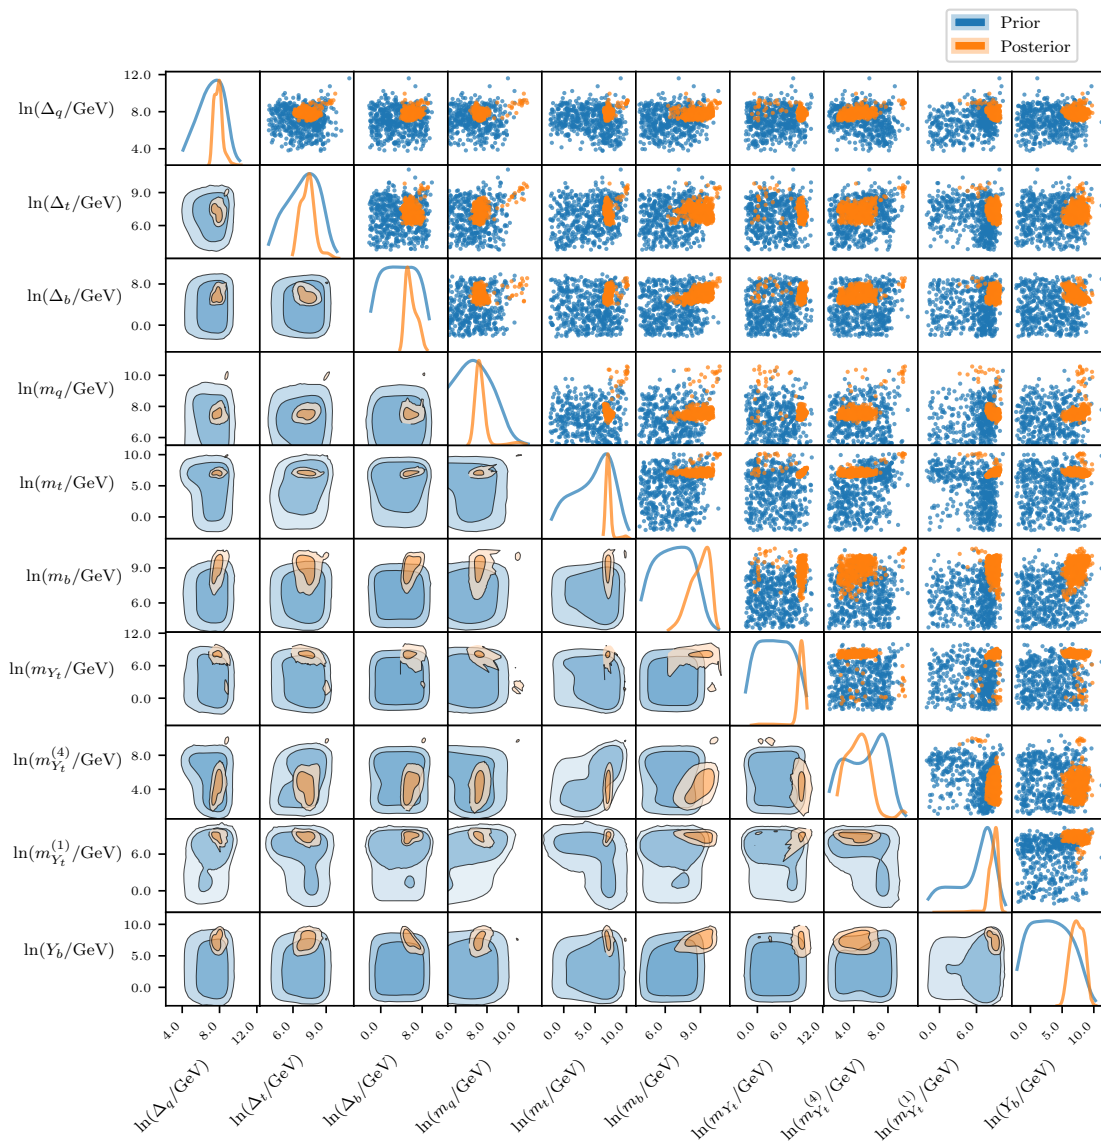


Figure 7.17: 1D and 2D marginalised priors and posteriors for the fermion sector parameters in the $M4DCHM^{14-14-10}$. The parameters $m_{Y_t}^{(4)}$ and $m_{Y_t}^{(1)}$ are defined by Equations (7.3) and (7.4).

The fermion parameters in the M4DCHM¹⁴⁻¹⁴⁻¹⁰ are seen from Figure 7.17 to be significantly more constrained than those in the M4DCHM⁵⁻⁵⁻⁵, with posteriors concentrated on considerably smaller regions in comparison to the priors. Being emblematic of higher fine-tuning, this might explain some of the difficulty in fitting this model. This is particularly evident for the on-diagonal mass parameter m_t , whose prior allows it to take effectively any value up to 20 TeV but whose posterior is localised between ~ 0.7 TeV and ~ 4 TeV. Mass constraints are the main factor here, with the Z decay and Higgs signal strength constraints also contributing to the upper bound somewhat, and the same is true for all of the other on-diagonal mass and mixing parameters as well. Those couplings for the left-handed quark partners, Δ_q and m_q , share similar posterior distributions, respectively ranging from ~ 1.2 TeV to ~ 3.3 TeV and from ~ 1.3 TeV to ~ 3.0 TeV. Having a slightly wider range is Δ_t , being mostly constrained between ~ 0.5 TeV and ~ 3.5 TeV, while Δ_b is able to take yet lower values from ~ 0.1 TeV to ~ 3.0 TeV. Despite its prior, m_b tends towards very large values with a posterior between ~ 4 TeV and ~ 25 TeV. Evidently, the right-handed bottom quark must possess a small compositeness Δ_b/m_b in this model, as is to be reasonably expected from its relatively small mass.

The remaining proto-Yukawa parameters display some interesting features. Save for the minor modes found in the second scan, the posterior for m_{Y_t} is concentrated between ~ 1.8 TeV and ~ 6.7 TeV, while

$$m_{Y_t}^{(4)} := m_{Y_t} + \frac{1}{2}Y_t \quad (7.3)$$

has a posterior constrained to lower values $\lesssim 0.4$ TeV. In other words, it must be that $Y_t \approx -2m_{Y_t}$. It is not clear exactly which experiments are to blame, but it seems that the posterior is suppressed at smaller values of m_{Y_t} and at greater values of $m_{Y_t}^{(4)}$ by large low-likelihood regions stemming from the Higgs signal strength constraints. Z decays also slightly favour the smaller values of $m_{Y_t}^{(4)}$. Meanwhile, both the prior and posterior for

$$m_{Y_t}^{(1)} := m_{Y_t} + 4Y_t/5 + 4\tilde{Y}_t/5 \quad (7.4)$$

favour large values, from around 2.5 TeV to 17 TeV, with smaller values again perhaps being suppressed in the posterior due to a volume effect from the Higgs signal strengths. From the previous observations, this implies \tilde{Y}_t must also be quite large, greater than around 4.5 TeV. Finally, SM mass constraints force Y_b to be larger than 150 GeV, even though smaller values are favoured by the prior.

None of these fermion parameters display any particularly strong correlations in their posteriors, although there does seem to be a mild negative correlation between Δ_b and Y_b , and a mild positive correlation between m_b and $m_{Y_t}^{(4)}$. Interestingly, the prior disfavours regions where both m_t and $m_{Y_t}^{(1)}$, and to a lesser extent both $m_{Y_t}^{(4)}$ and $m_{Y_t}^{(1)}$, are small.

7.2.2 Frequentist Fit

While our Bayesian fit of the M4DCHM¹⁴⁻¹⁴⁻¹⁰ is lacking in reproducibility somewhat, our frequentist fit has found more success. Having seen general consistency between separate scans with 20,000 and 50,000 points, we give the profile likelihood ratios for all of the parameters from the combined runs in Figures 7.18 to 7.22. The best-fit point found here has a likelihood of $\ln(\mathcal{L}) = -15.1$.




Figure 7.18: Profile likelihoods for the decay constants in the M4DCHM¹⁴⁻¹⁴⁻¹⁰.




Figure 7.19: Profile likelihoods for the gauge couplings in the M4DCHM¹⁴⁻¹⁴⁻¹⁰.




Figure 7.20: Profile likelihoods for the couplings linking the elementary and composite sectors in the M4DCHM^{14–14–10}.




Figure 7.21: Profile likelihoods for the on-diagonal mass parameters in the M4DCHM^{14–14–10}.




Figure 7.22: Profile likelihoods for the off-diagonal mass and proto-Yukawa parameters in the M4DCHM¹⁴⁻¹⁴⁻¹⁰.

Notice that the distributions in these plots are significantly more smooth than those for the M4DCHM⁵⁻⁵⁻⁵, lending more credibility to these results. We first observe here that the profile likelihood grows as f increases up to our cutoff of 5 TeV. From the confidence interval bounds shown on the plot in Figure 7.18, it is seen that $f \gtrsim 4.2$ TeV at the 1σ confidence level⁴, and $f \gtrsim 3.7$ TeV at the 2σ level. This range on f translates to a misalignment $s_{(h)} \lesssim 0.065$ at the 2σ level, giving a rather significant fine-tuning estimate of $\gtrsim 230$. These frequentist scans were able to produce a decent fit despite this tuning, and in contrast to the Bayesian scans, because here the decay constants were sampled uniformly over their ranges, with no disinclination for the larger values that optimise the likelihood. As for the other decay constants, f_1 tends to lie in the middle of its theoretical range, having a best fit value of ~ 6 TeV, while f_X tends towards larger values. Again, the heavy gluon decay constant f_G is not particularly constrained since we do not include flavour constraints in our scans.

Regarding the composite gauge couplings, g_ρ has a clear preference for smaller values. Our result of $g_\rho \lesssim 2.5$ fits quantitatively with Ref. [87], which found that under some mild assumptions the Higgs mass is expected to scale with g_ρ in this model, and a coupling of such a size is to be expected to reproduce the correct Higgs mass. However, it should be noted that we found many points across *all* values of g_ρ

⁴These bounds assume that our scans have found the true maximum likelihood point in the M4DCHM¹⁴⁻¹⁴⁻¹⁰. It seems likely that had we extended our bound on f to greater than 5 TeV, we would have found points that even better fit the data. However, as long as the best-fit point in our scan has a comparable likelihood to the true global optimum, these confidence intervals will be approximately accurate.

that give a reasonable Higgs mass, and it is strictly a combination of the constraints that results in our likelihood distribution for g_ρ . The other gauge couplings are under no such constraint, and are seen to heavily prefer larger values. Note that none of these distributions match with the posteriors given previously. Part of this difference can be attributed to the priors for the couplings, but it seems the more pressing factor is the LHC search constraints used in these scans that disfavour light gauge bosons and therefore small gauge couplings.

The remaining fermion parameters mostly follow the patterns seen in the Bayesian results, but typically at greater values than suggested by the posteriors. The marginalised posteriors and profile likelihoods do not need to be similar, of course, because of the prior-dependence of the posteriors and the difference between the marginalisation and profiling procedures, but other factors that could contribute to the differences here are the LHC constraints that encourage heavier fermions and therefore larger values of certain couplings, and also the poor quality of the Bayesian fit of this model. That the priors so disfavour the ranges highlighted here implies a considerable amount of fine-tuning for this best-fit region.

Figure 7.20 shows that the mixing parameters Δ_q and Δ_t lie above ~ 4.4 TeV at the 2σ level, with Δ_q tending to be the smaller of the two, being constrained below ~ 9.0 TeV, while Δ_t can reach upwards of 17 TeV and skews towards higher values. The right-handed bottom mixing Δ_b has a profile likelihood distribution that more closely matches the posterior found earlier, with a 2σ confidence interval from ~ 0.5 TeV to ~ 4.0 TeV. The mass parameter m_t is once more narrowly constrained, to within about 15% of 2 TeV. Larger values are taken by m_q , lying between ~ 2.8 TeV and ~ 5.0 TeV, and m_b , reaching up to ~ 8.1 TeV at the 2σ level. Again we see the top quark is considerably composite while the bottom quark is mostly elementary.

Where the frequentist and Bayesian results most differ is in the parameter m_{Y_t} , with the likelihood here favouring values less than ~ 0.7 TeV. Since the posterior was driven by a large low-likelihood region at small m_{Y_t} , these two sets of results are not contradictory. And note from Figure 4.3 that a large m_{Y_t} leads to a multitude of light fermions, while a small m_{Y_t} implies heavier fermions close to m_q and m_t in mass, so the LHC constraints in these scans are best satisfied by taking m_{Y_t} small. Better matching the Bayesian results are $m_{Y_t}^{(4)}$, being essentially negligible; $m_{Y_t}^{(1)}$, being constrained to values above ~ 4.0 TeV; and Y_b , again taking values greater than 150 GeV but now also constrained to less than around 1.3 TeV. See Equations (7.3) and (7.4) for the definitions of $m_{Y_t}^{(4)}$ and $m_{Y_t}^{(1)}$.

Interestingly, $m_{Y_t}^{(1)}$ is one of the few parameters whose profile likelihood is correlated with that of any other parameter in this model. All of the correlations present are shown in Figure 7.23. From this, we see that higher symmetry scales f can accommodate stronger Δ_t , m_q , and $m_{Y_t}^{(1)}$ couplings. Consequently, the allowed ranges of these parameters can be roughly estimated as the experimental lower bound on f further improves. In addition to these correlations, we also see from Figure 7.23 that g_X and f_X are negatively correlated, as are Y_b and Δ_b . The former correlation can




Figure 7.23: Profile likelihoods for parameters displaying slight correlations in the M4DCHM¹⁴⁻¹⁻¹⁰. Contours define the 1σ and 2σ confidence levels.

heuristically be explained by $g_X f_X$ being linked to the mass of a composite gauge boson, so large values are more desirable in light of the LHC search constraints. The latter correlation has already been seen in the posterior (see Figure 7.17), suggesting a constraint common to both fits, such as the bottom mass constraint, is the cause. Note that these correlations help to explain the differences in parameter scales between the Bayesian and frequentist results, since the Bayesian scans favour smaller f , and hence smaller Δ_t , m_q , and $m_{Y_t}^{(1)}$.

7.3 M4DCHM¹⁴⁻¹⁻¹⁰

7.3.1 Bayesian Fit

For this model we performed four Bayesian scans: the first three with 2000 live points, and the last with 4000 live points. Perhaps because of its smaller parameter space, this was by far the easiest of the three models to fit, with all four of these scans having good agreement in their results. Indeed, the posteriors of the 4000 point run and the first 2000 point run, for example, are displayed in Figures 7.24 and 7.25 and are seen to match each other closely. The evidences also match quite well across scans:

$$\begin{aligned}
 \ln(\mathcal{Z})_{\text{Run 1}} &= -37.54 \pm 0.12, \\
 \ln(\mathcal{Z})_{\text{Run 2}} &= -37.38 \pm 0.12, \\
 \ln(\mathcal{Z})_{\text{Run 3}} &= -37.51 \pm 0.09, \\
 \ln(\mathcal{Z})_{\text{Run 4}} &= -37.70 \pm 0.06.
 \end{aligned}
 \tag{7.5}$$

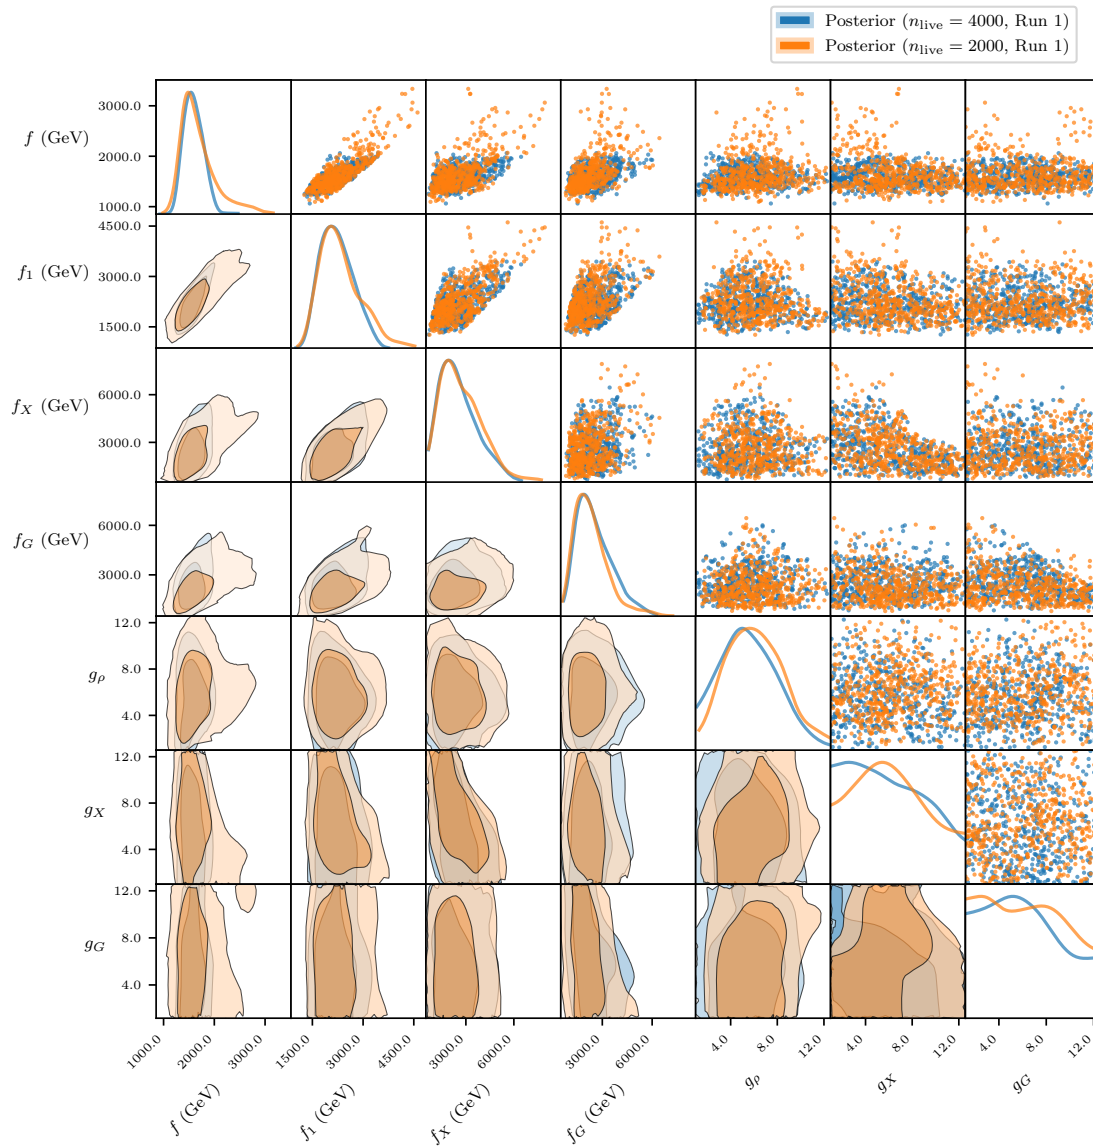


Figure 7.24: 1D and 2D marginalised posteriors for the gauge sector parameters in the M4DCHM¹⁴⁻¹⁻¹⁰ found in two different runs with 2000 and 4000 live points respectively.

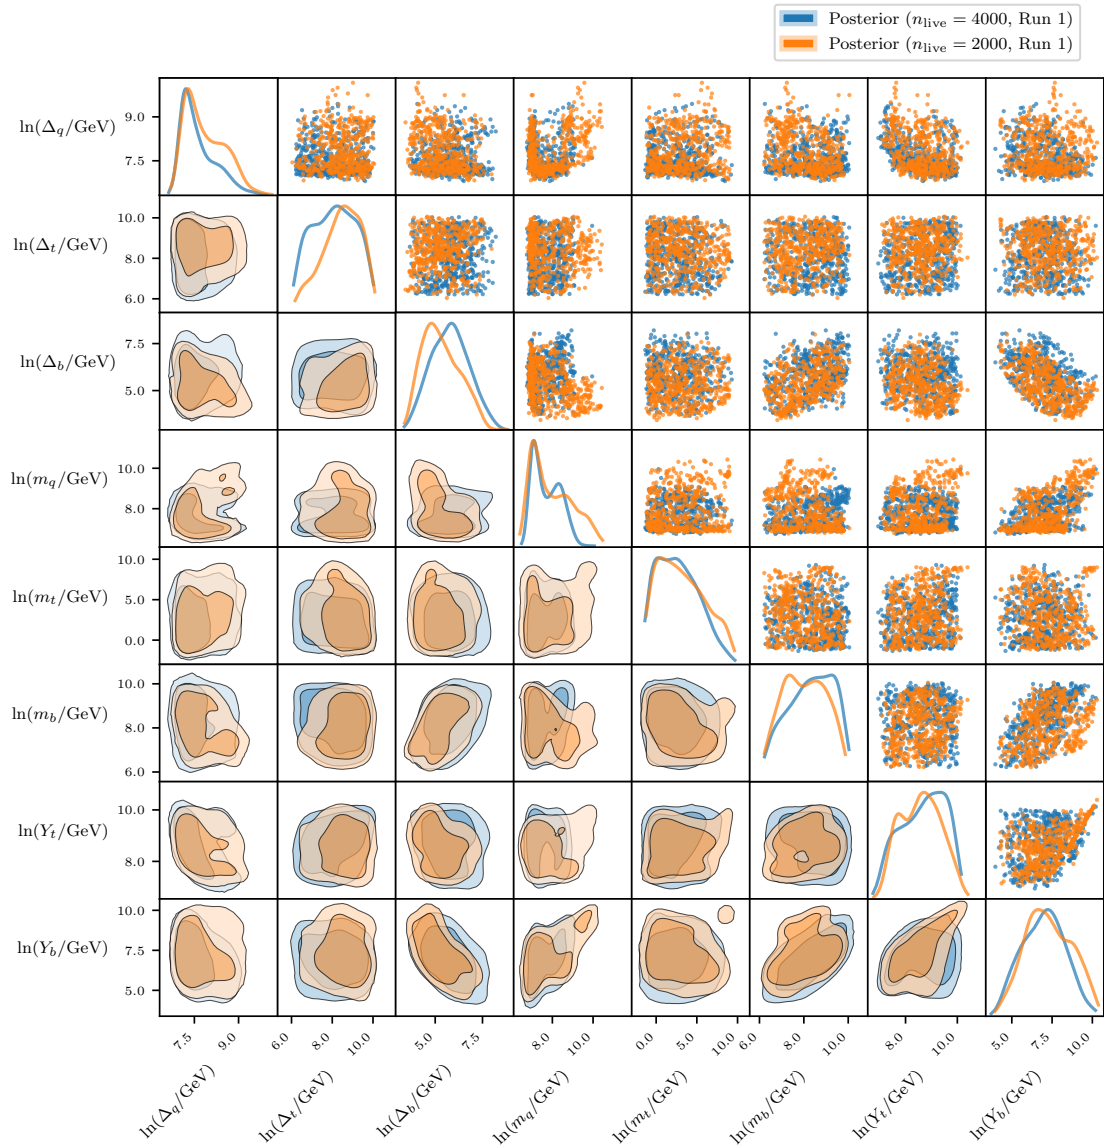


Figure 7.25: 1D and 2D marginalised posteriors for the fermion sector parameters in the M4DCHM¹⁴⁻¹⁻¹⁰ found in two different runs with 2000 and 4000 live points respectively.

As such, these results are quite reliable. The results from the combined runs are given in Figures 7.26 and 7.27.

At first glance of Figure 7.26, there is little qualitative difference between the priors and posteriors for the gauge sector couplings in this model and those in the previous models, but there are a few subtle differences upon closer inspection. Here the posterior for f is mostly contained between ~ 1.3 TeV and ~ 1.8 TeV, though having a tail reaching upwards of about 2.2 TeV. This approximately mirrors the result in the M4DCHM⁵⁻⁵⁻⁵, giving a relatively reasonable fine-tuning of $\sim 30-80$. And like in the M4DCHM¹⁴⁻¹⁴⁻¹⁰, f_1 tends to favour the upper end of its range, spanning between roughly 1.3 TeV and 3.4 TeV, and f_X tends towards the lower values, going from 0.5 TeV to around 5.8 TeV. But f_G is where the behaviour departs from that in the M4DCHM¹⁴⁻¹⁴⁻¹⁰, here having an affinity for *lower* values and ranging from 0.5 TeV to ~ 5.3 TeV, though again this difference is too slight to attribute a cause. The gauge coupling g_ρ behaves in the same way as in the other models, both in its preference for values from ~ 2 to ~ 8 , and in its relationships to the various constraints. However, the same is not true for the other gauge couplings. Surprisingly, in this model the posteriors for g_X and g_G almost exactly align with the priors, and indeed none of the χ^2 contributions from any of the constraints seem to depend on these parameters at all!

In contrast to the gauge parameters, there is a clear difference in the posteriors of the fermion parameters between the two minimally-tuned models: those in the M4DCHM¹⁴⁻¹⁻¹⁰, shown in Figure 7.27, cover a markedly larger fraction of volume of the total parameter space. So these parameters require less fine-tuning to satisfy the constraints than those in the M4DCHM¹⁴⁻¹⁴⁻¹⁰. This is taken to the extreme by the mass parameter m_t being almost entirely unconstrained in this model: it is able to take values up to ~ 10 TeV but with great preference towards values $\lesssim 0.5$ TeV purely stemming from the prior. Other parameters that require little fine-tuning, having posteriors well within the regions favoured by their priors, are Δ_q , Δ_b , and m_q . The posteriors of the first two of these respectively span from ~ 0.9 TeV to ~ 5.0 TeV, and from ~ 70 GeV to ~ 1.8 TeV, again reflecting the right-handed bottom quark being mostly elementary. The last of these, m_q , along with m_b , has a hard lower bound of 0.5 TeV from the resonance mass cutoffs, with the former ranging upwards of 8 TeV and the latter being drawn up to ~ 21 TeV by the constraints. If the LHC constraints were included in this fit, we might expect these lower bounds to be increased and larger parameter values to be favoured. Unlike the other models, here the right-handed top quark can mix very strongly with the composite sector, with the posterior for Δ_t going between ~ 0.5 TeV and ~ 17 TeV, being maximised at ~ 7.5 TeV despite the prior giving preferable weight to lower values. Similarly, the proto-Yukawa couplings Y_t and Y_b are drawn to unnaturally large values by the constraints, with posteriors ranging from ~ 1.6 TeV to ~ 22 TeV, and from ~ 50 GeV to ~ 63 TeV respectively. The posteriors for all of these parameters are shaped primarily by the SM mass constraints, along with some contributions from the Z decay constraints that slightly disfavour very large values of the couplings.




Figure 7.26: 1D and 2D marginalised priors and posteriors for the gauge sector parameters in the $M4DCHM^{14-1-10}$.

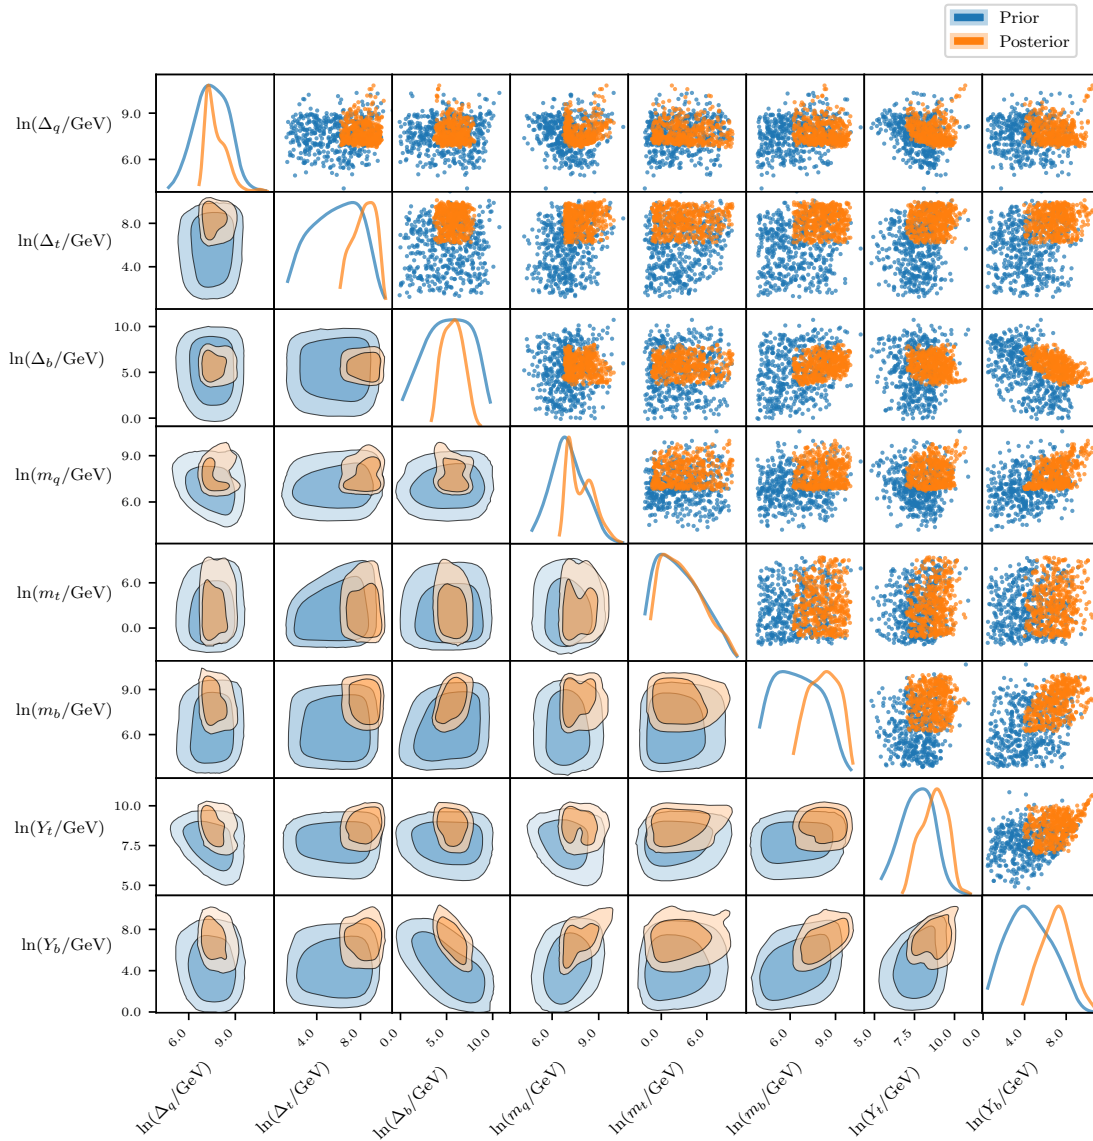


Figure 7.27: 1D and 2D marginalised priors and posteriors for the fermion sector parameters in the M4DCHM¹⁴⁻¹⁻¹⁰.

