

THE UNIVERSITY OF CHICAGO

LEPTON PRECISION EXPERIMENTS AS PROBES OF NEW PHYSICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
DUNCAN JAVIER MACKENZIE ROCHA

CHICAGO, ILLINOIS

AUGUST 2026

To my parents.

“Oh Grandma, you wouldn’t understand.”

– Barbara Mackenzie

TABLE OF CONTENTS

LIST OF FIGURES	viii
LIST OF TABLES	xiii
ACKNOWLEDGMENTS	xv
ABSTRACT	xvi
1 INTRODUCTION	1
2 REVIEW OF THE LEPTON SECTOR IN THE STANDARD MODEL	4
2.1 The Higgs and Electroweak Symmetry Breaking	4
2.2 The Leptons	5
2.3 Models of New Physics	7
3 LEPTON FLAVOR IN THE 2HDM	11
3.1 Introduction	11
3.2 Two Higgs Doublet Models	13
3.2.1 The Higgs Basis and Electroweak Symmetry Breaking	14
3.2.2 2HDM Yukawas	17
3.3 Lepton Flavor	19
3.3.1 Flavor Considerations in the 2HDM	20
3.3.2 Controlling $\mathcal{O}(1)$ Coefficients	22
3.3.3 The Inheritance of Flavor Structure in the 2HDM	25
3.4 Dipoles in the Low Energy EFT	28
3.4.1 Blind Spots in the Dipole Operator $\mu - e$ Block	31
3.5 Experimental Constraints	33
3.5.1 Electric Dipole Moment of the Electron	33
3.5.2 Flavor Violating μ Decays	34
3.5.3 Flavor Violating Higgs Decays at the LHC	36
3.5.4 Higgs Diagonal Couplings	37
3.5.5 Heavy Higgs Searches	38
3.6 Phenomenology of the Generic (Type-III) 2HDM	40
3.6.1 Cheng-Sher and Dipole Blind Spots	40
3.6.2 $\mu\tau$ -Specific Scenario	43
3.7 Phenomenology of the Type-II and Lepton-Specific 2HDMs	44
3.7.1 The MSSM	45
3.7.2 The Dark-side of κ_τ	46
3.7.3 The $\mu - \tau$ Specific Flavor Structure	49
3.8 Conclusion	49

4	PRECISION HIGGS MEASUREMENTS TO SEARCH FOR NEW ELECTROWEAK-SCALE PHYSICS	51
4.1	Introduction	51
4.2	Enhancing the electron Yukawa in SMEFT	52
4.2.1	Loop level connections to other SMEFT operators	53
4.2.2	Generating $\mathcal{O}_{eH}$ radiatively	59
4.2.3	Constraints on the flavour structure of $\mathcal{O}_{eH}$	64
4.3	Selected extensions which enhance the electron Yukawa	71
4.3.1	Model-dependent connections between Yukawas and dipoles	71
4.3.2	A heavy second Higgs doublet	75
4.3.3	Leptoquarks	78
4.3.4	Comments on fine-tuning	83
4.4	Discussion	85
5	NEUTRINO EXPERIMENTS AND LONG-BASELINE SIGNATURES OF HIDDEN SECTORS	87
5.1	Model Overview	88
5.2	Analysis	90
5.3	Results	96
5.4	Discussion and Outlook	98
6	NEW PHYSICS PROBES IN PRECISION MUON EXPERIMENTS	100
6.1	Simple Models of Vector Boson Mediators	101
6.2	Vector Boson Signatures at the <i>MUonE</i> Experiment	104
6.2.1	The Apparatus	104
6.2.2	Search Strategy	106
6.2.3	Signal Event Generation	108
6.2.4	Backgrounds	112
6.3	Analysis and Results	113
6.3.1	Dark photon model	114
6.3.2	The gauged $U(1)_{B-L}$ model	116
6.3.3	The gauged $U(1)_{L_e-L_\mu}$ model	118
6.4	Discussion	118
7	CONCLUSION	120
	APPENDIX A GENERATING YUKAWAS AT LOOP LEVEL IN THE MSSM	122
	APPENDIX B LOOP FUNCTIONS FOR THE DIPOLE OPERATORS	124
	APPENDIX C LOOP CORRECTIONS TO THE DARK-SIDE	127

APPENDIX D	THEORETICAL EXPRESSIONS FOR MESON DECAYS	128
D.1	$B_s \rightarrow ee$	128
D.2	$D_s \rightarrow e\nu$	130
APPENDIX E	$\mu \rightarrow e$ CONVERSION	132
APPENDIX F	PERTURBATIVITY BOUNDS	134
APPENDIX G	MORE DETAILS ON SCALAR LEPTOQUARKS	135
G.1	Matching expressions for scalar leptoquarks	135
G.2	Expressions for $\Delta\kappa_e$ and Δa_e	136
G.3	Top FCNCs	136
REFERENCES		139

LIST OF FIGURES

- 3.1 Diagrams which give the largest contribution to the dipole coefficients $L_{L/R}$. There are several additional subleading 2-loop diagrams, which are included in our analysis and found in [1]. Furthermore, the one-loop contribution can be relevant for τ decays, but these constraints are subdominant to LHC searches. 29
- 3.2 Plots of the 1 and 2 Loop contributions to $[L_R]_{e\mu}$. For these plots we take $[\overline{\mathcal{Y}}_2^f]_{ii} = m_i^f/v$ and $[\overline{\mathcal{Y}}_2^\ell]_{i\neq j} = \sqrt{m_i m_j}/v$. For the $L_R^{H^\pm\gamma}, L_R^{H^\pm Z}, L_R^{H^\pm Z}$ diagrams we also take $Z_3 = Z_7 = 1$ as a benchmark for evaluating their magnitudes. The full expressions for these diagrams is given in Ref. [1], with a conversion factor of $L_R^i = em_\mu/(16\pi^2 v^2)c_R^i$. Note that the three dominant contributions are those listed in further detail in the text: $L_R^{th}, L_R^{tH^\pm}$, and $L_R^{h,kin}$. Note as well that the signs of the amplitudes are depicted by solid/dashed styles - and a cancellation is occurring in certain ranges between the three dominant diagrams. 30
- 3.3 Contributions to the $\mu \rightarrow e$ conversion rate. The blob on the right side represents the sum of two loop diagrams which contribute to the effective dipole coefficients $L_{L/R}$ (Sec. 3.4). There are additional two loop contributions, such as box diagrams, whose analytic representations are not yet available in the literature and are therefore not included in this analysis. 35
- 3.4 Production modes for $H/A/H^\pm$ 38
- 3.5 A characterization of the $\mu \rightarrow e\gamma$ “Blind Spot,” where the loop contributions of the 2HDM to the dipole operators $L_{L/R}$ experience destructive interference. The dashed line is the analytical formula in Eq. 3.49. these plots, we assume the relation $\cos(\beta - \alpha) = Z_6 v^2/(m_H^2 - m_h^2)$. The dark blue excluded region assumes $\overline{\mathcal{Y}}_2 \sim \sqrt{m_i m_j}/v$, which is the Cheng-Sher [CS] structure. The purple constraint is from Heavy Higgs searches at the LHC with a top pair final state, and the gray is the constraints are from κ_t and κ_W measurements at the LHC. The $\mu \rightarrow e\gamma$ constraint is from current MEG-II results. 42
- 3.6 The parameter space for two case of the Type-II 2HDM, in the Cheng-Sher flavor scenario. In the left, the MSSM relation between $\cos(\beta - \alpha)$ and $\tan \beta$ is applied, resulting in a strong κ_τ bound. In the right, $\cos(\beta - \alpha) = 0$, and κ_τ (as well as certain heavy Higgs searches) does not apply. Future bounds from the HL-LHC projections, MEG-II, and Mu2e. Conservatively, we do not project the $H \rightarrow t\bar{t}$ since it is an interference search and the projections are nontrivial. 46
- 3.7 The current (left) and future (right) constraints for the “darkside” ($\kappa_\tau \approx -1$) scenario in a Type-II 2HDM. The dipole constraints disfavors larger heavy Higgs masses in the darkside. The collider searches $A \rightarrow Zh$ and $H \rightarrow \tau\tau$ complement the dipole searches, and jointly they disfavor much of the darkside parameter space. Furthermore, the $\mu \rightarrow e$ conversion bound is enhanced in the Type-II 2HDM due to the enhanced down-type quark couplings. The future bounds, in dashed lines, depicted on the right, are using projections for the HL-LHC, MEG-II, and Mu2e experiments. 47

3.8	The current (left) and future (right) constraints for the “darkside” ($\kappa_\tau \approx -1$) scenario in a Lepton-Specific (Type-X) 2HDM. The blind spot now extends into very large $\tan\beta$, because there is no enhancement of the b -quark Barr-Zee diagram. This region will be explored from above by Mu2e, and from below by the LHC $pp \rightarrow A \rightarrow Zh$ experiment. The future bounds (right side) are using projections for the HL-LHC, MEG-II, and Mu2e experiments.	48
4.1	An example two-loop Barr-Zee diagram contributing to Δa_ℓ . Crosses on scalar legs indicate Higgs components set to v	56
4.2	Current and projected upper limits on $ \Delta\kappa_e $ (defined as $ \kappa_e - 1 $) at 95% CL, assuming coefficients are generated at 2 TeV and only one SMEFT coefficient is non-zero at a time. Current LHC bounds on $[\mathcal{C}_{lequ}^{(1)}]_{1133}$ and $[\mathcal{C}_{lequ}^{(3)}]_{1133}$ are taken from [2], while those on $[\mathcal{C}_{lequ}^{(1)}]_{1122}$, $[\mathcal{C}_{lequ}^{(3)}]_{1122}$ and $[\mathcal{C}_{ledq}]_{1133}$ are from [3]. All HL-LHC and FCC-ee projections, shown by green bars, are taken from [4]. Constraints from the decays $D_s^+ \rightarrow e^+\nu$ and $B_s \rightarrow e^+e^-$ follow from App. D, where the hatching on the latter bar indicates the assumption that the coefficient is in the up mass basis. Future Δa_e sensitivity is assumed to reach $\Delta a_e^{\text{future}} < 5 \times 10^{-14}$	62
4.3	Schematic diagrams showing how the EW dipole operators O_{eW} and O_{eB} are generically generated at one loop order higher than the O_{eH} operator in UV models. The grey blob represents a diagram of arbitrary loop order involving exchange of heavy UV states. The left hand diagram matches to the dimension-six O_{eH} SMEFT operator, while the right hand diagram matches to the dimension-six dipole operators. The gauge boson line connects to any charged particle in the diagram, including the H loop and any charged particles within the grey blob.	72
4.4	Parameter space for a φ extension coupled to electrons. Blue shaded regions are consistent with the indicated Δa_e measurement at 95% C.L. Dashed lines show the κ_e enhancement corresponding to the given parameter values. Solid (dotted) lines represent constraints on the couplings from current (projected) measurements at 95% C.L. The R_e constraints are taken from [4], $\kappa_\lambda^{\text{LHC}}$ and $\kappa_\lambda^{\text{FCCee}}$ are from [5] and $\kappa_\lambda^{\text{FCChh}}$ assumes 6% uncertainty at 95% CL [6, 7].	77
4.5	Parameter space for a second Higgs doublet φ with Yukawa-like couplings to electrons and top quarks, and assuming $\lambda_\varphi = 1$. Shaded regions and line styles have the same meaning as in Fig. 4.4. The κ_t^{LHC} limit is from [8], and the κ_t^{FCCee} projection is from [9].	77
4.6	95% C.L constraints on $\Delta\kappa_e$ implied by current measurements of Δa_e , assuming a UV completion of either ω_1 or Π_7 with only electron-top couplings.	80

4.7	Parameter space for ω_1 scalar leptoquarks with both electron-top and electron-charm couplings assuming $\lambda_{\omega_1} = 0$. Red and blue shaded regions are compatible with the stated $\Delta\kappa_e$ enhancement and measurement of Δa_e respectively. The green dashed line shows the 95% upper limit on the electron-charm couplings derived from current constraints on $\text{BR}(D_s^+ \rightarrow e^+\nu)$. The darker blue region is not clearly visible in the left hand plot, but exists as a very narrow band at the top of the lighter blue band.	81
4.8	Parameter space for ω_1 (a) and Π_7 (b) leptoquarks assuming $M_{\omega_1/\Pi_7} = 2$ TeV and only $(y_{\omega_1}^{eu})_{13}(y_{\omega_1}^{ql})_{31}^*$ and λ_{ω_1} or $(y_{\Pi_7}^{eq})_{13}(y_{\Pi_7}^{lu})_{13}^*$ and λ_{Π_7} are non-zero.	83
5.1	Schematic cartoon of our MiniBooNE and MicroBooNE signal processes. Top: the 8 GeV proton beam impinges on a fixed Be target and produces mesons including charged kaons, whose decays yield a flux of muons and associated scalars S through $\pi^\pm, K^\pm \rightarrow \mu^\pm \nu_\mu S$ decays. Bottom: S may also be created via a spray of pions/kaons whose decays yield a secondary muon beam in the decay pipe. The muons then interact with the beam dump, where they produce long-lived muon-philic scalar particles S through bremsstrahlung-like processes. Note that the meson decay mode is independent of the beam dump, while the muon bremsstrahlung mode requires scattering in the beam dump. In both cases, for $m_S < 2m_\mu$, the scalars eventually decay to diphotons in the downstream detector. For details of the experimental setup and detector performance, see Ref. [10].	88
5.2	Left: Feynman diagram for scalar production from $K, \pi \rightarrow \mu\nu S$ decays. Middle: Scalar production via μ - N scattering in the beam dump. Right: Scalar decay through a muon loop using only the coupling in Eq. (5.1).	89
5.3	Heatmaps (arbitrary units) of the processes in the beam dump that produce muon-philic scalars, for $r_\perp < 0.9$ m, where $r_\perp$ is the cylindrical radius in the $x - y$ plane. This displays the scalar production inside the decay pipe, however the analysis considers production events that occur outside the pipe as well. The top (middle) diagram shows the vertices where the $\pi^\pm$ ($K^\pm$) decays, which may contain scalars in their decay products. The lower diagram shows the spatial distribution of scalar emission from muons scattering in the beam dump, specifically for $m_S = 100$ MeV, although this distribution does not vary significantly over the scalar mass. The plot has several features: the target can be seen on the far left, with the electromagnet adjacent, and the beam dump is visible on the right.	90
5.4	Energy distribution of charged pions (dark green), charged kaons (dark blue), and secondary muons (dark red) produced from proton-Be collisions at 8 GeV. Also shown are energy distributions for scalars produced from pion decays (light green), kaon decays (light blue) and from secondary muon scattering in the beam dump (orange). The scalar distributions are normalized to $y = 1$ to enhance visibility on the plot, but in our analysis we explore the region $y \in [10^{-5}, 10^{-2}]$. These curves were generated with a GEANT simulation utilizing the setup described in Refs. [10, 11].	91

5.5	Example signal distributions generated by scalar decays in the MiniBooNE and MicroBooNE detectors. The data points are measured $\gamma\gamma$ -like events, and the background prediction arises primarily from neutrino-induced $\pi^0 \rightarrow \gamma\gamma$ events, and digitized from [12, 13]. The MiniBooNE error bars include statistical and systematic uncertainties. The MicroBooNE error bars are statistical only. These particular signal distributions correspond to 20σ (top panel) and 6σ (bottom panel) discrepancies with the data.	92
5.6	Parameter space for the singlet scalar in Eq. (5.1), decaying to diphotons via Eq. (5.3). The red and purple shaded regions are excluded at 5σ by existing MiniBooNE and MicroBooNE data, respectively, based on the analysis in this paper and the magenta dashed curve represents a future 2σ projection for SBND sensitivity. The green band represents the parameter space where this scenario resolves the muon $g - 2$ anomaly. For $S > 2m_\mu$ we show constraints from BaBar and CMS based on the $S \rightarrow \mu^+\mu^-$ decay channel along with future projections for Belle-II [14]. For $S < 2m_\mu$ we also show projections for a dedicated muon beam-dump experiment [15] and for the M^3 muon missing momentum proposal [16]. Also shown are cosmological bounds for $m_S < 1$ MeV where S particles thermalize with the SM as a relativistic species during Big Bang nucleosynthesis [17].	98
6.1	Feynman diagrams of vector boson production at MUonE.	103
6.2	A schematic diagram for the MUonE apparatus with displaced vertex signature, borrowed from Fig. [18]. Orange line: incoming muon beam and outgoing scattered muon. Blue dashed line: displaced decaying vector boson. Green lines: decay products of the vector boson. Black squares: tracking stations. Gray squares: Beryllium targets. Blue and dark blue squares: the ECAL and the muon filter.	105
6.3	Signal event geometry in the corresponding volume, borrowed from Fig. [18]. We show the ECAL and muon filter after the module, but note that this can be a few meters away from the module that we consider.	107
6.4	Production cross sections per unit charge, normalized to $g_V^2 = 1$ (Table. 6.1). At low masses, production from scattering off electrons is much more likely, as the electron is more likely to radiate out a gauge boson than the heavy nucleus, leading to a larger cross section. At large vector boson masses, production from nuclear scattering dominates, extending the experimental reach of MUonE to ~ 100 - 150 MeV.	109
6.5	Energy and angle distribution for the produced vector boson and deflected muon for $m_{A'} = 12$ MeV.	111
6.6	Energy and angle distribution for the produced vector boson and deflected muon for $m_{A'} = 44$ MeV.	112

6.7	The result for dark photon search. MUonE sensitivity is shown with the red dotted curve (for μe scattering), the red dashed curve (for μN scattering), and the red solid curve (combined). Existing limits are shown in gray shaded regions. Selected future limits that are planned for a similar timeline as MUonE are also shown. Specifically, we add an arrow for the SeaQuest projection to indicate that it is expected to probe the parameter space below the curve.	115
6.8	The result for gauged $U(1)_{B-L}$ search. MUonE sensitivity is shown with the red dotted curve (for μe scattering), the red dashed curve (for μN scattering), and the red solid curve (combined). Existing limits are shown in gray shaded regions. Selected future limits that are planned for a similar timeline as MUonE are also shown. Specifically, we add an arrow for the SeaQuest projection to indicate that it is expected to probe the parameter space below the curve.	116
6.9	The result for gauged $U(1)_{L_e-L_\mu}$ search. MUonE sensitivity is shown with the red dotted curve (for μe scattering), the red dashed curve (for μN scattering), and the red solid curve (combined). Existing limits are shown in gray shaded regions. Selected future limits that are planned for a similar timeline as MUonE are also shown. Specifically, we add an arrow for the SeaQuest projection to indicate that it is expected to probe the parameter space below the curve. . . .	117
A.1	Loop diagram that leads to flavor violation the MSSM.	123

LIST OF TABLES

3.1	The extra Yukawa couplings (to H_2 in the Higgs basis) for each discrete symmetry case. Here, L-S stands for Lepton-Specific.	27
3.2	Current and future constraints in muon flavor violating decays and electron EDM.	34
3.3	LHC Constraints on the light Higgs couplings. We only list those used in our analysis. Note that we use the shorthand notation $\bar{\mathcal{Y}}_{2\langle ij \rangle} \equiv \sqrt{ \bar{\mathcal{Y}}_2]_{ij} ^2 + \bar{\mathcal{Y}}_2]_{ji} ^2}$. The Higgs FV decay limit are shown at 90%, while the intervals on the κ quantities are 1σ	37
3.4	List of analyses used in the recasting of heavy Higgs constatints. All are with a luminosity of 139 fb^{-1}	40
4.1	68% confidence intervals for κ_e derived from current and future constraints on Δa_e under different assumptions for the new physics scale Λ	59
4.2	Current and future 95% CL expected limits on charged LFV Higgs decays.	65
4.3	Summary of limits on $\mathcal{C}_{eH}$ originating from searches for rare muon decays, as well as projections for future searches, all reported at 90% CL. For compactness, we denote the flavour average $C_{\langle e\mu \rangle} \equiv \sqrt{ C_{e\mu} ^2 + C_{\mu e} ^2}$	66
4.4	Current constraints and future sensitivity (all at 68% CL) on Higgs-charged lepton coupling modifiers. The FCC-ee projection assumes a Higgs-pole run and there is no projection for FCC-hh available. Where necessary 95% CL projections have been halved to estimate the 68% CL sensitivity.	69
4.5	Single particle states which match at tree or one loop level onto $\mathcal{C}_{eH}$, with coefficients unsuppressed by y_e (or by any other SM Yukawa couplings smaller than y_b). In all cases, $\mathcal{C}_{e\gamma}$ is generated via the same product of couplings as $\mathcal{C}_{eH}$, with this coupling dependence given in the last column (in the notation of [19]). States labelled in blue match to $\mathcal{C}_{eH}$ (and $\mathcal{C}_{e\gamma}$) via diagrams involving a non-renormalisable interaction, meaning that the UV completion of these operators would require additional BSM states. The asterisks on ‘tree*’ for the $\mathcal{L}_1$ state indicate that this state matches at tree level onto these operators, but that the couplings involved cannot arise within a weakly coupled UV completion (for example it is known that $\mathcal{C}_{e\gamma}$ cannot be generated at tree level within a weakly coupled UV completion [20, 21]).	73
4.6	Estimated constraints at 95% confidence on $\Delta\kappa_e$ from current and future Δa_e measurements, considering different scenarios for the relative loop suppression of $\mathcal{C}_{e\gamma}$ compared to $\mathcal{C}_{eH}$. We are estimating the future precision of Δa_e to be $\Delta a_e^{\text{future}} \sim 5 \times 10^{-14}$	74
6.1	The vector boson mediator considered in this work. Columns show the notation of the coupling g_V and the charge of each SM fermions, respectively.	103
6.2	Event selection criteria for our proposed search. In our numerical results, these requirements are imposed on our signal and background MC events. These criteria ensure that SM backgrounds are negligible for our signal of interest (see Supplemental Material for a discussion).	107

E.1	Wavefunction overlap integrals for muons captured in the two relevant nuclei.	133
-----	---	-----

ACKNOWLEDGMENTS

I'm so incredibly grateful to find myself surrounded by such wonderful family and friends. To my parents, Laurie and Rick, who supported and encouraged me when I first aspired to be a scientist, I owe an unimaginable debt. They continue to be my biggest sources of support and regardless of what's to come, they define a place that I will always call home. To my grandma Barbara, who makes me feel so valued and always tells her story of asking me to explain my research, I will be sending a copy of my thesis shortly. To my aunts, uncles, cousins, and the rest of my family in California and Texas, I'm proud of all of you and particularly proud to bring the second/first doctoral degree to our family.

It's hard to understand how I found such amazing friends at each of the three major stages of my life. Ele, Zander, Cora, you are my older siblings who have supported me even from afar, and after all this time, and I owe you so much. Viv, Adrian, Mikey, Martha, you became my family in college and been amazing friends ever since. And Juan Pablo (Dr. Parallax/Trash Boy), Jordan, Paola (RAHBAR SABA ETEZAD-RAZAVI), Saba (Paolita), George, Matty B, RJ, Cesco, Sarah, Alya, Tong, and so many others, you all have become my family in Chicago, and I'm grateful to have shared this time with such amazing people who have helped me not only to see this chapter through but to grow tremendously as a scientist and as human being.

To my advisors Carlos and Gordan, who have shaped me as a young physicist and who I admire greatly, you both have amazing qualities that I really hope I have acquired while working with you. To you I owe my confidence, my love for research, and my appreciation for what makes doing physics so special – the community.

ABSTRACT

This thesis explores the duality between the precision and discovery experimental strategies, with focused case studies in the phenomenology of the lepton sector. The complementarity of the two approaches is highlighted through an exploration of flavor violation in scenarios of extended Higgs sectors, and of theories in which the electron Yukawa is enhanced. In these cases, precision experiments offer probes that can help explore well-motivated scenarios, broadening experimental coverage of parameter space. Later, we describe cases where the paradigm can be inverted, in particular when precision experiments may be used to conduct discovery searches for intermediate-scale hidden sectors. Together, these studies highlight the blurring of the line that separates precision and discovery experimental search strategies.

CHAPTER 1

INTRODUCTION

The Higgs boson was discovered in 2012 at the Large Hadron Collider (LHC), completing the expected particle content of the Standard Model. Since its discovery, no additional elementary particles have been observed at the scale of Electroweak Symmetry Breaking (EWSB). Still, it's known that the Standard Model is not a complete description of nature: it does not explain the observed dark matter density, nor the asymmetry between matter and antimatter in the universe, nor the oscillations of neutrinos, to name a few examples. Solutions to these problems could exist at a vast range of energy scales - some models are discoverable at the LHC, some contain lighter particles, and some at energy scales far higher than what could be probed in terrestrial experiments. Given the vast space of well-motivated theories, the next discovery in physics could occur in virtually any experiment. For this reason, it's vital to have a complete understanding of each experiment and characterize how different probes work together to effectively explore the theory landscape.

In the goal of searching for new physics, there are two broad experimental strategies: discovery and precision. Discovery machines explore the frontier of high-energy physics - the Tevatron, the LHC, and the next-generation Future Circular Collider are all examples. Historically, this approach is what drives the construction of ever larger circular colliders, and has led to many of the most important discoveries: the electroweak gauge bosons, the heavy quarks, and most recently, the Higgs. These experiments require a lot of resources and a large workforce, but provide a complete exploration of new kinematic regimes and offer direct evidence for new elementary particles. In the past decade, with the running of the LHC, particle physics has focused on discovery, necessitating a certain emphasis for particle phenomenologists: determining what kinds of new particles can be uncovered at the LHC, and which final states offer the best prospects for discovery.

As the LHC continues developing its planned physics program, the focus of experimental

physics is also shifting toward the precision frontier. This strategy focuses on using increasingly precise measurements of known processes to search for deviations from the Standard Model predictions, which could indicate effects of new physics. Although less directly interpretable than searches for new particles, precision measurements can provide extremely useful information about viable models. For one, precision probes can reach far higher than the energy scale accessible at the LHC. In particular, precision measurements of the properties of leptons, such as the electric dipole moment and flavor-violating decays, are extremely sensitive probes with robust and well-understood Standard Model predictions. Experiments that measure these quantities precisely are useful alternatives to collider physics, and use known techniques with relatively cheap experimental apparatuses. Still, they come with several important drawbacks. For one, many models exist that naturally avoid or weaken these constraints, or can be easily extended to do so. Secondly, if a discrepancy with the SM is detected at one of these experiments, the precise nature of the underlying physics could not be determined until a collider is used to explore such an energy threshold. With this in mind, it's clear that the most powerful approach to search for new physics combines precision probes and collider physics.

We will show in this thesis that in the modern age of particle physics, the boundary between the precision and discovery frontiers is increasingly blurred. Machines that are built for discovery, like the LHC, can be used for precision physics, and vice versa. We will explore this in two contexts - using high energy colliders for precision physics, and using SM precision experiments for discovering new particles. As the timescale and resource requirements for new high energy colliders continues to grow, it becomes increasingly important to exploit the duality of precision and discovery, by combining their strengths to deepen our understanding of theory landscape and to search for cases where their domains intersect. In this thesis we will encounter this in four different contexts. Each explores a different aspect of lepton precision experiments, but range in the exact purpose of said experiments. The first two

chapters treat lepton precision experiments as probes of new Electroweak-scale physics, while the latter two recast precision experiments as discovery machines.

From a theoretical perspective, our guiding principle for model-building will be simplicity. All of the extensions we consider in this thesis will be single-field, although in Chapter 4 we sometimes consider non-renormalizable interactions that would generically require more heavy states. Single field extensions are useful benchmarks because of their simplicity - connections between different observables are more tightly correlated and more predictive, and their parameter spaces are minimal which makes experimental constraints easier to portray.

The thesis begins with a review of relevant topics in the Standard Model, SMEFT, and models of hidden sectors. In Chapter 3, we explore the possibility of discovering electroweak-scale new physics that has yet evaded detection, where the complementarity of precision and discovery probes is particularly manifest. In Chapter 4, we show that the LHC and the FCC-ee can be used as precision probes of the value of the electron Yukawa, but there are competing precision experiments. In Chapter 5, we show how a precision neutrino experiment can in fact be used to discover new particles with existing data. In Chapter 6, we show another example of this: a SM precision probe that can also search for new gauge forces. And lastly, this thesis will be concluded in Chapter 7.

CHAPTER 2

REVIEW OF THE LEPTON SECTOR IN THE STANDARD MODEL

2.1 The Higgs and Electroweak Symmetry Breaking

The Standard Model is a quantum field theory that is essentially defined by the gauge symmetry group of its forces, a list of particles and their charges under this gauge group, and an experimental determination of its parameters. All except one of the renormalizable terms consistent with these gauge symmetries are present in the lagrangian,¹ a remarkable property that is consistent with its interpretation as a low-energy Effective Field Theory (EFT) of some underlying theory. The gauge group of the Standard Model is

$$\text{SM} : SU(3)_C \times SU(2)_L \times U(1)_Y \quad (2.1)$$

which stand for the color, weak, and hypercharge forces, respectively. The SM symmetry group is spontaneously broken, in a mechanism known as Electroweak Symmetry Breaking (EWSB). This breaking is induced by the Higgs (H), a field which is a **2** (doublet) under $SU(2)_L$ and has hypercharge $\frac{1}{2}$. The Higgs has a potential

$$V(H) = \frac{\lambda}{4} \left(|H|^2 - \frac{v^2}{2} \right)^2 \quad (2.2)$$

which is minimized for $\langle H \rangle = v$. This vev does not transform under the $SU(2)_L$ or $U(1)_Y$ symmetries in Eq. 2.1, and thus (some of) those symmetries are broken. However, there is a linear combination of generators under which the vev is invariant, which can be thought of as a cylindrical rotation in the doublet DOFs around the direction spontaneously chosen by the

1. The exception being, notably, the $G\tilde{G}$ term, which provides an important hint about the UV completion of the Standard Model.

Higgs vev. Thus, there is one $U(1)$ rotation that survives, so the SM gauge group is broken down to $SU(3)_C \times U(1)_{EM}$, with the remaining $U(1)_{EM}$ being the familiar electromagnetic gauge symmetry. We can always leverage a gauge transformation to bring the Higgs doublet to the parameterization

$$H = \begin{pmatrix} G^\pm \\ (v + h + iG_0)/\sqrt{2} \end{pmatrix} \quad (2.3)$$

where $G^\pm$ and G_0 become the longitudinal modes of the $W^\pm$ and Z bosons, giving them their masses. In this basis, the generators of $SU(2)_L$ are the Pauli matrices. The h DOF is the radial mode of the Higgs, the only massive scalar mode by gauge symmetry, and is consistent with the state that was discovered at the LHC in 2012.

2.2 The Leptons

The fermion (spin- $\frac{1}{2}$) particles in the SM are the up quarks, down quarks, and leptons. In this thesis, we focus on the leptons and so we offer a brief review of the lepton sector of the SM. The leptons consist of two left-handed Weyl spinors and one that is right handed. First, ℓ_L is a **2** under $SU(2)_L$ with hypercharge $-\frac{1}{2}$, which contains the left-handed charged lepton and neutrino DOFs. There is also ℓ_R , the hypercharge -1 right-handed charged lepton. With these two leptons, one can write down a Yukawa term with the Higgs doublet,

$$\mathcal{L} \supset - \sum_{i,j=1}^3 \bar{\ell}_L^i y_{ij} \Phi \ell_R^j + h.c. , \quad (2.4)$$

where the indices run over the three flavors of lepton (electron, muon, and tau), and y_{ij} is a matrix of parameters called the Yukawa matrix. The lepton doublet can be parameterized

as $\ell_L = (\nu_L, e_L)$, so that after electroweak symmetry breaking,

$$\mathcal{L} \supset -y_{ij} \begin{pmatrix} \bar{\nu}_L^i \\ \bar{\ell}_L^i \end{pmatrix} \cdot \begin{pmatrix} G^\pm \\ (v + h + iG_0)/\sqrt{2} \end{pmatrix} l_R^j + h.c. = -\frac{y_{ij}}{\sqrt{2}}(v + h)\bar{\ell}_L^i l_R^j + h.c. . \quad (2.5)$$

Thus the standard model predicts that the mass matrix of the leptons is

$$[M_\ell]_{ij} = \frac{vy_{ij}}{\sqrt{2}} \quad (2.6)$$

and the Yukawa couplings of the physical Higgs boson to the leptons are

$$Y_{ij}^h = \frac{y_{ij}}{\sqrt{2}} = \frac{1}{v}[M_\ell]_{ij} . \quad (2.7)$$

Here we encounter a very important property of the SM - the direct proportionality between the couplings of the Higgs and the masses of the charged leptons. To write the Lagrangian in terms of its observed physical states, we remove the lepton mass mixing terms by bringing it to a diagonal form. To do so we apply rotations in flavor space to the left and right handed charged leptons $\ell_L^i \rightarrow [U_L]_{ij}\ell_L^j$ and $\ell_R^i \rightarrow [U_R]_{ij}\ell_R^j$, under which the kinetic terms and gauge couplings are invariant. However, this transforms the Yukawas, and we can choose $U_{L/R}$ such that the mass matrix becomes diagonal:

$$M_\ell \rightarrow [U_L^\dagger M_\ell U_R]_{ij} = m_i \delta_{ij} . \quad (2.8)$$

As a consequence of the aforementioned proportionality, the Higgs couplings are also diagonalized,

$$Y_{ij}^h = \frac{m_i}{v} \delta_{ij} . \quad (2.9)$$

This relation is of utmost importance, as both the diagonal and off-diagonal entries of Y_{ij}^h can be tested at the LHC and in precision experiments. We will explore the consequences of this prediction in Chapters 3 and 4.

2.3 Models of New Physics

New physics essentially has two ways of manifesting: through slight modifications to observables from their SM-predicted values, or through direct observation of new particles being created and decaying.

Heavy New Physics If new physics is heavier than the scales accessible through experiments, their effects are well-described by Effective Field Theory (EFT). Effective Field Theory (EFT) encodes the effects of new physics through higher dimensional operators. This encoding can essentially be understood as a taylor-esque expansion in E/m , where E is the energy scale of the experiment and m is some heavy particle’s mass. For example, propagators with momenta much less than their masses can taylor expanded

$$\frac{i}{p^2 - m^2} = -\frac{i}{m^2} - \frac{ip^2}{m^4} - \frac{ip^4}{m^6} + \mathcal{O}(m^{-8}) . \quad (2.10)$$

Which is valid so long as $p^2 \leq m^2$. Let us sketch schematically the power of this approach. Consider a UV lagrangian that contains a heavy field ϕ and couples to some lighter particles through two operators $\mathcal{O}_1$, $\mathcal{O}_2$ (local functions of other light fields in the theory). This is captured through the “UV Completion” Lagrangian,

$$\mathcal{L}_{\text{UV}} \supset \phi(-\partial^2 - M_\phi^2)\phi - \lambda_1\phi\mathcal{O}_1 - \lambda_2\phi\mathcal{O}_2 \quad (2.11)$$

With this lagrangian we can compute, say, the amplitude of $\mathcal{O}_2 \rightarrow \mathcal{O}_1$ at tree level. In momentum space this is

$$\implies \mathcal{A}_{\text{UV}}(\mathcal{O}_1 \rightarrow \mathcal{O}_2) \sim \frac{i\lambda_1\lambda_2}{p^2 - M_\phi^2} \quad (2.12)$$

However, if we can only build experiments that probe $\mathcal{A}$ at momentum scales $p^2 \ll M_\phi^2$, then we can perform a Taylor expansion of the propagator as in Eq. 2.10 and the amplitude remains the same. Equivalently, one can start with a Lagrangian where the vertices encode this Taylor expansion,

$$\mathcal{L}_{\text{EFT}} \supset \mathcal{O}_1 \left[-\frac{1}{M_\phi^2} + \frac{\partial^2}{M_\phi^4} - \frac{\partial^4}{M_\phi^6} + \dots \right] \mathcal{O}_2 . \quad (2.13)$$

In this theory, ϕ is not present, but still the amplitude is

$$\mathcal{A}_{\text{EFT}}(\mathcal{O}_1 \rightarrow \mathcal{O}_2) \sim -i\lambda_1\lambda_2 \left[\frac{1}{M_\phi^2} + \frac{p^2}{M_\phi^4} + \frac{p^4}{M_\phi^6} + \dots \right] . \quad (2.14)$$

We have traded the field ϕ for an infinite tower of higher-dimensional terms in our lagrangian, but the amplitudes (for $E \ll M_\phi$) in either theory are exactly identical. This rather simple example illustrates a very powerful principle in QFT—at energy scales below some heavy particle mass, one can always pick the coefficients of the EFT such that the prediction in both theories match arbitrarily well. Perhaps more surprisingly, the exact equivalence between EFT and UV completion persists at loop level (again, so long as $E \leq M_\phi$).

This may appear to be a needless complication - we replaced a heavy field with two couplings with infinitely many new terms in the Lagrangian. However, in the context of experimental physics this is actually quite powerful. For one, experimental physics has limited precision - so in practice, we can limit our considerations to only the first few terms in the E/m expansion. Secondly, experimental physics starts at lower energy scales and

work upwards, so working in the EFT can capture the generic effects of new physics, without having to select a particular model. Because of this, EFT can provide a powerful, model independent framework for potential modifications to the Lagrangian of known physics.

In modern particle physics, Effective Field Theory is ubiquitous. One important place it arises is considering possible effects from new physics above the Electroweak scale, $M \gg \langle H \rangle$. This physics can be captured at leading order using the Standard Model Effective Field Theory (SMEFT), which is an extension of the SM that includes all gauge-invariant dimension-6 operators. This gives the lagrangian

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \sum_i \frac{c_i}{\Lambda^2} \mathcal{O}_i^{\text{D6}} \quad (2.15)$$

This expands the number of parameters in the theory significantly, from 19 to 2499. SMEFT is canonically written in the *Warsaw Basis*, which enumerates and organizes all operators [22]. However, most models of new physics only populate a smaller sector of the full SMEFT space, and this will occur in the extensions considered in this thesis as well. We will see explicit examples of this in Chapter 4. For further review, see Ref. [23].

New Higgs Doublets One scenario of heavy new physics where SMEFT breaks down is the Two Higgs Doublet Model (2HDM), which is a simple extension of the SM that includes a second source of Electroweak Symmetry Breaking. Although this theory maps onto SMEFT in certain limits, its better treated in the full picture accounting for all new states. The 2HDM is perhaps the most rich single-field extension of new physics, and is capable of many important phenomenological effects that guide experimental and cosmological studies. In Chapter 3 we explore this theory in both collider and precision experiments. For a more detailed review, see [24].

Light Gauge Bosons Turning to new physics below the EW scale, we will also consider new light particles in Chapters 5 and 6. In particular, we will consider the addition of new light gauge bosons to the SM. The addition of new gauge forces is heavily restricted in the Standard Model by anomaly cancellation. If there are no new fermions, then the couplings to the Standard Model determine the value of the so-called “trace anomaly”: a one loop amplitude that breaks gauge invariance, but can be nonzero if the charges of the fermions are not chosen correctly. If the charge assignments indeed lead to nonzero trace anomaly, then although the theory is gauge invariant at the level of the Lagrangian, the path integral, which is a full definition of the QFT, is not truly gauge invariant. Therefore, for the QFT to be consistent, charges must be assigned in a way to lead to zero trace anomaly. With this in mind, there are only a few possible extensions of the SM that avoid the trace anomaly. Of these, only three types are phenomenologically viable: the dark photon, the gauged $L_i - L_j$, and the $B - L$ gauge bosons. These three possibilities are reviewed in Section 6.1.

Light New Scalars Another scenario considered in this thesis is that of light scalars - who’s couplings can only be restricted by additional global symmetries. In the quark sector, global symmetries naively would enforce zeroes in the CKM matrix, but in the lepton sector they can be chosen freely. This global symmetry cannot be exact, because the PMNS matrix has off-diagonal terms. However, the neutrino masses are not generated by the SM and therefore this level of breaking is viable. Enforcing new global symmetries allows one to couple a new scalar to any of the new fermions. This theory further requires additional new physics above the Electroweak scale. Unless the scalar is a new Higgs doublet, it cannot form a renormalizable gauge-invariant combination with the leptons. Nonetheless, its a viable theory that we study in Chapter 5.

CHAPTER 3

LEPTON FLAVOR IN THE 2HDM

3.1 Introduction

In the first chapter of this thesis, we will explore the complementarity between precision and discovery experiments in the Two Higgs Doublet Model (2HDM). The 2HDM is a single-field extension of the SM that includes an additional Higgs doublet that breaks Electroweak symmetry, but includes many new parametric degrees of freedom. In particular, one aspect of the 2HDM is that it may be flavor violating - an extra Yukawa matrix introduces the possibility of $\mu \rightarrow e\gamma$, for example, whose branching ratio in the Standard Model is $\sim 10^{-51}$. This is far beyond the reach of any conceivable experiment, but can be easily raised to experimental levels in the 2HDM. This possibility is well-known in the 2HDM and the parameter space where it occurs in its most generic form is largely ruled out. However, the possibility of flavor violation at one loop remains open, and we will show that in this case the bounds placed by direct collider searches and precision experiments are currently very competitive with one another. In the next decade, advances at both the LHC and in precision observables will jointly cover a significant portion of the 2HDM parameter space.

As we saw in Section 2.1, a remarkable property of the SM is that one can assign arbitrary complex Yukawa couplings of the fermions to the Higgs field and, upon diagonalization to the fermion mass basis, no tree-level flavor violation effects are left in the Higgs and neutral gauge sectors. All flavor violating effects are induced by charged gauge currents and can be parametrized by the CKM and PMNS matrix elements in the quark and lepton sectors. It is clear that the above property is closely related to the fermion mass generation mechanism in the SM, where the fermion mass matrices are directly proportional to the Yukawa ones.

The phenomenology of flavor changes drastically if there are additional Higgs doublets at the EW scale, a class of scenarios broadly known as N Higgs doublet models (NHDM).

In these models, the SM relation between mass and coupling matrices is broken, allowing tree-level flavor violation in the Higgs sector. The simplest extensions of the Higgs sector of this kind is the Two Higgs doublet Model [24, 25, 26, 27, 28, 29, 30] (2HDM), which includes a single additional source of electroweak symmetry breaking. Flavor symmetries are usually invoked to ensure that each quark and lepton sector couples to only one of the two Higgs doublets, restoring the flavor conservation mechanism present in the SM. These flavor symmetries, which hold in the ultraviolet regime, may be violated in the infrared. This is the case, for instance, in the minimal supersymmetric extension of the SM (MSSM), in which flavor conservation is ensured by the holomorphicity of the superpotential and the structure of couplings, but is violated by supersymmetry breaking effects at low energies [31, 32, 33, 34, 35, 36].

In the absence of flavor symmetries, or in the case that these symmetries are violated, flavor violating effects are expected to be present, leading to observable effects in Higgs physics at colliders as well as in dedicated flavor violating searches. This raises several questions: Which patterns of flavor violation are theoretically motivated? Which experiments place the strongest restrictions on 2HDM lepton flavor violation, and what parameter space remains to be probed? In this article, we address these questions within the framework of the 2HDM, approaching the question of flavor from an EFT-minded perspective. We make contact with the MSSM when appropriate, but remain largely agnostic to the UV completion of the 2HDM, aside from several key assumptions.

For our analysis, we will assume that the most relevant flavor violating effects are described by the dimensionless Yukawa couplings in the low energy effective theory, while direct flavor violating effects, related to higher dimensional operators, are efficiently suppressed in the 2HDM low energy effective theory. We will concentrate on the lepton sector and therefore the quark sector structure, although relevant, will not offer any new observables in our analysis. For the aim of computing the relevant processes in the lepton sector, we will treat

the quark sector as following aligned flavor violation - that there is only one fundamental flavor matrix in each of the up and down sectors, so the CKM matrix is the only source of quark flavor violation.

3.2 Two Higgs Doublet Models

The 2HDM is an extension of the Standard Model, which includes an additional complex scalar with the same charge assignments as the SM Higgs, transforming as $(\mathbf{1}, \mathbf{2})_1$ under the SM $SU(3)_c \times SU(2)_L \times U(1)_Y$ gauge symmetry group. In addition to the kinetic term, the most general gauge invariant interactions of the two Higgs doublets are given by

$$\mathcal{L} \supset - \sum_{f \in \{u, d, \ell\}} \bar{f}_L^i ([y_1]_{ij} \Phi_1 + [y_2]_{ij} \Phi_2) f_R^j + h.c. - V(\Phi_1, \Phi_2). \quad (3.1)$$

Where $V(\Phi_1, \Phi_2)$ includes all possible gauge invariant renormalizable combinations of Φ_1, Φ_2 .