Several interesting features are present in the fermion parameter priors and posteriors in Figure 7.27. There is a slight positive correlation between the posteriors of m_b and Y_b , and a negative correlation in both the priors and posteriors of Δ_b and Y_b . There is also a notable effect wherein large values of Y_t are favoured in both the prior and posterior when Y_b is also large. The prior tends to disfavour both Δ_q and Y_t having small values, and to a lesser extent the same thing for both Δ_q and m_q , and the posteriors for these pairs seem to be constrained to regions for which $\Delta_q Y_t$ and $\Delta_q m_q$ lie above some respective fixed values. Unfortunately, analysing the causes of these features would be too involved a process to be viable for us.

As a final note, our result that Δ_t is usually much greater than m_t indicates that the right-handed top quark must be almost completely composite in this model given the constraints. It might therefore do well to consider a model with the same $14-1-10$ symmetry structure that incorporates t_R as an entirely composite state as a limiting case of the M4DCHM¹⁴⁻¹⁻¹⁰.

7.3.2 Frequentist Fit

As with the M4DCHM¹⁴⁻¹⁴⁻¹⁰, we have obtained generally consistent results for the M4DCHM¹⁴⁻¹⁻¹⁰ between two frequentist scans with populations of 20,000 and 50,000 points. Profile likelihood ratios for all of the parameters from the combined runs are presented in Figures 7.28 to 7.32. The best-fit point that was found here has a likelihood of $\ln(\mathcal{L}) = -15.4$.




Figure 7.28: Profile likelihoods for the decay constants in the M4DCHM¹⁴⁻¹⁻¹⁰.




Figure 7.29: Profile likelihoods for the gauge couplings in the M4DCHM¹⁴⁻¹⁻¹⁰.




Figure 7.30: Profile likelihoods for the couplings linking the elementary and composite sectors in the M4DCHM¹⁴⁻¹⁻¹⁰.




Figure 7.31: Profile likelihoods for the on-diagonal mass parameters in the M4DCHM¹⁴⁻¹⁻¹⁰.

Immediately in Figure 7.28 we see a distinction between this model and the previous two: the Higgs decay constant f has at the best-fit point a value *not* near the largest value allowed by the scans, but instead maximises the likelihood at ~ 2.8 TeV. At the 2σ confidence level, it ranges between approximately 2.0 TeV and 3.1 TeV, confining the misalignment to $0.08 \lesssim s_{\langle h \rangle} \lesssim 0.12$ and giving an estimated tuning of $\sim 65-160$. A fairly wide spread is present in the other decay constants, with f_1 , f_X , and f_G generally lying between roughly ~ 2 TeV and ~ 5 TeV at the 2σ level. The gauge coupling g_ρ is constrained between ~ 2.4 and ~ 3.3 , echoing the low range that was found in the M4DCHM¹⁴⁻¹⁴⁻¹⁰. No major preference for g_X is seen




Figure 7.32: Profile likelihoods for the proto-Yukawa parameters in the M4DCHM¹⁴⁻¹⁻¹⁰.

across its possible range, just as in the Bayesian fit, while g_G tends towards larger values $\gtrsim 5$, conflicting with the Bayesian results, presumably to best satisfy LHC search constraints on heavy gluon decays.

As happened in the M4DCHM¹⁴⁻¹⁴⁻¹⁰, the profile likelihood ratios for the fermion sector parameters in Figures 7.30 to 7.32 typically match the structure of the posteriors from the Bayesian fit, but at a higher scale. For example, Δ_q here takes values from ~ 2.0 TeV to ~ 5.4 TeV, while Δ_t is significantly larger, lying above ~ 24 TeV (discounting an outlier likely point at 6 TeV), and Δ_b goes between roughly 0.7 TeV and 10 TeV. The mass parameter m_q fits comfortably within the posterior range at the 2σ level, from ~ 1.8 TeV to ~ 2.7 TeV, as does m_t , which is constrained to be below around 3.0 TeV. Returning to larger values is m_b , which ranges from ~ 2.4 TeV to ~ 27 TeV.

It is interesting that in this fit the proto-Yukawa couplings Y_t and Y_b do *not* extend to the same large values that their posteriors reach. Respectively, they are constrained between approximately 2.3 TeV and 7.7 TeV, and between 0.1 TeV and 0.6 TeV, even though their posteriors both prefer values around 8.0 TeV in spite of their priors favouring lower values. Apparently, either the most likely region found in the frequentist fit is finely-tuned enough to be irrelevant in the Bayesian scans, or the LHC constraints in these frequentist scans pulled Y_t and Y_b back towards lower values. Recall from the discussion in Section 7.2.2 that very large proto-Yukawa couplings lead to extremely light composite resonances, so it makes sense that they would be disfavoured here.

Three correlations between parameters were found in this fit: two positive correlations between f and Δ_q and also Δ_b and m_b , and a negative correlation between Δ_b and Y_b , all shown in Figure 7.33. The latter was also present in the Bayesian fit, and is the same (negative) correlation seen in the M4DCHM¹⁴⁻¹⁴⁻¹⁰.




Figure 7.33: Profile likelihoods for parameters displaying slight correlations in the $M4DCHM^{14-1-10}$. Contours define the 1σ and 2σ confidence levels.

7.4 Discussion and Model Comparisons

Having seen the fit results for all of our models, we are now in a position to discuss the similarities and differences between the models, and compare them based on their fits.

We begin with the frequentist results, for which the only point of comparison is the likelihood of the best-fit region in each model. The likelihoods of the best-fit points mentioned previously are remarkably similar across models - all within a factor of 1.3 of each other - so on this front the models all seem about equally viable. More can be said by looking at the contributions to the total likelihood from each constraint for these points, which are listed in Table 7.1. Here, it is seen that not only are the *total* likelihoods of the points comparable, but most of the *individual* constraints are satisfied to a similar degree in all models. The SM mass constraints are generally satisfied very well, with the $M4DCHM^{14-14-10}$ slightly outperforming the other models in the bottom mass constraint. More difficulty was had in reproducing the correct EWSB scale as measured by the G_μ constraint, which is unsurprising in light of the unideal scanning method that was used (see Section 6.1). In any case, the best-fit point in the $M4DCHM^{5-5-5}$ gives the most realistic EWSB scale despite the (questionably relevant) double tuning of the Higgs potential in this model. The oblique S and T constraints that parameterise electroweak precision observables are satisfied fairly well in all models, while the Z decay constraints $R_{e,\mu,\tau,b}$ are by far the least well satisfied. For both of these sets of constraints, the best-fit point of the $M4DCHM^{5-5-5}$ has the poorest performance. Only very few points were found to satisfy the Z decay constraints to a considerably better degree than these best-fit points, possibly indicating intrinsic obstacles in our models inhibiting more realistic Z decays. The Higgs signal strengths were also found to be difficult to satisfy, particularly in the $M4DCHM^{14-14-10}$. Finally, the LHC search constraints prove mostly inconsequential for these points, whose high-mass resonances mostly evade the ranges probed by collider experiments. In all, the single highest-likelihood point found in the scans is from the $M4DCHM^{14-14-10}$, but it is no more likely than points in the other models by any significant margin.

Constraint	χ^2 Contribution		
	M4DCHM ⁵⁻⁵⁻⁵	M4DCHM ¹⁴⁻¹⁴⁻¹⁰	M4DCHM ¹⁴⁻¹⁻¹⁰
m_{bottom}	1.23	0.52	1.07
m_{top}	0.38	0.37	0.69
m_Z	0.01	0.01	0.02
m_W	0.04	0.04	0.05
m_H	0.04	0.00	0.00
G_μ	1.46	2.31	2.81
S, T	0.44	0.09	0.16
R_e	4.95	4.53	4.63
R_μ	7.67	6.88	7.06
R_τ	0.34	0.22	0.24
R_b	1.30	1.23	1.21
μ_{WW}^{gg}	2.06	2.37	1.89
μ_{ZZ}^{gg}	1.78	1.93	1.49
$\mu_{\gamma\gamma}^{gg}$	3.43	4.08	3.78
$\mu_{\tau\tau}^{gg}$	0.15	0.21	0.18
LHC (total)	0.00	0.00	0.01
Total	25.28	24.79	25.27

Table 7.1: χ^2 values contributed by each constraint to the total likelihood for the best-fit points in each model. The minimum possible χ^2 values have been subtracted from the Higgs signal strength contributions.

	$\ln(\mathcal{Z})$	D_{KL}	$\langle \ln(\mathcal{L}) \rangle_P$	$\max \ln(\mathcal{L})$	BMD
M4DCHM ⁵⁻⁵⁻⁵	-28.62 ± 0.04	15.71	-12.92	-8.62	7.23
M4DCHM ¹⁴⁻¹⁴⁻¹⁰	-41.08 ± 0.08	25.33	-15.75	-9.26	20.49
M4DCHM ¹⁴⁻¹⁻¹⁰	-37.58 ± 0.04	15.83	-21.75	-16.66	7.52

Table 7.2: Statistics from the combined Bayesian scans of each model.

What interests us most, however, are the Bayesian results, for they take into account the naturalness of the models. The Bayesian evidences \mathcal{Z} calculated from the combined samples from the scans of each model are given in Table 7.2, along with other quantities of interest such as the Kullback-Leibler (KL) divergence D_{KL} (see Section 5.1) and the posterior-averaged log-likelihood

$$\langle \ln(\mathcal{L}) \rangle_P = \ln(\mathcal{Z}) + D_{\text{KL}}, \quad (7.6)$$

as well as the maximum likelihoods found in the scans. Note that although the maximum likelihoods here are far better than those from the frequentist scans, these ones have the advantage of not taking the LHC search constraints into account. If the search constraints are included, they end up being the dominant contributions to the likelihoods of these points.

The most important point to note from Table 7.2 is that the M4DCHM⁵⁻⁵⁻⁵ has by far the greatest evidence of the three models, being ~ 7800 times greater than the evidence for the M4DCHM¹⁴⁻¹⁻¹⁰, which is itself ~ 12 times greater than that for the M4DCHM¹⁴⁻¹⁴⁻¹⁰. According to the traditional Jeffreys scale for interpreting evidences, the M4DCHM⁵⁻⁵⁻⁵ is then the “decisively” preferred model, while the M4DCHM¹⁴⁻¹⁻¹⁰ is “strongly” preferred over the M4DCHM¹⁴⁻¹⁴⁻¹⁰, granted we assign an equal prior likelihood to each of our models. Keep in mind, though, that these evidences are prior-dependent. We have not attempted fits using alternative priors, and so cannot comment on how strongly these evidences are coupled to the priors we have imposed on the spaces of these models.

The reasons for this hierarchy of preferences can be ascertained from Table 7.2. First, the M4DCHM¹⁴⁻¹⁴⁻¹⁰ is by far the least preferred model primarily because of its large KL divergence. Recall the KL divergence measures the information gained from going to the prior to the posterior probability, and hence is an indirect measure of fine-tuning. We have already recognised this fine-tuning in the M4DCHM¹⁴⁻¹⁴⁻¹⁰ by the narrowly constrained posteriors in Figures 7.16 and 7.17. On the other hand, the M4DCHM¹⁴⁻¹⁻¹⁰ is disfavoured mostly because it has a much smaller posterior-averaged log-likelihood, *i.e.* because it has a more difficult time satisfying the experimental constraints. Not only is this evident in the average log-likelihood, but also in the likelihood of the best-fit point found in this model, being drastically lower than those from the other models. The M4DCHM⁵⁻⁵⁻⁵ is the only model with both a low KL divergence (tuning) and relatively large average likelihood, which




Figure 7.34: Posteriors of the contributions to the χ^2 from the Higgs signal strengths in each model. The minimum possible χ^2 contributions have been subtracted from each constraint.

combine to give the greatest evidence. Note that fine-tuning has previously only been quantified in the literature of CHMs through the sensitivity of the observables with respect to the input parameters, either to first order with the Barbieri-Giudice measure [87, 100, 101, 117], or with a higher-order measure [30, 88, 136], and our approach of using the KL divergence as a fine-tuning measure has not been taken before.

It is of interest to determine the exact constraints that the M4DCHM¹⁴⁻¹⁻¹⁰ has trouble satisfying. This is most easily done by plotting the posteriors of the χ^2 contributions from the different constraints for each model. Most of the posteriors are essentially identical across models, with the only exceptions being those for the Higgs signal strength constraints, shown in Figure 7.34. Clearly, the M4DCHM⁵⁻⁵⁻⁵ can satisfy these constraints without problem, while the other models have more difficulty - particularly the M4DCHM¹⁴⁻¹⁻¹⁰. Indeed, the χ^2 contributions from the three signal strengths in Figure 7.34 are each larger in the M4DCHM¹⁴⁻¹⁻¹⁰ than those in the M4DCHM⁵⁻⁵⁻⁵ by ~ 6 , which roughly leads to an average log-likelihood difference of $(6 + 6 + 6)/2 = 9$ and thereby almost entirely accounts for the different average likelihoods listed in Table 7.2. Looking into the cause further, it is found that the M4DCHM¹⁴⁻¹⁻¹⁰ has such difficulty with the Higgs signal strengths due to tensions with the SM mass constraints. The most significant factor is the bottom quark mass, whose relation to the Higgs signal strengths is displayed in Figure 7.35: a realistic bottom mass typically entails undesirably small signal strengths. This effect is not present in the other models, or at least not to as great an extent as in the M4DCHM¹⁴⁻¹⁻¹⁰.

It also proves useful to analyse the priors for observables taken as constraints in each model, to see whether any model gives realistic predictions more naturally

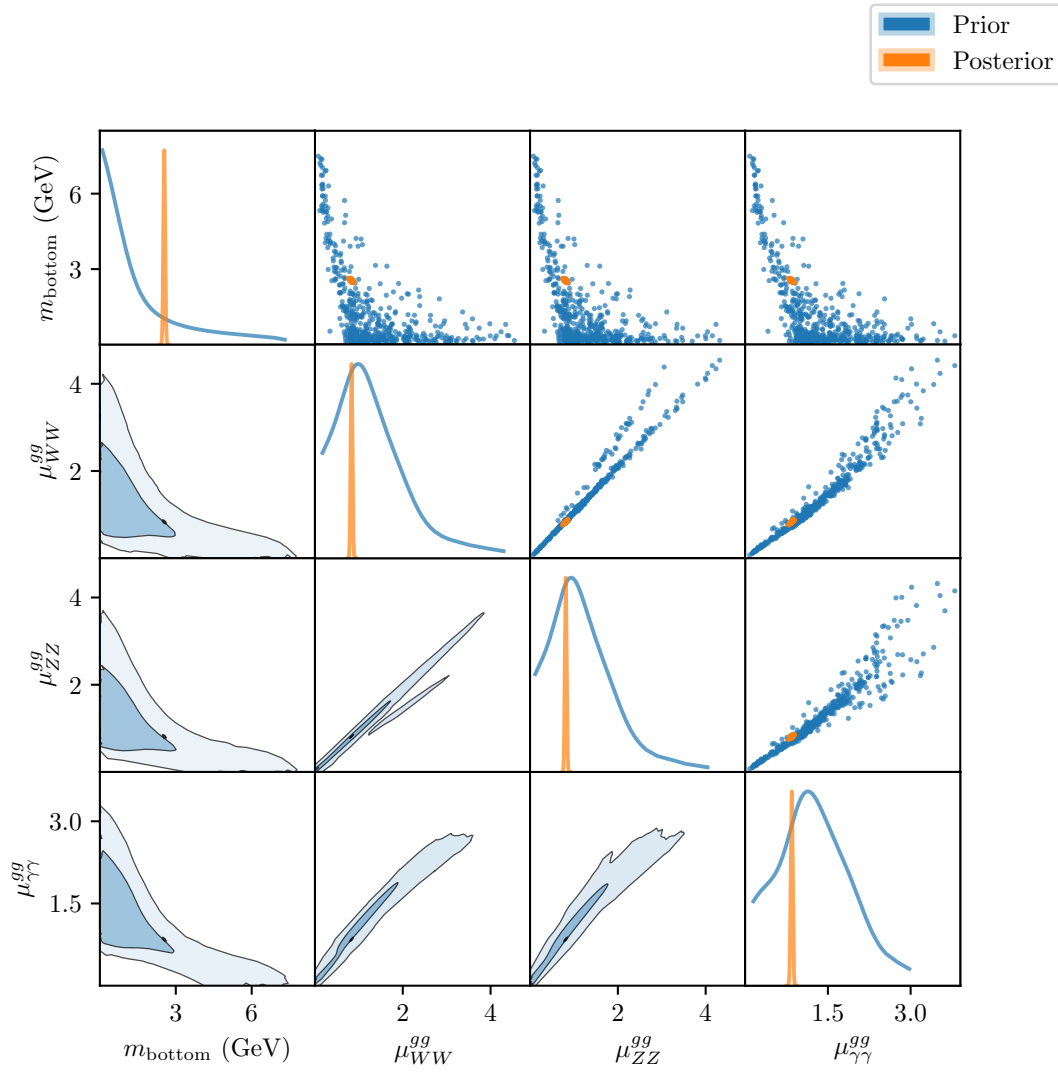


Figure 7.35: Priors and posteriors for the Higgs signal strengths and bottom mass in the M4DCHM¹⁴⁻¹⁻¹⁰.




Figure 7.36: Priors for the bottom, top, and Higgs masses in all models.

than others. We give the priors for the masses of the SM particles in Figure 7.36. As expected, the Higgs, being a pNGB, is generated most often with a low mass. But it also lies above 1 TeV relatively often, especially in the M4DCHM¹⁴⁻¹⁴⁻¹⁰, and to a lesser extent in the M4DCHM¹⁴⁻¹⁻¹⁰, both of which have been noted in other works to have a predisposition for producing too heavy a Higgs [87, 100]. The top and bottom quark masses, on the other hand, all tend to be generated relatively close to their experimental values⁵. The M4DCHM¹⁴⁻¹⁴⁻¹⁰ gives quite an unfavourable prior for the top mass, with very light values $\lesssim 80$ GeV being preferentially generated. Remarkably, the prior for the top mass in the M4DCHM⁵⁻⁵⁻⁵ has a maximum almost exactly at the experimental value. This would make it easier to reproduce the top mass in this model, leading to a lower fine-tuning and contributing to the model’s superiority. We are not aware of any literature where this has previously been noted.

We conclude this section with some general comments on the parameter spaces of our models. During our Bayesian scans, we noticed that the slice sampling steps take $\mathcal{O}(100)$ likelihood evaluations in each slice instead of the more typical $\mathcal{O}(5)$ evaluations, indicating that the likelihood is quite “sheety” in these models: it is only non-negligible across thin hypersurfaces in the parameter spaces. This notion can be quantified with the **Bayesian model dimensionality** (BMD) of each model, defined similarly to the KL divergence Equation (5.7),

$$\text{BMD} := 2 \int d^n p P(\mathbf{p}|\mathcal{D}, \mathcal{M}) \left(\ln \left(\frac{P(\mathbf{p}|\mathcal{D}, \mathcal{M})}{\pi(\mathbf{p}|\mathcal{M})} \right) - D_{\text{KL}} \right)^2, \quad (7.7)$$

⁵Those in the M4DCHM¹⁴⁻¹⁻¹⁰ span significantly shorter ranges than in the other models because we applied the hard mass cuts Equation (6.16) in this model, but not in the others (at least not in the Bayesian scans), on account of the computational expense required.

which measures the *effective* dimensionality of the region of parameter space that the posterior occupies [137]. This definition is used because it matches the *actual* dimensionality when the posterior is Gaussian. Broader posterior distributions will have lower BMDs, and narrower distributions higher BMDs. It can be difficult to intuit the BMD of a model based on the marginalised posteriors such as those given in Sections 7.1 to 7.3, since it is most often the case that the constraints restrict only certain non-trivial combinations of the input parameters, and these correlations would be obscured in a 2D plot by the marginalisation. In any case, these effective dimensionalities can be calculated with `anesthetic` and are listed for our models in Table 7.2.

Clearly, the $M4DCHM^{5-5-5}$ and $M4DCHM^{14-1-10}$ are indeed very sheety, with effective dimensionalities of ~ 7 in contrast to their respective true dimensionalities of 22 and 18 (including the SM gauge couplings). Positive constraints that restrict the space to reproduce particular values for the observables, such as the SM mass constraints⁶, are apparently not sufficient to localise the parameters completely, and leave these ~ 7 degrees of freedom within the spaces. These effective dimensionalities should be expected to be quite small on account of the fine-tuning required for EWSB and the effect of the constraints on the space, but it is difficult to say whether ~ 7 is a reasonable number of dimensions on purely theoretical grounds. In contrast, there seems to be little difference between the true and effective dimensionalities of the $M4DCHM^{14-14-10}$. This might be an artefact of the poor sampling of this model, indicating that our scans failed to locate the optimal ~ 7 dimensional surface for the posterior, assuming such a surface exists in this model. We aim to resolve this issue in our future scans of the $M4DCHM^{14-14-10}$.

It is further possible to calculate the principal directions in parameter space that point along these hypersurfaces with `anesthetic`. If these models were to be explored in future scans, it would be beneficial to use these directions to cleverly parameterise the spaces and scan along these optimal surfaces directly, in order to achieve greatly improved scanning efficiencies.

7.5 Experimental Signatures

In our fits, we have found a number of points in each model that satisfy all constraints at the 3σ level individually, which we regard as *valid* points. We have collected, in total, $\sim 14,000$ valid points in the $M4DCHM^{5-5-5}$, $\sim 77,000$ in the $M4DCHM^{14-14-10}$, and $\sim 280,000$ in the $M4DCHM^{14-14-10}$. Here we analyse the phenomenology of these experimentally viable points so that we may anticipate possible signatures of each model in future collider experiments.

⁶For example, it follows from Equations (H.23) to (H.25) that one degree of freedom is removed in order to reproduce each of the top and bottom masses.

Composite Resonances

We begin with the BSM particle spectra. Recall that these models contain several up-type (U) and down-type (D) composite fermions, exotic fermions Q_x with electric charges $x \in \{4/3, 5/3, 8/3\}$, as well as composite gauge bosons with unit charge (W) and neutral charge⁷ (Z). Our convention for distinguishing the different members of each species is to label them with numerical subscripts in order from lightest to heaviest, including SM states. For example, the lightest neutral composite gauge boson will be called Z_3 , being the lightest neutral boson after the photon (Z_1) and the electroweak Z boson (Z_2). The profile likelihood ratios⁸ of the masses of the lightest composite resonances of several of these species for each model are given in Figure 7.37. We see that broadly, the most likely points in all models give composite U and D resonances between 1.7 TeV and 4.5 TeV, and composite gauge bosons above 3 TeV. Let us discuss the mass spectra for each model in detail below.

Inasmuch as the frequentist results for the M4DCHM⁵⁻⁵⁻⁵ can be trusted, we see from Figure 7.37 that the lightest U and D resonances lie above 1.7 TeV at the 2σ level, with U_4 typically being the lighter of the two, being constrained below 3.7 TeV and D_4 below 4.5 TeV. The next three lightest of each of these species, and the lightest $Q_{4/3}$ and $Q_{5/3}$ resonances, range from around 2 TeV to 8 TeV, with best-fit values from 3 TeV to 6 TeV. The remaining fermions can be as light as 3 TeV, but generally range upwards of 30 TeV with best-fit values above 10 TeV. There are no apparent correlations between the masses across fermion species, aside from the ones that follow from Figure 4.3. As for the gauge bosons, Figure 7.37 shows that the Z_3 resonance is lighter than W_2 , lying above 3 TeV while W_2 lies above 6 TeV. The second lightest neutral composite boson, Z_4 , is degenerate with W_2 , with the remaining bosons being very heavy, having best-fit masses above 10 TeV.

Now for the M4DCHM¹⁴⁻¹⁴⁻¹⁰, recall from Section 7.2.2 that the parameter m_{Y_t} tends to be small in comparison to the on-diagonal masses. It then follows from Figure 4.3 that two pairs each of U, D, and $Q_{5/3}$ particles, and one pair each of $Q_{4/3}$ and $Q_{8/3}$ particles, will come with masses of approximately m_u and m_q . These are found to be the lightest composite resonances, so there is much approximate degeneracy among the lightest states. The three lightest U, D, and $Q_{5/3}$ resonances, and the lightest $Q_{4/3}$ and $Q_{8/3}$ particles have masses ranging from 1.8 TeV to 2.5 TeV at the 2σ confidence level, being 1.85 TeV at the best-fit point. These features can be seen in Figure 7.37. The remaining fermions (save for the three heaviest U resonances) all have best-fit masses at or below 4 TeV, but can range up to 8 TeV at the 2σ confidence level. Also degenerate are the lightest charged and neutral composite gauge bosons, as seen in Figure 7.37. The second lightest of these species act similarly, with each pair ranging from about 4 TeV to 10 TeV at the 2σ

⁷Not including the gluon sector.

⁸We do not analyse the masses in a Bayesian framework because our Bayesian scans did not include LHC search constraints, so many of the points found in them give resonances that are excluded by experiment.

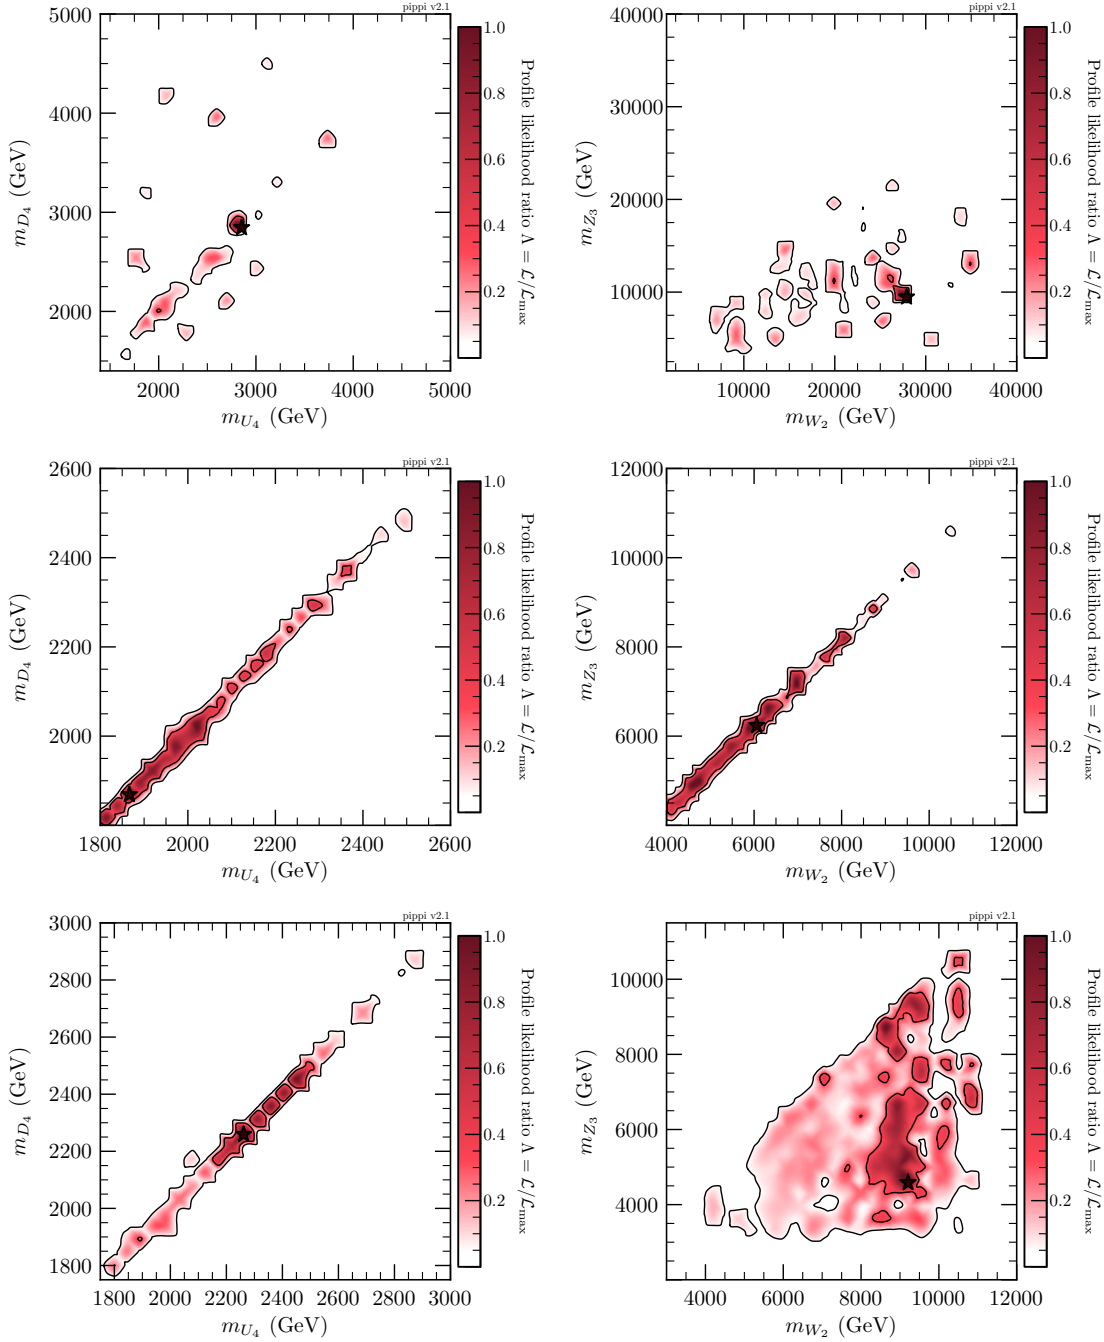


Figure 7.37: Profile likelihoods for the masses of the lightest up-type and down-type composite resonances U_4 and D_4 , as well as the lightest neutral and charged gauge bosons Z_3 and W_2 in (*top*) the $M4DCHM^{5-5-5}$, (*middle*) the $M4DCHM^{14-14-10}$, (*bottom*) the $M4DCHM^{14-1-10}$. Valid and non-valid points alike are taken into account in these plots.

confidence level and having masses between 6 TeV and 7 TeV at the best-fit point. The remaining gauge bosons all have best-fit masses above 10 TeV.

The analysis for the M4DCHM¹⁴⁻¹⁻¹⁰ is the most straightforward since the masses of all fermions are given by the simple formulae in Figure 4.3. It is found that the lightest fermions in this model are those with mass m_q , which recall from Section 7.3.2 spans 1.8 TeV to 2.7 TeV at the 2σ level. Among these are the three lightest U types, the two lightest D and $Q_{5/3}$ types, and the $Q_{4/3}$ and $Q_{8/3}$ resonances. The rest of the fermions have best-fit masses below 8 TeV but can range up to 25 TeV, except for U_{14} , which lies above 22 TeV. As in the M4DCHM⁵⁻⁵⁻⁵, the lightest composite boson in this model is Z_3 , ranging from 3 TeV to 10 TeV, while W_2 and W_3 are approximately degenerate with Z_4 and lie between 4 TeV and 11 TeV. The remaining bosons, including the heavy gluons, once again tend to be heavier than 10 TeV.