In its most generic form, it is written

$$\begin{aligned} V(\Phi_1, \Phi_2) = & \sum_{ij} m_{ij}^2 \Phi_i^\dagger \Phi_j + \frac{1}{2} \lambda_1 |\Phi_1|^4 + \frac{1}{2} \lambda_2 |\Phi_2|^4 + \lambda_3 |\Phi_1|^2 |\Phi_2|^2 + \lambda_4 |\Phi_1^\dagger \Phi_2|^2 \\ & + \left[\frac{1}{2} \lambda_5 (\Phi_1^\dagger \Phi_2)^2 + \left(\lambda_6 |\Phi_1|^2 + \lambda_7 |\Phi_2|^2 \right) \Phi_1^\dagger \Phi_2 + h.c. \right] \end{aligned} \quad (3.2)$$

In general, $m_{12} = m_{21}^*$; λ_5 , λ_6 , and λ_7 can be complex, and their phases explicitly introduce CP-violation. A sufficient but not necessary condition for CP symmetry to be preserved in the Higgs potential is given by [37, 38]

$$\text{Im}[\lambda_6 \lambda_7^*] = \text{Im}[\lambda_6^2 \lambda_5^*] = \text{Im}[\lambda_7^2 \lambda_5^*] = \text{Im}[m_{12}^4 \lambda_5^*] = 0, \quad (3.3)$$

In this case all parameters can be rendered real by a simple phase redefinition of one of the two Higgs doublets. In the following, we will assume that this condition is fulfilled and

that the potential is therefore CP invariant.

3.2.1 *The Higgs Basis and Electroweak Symmetry Breaking*

After Electroweak Symmetry Breaking (EWSB), the original parameters of the 2HDM map onto a new set of physical parameters, including e.g. the masses of the fermions, scalar Higgs bosons, and gauge bosons. As an intermediate step which simplifies many expressions, we first enter the “Higgs Basis,” and refer primarily to the parameters of the potential in this basis[25]. We arrive in this basis by performing a field redefinition of Φ_1 and Φ_2 such that the vev is entirely contained within one Higgs doublet, denoted as H_1 , with the other doublet, denoted as H_2 , having zero vev. This rotation is parameterized by the angle β , defined as

$$\tan(\beta) = \frac{\langle \Phi_2 \rangle}{\langle \Phi_1 \rangle}. \quad (3.4)$$

The relationship between H_i and Φ_i is written in terms of the the usual 2D rotation matrix,

$$H_i = R_{ij}(\beta)\Phi_j, \quad (3.5)$$

where

$$R(\beta) = \begin{pmatrix} \cos \beta & \sin \beta \\ -\sin \beta & \cos \beta \end{pmatrix}. \quad (3.6)$$

Performing this rotation induces a map from the Higgs potential to a potential of the same generic form (since it’s the most general possible potential) but with different parameters, $\{m_{ij}, \lambda_i\} \rightarrow \{Y_2, Y_3, v, Z_i\}$. This mapping is given e.g. in Ref. [39]. We write the potential,

in terms of the Higgs Basis parameters, as

$$V(H_1, H_2) = Y_2 |H_2|^2 - (Y_3 H_1^\dagger H_2 + h.c.) \quad (3.7)$$

$$+ \frac{1}{2} Z_1 \left(|H_1|^2 - \frac{v^2}{2} \right)^2 + \frac{1}{2} Z_2 |H_2|^4 + Z_3 |H_1|^2 |H_2|^2 \\ + Z_4 |H_1^\dagger H_2|^2 + \left[\frac{1}{2} Z_5 (H_1^\dagger H_2)^2 + \left(Z_6 |H_1|^2 + Z_7 |H_2|^2 \right) H_1^\dagger H_2 + h.c. \right] \quad (3.8)$$

Finally, we compute the theory after EWSB by writing the Higgs doublets in terms of their real scalar degrees of freedom:

$$H_1 = \begin{pmatrix} G^\pm \\ (v + H_1^0 + iG^0)/\sqrt{2} \end{pmatrix} \quad H_2 = \begin{pmatrix} H^\pm \\ (H_2^0 + ia^0)/\sqrt{2} \end{pmatrix}. \quad (3.9)$$

The goldstones, $G^\pm$ and G^0 , are eaten by the EW bosons, but H_1^0 and the four DOFs of H_2 are dynamical. By inserting these definitions into the Higgs potential, the mass of the charged Higgs is

$$M_\pm = Y_2 + \frac{1}{2} Z_3 v^2, \quad (3.10)$$

and the mass matrix for the physical neutral Higgses can be written

$$M_{H_i^0, a^0} = \begin{pmatrix} Z_1 v^2 & Z_6^R v^2 & -Z_6^I v^2 \\ Z_6^R v^2 & M_\pm^2 + \frac{1}{2}(Z_4 + Z_5^R) v^2 & -\frac{1}{2} Z_5^I v^2 \\ -Z_6^I v^2 & -\frac{1}{2} Z_5^I v^2 & M_\pm^2 + \frac{1}{2}(Z_4 - Z_5^R) v^2 \end{pmatrix}, \quad (3.11)$$

where the superscripts R and I refer to the real and imaginary parts of the corresponding quantities. Assuming all quartic couplings to be real, one recovers the CP-conservation expressions. In this case, the pseudoscalar a^0 does not mix with the other CP-even scalars, and so we can write the mass eigenstates h, H, A in terms of a single mixing angle $\beta - \alpha$, as

is standard in the 2HDM,

$$\begin{pmatrix} h \\ H \\ A \end{pmatrix} = \begin{pmatrix} \sin(\beta - \alpha) & \cos(\beta - \alpha) & 0 \\ \cos(\beta - \alpha) & -\sin(\beta - \alpha) & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} H_1^0 \\ H_2^0 \\ a^0 \end{pmatrix}. \quad (3.12)$$

To make contact with the computations in Ref. [1], which we use extensively, we can alternatively express the diagonalization matrix in terms of the q -matrix

$$\begin{pmatrix} \sin(\beta - \alpha) & \cos(\beta - \alpha) & 0 \\ \cos(\beta - \alpha) & -\sin(\beta - \alpha) & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} q_{11} & \text{Re}(q_{12}) & \text{Im}(q_{12}) \\ q_{21} & \text{Re}(q_{22}) & \text{Im}(q_{22}) \\ q_{31} & \text{Re}(q_{32}) & \text{Im}(q_{32}) \end{pmatrix} \quad (3.13)$$

The form of the mass matrix in Eq. (3.11) has two implications important for our work. The first is that in the CP-conservation case, the off diagonal terms imply a relation between the CP-even neutral Higgs bosons mixing angle in this basis, defined as $\beta - \alpha$, and the mass eigenstates, namely

$$\frac{-Z_6 v^2}{\sqrt{(m_H^2 - m_h^2)(m_H^2 - Z_1 v^2)}} = \cos(\beta - \alpha) \simeq \frac{-Z_6 v^2}{(m_H^2 - m_h^2)} \quad (3.14)$$

In certain parts of our analysis, we use this formula to approximately relate m_H to $\cos(\beta - \alpha)$, and to explore parameter space where the couplings can be safely treated perturbatively, e.g. parameter space where $|Z_6| \leq 4\pi$. Second, there is a constraint on the mass splitting between the heavy Higgs and the charged Higgs boson,

$$|M_{H,A}^2 - M_{\pm}^2| \simeq |Z_4 \pm Z_5| \frac{v^2}{2} < 4\pi v^2 \quad (3.15)$$

A more accurate expression of these bounds may be found in Refs. [40, 24, 39]. In this paper, to simplify the parameter space, we often will work in regions where $m_H \simeq m_H^\pm \simeq m_A$, which is an approximation that becomes more accurate for $m_H \gg v$ or for moderate values of the quartic couplings.

3.2.2 2HDM Yukawas

Turning to the fermion sectors, we compute the Yukawa couplings between the scalar mass eigenstates and the SM fermions. There are three new scalars (plus one charged scalar) which couple to the SM fermions. The coupling of the Higgs doublets $\Phi_{1,2}$ to the SM fermions is written as in Eq. (3.1)

$$\mathcal{L} = - \sum_{f \in \{u, d, \ell\}} \bar{f}_L^i ([y_1^f]_{ij} \Phi_1 + [y_2^f]_{ij} \Phi_2) f_R^j + h.c. \quad (3.16)$$

We then rotate the Higgs fields into the ‘‘Higgs Basis’’, and the Yukawa matrices corresponding to H_1, H_2 may be related to y_i via the same transformation, with a factor of $\sqrt{2}$ that will absorb that of the Higgs normalization,

$$\begin{aligned} \mathcal{Y}_1^f &= \frac{y_1^f \cos \beta + y_2^f \sin \beta}{\sqrt{2}}, \\ \mathcal{Y}_2^f &= \frac{-y_1^f \sin \beta + y_2^f \cos \beta}{\sqrt{2}}. \end{aligned} \quad (3.17)$$

In this basis, as emphasized before, the H_1 doublet acquires the full vev v , and the neutral component H_1^0 of H_1 behaves in an SM-like manner. Since H_2 acquires no vev, the masses arise entirely from the coupling to H_1 :

$$M_{ij}^f = v [\mathcal{Y}_1^f]_{ij} \quad (3.18)$$

which may be diagonalized by a rotation by unitary matrices,

$$m_i^f \delta_{ij} = [U_L M^f U_R^\dagger]_{ij} . \quad (3.19)$$

We denote Yukawa matrices written in the fermion mass basis by a bar:

$$\overline{\mathcal{Y}}_i^f = U_L \mathcal{Y}_i^f U_R^\dagger, \quad \overline{y}_i^f = U_L y_i^f U_R^\dagger, \quad \text{etc.} \quad (3.20)$$

The electroweak bosons also get their masses via their coupling to H_1 , in the same way as in the SM. In the end, there are five physical Higgs DOFs - the light Higgs (h), which is constrained to be “SM-like” by experimental observations; the heavy Higgs (H); the charged Higgs ($H^\pm$) with 2 DOFs; and the pseudoscalar Higgs (A). In terms of the q -matrix defined in the previous section, the low-energy lagrangian in the lepton sector $f = \ell$ can be written

$$\mathcal{L} \supset -q_{kr} h_r [\overline{\mathcal{Y}}_k^\ell]_{ij} \bar{e}_L^i e_R^j - \sqrt{2} [\overline{\mathcal{Y}}_2^\ell]_{ij} H^+ \bar{\nu}_L^i e_R^j + h.c. . \quad (3.21)$$

Or, in terms of the mixing angles, the Yukawa couplings of the lightest CP-even physical Higgs h can be written as

$$\mathcal{L} \supset - \left[\frac{m_i}{v} \delta_{ij} \sin(\beta - \alpha) + [\overline{\mathcal{Y}}_2^\ell]_{ij} \cos(\beta - \alpha) \right] h \bar{f}_L^i f_R^j + h.c. \quad (3.22)$$

while the heavier CP-even Higgs has couplings of the form

$$\mathcal{L} \supset - \left[\frac{m_i}{v} \delta_{ij} \cos(\beta - \alpha) - [\overline{\mathcal{Y}}_2^\ell]_{ij} \sin(\beta - \alpha) \right] H \bar{f}_L^i f_R^j + h.c. \quad (3.23)$$

Note that flavor violating couplings for the physical Higgses arise entirely from the off-diagonal elements of the $\overline{\mathcal{Y}}_2$ matrix, while the deviations from the SM-Higgs diagonal couplings come from the diagonal elements of $\overline{\mathcal{Y}}_2$. In the following sections, we will examine the

bounds on the generic $\overline{\mathcal{Y}}_2$ elements and motivate specific structures for $\overline{\mathcal{Y}}_2$.

3.3 Lepton Flavor

In the SM, the lepton and quark Yukawa matrices have the pleasant property of being technically natural: if we set the Yukawas to zero, we recover the global flavor symmetry

$$U(3)^5 = U(3)_\ell \times U(3)_e \times U(3)_Q \times U(3)_u \times U(3)_d . \quad (3.24)$$

Consequently, if any of the fermion masses are set to zero, we restore a subgroup of this symmetry. Therefore, although there are clearly unexplained hierarchies in the fermion masses, their technical naturalness is perfectly consistent with EFT principles. This hierarchy does not necessarily raise fine tuning issues in the same way as, say, the Electroweak Hierarchy Problem.

To perform a spurion analysis, we make this symmetry explicit by treating the fermions as transforming as

$$\ell_L \sim \mathbf{3}_\ell \quad , \quad e_R \sim \mathbf{3}_e \quad , \quad Q_L \sim \mathbf{3}_Q \quad , \quad u_R \sim \mathbf{3}_u \quad , \quad d_R \sim \mathbf{3}_d \quad (3.25)$$

and the SM Higgs Yukawas Y_{SM}^f as

$$Y_{SM}^\ell = \mathbf{3}_\ell \times \bar{\mathbf{3}}_e \quad , \quad Y_{SM}^u = \mathbf{3}_Q \times \bar{\mathbf{3}}_u \quad , \quad Y_{SM}^d = \mathbf{3}_Q \times \bar{\mathbf{3}}_d \quad (3.26)$$

Because the global $U(5)^5$ symmetry must be conserved, the Higgs couplings and the fermion masses are forced to inherit the same flavor matrix. Therefore, diagonalization of the masses also diagonalizes the Higgs couplings, and the SM robustly predicts no flavor changing neutral currents.

In the 2HDM, the two sets of Yukawa matrices of the 2HDM must have the same spurion

charges as in the SM case, Eq. (3.26) with $Y_{SM}^f \rightarrow y_i^f$. This raises two main problems. First, unless there is some correlation between the two sets of Yukawas (which requires additional structure in the EFT), setting one of the fermion masses to zero does not restore any global symmetry, since there will still be a Yukawa coupling to the other Higgs DOF (H_2 in the Higgs basis). Therefore, the fermion masses are no longer technically natural, and the smallest fermion masses generically require tuning in order to justify their size. Secondly, any linear combination of the two Yukawas will have the same flavor charge - meaning one is no longer required by symmetry to use the same matrix for the neutral Higgs couplings as for the lepton mass matrices, and thus in general they are not simultaneously diagonalizable. The loss of technical naturalness is an important problem from an EFT perspective - it's this problem that we address in this section.

3.3.1 Flavor Considerations in the 2HDM

If the 2HDM is to remain well-motivated from an EFT standpoint, then one must try to restore technical naturalness for the fermion mass hierarchies, which requires additional flavor symmetries to be invoked. To this end, there are essentially two options for making the lepton mass hierarchy compatible with EFT principles in the 2HDM:

1. Make a motivated assumption or choice of model for the flavor structure in the Yukawa matrices, inducing hierarchies in the entries of both Yukawas.
2. Impose a discrete symmetry that disallows one of the two sets of Yukawa matrices, reinstating default technical naturalness and forgoing questions about the origin of flavor.

In the first approach, one has to guess the EFT structure that is imposed in the full UV theory of flavor. This additional structure will be captured in the 2HDM Yukawas by some flavorful global symmetry G_F and a set of spurionic parameters. The left and right handed

leptons and the spurions are charged under G_F , and G_F -invariance dictates the Lagrangian of the EFT. This fixes the spurion dependence of the Yukawa matrix entries and induces some hierarchical structure in the fermion masses and Yukawas. The simplest choice is $G_F = U(1)_{FN}$ with one flavor spurion, which is known as the Froggatt-Nielsen (FN) model. The leptons and FN spurion are charged in such a way as to require several insertions of the spurion to set the masses of the lighter leptonic degrees of freedom. Then the value of this spurionic parameter is taken to be small, giving rise to a hierarchy in the lepton masses. Generically, if all states below the flavor UV completion scale also respect G_F , then the loop corrections that they generate will have the same spurion patterns as the tree-level Yukawas, albeit with new $\mathcal{O}(1)$ coefficients. Although the 2HDM does not appear technically natural in full generality, restricting to a specific G_F -spurion structure renders the flavor spurions themselves technically natural, and therefore the mass hierarchy as well.

Alternatively, to avoid making such a choice, discrete symmetries can be imposed. To achieve this, there are several ways that $\mathbb{Z}_2$ charges can be assigned to the fermions, resulting in 2HDM categories known as type-I, type-II, Lepton-Specific, and Flipped [24]. However, in certain scenarios, including the MSSM, simple extensions of the 2HDM introduce controlled flavor violation. In these scenarios, there are new $U(3)_\ell \times U(3)_e$ spurions appearing in other heavy sectors, which are transmitted to the 2HDM via loops. There must also be a soft- $\mathbb{Z}_2$ breaking spurion in the heavy sector. The MSSM is a canonical example of this, since the slepton masses present new flavorful matrices and the μ -term breaks $\mathbb{Z}_2$. Indeed, as outlined in Appendix A, both of these parameters appear in the $\mathbb{Z}_2$ -breaking loop correction. If flavor is broken at loop level in the 2HDM, we are again somewhat forced into the same situation as in the generic case - the question of flavor structure again becomes important.

In either of these cases, one would like to parameterize an assumption about flavor with a “flavor ansatz.” However, just where the flavor ansatz should be inserted is not immediately obvious—in some sense, there is a continuous family of flavor matrices parameterized by

$\beta - \alpha$, any one of which could naively be replaced with a flavor ansatz after EWSB. In the rest of the section, we motivate a framework for inserting flavor ansatz, and we find that there is a fairly unambiguous choice for where it must be inserted.

3.3.2 Controlling $\mathcal{O}(1)$ Coefficients

In this section we outline the process of matrix diagonalization, with the goal of informing the insertion of flavor ansatz in the 2HDM. Flavor violation specifically occurs because, although the spurion structure of any flavorful matrix is fixed and uniform amongst all flavorful matrices, the $\mathcal{O}(1)$ coefficients are not. That being said, flavor ansatz are generally defined as *expected* magnitudes of flavorful entries—but what exactly is meant by *expected* must be defined more concretely. In this regard, in many analyses, $\mathcal{O}(1)$ coefficients are sampled from unit gaussian distributions and the results are displayed as scatterplots. We construct a framework for this analytically - we treat the $\mathcal{O}(1)$ as sampled from a statistical distribution of unit variance and compute the resulting flavor structure after EWSB.

Consider two matrices A and B with the same (real-valued) spurionic structure X_{ij} , where for example, in Froggatt-Nielsen,

$$X_{ij} \stackrel{\text{FN}}{=} \left(\frac{\langle \phi_{\text{FN}} \rangle}{\Lambda_{\text{FN}}} \right)^{|q_{\ell_i} - q_{e_j}|} . \quad (3.27)$$

In general, other flavor structures can arise as well in this framework. Flavorful matrices will express the spurion structure, but with $\mathcal{O}(1)$ coefficients a_{ij} and b_{ij} :

$$A_{ij} = a_{ij} X_{ij} \quad , \quad B_{ij} = b_{ij} X_{ij} \quad . \quad (3.28)$$

To parameterize our ignorance of the coefficients, we consider an ensemble of such matrices, with the complex $\mathcal{O}(1)$ coefficients a_{ij} , b_{ij} sampled independently from a symmetric

distribution such that

$$\langle a_{ij} \rangle = \langle b_{ij} \rangle = 0 \quad , \quad \langle a_{ij} a_{kl}^* \rangle = \langle b_{ij} b_{kl}^* \rangle = \delta_{ik} \delta_{jl} \quad , \quad \langle a_{ij} b_{ij}^* \rangle = 0 \quad . \quad (3.29)$$

In this case, because $\langle A_{ij} B_{ij}^* \rangle \approx 0$, A and B will be said to be *statistically independent*. The concept of statistical independence of $\mathcal{O}(1)$ coefficients arises in many UV completions of flavor, especially when there are many states that are integrated out to set the spurion structure of the flavor EFT. It is not fully generic, but we assume that all flavorful matrices are generated with this property at the scale of flavor completion.

Additionally, because $\langle |A_{ij}|^2 \rangle = \langle |B_{ij}|^2 \rangle = X_{ij}^2$, they will be said to have *identical flavor structures*. When A is diagonalized,

$$\bar{A}_{ij} \equiv [U_L A U_R^\dagger]_{ij} = \lambda_i \delta_{ij} \quad , \quad (3.30)$$

the eigenvalues are predicted by the underlying theory up to $\mathcal{O}(1)$, $\langle \lambda_i \rangle = e_i$. Then, so long as B is statistically independent from A and has the same flavor structure, applying the same transformation to the other matrix will produce a “structure matrix,”

$$\bar{B}_{ij} \equiv [U_L B U_R^\dagger]_{ij} = b_{ij} S_{ij}^F \quad . \quad (3.31)$$

Here, S_{ij}^F is a fixed matrix that can be predicted with knowledge of the spurion structure X_{ij} , along with the induced diagonalization matrices U_L and U_R . S_F is computed as the average over an ensemble of $\mathcal{O}(1)$ coefficients. However, deriving S^F analytically from a specific UV completion is often nontrivial, and typically numerical codes are used to display the distribution of flavorful observables.

For example, suppose A and B both have the hierarchical structure of the Fritzsch

textures [41]

$$X = \begin{pmatrix} f & d & 0 \\ d^* & 0 & c \\ 0 & c^* & 0 \end{pmatrix}, \quad (3.32)$$

Then the average eigenvalues and structure matrix are

$$e_i = \{f, d^2/f, c^2 f/d^2\} \quad , \quad S_{ij}^F = \begin{pmatrix} e_1 & \sqrt{e_1 e_2} & \sqrt{e_1 e_3} \\ \sqrt{e_1 e_2} & e_2 & \sqrt{e_2 e_3} \\ \sqrt{e_1 e_3} & \sqrt{e_2 e_3} & e_3 \end{pmatrix}. \quad (3.33)$$

This kind of structure is known as the **Cheng-Sher ansatz** and is one of the most common ansatze for flavor [42]. It also arises in many Froggatt-Nielsen models.

Another interesting possibility is that an flavor-specific $\mathbb{Z}_2$ may emerge as part of the flavor symmetry group $G_F \supset \mathbb{Z}_{2,e^i}$. This would render one of the leptons flavor-aligned, and disallow any flavor violation involving that lepton.¹ This type of symmetry may arise accidentally from the Froggatt-Nielsen charge assignments, or through some other mechanism. We will consider this model under the assumption that the electron is the flavor-aligned lepton, which disallows flavor violation in the μe and τe channels. As such, because the flavor structure is “aligned” for the electron couplings, but Cheng-Sher-like for the $\mu\tau$ sector, we call this model the **$\mu\tau$ -Specific Ansatz**, and assume the following structure

$$S^F = \begin{pmatrix} e_1 & \sqrt{e_1 e_2} & 0 \\ \sqrt{e_1 e_2} & e_2 & 0 \\ 0 & 0 & e_3 \end{pmatrix}. \quad (3.34)$$

It should be noted that these zeros survive RG flow since they are enforced by global

1. This flavor-specific symmetry must be broken at some scale to reproduce the observed PMNS structure.

symmetry. In contrast, the Fritzsch texture zeros do not correspond to a particular symmetry, and so the zeros might deviate slightly under RG flow, an effect which we neglect since it will be proportional to loop-suppressed powers of the lepton masses or weak gauge couplings.

Lastly, we emphasize that this line of reasoning relies on A and B above having *identical flavor structures* and being *statistically independent*. In the generic 2HDM, we assume the Yukawa matrices of the two Higgs doublets satisfy these criteria, so that the insertion of a flavor ansatz S^F is straightforward. For the 2HDM with tree-level discrete symmetries, neither condition will in general be satisfied by the two Yukawa matrices. We take care to ensure proper treatment of the flavor ansatz in the next section.

3.3.3 The Inheritance of Flavor Structure in the 2HDM

Motivated by the discussion in Section 3.3.1, we consider two scenarios: when lepton Yukawas to both Higgs doublets are both generated at tree level, or when there is a loop factor between them. Assuming they are generated with essentially random $\mathcal{O}(1)$ coefficients, the 2HDM in full generality can be expressed

$$[y_1^\ell]_{ij} = a_{ij}X_{ij} \quad , \quad [y_2^\ell]_{ij} = \varepsilon b_{ij}X_{ij} \quad (3.35)$$

where for example $\varepsilon = 1$ is the generic type-III 2HDM and $\varepsilon = 0$ is the 2HDM with an exact $\mathbb{Z}_2$ symmetry. In the intermediate case, nonzero ε can represent a loop factor between y_2 and y_1 , e.g. as occurs in the MSSM. To utilize the flavor framework constructed in the previous section, we leverage the assumption that the coefficients are random, $\langle a_{ij} \rangle = \langle b_{ij} \rangle = 0$, $\langle a_{ij}^2 \rangle = \langle b_{ij}^2 \rangle = 1$. It should be mentioned that, even if $\langle [y_1]_{ij}^* [y_2]_{ij} \rangle = 0$ at the scale that y_2 is generated (in the generic 2HDM, this would be the flavor UV completion scale, while in the MSSM it would be the slepton mass scale), this exact relation is broken by RG effects. Even if the top Yukawa is involved, this is only relevant for very high flavor UV completion

scales, e.g. near the GUT scale. We neglect the effects of RG.

Although y_1 and y_2 are statistically independent, they only share a flavor structure *up to a constant rescaling factor* ε , due to e.g. a loop suppression. First, we extract the rescaling factor ε , so that $y_2 = \varepsilon z_2$, so that now z_2 and y_1 have *identical* flavor structures. Now writing the Yukawas in the Higgs basis,

$$\begin{aligned}\mathcal{Y}_1 &= \frac{1}{v} M_\ell = (y_1 \cos \beta + \varepsilon z_2 \sin \beta) / \sqrt{2} \\ \mathcal{Y}_2 &= (-y_1 \sin \beta + \varepsilon z_2 \cos \beta) / \sqrt{2} .\end{aligned}\tag{3.36}$$

In order apply the framework of the previous section, we first note that generally $\mathcal{Y}_2$ is not statistically independent from M_ℓ ,

$$\frac{1}{v} \langle [M_\ell]_{ij} [\mathcal{Y}_2]_{ij} \rangle = \frac{1}{2} \left[\langle [y_2]_{ij}^2 \rangle - \langle [y_1]_{ij}^2 \rangle \right] \cos \beta \sin \beta + \frac{1}{2} \langle [y_1]_{ij} [y_2]_{ij} \rangle (\cos^2 \beta - \sin^2 \beta) .\tag{3.37}$$

The last term is zero by assumption, but the first is only zero if $\beta = 0, \pi/2$ or in the generic 2HDM when y_2 and y_1 are induced at the same order ($\varepsilon = 1$). To address this, we construct two matrices A and B , which have identical flavor structures with y_1 , and are statistically independent from one another. These are

$$A = \frac{1}{Nv} M_\ell \quad , \quad B = \frac{1}{N} [-\varepsilon y_1 \sin \beta + z_2 \cos \beta] \quad , \quad N = \sqrt{\cos^2 \beta + \varepsilon^2 \sin^2 \beta} .\tag{3.38}$$

The normalization N is included to ensure $\langle B_{ij} \rangle = \langle A_{ij} \rangle = \langle [y_1]_{ij} \rangle$. The lepton masses are expected to be $\langle m_i \rangle = N e_i$, where e_i are the flavor eigenvalues predicted by the underlying theory. Furthermore, as described in the previous section, when A is diagonalized, B will manifest as the flavor structure matrix $\bar{B}_{ij} = a_{ij} S_{ij}^F$ where a_{ij} are $\mathcal{O}(1)$ numbers. With this, we can solve for $\bar{\mathcal{Y}}_2$ in terms of the now-diagonal mass matrix $[\bar{M}_\ell]_{ij} = m_i \delta_{ij}$ and the flavor

	$[\overline{\mathcal{Y}}_2^u]_{ij}$	$[\overline{\mathcal{Y}}_2^d]_{ij}$	$[\overline{\mathcal{Y}}_2^\ell]_{i \neq j}$
Type-I	$\frac{m_i^u}{v} \delta_{ij} \cot \beta$	$\frac{m_i^d}{v} \delta_{ij} \cot \beta$	$\frac{1}{1+\varepsilon^2 \cot^2 \beta} \left[\frac{m_i}{v} \delta_{ij} (1 - \varepsilon^2) \cot \beta + \frac{\varepsilon}{\sin^2 \beta} a_{ij} N S_{ij}^F \right]$
Type-II	$\frac{m_i^u}{v} \delta_{ij} \cot \beta$	$-\frac{m_i^d}{v} \delta_{ij} \tan \beta$	$\frac{1}{1+\varepsilon^2 \tan^2 \beta} \left[-\frac{m_i}{v} \delta_{ij} (1 - \varepsilon^2) \tan \beta + \frac{\varepsilon}{\cos^2 \beta} a_{ij} N S_{ij}^F \right]$
L-S	$\frac{m_i^u}{v} \delta_{ij} \cot \beta$	$\frac{m_i^d}{v} \delta_{ij} \cot \beta$	$\frac{1}{1+\varepsilon^2 \tan^2 \beta} \left[-\frac{m_i}{v} \delta_{ij} (1 - \varepsilon^2) \tan \beta + \frac{\varepsilon}{\cos^2 \beta} a_{ij} N S_{ij}^F \right]$
Flipped	$\frac{m_i^u}{v} \delta_{ij} \cot \beta$	$-\frac{m_i^d}{v} \delta_{ij} \tan \beta$	$\frac{1}{1+\varepsilon^2 \cot^2 \beta} \left[\frac{m_i}{v} \delta_{ij} (1 - \varepsilon^2) \cot \beta + \frac{\varepsilon}{\sin^2 \beta} a_{ij} N S_{ij}^F \right]$

Table 3.1: The extra Yukawa couplings (to H_2 in the Higgs basis) for each discrete symmetry case. Here, L-S stands for Lepton-Specific.

structure matrix S^F , to arrive at the main result of this section,

$$[\overline{\mathcal{Y}}_2]_{ij} = \frac{1}{1 + \varepsilon^2 \tan^2 \beta} \left[-\frac{1}{v} m_i \delta_{ij} (1 - \varepsilon^2) \tan \beta + \frac{\varepsilon}{\cos^2 \beta} a_{ij} N S_{ij}^F \right]. \quad (3.39)$$

The N will drop when substituting the flavor eigenvalues $\{e_i\}$ for the lepton masses $\{m_i\}$. For example, the Cheng-Sher ansatz, $S_{ij}^F = \sqrt{e_i e_j} \sim \sqrt{m_i m_j} / N v$. As a consistency check, we take the limiting cases in this formalism. In the limit $\varepsilon \rightarrow 0$, we recover the type-II 2HDM expression. For $\varepsilon \rightarrow 1$ we are in a generic 2HDM with no discrete symmetries and $\tan \beta$ is phenomenologically irrelevant, and indeed it drops out of Eq. 3.39. For $\tan \beta = 0$, we are diagonalizing entirely y_1 and so $\overline{\mathcal{Y}}_2 \sim \varepsilon N S^F$. For $\tan \beta \rightarrow \infty$, the mass matrix is fed only by the loop-induced Yukawa, and so the orthogonal component is y_1 and is effectively ε^{-1} -enhanced, $\overline{\mathcal{Y}}_2 \sim N S^F / \varepsilon$. In the intermediate range $1 < \tan \beta < \varepsilon^{-1}$, there is a $\tan^2 \beta$ enhancement of the flavor violating terms as we interpolate between the ε and ε^{-1} regimes, and this will have important phenomenological consequences which we explore in Sec. 3.7.

Although we performed this analysis in the Type-II case, this discussion generalizes to any of the discrete symmetry cases. With the two Higgs doublets assigned opposite charges under a discrete symmetry, there are essentially four possibilities for assigning charges to u_R , d_R , and ℓ_L .² We enumerate the resulting β -dependence of the extra Yukawas in Table 3.1.

². Picking the up quarks to couple to Φ_1 , then the enumeration is two possibilities each for d/ℓ assignments.

3.4 Dipoles in the Low Energy EFT

In many cases, the predominant constraints on the 2HDM lepton Yukawas come from the loop-induced lepton dipole operators. These effective operators contribute to the electric and magnetic dipole moments of the leptons and to their flavor-violating decay widths. In this section, we review the 2HDM contributions to dipoles, and later we evaluate the experimental constraints. In the following, we consider the 2HDM as defined at the scale of the EWSB vev v , and present the matching contributions to the Low Energy EFT.³ There are two dipole operators in the low energy EFT. In a chiral operator basis, the dipole operators are written in the form

$$\mathcal{L}_{\text{eff}} \supset -\bar{e}^i \sigma^{\mu\nu} \left([L_L]_{ij} P_L + [L_R]_{ij} P_R \right) e^j F_{\mu\nu} + h.c. , \quad (3.40)$$

where $P_{L/R} = (1 \mp \gamma^5)/2$. The hermitian conjugate implies that the physical parameters are the projections onto the components $L_L \rightarrow \frac{1}{2}(L_L + L_R^\dagger)$ and $L_R \rightarrow \frac{1}{2}(L_L^\dagger + L_R)$. Although the 2HDM contributions to the dipole operators begin at 1-loop, the 2-loop diagrams typically dominate, especially for the $\mu - e$ block. This is primarily because the one-loop diagram has an extra $Y_{\ell\ell}$ and requires a chirality flip, both of which are avoided in the 2 loop diagrams. Still, the two loop diagrams include an additional loop factor and some extra gauge couplings, but it turns out that the suppression of the 1-loop diagram is very strong.

The diagrams that dominate $L_{L/R}$ in most of the available parameter space are the 2-loop Barr-Zee diagrams, displayed in Fig. 3.1. There are two varieties, with either a top or $W^\pm$ running in the loop. The contribution to L_R of the top Barr-Zee diagram using the

3. Due to the small couplings, we neglect effects of RG between the EW scale and the scale of the precision experiments. A full treatment of the running requires matching and running involving multiple scales, see e.g. [43].




Figure 3.1: Diagrams which give the largest contribution to the dipole coefficients $L_{L/R}$. There are several additional subleading 2-loop diagrams, which are included in our analysis and found in [1]. Furthermore, the one-loop contribution can be relevant for τ decays, but these constraints are subdominant to LHC searches.

three neutral scalars can be written (re-expressing the results of Ref. [1])

$$[L_R^{u_k}]_{ij} = \frac{3e\alpha Q_u^2}{32\pi^3} \sum_{n=1}^3 \frac{m_{u_k}}{M_{h_n}^2} [\bar{Y}_n^\ell]_{ij} \left[f_2(x_{u_k h_n}) \text{Re} [[\bar{Y}_n^u]_{kk}] + i f_1(x_{u_k h_n}) \text{Im} [[\bar{Y}_n^u]_{kk}] \right] \quad (3.41)$$

$$[\bar{Y}_n^\ell]_{ij} = \sum_r q_{nr} [\bar{\mathcal{Y}}_r^u]_{ij} \quad , \quad [\bar{Y}_n^u]_{ij} = \sum_r q_{nr}^* [\bar{\mathcal{Y}}_r^u]_{ij}$$

Where $\bar{Y}_k^f$ are the Yukawa matrices of the neutral mass eigenstates to a fermion f , indexed by $h_k = \{h, H, A\}$. Note that for the down-type fermions, the expression changes slightly because $\bar{Y}_n^{d/\ell} = \sum_r q_{nr} \mathcal{Y}_n^{d/\ell}$. The loop functions f_i are given throughout [1], and we list the relevant ones in our Appendix B. A second, related loop is given by using the charged $H^\pm$ boson, which contributes

$$[L_R^{tH^\pm}]_{ij} = \frac{e\alpha}{128\pi^3 s_w^2} \frac{m_t}{M_{H^\pm}^2} [\bar{\mathcal{Y}}_2]_{ij} \left[[\bar{\mathcal{Y}}_2^u]_{tt} f_5(x_{tH^\pm}, x_{tW}) + \frac{m_b}{m_t} [\bar{\mathcal{Y}}_2^d]_{bb} f_6(x_{tH^\pm}, x_{tW}) \right]. \quad (3.42)$$

The last important contribution is the Barr-Zee W loop digram, and is written

$$[L_R^{h,\text{kin}}]_{ij} = \frac{e\alpha}{128\pi^3} \frac{1}{v} \sum_{k=2}^3 q_{kr} [\mathcal{Y}_r]_{ij} q_{k1} F_{\text{kin}}(m_{h_k}, m_H^\pm, m_h, s_w) \quad (3.43)$$

where F_{kin} is defined in Appendix B. Lastly, for the tau decay modes, the one-loop diagram




Figure 3.2: Plots of the 1 and 2 Loop contributions to $[L_R]_{e\mu}$. For these plots we take $[\bar{\mathcal{Y}}_2^f]_{ii} = m_i^f/v$ and $[\bar{\mathcal{Y}}_2^\ell]_{i \neq j} = \sqrt{m_i m_j}/v$. For the $L_R^{H^\pm \gamma}, L_R^{H^\pm Z}, L_R^{H^\pm Z}$ diagrams we also take $Z_3 = Z_7 = 1$ as a benchmark for evaluating their magnitudes. The full expressions for these diagrams is given in Ref. [1], with a conversion factor of $L_R^i = em_\mu/(16\pi^2 v^2)c_R^i$. Note that the three dominant contributions are those listed in further detail in the text: L_R^{th} , $L_R^{tH^\pm}$, and $L_R^{h,kin}$. Note as well that the signs of the amplitudes are depicted by solid/dashed styles - and a cancellation is occurring in certain ranges between the three dominant diagrams.

will be important. It is

$$[L_R^{1 \text{ loop}}]_{rs} = \frac{e}{64\pi^2} \sum_{k=1}^3 \sum_{i=1}^3 \frac{1}{M_{h_k}^2} \left[m_i^\ell [\bar{Y}_k^\ell]_{ri} [\bar{Y}_k^\ell]_{is} \left(3 + 2 \log(x_{\ell_i h_k}) \right) - \frac{1}{3} m_r^\ell [\bar{Y}_k^\ell]_{ri}^* [\bar{Y}_k^\ell]_{is} \right] \quad (3.44)$$

Where as above, $[\bar{Y}_k^\ell]_{ij} = \sum_p q_{kp} [\bar{\mathcal{Y}}_p^\ell]_{ij}$.

There are several other loops that are subleading in magnitude, including e.g. Barr-Zee diagrams with the lighter fermions. To display the relative magnitudes of these diagrams, we plot their contributions to $L_{\mu e}$ for $Z_6 = \pm 1$ as a function of the Heavy Higgs and charged Higgs masses in Fig. 3.2, assuming the heavy and charged Higgs masses are degenerate. In this plot, we take the diagonal couplings to follow the natural pattern $[\bar{\mathcal{Y}}_2^f]_{ii} = \frac{m_i^f}{v}$, and the top loop changes sign if the sign of $[\bar{\mathcal{Y}}_2^u]_{tt}$ is flipped. The relative magnitudes of the 2-loop diagrams are independent of the value of $[\bar{\mathcal{Y}}_2^\ell]_{e\mu}$, but the 1-loop diagram instead contains the factor $[\bar{\mathcal{Y}}_2^\ell]_{\mu\tau}^* [\bar{\mathcal{Y}}_2^\ell]_{\tau e}$. Thus, to standardize the comparison, we take the Cheng-Sher Ansatz $[\bar{\mathcal{Y}}_2^\ell]_{ij} = \sqrt{m_i m_j}/v$ as a benchmark. There are subleading diagrams that depend

on Z_3 and Z_7 parameters of the 2HDM potential. We take these to be unity in these plots, although increasing them to the perturbative limit $Z_{3,7} = 4\pi$ still results in these diagrams being subdominant. Fig. 3.2 motivates a more focused analysis using only the $L_{L/R}^{th}$, $L_{L/R}^{tH^\pm}$, and $L_{L/R}^{h,kin}$ diagrams, since these contributions are typically at least an order of magnitude larger than the others. We do not show the τ LFV modes because they provide constraints subdominant to the LHC constraints, but for the τ , the 1-loop contribution can be comparable to the 2-loop. In the rest of our plots, we use all contributions listed in Ref. [1], except for the top loop, for which we instead use a generalized version of Ref. [44].

3.4.1 *Blind Spots in the Dipole Operator $\mu - e$ Block*

In the $\mu - e$ flavor entries of the dipole operators, there are several diagrams of the same order, with mixed signs (Fig. 3.2). This suggests that there may be regions of parameter space with destructive interference between the largest diagrams, relaxing the bounds from $\mu \rightarrow e\gamma$ and electron EDM. To explore the cancellation regions, we will assume that the quark sector follows aligned flavor violation—that is, there is only one flavor spurion that sets the quark masses, and therefore the Yukawas of H_2 will be $\overline{\mathcal{Y}}_2^u \propto M^u$. In particular, this additional assumption implies that there are no new CP-violating phases in the up-quark sector, and so the phase that generates the electron EDM will be generated purely by the lepton Yukawas. Consequently, the cancellation conditions for electron EDM and flavor violating observables will be the same, and to compute the condition we must only consider the terms in the dipole operators that are proportional to $\overline{\mathcal{Y}}_2^\ell$. More concretely, we first break the dipole operators into

$$[L_R]_{ij} = \frac{m_i}{v} \delta_{ij} I_1 + [\overline{\mathcal{Y}}_2^\ell]_{ij} I_2 + \mathcal{O}(m_\ell^2/v^2) + \dots \quad (3.45)$$

Where $I_{1,2} \equiv I_{1,2}(m_t, \{m_{hk}\}, s_w, \beta - \alpha)$ are functions of the mixing angles, heavy fermion masses, Higgs masses, and electroweak parameters. The terms on the right side are higher order in the lepton masses (but, especially in the case of the hyper-sensitive electron EDM, possibly still important).

Because the top quark and the W -bosons do not introduce a CP-violating phase, I_1 will be mostly real (although suleading one loop effects can introduce a phase) and the electron EDM will have small dependence on the first term. Now expanding I_2 to make the dependence on $[\overline{\mathcal{Y}}_2^u]_{tt}$ explicit,

$$[L_R]_{ij} = \frac{m_i}{v} \delta_{ij} I_1 + \frac{e\alpha}{24\pi^3 v} [\overline{\mathcal{Y}}_2]_{ij} \left[F_0 + \frac{[\overline{\mathcal{Y}}_2]_{tt}}{m_t/v} F_{tt} \right], \quad (3.46)$$

where $F_{0/tt} \equiv F_{0/tt}(m_t, \{m_{hk}\}, s_w, \beta - \alpha)$. The terms proportional to $[\overline{\mathcal{Y}}_2^u]_{tt}$ are contained in the top loop Barr-Zee diagrams, and thus F_{tt}

$$\begin{aligned} F_{tt} = & \frac{m_t^2}{m_h^2} f_2(x_{th}) \cos^2(\beta - \alpha) + \frac{m_t^2}{m_H^2} f_2(x_{tH}) \sin^2(\beta - \alpha) + \frac{m_t^2}{m_A^2} f_1(x_{th}) \\ & + \frac{3}{16s_w^2} \frac{m_t^2}{m_{H^\pm}} f_5(x_{tH^\pm}, x_{tW}) \end{aligned} \quad (3.47)$$

The diagrams contributing to F_0 are dominated by the top loop and W loops, as can be seen in Fig. 3.2. Because these diagrams dominate, it's a good approximation that

$$F_0 \approx \cos(\beta - \alpha) \sin(\beta - \alpha) \left(\frac{m_t^2}{m_h^2} f_2(x_{th}) - \frac{m_t^2}{m_H^2} f_2(x_{tH}) - \frac{3}{16} F_{h,\text{kin}}(\{m_{hk}\}, s_w) \right). \quad (3.48)$$

The cancellation condition can be expressed as

$$\frac{[\overline{\mathcal{Y}}_2^u]_{tt}}{m_t/v} \stackrel{\text{cancel}}{\approx} - \frac{F_0(m_t, \{m_h\}, s_w, \beta - \alpha)}{F_{tt}(m_t, \{m_h\}, s_w, \beta - \alpha)}, \quad (3.49)$$

where the functions on the right hand depend on $\log$ or $\log^2$ of the large heavy Higgs masses.

Smaller loop corrections can shift the blind spot slightly by perturbing the numerator in Eq. 3.49. This expression can be further simplified in the $m_H \gg v$ limit. In this limit, we can take the approximation $\cos(\beta - \alpha) \approx Z_6 v^2 / m_H^2$, and the heavy Higgs masses are constrained to be nearly degenerate. Therefore,

$$\frac{[\overline{\mathcal{Y}}_2^u]_{tt}}{m_t/v} \stackrel{\text{cancel}}{\approx} Z_6 \frac{v^2}{m_t^2} \times \frac{\frac{3}{16} F_{h,\text{kin}}(\{m_{hk}\}, s_w) - \frac{m_t^2}{m_h^2} f_2(x_{th})}{f_2(x_{tH}) + f_1(x_{tH}) + \frac{3}{16s_w^2} f_5(x_{tH}, x_{tW})} \quad (m_H \gg v). \quad (3.50)$$

The ratio of loop factors on the right hand side is $\mathcal{O}(0.5)$ for most heavy Higgs masses. We provide a numerical test of Eq. 3.49 in Fig. 3.5.