Predicted Cross Sections

Such resonances would be discovered in collider experiments, similarly to the Higgs boson, through the observation of excesses in the invariant mass distributions of the resonance decay products. For the valid points we have found, we give the cross sections of such processes involving the lightest composite resonances that would be measured at the 13 TeV LHC, calculated through procedures outlined in Sections 6.2 and 6.3, in Figures 7.38 to 7.41. Note that many points in Figures 7.38 to 7.41 lie at resonance masses below the 2σ lower bounds mentioned above, especially in the M4DCHM⁵⁻⁵⁻⁵, whose double tuning favours lighter composite partners. This is because the 2σ bounds are shaped by the regions of highest total likelihood, but the valid points only need to satisfy every constraint to a reasonable degree *separately*, and are not guaranteed to possess high *total* likelihoods.

The most promising channels for probing a particular model are those for which the model predicts cross sections within a narrow range that lies close to the current experimental upper bounds. For the U_4 decays in Figure 7.38, we see this is the case for the $U_4 \rightarrow bW^+$ decay channel in the M4DCHM⁵⁻⁵⁻⁵, and the $U_4 \rightarrow tZ$ and $U_4 \rightarrow tH$ channels in the M4DCHM¹⁴⁻¹⁻¹⁰. Unfortunately, the M4DCHM¹⁴⁻¹⁴⁻¹⁰ predicts heavier U_4 resonances, mostly above 1.75 TeV, that are more difficult to detect. Most of the points in Figure 7.38 predict cross sections for D_4 decays that are well out of reach of experiments for the near future, with the exception of the $D_4 \rightarrow bZ$ channel for the lower-mass resonances in the M4DCHM⁵⁻⁵⁻⁵. On the other hand, exotic fermion decays in Figure 7.39 display the ideal behaviour for probing the M4DCHMs: all three models make remarkably precise predictions for the $Q_{4/3_1} \rightarrow \bar{b}W^+$ and $Q_{5/3_1} \rightarrow tW^+$ decays, with cross sections comparable to the upper bounds towards the lighter end of the mass spectrum. Probing these channels further can easily confirm or rule out exotic resonances of certain masses and thereby constrain the models rather straightforwardly, although the cross sections fall off rapidly at higher masses, so much of the space remains out of reach of the LHC.

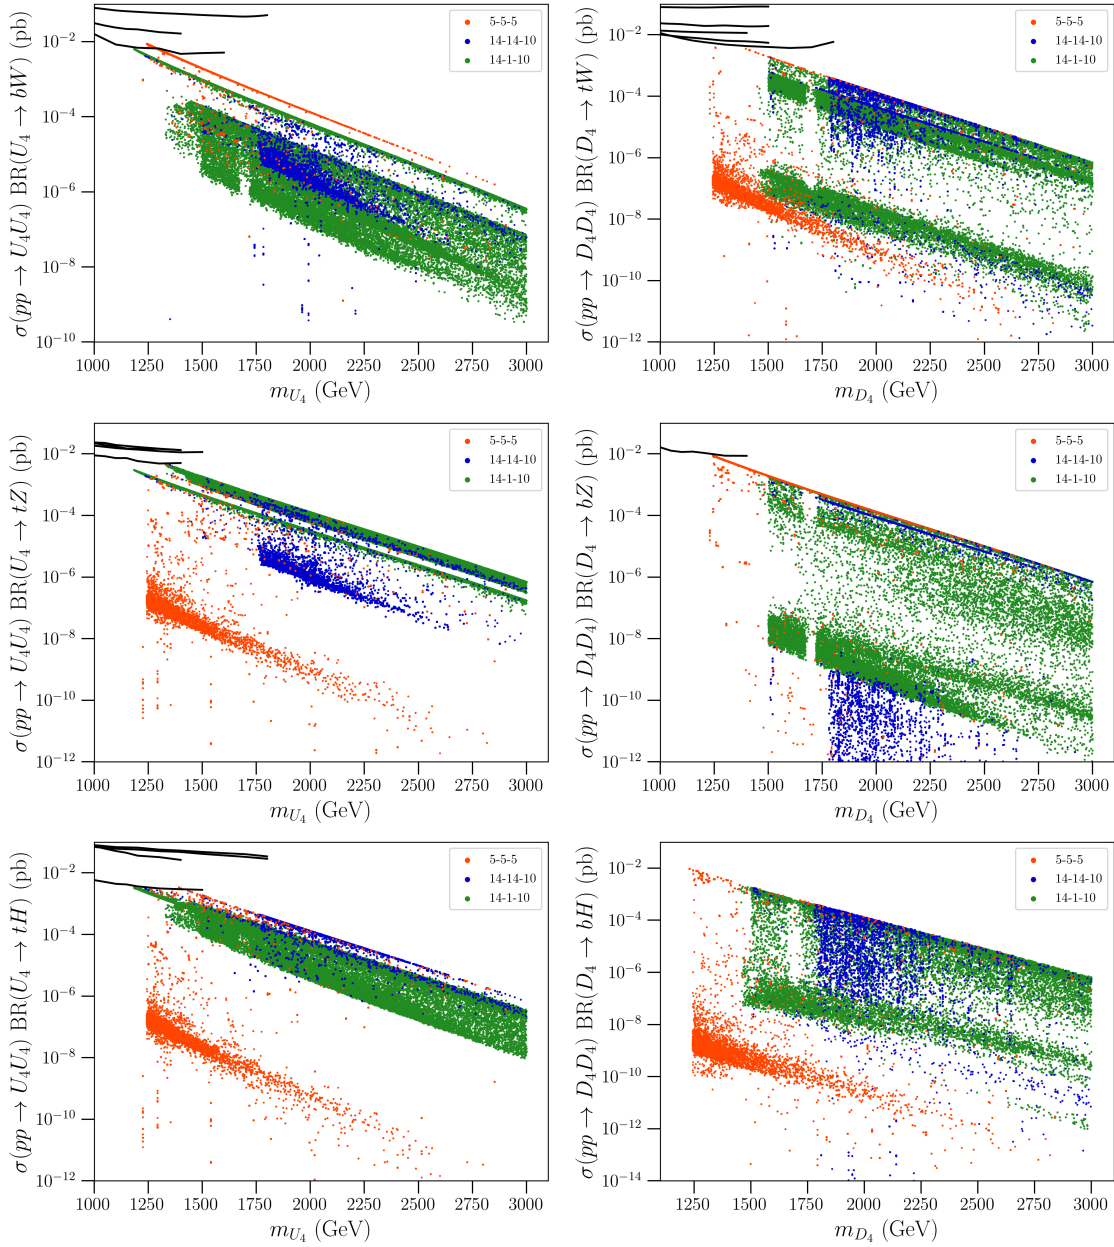


Figure 7.38: Cross sections for decays of the lightest up-type and down-type resonances U_4 and D_4 into SM final states at $\sqrt{s} = 13$ TeV for valid points in each model. Black lines mark the 95% CL upper bounds taken as constraints.




Figure 7.39: Cross sections for decays of the lightest resonances Q_x with exotic electric charge x into SM final states at $\sqrt{s} = 13$ TeV for valid points in each model. Most of the points in each model lie on the prominent lines in these plots. Black lines mark the 95% CL upper bounds taken as constraints.

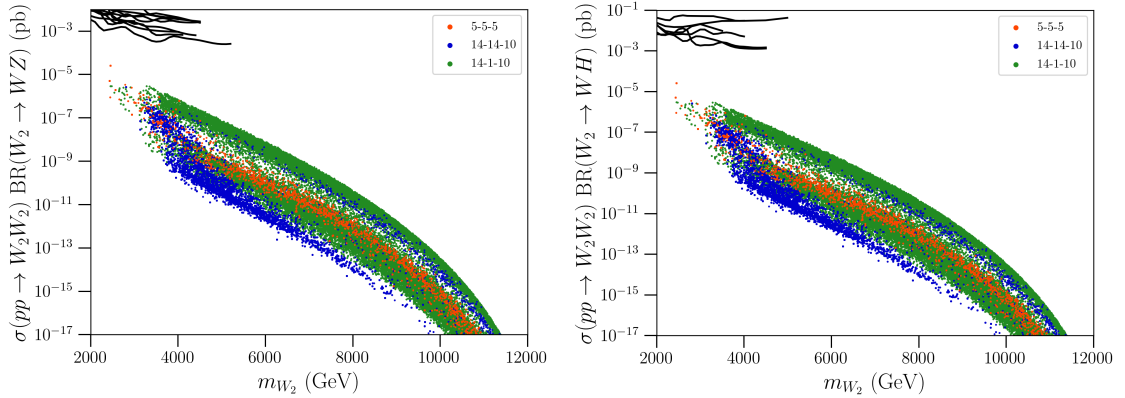


Figure 7.40: Cross sections for decays of the lightest charged composite vector boson W_2 into SM final states at $\sqrt{s} = 13$ TeV for valid points in each model. Black lines mark the 95% CL upper bounds taken as constraints.

Based on Figure 7.40, we would not expect to detect any charged composite gauge boson decays in the analysed channels under any of the models, as their predicted cross sections are well below what can be probed with current collider sensitivity. In addition to the channels in Figure 7.40, we also considered $W_2 \rightarrow t\bar{b}$ decays, but the predicted cross sections for this process were all negligible and not worth presenting. There is better hope for the neutral gauge boson decays in Figure 7.41, with all models predicting similar cross sections. Promising channels here are $Z_3 \rightarrow t\bar{t}$, and to a lesser extent $Z_3 \rightarrow W^+W^-$. However, many points predict resonance masses too large to be probed in the near future, so failure to detect any resonances at the LHC, for example, would not be sufficient to rule out any of these models.




Figure 7.41: Cross sections for decays of the lightest neutral composite vector boson Z_3 into SM final states at $\sqrt{s} = 13$ TeV for valid points in each model. Black lines mark the 95% CL upper bounds taken as constraints.

Higgs Signal Strengths

Other important theoretical predictions of the M4DCHMs that we analyse are the (gluon-fusion produced) Higgs signal strengths. Recall from Section 6.3 that the signal strength μ_{XX}^{gg} is the ratio of the $gg \rightarrow H \rightarrow XX$ cross section to the value predicted by the SM.

Modifications to the Higgs couplings in a composite Higgs framework have been subject to considerable analysis. To summarise the points that concern us here, in CHMs there are three sources of new physics that can modify the gluon-fusion production of a Higgs boson compared to the SM: non-linearities of the pNGB Higgs (dependent only on v/f), modified Yukawa couplings of SM particles, and loop contributions from composite resonances [138]. In our models, contributions from the latter two factors mostly cancel [139], so the contributions from the Higgs non-linearities are expected to dominate. In the limit where these are the only contributions, the signal strengths will be of the form

$$\mu_{XX}^{gg} = 1 - C_x \frac{v^2}{f^2} + \mathcal{O}\left(\frac{v^4}{f^4}\right), \quad (7.8)$$

for some channel-dependent (and model-dependent) constants C_x [138]. It is possible for Equation (7.8) to be significantly violated if the light quarks are considerably composite [140], but our models do not allow for this case. Further discussion on this analytic approximation and the subleading contributions to the signal strengths can be found in Ref. [100].

We give the signal strengths⁹ of the valid points in each of our models, along with their profile likelihood ratios from the frequentist scans (which include valid and non-valid points), in Figures 7.42 to 7.44. The signal strengths μ_{ZZ}^{gg} and μ_{WW}^{gg} are approximately equal due to the custodial symmetry of the M4DCHMs. It is immediately apparent in Figures 7.42 to 7.44 that the approximation Equation (7.8) does not hold well, especially in the M4DCHM⁵⁻⁵⁻⁵, but there do seem to be some structures of points that loosely follow the $-C_x v^2/f^2$ deviation from unity, which is most clearly seen for $\mu_{\gamma\gamma}^{gg}$ in the minimally-tuned models. This is in contrast to Ref. [101], where clear curves of points were found. It is possible that our points do not follow this distribution cleanly because we consider larger values of f , so the curves are less pronounced, and because we have heavier composite resonances that contribute to greater deviations from Equation (7.8).

It should also be noted that the frequentist and Bayesian scans found separate regions of points that form quite distinct structures in Figures 7.42 to 7.44, such as the regions of points in the M4DCHM⁵⁻⁵⁻⁵ that lie around the experimental values and those that lie mostly below unity, or the narrow bands of points that extend to large f and the wide bands of points constrained to $1 \text{ TeV} \lesssim f \lesssim 2 \text{ TeV}$ in the M4DCHM¹⁴⁻¹⁻¹⁰. Since the frequentist scans explored the most *likely* regions, and the Bayesian scans the most *natural* regions of parameter space, it is not too surprising that the two approaches found different behaviours for the observables. The regions can be distinguished by the profile likelihood plots, since these are sensitive to only the most likely points from the frequentist scans.

Among the valid points in the M4DCHM⁵⁻⁵⁻⁵, the predicted signal strengths in Figure 7.42 span wide ranges of values between 0.8 and 1.3. It is therefore unlikely that further probing the signal strengths could rule out or provide evidence for the M4DCHM⁵⁻⁵⁻⁵. Interestingly, the profile likelihood plots in Figure 7.42 show that the best-fit points in this model give signal strengths that are significantly lower than the individually optimal experimental values, so there must be some tension between satisfying the signal strength constraints and one or more other constraints. This is also the case for the minimally-tuned models, which give similar predictions to each other as seen in Figures 7.43 and 7.44. All the valid points in these two models give signal strengths that are less than the SM prediction of unity, excepting a small portion of points for μ_{WW}^{gg} and μ_{ZZ}^{gg} . Furthermore, we see a remarkable prediction of these models: that the signal strength $\mu_{\gamma\gamma}^{gg}$ lies in a narrow range from ~ 0.8 to ~ 0.9 . These clear predictions may serve as tests of the M4DCHM¹⁴⁻¹⁴⁻¹⁰ and M4DCHM¹⁴⁻¹⁻¹⁰ upon more precise measurements of Higgs decays. It is projected

⁹The signal strengths for the $h \rightarrow \tau\tau$ decay channel are not presented (even though this process is taken as a constraint) because this channel is subject to large uncertainties. In any case, the predicted values tend to be similar to those for the $h \rightarrow \gamma\gamma$ channel.

that $\mu_{\gamma\gamma}^{gg}$ will be measured to an uncertainty of around 5% by each ATLAS and CMS once they achieve integrated luminosities of 3000 fb^{-1} in the future high-luminosity run of the LHC [141, 142], so it is highly likely, based on current measurements, that this channel will disfavour either both of these models or the SM in the coming years. At the same time, μ_{ZZ}^{gg} will see a similar improvement in its measurement, while μ_{WW}^{gg} will have only modest improvements. But μ_{ZZ}^{gg} is not so narrowly predicted, so the diphoton decay channel will indeed serve as the best test of the M4DCHMs.

Curiously, although the two minimally-tuned models give the same predictions for $\mu_{\gamma\gamma}^{gg}$, they arrive at these predictions for different reasons. In the case of the M4DCHM¹⁴⁻¹⁴⁻¹⁰, it is the direct collider search constraints that seemingly require $\mu_{\gamma\gamma}^{gg}$ to be significantly less than unity; if these constraints were to be discounted, there would be a sizeable number of valid points with signal strengths $\mu_{\gamma\gamma}^{gg} \in [0.8, 1.2]$. The main search constraints that are violated for larger values of $\mu_{\gamma\gamma}^{gg}$ are from the fermion decays $F \rightarrow tW$, bW , qW , tH , bH , tZ , and jZ , where q refers to a light quark jet and j refers to a light quark or b jet. Each of these constraints is violated at the 3σ level for about 30% of points for which the $\mu_{\gamma\gamma}^{gg}$ constraint is satisfied to within 2σ . The combined effect is that all valid points have a $\mu_{\gamma\gamma}^{gg}$ below around 0.9.

For the M4DCHM¹⁴⁻¹⁻¹⁰, on the other hand, it is the SM mass constraints that favour the lower values of $\mu_{\gamma\gamma}^{gg}$. We have seen as much for the bottom mass in Figure 7.35. The same is true to a lesser degree for the top quark and Higgs masses, and satisfying all three mass constraints at once is what results in all valid points having $\mu_{\gamma\gamma}^{gg} \lesssim 0.9$. Note also from Figure 7.35 that the Higgs signal strengths are all highly correlated - a property shared by the other models as well. The fact that the SM mass constraints were included in the Bayesian scans, while the direct search constraints were not, might explain why the M4DCHM¹⁴⁻¹⁻¹⁰ had more difficulty satisfying the signal strength constraints than the M4DCHM¹⁴⁻¹⁴⁻¹⁰. But note that if this unfair advantage were to be fixed, the disparity between the evidences of the M4DCHM¹⁴⁻¹⁴⁻¹⁰ and the M4DCHM¹⁴⁻¹⁻¹⁰ would only be exacerbated, and the M4DCHM⁵⁻⁵⁻⁵ would remain the superior model from a Bayesian standpoint.

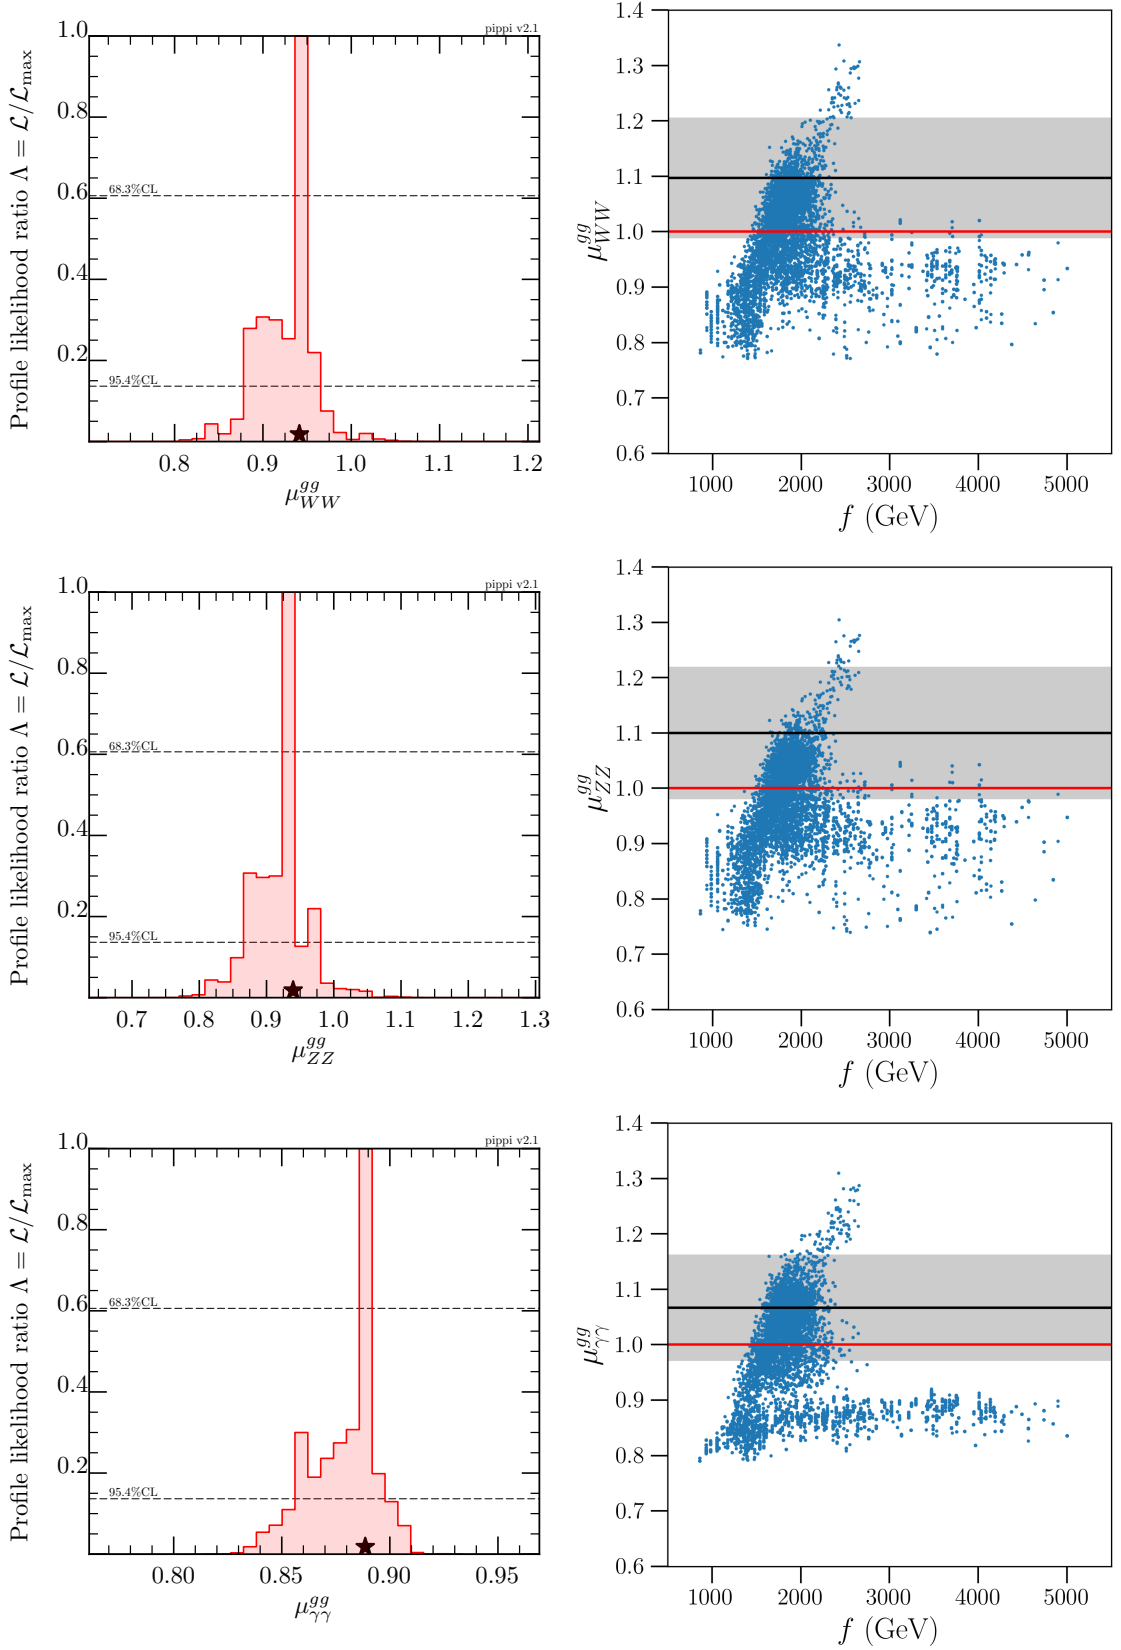


Figure 7.42: *Left:* profile likelihood ratios for the Higgs signal strengths in the M4DCHM⁵⁻⁵⁻⁵. *Right:* Higgs signal strengths for valid points in the M4DCHM⁵⁻⁵⁻⁵. Red lines show the SM predictions $\mu_{XX}^{gg} = 1$, while black lines mark the best-fit values, and the shaded regions the 1σ uncertainties, of the combined measurements of these processes.

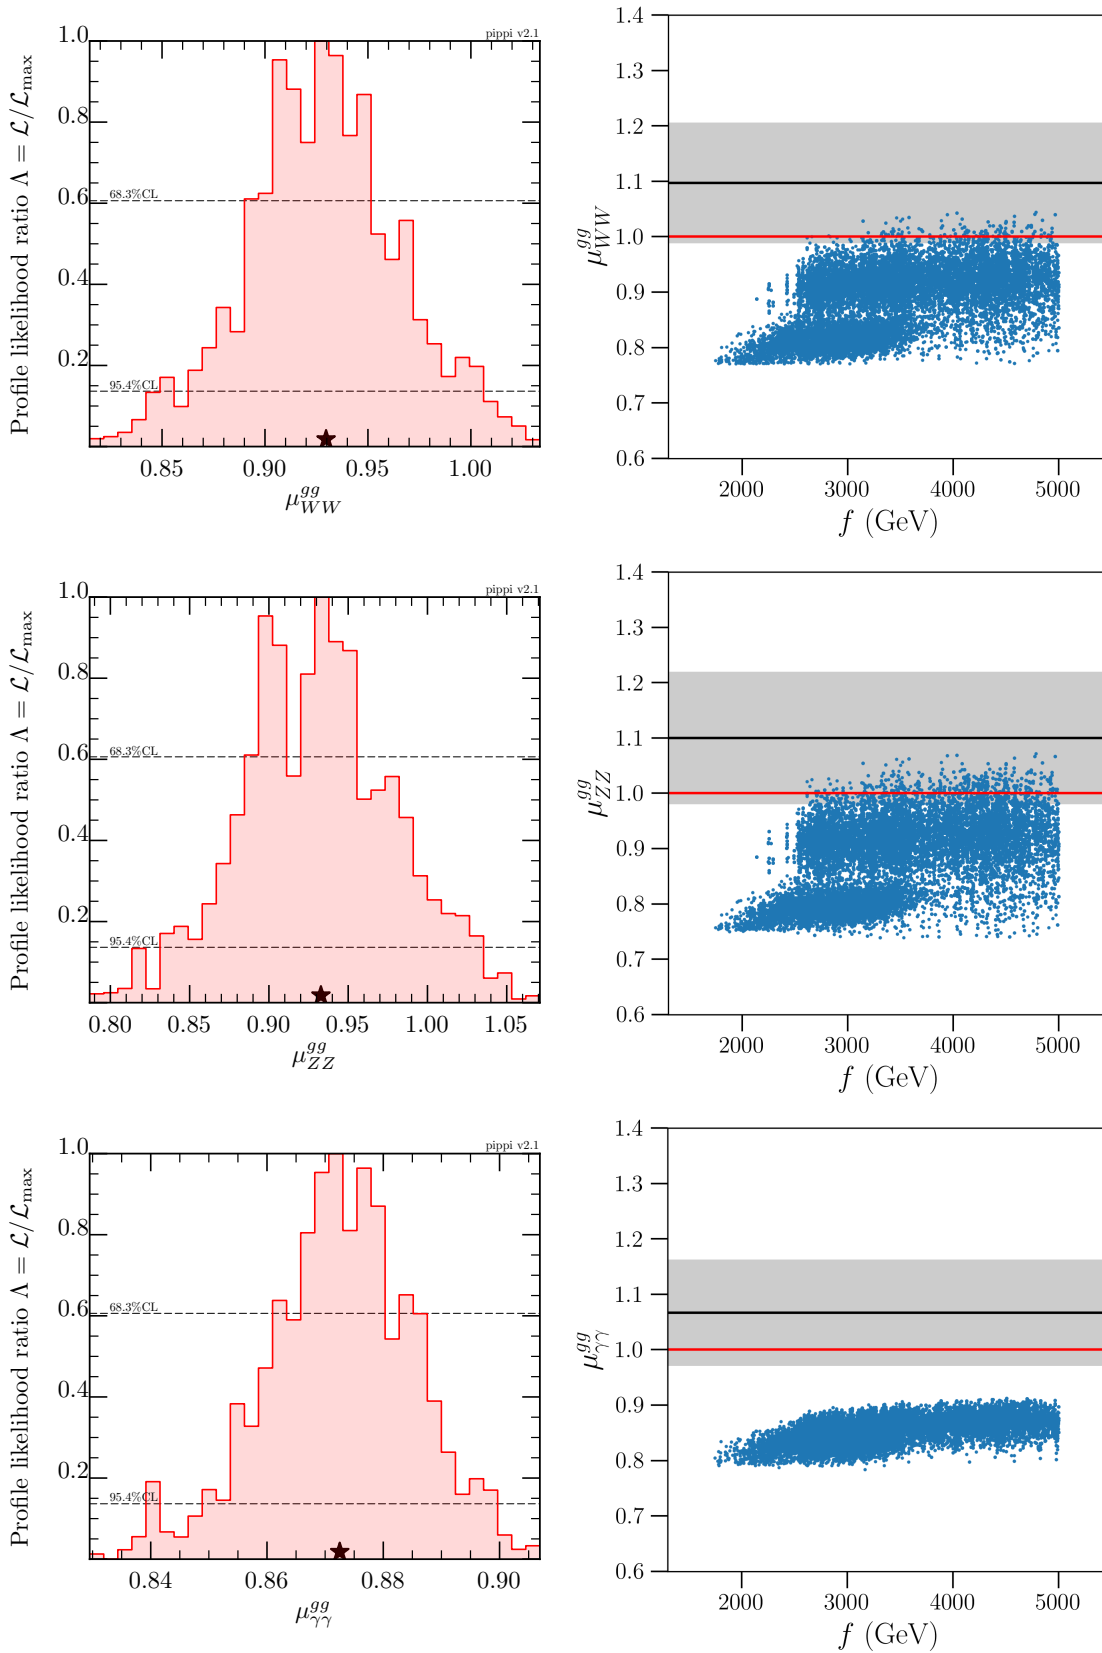


Figure 7.43: *Left:* profile likelihood ratios for the Higgs signal strengths in the M4DCHM¹⁴⁻¹⁴⁻¹⁰. *Right:* Higgs signal strengths for valid points in the M4DCHM¹⁴⁻¹⁴⁻¹⁰. Red lines show the SM predictions $\mu_{XX}^{gg} = 1$, while black lines mark the best-fit values, and the shaded regions the 1σ uncertainties, of the combined measurements of these processes.

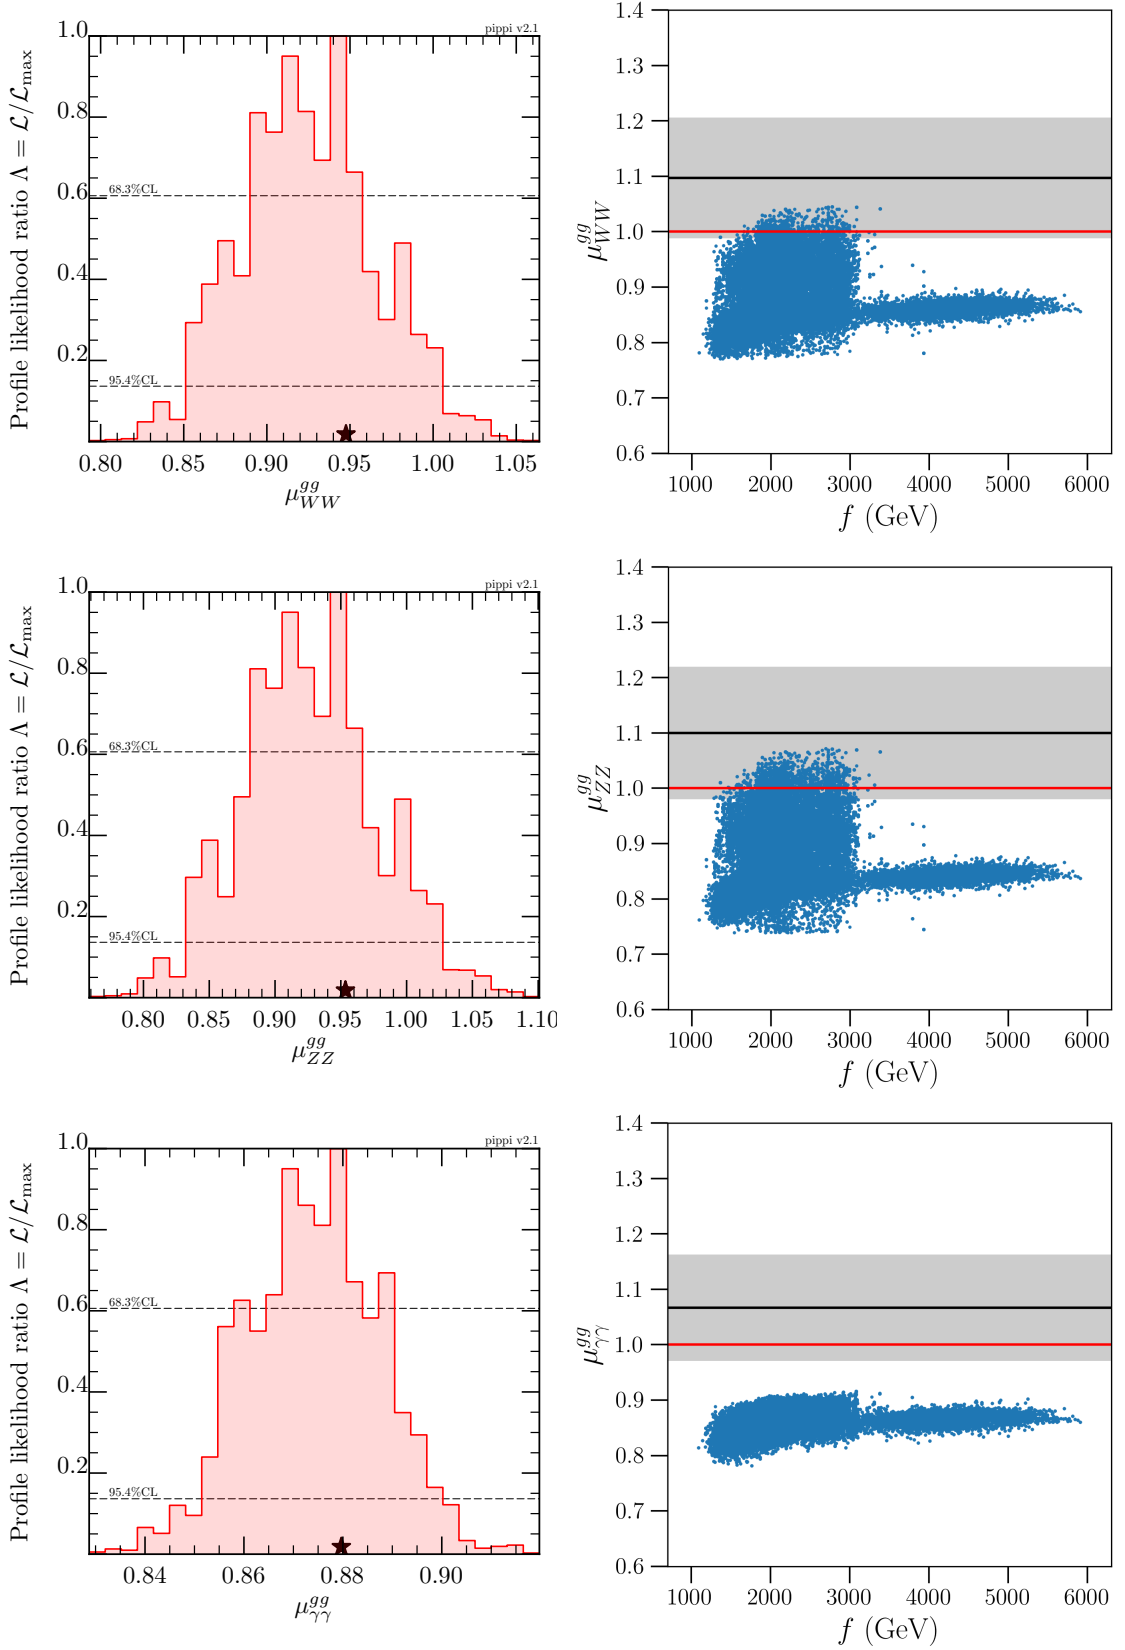


Figure 7.44: *Left:* profile likelihood ratios for the Higgs signal strengths in the M4DCHM¹⁴⁻¹⁻¹⁰. *Right:* Higgs signal strengths for valid points in the M4DCHM¹⁴⁻¹⁻¹⁰. Red lines show the SM predictions $\mu_{XX}^{gg} = 1$, while black lines mark the best-fit values, and the shaded regions the 1σ uncertainties, of the combined measurements of these processes.

Conclusion

We are in a position today where the prevailing theory of particle physics - the Standard Model - is known to be incomplete, but is so successful at describing most accessible phenomena that the correct extension of the theory remains a mystery. One of the few objectionable aspects of the SM that can serve as motivation for proposed extensions is that it offers no protection to the Higgs mass from large quantum corrections, and so can only accommodate the relatively low mass of the Higgs boson through an incredible fine-tuning of its model parameters. In much the same way that theoretical considerations provided hints of interesting physics around the electroweak scale that ultimately led to the discovery the Higgs boson, this so-called Hierarchy Problem suggests there is some new physics at the TeV scale shielding the Higgs mass from large corrections that could be discovered at the LHC. One attractive possibility is that the Higgs boson is in fact a composite state of some new TeV-scale strong dynamics, for then it could not exist at arbitrarily high energies and so would be impervious to high-energy quantum corrections. But so far experiments at the LHC have found no evidence for any such composite Higgs model, or indeed for any proposed solution to the Hierarchy Problem, and as the excluded energy range on their new physics grows ever larger, such theories are seeming ever less natural. We have offered in this thesis comprehensive numerical explorations of several prospective CHMs to assess the viability of such models in light of these recent collider searches and other experimental constraints.

Our models have been chosen from the vast landscape of CHMs as those that minimally extend the SM, with structures motivated by strict phenomenological and practical demands expounded upon in Chapter 3. For example, to account for the fact that no other composite resonances have yet been found in collider experiments, these models suppose the Higgs to be a pseudo-Nambu-Goldstone boson so that it is naturally lighter than the rest of the composite sector. This implies an approximate symmetry of the composite sector that is spontaneously broken, and it can be reasoned that the $SO(5) \rightarrow SO(4)$ symmetry breaking pattern is the minimal group structure that gives rise to the Higgs doublet in the NGB spectrum while ensuring custodial symmetry. The $SO(5)$ symmetry is explicitly broken by linear couplings between the SM-like elementary sector and the composite sector, which therefore provide a quantum effective potential for the Higgs through the Coleman-Weinberg mechanism. Such linear couplings are favourable because they can accommodate known flavour physics and explain the hierarchy of fermion masses, and they are

natural features in the 4D models dual to warped 5D theories that also attempt to solve the Hierarchy Problem. The effective potential appoints a non-zero vev to the Higgs field that misaligns the vacuum from the $SO(4)$ -symmetric direction, spontaneously breaking the embedded electroweak subgroup and giving mass to the SM particles. But ordinarily, since the theory is non-renormalisable, the effective potential is not calculable. In order to make predictions in such models, we demand a sufficiently large symmetry structure to render the potential calculable and finite, making use of the notion of multi-site models - deconstructions of the dual 5D theories - to arrive at the theory with the minimal viable calculable structure: the Two-Site Minimal 4D CHM.

Within the Two-Site M4DCHM, some freedom exists in the choice of $SO(5)$ representations for the composite fermion fields. We consider three versions of the M4DCHM that differ in their fermion sector embeddings, taking the limiting case where only the third generation quarks couple to the composite sector. Our models embed the composite partners of the left-handed quark doublet, right-handed top quark, and right-handed bottom quark into the respective representations $(\mathbf{5}, \mathbf{5}, \mathbf{5})$, $(\mathbf{14}, \mathbf{14}, \mathbf{10})$, and $(\mathbf{14}, \mathbf{1}, \mathbf{10})$. The first of these models, with all partners in the fundamental representation, is naively expected to suffer from a double tuning in the Higgs potential, while the others are minimally tuned. The second model offers an extra Yukawa structure not present in the others, and the third can entertain the possibility of an entirely composite right-handed top quark. Other viable embeddings display qualitatively similar behaviour to the three considered here [100].

We have performed extensive numerical explorations of these three models that constitute the first convergent global fits of realistic CHMs. Our fits constrain the parameter spaces into those regions that best reproduce, collectively, the SM masses, the electroweak scale, electroweak precision observables, various Z boson decay ratios and Higgs signal strengths, and that best satisfy heavy resonance production bounds from collider experiments. Software developed for Ref. [101] was employed to calculate these observables, which we modified with updated experimental values, including 40 new heavy resonance searches from Run II of the LHC. The fits are carried out under both frequentist and Bayesian statistical frameworks, aided by advanced scanning algorithms. A differential evolution algorithm provided by the `Diver` package, λ jDE, was used to find the best-fit regions of parameter space for the frequentist fits, while the Bayesian fits were instead performed with the nested sampling program `PolyChord`, which uses slice sampling to explore the spaces in search of the most natural viable regions. Adequately fitting these models proved to be quite difficult on account of their large parameter spaces and the highly non-trivial effects of the elementary-composite couplings, and indeed because of this we consider our results from the frequentist fit of the M4DCHM⁵⁻⁵⁻⁵ and Bayesian fit of the M4DCHM¹⁴⁻¹⁴⁻¹⁰ to be only preliminary.

All three models were found to have comparable likelihoods at their best-fit points in the frequentist fits, with each further satisfying the individual constraints to a similar degree. Profile likelihood ratios were used to approximate confidence

intervals for the input parameters, and it was found that the Higgs boson decay constant f - a rough measure of the scale of new physics - prefers quite large values in the M4DCHM⁵⁻⁵⁻⁵ and M4DCHM¹⁴⁻¹⁴⁻¹⁰, where it respectively lies above 2.1 TeV and 3.7 TeV at the 2σ confidence level, but is constrained between 2.0 TeV and 3.1 TeV in the M4DCHM¹⁴⁻¹⁻¹⁰. These ranges correspond to estimated fine-tunings of at least ~ 70 , ~ 230 , and ~ 65 in the respective models. Regions with lower tuning are preferentially weighted in the Bayesian fits, which accordingly gave marginalised posteriors for f over lower ranges between 1.2 TeV and 2.2 TeV.

The most notable result from the Bayesian fits is that the M4DCHM⁵⁻⁵⁻⁵ was found to have a Bayesian evidence many orders of magnitude greater than those of the other models with the imposed priors, and so is decisively superior from a Bayesian standpoint. Tension between the Higgs signal strength and SM mass constraints was the main obstruction to a higher evidence for the M4DCHM¹⁴⁻¹⁻¹⁰, while a considerable fine-tuning resulted in the M4DCHM¹⁴⁻¹⁴⁻¹⁰ having the lowest evidence by an order of magnitude. We took a novel approach in measuring the fine-tuning within a composite Higgs framework, quantifying it with the Kullback-Leibler divergence between the posterior and the prior. This viewpoint led us to discover that the top quark mass is incorporated significantly more naturally in the M4DCHM⁵⁻⁵⁻⁵ than in the other models. Interestingly, it was found that the posterior occupies surfaces of only ~ 7 effective dimensions in the (18+ dimensional) parameter spaces of the M4DCHM⁵⁻⁵⁻⁵ and M4DCHM¹⁴⁻¹⁻¹⁰, illustrating why these models are so difficult to fit. A similar result was not found for the M4DCHM¹⁴⁻¹⁴⁻¹⁰, but this might simply be a reflection of its poor fit.

Each model was found to be capable of satisfying all constraints individually to within 3σ , and their collider phenomenology was analysed in these viable regions. Our results suggest that the lowest-lying fermionic resonances in these models are excluded below ~ 1 TeV, generally lying in the ~ 1.1 TeV to ~ 3.0 TeV mass range. We expect the models to have different signals for the various decay modes of the up-type and down-type resonances at the $\sqrt{s} = 13$ TeV LHC, but remarkably all models give the same precise predictions for the exotic fermion decays $Q_{5/3} \rightarrow tW^+$ and $Q_{4/3} \rightarrow \bar{b}W^+$. These exotic decays constitute the most promising channels for probing these models in future collider searches. The lightest composite vector bosons range in mass from ~ 1 TeV to around 12 TeV, and neutral bosons on the lighter end of the spectrum can be probed effectively in the $Z_3 \rightarrow t\bar{t}$ decay channel. Finally, the Higgs signal strengths predicted by the M4DCHMs were analysed, and the two minimally-tuned models were found to have a clear signature: they predict the $gg \rightarrow H \rightarrow \gamma\gamma$ cross section to be between $\sim 80\%$ and $\sim 90\%$ of the value predicted by the SM. This is already in tension with experiment at around the 2σ level, and has the potential of being strongly ruled out in the future high-luminosity run of the LHC.

There are many avenues of research that can be taken to extend this work. For example, one might try to extend our models by introducing composite couplings to the third generation leptons, which has been shown to produce interesting fine-

tuning effects [30], or even across all fermions - a case that leads to rich flavour physics. Fitting these models would undoubtedly be challenging, but our results could aid future explorations in targeting the most viable regions of parameter space in the top-bottom sector. It would also be interesting to fit the next-to-minimal 4DCHM, based on the $SO(6) \rightarrow SO(5)$ symmetry breaking pattern, as it offers an extra neutral pNGB that may act as a dark matter candidate.

If the discovery of the Higgs boson is anything to go by, then theoretical hints that new physics exists around a particular energy scale should be taken seriously. And although the current theory has proven to be more robust at the energies probed by the LHC than was perhaps optimistically expected from the hints provided by the Hierarchy Problem, there is still plenty of room to discover evidence of composite Higgs models in the coming years and finally go beyond the Standard Model.

Representation Theory

Here we provide the prerequisite details on representation theory that will be used throughout this work. In a sentence, the point of representation theory is to understand the elements of a group G by representing them as linear transformations on some vector space V . The results we state in this section can be found in any good introductory textbook on representation theory, *e.g.* Ref. [16].

A.1 Groups and Lie Groups

Groups are designed to capture the notion of symmetry, with group elements to be regarded as abstract transformations of some object. For example, the group $SO(2)$ of orientation-preserving isometries of the circle contains the element

[rotate clockwise by 80°].

For a collection G of transformations to be classified as a group, the following three properties must be satisfied:

1. G must contain the [do nothing] transformation e , called the *identity* element.
2. If two transformations g_1, g_2 are in G , then G must also contain the composition g_2g_1 , defined as the result of performing g_1 followed by g_2 . Note this implies the group composition rule is associative.
3. The transformation that undoes $g \in G$, denoted g^{-1} and defined so that $gg^{-1} = g^{-1}g = e$, must also be in G .

Some groups may be endowed with extra structure. **Lie groups**, for example, are those whose elements can be understood as points on a smooth manifold such that, roughly speaking, points that are close together are similar as transformations. Specifically, they are groups with smooth group multiplication and inversion maps

$$(g_1, g_2) \mapsto g_1g_2, \tag{A.1}$$

$$g \mapsto g^{-1}. \tag{A.2}$$

Take $SO(2)$, for example. It can be understood as a circle: if the circle is embedded in the complex plane as

$$S^1 = \{e^{i\theta} \mid \theta \in \mathbb{R}\}, \quad (\text{A.3})$$

each point $e^{i\theta} \in S^1$ can be identified with the transformation

$$g_\theta = [\text{rotate counter-clockwise by } \theta \text{ radians}] \in SO(2). \quad (\text{A.4})$$

Clearly the maps

$$(g_{\theta_1}, g_{\theta_2}) \mapsto g_{\theta_1} g_{\theta_2} = g_{(\theta_1 + \theta_2)}, \quad (\text{A.5})$$

$$g_\theta \mapsto g_\theta^{-1} = g_{(-\theta)}, \quad (\text{A.6})$$

are smooth when $SO(2)$ is imbued with the induced topology from the circle, demonstrating that $SO(2)$ is a Lie group.

More generally, all closed subgroups of $GL(n, \mathbb{R})$ or $GL(n, \mathbb{C})$ (the invertible $n \times n$ real or complex matrices) are Lie groups. Not all Lie groups are of this form, but all of the ones we will be interested in are. This includes the *unitary* groups

$$U(n) = \{U \in GL(n, \mathbb{C}) \mid U^\dagger U = \mathbb{1}\}, \quad (\text{A.7})$$

and the *orthogonal* groups

$$O(n) = \{O \in GL(n, \mathbb{R}) \mid O^\top O = \mathbb{1}\}, \quad (\text{A.8})$$

as well as their “special” subgroups $SU(n)$ and $SO(n)$, containing only those elements with unit determinant. These groups usually arise in physics because their defining properties are such that their elements preserve some sort of norm or length.

A.2 Representations

We are not so much interested in the group elements themselves; what interests us is how objects (specifically quantum fields) are affected when these transformations are performed on them. To this end, we assign an operator S_g to each element g of a group G that enacts the transformation g . How can we tell that S_g actually does what we want it to do? Well, g is entirely defined by its compositions with the other group elements. As long as the operators are assigned so that compositions are preserved,

$$S_{g_2} S_{g_1} = S_{g_2 g_1} \quad \forall g_1, g_2 \in G, \quad (\text{A.9})$$

then S_g can honestly be said to enact the transformation g . A map between group elements and a set of operators that satisfies Equation (A.9),

$$\begin{aligned} \rho : G &\rightarrow \{S_g \mid g \in G\} \\ g &\mapsto \rho(g) = S_g, \end{aligned} \tag{A.10}$$

is called a **representation** of G . Perhaps surprisingly, there is no unique representation of any non-trivial group. Almost always the label “representation” actually refers to only a *linear* representation - that is, to a representation that maps to linear operators on some vector space V ,

$$\rho(g) : V \rightarrow V, \tag{A.11}$$

simply because linear operators are easier to handle. We will adopt this terminology. The **dimension** of a representation is the dimension of V , and can often be used to distinguish the different representations of a given group.

Note that the distinction between g and $\rho(g)$ is usually unimportant if it is understood the representation ρ is being used, and indeed we have already confused the two by defining the unitary and orthogonal groups as sets of matrices rather than as sets of abstract transformations. Those matrices of Equations (A.7) and (A.8) form the so-called **defining representations** of the unitary and orthogonal groups.

Representations must assign $\rho(e) = \mathbb{1}$ for the group structure to be preserved. For a Lie group G such as those we are interested in, since the manifold structure can be parameterised by analytic functions, the group elements infinitesimally close to the identity may be found by first-order Taylor expansion about the identity $e \equiv \mathbb{1}$:

$$\rho(g(\vec{\varepsilon})) = \mathbb{1} + i\varepsilon^a T_a^\rho \tag{A.12}$$

in terms of infinitesimal coordinates ε^a , equal in number to the dimension of the group, and some linear operators T_a^ρ , known as **generators** of G for the representation ρ . This is simply the infinitesimal version of the relation

$$\rho(g(\vec{\theta})) = \exp[i\theta^a T_a^\rho] = \sum_{n=0}^{\infty} \frac{(i\theta^a T_a^\rho)^n}{n!} \tag{A.13}$$

for group elements a finite distance from the identity, where $\vec{\theta}$ is a collection of finite coordinates. Not all group elements can be written as the exponential of a matrix, however. In general, only a connected subgroup of G can be expressed in this way. But if G is compact and connected, such as $SU(n)$ or $SO(n)$, then every group element *is* the exponential of some linear combination of the generators.

Example: Representations of $U(1)$

We have already come close to representing the elements of a Lie group through exponentials, in Equation (A.4). That identification between elements of $SO(2)$

and points on the complex unit circle is actually an isomorphism between $SO(2)$ and the group $U(1)$ of rotations of the complex plane about the origin, since such rotations are enacted by multiplication by phases $e^{i\theta}$. An appropriate representation of $U(1)$ therefore uses the 1×1 complex matrices

$$\rho(\text{counter-clockwise rotation of } \theta \text{ radians}) = [e^{i\theta}] = \exp(i\theta[1]), \quad (\text{A.14})$$

the sole generator here being the matrix $[1]$. A different representation could be constructed by taking the generator to be, say, $[-3]$, or any other 1×1 matrix with integer entry¹. These representations would respectively be labelled $+1$ or -3 .

In QFT, the different representations of fields under a $U(1)$ symmetry are accounted for by writing all of the representations as

$$\rho(\text{counter-clockwise rotation of } \theta \text{ radians}) = \exp[i\theta X], \quad (\text{A.15})$$

where X is an operator that yields the appropriate generator when acting on each field. In this case the group would be labelled $U(1)_X$ to signify that fields furnish representations according to their X charges.

These basic representations of $U(1)$ are used extensively throughout this work. Below we will detail other types of representations that will also be used.

Reducible and Irreducible Representations

First we discuss how certain representations may be understood in terms of others. Suppose we have a representation ρ of a group G , acting on a vector space V . Sometimes it is the case that there is a basis of V in which ρ can be expressed in block diagonal form:

$$\rho(g) := \begin{pmatrix} \rho_1(g) & 0 \\ 0 & \rho_2(g) \end{pmatrix}, \quad \forall g \in G. \quad (\text{A.16})$$

From this form it follows that the functions ρ_1 and ρ_2 are also representations of G , acting on complementary subspaces of V . A representation ρ for which such a basis exists is called **reducible**, and is denoted by the **direct sum**

$$\rho = \rho_1 \oplus \rho_2. \quad (\text{A.17})$$

The representations ρ_1 and/or ρ_2 may themselves be reducible, in which case ρ decomposes as the direct sum of more representations, having further block diagonal structure. Representations that cannot be written as a direct sum are **irreducible**. For compact groups, every reducible representation is the direct sum of some number

¹The generator must be an integer because the circle identifies $\theta \sim \theta + 2\pi$, and this must be reflected in the operators assigned by the representation. For *projective* representations that are used in QFT, where operators need only preserve the group structure up to a phase, *any* real number can be used as a generator of $U(1)$.

of irreducible representations. Typically in the context of QFT the subspaces on which the irreducible constituent representations of a reducible representation act are interpreted as independent multiplets of quantum fields, so to describe a single quantum field multiplet one would only need to use an irreducible representation.

This reasoning can be applied in the other direction, of course, forming a new representation through the direct sum of two known representations through the definition Equation (A.16), although we shall not do so explicitly in this work.

Trivial Representation

All groups have at least one representation, defined by

$$\rho(g) = \mathbb{1} \quad \forall g \in G. \tag{A.18}$$

It is trivial, of course, because nothing changes under its action. The generator of this representation is simply the zero operator, for $\exp[i\theta \cdot 0] = \mathbb{1}$. This representation is denoted by $\mathbf{1}$ to evoke imagery of the identity operator.

Adjoint Representations

In addition to the trivial representation, every Lie group G also has an **adjoint** representation \mathbf{Ad}_G . Its generators are defined using the structure constants of the Lie algebra of G (see Equation (2.8)) as

$$(T_a^{\mathbf{Ad}})_b^c := if_{ab}^c. \tag{A.19}$$

Note the dimension of the adjoint representation of a group is equal to the dimension of the group itself. For $SU(n)$, the adjoint is $\mathbf{n}^2 - \mathbf{1}$, while for $SO(n)$ it is $\mathbf{n}(\mathbf{n} - \mathbf{1})/2$.

The definition Equation (A.19) actually allows the adjoint transformation to be given in terms of the generators of any other representation. Indeed, consider a vector A^a transforming in the adjoint representation. It must have $\dim(G)$ -many components, and so we may form the object

$$A^\rho := A^a T_a^\rho \tag{A.20}$$

for generators T_a^ρ of any representation ρ . (The superscript ρ is a label, not an index). Under an infinitesimal adjoint transformation, then,

$$\begin{aligned}
A^\rho &\rightarrow \left(e^{i\varepsilon^a T_a^{\mathbf{Ad}}} A \right)^c T_c^\rho \\
&= A^\rho + (i\varepsilon^a (T_a^{\mathbf{Ad}})_b^c A^b) T_c^\rho \\
&= A^\rho + i\varepsilon^a i f_{ab}^c A^b T_c^\rho \\
&= A^\rho + i\varepsilon^a A^b [T_a^\rho, T_b^\rho] \\
&= A^\rho + i\varepsilon^a [T_a^\rho, A^\rho] \\
&= \exp [i\varepsilon^a T_a^\rho] A^\rho \exp [-i\varepsilon^a T_a^\rho],
\end{aligned} \tag{A.21}$$

where we have made extensive use of Equation (2.8). This extends to finite transformations as

$$A^\rho \xrightarrow{g \in G} \rho(g) A^\rho \rho(g)^{-1}. \tag{A.22}$$

Anything transforming with the same pattern as Equation (A.22) is said to transform adjointly. This property proves useful because if a vector $\vec{\varphi}$ transforms under representation ρ , then so too does $A^\rho \vec{\varphi}$:

$$A^\rho \vec{\varphi} \rightarrow [\rho(g) A^\rho \rho(g)^{-1}] [\rho(g) \vec{\varphi}] = \rho(g) A^\rho \vec{\varphi}, \tag{A.23}$$

making it straightforward to construct invariant interactions.

Fundamental Representations

As we have seen with the unitary and orthogonal groups, some groups are explicitly defined in terms of particular matrices. Such defining representations are often called **fundamental representations** in the physics literature, although it should be noted that “fundamental representation” has a technical meaning that coincidentally overlaps with the defining representation in certain cases. For unitary and orthogonal groups, the defining representations are indeed fundamental representations.

The defining representations of $U(n)$ and $O(n)$ act on n -dimensional vectors $\vec{\varphi}$:

$$\varphi_i \xrightarrow{g \in G} \varphi'_i = g_i^j \varphi_j, \tag{A.24}$$

and are accordingly labelled as \mathbf{n} . Here we do not distinguish between the group element g and the operator $\rho(g)$, because of course the group elements are defined by this representation.

What is special about these defining representations is that they are the lowest-dimensional representations that assign group elements to operators injectively. For semisimple Lie groups such as $SU(n)$ and $SO(n)$, it turns out that this allows

all other finite-dimensional irreducible representations to be constructed from the defining representation.

Symmetric and Antisymmetric Representations

In this section we focus specifically on representations of $SO(n)$, which is all we will be using in this work. To show how other representations can be constructed from a vector $\vec{\varphi}$ in the fundamental representation \mathbf{n} , we follow Ref. [52] and consider the object defined by

$$T_{ij} := \varphi_i \varphi_j. \quad (\text{A.25})$$

What is important about this object is not its definition; we are only interested in its transformation properties. Based on Equation (A.24), T will transform as

$$T_{ij} \xrightarrow{g \in G} \varphi'_i \varphi'_j = g_i^k \varphi_k g_j^m \varphi_m = g_i^k g_j^m T_{km}, \quad (\text{A.26})$$

or in matrix notation,

$$T \xrightarrow{g \in G} T' = g T g^\top. \quad (\text{A.27})$$

This type of transformation forms a representation of $SO(n)$, for it preserves the group structure². There are n^2 components of T , so this representation is labelled \mathbf{n}^2 .

But notice this is a reducible representation: it does not mix the symmetric and antisymmetric subspaces together. That is, the antisymmetric and symmetric parts of T ,

$$A := \frac{1}{2}(T - T^\top), \quad S := \frac{1}{2}(T + T^\top), \quad (\text{A.28})$$

remain respectively antisymmetric and symmetric under Equation (A.27)

$$A \xrightarrow{g \in G} A' = g A g^\top, \quad S \xrightarrow{g \in G} S' = g S g^\top. \quad (\text{A.29})$$

In addition, the transformation does not mix the trace of T with anything:

$$\text{Tr}[T] \xrightarrow{g \in G} \text{Tr}[T'] = \text{Tr}[g T g^\top] = \text{Tr}[T], \quad (\text{A.30})$$

using the defining property of $g \in SO(n)$ in Equation (A.8). We might as well then redefine S to be traceless (A is automatically traceless) and include a pure trace component in the decomposition of T ,

$$T = S + A + \frac{1}{n} \text{Tr}[T] \mathbf{1}. \quad (\text{A.31})$$

²Performing transformation g_1 followed by g_2 yields $g_2 g_1 T g_1^\top g_2^\top = (g_2 g_1) T (g_2 g_1)^\top \checkmark$.

Schematically, we have decomposed the \mathbf{n}^2 representation into three irreducible representations: the traceless symmetric, the antisymmetric, and the pure trace singlet

$$\underbrace{T}_{\mathbf{n}^2} = \underbrace{S}_{\frac{\mathbf{n}(\mathbf{n}+1)}{2}-1} \oplus \underbrace{A}_{\frac{\mathbf{n}(\mathbf{n}-1)}{2}} \oplus \underbrace{\text{Tr}[T]}_{1}. \quad (\text{A.32})$$

The dimensions of the irreducible representations can be understood by simply counting the respective independent components in Equation (A.28) and accounting for the one degree of freedom in S that was removed by making it traceless.

Note that even though these irreducible representations were found by considering the tensor product of two fundamental representations, multiplets in these representations do not have to be constructed this way; any antisymmetric or traceless symmetric matrix transforming as in Equation (A.29) will do.

Charges and Generators of a Symmetry

It will prove useful to investigate the connection between the symmetry group G and the conserved charges Q_a that it generates. Ultimately we shall show the charge associated to a generator T_a is the generator itself. We proceed with an argument largely based on Ref. [19].