3.5 Experimental Constraints

3.5.1 Electric Dipole Moment of the Electron

The electric dipole moment (EDM) of the electron has been measured very precisely and provides constraints on new CP-violating phases in theories beyond the SM. Most recently, the EDM of the electron was measured to be

$$|d_e| \leq 4.1 \times 10^{-30} e \text{ cm} [45]. \quad (3.51)$$

And in the 2HDM, d_e is related to $L_{L/R}$ through

$$d_e = \text{Im}[L_R]_{ee} - \text{Im}[L_L]_{ee} \quad (3.52)$$

The implication of this constraint depends on the loop level at which the dipole operator is generated. In the 2HDM, it is generated at 2 loops, providing constraints in the 10s of TeVs on the masses of heavy states in the CP-violating 2HDM. In most CP-violating scenarios, this constraint reaches far beyond the scales probed at the LHC or through other precision

Observable	Limit	Experiment	Year
BR(μ Au $\rightarrow$ e Au)	$\leq 7 \times 10^{-13}$ [46]	SINDRUM-II	Current
BR($\mu \rightarrow e\gamma$)	$\leq 1.5 \times 10^{-13}$ [47]	MEG-II	Current
BR($\mu \rightarrow e\gamma$)	$\leq 6 \times 10^{-14}$ [47]	MEG-II	Late 2020s
BR(μ Al $\rightarrow$ e Al)	$\leq 6 \times 10^{-16}$ [48]	Mu2e	Late 2020s
BR($\mu \rightarrow 3e$)	$\lesssim \mathcal{O}(10^{-16})$	Mu3e	~ 2030 s
d_e	$\leq 4.1 \times 10^{-30} e \text{ cm}$ [45]		Current

Table 3.2: Current and future constraints in muon flavor violating decays and electron EDM.

measurements. Continuing under the assumption that the quark sector follows aligned flavor violation, the dipole operators constrain the imaginary component of $[\overline{\mathcal{Y}}_2^\ell]_{ee}$.

In order to explore parameter space that avoids the electron EDM bound, there is essentially only two options. First, since we have assumed aligned flavor violation in the quark sector, one could additionally impose CP conservation in the lepton sector. Second, one could impose hermiticity on the Yukawas. In this scenario, all the Yukawas are near-hermitian, which is a relic of a generalized parity symmetry that is broken at a higher scale. As seen in Eq. 3.52, hermiticity of $L_{L/R}$ imposes zero electron EDM, but the combination of the diagonalization matrices for the up and down quark DOFs, $V^{\text{CKM}} = U_L^u (U_L^d)^\dagger$, can still be arbitrary, allowing for the CKM to be the only source of CP-violation. Hermiticity can realistically arise in UV completions of the 2HDM, in particular in Left-Right symmetric models. Here, the generalized parity symmetry $\ell_L \leftrightarrow \ell_R^*$, $\Phi \rightarrow \Phi^\dagger$ imposes hermiticity on the Yukawas, which remains after the parity-symmetric phase is broken down to the SM. Loop corrections can violate the hermiticity of the Yukawas, but in a way that is suppressed both by a loop factor and by several insertions of the heaviest lepton mass m_τ .

3.5.2 Flavor Violating μ Decays

Flavor violating muon decays are a powerful probe of new physics at the EW scale. The long lifetime of the muon allows for uniquely high sensitivity, especially to new particles




Figure 3.3: Contributions to the $\mu \rightarrow e$ conversion rate. The blob on the right side represents the sum of two loop diagrams which contribute to the effective dipole coefficients $L_{L/R}$ (Sec. 3.4). There are additional two loop contributions, such as box diagrams, whose analytic representations are not yet available in the literature and are therefore not included in this analysis.

with EW-scale masses. In Table 3.2 we present the current and projected future bounds for flavor violating muon decays. In the scenarios we consider, $\mu \rightarrow e\gamma$ tends to be the most stringent constraint. The 2HDM induces $\mu \rightarrow e\gamma$ through the effective flavor violating coefficients $[L_{L/R}]_{e\mu}$. The current experimental constraint is from the second run of the MEG experiment (MEG-II), which placed a bound of $\text{BR}(\mu \rightarrow e\gamma) \leq 1.5 \times 10^{-13}$ (90 %). The upcoming MEG-II run is expected to improve this limit to $\leq 6 \times 10^{-14}$ [49]. In terms of the coefficients of the dipole operators $L_{L/R}$, the decay rate for $\mu \rightarrow e\gamma$ is given by

$$\Gamma(\mu \rightarrow e\gamma) = \frac{m_\mu^3}{4\pi} \left(|[L_L]_{e\mu}|^2 + |[L_R]_{\mu e}|^2 \right) \quad (3.53)$$

A closely related decay channel is $\mu \rightarrow 3e$, since the dominant contribution comes from attaching the γ leg of the dipole vertex to a pair of fermions. The rate for $\mu \rightarrow 3e$ is

$$\Gamma(\mu \rightarrow 3e) = \frac{\alpha m_\mu^3}{12\pi^2} \left| \log \frac{m_e^2}{m_\mu^2} - \frac{11}{4} \right| \times \left(|[L_L]_{\mu e}|^2 + |[L_R]_{e\mu}|^2 \right) \quad (3.54)$$

which is suppressed compared the $\mu \rightarrow e\gamma$ rate. Still, in the coming years Mu3e will increase sensitivity in the $\mu \rightarrow 3e$ channel, likely surpassing the competing bound from $\mu \rightarrow e\gamma$. In our projected bounds, we include the nearer-term bounds from Mu2e and Meg-II, but we quote the increased sensitivity at Mu3e in Table 3.2.

The $\mu \rightarrow e$ conversion search channel currently provides a weaker constraint on the flavor violating couplings of the Higgs, but the hierarchy will flip with the Mu2e experiment. As opposed to $\mu \rightarrow e\gamma$, which searches for the decays of free muons, $\mu \rightarrow e$ conversion first traps the in bound states around ionized heavy nuclei. The muon can then decay via photon or higgs exchange with the quarks in the nucleus, which allows for $\mu \rightarrow e$ conversion to be kinematically viable. This process includes a contribution from the same effective dipole operator which generates $\mu \rightarrow e\gamma$, but has additional contributions from virtual Higgs exchange with the nucleus, displayed in 3.3. The current experimental limit is from muons captured by gold atoms in the SINDRUM II experiment, which constrains the branching ratio for muon capture to $\text{BR}(\mu + Au \rightarrow e + Au) \leq 7 \times 10^{-13}$ (90%) [46]. In the coming years, the Mu2e experiment is expected to greatly improve on this result. They project a sensitivity of $\text{BR}(\mu + Al \rightarrow e + Al) \leq 6 \times 10^{-16}$ (90%), an improvement of, naively, about four orders of magnitude. We include a brief summary of the $\mu \rightarrow e$ rates in Appendix E.

3.5.3 Flavor Violating Higgs Decays at the LHC

The observation of the Higgs decay channel $h \rightarrow \mu\tau$ would be a direct, unambiguous measurement of flavor violation in the Higgs sector. For the flavor violating modes $h \rightarrow \mu\tau$ and $h \rightarrow \mu e$, the constraints on the branching ratios are a factor of ~ 4 weaker than for the $h \rightarrow \mu\mu$ mode, primarily due to increased difficulty of identifying and characterizing τ s and electrons and the differing backgrounds. Still, ATLAS and CMS searches for flavor violating Higgs decays provide the most direct probe of its flavor violating couplings. Notably, in the most recent ATLAS searches for these effects, there are slight excesses ($\sim 2\sigma$) in both the $h \rightarrow \tau\mu$ and $h \rightarrow \tau e$ channels [50]. Equivalent CMS searches have placed limits that are similar, however they do not observe any excess [51]. We do not attempt to fit for these excesses, although we note that they are in a region of parameter space loosely compatible with the 2HDM under the assumptions we make in this paper. In particular, certain regimes

Observable	Current Limit	HL-LHC Limit	Constraint
$\text{BR}(h \rightarrow \mu\tau)$	$\leq 0.18\%$ [50]		$\cos(\beta - \alpha)\overline{\mathcal{Y}}_{2\langle\mu\tau\rangle} \leq 9.5 \times 10^{-4}$
$\text{BR}(h \rightarrow \tau e)$	$\leq 0.20\%$ [50]		$\cos(\beta - \alpha)\overline{\mathcal{Y}}_{2\langle\tau e\rangle} \leq 9.5 \times 10^{-4}$
$\text{BR}(h \rightarrow \mu e)$	$\leq 6.2 \times 10^{-5}$ [52]		$\cos(\beta - \alpha)\overline{\mathcal{Y}}_{2\langle\mu e\rangle} \leq 1.8 \times 10^{-4}$
κ_W	1.00 ± 0.05 [53]	± 0.016 [54]	$ \cos(\beta - \alpha) \leq 0.40$
κ_τ	0.94 ± 0.06 [53]	± 0.019 [54]	
κ_b	$0.89^{+0.10}_{-0.09}$ [53]	± 0.036 [54]	

Table 3.3: LHC Constraints on the light Higgs couplings. We only list those used in our analysis. Note that we use the shorthand notation $\overline{\mathcal{Y}}_{2\langle ij \rangle} \equiv \sqrt{|\overline{\mathcal{Y}}_{2ij}|^2 + |\overline{\mathcal{Y}}_{2ji}|^2}$. The Higgs FV decay limit are shown at 90%, while the intervals on the κ quantities are 1σ .

of the $\mu\tau$ -Specific ansatz with larger $\mathcal{O}(1)$ coefficients may fit this excess well. It's possible in both the Type-III and loop induced FV scenarios.

In the 2HDM, flavor violating Higgs decays enter at tree-level, producing potentially large deviations from the SM expectations. The branching ratio of $h \rightarrow \ell_i \ell_j$ is given by

$$\frac{\text{BR}(h \rightarrow \ell_i \ell_j)}{\text{BR}(h \rightarrow b\bar{b})} = \frac{v^2}{m_b^2} \cos^2(\beta - \alpha) \left(|\overline{\mathcal{Y}}_{2ij}^\ell|^2 + |\overline{\mathcal{Y}}_{2ji}^\ell|^2 \right), \quad (3.55)$$

and $\text{BR}(h \rightarrow b\bar{b})$ is measured with a central value of 58% . The current and future HL-LHC bounds from ATLAS are given in Table 3.3.

3.5.4 Higgs Diagonal Couplings

To measure the observations at the LHC against the predictions of the SM, the couplings of the Higgs are typically quantified via the κ values, defined for fermions and vectors as

$$\kappa_f = \frac{y_f^{\text{obs}} v}{m_f} \quad , \quad \kappa_V = \frac{g_{hVV}^{\text{obs}} v}{2m_V^2} \quad , \quad (3.56)$$




Figure 3.4: Production modes for $H/A/H^\pm$.

where y_f^{obs} and g_{hVV}^{obs} are the observed couplings for the 125 GeV boson at the LHC. The SM predicts $\kappa_{f/V} = 1$ for all species. In the 2HDM, there are only special cases in which $\kappa_{f/V} = 1$, such as the alignment limit. For the EW bosons, for example,

$$\kappa_V^{2\text{HDM}} = \sin(\beta - \alpha) , \quad (3.57)$$

so the SM prediction is only recovered if $\cos(\beta - \alpha) = 0$. The 2HDM predictions for κ_f depend on both $\cos(\beta - \alpha)$ and $(\overline{\mathcal{Y}}_2)_{ii}$. We can express the κ_f by employing the parametrization in Eq. 3.22:

$$\kappa_i = \sin(\beta - \alpha) + \frac{v}{m_i} [\overline{\mathcal{Y}}_2^f]_{ii} \cos(\beta - \alpha) . \quad (3.58)$$

The constraints on κ_τ , κ_b , and κ_W will be most relevant and their experimental limits are listed in Table 3.3.

3.5.5 Heavy Higgs Searches

Experiments at the LHC have not yet observed evidence for fundamental particles outside of the SM. This places constraints on the 2HDM, which is expected to include several new particles with masses not too much higher than the electroweak scale. The collider implications of the 2HDM are highly dependent on the scenario - the four discrete symmetry cases and the generic case all have distinct phenomenologies. In the generic scenario, the diagonal

elements of the $\overline{\mathcal{Y}}_2^f$ matrices are expected to exhibit a similar hierarchy as the lepton masses, so ggF with decays to tops are the most constraining search channel. These limits are set by interference pattern searches at the LHC [55].

In the discrete symmetry cases, the collider constraints follow the typical 2HDM constraints, with $\tan\beta$ -enhancements facilitating the searches and allowing for better constraints on heavy Higgses. There are a number relevant searches - we include those that set important or leading limits for our plots. In addition to the $H \rightarrow tt$ search described above, the $pp \rightarrow H/A \rightarrow \tau\tau$ channel provides important constraints in the large $\tan\beta$ regime of the Type-II 2HDM. In this case the $b\bar{b}H$ production mode is enhanced by $\tan\beta$, making this analysis especially useful. The other important search channel is $pp \rightarrow A \rightarrow Zh$, which gains strength in the low $\tan\beta$ regime of both type-II and Lepton-Specific theories. However, this decay channel has zero width in the alignment limit, it originates from the kinetic terms of the Higgs doublets.

To construct the bounds, we run 2HDMC [56] and SuShi [57, 58] over a grid of points in our slices of parameter space. We compute the ggf and $b\bar{b}H$ production cross sections (diagrams shown in Fig. 3.4). We then use the tree-level branching ratios computed with 2HDMC to recast the experimental limits into our parameter space. For the projected improvements at the HL-LHC, we make the optimistic assumption that statistical uncertainties dominate, and rescale cross section limits by $\sqrt{(139 \text{ fb}^{-1})/(3 \text{ ab}^{-1})} = 0.22$. In Table 3.4 we list the search channels which we use as bounds in our studies.

Notably, a 14 TeV muon collider would offer a Higgs-pair-production channel that is dominated by the kinetic term of the second Higgs. This removes potential suppression from the production vertex, providing much stronger bounds that are largely universal across the generic and discrete symmetry 2HDMs. A proposal for such an experiment is given in Ref. [59], and although we do not include these projections they would vastly exceed the LHC sensitivities and provide coverage for almost any $\tan\beta$.

Channel	Production Channels	Ref.
$pp \rightarrow H/A \rightarrow \tau\tau$	ggf/bbh	[60]
$pp \rightarrow H/A \rightarrow tt$	ggf	[55]
$pp \rightarrow A \rightarrow Zh$	ggf/bbh	[61]
$pp \rightarrow H \rightarrow ZZ$	ggf	[62]

Table 3.4: List of analyses used in the recasting of heavy Higgs constatints. All are with a luminosity of 139 fb^{-1} .

3.6 Phenomenology of the Generic (Type-III) 2HDM

First, we explore the case when the Lepton Yukawas are induced to the same order for both Higgs doublets, or when $\varepsilon = 1$ in Eq. 3.39. This is the most generic scenario in the 2HDM, although one must of course still make an ansatz for the flavor structures. In this section we apply the experimental constraints outlined in Section 3.5 to the generic 2HDM, and compare the role of dipole observables, direct Higgs decay channels, and LHC searches.

Ultimately, the generic 2HDM is disfavored by null results in $\mu \rightarrow e\gamma$ and, in CP-violating scenarios, by $e\text{EDM}$. In the CP-conserving case with the $\mu\tau$ -Specific ansatz, where both $\mu \rightarrow e\gamma$ and $e\text{EDM}$ are zero by symmetry, the LHC $h \rightarrow \mu\tau$ searches are approaching a well-motivated benchmark of $[\overline{\mathcal{Y}}_2^\ell]_{\langle e\mu \rangle} = \sqrt{2m_\mu m_\tau}/v$.

3.6.1 Cheng-Sher and Dipole Blind Spots

Although the generic Cheng-Sher scenario of the 2HDM most effectively probed by dipole observables, blind spots appear in the parameter space where the bounds from both $e\text{EDM}$ and $\mu \rightarrow e\gamma$ relax. The conditions for this cancellation were computed in Section 3.4.1, and here we present a phenomenological study of the possibility of living in such a blind spot.

The first natural question is the degree of tuning required to enter the blind spot. The tuning will be performed on the top coupling in the generic 2HDM, and the story changes considerably in the case of discrete symmetries. In the generic 2HDM, the experimental

constraints will impose $[\overline{\mathcal{Y}}_2^u]_{tt} - [\overline{\mathcal{Y}}_2^u]_{tt}^{\text{cancel}} \lesssim \Delta[\overline{\mathcal{Y}}_2^u]_{tt}$. The expected value of the top coupling is $[\overline{\mathcal{Y}}_2^u]_{tt} \sim \mathcal{O}(m_t/v)$, and so we select the tuning measure

$$\Delta_t \equiv \frac{\Delta[\overline{\mathcal{Y}}_2]_{tt}}{m_t/v} . \quad (3.59)$$

For electric dipole moment, expanding around the cancellation point means

$$\Delta_t = \frac{\Delta[\overline{\mathcal{Y}}_2^u]_{tt}}{m_t/v} = \frac{d_e^{\text{exp.}}}{\frac{e\alpha}{12\pi^3 v} F_{tt} \text{Im}([\overline{\mathcal{Y}}_2^\ell]_{ee})} = \frac{0.0017}{F_{tt}} \times \frac{m_e/v}{\text{Im}([\overline{\mathcal{Y}}_2^\ell]_{ee})} \quad (\text{eEDM}) . \quad (3.60)$$

Since F_{tt} is $\mathcal{O}(1)$, this result shows that avoiding electron EDM requires a very precise tuning. However, as mentioned previously, if the lepton Yukawas are CP-conserving or Hermitian, then the electron EDM will be avoided. However, there is still a (milder) tuning required to avoid $\mu \rightarrow e\gamma$ in the Cheng-Sher ansatz. This tuning can be expressed analytically as

$$\Delta_t = \frac{\Delta[\overline{\mathcal{Y}}_2^u]_{tt}}{m_t} = \frac{\sqrt{4\pi\Gamma_\mu \text{Br}(\mu \rightarrow e\gamma)^{\text{limit}}}}{\frac{(m_\mu)^{\frac{3}{2}}}{24\pi^3 v} F_{tt} [\overline{\mathcal{Y}}_2^\ell]_{\langle e\mu \rangle}} = \frac{0.04}{F_{tt}} \times \frac{\sqrt{2m_e m_\mu}/v}{\sqrt{|[\overline{\mathcal{Y}}_2^\ell]_{e\mu}|^2 + |[\overline{\mathcal{Y}}_2^\ell]_{\mu e}|^2}} \quad (\text{MEG-II}) \quad (3.61)$$

Which is in the range $\mathcal{O}(0.05)$ for most of our parameter space.

If one accepts this level of tuning and enters the blind spot, $\mu \rightarrow e$ conversion offers an important alternative constraint that receives contributions not only from the dipole effective operator but also from tree-level Higgs exchange (Fig. 3.3). In the limit that $[\overline{\mathcal{Y}}_2^{u/d}]_{qq} \ll m_q/v$ and $L_{L/R} \approx 0$, the $\mu \rightarrow e$ conversion rate is

$$\Gamma_{\mu \rightarrow e} \approx \frac{m_\mu^3 m_p^2}{4m_h^4 v^2} \cos^2(\beta - \alpha) \left[\sum_{q \in \{u,d,s\}} (f_q^p + f_q^n) \right]^2 \left[|[\overline{\mathcal{Y}}_2]_{e\mu}|^2 + |[\overline{\mathcal{Y}}_2]_{\mu e}|^2 \right] . \quad (3.62)$$




Figure 3.5: A characterization of the $\mu \rightarrow e\gamma$ “Blind Spot,” where the loop contributions of the 2HDM to the dipole operators $L_{L/R}$ experience destructive interference. The dashed line is the analytical formula in Eq. 3.49. In these plots, we assume the relation $\cos(\beta - \alpha) = Z_6 v^2 / (m_H^2 - m_h^2)$. The dark blue excluded region assumes $\bar{\mathcal{Y}}_2 \sim \sqrt{m_i m_j} / v$, which is the Cheng-Sher [CS] structure. The purple constraint is from Heavy Higgs searches at the LHC with a top pair final state, and the gray is the constraints from κ_t and κ_W measurements at the LHC. The $\mu \rightarrow e\gamma$ constraint is from current MEG-II results.

This can be converted into the constraint

$$\cos(\beta - \alpha) \sqrt{|[\bar{\mathcal{Y}}_2]_{e\mu}|^2 + |[\bar{\mathcal{Y}}_2]_{\mu e}|^2} \lesssim 1.8 \times 10^{-5} \quad (\text{SINDRUM-II, } L_L = L_R = 0) . \quad (3.63)$$

This is an order of magnitude stronger than the bounds on the μe couplings from direct measurements at the LHC. Furthermore, in the discrete symmetry cases of the 2HDM, this rate can be greatly enhanced. The extra quark couplings, which we have assumed small in the above equation, can receive contributions enhanced by the ratio of the Higgs vevs. This effect will be encountered in the next section.

If $\bar{\mathcal{Y}}_2$ has Cheng-Sher structure, then the $h \rightarrow \tau\mu$ branching ratio is enhanced over $h \rightarrow \mu e$ by a factor m_τ/m_e , and this is strong enough that the $\mu\tau$ channel provides another probe of the dipole blind spot. Specifically, this bound from inferring the hierarchy between $[\bar{\mathcal{Y}}_2]_{e\mu}$

and $[\overline{\mathcal{Y}}_2]_{\tau\mu}$ is

$$\cos(\beta - \alpha) \sqrt{|[\overline{\mathcal{Y}}_2]_{e\mu}|^2 + |[\overline{\mathcal{Y}}_2]_{\mu e}|^2} \lesssim 1.2 \times 10^{-5} \quad (\text{LHC, } L_L = L_R = 0, \text{ Cheng-Sher } \overline{\mathcal{Y}}_2) \quad (3.64)$$

provided by $h \rightarrow \mu\tau$ measurements at the LHC. This is slightly stronger than the competing bound offered by $\mu \rightarrow e$ conversion, although it depends additionally on the $\mathcal{O}(1)$ coefficients of the entries of the $\mu\tau$ entries of $\overline{\mathcal{Y}}_2^\ell$.

In Fig. 3.5, we plot the parameter space where the cancellation becomes strong enough to avoid $\mu \rightarrow e\gamma$ constraints for the Cheng-Sher flavor violating structure. At all heavy Higgs masses, a cancellation is possible - at larger masses the Yukawa tuning requirement grows slightly less strict. There is a complementarity in the different probes of flavor violation, and in particular, the dipole probes are expected to greatly increase sensitivity in the next decade. Together these results will eventually outpace progress at the LHC. They will rule out blind spots in parameter space where cancellations can occur, with the final closure of hyper-tuned cancellations regions provided by direct, tree-level probes of the Higgs flavor violating couplings. The phenomenology of the 2HDM changes dramatically in the discrete symmetry case, with colliders providing constraints that compete with dipole observables. In the following sections we explore these scenarios.