Suppose a multiplet of canonical fields $\vec{\phi} = (\phi^1, \phi^2, \dots)$ transforms in a representation of G that has generators T_a . Then an infinitesimal symmetry $S \in G$ is of the form

$$S = \mathbb{1} + i\varepsilon^a T_a, \tag{B.1}$$

and the fields transform as

$$\phi^i \mapsto \phi^i + i\varepsilon^a (T_a)^i_j \phi^j. \tag{B.2}$$

$$\tag{B.3}$$

This holds for arbitrary small ε^a , so it must be that

$$\frac{\delta\phi^i}{\delta\varepsilon^a} = i (T_a)^i_j \phi^j. \tag{B.4}$$

$$\tag{B.5}$$

Now to look at the conserved charges. From Nöther's Theorem, the conserved currents are given by

$$J_a^\mu = -\frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi^i)} \frac{\delta\phi^i}{\delta\varepsilon^a}, \tag{B.6}$$

giving conserved charges

$$Q_a = \int d^3x J_a^0(\vec{x}) = -i \int d^3x \pi_i(\vec{x}) (T_a)^i_j \phi^j, \tag{B.7}$$

where $\pi_i = \frac{\partial \mathcal{L}}{\partial \dot{\phi}^i}$ is the momentum density conjugate to ϕ^i . But now using the canonical commutation relations

$$[\phi^i(\vec{x}), \phi^k(\vec{y})] = 0, \quad (\text{B.8})$$

$$[\pi_i(\vec{x}), \phi^k(\vec{y})] = -i\delta_i^k \delta^3(\vec{x} - \vec{y}), \quad (\text{B.9})$$

we calculate

$$[Q_a, \phi^k(\vec{y})] = -i \int d^3x [\pi_i(\vec{x}), \phi^k(\vec{y})] (T_a)^i_j \phi^j(\vec{x}) = - (T_a)^k_j \phi^j(\vec{y}), \quad (\text{B.10})$$

$$[Q_a, \pi_k(\vec{y})] = -i \int d^3x \pi_i(\vec{x}) (T_a)^i_j [\phi^j(\vec{x}), \pi_k(\vec{y})] = (T_a)^i_k \pi_i(\vec{y}). \quad (\text{B.11})$$

$$(\text{B.12})$$

With these, we get

$$[Q_a, Q_b] = -i \int d^3x \pi_i(\vec{x}) [T_a, T_b]^i_j \phi^j(\vec{x}) = i f_{ab}{}^c Q_c. \quad (\text{B.13})$$

Thus, the charges Q_a form the same Lie algebra as the generators T_a of the symmetry, and so we may choose $T_a = Q_a$. In other words, the charge generated by a continuous symmetry is itself the generator of that symmetry.

Quantum Effective Action

In classical field theory, field dynamics are completely specified by the action

$$S[\phi, \partial\phi, \dots] = \int d^4x \mathcal{L}[\phi(x), \partial\phi(x), \dots], \quad (\text{C.1})$$

with the equations of motion given by the principle of stationary action $\delta S = 0$. This is not true in quantum field theory because of the non-commutative operator nature of the fields. Instead, the Schrödinger-like time evolution equation needs to be solved through the use of a Dyson series, whose various terms can be interpreted pictorially as Feynman diagrams of interacting off-shell particles. Internal loops in these Feynman diagrams lead to the difference between the classical and quantum regimes.

What is remarkable, though, is that the effects of loop diagrams can automatically be taken into account if one uses instead the *quantum effective action*, Γ , in place of the regular action. The principle of stationary *effective* action then gives the equations of motion for the vacuum expectation values of the quantum fields. Here we summarise the treatment of the effective potential that can be found in any introductory textbook on quantum field theory, *e.g.* Ref. [52, 53, 143].

The quantum effective action is perhaps most easily arrived at through the path integral formulation of quantum field theory. In classical field theory, the vacuum expectation value of any operator \mathcal{O} is given by

$$\langle \mathcal{O} \rangle = \frac{\int \mathcal{D}\Phi \mathcal{O} e^{\frac{i}{\hbar} S[\Phi]}}{\int \mathcal{D}\Phi e^{\frac{i}{\hbar} S[\Phi]}}, \quad (\text{C.2})$$

where Φ stands for all fields in the theory and $S[\Phi]$ is the action. The situation is subtly different in quantum field theory due to loop diagrams. To account for this, an operator's vev is given by its expectation value in the presence of classical source fields (complex-valued functions), in the limit that the source fields vanish. To make the notation clear, quantum fields will be hatted and classical fields unhatted. What happens is that for every quantum field $\hat{\phi}^i$ in the theory (not necessarily elementary), a classical source field J_i is introduced, effecting a change in the action

$$S[\hat{\Phi}] \rightarrow S_J[\hat{\Phi}] := S[\hat{\Phi}] + \int d^4x J(x) \cdot \hat{\Phi}(x). \quad (\text{C.3})$$

Here we have introduced the notation $\hat{\Phi} = (\hat{\phi}^1, \hat{\phi}^2, \dots)$, $J = (J_1, J_2, \dots)$, and the dot denotes the naïve dot product. The source fields have the effect of exciting their associated field from the vacuum. The expectation value Φ_J of $\hat{\Phi}$ in the presence of currents J is then

$$\Phi_J(x) := \langle J | \hat{\Phi}(x) | J \rangle = \frac{\int \mathcal{D}\hat{\Phi} \hat{\Phi}(x) e^{\frac{i}{\hbar} S_J[\hat{\Phi}]}}{\int \mathcal{D}\hat{\Phi} e^{\frac{i}{\hbar} S_J[\hat{\Phi}]}} \quad (\text{C.4})$$

while the vevs Φ_0 are found by

$$\Phi_0(x) := \lim_{J \rightarrow 0} \Phi_J(x). \quad (\text{C.5})$$

Now notice that if we define the so-called Schwinger functional

$$W[J] := \frac{\hbar}{i} \ln \left(\int \mathcal{D}\hat{\Phi} e^{\frac{i}{\hbar} S_J[\hat{\Phi}]} \right), \quad (\text{C.6})$$

we have the formal relation

$$\phi_J^i(x) = \frac{\delta W[J]}{\delta J_i(x)}, \text{ or more compactly, } \Phi_J(x) = \frac{\delta W[J]}{\delta J(x)}. \quad (\text{C.7})$$

In many physical cases this relation is invertible, allowing the current J to be found in terms of Φ_J . To make this functional dependence more transparent, let J^Φ denote this set of currents for a given classical field configuration Φ . That is, J^Φ is the current that satisfies

$$\Phi(x) = \frac{\delta W[J^\Phi]}{\delta J^\Phi(x)}. \quad (\text{C.8})$$

These notions now allow the quantum effective action Γ to be defined as the Legendre transformation of W :

$$\Gamma[\Phi] := W[J^\Phi] - \int d^4x J^\Phi(x) \cdot \Phi(x). \quad (\text{C.9})$$

By virtue of Equation (C.8), this satisfies

$$\frac{\delta \Gamma[\Phi]}{\delta \phi^i(x)} = \int d^4y \frac{\delta W[J^\Phi]}{\delta J_j^\Phi(y)} \frac{\delta J_j^\Phi(y)}{\delta \phi^i(x)} - J_i^\Phi(x) - \int d^4y \frac{\delta J_j^\Phi(y)}{\delta \phi^i(x)} \phi^j(x) \quad (\text{C.10})$$

$$= -J_i^\Phi(x). \quad (\text{C.11})$$

Taking the limit as $J \rightarrow 0$, the classical fields Φ take on their quantum vacuum expectation values Φ_0 and we are left with

$$\frac{\delta \Gamma[\Phi]}{\delta \phi^i(x)} \Big|_{\Phi=\Phi_0} = 0. \quad (\text{C.12})$$

In other words, the effective action is stationary when the fields all achieve their vev. This condition becomes clearer when the effective potential is expanded in powers of momentum (or derivatives in position space):

$$\Gamma[\Phi] = \int d^4x \left(-V_{\text{eff}}(\Phi) + \frac{1}{2}(\partial_\mu \Phi(x))(\partial^\mu \Phi(x))Z(\Phi) + \dots \right) \quad (\text{C.13})$$

for some functions V_{eff} and Z . The first term, V_{eff} , is named the quantum effective potential by analogy with the classical potential being the negative density of the non-derivative terms in the classical action. When the fields all achieve their vevs, Φ_0 will be constant throughout spacetime on account of the Poincaré invariance of the vacuum, so the effective action becomes simply

$$\Gamma[\Phi_0] = -VT V_{\text{eff}}(\Phi_0), \quad (\text{C.14})$$

where VT is the volume of spacetime¹. Restricted to constant field configurations, the vev condition Equation (C.12) then reduces to a simple derivative relation

$$\frac{\partial V_{\text{eff}}}{\partial \phi^i(x)} \Big|_{\Phi=\Phi_0} = 0. \quad (\text{C.15})$$

For the vevs to be stable under small perturbations, Equation (C.15) must represent a minimum of the effective potential rather than any stationary point. This fact - that the possible vevs are given by the minima of the effective potential - is used extensively throughout this work.

C.1 Interpretation of the Effective Potential

Of course, it is no good to know this fact without having any way to calculate the effective potential. The key to its calculation lies in a physical interpretation of the Schwinger functional of Equation (C.6). Notice that the Schwinger functional, when exponentiated, gives the amplitude of transitioning from the vacuum in the far past to the vacuum in the far future in the presence of sources J :

$$\exp\left(\frac{i}{\hbar}W[J]\right) = \int \mathcal{D}\hat{\Phi} e^{\frac{i}{\hbar}S_J[\hat{\Phi}]} = \langle \Omega(t = \infty) | \Omega(t = -\infty) \rangle_J. \quad (\text{C.16})$$

But this amplitude is simply the sum of all (connected and disconnected) vacuum \rightarrow vacuum Feynman diagrams, and it is an elementary result in quantum field theory that this sum is the exponential of the sum of only the connected vacuum \rightarrow vacuum diagrams. So $iW[J]/\hbar$ is the sum of all connected vacuum diagrams, any of which has the form of a tree diagram (one with no loops), with vertices that are one-particle-irreducible (1PI) diagrams (ones that cannot be disconnected with the removal of a line).

¹This infinite factor will cancel out in calculations.

The trick now is to consider what happens when the action S is replaced by a new function \tilde{S} whose vertices contain all 1PI diagrams. By the same reasoning above, the new Schwinger functional \tilde{W} will be the sum of all connected diagrams whose vertices are 1PI diagrams. So the original Schwinger functional, having only tree-level diagrams, can equivalently be defined as the classical limit of \tilde{W} :

$$W[J] = \lim_{\hbar \rightarrow 0} \tilde{W}[J] = \lim_{\hbar \rightarrow 0} \frac{\hbar}{i} \ln \left(\int \mathcal{D}\hat{\Phi} e^{\frac{i}{\hbar} \tilde{S}_J[\hat{\Phi}]} \right). \quad (\text{C.17})$$

The $\hbar \rightarrow 0$ limit picks out the configuration that extremises \tilde{S}_J , giving the condition

$$\frac{\delta \tilde{S}[\Phi_J]}{\delta \Phi_J} = -J \quad (\text{C.18})$$

and reducing Equation (C.17) to

$$W[J] = \tilde{S}[\Phi_J] + \int d^4x J(x) \cdot \Phi_J(x). \quad (\text{C.19})$$

But this is exactly the inverse Legendre transformation of Equation (C.9), with \tilde{S} being the quantum effective potential! We now see that Γ generates the sum of all 1PI diagrams, and so may be written as

$$\Gamma[\Phi] = \sum_{N=1}^{\infty} \frac{1}{N!} \int d^4x_1 \cdots d^4x_N \Gamma^{(N)}(x_1, \dots, x_N) \Phi(x_1) \cdots \Phi(x_N), \quad (\text{C.20})$$

where $\Gamma^{(N)}$ is the sum of all 1PI diagrams with N external lines. Here we are treating Φ as a single field for notational convenience, but it should be understood that for multiple fields all of the different possible combinations of N fields should be included. To make the connection with the effective potential, we now go to momentum space (following Ref. [83])

$$\Gamma^{(N)}(x_1, \dots, x_N) = \int \frac{d^4p_1}{(2\pi)^4} \cdots \frac{d^4p_N}{(2\pi)^4} e^{i \sum_i p_i \cdot x_i} \Gamma^{(N)}(p_1, \dots, p_N) (2\pi)^4 \delta^4\left(\sum_i p_i\right), \quad (\text{C.21})$$

resulting in

$$\Gamma[\Phi] = \sum_{N=1}^{\infty} \frac{1}{N!} \int d^4x \Gamma^{(N)}(p_1, \dots, p_N) (\Phi(x))^N. \quad (\text{C.22})$$

Expanding $\Gamma^{(N)}(p_1, \dots, p_N)$ around zero momentum and comparing with Equation (C.13), we see that

$$V_{\text{eff}}(\Phi) = - \sum_{N=1}^{\infty} \frac{1}{N!} \Gamma^{(N)}(0, \dots, 0) (\Phi(x))^N. \quad (\text{C.23})$$

In other words, the effective potential is the (negative) sum of all 1PI diagrams where the external legs have vanishing momenta. This fact makes V_{eff} extraordinarily useful: for one, if loop diagrams are ignored then the effective potential is the negative sum of the non-derivative terms in the Lagrangian - that is to say, it is the classical potential. Hence, V_{eff} is just the classical potential with quantum corrections. Furthermore, if the squared mass of a field is defined as the inverse of its propagator at zero momentum, as it is for scalar fields, then the (renormalised!) mass matrix of the fields is simply

$$(M^2)_{ab} = \frac{\partial^2 V_{\text{eff}}}{\partial \phi^b \partial \phi^a} \Big|_{\Phi=\Phi_0}. \quad (\text{C.24})$$

Similar considerations hold for other coupling constants, but this is the only one needed for this work.

Isomorphism Between $SO(4)$ and $SU(2) \times SU(2)$

A fact that is exploited throughout this work is that $SO(4)$ is locally isomorphic to $SU(2)_L \times SU(2)_R$, meaning the two groups share the same Lie algebra. This follows because there is a certain way to embed the four components of a 2×2 matrix Σ (of a certain class) into a vector $\vec{\varphi}$ such that an infinitesimal $SU(2)_L \times SU(2)_R$ transformation of Σ manifests as an infinitesimal $SO(4)$ transformation of $\vec{\varphi}$.

Specifically, Σ is to be taken as a pseudo-real matrix, *i.e.* as a matrix satisfying

$$\Sigma^* = \sigma^2 \Sigma \sigma^2, \tag{D.1}$$

where σ^2 is the second Pauli sigma matrix. The set of all pseudo-real matrices furnishes a representation of $SU(2)_L \times SU(2)_R$, transforming under the group as

$$\Sigma \rightarrow g_L \Sigma g_R^\dagger, \tag{D.2}$$

on account of the properties of the Pauli matrices. This is the **bidoublet** representation, denoted $(\mathbf{2}, \mathbf{2})$.

The condition Equation (D.1) allows Σ to be written in terms of four real numbers φ^i as

$$\Sigma = \frac{1}{\sqrt{2}} \begin{pmatrix} \varphi^4 + i\varphi^3 & \varphi^2 + i\varphi^1 \\ -\varphi^2 + i\varphi^1 & \varphi^4 - i\varphi^3 \end{pmatrix}. \tag{D.3}$$

Now notice that the chiral transformation Equation (D.2) leaves invariant

$$\text{Tr} [\Sigma^\dagger \Sigma] = |\vec{\varphi}|^2, \tag{D.4}$$

and so acts as an $SO(4)$ transformation on $\vec{\varphi} = (\varphi^1, \varphi^2, \varphi^3, \varphi^4)$. Hence the algebra of $SU(2)_L \times SU(2)_R$ is contained in the algebra of $SO(4)$, and since the algebras have equal dimensions they must be the same. Note this local isomorphism does not mean $SU(2)_L \times SU(2)_R$ and $SO(4)$ are the same groups; in fact there is a 2 : 1 correspondence between elements of $SU(2)_L \times SU(2)_R$ and elements of $SO(4)$.

Equation (D.3) allows Σ to be expressed in terms of the $SU(2)_L$ doublet

$$\Phi := \frac{1}{\sqrt{2}} \begin{pmatrix} \varphi^2 + i\varphi^1 \\ \varphi^4 - i\varphi^3 \end{pmatrix} \quad (\text{D.5})$$

as

$$\Sigma = (\Phi^c, \Phi), \quad (\text{D.6})$$

using the conjugate doublet $\Phi^c \equiv i\sigma_2\Phi^*$. The notation here is no coincidence. Φ acts entirely like the SM complex Higgs doublet. If Φ is assigned a hypercharge of $+1/2$, then a $U(1)_Y$ transformation $\Phi \rightarrow e^{i\frac{\alpha}{2}}\Phi$ manifests as an $SU(2)_R$ transformation $\Sigma \rightarrow \Sigma e^{-i\alpha T_R^3}$, where $T_R^3 = \sigma^3$ is the third generator of $SU(2)_R$. For this reason T_R^3 is identified as the hypercharge operator. From Equation (D.5) it can be seen how the components of the Higgs doublet map to a vector under $SO(4)$.

Custodial Symmetry

In Section 2.4, it was shown that EWSB in the SM predicts the W^\pm and Z boson masses are related at tree level by

$$\frac{m_W^2}{m_Z^2} = \frac{g^2}{g^2 + g'^2}, \quad (\text{E.1})$$

where g and g' are the coupling strengths of the $SU(2)_L$ and $U(1)_Y$ gauge bosons. This relation has been experimentally verified to better than 1% accuracy, so any other proposed model of EWSB would need to reproduce this prediction to be phenomenologically viable.

A nice argument is presented in Ref. [24] showing that this relation is always obtained in theories where

1. The photon is massless and the electric charge operator is given by $Q = T_L^3 + Y$, and
2. When $g' = 0$, there is a global $SU(2)$ symmetry under which the $SU(2)_L$ gauge bosons transform as a triplet. The global $SU(2)$ symmetry is known as a “custodial” symmetry on account of its role in protecting the relation Equation (E.1).

The proof is fairly straightforward.

The custodial symmetry imposes equal mass terms on all the $SU(2)_L$ bosons, meaning the electroweak mass matrix must take the form

$$m_{\text{gauge}}^2 = \left(\begin{array}{c|cccc} & W_\mu^1 & W_\mu^2 & W_\mu^3 & B_\mu \\ \hline W^{1\mu} & g^2 u^2 & 0 & 0 & \cdot \\ W^{2\mu} & 0 & g^2 u^2 & 0 & \cdot \\ W^{3\mu} & 0 & 0 & g^2 u^2 & \cdot \\ B^\mu & \cdot & \cdot & \cdot & \cdot \end{array} \right), \quad (\text{E.2})$$

for some parameter u . Note this ostensibly does not agree with the mass matrix in Section 2.4, but the off-diagonal terms in that matrix give interactions that cancel each other out and so can equivalently be set to zero, resulting in a matrix adhering to this form.

That the photon is massless is another way of stating

$$Q\Phi_0 = 0, \quad \text{or} \quad T_L^3\Phi_0 = -Y\Phi_0, \quad (\text{E.3})$$

which leads to the conditions on the mass matrix (see Equation (2.21))

$$g'm_{a3}^2 = -gm_{a4}^2, \quad (\text{E.4})$$

forcing

$$m_{\text{gauge}}^2 = \left(\begin{array}{c|cccc} & W_\mu^1 & W_\mu^2 & W_\mu^3 & B_\mu \\ \hline W^{1\mu} & g^2u^2 & 0 & 0 & 0 \\ W^{2\mu} & 0 & g^2u^2 & 0 & 0 \\ W^{3\mu} & 0 & 0 & g^2u^2 & -g'gu^2 \\ B^\mu & 0 & 0 & -g'gu^2 & g'^2u^2 \end{array} \right). \quad (\text{E.5})$$

This is proportional to the SM electroweak mass matrix, and so reproduces the ratio Equation (E.1).

E.1 Custodial Symmetry in the SM

The reason Equation (E.1) holds in the SM is because the vector subgroup $SU(2)_{L+R}$ of $SU(2)_L \times SU(2)_R$ acts as a custodial symmetry. $SU(2)_{L+R}$ is a good approximate symmetry of the SM, being exact in the limit of equal u and d quark masses in each generation. Technically on account of the approximate nature of $SU(2)_{L+R}$, Equation (E.1) does receive small corrections in the SM but these are negligible for our purposes¹.

Chiral symmetry is exact in the Higgs potential terms, for the potential is not just invariant under $SU(2)_L \times U(1)_Y$ but has a larger accidental $SO(4)$ symmetry:

$$\mathcal{V}(\Phi) = \mathcal{V}(\Phi^\dagger\Phi) = \mathcal{V}(\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2). \quad (\text{E.6})$$

Recalling Appendix D, the potential therefore has a chiral $SU(2)_L \times SU(2)_R$ symmetry that is readily apparent when expressed in terms of the bidoublet $\Sigma = (\Phi^c, \Phi)$ as

$$\mathcal{V}(\Sigma) = -\mu^2\text{Tr} [\Sigma^\dagger\Sigma] + \lambda\text{Tr} [\Sigma^\dagger\Sigma]^2. \quad (\text{E.7})$$

In this form, the Higgs kinetic term is

$$\frac{v^2}{4}\text{Tr} [(D_\mu\Sigma)^\dagger(D^\mu\Sigma)]. \quad (\text{E.8})$$

¹Since the t quark is much heavier than the others, Equation (E.1) receives a relative correction of approximately $\frac{3G_F m_t^2}{8\sqrt{2}\pi^2} \approx 1\%$ [144].

This does not exhibit chiral symmetry since the hypercharge generator is gauged, explicitly breaking the $SU(2)_R$ symmetry (recall from Appendix D the hypercharge is the third generator of $SU(2)_R$). But the limit $g' \rightarrow 0$ decouples the hypercharge generator, restoring chiral symmetry and therefore also $SU(2)_{L+R}$ symmetry.

The quark-Higgs sector also displays this symmetry in the limit of quarks in each generation having equal masses. In this limit, the Yukawa interactions of the i^{th} generation are given by

$$\mathcal{L}_{\text{Yukawa}} = -y_i \bar{q}_L^i \Sigma q_R^i + h.c. \quad (\text{E.9})$$

where the right-handed doublets $q_R^i = (u_R^i, d_R^i)^\top$ transforming as a fundamental of $SU(2)_R$ have been introduced. Such terms are obviously chirally invariant due to the Higgs transformation property Equation (D.2). But of course the quarks, having unequal masses, must have different Yukawa couplings, making this symmetry only approximate.

A global $SU(2)_{L+R}$ transformation enacts a global $SU(2)_L$ transformation, under which the $SU(2)_L$ gauge bosons transform as a triplet on account of Equation (2.13), fulfilling the condition for a custodial symmetry.

Computation of the Nambu-Goldstone Boson Matrix

Here we include the computation of the NGB matrix

$$U[\vec{\theta}] = e^{i\frac{\sqrt{2}}{f}\theta^a \hat{T}_a} \quad (\text{F.1})$$

for broken generators \hat{T}^a of $SO(n) \rightarrow SO(n-1)$. We use the convenient basis given in Appendix G (generalised to n dimensions), consistent with the choice of vacuum vector $\Phi_0 = (0, 0, \dots, 0, f)^\top$, in which

$$(\hat{T}_a)_{ij} = -\frac{i}{\sqrt{2}} (\delta_i^a \delta_j^n - \delta_j^a \delta_i^n), \quad a = 1, \dots, n-1. \quad (\text{F.2})$$

That is, we are calculating the exponential of

$$M := i\frac{\sqrt{2}}{f}\theta^a \hat{T}_a = \left(\begin{array}{c|c} 0 & \vec{v} \\ \hline -\vec{v}^\top & 0 \end{array} \right), \quad (\text{F.3})$$

where $\vec{v} := (\theta^1, \dots, \theta^{n-1})^\top / f$. To begin, note

$$M^2 = (-1) \left(\begin{array}{c|c} \vec{v}\vec{v}^\top & 0 \\ \hline 0 & \vec{v}^\top\vec{v} \end{array} \right), \quad (\text{F.4})$$

$$\implies M^{2k} = (-1)^k \left(\begin{array}{c|c} (\vec{v}\vec{v}^\top)^k & 0 \\ \hline 0 & (\vec{v}^\top\vec{v})^k \end{array} \right), \quad (\text{F.5})$$

$$\implies M^{2k+1} = (-1)^k \left(\begin{array}{c|c} 0 & (\vec{v}\vec{v}^\top)^k \vec{v} \\ \hline -(\vec{v}^\top\vec{v})^k \vec{v}^\top & 0 \end{array} \right). \quad (\text{F.6})$$

Using the fact that $(\vec{v}\vec{v}^\top)^k = \vec{v}(\vec{v}^\top\vec{v})^{k-1}\vec{v}^\top$ for $k \geq 1$, and $\vec{v}^\top\vec{v} = |\vec{v}|^2$, these simplify to

$$M^{2k} = (-1)^k |\vec{v}|^{2k} \left(\begin{array}{c|c} \frac{\vec{v}\vec{v}^\top}{|\vec{v}|} & 0 \\ \hline 0 & 1 \end{array} \right) =: (-1)^k |\vec{v}|^{2k} M_0 \quad (\text{for } k \geq 1), \quad (\text{F.7})$$

$$M^{2k+1} = (-1)^k |\vec{v}|^{2k+1} \left(\begin{array}{c|c} 0 & \frac{\vec{v}}{|\vec{v}|} \\ \hline -\frac{\vec{v}^\top}{|\vec{v}|} & 0 \end{array} \right) =: (-1)^k |\vec{v}|^{2k+1} M_1. \quad (\text{F.8})$$

Hence, the exponential is

$$\begin{aligned} e^M &= \sum_{k=0}^{\infty} \frac{1}{(2k)!} M^{2k} + \sum_{k=0}^{\infty} \frac{1}{(2k+1)!} M^{2k+1} \\ &= M^0 + \sum_{k=1}^{\infty} \frac{(-1)^k}{(2k)!} |\vec{v}|^{2k} M_0 + \sum_{k=0}^{\infty} \frac{(-1)^k}{(2k+1)!} |\vec{v}|^{2k+1} M_1 \\ &= \mathbb{1} + (\cos |\vec{v}| - 1) M_0 + \sin |\vec{v}| M_1 \\ &= \begin{pmatrix} \mathbb{1} - \frac{\vec{v}\vec{v}^\top}{|\vec{v}|^2} (1 - \cos |\vec{v}|) & \frac{\vec{v}}{|\vec{v}|} \sin |\vec{v}| \\ -\frac{\vec{v}^\top}{|\vec{v}|} \sin |\vec{v}| & \cos |\vec{v}| \end{pmatrix}. \end{aligned} \quad (\text{F.9})$$

Substituting $\vec{v} = \vec{\theta}/f$ yields

$$U[\vec{\theta}] = \begin{pmatrix} \mathbb{1} - \frac{\vec{\theta}\vec{\theta}^\top}{|\vec{\theta}|^2} \left(1 - \cos \frac{|\vec{\theta}|}{f}\right) & \frac{\vec{\theta}}{|\vec{\theta}|} \sin \frac{|\vec{\theta}|}{f} \\ -\frac{\vec{\theta}^\top}{|\vec{\theta}|} \sin \frac{|\vec{\theta}|}{f} & \cos \frac{|\vec{\theta}|}{f} \end{pmatrix}. \quad (\text{F.10})$$

$SO(5)$ Generators and Multiplets

Here we give the conventions for $SO(5)$ generators and embeddings we use in this work. This section makes extensive use of Appendix D, since the local isomorphism $SO(4) \cong SU(2)_L \times SU(2)_R$ allows representations of $SO(5) \cong SO(4) \times SO(5)/SO(4)$ to be decomposed into those of $SU(2)_L \times SU(2)_R$, making the embedded fields' quantum numbers easily identifiable.

G.1 Generators

To exploit the local isomorphism, the generators are conveniently split into $(T_{L,R}^a)_{a=1}^3$, which generate the subgroup locally isomorphic to $SU(2)_{L,R}$, and the broken generators $(\hat{T}^b)_{b=1}^4$ associated with the $SO(5)/SO(4)$ coset. In the fundamental representation they are given by

$$(T_L^a)_{ij} = -\frac{i}{2} (\epsilon^{abc} \delta_i^b \delta_j^c + (\delta_i^a \delta_j^4 - \delta_i^4 \delta_j^a)), \quad (\text{G.1})$$

$$(T_R^a)_{ij} = -\frac{i}{2} (\epsilon^{abc} \delta_i^b \delta_j^c - (\delta_i^a \delta_j^4 - \delta_i^4 \delta_j^a)), \quad (\text{G.2})$$

$$(\hat{T}^b)_{ij} = -\frac{i}{\sqrt{2}} (\delta_i^b \delta_j^5 - \delta_j^b \delta_i^5). \quad (\text{G.3})$$

Note this is consistent with our choice of the $SO(4)$ -invariant vector $\Phi_0 = (0, 0, 0, 0, 1)^T$ that implicitly embeds $SO(4)$ into the top left 4×4 block of $SO(5)$ matrices.

G.2 Field Embeddings

The representations of $SO(5)$ that we are interested in decompose under $SO(4)$ and subsequently under $SU(2)_L \times SU(2)_R$ as

$$\begin{aligned} \mathbf{5} &\rightarrow \mathbf{4} \oplus \mathbf{1} && \cong (\mathbf{2}, \mathbf{2}) \oplus (\mathbf{1}, \mathbf{1}), \\ \mathbf{10} &\rightarrow \mathbf{6} \oplus \mathbf{4} && \cong (\mathbf{3}, \mathbf{1}) \oplus (\mathbf{1}, \mathbf{3}) \oplus (\mathbf{2}, \mathbf{2}), \\ \mathbf{14} &\rightarrow \mathbf{9} \oplus \mathbf{4} \oplus \mathbf{1} && \cong (\mathbf{3}, \mathbf{3}) \oplus (\mathbf{2}, \mathbf{2}) \oplus (\mathbf{1}, \mathbf{1}). \end{aligned} \quad (\text{G.4})$$

Multiplets in the representation \mathbf{r} will be denoted $\Psi_{\mathbf{r}}$. Representations of $SO(4)$ are embedded inside those of $SO(5)$ as

$$\Psi_{\mathbf{5}} = \begin{pmatrix} \Psi_4 \\ \Psi_1 \end{pmatrix}, \quad \Psi_{\mathbf{10}} = \left(\begin{array}{c|c} \Psi_6 & \frac{1}{\sqrt{2}}\Psi_4 \\ \hline -\frac{1}{\sqrt{2}}\Psi_4^\top & 0 \end{array} \right), \quad \Psi_{\mathbf{14}} = \left(\begin{array}{c|c} \Psi_9 - \frac{1}{2\sqrt{5}}\mathbb{1}\Psi_1 & \frac{1}{\sqrt{2}}\Psi_4 \\ \hline \frac{1}{\sqrt{2}}\Psi_4^\top & \frac{2}{\sqrt{5}}\Psi_1 \end{array} \right). \quad (\text{G.5})$$

Here we are treating $\Psi_{\mathbf{10,14}}$ as matrices, which are acted upon by the generators through the commutator,

$$T^A \Psi_{\mathbf{10,14}} \equiv [T^A, \Psi_{\mathbf{10,14}}], \quad (\text{G.6})$$

in accordance with their adjoint transformation rules. The $SO(4)$ multiplets are populated with fields Ψ^{n_L, n_R} having $SU(2)_L \times SU(2)_R$ quantum numbers¹ (n_L, n_R) according to [145]

$$\begin{aligned} \Psi_{\mathbf{1}} &= \Psi^{0,0}, \\ \Psi_{\mathbf{4}} &= \frac{1}{\sqrt{2}} \begin{pmatrix} i\Psi^{-,-} - i\Psi^{+,+} \\ \Psi^{-,-} + \Psi^{+,+} \\ i\Psi^{+,-} - \Psi^{-,+} \\ i\Psi^{-,+} - \Psi^{+,-} \end{pmatrix}, \\ \Psi_{\mathbf{6}} &= \frac{1}{2} \begin{pmatrix} 0 & \Psi_+^{0,0} & -i(\Psi_-^{\pm,0} + \Psi_-^{0,\pm}) & \Psi_+^{\pm,0} - \Psi_+^{0,\pm} \\ & 0 & \Psi_+^{\pm,0} + \Psi_+^{0,\pm} & i(\Psi_-^{\pm,0} - \Psi_-^{0,\pm}) \\ & & 0 & -i\Psi_-^{0,0} \\ & & & 0 \end{pmatrix}, \\ \Psi_{\mathbf{9}} &= \frac{1}{2} \begin{pmatrix} \Psi_+^{\pm,\pm} - \Psi^{0,0} & & i\Psi_+^{\pm,\pm} & \Psi_+^{\pm,0} + i\Psi_+^{0,\pm} & i\Psi_-^{\pm,0} - \Psi_+^{0,\pm} \\ & & -(\Psi_+^{\pm,\pm} + \Psi^{0,0}) & i\Psi_-^{\pm,0} - \Psi_-^{0,\pm} & i\Psi_+^{0,\pm} - \Psi_+^{\pm,0} \\ & & & \Psi_-^{0,0} - \Psi_-^{\mp,\pm} & i\Psi_+^{\mp,\pm} \\ & & & & \Psi_-^{0,0} + \Psi_-^{\mp,\pm} \end{pmatrix}, \end{aligned} \quad (\text{G.7})$$

where $\Psi_{\mathbf{6}}$ is antisymmetric and $\Psi_{\mathbf{9}}$ symmetric. In a questionably successful effort to reduce clutter, the superscript \pm denotes $\pm 1/2$ in $\Psi_{\mathbf{4}}$ and ± 1 in the other multiplets, and we have defined²

$$\begin{aligned} \Psi_{\pm}^{0,0} &= \Psi_1^{0,0} \pm \Psi_2^{0,0}, & \Psi_{\pm}^{\pm,0} &= (\Psi^{+,0} \pm \Psi^{-,0})/\sqrt{2}, \\ \Psi_{\pm}^{\mp,\pm} &= \Psi^{-,+} \pm \Psi^{+,-}, & \Psi_{\pm}^{0,\pm} &= (\Psi^{0,+} \pm \Psi^{0,-})/\sqrt{2}, \\ \Psi_{\pm}^{\pm,\pm} &= \Psi^{+,+} \pm \Psi^{-,-}. \end{aligned} \quad (\text{G.8})$$

¹Essentially the $SO(5)$ multiplets are being expressed as linear combinations of eigenvectors of $T_{L,R}^3$ with eigenvalues $n_{L,R}$, and Ψ^{n_L, n_R} are the coefficients.

²Two different fields have $(n_L, n_R) = (0, 0)$ in $\Psi_{\mathbf{6}}$.

Mass Matrices in the M4DCHM

Here we give the explicit mass matrices for each model of the M4DCHM. They are given in the Site 0 holographic gauge, in bases consisting of the fields with definite $SU(2)_L \times SU(2)_R$ quantum numbers presented in Appendix G. Only particles of the same electric charge can mix, so the mass matrices are separated according to the particles' charges.

H.1 Boson Sector

The mixing terms between the gauge bosons are independent of the fermion representations, so they are the same across all models considered in this work. Their mixings are represented by symmetric matrices, whose singular values are the squared masses of the resonances.

Expansion of the Lagrangian in Equation (4.6) reveals the $SU(3)$ resonances only mix among themselves, through the matrix

$$M_{\text{gluon}}^2 = \left(\begin{array}{c|cc} & G_\mu^0 & \rho_{G_\mu} \\ \hline G_\mu^0 & \frac{1}{2}(g_s^0)^2 f_G^2 & -\frac{1}{2}g_s^0 g_G f_G^2 \\ \rho_{G_\mu} & & \frac{1}{2}g_G^2 f_G^2 \end{array} \right). \quad (\text{H.1})$$

Note one singular value of this matrix is 0, as expected for the massless SM gluons.

Of the remaining bosons, four are charged and seven are uncharged. The charged bosons $\{W_\mu^{0\pm}, \rho_{L_\mu}^\pm, \rho_{R_\mu}^\pm, \mathbf{a}_\mu^\pm\}$ are linear combinations of the gauge fields that come paired with the generators (using our choice of generators) given by the form

$$X_\mu^\pm := \frac{X_\mu^1 \mp iX_\mu^2}{\sqrt{2}}. \quad (\text{H.2})$$

Their mass matrix is given by Equation (H.3), and the uncharged bosons have the mass matrix Equation (H.4).

$$M_{\text{charged}}^2 = \begin{pmatrix} W_{\mu}^{0+} & \rho_{L\mu}^+ & \rho_{R\mu}^+ & \alpha_{\mu}^+ \\ \hline W_{\mu}^{0-\mu} & -\frac{1}{2}g_0g_{\rho}f_1^2 \cos^2\left(\frac{h}{2f}\right) & -\frac{1}{2}g_0g_{\rho}f_1^2 \sin^2\left(\frac{h}{2f}\right) & -\frac{1}{2\sqrt{2}}g_0g_{\rho}f_1^2 \sin\left(\frac{h}{f}\right) \\ \rho_L^{-\mu} & \frac{1}{2}g_{\rho}^2 f_1^2 & 0 & 0 \\ \rho_R^{-\mu} & \frac{1}{2}g_{\rho}^2 f_1^2 & 0 & 0 \\ \alpha^{-\mu} & \frac{1}{2}g_{\rho}^2 f_1^2 & \frac{1}{2}g_{\rho}^2 f_1^2 & \frac{1}{2}g_{\rho}^2 \frac{f_1^4}{f_1^2 - f^2} \end{pmatrix}, \quad (\text{H.3})$$

$$M_{\text{neutral}}^2 = \begin{pmatrix} W_{\mu}^{03} & B_{\mu}^0 & \rho_{L\mu}^3 & \rho_{R\mu}^3 & \alpha_{\mu}^3 & \alpha_{\mu}^4 \\ \hline W_{\mu}^{03\mu} & 0 & -\frac{1}{2}g_0g_{\rho}f_1^2 \cos^2\left(\frac{h}{2f}\right) & -\frac{1}{2}g_0g_{\rho}f_1^2 \sin^2\left(\frac{h}{2f}\right) & -\frac{1}{2\sqrt{2}}g_0g_{\rho}f_1^2 \sin\left(\frac{h}{f}\right) & 0 \\ B^{0\mu} & \frac{1}{2}g_0^2(f_1^2 + f_x^2) & -\frac{1}{2}g_0'g_{\rho}f_1^2 \sin^2\left(\frac{h}{2f}\right) & -\frac{1}{2}g_0'g_{\rho}f_1^2 \cos^2\left(\frac{h}{2f}\right) & -\frac{1}{2}g_0'g_X f_X^2 & 0 \\ \rho_L^{3\mu} & \frac{1}{2}g_{\rho}^2 f_1^2 & 0 & 0 & 0 & 0 \\ \rho_R^{3\mu} & \frac{1}{2}g_{\rho}^2 f_1^2 & 0 & 0 & 0 & 0 \\ \alpha^{3\mu} & \frac{1}{2}g_{\rho}^2 f_1^2 & 0 & 0 & 0 & 0 \\ \rho_X^{\mu} & \frac{1}{2}g_X^2 f_X^2 & \frac{1}{2}g_X^2 f_X^2 & \frac{1}{2}g_X^2 f_X^2 & \frac{1}{2}g_X^2 f_X^2 & \frac{1}{2}g_X^2 f_X^2 \\ \alpha^{4\mu} & \frac{1}{2}g_{\rho}^2 f_1^2 & \frac{1}{2}g_{\rho}^2 f_1^2 & \frac{1}{2}g_{\rho}^2 f_1^2 & \frac{1}{2}g_{\rho}^2 f_1^2 & \frac{1}{2}g_{\rho}^2 \frac{f_1^4}{f_1^2 - f^2} \end{pmatrix}. \quad (\text{H.4})$$

H.2 Fermion Sector

The Lagrangians Equations (4.22), (4.25) and (4.27) lead to the quark partner mass matrices for each model below. Subscripts denote the matrices for the up-type (U) and down-type (D) fields, as well as the fields of exotic charges 4/3, 5/3, and 8/3. The subscripts/tildes of the fields match those of the multiplets in which they are contained. Where necessary to avoid ambiguity, subscripts denote the representations of $SO(5)$ or $SO(4)$ in which the fields reside. Fields may share a label with those in other matrices due to the different definitions of \pm in superscripts for different $SO(4)$ representations (see discussion after Equation (G.8)), but the electric charges distinguish them and remove ambiguity about the representation each comes from.

To save space in writing the matrices, we define the quantities

$$s_{xh} = \sin\left(x \frac{h}{f}\right), \quad c_{xh} = \cos\left(x \frac{h}{f}\right) \quad \text{for } x \in \mathbb{R}, \quad (\text{H.5})$$

and

$$c_{\pm} = \frac{c_h \pm c_{2h}}{2}, \quad \tilde{c} = \frac{3 + 5c_{2h}}{8}, \quad \text{and} \quad \tilde{s} = \frac{\sqrt{5}}{4}s_{2h}. \quad (\text{H.6})$$

H.2.1 M4DCHM⁵⁻⁵⁻⁵

A caveat: in this model a different convention was used for the field embeddings than was presented in Appendix G. The fields in Ψ_4 in Appendix G have been redefined here as $\Psi^{-,-} \rightarrow -i\Psi^{-,-}$, $\Psi^{+,+} \rightarrow i\Psi^{+,+}$, and $\Psi^{+,-} \rightarrow -i\Psi^{+,-}$.

$$M_{\frac{4}{3}}^{5-5-5} = \left(\begin{array}{c|cc} & \Psi_{bR}^{-,-} & \tilde{\Psi}_{bR}^{-,-} \\ \hline \bar{\Psi}_{bL}^{-,-} & m_b & m_{Y_b} \\ \tilde{\Psi}_{bL}^{-,-} & 0 & m_{\tilde{b}} \end{array} \right), \quad (\text{H.7})$$

$$M_{\frac{5}{3}}^{5-5-5} = \left(\begin{array}{c|cc} & \Psi_{tR}^{+,+} & \tilde{\Psi}_{tR}^{+,+} \\ \hline \bar{\Psi}_{tL}^{+,+} & m_t & m_{Y_t} \\ \tilde{\Psi}_{tL}^{+,+} & 0 & m_{\tilde{t}} \end{array} \right), \quad (\text{H.8})$$

$$\begin{aligned}
M_U^{5-5-5} &= \begin{pmatrix} t_L^0 & \Psi_{tR}^{+,-} & \tilde{\Psi}_{tR}^{+,-} & \Psi_{tR}^{-,+} & \tilde{\Psi}_{tR}^{-,+} & \Psi_{bR}^{+,+} & \tilde{\Psi}_{bR}^{+,+} & \Psi_{tR}^{0,0} & \tilde{\Psi}_{tR}^{0,0} \\ \Psi_{tL}^{+,-} & -c_{h/2}^2 \Delta_{tL} & 0 & s_{h/2}^2 \Delta_{tL} & 0 & -\Delta_{bL} & 0 & \frac{i}{\sqrt{2}} s_h \Delta_{tL} & 0 \\ \Psi_{tL}^{+,-} & m_t & m_{Y_t} & 0 & 0 & 0 & 0 & 0 & 0 \\ \Psi_{tL}^{+,-} & 0 & m_{\tilde{t}} & 0 & 0 & 0 & 0 & 0 & 0 \\ \Psi_{tL}^{-,+} & 0 & 0 & m_t & m_{Y_t} & 0 & 0 & 0 & 0 \\ \Psi_{tL}^{-,+} & 0 & 0 & 0 & m_{\tilde{t}} & 0 & 0 & 0 & 0 \\ \Psi_{bL}^{+,+} & 0 & 0 & 0 & 0 & m_b & m_{Y_b} & 0 & 0 \\ \Psi_{bL}^{+,+} & 0 & 0 & 0 & 0 & 0 & m_{\tilde{b}} & 0 & 0 \\ \Psi_{tL}^{0,0} & 0 & 0 & 0 & 0 & 0 & 0 & m_t & m_{Y_t} + Y_t \\ \Psi_{tL}^{0,0} & -c_h \Delta_{tR}^\dagger & 0 & 0 & 0 & 0 & 0 & 0 & m_{\tilde{t}} \end{pmatrix}, \tag{H.9}
\end{aligned}$$

$$\begin{aligned}
M_D^{5-5-5} &= \begin{pmatrix} \bar{b}_L^0 & \Psi_{bR}^{+,-} & \tilde{\Psi}_{bR}^{+,-} & \Psi_{bR}^{-,+} & \tilde{\Psi}_{bR}^{-,+} & \Psi_{tR}^{+,-} & \tilde{\Psi}_{tR}^{+,-} & \Psi_{bR}^{0,0} & \tilde{\Psi}_{bR}^{0,0} \\ \Psi_{bL}^{+,-} & s_{h/2}^2 \Delta_{bL} & 0 & -c_{h/2}^2 \Delta_{bL} & 0 & -\Delta_{tL} & 0 & \frac{i}{\sqrt{2}} s_h \Delta_{bL} & 0 \\ \Psi_{bL}^{+,-} & m_b & m_{Y_b} & 0 & 0 & 0 & 0 & 0 & 0 \\ \Psi_{bL}^{-,+} & 0 & m_{\tilde{b}} & 0 & 0 & 0 & 0 & 0 & 0 \\ \Psi_{bL}^{-,+} & 0 & 0 & m_b & m_{Y_b} & 0 & 0 & 0 & 0 \\ \Psi_{dL}^{-,+} & 0 & 0 & 0 & m_{\tilde{b}} & 0 & 0 & 0 & 0 \\ \Psi_{tL}^{+,-} & 0 & 0 & 0 & 0 & m_t & m_{Y_t} & 0 & 0 \\ \Psi_{tL}^{+,-} & 0 & 0 & 0 & 0 & 0 & m_{\tilde{t}} & 0 & 0 \\ \Psi_{bL}^{0,0} & 0 & 0 & 0 & 0 & 0 & 0 & m_b & m_{Y_b} + Y_b \\ \Psi_{bL}^{0,0} & -c_h \Delta_{bR}^\dagger & 0 & 0 & 0 & 0 & 0 & 0 & m_{\tilde{b}} \end{pmatrix}. \tag{H.10}
\end{aligned}$$

H.2.2 M4DCHM¹⁴⁻¹⁴⁻¹⁰

$$M_{\frac{4}{3}}^{14-14-10} = \left(\begin{array}{c|cc} & \Psi_L^{-,-} & \tilde{\Psi}_{14L}^{-,-} \\ \hline \bar{\Psi}_R^{-,-} & m_q & m_{Y_t} \\ \tilde{\Psi}_{14R}^{-,-} & 0 & m_t \end{array} \right), \quad M_{\frac{8}{3}}^{14-14-10} = \left(\begin{array}{c|cc} & \Psi_L^{+,+} & \tilde{\Psi}_{14L}^{+,+} \\ \hline \bar{\Psi}_R^{+,+} & m_q & m_{Y_t} \\ \tilde{\Psi}_{14R}^{+,+} & 0 & m_t \end{array} \right), \quad (\text{H.11})$$

$$M_{\frac{5}{3}}^{14-14-10} = \left(\begin{array}{c|cccccccccc} & \tilde{\Psi}_{10L}^{+,+} & \Psi_L^{+,+} & \tilde{\Psi}_{14L}^{+,+} & \Psi_L^{+,0} & \tilde{\Psi}_{14L}^{+,0} & \Psi_L^{0,+} & \tilde{\Psi}_{14L}^{0,+} & \tilde{\Psi}_{10L}^{+,0} & \tilde{\Psi}_{10L}^{0,+} \\ \hline \tilde{\Psi}_{10R}^{+,+} & m_b & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \bar{\Psi}_R^{+,+} & \frac{1}{2}Y_b & m_q & m_{Y_t} + \frac{1}{2}Y_t & 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{14R}^{+,+} & 0 & 0 & m_t & 0 & 0 & 0 & 0 & 0 & 0 \\ \bar{\Psi}_R^{+,0} & 0 & 0 & 0 & m_q & m_{Y_t} & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{14R}^{+,0} & 0 & 0 & 0 & 0 & m_t & 0 & 0 & 0 & 0 \\ \bar{\Psi}_R^{0,+} & 0 & 0 & 0 & 0 & 0 & m_q & m_{Y_t} & 0 & 0 \\ \tilde{\Psi}_{14R}^{0,+} & 0 & 0 & 0 & 0 & 0 & 0 & m_t & 0 & 0 \\ \tilde{\Psi}_{10R}^{+,0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_b & 0 \\ \tilde{\Psi}_{10R}^{0,+} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_b \end{array} \right), \quad (\text{H.12})$$

$$M_D^{14-14-10} = \left(\begin{array}{c|ccccccccccc} & b_L^0 & \tilde{\Psi}_{10L}^{-,-} & \Psi_L^{-,-} & \tilde{\Psi}_{14L}^{-,-} & \Psi_L^{-,0} & \tilde{\Psi}_{14L}^{-,0} & \Psi_L^{0,-} & \tilde{\Psi}_{14L}^{0,-} & \tilde{\Psi}_{10L}^{-,0} & \tilde{\Psi}_{10L}^{0,-} \\ \hline \bar{b}_R^0 & 0 & 0 & -c_h \Delta_q & 0 & -\frac{1}{\sqrt{2}}s_h \Delta_q & 0 & \frac{i}{\sqrt{2}}s_h \Delta_q & 0 & 0 & 0 \\ \tilde{\Psi}_{10R}^{-,-} & -\frac{i}{\sqrt{2}}s_h \Delta_b^\dagger & m_b & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \bar{\Psi}_R^{-,-} & 0 & \frac{1}{2}Y_b & m_q & m_{Y_t} + \frac{1}{2}Y_t & 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{14R}^{-,-} & 0 & 0 & 0 & m_t & 0 & 0 & 0 & 0 & 0 & 0 \\ \bar{\Psi}_R^{-,0} & 0 & 0 & 0 & 0 & m_q & m_{Y_t} & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{14R}^{-,0} & 0 & 0 & 0 & 0 & 0 & m_t & 0 & 0 & 0 & 0 \\ \bar{\Psi}_R^{0,-} & 0 & 0 & 0 & 0 & 0 & 0 & m_q & m_{Y_t} & 0 & 0 \\ \tilde{\Psi}_{14R}^{0,-} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_t & 0 & 0 \\ \tilde{\Psi}_{10R}^{-,0} & -s_{h/2}^2 \Delta_b^\dagger & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_b & 0 \\ \tilde{\Psi}_{10R}^{0,-} & -c_{h/2}^2 \Delta_b^\dagger & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_b \end{array} \right), \quad (\text{H.13})$$

H.2.3 M4DCHM¹⁴⁻¹⁻¹⁰

$$M_{\frac{4}{3}}^{14-1-10} = \left(\frac{\Psi_L^{-,-}}{\bar{\Psi}_R^{-,-}} \middle| m_q \right), \quad M_{\frac{8}{3}}^{14-1-10} = \left(\frac{\Psi_L^{+,+}}{\bar{\Psi}_R^{+,+}} \middle| m_q \right), \quad (\text{H.15})$$

$$M_{\frac{5}{3}}^{14-1-10} = \left(\begin{array}{c|cccccc} & \tilde{\Psi}_L^{+,+} & \Psi_L^{+,+} & \Psi_L^{+,0} & \Psi_L^{0,+} & \tilde{\Psi}_L^{+,0} & \tilde{\Psi}_L^{0,+} \\ \hline \tilde{\Psi}_R^{+,+} & m_b & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_R^{+,-} & \frac{1}{2}Y_b & m_q & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_R^{+,0} & 0 & 0 & m_q & 0 & 0 & 0 \\ \tilde{\Psi}_R^{0,+} & 0 & 0 & 0 & m_q & 0 & 0 \\ \tilde{\Psi}_R^{+,0} & 0 & 0 & 0 & 0 & m_b & 0 \\ \tilde{\Psi}_R^{0,+} & 0 & 0 & 0 & 0 & 0 & m_b \end{array} \right), \quad (\text{H.16})$$

$$M_D^{14-1-10} = \left(\begin{array}{c|cccccc} & b_L^0 & \tilde{\Psi}_L^{-,-} & \Psi_L^{-,-} & \Psi_L^{0,-} & \Psi_L^{0,-} & \tilde{\Psi}_L^{-,0} & \tilde{\Psi}_L^{0,-} \\ \hline \bar{b}_R^0 & 0 & 0 & -c_h \Delta_q & -\frac{1}{\sqrt{2}}s_h \Delta_q & \frac{i}{\sqrt{2}}s_h \Delta_q & 0 & 0 \\ \tilde{\Psi}_R^{-,-} & -\frac{i}{\sqrt{2}}s_h \Delta_b^\dagger & m_b & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_R^{-,-} & 0 & \frac{1}{2}Y_b & m_q & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_R^{-,0} & 0 & 0 & 0 & m_q & 0 & 0 & 0 \\ \tilde{\Psi}_R^{0,-} & 0 & 0 & 0 & 0 & m_q & 0 & 0 \\ \tilde{\Psi}_R^{-,0} & -s_{h/2}^2 \Delta_b^\dagger & 0 & 0 & 0 & 0 & m_b & 0 \\ \tilde{\Psi}_R^{0,-} & -c_{h/2}^2 \Delta_b^\dagger & 0 & 0 & 0 & 0 & 0 & m_b \end{array} \right), \quad (\text{H.17})$$

$$M_U^{14-1-10} = \left(\begin{array}{c|cccccccccccc} & t_L^0 & \Psi_{1L}^{0,0} & \tilde{\Psi}_L^{0,0} & \tilde{\Psi}_L^{+,-} & \Psi_{4L}^{+,-} & \tilde{\Psi}_L^{-,+} & \Psi_{4L}^{-,+} & \Psi_{9L}^{-,+} & \Psi_{9L}^{0,0} & \Psi_{9L}^{+,-} & \tilde{\Psi}_{1L}^{0,0} & \tilde{\Psi}_{2L}^{0,0} \\ \hline \bar{t}_R^0 & 0 & \tilde{s} \Delta_q & 0 & 0 & -c_+ \Delta_q & 0 & -ic_- \Delta_q & s_h s_{h/2}^2 \Delta_q & -\frac{s_{2h}}{4} \Delta_q & s_h c_{h/2}^2 \Delta_q & 0 & 0 \\ \tilde{\Psi}_{1R}^{0,0} & 0 & m_q & \frac{2}{\sqrt{5}}Y_t & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{1R}^{0,0} & -\Delta_t^\dagger & 0 & m_t & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_R^{+,-} & 0 & 0 & 0 & m_b & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{4R}^{+,-} & 0 & 0 & 0 & \frac{1}{2}Y_b & m_q & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_R^{-,+} & 0 & 0 & 0 & 0 & 0 & m_b & 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{4R}^{-,+} & 0 & 0 & 0 & 0 & 0 & \frac{1}{2}Y_b & m_q & 0 & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{9R}^{-,+} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_q & 0 & 0 & 0 & 0 \\ \tilde{\Psi}_{9R}^{0,0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_q & 0 & 0 & 0 \\ \tilde{\Psi}_{9R}^{+,-} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_q & 0 & 0 \\ \tilde{\Psi}_{1R}^{0,0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_b & 0 \\ \tilde{\Psi}_{2R}^{0,0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & m_b \end{array} \right). \quad (\text{H.18})$$

H.2.4 Calculation of the Top and Bottom Masses

Fermion masses are the singular values of the Higgs-dependent mass matrices $M(s_h)$ given above. Unfortunately, many of the singular values, including those corresponding to the top and bottom quark masses, cannot be expressed in a reasonably understandable form. But a type of perturbation theory may be employed to find the SM quark masses as power series in $s_h \equiv \sin(h/f)$.

The calculation takes advantage of the fact that the squared quark (and quark partner) masses are the eigenvalues of

$$M^2(s_h) := M^\dagger(s_h)M(s_h). \quad (\text{H.19})$$

Let \vec{v} denote the eigenvector corresponding to the appropriate SM quark for the matrix $M^2(s_h)$, so that

$$M^2(s_h)\vec{v} \equiv m^2\vec{v}, \quad (\text{H.20})$$

where m is the quark mass. We expand these quantities in powers of s_h :

$$\begin{aligned} M^2(s_h) &= \sum_{k=0}^{\infty} \frac{1}{k!} \frac{\partial^k M^2}{\partial s_h^k} \Big|_{s_h=0} s_h^k, \\ \vec{v} &= \sum_{k=0}^{\infty} \vec{v}_k s_h^k, \\ m^2 &= \sum_{k=2}^{\infty} m_k^2 s_h^k, \end{aligned} \quad (\text{H.21})$$

where crucially $m^2 = \mathcal{O}(s_h^2)$ because the SM quarks are massless before EWSB. Equating the coefficients of s_h^n in Equation (H.20) gives the general relation

$$\sum_{k=0}^n \frac{1}{k!} \frac{\partial^k M^2}{\partial s_h^k} \Big|_{s_h=0} \vec{v}_{n-k} = \sum_{k=2}^n m_k^2 \vec{v}_{n-k}, \quad (\text{H.22})$$

where the right-hand side vanishes for $n < 2$. It is possible to solve Equation (H.22) (for $n = 0, 1, 2, 3, \dots$) simultaneously for the coefficients m_k^2 . Doing so for the models considered here yields the first-order formulae for the top and bottom masses given below in Equations (H.23) to (H.25).

M4DCHM⁵⁻⁵⁻⁵

$$m_{\text{top}}^2 = \frac{m_b^2 m_t^2 m_q^2 m_t^2 Y_t^2 \Delta_{tL}^2 \Delta_{tR}^2}{2 \left(m_t^2 m_t^2 (m_b^2 + m_{Y_b}^2) \Delta_{bL}^2 + m_b^2 m_b^2 (m_t^2 m_t^2 + (m_t^2 + m_{Y_t}^2) \Delta_{tL}^2) \right) \left((m_{Y_t} + Y_t)^2 \Delta_{tR}^2 + m_t^2 (m_t^2 + \Delta_{tR}^2) \right)} s_{\langle h \rangle}^2 + \mathcal{O}(s_{\langle h \rangle}^4),$$

$$m_{\text{bottom}}^2 = \text{same as } m_{\text{top}}^2 \text{ with } (t, \tilde{t}) \text{ and } (b, \tilde{b}) \text{ interchanged.} \quad (\text{H.23})$$

M4DCHM¹⁴⁻¹⁴⁻¹⁰

$$m_{\text{top}}^2 = \frac{m_q^2 m_t^2 m_b^2 \left(\frac{3}{4} Y_t + 2 \tilde{Y}_t \right)^2 \Delta_q^2 \Delta_t^2}{5 \left(\frac{1}{4} m_t^2 Y_b^2 \Delta_q^2 + m_b^2 \left(m_q^2 m_t^2 + \Delta_q^2 \left(m_t^2 + (m_{Y_t} + \frac{1}{2} Y_t)^2 \right) \right) \right) \left(m_q^2 m_t^2 + \Delta_t^2 \left(m_q^2 + \left(m_{Y_t} + \frac{4}{5} Y_t + \frac{4}{5} \tilde{Y}_t \right)^2 \right) \right)} s_{\langle h \rangle}^2 + \mathcal{O}(s_{\langle h \rangle}^3),$$

$$m_{\text{bottom}}^2 = \frac{m_b^2 m_t^2 Y_b^2 \Delta_q^2 \Delta_b^2}{8 \left(m_b^2 + \Delta_b^2 \right) \left(\frac{1}{4} m_t^2 Y_b^2 \Delta_q^2 + m_b^2 \left(m_q^2 m_t^2 + \Delta_q^2 \left(m_t^2 + (m_{Y_t} + \frac{1}{2} Y_t)^2 \right) \right) \right)} s_{\langle h \rangle}^2 + \mathcal{O}(s_{\langle h \rangle}^3). \quad (\text{H.24})$$

M4DCHM¹⁴⁻¹⁻¹⁰

$$m_{\text{top}}^2 = \frac{m_q^2 m_b^2 Y_t^2 \Delta_q^2 \Delta_t^2}{\left(\frac{1}{4} Y_b^2 \Delta_q^2 + m_b^2 (m_q^2 + \Delta_q^2) \right) \left(m_q^2 m_t^2 + \Delta_t^2 (m_q^2 + \frac{4}{5} Y_t^2) \right)} s_{\langle h \rangle}^2 + \mathcal{O}(s_{\langle h \rangle}^3),$$

$$m_{\text{bottom}}^2 = \frac{m_b^2 Y_b^2 \Delta_q^2 \Delta_b^2}{8 \left(m_b^2 + \Delta_b^2 \right) \left(\frac{1}{4} Y_b^2 \Delta_q^2 + m_b^2 (m_q^2 + \Delta_q^2) \right)} s_{\langle h \rangle}^2 + \mathcal{O}(s_{\langle h \rangle}^3). \quad (\text{H.25})$$

Fermion Form Factors

Here we provide expressions for the form factors $\Pi_{t,b_{L,R}}$ and $M_{t,b}$ for each model of the M4DCHM, originally given in Ref. [100]. See Equation (4.33) for the definitions of these functions. In all models, the form factors are given in terms of the functions

$$\begin{aligned}
 A_R(m_1, m_2, m_3, m_4, \Delta) &:= \Delta^2(m_1^2 m_2^2 + m_2^2 m_3^2 - p^2(m_1^2 + m_2^2 + m_3^2 + m_4^2) + p^4), \\
 A_L(m_1, m_2, m_3, m_4, \Delta) &:= \Delta^2 m_1^2 m_4^2 + A_R(m_1, m_2, m_3, m_4, \Delta), \\
 A_M(m_1, m_2, m_3, m_4, \Delta_1, \Delta_2) &:= \Delta_1 \Delta_2 m_1 m_2 m_4 (m_3^2 - p^2), \\
 B(m_1, m_2, m_3, m_4, m_5) &:= m_1^2 m_2^2 m_3^2 - p^2(m_1^2 m_2^2 + m_1^2 m_3^2 + m_2^2 m_3^2 + m_2^2 m_5^2 + m_3^2 m_4^2) \\
 &\quad + p^4(m_1^2 + m_2^2 + m_3^2 + m_4^2 + m_5^2) - p^6, \tag{I.1}
 \end{aligned}$$

in Minkowski space.