3.6.2 $\mu\tau$ -Specific Scenario

In the CP-conserving $\mu\tau$ -Specific Scenario, the $\mu \rightarrow e\gamma$ constraint does not apply, and the $h \rightarrow \mu\tau$ constraint becomes dominant. In this case, the branching ratio of $h \rightarrow \mu\tau$ is

$$\text{BR}(h \rightarrow \mu\tau) = 0.2\% \times \left(\frac{\cos(\beta - \alpha)}{0.4} \right)^2 \left(\frac{|a_{e\mu}|^2 + |a_{\mu e}|^2}{2} \right) \quad (3.65)$$

To reach the level of what's observed at the LHC, one requires either a large value of $\cos(\beta - \alpha)$ (appealing to tuning to avoid all of the κ constraints), or if $\cos(\beta - \alpha)$ is relaxed to ~ 0.1 , $\sqrt{|a_{e\mu}|^2 + |a_{\mu e}|^2} \gtrsim 4\sqrt{2}$. Since both the κ and $h \rightarrow \mu\tau$ constraints will increase in sensitivity by the roughly the same amount at the HL-LHC, this conclusion will not change dramatically: $h \rightarrow \mu\tau$ does probe the $\mathcal{O}(1)$ parameter space of the Cheng-Sher ansatz, but is strongest constraint comes from the κ probes.

3.7 Phenomenology of the Type-II and Lepton-Specific 2HDMs

In this section, we shift our focus from the generic 2HDM to models with tree-level discrete symmetries, in the scenario where loop corrections violate the $\mathbb{Z}_2$ and introduce new flavor spurions. These loop effects are most important when we require the vev to be generated mainly for the Higgs doublet to which the lepton Yukawas couple at loop-level—in this case, flavor violating observables are enhanced by the ratio of the vevs. As noted in the Sec. 3.3, the flavor violating entries grow like $\tan^2 \beta$, and thus all of the dipole and LHC $h \rightarrow \ell_i \ell_j$ constraints provide an upper bound on $\tan \beta$, dependent on the precise values of $\cos(\beta - \alpha)$ and the heavy Higgs masses.

In this regime, however, $|\kappa_\tau - 1|$ will also be enhanced by the same factor, making it one of the most important constraints. In fact, the κ_τ constraint dominates most of the parameter space. In the limit of large $\tan \beta$ (while still $\tan \beta < \varepsilon^{-1}$), the $[L_R]_{\mu e} \propto \cos(\beta - \alpha) \times \varepsilon \tan^2 \beta$, but $|\kappa_\tau - 1| \propto \cos(\beta - \alpha) \tan \beta$. Although L_R has a faster scaling with $\tan \beta$, it turns out that the constraint from κ_τ is much stronger, and this hierarchy is largely independent from choice of $\cos(\beta - \alpha)$ or Z_6 for the aforementioned reason. However, there are several cases where this line of reasoning breaks down.

First, for $\tan \beta \sim \varepsilon^{-1}$, the numerical coefficient of loop correction to $[\overline{\mathcal{Y}}_2^\ell]_{\tau\tau}$ can be chosen to cancel κ_τ . This leads to yet another blind spot in parameter space, although the location of the κ_τ blind spot is extremely sensitive to the precise value of $[\overline{\mathcal{Y}}_2^\ell]_{\tau\tau}$. We do not pursue

this option further but note it is a possibility. Secondly, the constraints from κ_τ are weakened if there are correlations between $\cos(\beta - \alpha)$ and $\tan \beta$. We consider two such scenarios in the rest of this section.

We also note that the results of this section are sensitive to a number of generic $\mathcal{O}(1)$ coefficients. All of the ansatz coefficients we have chosen to be unity, so the remaining choice is the loop factor ε . We choose it in the range $(0.1 - 1)/(4\pi)^2$, very close to what is expected from a 1-loop contribution, $\varepsilon \sim \mathcal{O}(1) \times (\text{couplings})/(4\pi)^2$. Still, due to the subjectiveness of each choice, if ε is chosen larger, the flavor violating constraints will be more dominant, and vice versa.

3.7.1 The MSSM

First, we present bounds on the MSSM-like case in Fig. 3.6, where we are in the Type-II 2HDM and α is predicted by the tree-level relation

$$\sin \alpha = \frac{-(m_A^2 + m_Z^2) \sin \beta \cos \beta}{\sqrt{((m_A^2 + m_Z^2) \sin \beta \cos \beta)^2 + (m_H^2 \sin^2 \beta + m_Z^2 \cos^2 \beta - m_h^2)^2}} \quad (3.66)$$

In this case, regions below $m_H \sim 700 \text{ GeV}$ are excluded by measurements of κ_τ . All of the usual Type-II LHC constraints apply, however we focus on the two that are most important for heavy Higgs masses above the κ_τ threshold, namely $pp \rightarrow H/A \rightarrow \tau\tau$ and $pp \rightarrow H/A \rightarrow tt$. The dipole couplings impose $\tan \beta \lesssim 40$, however the future dipole constraints will outpace the projected bounds of the HL-LHC.

Related to this scenario is the alignment limit $\cos(\beta - \alpha) = 0$, where the κ_τ constraint is avoided entirely. This limit is achievable also within the context of the MSSM, due to loop corrections that are enhanced with $\tan \beta$ [63]. In the large $\tan \beta$ limit, $\cos(\beta - \alpha)$ is




Figure 3.6: The parameter space for two case of the Type-II 2HDM, in the Cheng-Sher flavor scenario. In the left, the MSSM relation between $\cos(\beta - \alpha)$ and $\tan \beta$ is applied, resulting in a strong κ_τ bound. In the right, $\cos(\beta - \alpha) = 0$, and κ_τ (as well as certain heavy Higgs searches) does not apply. Future bounds from the HL-LHC projections, MEG-II, and Mu2e. Conservatively, we do not project the $H \rightarrow t\bar{t}$ since it is an interference search and the projections are nontrivial.

described by

$$t_\beta c_{\beta-\alpha} \simeq \frac{-1}{m_H^2 - m_h^2} \left[m_h^2 + m_Z^2 + \frac{3m_t^4}{4\pi^2 v^2 M_S^2} \left\{ A_t \mu t_\beta \left(1 - \frac{A_t^2}{4M_S^2} \right) - \mu^2 \left(1 - \frac{A_t^2}{2M_S^2} \right) \right\} \right], \quad (3.67)$$

where $t_\beta \equiv \tan \beta$ and $c_{\beta-\alpha} \equiv \cos(\beta - \alpha)$. In the alignment limit, the light CP-even Higgs behaves exactly like the SM Higgs, so only limits that apply are direct heavy Higgs searches and modifications to the lepton dipoles.

3.7.2 The Dark-side of κ_τ

Despite the importance of the κ_τ , there exist small regions of parameter space with sizeable $\cos(\beta - \alpha)$ for which the κ constraints are effectively blind, which we will call the “dark-side” scenario of the 2HDM. This occurs when the fermion couplings κ_f are near -1 , which is




Figure 3.7: The current (left) and future (right) constraints for the “darkside” ($\kappa_\tau \approx -1$) scenario in a Type-II 2HDM. The dipole constraints disfavor larger heavy Higgs masses in the darkside. The collider searches $A \rightarrow Zh$ and $H \rightarrow \tau\tau$ complement the dipole searches, and jointly they disfavor much of the darkside parameter space. Furthermore, the $\mu \rightarrow e$ conversion bound is enhanced in the Type-II 2HDM due to the enhanced down-type quark couplings. The future bounds, in dashed lines, depicted on the right, are using projections for the HL-LHC, MEG-II, and Mu2e experiments.

consistent with LHC bounds because the sign of the Yukawas cannot be probed by tree-level searches. The dark-side is more easily accessible in the Type-II and Lepton-Specific 2HDMs, where the top Yukawa scales oppositely as the lepton Yukawas with $\tan \beta$. In the Type-I and Flipped 2HDMs, the top coupling is enhanced, and so the dipole constraints are greatly enhanced and this rules out the majority of the parameter space. In the Type-II and Lepton-Specific 2HDMs, the wrong-sign region is defined at tree-level and for $1 \ll \tan \beta \ll \varepsilon^{-1}$ to be

$$\cos(\beta - \alpha) = \frac{2}{\tan \beta} \implies \kappa_\tau \approx -1 \quad (\text{dark-side}) . \quad (3.68)$$

A discussion of the loop corrections to this condition are given in Appendix. C.

By coincidence, this relation also causes the blind spot of the dipoles to appear for moderate ($\lesssim$ TeV) heavy Higgs masses. Furthermore, the blind spot is largely insensitive to




Figure 3.8: The current (left) and future (right) constraints for the “darkside” ($\kappa_\tau \approx -1$) scenario in a Lepton-Specific (Type-X) 2HDM. The blind spot now extends into very large $\tan \beta$, because there is no enhancement of the b -quark Barr-Zee diagram. This region will be explored from above by Mu2e, and from below by the LHC $pp \rightarrow A \rightarrow Zh$ experiment. The future bounds (right side) are using projections for the HL-LHC, MEG-II, and Mu2e experiments.

$\tan \beta$, since both $[\bar{\mathcal{Y}}_2^u]_{tt} \propto \cot \beta$ and $F_0 \propto \cos(\beta - \alpha) \propto \cot \beta$, and both factors drop out of Eq. 3.49. In Type-II the blind spot disappears once the bottom quark becomes the dominant diagram around $\tan \beta \sim 30$.

We display the current and future projected bounds in the dark-side for the Type-II and Lepton-Specific 2HDMs in Figs. 3.7 and 3.8. The blind spot is visible in both figures as a weakening of the $\mu \rightarrow e\gamma$ bound. In the Type-II 2HDM (Fig. 3.7), the blind spot is partially covered by the $\mu \rightarrow e$ conversion bound, due to the enhanced down-type couplings of the heavy Higgs. This tree-level diagram falls like m_H^{-2} , which is why the $\mu \rightarrow e$ conversion bound weakens at larger heavy Higgs masses. The region where the $\mu \rightarrow e\gamma$ probe is weak is precisely collider constraints and $\mu \rightarrow e$ conversion are strongest. In the coming years, the constraints will exclude a significant portion of the open parameter space for loop-induced lepton flavor violation in the darkside.

The Lepton-Specific 2HDM (Fig. 3.8) lacks the enhancement of the down-type couplings.

This completely eliminates the bound from $pp \rightarrow H \rightarrow \tau\tau$, which relies on the enhancement of bottom quark couplings in the production process. It also slightly modifies the shape of the $pp \rightarrow A \rightarrow Zh$ bound, enhancing it due to the suppression of the $A \rightarrow bb$ decay mode. Similarly, the $\mu \rightarrow e$ conversion bound is greatly weakened, as all quark couplings are now suppressed with $\tan\beta$. Finally, the dipole cancellation is of especially high quality and is not cut off by the contribution from bottom quarks as in the Type-II 2HDM. the blind spot parameter space remains open, with heavy Higgs masses to which the LHC $pp \rightarrow A \rightarrow Zh$ search is uniquely sensitive. In the coming years, this region will be largely covered by $\mu \rightarrow e$ conversion and the HL-LHC.

3.7.3 The $\mu - \tau$ Specific Flavor Structure

In the $\mu\tau$ -Specific scenario, $h \rightarrow \mu\tau$ offers an upper bound on $\tan\beta$. Typically, the κ constraints impose $|\cos(\beta - \alpha) \tan\beta| \lesssim 0.1$, and the branching ratio is far below the LHC limits. In the dark-side, however, $\cos(\beta - \alpha)$ is enhanced and

$$\text{BR}(h \rightarrow \mu\tau) = 0.2\% \times (16\pi^2\varepsilon)^2 \left(\frac{\tan\beta}{32}\right)^2 \frac{|a_{\mu\tau}|^2 + |a_{\tau\mu}|^2}{2} \quad (3.69)$$

Where the a_{ij} are the $\mathcal{O}(1)$ numerical coefficients, and we have normalized to the bounds set by ATLAS and CMS. The other bounds from the LHC are the exact same as those in Figs. 3.7 and 3.8. Therefore, in the $\mu\tau$ -Specific scenario, $h \rightarrow \mu\tau$ provides an important upper bound on $\tan\beta$ in the dark-side. The $h \rightarrow \mu\tau$ bound is superseded by the dipoles in the standard Cheng-Sher ansatz, but is irreducible.

3.8 Conclusion

The flavor structure of a second Higgs doublet is not completely free—it is constrained by the same EFT logic that governs the SM, and we have shown that taking this seriously selects

the Cheng-Sher and $\mu\tau$ -Specific ansatzes as natural benchmarks. The coincidence that the dark-side ($\kappa_\tau = -1$) is precisely where the $\mu \rightarrow e\gamma$ blind spot is located provides a target for future experiments to explore. The dark-side remains as a region that cannot be probed by Higgs couplings at the LHC, so heavy Higgs searches and lepton flavor violating probes play a leading role in exploring this possibility. The most relevant scenarios in this regard are the Type-II and Lepton-Specific 2HDMs, and Figs. 3.7 and 3.8 show that for $m_H \lesssim 2$ TeV, the HL-LHC, Mu2e, and MEG-II will jointly close this region. Exploiting their complementarity is key to decisively testing flavor violation in the dark-side of the 2HDM. As we have shown, lepton flavor violation could realistically be the first hint of an extended Higgs sector.

CHAPTER 4

PRECISION HIGGS MEASUREMENTS TO SEARCH FOR NEW ELECTROWEAK-SCALE PHYSICS

4.1 Introduction

In this Chapter, we explore the usage of high energy colliders as precision machines. Typically, high energy colliders are thought of as machines with the purpose of discovery: the Tevatron discovered the top quark, and the LHC discovered the Higgs boson. However, in many cases high energy colliders have provided crucial precision measurements, such as LEP proving there are only three generations of neutrinos through precision Z measurements. This is also the current direction of particle physics experiments: with the SM complete, and no signs of new physics yet, the experimental program of the LHC is shifting towards a goal of a high-precision characterization of the couplings of the Higgs.

As we saw in the last chapter, Higgs couplings to the lepton sector are an important testing ground for the Standard Model, in part because the SM predictions for the leptons are so simple. With the precision of such measurements set to increase significantly at the HL-LHC and Future Circular Collider (FCC), it is crucial to quantify the effects of heavy new physics on collider observables. In particular, we focus on a very specific search channel: the electron Yukawa coupling. This channel is currently very relevant in the planning of the FCC - a strong bound can be placed, but it would require a year of dedicated runtime, during which time not much other physics could be studied. Given this, it is worth assessing the full space of beyond the Standard Model (BSM) theories that could be tested by this measurement, asking which models would be first seen here, how theoretically plausible these models are, and whether future measurements of other observables could test the same theory space more easily. There are many ways of defining ‘plausibility’ of a model, all of which are subjective. We explore a few different criteria, namely: how aligned and electron-specific

does the flavour structure need to be? Is there a large hierarchy of couplings needed in any particular model? How fine-tuned and radiatively stable would such a model be?

Considering the lack of new EW-scale particles, its a well-motivated assumption that the scale of new physics is well above the Electroweak scale, and so the natural framework to use is SMEFT (Sec. 2.3). In SMEFT, there is just one operator that can lead to large enhancements of the electron Yukawas, namely,

$$\mathcal{O}_{eH} = |H|^2 H \bar{l} e. \quad (4.1)$$

Using a spurion analysis one can connect the above operator to other related operators in SMEFT. These operators form a set that all mix under RG flow, and therefore they all generically appear in any UV completion capable of enhancing the electron Yukawa. The main results of this chapter are a characterization of the connections between the electron Yukawa and other observables in a pure SMEFT framework, a presentation of the landscape of models that offer this feature, and a few model-specific case studies. This Chapter is based on the published work [64].

4.2 Enhancing the electron Yukawa in SMEFT

At the electroweak (EW) scale, the effects of decoupled new physics are well described by higher-dimensional operators in the SMEFT framework. In the Warsaw basis [22], there are several operators which can shift the electron Yukawa, however we focus on $\mathcal{O}_{eH}$, defined in Eq. (4.1), as it allows for a non-universal Yukawa coupling modification. Thus we consider the following Lagrangian for the electron Yukawa coupling up to dimension six:

$$-\mathcal{L}_{\text{Yuk}} = y_e H \bar{l} e - \underbrace{\mathcal{C}_{eH} |H|^2 H \bar{l} e}_{\mathcal{O}_{eH}} + \dots, \quad (4.2)$$

where l and e are here the first generation left-handed lepton doublet and right-handed lepton singlet respectively.

Importantly, this new dimension-6 term modifies the relationship between fermion mass and Higgs coupling from the SM prediction. Assuming no operators of dimension higher than those in Eq. (4.2) are important, we may proceed to calculate the ratio of the Yukawa coupling to its SM value:

$$\kappa_e = 1 - \frac{v^3}{\sqrt{2} m_e} \mathcal{C}_{eH}. \quad (4.3)$$

Here and throughout, we assume that the $\mathcal{C}_{eH}$ coefficient is purely real. Imaginary parts have strong indirect constraints from electric dipole moments [65, 66], and are not necessary to generate the Higgs coupling deviations that we focus on. In the SM, the electron Yukawa y_e is very small, of $\mathcal{O}(10^{-6})$ in natural units. It appears that the effects of new physics could easily dominate such a small value, even with a moderate or large separation between v and the scale of new physics Λ . However, understanding both the fine-tuning and RG flow of such a theory leads to theoretical insights on the parameter space. This will be discussed in Sec. 4.3.4. Remaining within the context of the EFT description, in this section we investigate the general implications of the $\mathcal{O}_{eH}$ operator; specifically, how it can be connected to other operators via RG flow, and the model-independent constraints or future tests this can enable for κ_e . We also consider various assumptions concerning the EFT flavour structure, with varying degrees of plausibility and motivation, and how they may be constrained by complementary experimental probes. In Sec. 4.3 we will go on to consider explicit models under certain flavour structure assumptions and the additional EFT operators they give rise to.

4.2.1 Loop level connections to other SMEFT operators

The operator $\mathcal{O}_{eH}$ will mix with other dimension-6 SMEFT operators under RG, producing a set of operators emergent in the IR. Spurion analysis allows us to quickly determine the

operators in this set by considering the flavour symmetries of the SM. Under the global $U(3)^5$ flavour group, the $\mathcal{O}_{eH}$ operator is charged as:

$$\mathcal{O}_{eH} \sim \bar{\mathbf{3}}_L \times \mathbf{3}_e. \quad (4.4)$$

In general, it can mix under RG with other operators with the same flavour charge, or with different flavour charges as long as corresponding spurionic factors of SM Yukawas enter the running. This immediately restricts the number of dimension-6 SMEFT operators which can be connected to $\mathcal{O}_{eH}$ via renormalisation group equations (RGEs), at any loop order. Specifically, the complete list of operators with the same flavour charge as $\mathcal{O}_{eH}$ consists only of the leptonic electroweak dipoles:

$$\mathcal{O}_{eB} = (\bar{\ell} \sigma_{\mu\nu} e) H B^{\mu\nu} \sim \bar{\mathbf{3}}_L \times \mathbf{3}_e, \quad (4.5)$$

$$\mathcal{O}_{eW} = (\bar{\ell} \sigma_{\mu\nu} e) \tau^I H W_I^{\mu\nu} \sim \bar{\mathbf{3}}_L \times \mathbf{3}_e. \quad (4.6)$$

An additional four operators carry the same charge under $U(3)_L \times U(3)_e$ as $\mathcal{O}_{eH}$, whilst also carrying additional SM flavour charges:

$$\mathcal{O}_{lequ}^{(1)} = (\bar{\ell}^\alpha e) \epsilon_{\alpha\beta} (\bar{q}^\beta u) \sim \bar{\mathbf{3}}_L \times \mathbf{3}_e \times \bar{\mathbf{3}}_Q \times \mathbf{3}_u, \quad (4.7)$$

$$\mathcal{O}_{lequ}^{(3)} = (\bar{\ell}^\alpha \sigma_{\mu\nu} e) \epsilon_{\alpha\beta} (\bar{q}^\beta \sigma^{\mu\nu} u) \sim \bar{\mathbf{3}}_L \times \mathbf{3}_e \times \bar{\mathbf{3}}_Q \times \mathbf{3}_u, \quad (4.8)$$

$$\mathcal{O}_{ledq} = (\bar{\ell}^\alpha e) (\bar{d} q^\alpha) \sim \bar{\mathbf{3}}_L \times \mathbf{3}_e \times \bar{\mathbf{3}}_d \times \mathbf{3}_Q, \quad (4.9)$$

$$\mathcal{O}_{le} = (\bar{\ell} \gamma^\mu \ell) (\bar{e} \gamma^\mu e) \sim \bar{\mathbf{3}}_L \times \mathbf{3}_e \times \bar{\mathbf{3}}_e \times \mathbf{3}_L. \quad (4.10)$$

The SM Yukawas Y_u , Y_d and Y_e , can be thought of as having the spurionic flavour charges $Y_u \sim \mathbf{3}_Q \times \bar{\mathbf{3}}_u$, $Y_d \sim \mathbf{3}_Q \times \bar{\mathbf{3}}_d$ and $Y_e \sim \mathbf{3}_L \times \bar{\mathbf{3}}_e$ respectively, meaning that each of the four operators above has the same flavour charge as one of $Y_u^\dagger \mathcal{O}_{eH}$, $Y_d \mathcal{O}_{eH}$ or $Y_e \mathcal{O}_{eH}$. If the operators involve first generation lepton indices together with third generation quark or

lepton flavour indices, they can be connected via running to $\mathcal{O}_{eH}$ by amounts proportional to the large top Yukawa or (much smaller) bottom or tau Yukawas.