I.1 M4DCHM⁵⁻⁵⁻⁵

In this model,

$$\begin{aligned}
 \Pi_{t_L} &= \Pi_{q_t}^{(4)} + \Pi_{q_b}^{(4)} + \frac{1}{2} (\Pi_{q_t}^{(1)} - \Pi_{q_t}^{(4)}) s_h^2, & M_t &= (M_t^{(1)} - M_t^{(4)}) \sqrt{\frac{1 - s_h^2}{2}} s_h, \\
 \Pi_{t_R} &= \Pi_t^{(1)} - (\Pi_t^{(1)} - \Pi_t^{(4)}) s_h^2, \tag{I.2}
 \end{aligned}$$

where

$$\begin{aligned}
 \Pi_{q_t}^{(1)} &= \frac{A_L(m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}} + Y_t, 0, \Delta_{tL})}{B(m_t, m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}} + Y_t, 0)}, & \Pi_{q_t}^{(4)} &= \frac{A_L(m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}}, 0, \Delta_{tL})}{B(m_t, m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}}, 0)}, \\
 \Pi_t^{(1)} &= \frac{A_L(m_t, 0, m_{Y_t} + Y_t, 0, \Delta_{tR})}{B(m_t, m_{\tilde{t}}, 0, m_{Y_t} + Y_t, 0)}, & \Pi_t^{(4)} &= \frac{A_L(m_t, 0, m_{Y_t}, 0, \Delta_{tR})}{B(m_t, m_{\tilde{t}}, 0, m_{Y_t}, 0)}, \\
 M_t^{(1)} &= \frac{A_M(m_t, m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}} + Y_t, \Delta_{tL}, \Delta_{tR})}{B(m_t, m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}} + Y_t, 0)}, & M_t^{(4)} &= \frac{A_M(m_t, m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}}, \Delta_{tL}, \Delta_{tR})}{B(m_t, m_{\tilde{t}}, 0, m_{Y_{\tilde{t}}}, 0)}. \tag{I.3}
 \end{aligned}$$

The form factors for the bottom sector are obtained by interchanging all t and \tilde{t} subscripts with b and \tilde{b} .

I.2 M4DCHM¹⁴⁻¹⁴⁻¹⁰

In this model,

$$\begin{aligned}\Pi_{tL} &= \Pi_q^{(4)} - (\Pi_q^{(4)} - \Pi_q^{(9)}) \frac{s_h^2}{2} + (5\Pi_q^{(1)} - 8\Pi_q^{(4)} + 3\Pi_q^{(9)}) \frac{s_h^2(1 - s_h^2)}{4}, \\ \Pi_{tR} &= \frac{1}{5} (8\Pi_t^{(4)} - 3\Pi_t^{(9)}) - \frac{3}{2} (\Pi_t^{(4)} - \Pi_t^{(9)}) s_h^2 + (5\Pi_t^{(1)} - 8\Pi_t^{(4)} + 3\Pi_t^{(9)}) \frac{(4 - 5s_h^2)^2}{5}, \\ M_t &= i \frac{\sqrt{5}}{2} \left((M_t^{(1)} - M_t^{(4)}) - (5M_t^{(1)} - 8M_t^{(4)} + 3M_t^{(9)}) \frac{s_h^2}{4} \right) s_h \sqrt{1 - s_h^2},\end{aligned}\quad (\text{I.4})$$

and

$$\begin{aligned}\Pi_{bL} &= \Pi_q^{(4)} - (\Pi_q^{(4)} - \Pi_q^{(9)}) s_h^2, & M_b &= -M_b^{(4)} s_h \sqrt{\frac{1 - s_h^2}{2}}, \\ \Pi_{bR} &= \Pi_b^{(6)} + (\Pi_b^{(4)} - \Pi_b^{(6)}) \frac{s_h^2}{2},\end{aligned}\quad (\text{I.5})$$

where

$$\begin{aligned}\Pi_q^{(1)} &= \frac{A_L(m_t, 0, m_{Y_t}^{(1)}, 0, \Delta_q)}{B(m_q, m_t, 0, m_{Y_t}^{(1)}, 0)}, & \Pi_t^{(1)} &= \frac{A_R(m_q, 0, m_{Y_t}^{(1)}, 0, \Delta_t)}{B(m_q, m_t, 0, m_{Y_t}^{(1)}, 0)}, \\ \Pi_q^{(4)} &= \frac{A_L(m_t, m_b, m_{Y_t}^{(4)}, Y_b/2, \Delta_q)}{B(m_q, m_t, m_b, m_{Y_t}^{(4)}, Y_b/2)}, & \Pi_t^{(4)} &= \frac{A_R(m_q, m_b, m_{Y_t}^{(4)}, Y_b/2, \Delta_t)}{B(m_q, m_t, m_b, m_{Y_t}^{(4)}, Y_b/2)}, \\ \Pi_q^{(9)} &= \frac{A_L(m_t, 0, m_{Y_t}, 0, \Delta_q)}{B(m_q, m_t, 0, m_{Y_t}, 0)}, & \Pi_t^{(9)} &= \frac{A_R(m_q, 0, m_{Y_t}, 0, \Delta_t)}{B(m_q, m_t, 0, m_{Y_t}, 0)}, \\ \Pi_b^{(4)} &= \frac{A_R(m_q, m_t, Y_b/2, m_{Y_t}^{(4)}, \Delta_b)}{B(m_q, m_t, m_b, m_{Y_t}^{(4)}, Y_b/2)}, & M_t^{(1)} &= \frac{A_M(m_q, m_t, 0, m_{Y_t}^{(1)}, \Delta_q, \Delta_t)}{B(m_q, m_t, 0, m_{Y_t}^{(1)}, 0)}, \\ \Pi_b^{(6)} &= \frac{A_R(m_q, 0, 0, 0, \Delta_b)}{B(m_q, m_b, 0, 0, 0)}, & M_t^{(4)} &= \frac{A_M(m_q, m_t, m_b, m_{Y_t}^{(4)}, \Delta_q, \Delta_t)}{B(m_q, m_t, m_b, m_{Y_t}^{(4)}, Y_b/2)}, \\ M_b^{(4)} &= -i \frac{A_M(m_q, m_b, m_t, Y_b/2, \Delta_q, \Delta_b)}{B(m_q, m_t, m_b, m_{Y_t}^{(4)}, Y_b/2)}, & M_t^{(9)} &= \frac{A_M(m_q, m_t, 0, m_{Y_t}, \Delta_q, \Delta_t)}{B(m_q, m_t, 0, m_{Y_t}, 0)},\end{aligned}\quad (\text{I.6})$$

having defined

$$\begin{aligned}m_{Y_t}^{(1)} &:= m_{Y_t} + 4Y_t/5 + 4\tilde{Y}_t/5, \\ m_{Y_t}^{(4)} &:= m_{Y_t} + \frac{1}{2}Y_t,\end{aligned}\quad (\text{I.7})$$

for convenience.

I.3 M4DCHM¹⁴⁻¹⁻¹⁰

In this model,

$$\begin{aligned}\Pi_{tL} &= \Pi_q^{(4)} - (\Pi_q^{(4)} - \Pi_q^{(9)}) \frac{s_h^2}{2} + (5\Pi_q^{(1)} - 8\Pi_q^{(4)} + 3\Pi_q^{(9)}) \frac{s_h^2(1-s_h^2)}{4}, \\ \Pi_{tR} &= \Pi_t^{(1)}, \\ M_t &= -\frac{\sqrt{5}}{4} M_t^{(1)} s_h,\end{aligned}\tag{I.8}$$

and

$$\begin{aligned}\Pi_{bL} &= \Pi_q^{(4)} - (\Pi_q^{(4)} - \Pi_q^{(9)}) s_h^2, & M_b &= -i M_b^{(4)} s_h \sqrt{\frac{1-s_h^2}{2}}, \\ \Pi_{bR} &= \Pi_b^{(6)} + (\Pi_b^{(4)} - \Pi_b^{(6)}) \frac{s_h^2}{2},\end{aligned}\tag{I.9}$$

where

$$\begin{aligned}\Pi_q^{(1)} &= \frac{A_L(m_t, 0, 2Y_t/\sqrt{5}, 0, \Delta_q)}{B(m_q, m_t, 0, 2Y_t/\sqrt{5}, 0)}, & \Pi_t^{(1)} &= \frac{A_R(m_q, 0, 2Y_t/\sqrt{5}, 0, \Delta_t)}{B(m_q, m_t, 0, 2Y_t/\sqrt{5}, 0)}, \\ \Pi_q^{(4)} &= \frac{A_L(0, m_b, 0, Y_b/2, \Delta_q)}{B(m_q, 0, m_b, 0, Y_b/2)}, & \Pi_b^{(4)} &= \frac{A_R(m_q, 0, Y_b/2, 0, \Delta_b)}{B(m_q, 0, m_b, 0, Y_b/2)}, \\ \Pi_q^{(9)} &= \frac{A_L(0, 0, 0, 0, \Delta_q)}{B(m_q, 0, 0, 0, 0)}, & \Pi_b^{(6)} &= \frac{A_R(0, 0, 0, 0, \Delta_b)}{B(0, m_b, 0, 0, 0)}, \\ M_t^{(1)} &= -\frac{A_M(m_q, m_t, 0, 2Y_t/\sqrt{5}, \Delta_q, \Delta_t)}{B(m_q, m_t, 0, 2Y_t/\sqrt{5}, 0)}, & M_b^{(4)} &= -i \frac{A_M(m_q, m_b, 0, Y_b/2, \Delta_q, \Delta_b)}{B(m_q, 0, m_b, 0, Y_b/2)}.\end{aligned}\tag{I.10}$$

Collider Search Constraints

The complete list of collider searches we use to constrain our models in our global fits is given in the following tables.

Decay	Experiment	\sqrt{s} (TeV)	Analysis	Ref.
$V \rightarrow \ell\nu$	ATLAS	7	EXOT-2012-02	[146]
$V \rightarrow e\nu$	ATLAS	7	EXOT-2012-02	[146]
	CMS	13	PAS-EXO-15-006	[147]
	ATLAS	13	CONF-2016-061	[148]
	ATLAS	13	CONF-2018-017	[149]
$V \rightarrow \mu\nu$	ATLAS	7	EXOT-2012-02	[146]
	CMS	13	PAS-EXO-15-006	[147]
	ATLAS	13	CONF-2016-061	[148]
	ATLAS	13	CONF-2018-017	[149]
$V \rightarrow \tau\nu$	CMS	8	EXO-12-011	[150]
	CMS	13	PAS-EXO-16-006	[151]
	CMS	13	PAS-EXO-16-006	[151]
$V \rightarrow ee$	ATLAS	8	EXOT-2012-23	[152]
	CMS	8	EXO-12-061	[153]
	CMS	13	EXO-18-006	[154]
	CMS	13	PAS-EXO-16-031	[155]
	ATLAS	13	CONF-2016-045	[156]
	ATLAS	13	EXOT-2016-05	[157]
	ATLAS	13	EXOT-2018-08	[158]
	ATLAS	13	EXOT-2018-08	[158]
$V \rightarrow \mu\mu$	ATLAS	8	EXOT-2012-23	[152]
	CMS	8	EXO-12-061	[153]
	CMS	13	EXO-16-047	[159]
	CMS	13	PAS-EXO-16-031	[155]
	ATLAS	13	EXOT-2016-05	[157]
	ATLAS	13	CONF-2016-045	[156]
$V \rightarrow \tau\tau$	ATLAS	8	EXOT-2014-05	[160]
	CMS	8	EXO-12-046	[161]
	CMS	13	PAS-EXO-16-008	[162]

Table J.1: Experimental analyses of leptonic decay channels of vector resonances included as constraints in our scans.

Decay	Experiment	\sqrt{s} (TeV)	Analysis	Ref.
$V \rightarrow qq$	CMS	13	PAS-EXO-16-032	[163]
	CMS	13	PAS-EXO-16-032	[163]
$V \rightarrow jj$	CMS	13	EXO-15-001	[164]
	CMS	13	EXO-16-056	[165]
	ATLAS	13	EXOT-2015-02	[166]
	ATLAS	13	EXOT-2018-05	[167]
$V \rightarrow tb$	CMS	8	B2G-12-010	[168]
	CMS	8	B2G-12-009	[169]
	CMS	13	PAS-B2G-16-009	[170]
	CMS	13	PAS-B2G-16-017	[171]
	ATLAS	13	EXOT-2017-02	[172]
$V \rightarrow tt$	ATLAS	8	CONF-2015-009	[173]
	CMS	8	B2G-13-008	[174]
	CMS	13	PAS-B2G-15-003	[175]
	CMS	13	PAS-B2G-15-002	[176]
	CMS	13	B2G-17-017	[177]
	ATLAS	13	EXOT-2015-04	[178]
	ATLAS	13	EXOT-2016-24	[179]

Table J.2: Experimental analyses of hadronic decay channels of vector resonances included as constraints in our scans. Here, j refers to a light quark or b jet, and q refers to a light quark jet.

Decay	Experiment	\sqrt{s} (TeV)	Analysis	Ref.
$V \rightarrow ZH$	ATLAS	8	EXOT-2013-23	[180]
	CMS	8	EXO-13-007	[181]
	CMS	13	PAS-B2G-16-003	[182]
	CMS	13	B2G-17-006	[183]
	CMS	13	B2G-17-002	[184]
	CMS	13	B2G-17-004	[185]
	ATLAS	13	EXOT-2015-18	[186]
	ATLAS	13	CONF-2015-074	[187]
	ATLAS	13	CONF-2016-083	[188]
$V \rightarrow WZ$	ATLAS	8	EXOT-2013-08	[189]
	ATLAS	8	EXOT-2013-01	[190]
	ATLAS	8	EXOT-2013-07	[191]
	CMS	8	EXO-12-024	[192]
	CMS	13	PAS-EXO-15-002	[193]
	CMS	13	PAS-B2G-16-020	[194]
	CMS	13	B2G-16-029	[195]
	CMS	13	B2G-17-001	[196]
	CMS	13	B2G-17-005	[197]
	CMS	13	B2G-17-013	[198]
	CMS	13	B2G-18-002	[199]
	ATLAS	13	CONF-2016-055	[200]
	ATLAS	13	CONF-2016-062	[201]
ATLAS	13	CONF-2016-082	[202]	
$V \rightarrow WH$	ATLAS	8	EXOT-2013-23	[180]
	CMS	8	EXO-14-010	[203]
	CMS	13	PAS-B2G-16-003	[182]
	CMS	13	B2G-17-006	[183]
	CMS	13	B2G-17-002	[184]
	CMS	13	B2G-17-004	[185]
	ATLAS	13	EXOT-2015-18	[186]
	ATLAS	13	CONF-2016-083	[188]
$V \rightarrow WW$	ATLAS	8	EXOT-2013-01	[190]
	CMS	8	EXO-13-009	[204]
	CMS	13	B2G-17-001	[196]
	CMS	13	B2G-18-002	[199]
	ATLAS	13	CONF-2016-062	[201]
$V \rightarrow WW + ZH$	CMS	13	PAS-B2G-16-007	[205]

Table J.3: Experimental analyses of bosonic decay channels of vector resonances included as constraints in our scans.

Decay	Experiment	\sqrt{s} (TeV)	Analysis	Ref.
$F \rightarrow jW$	CDF	1.96	10110	[206]
	ATLAS	7	EXOT-2011-28	[207]
$F \rightarrow qW$	ATLAS	8	EXOT-2014-10	[208]
	CMS	8	B2G-12-017	[209]
$F \rightarrow bW$	ATLAS	7	EXOT-12-07	[210]
	CMS	7	EXO-11-050	[211]
	CMS	7	EXO-11-099	[212]
	ATLAS	8	CONF-2015-012	[213]
	CMS	8	B2G-12-017	[209]
	CMS	8	B2G-13-005	[214]
	ATLAS	13	CONF-2016-102	[215]
	CMS	13	B2G-17-003	[216]
	CMS	13	B2G-16-024	[217]
$F \rightarrow tW$	CDF	1.96	2009	[218]
	CMS	7	B2G-12-004	[219]
	CMS	8	B2G-12-012	[220]
	CMS	8	B2G-13-003	[221]
	CMS	8	B2G-13-006	[222]
	ATLAS	8	EXOT-2013-16	[223]
	ATLAS	8	EXOT-2014-17	[224]
	ATLAS	13	EXOT-2016-16	[225]
	ATLAS	13	EXOT-2017-34	[226]
	CMS	13	PAS-B2G-15-006	[227]
	CMS	13	B2G-16-019	[228]
CMS	13	B2G-17-014	[229]	

Table J.4: Experimental analyses of heavy quark decays included as constraints in our scans. Here, j refers to a light quark or b jet, and q refers to a light quark jet.

Decay	Experiment	\sqrt{s} (TeV)	Analysis	Ref.
$F \rightarrow jZ$	CDF	1.96	2006	[230]
$F \rightarrow bZ$	CMS	7	EXO-11-066	[231]
	CMS	8	B2G-13-003	[221]
	CMS	8	B2G-13-006	[222]
	ATLAS	13	EXOT-2016-35	[232]
$F \rightarrow tZ$	CMS	7	B2G-12-004	[219]
	CMS	7	EXO-11-005	[233]
	CMS	8	B2G-13-005	[214]
	ATLAS	13	CONF-2016-101	[234]
	ATLAS	13	EXOT-2016-15	[235]
	ATLAS	13	EXOT-2016-13	[236]
	ATLAS	13	EXOT-2016-35	[232]
$F \rightarrow bH$	ATLAS	8	CONF-2015-012	[213]
	CMS	8	B2G-12-019	[237]
	CMS	8	B2G-13-006	[222]
	CMS	8	B2G-14-001	[238]
$F \rightarrow tH$	CMS	8	B2G-13-005	[214]
	CMS	13	PAS-B2G-16-011	[239]
	CMS	13	B2G-16-024	[217]
	ATLAS	13	CONF-2016-013	[240]
	ATLAS	13	EXOT-2016-13	[236]

Table J.5: (*cont.*) Experimental analyses of heavy quark decays included as constraints in our scans. Here, j refers to a light quark or b jet, and q refers to a light quark jet.

Bibliography

- [1] GARGAMELLE NEUTRINO collaboration, *Search for elastic muon-neutrino electron scattering*, *Physics Letters B* **46** (1973) 121 – 124.
- [2] GARGAMELLE NEUTRINO collaboration, *Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment*, *Physics Letters B* **46** (1973) 138–140.
- [3] M. Gell-Mann, *A schematic model of baryons and mesons*, *Physics Letters* **8** (1964) 214 – 215.
- [4] G. Zweig, *An $SU(3)$ model for strong interaction symmetry and its breaking. Version 2*, in *Developments in the Quark Theory of Hadrons, Volume 1* (D. Lichtenberg and S. P. Rosen, eds.), pp. 22–101. 1964.
- [5] P. Higgs, *Broken symmetries, massless particles and gauge fields*, *Physics Letters* **12** (1964) 132 – 133.
- [6] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Physical Review Letters* **13** (1964) 508–509.
- [7] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Physical Review Letters* **13** (1964) 321–323.
- [8] P. W. Anderson, *Plasmons, Gauge Invariance, and Mass*, *Physical Review* **130** (1963) 439–442.
- [9] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, *Physical Review Letters* **13** (1964) 585–587.
- [10] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Physics Letters* **B716** (2012) 1–29, [[1207.7214](#)].
- [11] CMS collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Physics Letters* **B716** (2012) 30–61, [[1207.7235](#)].
- [12] R. Contino, *The Higgs as a Composite Nambu-Goldstone Boson*, in *Physics of the Large and the Small: TASI 2009*, pp. 235–306, 2011, [1005.4269](#), DOI.

- [13] S. Bifani, S. Descotes-Genon, A. Romero Vidal and M.-H. Schune, *Review of Lepton Universality tests in B decays*, *Journal of Physics G* **G46** (2019) 023001, [[1809.06229](#)].
- [14] R. Storn and K. Price, *Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces*, *Journal of Global Optimization* **11** (1997) 341–359.
- [15] G. Folland, *Quantum Field Theory: A Tourist Guide for Mathematicians*. Mathematical Surveys and Monographs. American Mathematical Society, 2008.
- [16] B. Hall, *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*. Graduate Texts in Mathematics. Springer, 2003.
- [17] S. Weinberg, *The Quantum Theory of Fields. Vol. 2: Modern applications*. Cambridge University Press, 2013.
- [18] B. Gripaios, A. Pomarol, F. Riva and J. Serra, *Beyond the Minimal Composite Higgs Model*, *Journal of High Energy Physics* **04** (2009) 070, [[0902.1483](#)].
- [19] S. Weinberg, *The Quantum Theory of Fields, Vol. 1: Foundations*. Cambridge University Press, 2005.
- [20] J. M. Cornwall, D. N. Levin and G. Tiktopoulos, *Derivation of gauge invariance from high-energy unitarity bounds on the S matrix*, *Physical Review D* **10** (1974) 1145–1167.
- [21] S. L. Glashow, *Partial-symmetries of weak interactions*, *Nuclear Physics* **22** (1961) 579 – 588.
- [22] A. Salam and J. Ward, *Electromagnetic and weak interactions*, *Physics Letters* **13** (1964) 168 – 171.
- [23] S. Weinberg, *A Model of Leptons*, *Physical Review Letters* **19** (1967) 1264–1266.
- [24] P. Sikivie, L. Susskind, M. B. Voloshin and V. I. Zakharov, *Isospin Breaking in Technicolor Models*, *Nuclear Physics* **B173** (1980) 189–207.
- [25] L. Susskind, *Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory*, *Physical Review D* **20** (1979) 2619–2625.
- [26] M. J. G. Veltman, *The Infrared - Ultraviolet Connection*, *Acta Physica Polonica B* **B12** (1981) 437.
- [27] G. 't Hooft, *Naturalness, Chiral Symmetry, and Spontaneous Chiral Symmetry Breaking*, *NATO Science Series B* **59** (1980) 135–157.

-
- [28] C. Csáki and P. Tanedo, *Beyond the Standard Model*, in *Proceedings, 2013 European School of High-Energy Physics*, pp. 169–268, 2015, [1602.04228](#), [DOI](#).
- [29] M. van Beekveld, S. Caron and R. Ruiz de Austri, *The current status of fine-tuning in supersymmetry*, *Journal of High Energy Physics* **01** (2020) 147, [[1906.10706](#)].
- [30] J. Barnard, D. Murnane, M. White and A. G. Williams, *Constraining fine tuning in Composite Higgs Models with partially composite leptons*, *Journal of High Energy Physics* **09** (2017) 049, [[1703.07653](#)].
- [31] N. Arkani-Hamed, S. Dimopoulos and G. R. Dvali, *The Hierarchy Problem and New Dimensions at a Millimeter*, *Physics Letters B* **429** (1998) 263–272, [[hep-ph/9803315](#)].
- [32] L. Randall and R. Sundrum, *A Large Mass Hierarchy from a Small Extra Dimension*, *Physical Review Letters* **83** (1999) 3370–3373, [[hep-ph/9905221](#)].
- [33] S. Dimopoulos and L. Susskind, *Mass without scalars*, *Nuclear Physics B* **155** (1979) 237 – 252.
- [34] E. Eichten and K. Lane, *Dynamical breaking of weak interaction symmetries*, *Physics Letters B* **90** (1980) 125 – 130.
- [35] S. Weinberg, *Implications of dynamical symmetry breaking*, *Physical Review D* **13** (1976) 974–996.
- [36] H. Terazawa, Y. Chikashige and K. Akama, *Unified model of the Nambu-Jona-Lasinio type for all elementary-particle forces*, *Physical Review D* **15** (1977) 480–487.
- [37] S. Dimopoulos and J. Ellis, *Challenges for extended technicolour theories*, *Nuclear Physics B* **182** (1981) 505 – 528.
- [38] D. B. Kaplan and H. Georgi, *$SU(2) \times U(1)$ breaking by vacuum misalignment*, *Physics Letters B* **136** (1984) 183–186.
- [39] D. B. Kaplan, H. Georgi and S. Dimopoulos, *Composite Higgs scalars*, *Physics Letters B* **136** (1984) 187–190.
- [40] M. J. Dugan, H. Georgi and D. B. Kaplan, *Anatomy of a Composite Higgs Model*, *Nuclear Physics B* **254** (1985) 299–326.
- [41] J. M. Maldacena, *The Large N limit of superconformal field theories and supergravity*, *International Journal of Theoretical Physics* **38** (1999) 1113–1133, [[hep-th/9711200](#)].

-
- [42] N. Arkani-Hamed, A. G. Cohen and H. Georgi, *Electroweak symmetry breaking from dimensional deconstruction*, *Physics Letters* **B513** (2001) 232–240, [[hep-ph/0105239](#)].
- [43] R. Rattazzi and A. Zaffaroni, *Comments on the holographic picture of the Randall-Sundrum model*, *Journal of High Energy Physics* **04** (2001) 021, [[hep-th/0012248](#)].
- [44] K. Agashe, R. Contino and A. Pomarol, *The Minimal Composite Higgs Model*, *Nuclear Physics* **B719** (2005) 165–187, [[hep-ph/0412089](#)].
- [45] R. Contino, Y. Nomura and A. Pomarol, *Higgs as a holographic pseudo-Goldstone boson*, *Nuclear Physics* **B671** (2003) 148–174, [[hep-ph/0306259](#)].
- [46] G. Panico and A. Wulzer, *Effective action and holography in 5D gauge theories*, *Journal of High Energy Physics* **05** (2007) 060, [[hep-th/0703287](#)].
- [47] R. Contino and A. Pomarol, *The holographic composite Higgs*, *Comptes Rendus Physique* **8** (2007) 1058–1067.
- [48] O. Witzel, *Review on Composite Higgs Models*, *Proceedings of Science LATTICE2018* (2019) 006, [[1901.08216](#)].
- [49] S. Elitzur, *Impossibility of spontaneously breaking local symmetries*, *Physical Review D* **12** (1975) 3978–3982.
- [50] J. Goldstone, *Field Theories with “Superconductor” Solutions*, *Il Nuovo Cimento (1955-1965)* **19** (1961) 154–164.
- [51] J. Goldstone, A. Salam and S. Weinberg, *Broken Symmetries*, *Phys. Rev.* **127** (1962) 965–970.
- [52] A. Zee, *Quantum Field Theory in a Nutshell*. Nutshell Handbook. Princeton University Press, Princeton, NJ, 2 ed., 2010.
- [53] M. D. Schwartz, *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2014.
- [54] E. Fabri and L. E. Picasso, *Quantum Field Theory and Approximate Symmetries*, *Physical Review Letters* **16** (1966) 408–410.
- [55] C. E. Vayonakis, *Born Helicity Amplitudes and Cross-Sections in Nonabelian Gauge Theories*, *Lettere al Nuovo Cimento (1971-1985)* **17** (1976) 383.
- [56] M. S. Chanowitz and M. K. Gaillard, *The TeV physics of strongly interacting W’s and Z’s*, *Nuclear Physics B* **261** (1985) 379 – 431.

-
- [57] M. Gell-Mann and M. Lévy, *The axial vector current in beta decay*, *Il Nuovo Cimento (1955-1965)* **16** (1960) 705–726.
- [58] B. Gripaios, *Lectures on Effective Field Theory*, [1506.05039](#).
- [59] S. Weinberg, *Approximate Symmetries and Pseudo-Goldstone Bosons*, *Physical Review Letters* **29** (1972) 1698–1701.
- [60] M. Gell-Mann, R. J. Oakes and B. Renner, *Behavior of current divergences under $SU(3)\times SU(3)$* , *Physical Review* **175** (1968) 2195–2199.
- [61] S. Coleman, J. Wess and B. Zumino, *Structure of Phenomenological Lagrangians. I*, *Physical Review* **177** (1969) 2239–2247.
- [62] C. G. Callan, S. Coleman, J. Wess and B. Zumino, *Structure of Phenomenological Lagrangians. II*, *Physical Review* **177** (1969) 2247–2250.
- [63] H. Georgi, *A tool kit for builders of composite models*, *Nuclear Physics B* **266** (1986) 274 – 284.
- [64] J. Thaler, *Little technicolor*, *Journal of High Energy Physics* **2005** (2005) 024024.
- [65] P. P. Stangl, *Direct Constraints, Flavor Physics, and Flavor Anomalies in Composite Higgs Models*, Ph.D. thesis, Technical University of Munich, 2018. [1811.11750](#).
- [66] J. Wess and B. Zumino, *Consequences of Anomalous Ward Identities*, *Physics Letters B* **37** (1971) 95 – 97.
- [67] E. Witten, *Global aspects of current algebra*, *Nuclear Physics B* **223** (1983) 422 – 432.
- [68] C.-S. Chu, P.-M. Ho and B. Zumino, *NonAbelian anomalies and effective actions for a homogeneous space G/H* , *Nuclear Physics* **B475** (1996) 484–504, [[hep-th/9602093](#)].
- [69] G. Panico and A. Wulzer, *The Composite Nambu-Goldstone Higgs*, *Lecture Notes in Physics* (2016) .
- [70] E. Witten, *Some Inequalities Among Hadron Masses*, *Physical Review Letters* **51** (1983) 2351.
- [71] J. Preskill, *Subgroup Alignment in Hypercolor Theories*, *Nuclear Physics* **B177** (1981) 21–59.
- [72] M. E. Peskin, *The Alignment of the Vacuum in Theories of Technicolor*, *Nuclear Physics* **B175** (1980) 197–233.

- [73] N. Arkani-Hamed, A. G. Cohen, E. Katz, A. E. Nelson, T. Gregoire and J. G. Wacker, *The Minimal moose for a little Higgs*, *Journal of High Energy Physics* **08** (2002) 021, [[hep-ph/0206020](#)].
- [74] N. Arkani-Hamed, A. G. Cohen, E. Katz and A. E. Nelson, *The Littlest Higgs*, *Journal of High Energy Physics* **07** (2002) 034, [[hep-ph/0206021](#)].
- [75] B. Bellazzini, C. Cski and J. Serra, *Composite Higgses*, *European Physical Journal C* **C74** (2014) 2766, [[1401.2457](#)].
- [76] R. Rattazzi, V. S. Rychkov, E. Tonni and A. Vichi, *Bounding scalar operator dimensions in 4D CFT*, *Journal of High Energy Physics* **12** (2008) 031, [[0807.0004](#)].
- [77] T. Appelquist and L. C. R. Wijewardhana, *Chiral hierarchies from slowly running couplings in technicolor theories*, *Physical Review D* **36** (1987) 568–580.
- [78] D. B. Kaplan, *Flavor at SSC energies: a new mechanism for dynamically generated fermion masses*, *Nuclear Physics B* **365** (1991) 259–278.
- [79] R. Contino and A. Pomarol, *Holography for fermions*, *Journal of High Energy Physics* **11** (2004) 058, [[hep-th/0406257](#)].
- [80] K. Agashe, R. Contino, L. Da Rold and A. Pomarol, *A custodial symmetry for $Zb\bar{b}$* , *Physics Letters* **B641** (2006) 62–66, [[hep-ph/0605341](#)].
- [81] S. Coleman and E. Weinberg, *Radiative corrections as the origin of spontaneous symmetry breaking*, *Physical Review D* **7** (1973) 1888–1910.
- [82] R. Jackiw, *Functional evaluation of the effective potential*, *Physical Review D* **9** (1974) 1686–1701.
- [83] S. Bhattacharjee, *Quantum Infrared Instabilities of gauge and gravity coupled Higgs Fields*, Ph.D. thesis, Homi Bhabha National Institute, 2013.
- [84] R. Barbieri and G. F. Giudice, *Upper Bounds on Supersymmetric Particle Masses*, *Nuclear Physics* **B306** (1988) 63–76.
- [85] G. Panico and A. Wulzer, *The Discrete Composite Higgs Model*, *JHEP* **09** (2011) 135, [[1106.2719](#)].
- [86] O. Matsedonskyi, G. Panico and A. Wulzer, *Light Top Partners for a Light Composite Higgs*, *Journal of High Energy Physics* **01** (2013) 164, [[1204.6333](#)].
- [87] G. Panico, M. Redi, A. Tesi and A. Wulzer, *On the Tuning and the Mass of the Composite Higgs*, *Journal of High Energy Physics* **03** (2013) 051, [[1210.7114](#)].

-
- [88] J. Barnard and M. White, *Collider constraints on tuning in composite Higgs models*, *Journal of High Energy Physics* **10** (2015) 072, [[1507.02332](#)].
- [89] M. Bando, T. Kugo and K. Yamawaki, *Nonlinear realization and hidden local symmetries*, *Physics Reports* **164** (1988) 217 – 314.
- [90] M. Bando, T. Kugo, S. Uehara, K. Yamawaki and T. Yanagida, *Is the ρ Meson a Dynamical Gauge Boson of Hidden Local Symmetry?*, *Physical Review Letters* **54** (1985) 1215–1218.
- [91] M. Bando, T. Fujiwara and K. Yamawaki, *Generalized Hidden Local Symmetry and the A_1 Meson*, *Progress of Theoretical Physics* **79** (1988) 1140–1166.
- [92] G. Ecker, J. Gasser, H. Leutwyler, A. Pich and E. de Rafael, *Chiral Lagrangians for Massive Spin 1 Fields*, *Physics Letters* **B223** (1989) 425–432.
- [93] N. Arkani-Hamed, A. G. Cohen and H. Georgi, *(De)constructing dimensions*, *Physical Review Letters* **86** (2001) 4757–4761, [[hep-th/0104005](#)].
- [94] R. Contino, T. Kramer, M. Son and R. Sundrum, *Warped/Composite Phenomenology Simplified*, *Journal of High Energy Physics* **05** (2007) 074, [[hep-ph/0612180](#)].
- [95] D. Becciolini, M. Redi and A. Wulzer, *AdS/QCD: The Relevance of the Geometry*, *Journal of High Energy Physics* **01** (2010) 074, [[0906.4562](#)].
- [96] Y. Kahn and J. Thaler, *Locality in Theory Space*, *Journal of High Energy Physics* **07** (2012) 007, [[1202.5491](#)].
- [97] R. Brauer and H. Weyl, *Spinors in n Dimensions*, *American Journal of Mathematics* **57** (1935) 425–449.
- [98] D. Marzocca, M. Serone and J. Shu, *General Composite Higgs Models*, *Journal of High Energy Physics* **08** (2012) 013, [[1205.0770](#)].
- [99] S. De Curtis, M. Redi and A. Tesi, *The 4D Composite Higgs*, *Journal of High Energy Physics* **04** (2012) 042, [[1110.1613](#)].
- [100] M. Carena, L. Da Rold and E. Pontón, *Minimal Composite Higgs Models at the LHC*, *Journal of High Energy Physics* **06** (2014) 159, [[1402.2987](#)].
- [101] C. Niehoff, P. Stangl and D. M. Straub, *Direct and indirect signals of natural composite Higgs models*, *Journal of High Energy Physics* **01** (2016) 119, [[1508.00569](#)].
- [102] D. Murnane, *The Landscape of Composite Higgs Models*, Ph.D. thesis, University of Adelaide, 2019.

-
- [103] H. Jeffreys, *The Theory of Probability*. Oxford Classic Texts in the Physical Sciences. Oxford University Press, 1998.
- [104] D. J. C. MacKay, *Bayesian Methods for Adaptive Models*, Ph.D. thesis, California Institute of Technology, 1992.
- [105] S. Kullback and R. A. Leibler, *On Information and Sufficiency*, *Annals of Mathematical Statistics* **22** (1951) 79–86.
- [106] C. Heymans, T. Tröster, M. Asgari, C. Blake, H. Hildebrandt, B. Joachimi et al., *KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints*, [2007.15632](#).
- [107] L. T. Hergt, W. J. Handley, M. P. Hobson and A. N. Lasenby, *Bayesian evidence for the tensor-to-scalar ratio r and neutrino masses m_ν : Effects of uniform vs logarithmic priors*, (Awaiting submission) .
- [108] J. Skilling, *Nested sampling for general Bayesian computation*, *Bayesian Analysis* **1** (2006) 833–859.
- [109] W. J. Handley, M. P. Hobson and A. N. Lasenby, *PolyChord: next-generation nested sampling*, *Monthly Notices of the Royal Astronomical Society* **453** (2015) 4385–4399.
- [110] W. J. Handley, M. P. Hobson and A. N. Lasenby, *PolyChord: nested sampling for cosmology*, *Monthly Notices of the Royal Astronomical Society* **450** (2015) L61–L65, [[1502.01856](#)].
- [111] D. J. Venzon and S. H. Moolgavkar, *A Method for Computing Profile-Likelihood-Based Confidence Intervals*, *Journal of the Royal Statistical Society* **37** (1988) 87–94.
- [112] Y. Akrami, C. Savage, P. Scott, J. Conrad and J. Edsjo, *Statistical coverage for supersymmetric parameter estimation: a case study with direct detection of dark matter*, *Journal of Cosmology and Astroparticle Physics* **07** (2011) 002, [[1011.4297](#)].
- [113] C. Strege, R. Trotta, G. Bertone, A. H. Peter and P. Scott, *Fundamental statistical limitations of future dark matter direct detection experiments*, *Physical Review D* **86** (2012) 023507, [[1201.3631](#)].
- [114] PARTICLE DATA GROUP collaboration, *Review of particle physics*, *Physical Review D* **98** (2018) 030001.
- [115] J. Brest, S. Greiner, B. Boskovic, M. Mernik and V. Zumer, *Self-Adapting Control Parameters in Differential Evolution: A Comparative Study on Numerical Benchmark Problems*, *IEEE Transactions on Evolutionary Computation* **10** (2006) 646–657.

-
- [116] GAMBIT collaboration, G. D. Martinez, J. McKay, B. Farmer, P. Scott, E. Roebber, A. Putze et al., *Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module*, *European Physical Journal* **C77** (2017) 761, [[1705.07959](#)].
- [117] C. Niehoff, P. Stangl and D. M. Straub, *Electroweak symmetry breaking and collider signatures in the next-to-minimal composite Higgs model*, *Journal of High Energy Physics* **04** (2017) 117, [[1611.09356](#)].
- [118] C. Niehoff, *Indirect Constraints on Composite Higgs Models and Leptoquarks*, Ph.D. thesis, Technical University of Munich, 2017.
- [119] G. Eilam, J. L. Hewett and A. Soni, *Rare decays of the top quark in the standard and two Higgs doublet models*, *Physical Review* **D44** (1991) 1473–1484.
- [120] A. H. Hoang, *The Top Mass: Interpretation and Theoretical Uncertainties*, in *Proceedings, 7th International Workshop on Top Quark Physics (TOP2014)*, 2014, [1412.3649](#).
- [121] ALEPH, DELPHI, L3, OPAL, SLD, LEP ELECTROWEAK WORKING GROUP, SLD ELECTROWEAK GROUP, SLD HEAVY FLAVOUR GROUP collaboration, S. Schael et al., *Precision electroweak measurements on the Z resonance*, *Physics Reports* **427** (2006) 257–454, [[hep-ex/0509008](#)].
- [122] ATLAS, CMS collaboration, *Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV*, *Journal of High Energy Physics* **08** (2016) 045, [[1606.02266](#)].
- [123] CMS collaboration, *Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13$ TeV*, *European Physical Journal* **C79** (2019) 421, [[1809.10733](#)].
- [124] ATLAS collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb⁻¹ of proton–proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, Tech. Rep. ATLAS-CONF-2018-031, CERN, Geneva, 2018.
- [125] CMS collaboration, *Measurements of Higgs boson production via gluon fusion and vector boson fusion in the diphoton decay channel at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-HIG-18-029, Geneva, 2019.
- [126] K. G. Chetyrkin, J. H. Kuhn and M. Steinhauser, *RunDec: A Mathematica package for running and decoupling of the strong coupling and quark masses*, *Computer Physics Communications* **133** (2000) 43–65, [[hep-ph/0004189](#)].

- [127] M. E. Peskin and T. Takeuchi, *Estimation of oblique electroweak corrections*, *Physical Review* **D46** (1992) 381–409.
- [128] R. Barbieri, A. Pomarol, R. Rattazzi and A. Strumia, *Electroweak symmetry breaking after LEP-1 and LEP-2*, *Nuclear Physics* **B703** (2004) 127–146, [[hep-ph/0405040](#)].
- [129] A. Freitas, *Higher-order electroweak corrections to the partial widths and branching ratios of the Z boson*, *Journal of High Energy Physics* **04** (2014) 070, [[1401.2447](#)].
- [130] J. D. Wells and Z. Zhang, *Precision Electroweak Analysis after the Higgs Boson Discovery*, *Physical Review* **D90** (2014) 033006, [[1406.6070](#)].
- [131] M. Aliev, H. Lacker, U. Langenfeld, S. Moch, P. Uwer and M. Wiedermann, *HATHOR: HAdronic Top and Heavy quarks crOss section calculator*, *Comput. Phys. Commun.* **182** (2011) 1034–1046, [[1007.1327](#)].
- [132] R. S. Chivukula, P. Ittisamai, K. Mohan and E. H. Simmons, *Simplified Limits on Resonances at the LHC*, *Physical Review* **D94** (2016) 094029, [[1607.05525](#)].
- [133] C. F. Uhlemann and N. Kauer, *Narrow-width approximation accuracy*, *Nuclear Physics* **B814** (2009) 195–211, [[0807.4112](#)].
- [134] W. Handley, *anesthetic: nested sampling visualisation*, *The Journal of Open Source Software* **4** (Jun, 2019) 1414.
- [135] P. Scott, *Pippi - painless parsing, post-processing and plotting of posterior and likelihood samples*, *European Physical Journal Plus* **127** (2012) 138, [[1206.2245](#)].
- [136] D. Murnane, M. J. White and A. G. Williams, *Exploring Fine-tuning of the Next-to-Minimal Composite Higgs Model*, *Journal of High Energy Physics* **04** (2019) 076, [[1810.08355](#)].
- [137] W. Handley and P. Lemos, *Quantifying dimensionality: Bayesian cosmological model complexities*, *Physical Review D* **100** (2019) 023512, [[1903.06682](#)].
- [138] G. F. Giudice, C. Grojean, A. Pomarol and R. Rattazzi, *The Strongly-Interacting Light Higgs*, *Journal of High Energy Physics* **06** (2007) 045, [[hep-ph/0703164](#)].
- [139] I. Low and A. Vichi, *On the production of a composite Higgs boson*, *Physical Review D* **84** (2011) 045019, [[1010.2753](#)].

-
- [140] C. Delaunay, C. Grojean and G. Perez, *Modified Higgs Physics from Composite Light Flavors*, *Journal of High Energy Physics* **09** 090, [[1303.5701](#)].
- [141] CMS collaboration, *Updates on Projections of Physics Reach with the Upgraded CMS Detector for High Luminosity LHC*, Tech. Rep. CMS-DP-2016-064, 2016.
- [142] ATLAS, CMS collaboration, A. A. Savin, *Prospects for Higgs and SM measurements at the HL-LHC*, *European Physical Journal Web of Conferences* **95** (2016) 04060.
- [143] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*. Addison-Wesley, Reading, USA, 1995.
- [144] R. A. Diaz and R. Martinez, *The Custodial Symmetry*, *Revista Mexicana Fisica* **47** (2001) 489–492, [[hep-ph/0302058](#)].
- [145] M. Montull, F. Riva, E. Salvioni and R. Torre, *Higgs Couplings in Composite Models*, *Physical Review* **D88** (2013) 095006, [[1308.0559](#)].
- [146] ATLAS collaboration, *ATLAS search for a heavy gauge boson decaying to a charged lepton and a neutrino in pp collisions at $\sqrt{s} = 7$ TeV*, *European Physical Journal C* **72** (2012) 2241, [[1209.4446](#)].
- [147] CMS collaboration, *Search for SSM W' production, in the lepton+MET final state at a center-of-mass energy of 13 TeV*, Tech. Rep. CMS-PAS-EXO-15-006, 2015.
- [148] ATLAS collaboration, *Search for new resonances decaying to a charged lepton and a neutrino in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2016-061, 2016.
- [149] ATLAS collaboration, *Search for a new heavy gauge boson resonance decaying into a lepton and missing transverse momentum in 79.8 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS experiment*, Tech. Rep. ATLAS-CONF-2018-017, 2018.
- [150] CMS collaboration, *Search for W' decaying to tau lepton and neutrino in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Physics Letters* **B755** (2016) 196–216, [[1508.04308](#)].
- [151] CMS collaboration, *Search for W' decaying to tau lepton and neutrino in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-EXO-16-006, 2016.
- [152] ATLAS collaboration, *Search for high-mass dilepton resonances in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Physical Review* **D90** (2014) 052005, [[1405.4123](#)].

-
- [153] CMS collaboration, *Search for physics beyond the standard model in dilepton mass spectra in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Journal of High Energy Physics* **04** (2015) 025, [[1412.6302](#)].
- [154] CMS collaboration, *Search for high mass resonances in dielectron final state*, Tech. Rep. CMS-PAS-EXO-18-006, 2018.
- [155] CMS collaboration, *Search for a high-mass resonance decaying into a dilepton final state in 13 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-EXO-16-031, 2016.
- [156] ATLAS collaboration, *Search for new high-mass resonances in the dilepton final state using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2016-045, 2016.
- [157] ATLAS collaboration, *Search for new high-mass phenomena in the dilepton final state using 36 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **10** (2017) 182, [[1707.02424](#)].
- [158] ATLAS collaboration, *Search for high-mass dilepton resonances using 139 fb^{-1} of pp collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Physics Letters* **B796** (2019) 68–87, [[1903.06248](#)].
- [159] CMS collaboration, *Search for high-mass resonances in dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Journal of High Energy Physics* **06** (2018) 120, [[1803.06292](#)].
- [160] ATLAS collaboration, *A search for high-mass resonances decaying to $\tau^+\tau^-$ in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **07** (2015) 157, [[1502.07177](#)].
- [161] CMS collaboration, *Z' to $\tau\tau - e\mu$ final state*, Tech. Rep. CMS-PAS-EXO-12-046, CERN, Geneva, 2015.
- [162] CMS collaboration, *Search for new physics with high-mass tau lepton pairs in pp collisions at $\sqrt{s} = 13$ TeV with the CMS detector*, Tech. Rep. CMS-PAS-EXO-16-008, 2016.
- [163] CMS collaboration, *Searches for narrow resonances decaying to dijets in proton-proton collisions at 13 TeV using 12.9 inverse femtobarns.*, Tech. Rep. CMS-PAS-EXO-16-032, 2016.
- [164] CMS collaboration, *Search for narrow resonances decaying to dijets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Physical Review Letters* **116** (2016) 071801, [[1512.01224](#)].

-
- [165] CMS collaboration, *Search for narrow and broad dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter mediators and other new particles*, *Journal of High Energy Physics* **08** (2018) 130, [[1806.00843](#)].
- [166] ATLAS collaboration, *Search for new phenomena in dijet mass and angular distributions from pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Physics Letters* **B754** (2016) 302–322, [[1512.01530](#)].
- [167] ATLAS collaboration, *Search for low-mass resonances decaying into two jets and produced in association with a photon using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Physics Letters* **B795** (2019) 56–75, [[1901.10917](#)].
- [168] CMS collaboration, *Search for $W' \rightarrow tb$ decays in the lepton + jets final state in pp collisions at $\sqrt{s} = 8$ TeV*, *Journal of High Energy Physics* **05** (2014) 108, [[1402.2176](#)].
- [169] CMS collaboration, *Search for $W' \rightarrow tb$ in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Journal of High Energy Physics* **02** (2016) 122, [[1509.06051](#)].
- [170] CMS collaboration, *Search for W' to tb in pp collisions at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-B2G-16-009, 2016.
- [171] CMS collaboration, *Search for W' boson resonances decaying into a top quark and a bottom quark in the leptonic final state at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-B2G-16-017, 2016.
- [172] ATLAS collaboration, *Search for $W' \rightarrow tb$ decays in the hadronic final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Physics Letters* **B781** (2018) 327–348, [[1801.07893](#)].
- [173] ATLAS collaboration, *A search for $t\bar{t}$ resonances using lepton plus jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2015-009, 2015.
- [174] CMS collaboration, *Search for resonant $t\bar{t}$ production in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Physical Review* **D93** (2016) 012001, [[1506.03062](#)].
- [175] CMS collaboration, *Search for top quark-antiquark resonances in the all-hadronic final state at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-B2G-15-003, 2016.
- [176] CMS collaboration, *Search for $t\bar{t}$ resonances in boosted semileptonic final states in pp collisions at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-B2G-15-002, 2016.

-
- [177] CMS collaboration, *Search for resonant $t\bar{t}$ production in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Journal of High Energy Physics* **04** (2019) 031, [[1810.05905](#)].
- [178] ATLAS collaboration, *Search for heavy particles decaying into top-quark pairs using lepton-plus-jets events in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *European Physical Journal* **C78** (2018) 565, [[1804.10823](#)].
- [179] ATLAS collaboration, *Search for heavy particles decaying into a top-quark pair in the fully hadronic final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Physical Review* **D99** (2019) 092004, [[1902.10077](#)].
- [180] ATLAS collaboration, *Search for a new resonance decaying to a W or Z boson and a Higgs boson in the $\ell\ell/\ell\nu/\nu\nu + b\bar{b}$ final states with the ATLAS detector*, *European Physical Journal* **C75** (2015) 263, [[1503.08089](#)].
- [181] CMS collaboration, *Search for Narrow High-Mass Resonances in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV Decaying to a Z and a Higgs Boson*, *Physics Letters* **B748** (2015) 255–277, [[1502.04994](#)].
- [182] CMS collaboration, *Search for heavy resonances decaying into a vector boson and a Higgs boson in the $(\ell\ell, \ell\nu, \nu\nu)$ $b\bar{b}$ final state*, Tech. Rep. CMS-PAS-B2G-16-003, 2016.
- [183] CMS collaboration, *Search for heavy resonances decaying into two Higgs bosons or into a Higgs boson and a W or Z boson in proton-proton collisions at 13 TeV*, *Journal of High Energy Physics* **01** (2019) 051, [[1808.01365](#)].
- [184] CMS collaboration, *Search for heavy resonances that decay into a vector boson and a Higgs boson in hadronic final states at $\sqrt{s} = 13$ TeV*, *European Physical Journal* **C77** (2017) 636, [[1707.01303](#)].
- [185] CMS collaboration, *Search for heavy resonances decaying into a vector boson and a Higgs boson in final states with charged leptons, neutrinos and b quarks at $\sqrt{s} = 13$ TeV*, *Journal of High Energy Physics* **11** (2018) 172, [[1807.02826](#)].
- [186] ATLAS collaboration, *Search for new resonances decaying to a W or Z boson and a Higgs boson in the $\ell^+\ell^-\bar{b}\bar{b}$, $\ell\nu\bar{b}\bar{b}$, and $\nu\bar{\nu}\bar{b}\bar{b}$ channels with pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Physics Letters* **B765** (2017) 32–52, [[1607.05621](#)].
- [187] ATLAS collaboration, *Search for new resonances decaying to a W or Z boson and a Higgs boson in the $\ell\bar{b}\bar{b}$, $\ell\nu\bar{b}\bar{b}$, and $\nu\bar{\nu}\bar{b}\bar{b}$ channels in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2015-074, 2015.

-
- [188] ATLAS collaboration, *A Search for Resonances Decaying to a W or Z Boson and a Higgs Boson in the $q\bar{q}^{(\prime)}b\bar{b}$ Final State*, Tech. Rep. ATLAS-CONF-2016-083, 2016.
- [189] ATLAS collaboration, *Search for high-mass diboson resonances with boson-tagged jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **12** (2015) 055, [[1506.00962](#)].
- [190] ATLAS collaboration, *Search for production of WW/WZ resonances decaying to a lepton, neutrino and jets in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *European Physical Journal* **C75** (2015) 209, [[1503.04677](#)].
- [191] ATLAS collaboration, *Search for WZ resonances in the fully leptonic channel using pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Physics Letters* **B737** (2014) 223–243, [[1406.4456](#)].
- [192] CMS collaboration, *Search for massive resonances in dijet systems containing jets tagged as W or Z boson decays in pp collisions at $\sqrt{s} = 8$ TeV*, *Journal of High Energy Physics* **08** (2014) 173, [[1405.1994](#)].
- [193] CMS collaboration, *Search for massive resonances decaying into pairs of boosted W and Z bosons at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-EXO-15-002, 2015.
- [194] CMS collaboration, *Search for new resonances decaying to WW/WZ $\rightarrow \ell\nu qq$* , Tech. Rep. CMS-PAS-B2G-16-020, 2016.
- [195] CMS collaboration, *Search for a heavy resonance decaying to a pair of vector bosons in the lepton plus merged jet final state at $\sqrt{s} = 13$ TeV*, *Journal of High Energy Physics* **05** (2018) 088, [[1802.09407](#)].
- [196] CMS collaboration, *Search for massive resonances decaying into WW, WZ, ZZ, qW, and qZ with dijet final states at $\sqrt{s} = 13$ TeV*, *Physical Review* **D97** (2018) 072006, [[1708.05379](#)].
- [197] CMS collaboration, *Search for a heavy resonance decaying into a Z boson and a vector boson in the $\nu\bar{\nu}q\bar{q}$ final state*, *Journal of High Energy Physics* **07** (2018) 075, [[1803.03838](#)].
- [198] CMS collaboration, *Search for a heavy resonance decaying into a Z boson and a Z or W boson in $2\ell 2q$ final states at $\sqrt{s} = 13$ TeV*, *Journal of High Energy Physics* **09** (2018) 101, [[1803.10093](#)].
- [199] CMS collaboration, *A multi-dimensional search for new heavy resonances decaying to boosted WW, WZ, or ZZ boson pairs in the dijet final state at 13 TeV*, *European Physical Journal C* **80** (2019) 237, [[1906.05977](#)].

-
- [200] ATLAS collaboration, *Search for resonances with boson-tagged jets in 15.5 fb⁻¹ of pp collisions at $\sqrt{s} = 13$ TeV collected with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2016-055, 2016.
- [201] ATLAS collaboration, *Search for diboson resonance production in the $\ell\nu qq$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector at the LHC*, Tech. Rep. ATLAS-CONF-2016-062, 2016.
- [202] ATLAS collaboration, *Searches for heavy ZZ and ZW resonances in the ℓqq and $\nu\nu qq$ final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2016-082, 2016.
- [203] CMS collaboration, *Search for massive WH resonances decaying into the $\ell\nu b\bar{b}$ final state at $\sqrt{s} = 8$ TeV*, *European Physical Journal C* **76** (2016) 237, [[1601.06431](#)].
- [204] CMS collaboration, *Search for massive resonances decaying into pairs of boosted bosons in semi-leptonic final states at $\sqrt{s} = 8$ TeV*, *Journal of High Energy Physics* **08** (2014) 174, [[1405.3447](#)].
- [205] CMS collaboration, *Combination of diboson resonance searches at 8 and 13 TeV*, Tech. Rep. CMS-PAS-B2G-16-007, 2016.
- [206] CDF collaboration, *Search for Heavy Top $t' \rightarrow Wq$ in Lepton Plus Jets Events in $\int \mathcal{L} dt = 4.6$ fb⁻¹*, Tech. Rep. 10110, 2010.
- [207] ATLAS collaboration, *Search for pair-produced heavy quarks decaying to Wq in the two-lepton channel at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Physical Review D* **86** (2012) 012007, [[1202.3389](#)].
- [208] ATLAS collaboration, *Search for pair production of a new heavy quark that decays into a W boson and a light quark in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Physical Review D* **92** (2015) 112007, [[1509.04261](#)].
- [209] CMS collaboration, *Search for vector-like quarks in final states with a single lepton and jets in pp collisions at $\sqrt{s} = 8$ TeV*, Tech. Rep. CMS-PAS-B2G-12-017, 2014.
- [210] ATLAS collaboration, *Search for pair production of heavy top-like quarks decaying to a high-pT W boson and a b quark in the lepton plus jets final state at $\sqrt{s}=7$ TeV with the ATLAS detector*, *Physics Letters B* **718** (2013) 1284–1302, [[1210.5468](#)].
- [211] CMS collaboration, *Search for heavy, top-like quark pair production in the dilepton final state in pp collisions at $\sqrt{s} = 7$ TeV*, *Physics Letters B* **716** (2012) 103–121, [[1203.5410](#)].

-
- [212] CMS collaboration, *Search for Pair Produced Fourth-Generation Up-Type Quarks in pp Collisions at $\sqrt{s} = 7$ TeV with a Lepton in the Final State*, *Physics Letters* **B718** (2012) 307–328, [[1209.0471](#)].
- [213] ATLAS collaboration, *Search for production of vector-like quark pairs and of four top quarks in the lepton plus jets final state in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2015-012, 2015.
- [214] CMS collaboration, *Search for vector-like charge 2/3 T quarks in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Physical Review* **D93** (2016) 012003, [[1509.04177](#)].
- [215] ATLAS collaboration, *Search for pair production of heavy vector-like quarks decaying to high- p_T W bosons and b quarks in the lepton-plus-jets final state in pp collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2016-102, 2016.
- [216] CMS collaboration, *Search for pair production of vector-like quarks in the $bW\bar{b}W$ channel from proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Physics Letters* **B779** (2018) 82–106, [[1710.01539](#)].
- [217] CMS collaboration, *Search for pair production of vector-like T and B quarks in single-lepton final states using boosted jet substructure in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Journal of High Energy Physics* **11** (2017) 085, [[1706.03408](#)].
- [218] CDF collaboration, *Search for New Bottomlike Quark Pair Decays $Q\bar{Q} \rightarrow (tW^{mp})(\bar{t}W^{\pm})$ in Same-Charge Dilepton Events*, *Physical Review Letters* **104** (2010) 091801, [[0912.1057](#)].
- [219] CMS collaboration, *Search for Heavy Quarks Decaying into a Top Quark and a W or Z Boson using Lepton + Jets Events in pp Collisions at $\sqrt{s} = 7$ TeV*, *Journal of High Energy Physics* **01** (2013) 154, [[1210.7471](#)].
- [220] CMS collaboration, *Search for top-quark partners with charge 5/3 in the same-sign dilepton final state*, *Physical Review Letters* **112** (2014) 171801, [[1312.2391](#)].
- [221] CMS collaboration, *Search for Vector-Like b' Pair Production with Multilepton Final States in pp collisions at $\sqrt{s} = 8$ TeV*, Tech. Rep. CMS-PAS-B2G-13-003, 2013.
- [222] CMS collaboration, *Search for pair-produced vectorlike B quarks in proton-proton collisions at $\sqrt{s}=8$ TeV*, *Physical Review* **D93** (2016) 112009, [[1507.07129](#)].

-
- [223] ATLAS collaboration, *Analysis of events with b-jets and a pair of leptons of the same charge in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **10** (2015) 150, [[1504.04605](#)].
- [224] ATLAS collaboration, *Search for vector-like B quarks in events with one isolated lepton, missing transverse momentum and jets at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Physical Review* **D91** (2015) 112011, [[1503.05425](#)].
- [225] ATLAS collaboration, *Search for new phenomena in events with same-charge leptons and b-jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **12** (2018) 039, [[1807.11883](#)].
- [226] ATLAS collaboration, *Search for pair production of heavy vector-like quarks decaying into high- p_T W bosons and top quarks in the lepton-plus-jets final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **08** (2018) 048, [[1806.01762](#)].
- [227] CMS collaboration, *Search for top quark partners with charge 5/3 at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-B2G-15-006, 2015.
- [228] CMS collaboration, *Search for heavy vector-like quarks decaying to same-sign dileptons*, Tech. Rep. CMS-PAS-B2G-16-019, 2017.
- [229] CMS collaboration, *Search for top quark partners with charge 5/3 in the same-sign dilepton and single-lepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Journal of High Energy Physics* **03** (2019) 082, [[1810.03188](#)].
- [230] CDF collaboration, *Search for New Particles Leading to Z+ jets Final States in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ -TeV*, *Physical Review* **D76** (2007) 072006, [[0706.3264](#)].
- [231] CMS collaboration, *Search B' to bZ*, Tech. Rep. CMS-PAS-EXO-11-066, 2012.
- [232] ATLAS collaboration, *Search for pair- and single-production of vector-like quarks in final states with at least one Z boson decaying into a pair of electrons or muons in pp collision data collected with the ATLAS detector at $\sqrt{s} = 13$ TeV*, *Physical Review* **D98** (2018) 112010, [[1806.10555](#)].
- [233] CMS collaboration, *Search for a Vector-like Quark with Charge 2/3 in t + Z Events from pp Collisions at $\sqrt{s} = 7$ TeV*, *Physical Review Letters* **107** (2011) 271802, [[1109.4985](#)].
- [234] ATLAS collaboration, *Search for pair production of vector-like top partners in events with exactly one lepton and large missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2016-101, 2016.

-
- [235] ATLAS collaboration, *Search for pair production of vector-like top quarks in events with one lepton, jets, and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Journal of High Energy Physics* **08** (2017) 052, [[1705.10751](#)].
- [236] ATLAS collaboration, *Search for pair production of up-type vector-like quarks and for four-top-quark events in final states with multiple b-jets with the ATLAS detector*, *Journal of High Energy Physics* **07** (2018) 089, [[1803.09678](#)].
- [237] CMS collaboration, *Search for pair-produced vector-like quarks of charge $-1/3$ in lepton+jets final state in pp collisions at $\sqrt{s} = 8$ TeV*, Tech. Rep. CMS-PAS-B2G-12-019, 2012.
- [238] CMS collaboration, *Search for pair-produced vector-like quarks of charge $-1/3$ decaying to bH using boosted Higgs jet-tagging in pp collisions at $\sqrt{s} = 8$ TeV*, Tech. Rep. CMS-PAS-B2G-14-001, 2014.
- [239] CMS collaboration, *Search for vector-like quark pair production in final states with leptons and boosted Higgs bosons at $\sqrt{s} = 13$ TeV*, Tech. Rep. CMS-PAS-B2G-16-011, 2016.
- [240] ATLAS collaboration, *Search for production of vector-like top quark pairs and of four top quarks in the lepton-plus-jets final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2016-013, 2016.