The isolation of $\mathcal{O}_{eH}$ is exacerbated further by helicity non-renormalisation theorems [67], which disallow its running at one loop into any operator other than the six-Higgs operator $\mathcal{O}_H$ (for which running is anyway suppressed by tiny lepton Yukawas, by the flavour arguments above). The leading RGEs describing the running from $\mathcal{O}_{eH}$ into the EW dipoles therefore arise at the two-loop level [66, 68, 65, 69]:¹

$$\frac{d}{d \ln \mu} \begin{pmatrix} \mathcal{C}_{eB} \\ \mathcal{C}_{eW} \end{pmatrix} = \frac{g^3}{(16\pi^2)^2} \frac{3}{4} \begin{pmatrix} \tan \theta_W \mathcal{Y}_H + 4 \tan^3 \theta_W \mathcal{Y}_H^2 (\mathcal{Y}_L + \mathcal{Y}_e) \\ \frac{1}{2} + \frac{2}{3} \tan^2 \theta_W \mathcal{Y}_H (\mathcal{Y}_L + \mathcal{Y}_e) \end{pmatrix} \mathcal{C}_{eH}, \quad (4.11)$$

where θ_W is the Weinberg angle and $\mathcal{Y}_X$ is the hypercharge of the field X . The covariant derivative we use is defined $D_\mu = \partial_\mu + ig \frac{\tau^I}{2} W_\mu^I + ig' \mathcal{Y} B_\mu$. Since some of the phenomenology of the leptonic dipoles is measured very precisely (e.g. Δa_e , $\mathcal{B}(\mu \rightarrow e\gamma)$), it is worth investigating whether this mixing, even at the two loop level, can provide relevant constraints on the coefficient of $\mathcal{O}_{eH}$. In particular, through this mixing a contribution to κ_e implies a shift in the anomalous magnetic moment of the electron, Δa_e .

Indirect constraints from Δa_e

To connect to measurements of Δa_e , we use Eq. (4.11) to evolve from the scale of new physics Λ to the EW scale μ_W using the leading-log approximation. At the EW scale, we match onto the coefficient $[L_{e\gamma}]_{ij}$ of the dipole operator $\bar{e}_L^i \sigma_{\mu\nu} e_R^j F^{\mu\nu}$ of the low energy effective field theory (LEFT). The tree level matching gives [70]

$$L_{e\gamma}(\mu_W) = \frac{v}{\sqrt{2}} (\cos \theta_W \mathcal{C}_{eB}(\mu_W) - \sin \theta_W \mathcal{C}_{eW}(\mu_W)). \quad (4.12)$$

1. The overall sign of this RGE is inconsistent between Ref. [69] and Refs. [66, 68, 65]. We use the sign from the latter references.

replaced by a W -boson loop. Given that the anomalous dimensions in (4.11) correspond to the divergent part of this diagram, the finite matching part is scheme-dependent and formally part of the next logarithmic order. We therefore do not include this in the calculation (consistent with the approach in Ref. [65]).

Summing Eqns. (4.13) and (4.14) we find²

$$L_{e\gamma}(m_h) = L_{e\gamma}^{\text{Matching}}(m_h) + L_{e\gamma}^{\text{Running}}(m_h) \quad (4.16)$$

$$= \frac{g^3 \sin \theta_W v}{12\sqrt{2}(16\pi^2)^2} \left(64 \sin^2 \theta_W f\left(\frac{m_t^2}{m_h^2}\right) - 9 \tan^2 \theta_W \ln\left(\frac{m_h}{\Lambda}\right) \right) \mathcal{C}_{eH}(\Lambda). \quad (4.17)$$

Finally, neglecting running of $L_{e\gamma}$ below the EW scale,³ we can relate this back to Δa_ℓ :

$$\begin{aligned} \Delta a_\ell &= \frac{4m_\ell}{e} \text{Re}([L_{e\gamma}]_{\ell\ell}) \\ &= -\frac{g^2 m_\ell^2}{3v^2(16\pi^2)^2} \left(64 \sin^2 \theta_W f\left(\frac{m_t^2}{m_h^2}\right) - 9 \tan^2 \theta_W \ln\left(\frac{m_h}{\Lambda}\right) \right) \Delta\kappa_\ell, \end{aligned} \quad (4.18)$$

where we define $\Delta\kappa_\ell = \kappa_\ell - 1$. Current constraints on Δa_e are dependent on the experimental value of the fine structure constant α . However, there is an existing tension between the two most precise experimental determinations of α , obtained using matter-wave interferometry of atomic recoils with either Caesium or Rubidium atoms [78, 79]. Comparing the resulting SM predictions obtained with the two inputs, to the most recent measurement [80], yields

2. For the sake of compactness, we have omitted the running of $\mathcal{C}_{eH}$ between Λ and the electroweak scale in the matching part of this analytical expression (i.e. we have set $\mathcal{C}_{eH}(m_h) = \mathcal{C}_{eH}(\Lambda)$ in the first term when going from (4.16) to (4.17)). This approximation represents at most a 10% correction to the constraints, but our numerical analysis includes the full running.

3. The self-renormalisation scaling of this coefficient, within the five-flavour LEFT, is small: $L_{e\gamma}(m_b) \approx 0.97 L_{e\gamma}(m_h)$ [77].

[81]⁴

$$\Delta a_e^{\text{Rb}} = (33.8 \pm 16.1) \times 10^{-14}, \quad (4.19)$$

$$\Delta a_e^{\text{Cs}} = (-102 \pm 26.4) \times 10^{-14}. \quad (4.20)$$

Clearly, (4.19) and (4.20) disagree by more than 5σ . In our study, we will take two approaches to dealing with this uncertain situation: (a) taking Δa_e^{Rb} as the current measurement (choosing Rb rather than Cs since it is less discrepant with the SM), and (b) inflating the error bars such that the 2σ region for Δa_e spans the 2σ region for both determinations. Specifically, for option (b) we define the following 2σ region:

$$(\Delta a_e)_{\text{Rb+Cs}} = [-154.8, 66.0] \times 10^{-14} \ (2\sigma). \quad (4.21)$$

Our qualitative conclusions will generally not depend on which of these approaches we take, but by doing both we can show the impact of the current inconsistency in α .

It is expected that by the time of a direct electron Yukawa measurement at FCC-ee there will be roughly an order of magnitude improvement in precision of Δa_e . With this in mind, we take $\Delta a_e^{\text{future}} < 5 \times 10^{-14}$ as a rough projection [81, 85]. Of course, any improvement in precision on Δa_e will rely on resolving the discrepancy in α . Using Eq. (4.18), we calculate the current and future constraints on κ_e from Δa_e assuming two different new physics scales ($\Lambda = 2, 10$ TeV) and list these in Tab. 4.1. These are to be compared to the current and future direct 68% limits on κ_e in Tab. 4.4. It can be seen that current constraints on κ_e from Δa_e are very weak, but if a future bound of $\Delta a_e^{\text{future}} < 5 \times 10^{-14}$ were to be achieved, it could outperform HL-LHC sensitivity.

4. The uncertainties on these values, whilst dominated by experimental errors from α determinations and the a_e measurement, receive a non-negligible component from the current 5σ discrepancy between the two calculations of the contributing 5-loop QED Feynman diagrams without lepton loops [82, 83, 84].

Λ (TeV)	κ_e (from Δa_e^{Rb})	κ_e (from $\Delta a_e^{\text{Rb+Cs}}$)	$ \kappa_e $ (projected future Δa_e)
2	[-890,-320]	[-890, 2300]	< 46
10	[-710, -250]	[-710, 1800]	< 37

Table 4.1: 68% confidence intervals for κ_e derived from current and future constraints on Δa_e under different assumptions for the new physics scale Λ .

4.2.2 Generating $\mathcal{O}_{eH}$ radiatively

Whilst UV physics could generate $\mathcal{O}_{eH}$ and therefore modify the electron Yukawa coupling at tree level, the potential precision of future measurements allows us to probe scenarios in which $\mathcal{O}_{eH}$ arises at loop level. As we found in Sec. 4.2.1, $\mathcal{O}_{eH}$ is an element in a set of operators which all mix under RG flow. If any of these operators is generated by integrating out UV physics, the rest of the operators (including $\mathcal{O}_{eH}$) will be populated after RG flow to the IR. In particular, we identify three semileptonic operators and one four-lepton operator as candidates to radiatively generate $\mathcal{O}_{eH}$: $\mathcal{O}_{lequ}^{(1)}$, $\mathcal{O}_{lequ}^{(3)}$, $\mathcal{O}_{ledq}$ and $\mathcal{O}_{le}$. The dipole operators are the only other possibility, but this possibility is strongly excluded by constraints from Δa_e .

The coefficients $\mathcal{C}_{lequ}^{(1)}$, $\mathcal{C}_{ledq}$ and $\mathcal{C}_{le}$ directly enter the one-loop RGE of $\mathcal{C}_{eH}$ [86, 87, 88]:

$$\begin{aligned} \frac{d}{d \ln \mu} \mathcal{C}_{eH} = & \frac{4N_c}{16\pi^2} (Y_u^\dagger Y_u Y_u^\dagger - \lambda Y_u^\dagger) \mathcal{C}_{lequ}^{(1)} - \frac{4N_c}{16\pi^2} (Y_d^\dagger Y_d Y_d^\dagger - \lambda Y_d) \mathcal{C}_{ledq} \\ & + \frac{8}{16\pi^2} (Y_e^\dagger Y_e Y_e^\dagger - \lambda Y_e) \mathcal{C}_{le}, \end{aligned} \quad (4.22)$$

where flavour indices are suppressed. The coefficient $\mathcal{C}_{lequ}^{(3)}$ first mixes into $\mathcal{C}_{lequ}^{(1)}$ through

$$\frac{d}{d \ln \mu} \mathcal{C}_{lequ}^{(1)} = \frac{1}{16\pi^2} (18g^2 + 30g'^2) \mathcal{C}_{lequ}^{(3)}, \quad (4.23)$$

and subsequently into $\mathcal{C}_{eH}$. Noting the Yukawa dependence of Eq. (4.22), we study the

following coefficient flavour indices:

$$[\mathcal{C}_{lequ}^{(1)}]_{1133}, [\mathcal{C}_{lequ}^{(3)}]_{1133}, [\mathcal{C}_{lequ}^{(1)}]_{1122}, [\mathcal{C}_{lequ}^{(3)}]_{1122}, [\mathcal{C}_{ledq}]_{1133}, [\mathcal{C}_{le}]_{1331}, \quad (4.24)$$

where the indices 1, 2, 3 correspond to first, second and third generation quarks and leptons respectively. The flavour charges of these operators, as discussed in Sec. 4.2.1, indicate we should also expect to generate contributions to the EW dipole operators. Indeed, $\mathcal{C}_{lequ}^{(3)}$ mixes into $\mathcal{C}_{eB}$ and $\mathcal{C}_{eW}$ at one-loop, whilst $\mathcal{C}_{lequ}^{(1)}$, $\mathcal{C}_{ledq}$ and $\mathcal{C}_{le}$ contribute at the two-loop level. So we can again study indirect constraints on $\Delta\kappa_e$ from Δa_e , assuming that both are generated radiatively from one of the operator coefficients in Eq. (4.24).

Assuming SMEFT coefficients are generated at a new physics scale $\Lambda = 2$ TeV, we compute the induced $\mathcal{C}_{eH}$ through numerically solving the one-loop RGE equations with the **DsixTools** package [89, 77]. To calculate contributions to Δa_e , we perform a resummed one-loop running and one-loop matching analysis, making use of the formulae provided in [90],⁵ supplemented by fixed-order contributions from two-loop anomalous dimensions where available in the literature. In particular, the calculation of the leading contributions from $[\mathcal{C}_{ledq}]_{1133}$ and $[\mathcal{C}_{le}]_{1331}$ to Δa_e requires the use of two-loop anomalous dimensions above

5. The contributions of $[\mathcal{C}_{lequ}^{(1)}]_{1122}$ and $[\mathcal{C}_{lequ}^{(3)}]_{1122}$ to Δa_e depend weakly on unknown non-perturbative charm loop effects, through the parameter denoted $c_T^{(c)}$ in Ref. [90]. We set this to zero here for simplicity, but if we were to take it to be $O(1)$, our constraints on these coefficients would change by $O(10\%)$.

and below the EW scale. Within the SMEFT, the relevant two-loop RGEs are [66]:

$$\begin{aligned} \frac{d}{d \ln \mu} \begin{pmatrix} \mathcal{C}_{eB} \\ \mathcal{C}_{eW} \end{pmatrix} &= \frac{Y_d g^3}{(16\pi^2)^2} \frac{N_c}{4} \begin{pmatrix} 3 \tan \theta_W \mathcal{Y}_Q + 4 \tan^3 \theta_W (\mathcal{Y}_L + \mathcal{Y}_e)(\mathcal{Y}_Q^2 + \mathcal{Y}_d^2) \\ \frac{1}{2} + 2 \tan^2 \theta_W \mathcal{Y}_Q (\mathcal{Y}_L + \mathcal{Y}_e) \end{pmatrix} \mathcal{C}_{ledq} \\ &\quad - \frac{Y_e g^3}{(16\pi^2)^2} \frac{1}{2} \begin{pmatrix} 3 \tan \theta_W \mathcal{Y}_L + 4 \tan^3 \theta_W (\mathcal{Y}_L + \mathcal{Y}_e)(\mathcal{Y}_L^2 + \mathcal{Y}_e^2) \\ \frac{1}{2} + 2 \tan^2 \theta_W \mathcal{Y}_L (\mathcal{Y}_L + \mathcal{Y}_e) \end{pmatrix} \mathcal{C}_{le}. \end{aligned} \quad (4.25)$$

Analogous to the two-loop running calculation in Sec. 4.2.1, these RGEs are used to calculate the EW dipole coefficients at the electroweak scale, which are then matched to $L_{e\gamma}$ at tree- and one-loop level. $\mathcal{C}_{ledq}$ and $\mathcal{C}_{le}$ also have one-to-one tree level matching relations with the four-fermion LEFT coefficients $L_{ed}^{S,RL}$ and $L_{ee}^{V,LR}$ respectively [70]. For evolution below the EW scale, the two-loop anomalous dimensions for the mixing of these four-fermion LEFT coefficients into $L_{e\gamma}$, have recently appeared in Ref. [91]:

$$\frac{d}{d \ln \mu} [L_{e\gamma}] = -\frac{4e^3}{3(16\pi^2)^2} [L_{ed}^{S,RL}] [M_d^\dagger] + \frac{50e^3}{9(16\pi^2)^2} [L_{ee}^{V,LR}] [M_e^\dagger] \quad (4.26)$$

where M_d and M_e are the diagonal mass matrices for down-type quarks and leptons respectively. We use these to calculate the additional fixed-order contribution to Δa_e which is induced from running to the b mass scale.

Putting this all together, assuming only one non-zero SMEFT coefficient at a time, Fig. 4.2 shows their contributions to $\Delta \kappa_e$, as constrained by current and projected bounds on Δa_e , as well as by complementary searches at existing and proposed colliders. We observe a strong degree of complementarity between collider and Δa_e constraints on the relevant SMEFT coefficients. Under current collider limits the coefficients $[\mathcal{C}_{lequ}^{(1)}]_{1133}$ and $[\mathcal{C}_{lequ}^{(3)}]_{1133}$ can accommodate κ_e enhancements at the level of HL-LHC sensitivity; however, present Δa_e




Figure 4.2: Current and projected upper limits on $|\Delta\kappa_e|$ (defined as $|\kappa_e - 1|$) at 95% CL, assuming coefficients are generated at 2 TeV and only one SMEFT coefficient is non-zero at a time. Current LHC bounds on $[C_{lequ}^{(1)}]_{1133}$ and $[C_{lequ}^{(3)}]_{1133}$ are taken from [2], while those on $[C_{lequ}^{(1)}]_{1122}$, $[C_{lequ}^{(3)}]_{1122}$ and $[C_{ledq}]_{1133}$ are from [3]. All HL-LHC and FCC-ee projections, shown by green bars, are taken from [4]. Constraints from the decays $D_s^+ \rightarrow e^+\nu$ and $B_s \rightarrow e^+e^-$ follow from App. D, where the hatching on the latter bar indicates the assumption that the coefficient is in the up mass basis. Future Δa_e sensitivity is assumed to reach $\Delta a_e^{\text{future}} < 5 \times 10^{-14}$.

measurements already exclude this possibility. If the HL-LHC were to observe a deviation $\Delta\kappa_e \sim 120$, current Δa_e bounds would therefore point to UV physics matching directly to $[\mathcal{C}_{eH}]_{11}$.

For some of the coefficients in the list (4.24), specifically ones which do not involve a top quark, there are additional constraints from meson decays to electrons, to which they contribute with a large chiral enhancement. This applies to the coefficient $[C_{ledq}]_{1133}$ which can mediate the $B_s \rightarrow e^+e^-$ decay (if the quark doublet flavour index is in the up mass basis), and the coefficient $[C_{lequ}^{(1)}]_{1122}$ which can mediate the $D_s^+ \rightarrow e^+\nu$ decay. Neither of these decay modes have been measured, and the experimental upper limits on them are orders of magnitude above the SM prediction [92, 93], but the chiral enhancement of the NP contribution from these coefficients is so large that meaningful constraints can nevertheless be set from these decays. The details of the calculation of these constraints are given in App. D. In Fig. 4.2, we show the corresponding upper limits on $\Delta\kappa_e$ in orange, assuming that a deviation in κ_e arises radiatively from the given coefficient. The crosshatching on the bar for $[C_{ledq}]_{1133}$ is to indicate that this bound depends on the flavour alignment of the operator indices; if the quark doublet flavour index is defined in the up type quark mass basis, then the constraint applies, but if it is defined in the down type quark mass basis, then there is no down-type flavour change and the constraint is entirely evaded. By contrast, the bound on $[C_{lequ}^{(1)}]_{1122}$ involves a charged current, and is almost independent of whether the indices are defined in the up or down mass basis (the difference is a factor of $V_{cs} \approx 0.97$).

Two coefficients remain viable candidates for radiatively inducing $\Delta\kappa_e$ above the FCC-ee Higgs-pole sensitivity without violating current bounds: $[\mathcal{C}_{lequ}^{(1)}]_{1133}$ and $[C_{ledq}]_{1133}$ (if aligned in the down-type quark mass basis). The parameter space for the latter could be fully tested by measurements of hadronic ratios above the Z -pole [4]. This leaves $[\mathcal{C}_{lequ}^{(1)}]_{1133}$, where the relevant parameter space for generating $\Delta\kappa_e \gtrsim 0.6$ can be largely probed by a future measurement of $\Delta a_e \lesssim 5 \times 10^{-14}$. We conclude that at the time of an FCC-ee Higgs-pole

run, a κ_e enhancement at this level, in the absence of anomalies in other processes, would again favour UV completions generating $[\mathcal{C}_{eH}]_{11}$ at leading order.

We note the simplification of studying only one SMEFT operator non-zero at any time could obscure scenarios in which large κ_e enhancements are allowed through the interplay of multiple SMEFT coefficients. Furthermore, in complete models one expects finite matching contributions at Λ to arise in similar places to those arising from RG evolution, thus an overly strict numerical interpretation of the results is to be discouraged. We study motivated multi-coefficient scenarios in the form of single field SM extensions in Sec. 4.3.

4.2.3 Constraints on the flavour structure of $\mathcal{O}_{eH}$

Although our goal is to specifically enhance the electron Yukawa by imposing a non-zero $[\mathcal{C}_{eH}]_{11}$, we are still required to make some assumption about the rest of the flavour entries of the coefficient, which are not restricted by any symmetry *a priori*. Obviously, the least constrained flavour structure would be one in which all of the other entries of $\mathcal{C}_{eH}$ are zero, which we denote the ‘electrophilic’ structure. However, such a scenario appears contrived from a UV perspective and so we ought to consider a range of possible UV scenarios.

In this section we review a few examples of the various flavour structures that can arise. We start with generic flavour constraints, without making any assumptions about $\mathcal{C}_{eH}$, and then examine more specific structures and in particular ones which avoid flavour-changing neutral currents (FCNCs) and correlate the diagonal components of $\mathcal{C}_{eH}$.

To review, both the lepton mass matrix and the CP-even Higgs Yukawas receive contributions from $\mathcal{O}_{eH}$. We may diagonalise the lepton mass matrix with unitary flavour rotations U_l and U_e ,

$$l \rightarrow U_l l \quad , \quad e \rightarrow U_e e \quad , \quad (4.27)$$

resulting in Higgs couplings in the mass basis

$$\mathcal{L}_h \supset -[y_e^{\text{eff}}]_{ij} h \bar{\ell}_L^i \ell_R^j, \quad (4.28)$$

$$y_e^{\text{eff}} = \frac{M_{ll}}{v} - \frac{v^2}{\sqrt{2}} U_l^\dagger \mathcal{C}_{eH} U_e, \quad (4.29)$$

where we have used ℓ to indicate the charged lepton after electroweak symmetry breaking, and M_{ll} is the diagonalised charged lepton mass matrix. Note that unless y_e and $\mathcal{C}_{eH}$ were already aligned in flavour space there is no reason to expect the combination $U_\ell^\dagger \mathcal{C}_{eH} U_e$ to be a diagonal matrix.

Experimental probes of off-diagonal elements of $\mathcal{O}_{eH}$

There are two experimental channels that constrain the lepton flavour violating (LFV) Higgs couplings, and they offer complementary bounds. For one, there are the LHC searches for LFV Higgs decays, where a Higgs from a traditional production channel decays to a pair of leptons of differing flavour. This search channel offers a direct probe of the lepton Yukawas at tree-level, and the subsequent runs at the LHC will steadily improve this bound. These bounds are most important in probing the LFV τ couplings.

	Current (PDG)	HL-LHC [94]	FCC [95]
$\text{BR}(h \rightarrow \mu e)$	4.4×10^{-5}	-	1.2×10^{-5}
$\text{BR}(h \rightarrow \tau e)$	2.0×10^{-3}	2.4×10^{-4}	1.6×10^{-4}
$\text{BR}(h \rightarrow \tau \mu)$	1.5×10^{-3}	2.4×10^{-4}	1.4×10^{-4}

Table 4.2: Current and future 95% CL expected limits on charged LFV Higgs decays.

The branching ratio of $h \rightarrow \ell_i \ell_j$ in the SMEFT, with $i \neq j$, is given by

$$\mathcal{B}(h \rightarrow \ell_i \ell_j) = \frac{1}{\Gamma_h} \frac{m_h}{16\pi} v^4 \left(|[\mathcal{C}_{eH}]_{ij}|^2 + |[\mathcal{C}_{eH}]_{ji}|^2 \right), \quad (4.30)$$

where we have neglected lepton masses in the kinematics. The current and future experi-

Experiment	Limit	Date	$\frac{v^2}{y_e^{\text{SM}}}[\mathcal{C}_{eH}(m_h)]_{\langle e\mu \rangle}$ Limit	
			$\Lambda = 2 \text{ TeV}$	$\Lambda = 10 \text{ TeV}$
Sundrum-II [46]	$\text{BR}(\mu \text{ Au} \rightarrow e \text{ Au}) < 7 \times 10^{-13}$	Current	6.6	6.0
MEG-II [47]	$\text{BR}(\mu \rightarrow e\gamma) < 1.5 \times 10^{-13}$	Current	0.39	0.33
MEG-II [47]	$\text{BR}(\mu \rightarrow e\gamma) < 6 \times 10^{-14}$	Late 2020s	0.25	0.21
Mu2e [48]	$\text{BR}(\mu \text{ Al} \rightarrow e \text{ Al}) < 6 \times 10^{-16}$	Late 2020s	0.22	0.20
Mu3e [96]	$\text{BR}(\mu \rightarrow 3e) \lesssim \mathcal{O}(10^{-16})$	~ 2030	~ 0.10	~ 0.08
COMET [97]	$\text{BR}(\mu \text{ Al} \rightarrow e \text{ Al}) \lesssim \mathcal{O}(10^{-17})$	Late 2030s	~ 0.03	~ 0.03

Table 4.3: Summary of limits on $\mathcal{C}_{eH}$ originating from searches for rare muon decays, as well as projections for future searches, all reported at 90% CL. For compactness, we denote the flavour average $C_{\langle e\mu \rangle} \equiv \sqrt{|C_{e\mu}|^2 + |C_{\mu e}|^2}$.

mental constraints are summarised in Table 4.2.

Alternatively, radiative LFV decays of leptons receive loop contributions involving the Higgs, thereby accessing the coefficients $\mathcal{C}_{eH}$. Because τ s are somewhat difficult to produce and observe, the τ LFV decay searches are subdominant to the relevant LHC constraints. In contrast, the naturally long lifetime of the muon makes it an extremely sensitive probe of EW-scale physics, and in turn provides a strong constraint on flavour violation in the μe sector. The sensitivity of these searches will improve considerably by the time the FCC becomes operational. We therefore review these constraints in brief, and then use them to evaluate a simple flavour scenario in the following section. The outlook of the future of muonic decay experiments, as well as their projected constraining power, is summarised in Table 4.3.

As discussed in Section 4.2.1, the photonic dipole coefficient $L_{e\gamma}$ is induced at two loops through running and matching contributions from $\mathcal{C}_{eH}$. The LFV decay probes are also somewhat sensitive to κ_t and κ_W , but motivated by the relevant LHC constraints we assume that they are effectively unity. The rate for the LFV decay $\ell_i \rightarrow \ell_j \gamma$ is [44]

$$\Gamma(\ell_i \rightarrow \ell_j \gamma) = \frac{m_i^3}{4\pi} \left(|[L_{e\gamma}(m_h)]_{ij}|^2 + |[L_{e\gamma}(m_h)]_{ji}|^2 \right). \quad (4.31)$$

Current searches for the rare decay $\mu \rightarrow e\gamma$ offer the strongest bound on the μe and $e\mu$ entries of $L_{e\gamma}$, set by the MEG-II experiment at $\text{BR}(\mu \rightarrow e\gamma) < 1.5 \times 10^{-13}$ [47]. In future years, other search channels will overtake this bound, so we briefly describe these decay modes.

Another important search channel is the $\ell_i \rightarrow 3\ell_j$ decay mode, which is mediated by a virtual photon or Higgs. The Higgs-mediated diagram is subdominant in comparison to the photon-mediated diagram due to the small electron Yukawa and it can be omitted at leading order. The photon-mediated diagram contains the same effective $L_{e\gamma}$ coupling, and the rate is approximately [44]

$$\Gamma(\ell_i \rightarrow 3\ell_j) \simeq \frac{\alpha m_i^3}{12\pi^2} \left| \log \frac{m_j^2}{m_i^2} - \frac{11}{4} \right| \times \left(\left| [L_{e\gamma}(m_h)]_{\mu e} \right|^2 + \left| [L_{e\gamma}(m_h)]_{e\mu} \right|^2 \right). \quad (4.32)$$

Note that this rate is suppressed over $\mu \rightarrow e\gamma$ by an additional factor of α . Nevertheless, the upcoming Mu3e experiment is expected to probe branching ratios for $\mu \rightarrow 3e$ down to $\lesssim 10^{-16}$ [96], improving the current bound on the μe couplings of $\mathcal{C}_{eH}$ by a factor of a few.

Lastly, $\mu \rightarrow e$ conversion is expected to compete with $\mu \rightarrow 3e$ and potentially result in the strongest bound on μe LFV, improving current constraints by up to two orders of magnitude. Both Higgs- and photon-mediated diagrams contribute to $\mu \rightarrow e$ conversion but, in contrast to the $\mu \rightarrow 3e$ case, we find the Higgs-mediated diagram is numerically larger. This is because the effective Higgs coupling to quarks in the nucleus is much larger than the electron Yukawa, which relatively suppresses $\mu \rightarrow 3e$. For further details on this calculation see App. E.

To summarise, we compile the current and projected strongest constraints on $\mathcal{C}_{eH}$ (from both Higgs decays in Tab. 4.2 and Tab. 4.4, and $\mu \rightarrow e$ transitions in Tab. 4.3) into the

following matrices (at 95% CL), with symmetric off-diagonal entries:

$$|\mathcal{C}_{eH}(m_h)| \times \text{TeV}^2 \lesssim \left\{ \begin{array}{l} \begin{pmatrix} 1.2 \times 10^{-2} & 1.1 \times 10^{-5} & 3.0 \times 10^{-2} \\ & 6.1 \times 10^{-3} & 2.6 \times 10^{-2} \\ & & 2.3 \times 10^{-2} \end{pmatrix} \quad (\text{current}) \\ \begin{pmatrix} 7.4 \times 10^{-5} & 9.2 \times 10^{-7} & 8.5 \times 10^{-3} \\ & 7.7 \times 10^{-4} & 7.9 \times 10^{-3} \\ & & 2.2 \times 10^{-3} \end{pmatrix} \quad (\text{projected}) \end{array} \right\} \quad (4.33)$$

where flavour indices run over e, μ, τ in that order from left to right and top to bottom. In the $e\mu$ entry, we have calculated the indirect limits assuming a new physics scale $\Lambda = 10 \text{ TeV}$. All other entries are obtained from direct bounds on $h \rightarrow ff'$ or $e^+e^- \rightarrow h$.

Examples of possible flavour structures

Anarchic Flavour The first and most straightforward UV possibility is one of full flavour anarchy, where all the entries of $\mathcal{C}_{eH}$ are of the same order. We assume that the anarchic structure of $\mathcal{C}_{eH}$ survives lepton mass diagonalisation, which is natural if the lepton mass matrix is mainly set by the renormalisable Higgs Yukawas. In this scenario, the strongest constraint on any entry of $\mathcal{C}_{eH}$ is set by the current bound on the branching ratio of $\mu \rightarrow e\gamma$ from Meg-II (Table 4.3), leading to a rough 95% CL indirect bound on κ_e of

$$|\Delta\kappa_e| = \frac{v^3}{\sqrt{2}m_e} |\mathcal{C}_{eH}]_{ee}| \sim \frac{v^3}{\sqrt{2}m_e} |\mathcal{C}_{eH}]_{e\mu}| \lesssim 0.3 \quad (\text{Anarchic flavour}). \quad (4.34)$$

An $\mathcal{O}(1)$ shift in κ_e is therefore allowed under an anarchic scenario without violating current flavour constraints or requiring excessive tuning. However, future improvements in muonic decay searches may improve this bound significantly prior to the start of the FCC. In par-

Coupling modifier	Current	HL-LHC	FCC-ee	FCC-hh
$ \kappa_e $	$\lesssim 120$ [99]	$\lesssim 60$ [100]	$\lesssim 1.3$ [6]	—
κ_μ	$1.07^{+0.25}_{-0.31}$ [8]	± 0.028 [7]	± 0.023 [7]	± 0.004 [7]
κ_τ	0.93 ± 0.07 [8]	± 0.016 [6]	± 0.0046 [6]	± 0.004 [6]

Table 4.4: Current constraints and future sensitivity (all at 68% CL) on Higgs-charged lepton coupling modifiers. The FCC-ee projection assumes a Higgs-pole run and there is no projection for FCC-hh available. Where necessary 95% CL projections have been halved to estimate the 68% CL sensitivity.

ticular, if COMET achieves its long-term sensitivity goals, the corresponding bound on $\Delta\kappa_e$ could be tightened to $\lesssim 0.02$, effectively excluding the possibility of observing it at the FCC-ee Higgs pole run under the anarchic flavour assumption.

Aligned flavour In the absence of neutrino masses, the SM possesses an exact vectorial $U(1)_e \times U(1)_\mu \times U(1)_\tau$ lepton flavour symmetry. It also possesses an approximate larger vectorial $U(3)_D$ symmetry which is broken only by the differences between the lepton Yukawa couplings. We can imagine that BSM physics also obeys these symmetries, giving flavour-aligned Wilson coefficients in both cases, and universal (and diagonal) Wilson coefficients in the latter. In the presence of neutrino masses, the misalignment between the neutrino mass basis and the lepton mass basis is physical, and it is reasonable to ask why effects of heavy BSM should be aligned with the lepton mass basis. Mechanisms such as a leptonic version of Spontaneous Flavour Violation (SFV) [98] may exist to justify this. Such an alignment in the lepton sector is radiatively stable (up to tiny neutrino mass effects), and ensures that charged lepton flavour violating operators are negligible, being suppressed by powers of m_ν/v .

Among the possible flavour alignments we wish to highlight a few interesting options.

- **Electrophilic** The simplest possibility is that the UV serendipitously features a form of flavour alignment and ‘diagonality’ such that in the basis in which the charged lepton Yukawa couplings have been diagonalised the $\mathcal{O}_{eH}$ operator generated in the UV lies entirely in the electron flavour direction. In the absence of neutrino masses, this can

be achieved by charging the BSM physics under $U(1)_e$ only.

In this case, within the EFT the only additional constraint on C_{eH} other than κ_e itself is the anomalous magnetic moment of the electron, as discussed in Sec. 4.2.1.

- **Universal** This is the case for which $U_l^\dagger \mathcal{C}_{eH} U_e \propto I_3$, i.e. the UV breaks the flavour symmetry as $SU(3)_L \times SU(3)_E \rightarrow SU(3)_D$, i.e. to the diagonal subgroup, predicting

$$\Delta\kappa_e = \Delta\kappa_\mu \frac{m_\mu}{m_e} = \Delta\kappa_\tau \frac{m_\tau}{m_e} . \quad (4.35)$$

Clearly the electron Yukawa modifications can, in this instance, be significant as compared to the heavier leptons. However note in Tab. 4.4 that at HL-LHC (FCC-hh) the projected future muon Yukawa constraints are at the 4% (0.4%) level. The corresponding indirect limit on the electron Yukawa modification, under this ansatz, would then be $\Delta\kappa_e \lesssim 8$ (0.8) at 68% CL.

- **Minimal flavour violation (MFV)** In this case $U_l^\dagger \mathcal{C}_{eH} U_e \propto \hat{M}$, such that

$$\Delta\kappa_e = \Delta\kappa_\mu = \Delta\kappa_\tau . \quad (4.36)$$

With κ_τ presently constrained as in Tab. 4.4, any possible deviations in the electron Yukawa would be correspondingly small. If the UV is MFV-like it is unlikely that electron Yukawa measurements would ever be competitive with the respective τ constraints.

To summarise, under the well-motivated Anarchic, MFV and Universal UV flavour assumptions we find that $\mu \rightarrow e\gamma$, HL-LHC and FCC-hh, would provide superior or competitive indirect bounds on the electron Yukawa coupling, respectively, as compared to an FCC-ee Higgs pole run. A more specific electrophilic scenario must be envisaged in order for a significantly modified electron Yukawa coupling to evade indirect constraints achievable

elsewhere.

4.3 Selected extensions which enhance the electron Yukawa

We now move from the EFT description to discuss simple extensions of the SM that can enhance the electron Yukawa. Within explicit models, more operators are generated than just $\mathcal{O}_{eH}$, meaning that additional tests and constraints can be envisaged, dependent on the individual model. However, we can again use the fact that the new physics must provide an additional source of electron chiral symmetry breaking to make some general statements.

In particular, if the contribution to the electron Yukawa from some model is dependent on a particular product of BSM couplings, then that same product of couplings can only appear in other observables that also break electron chiral symmetry. This can be understood through spurionic or flavour symmetry arguments similar to those in Sec. 4.2.1. Observables in which a vector current of leptons appears in the amplitude will instead be dependent on a different product of couplings. So, even within explicit models in which other effects are generated, the electron magnetic dipole moment is the only observable we can expect to be entirely correlated with the electron Yukawa.

For this reason, we begin this section with a discussion of the expected loop suppression of the dipole operators relative to the Yukawa operators in BSM models. We then explore simple single-particle extensions in more detail; a second Higgs doublet φ , and scalar leptoquarks.

4.3.1 *Model-dependent connections between Yukawas and dipoles*

As noted in Sec. 4.2.1, there is a 2-loop connection within the EFT between the Yukawa-like operators and the dipole operators, which allows weak but model-independent constraints to be put on deviations in κ_e from measurements of Δa_e . However, in almost every UV completion, this connection arises at lower loop order. Schematically, this can be understood




Figure 4.3: Schematic diagrams showing how the EW dipole operators O_{eW} and O_{eB} are generically generated at one loop order higher than the O_{eH} operator in UV models. The grey blob represents a diagram of arbitrary loop order involving exchange of heavy UV states. The left hand diagram matches to the dimension-six O_{eH} SMEFT operator, while the right hand diagram matches to the dimension-six dipole operators. The gauge boson line connects to any charged particle in the diagram, including the H loop and any charged particles within the grey blob.

by the diagrams in Fig. 4.3; for any UV completion it is generally possible to close a Higgs loop and radiate a gauge boson to obtain the electroweak dipole operators at one loop order higher than the $\mathcal{C}_{eH}$ operator.

For individual states which generate $\mathcal{C}_{eH}$ at tree or 1 loop level, this can be seen explicitly model-by-model, as summarised in Tab. 4.5. Here we have defined $\mathcal{C}_{e\gamma}$ as the coefficient of the SMEFT photonic dipole operator $(\bar{\ell}\sigma_{\mu\nu}e)HF^{\mu\nu}$, which is given in terms of $\mathcal{C}_{eB}$ and $\mathcal{C}_{eW}$ as $\mathcal{C}_{e\gamma} = \cos\theta_W\mathcal{C}_{eB} - \sin\theta_W\mathcal{C}_{eW}$. In all but one case listed in this table, $\mathcal{C}_{e\gamma}$ is generated at one loop order higher than $\mathcal{C}_{eH}$, or they are both generated at the same loop order. The exception is the $SU(2)_L$ doublet scalar φ , which has a 2 loop suppression of $\mathcal{C}_{e\gamma}$ relative to $\mathcal{C}_{eH}$ (in the case where it has no coupling to tops). This can be understood by the fact that, since all the Higgs legs in the tree level diagram generating $\mathcal{C}_{eH}$ come from the same vertex, closing a Higgs loop in this case gives a zero amplitude by the Ward identity:

$$\begin{array}{c}
 \begin{array}{c}
 \text{Diagram: A dashed line labeled } \varphi \text{ enters from the left, connects to a vertex, then forms a loop with a dashed line labeled } H \text{ exiting to the right. The loop contains a wavy line labeled } q \text{ and two curved arrows labeled } k \text{ and } k-q. \end{array}
 \end{array}
 \propto q \cdot \epsilon^*(q) = 0. \quad (4.37)$$

The only particles which have a renormalisable and gauge invariant interaction with three H bosons are the $SU(2)_L$ doublet scalar φ and the $SU(2)_L$ quadruplet scalars Θ_1 and Θ_3 .

State	Spin	SM charges	$\mathcal{C}_{eH}$	$\mathcal{C}_{e\gamma}$	Coupling dependence
$\mathcal{S}$	0	(1, 1, 0)	tree	1 loop	$\kappa_{\mathcal{S}}(\tilde{y}_{\mathcal{S}}^e)_{11}^*/f$
φ	0	(1, 2, $\frac{1}{2}$)	tree	2 loop	$\lambda_{\varphi}(y_{\varphi}^e)_{11}^*$
Ξ	0	(1, 3, 0)	tree	1 loop	$\kappa_{\Xi}(\tilde{y}_{\Xi}^e)_{11}^*/f$
Ξ_1	0	(1, 3, 1)	tree	1 loop	$\kappa_{\Xi_1}(\tilde{y}_{\Xi_1}^e)_{11}^*/f$
E	$\frac{1}{2}$	(1, 1, -1)	tree	1 loop	$(\tilde{\lambda}_E^e)_1(\lambda_E^e)_1^*/f$
Δ_1	$\frac{1}{2}$	(1, 2, $-\frac{1}{2}$)	tree	1 loop	$(\tilde{\lambda}_{\Delta_1}^l)_1^*(\lambda_{\Delta_1})_1/f$
Δ_3	$\frac{1}{2}$	(1, 2, $-\frac{3}{2}$)	tree	1 loop	$(\tilde{\lambda}_{\Delta_3}^l)_1^*(\lambda_{\Delta_3})_1/f$
Σ	$\frac{1}{2}$	(1, 3, 0)	tree	1 loop	$(\tilde{\lambda}_{\Sigma}^e)_1(\lambda_{\Sigma})_1^*/f$
Σ_1	$\frac{1}{2}$	(1, 3, -1)	tree	1 loop	$(\tilde{\lambda}_{\Sigma_1}^e)_1(\lambda_{\Sigma_1})_1^*/f$
$\mathcal{L}_1$	1	(1, 2, $\frac{1}{2}$)	tree*	tree*	$\{(\tilde{g}_{\mathcal{L}_1}^{Del})_{11}^*, (\tilde{g}_{\mathcal{L}_1}^{Dl})_{11}^*\} \gamma_{\mathcal{L}_1}/f$
φ	0	(1, 2, $\frac{1}{2}$)	1 loop	2 loop	$(y_{\varphi}^e)_{11}^*(y_{\varphi}^u)_{33}$
ω_1	0	(3, 1, $-\frac{1}{3}$)	1 loop	1 loop	$(y_{\omega_1}^{eu})_{13}(y_{\omega_1}^{ql})_{31}^*$
Π_7	0	(3, 2, $\frac{7}{6}$)	1 loop	1 loop	$(y_{\Pi_7}^{eq})_{13}^*(y_{\Pi_7}^{lu})_{13}$
$\mathcal{U}_2$	1	(3, 1, $\frac{2}{3}$)	1 loop	1 loop	$(g_{\mathcal{U}_2}^{ed})_{13}^*(g_{\mathcal{U}_2}^{lq})_{13}$
$\mathcal{Q}_5$	1	(3, 2, $-\frac{5}{6}$)	1 loop	1 loop	$(g_{\mathcal{Q}_5}^{eq})_{13}^*(g_{\mathcal{Q}_5}^{dl})_{31}$

Table 4.5: Single particle states which match at tree or one loop level onto $\mathcal{C}_{eH}$, with coefficients unsuppressed by y_e (or by any other SM Yukawa couplings smaller than y_b). In all cases, $\mathcal{C}_{e\gamma}$ is generated via the same product of couplings as $\mathcal{C}_{eH}$, with this coupling dependence given in the last column (in the notation of [19]). States labelled in blue match to $\mathcal{C}_{eH}$ (and $\mathcal{C}_{e\gamma}$) via diagrams involving a non-renormalisable interaction, meaning that the UV completion of these operators would require additional BSM states. The asterisks on ‘tree*’ for the $\mathcal{L}_1$ state indicate that this state matches at tree level onto these operators, but that the couplings involved cannot arise within a weakly coupled UV completion (for example it is known that $\mathcal{C}_{e\gamma}$ cannot be generated at tree level within a weakly coupled UV completion [20, 21]).

Therefore only models including these states can allow for $\mathcal{O}_{e\gamma}$ to be generated at 2 loop orders higher than $\mathcal{O}_{eH}$. Every other UV completion will match to $\mathcal{O}_{eH}$ with a diagram in which at least one H leg will originate from a different vertex to the $H^\dagger$ leg, thereby avoiding the kinematic situation which gives zero in $\mathcal{O}_{e\gamma}$ upon closing a Higgs loop.

Schematic indirect bounds on $\Delta\kappa_e$ from current and future measurements of Δa_e are shown in Tab. 4.6, for the three possible options for the relative loop suppression of $\mathcal{C}_{e\gamma}$ compared to $\mathcal{C}_{eH}$. These are calculated under the simple assumption that

$$C_{e\gamma} = \frac{e}{16\pi^2} \left(\frac{g^2}{16\pi^2} \right)^{n-1} C_{eH} \quad , \quad (4.38)$$

Relative loop suppression	$\Delta\kappa_e ((\Delta a_e^{\text{Rb}}))$	$\Delta\kappa_e ((\Delta a_e^{\text{Rb+Cs}}))$	$ \Delta\kappa_e (\Delta a_e^{\text{future}})$
0	[-0.004, -0.20]	[-0.20, 0.40]	< 0.01
1	[-0.6, -20]	[-20, 60]	< 2
2	[-90, -4000]	[-4000, 9000]	< 300

Table 4.6: Estimated constraints at 95% confidence on $\Delta\kappa_e$ from current and future Δa_e measurements, considering different scenarios for the relative loop suppression of $\mathcal{C}_{e\gamma}$ compared to $\mathcal{C}_{eH}$. We are estimating the future precision of Δa_e to be $\Delta a_e^{\text{future}} \sim 5 \times 10^{-14}$.

where $n = 0, 1, 2$ is the relative loop suppression, and g is some SM or BSM coupling which we take to be $O(1)$. In many realistic cases there would also be a logarithmic enhancement, which may render these estimates overly conservative.

We can see from this table that if a model generates $\Delta\kappa_e$ and Δa_e at the same loop order, then, given the current constraints on Δa_e , its contributions to electron Yukawa shifts are expected to be below $O(1)$ and hence unobservable at an FCC-ee Higgs pole run. This conclusion holds unless there is some fine-tuning or cancellation in the model which would alter this naive expectation. If instead a model generates Δa_e with a 1-loop suppression relative to $\Delta\kappa_e$, then current constraints allow an $O(10)$ — $O(100)$ deviation in the electron Yukawa, but future measurements of Δa_e could reduce this to $O(1)$. In this case, which applies to the majority of the new states listed in Tab. 4.5, the exact size of expected deviations in $\Delta\kappa_e$ will depend on specifics of the model, demanding a more complete analysis. Notice that many such states (specifically the ones inducing a tree level effect in $\mathcal{C}_{eH}$) are highlighted in blue in Tab. 4.5, signalling that the matching involves non-renormalisable interactions, or equivalently a UV completion involving an additional particle (or more), which can be from the same list above. Such two-particle models have been investigated in Refs. [101, 85], finding that indeed future projections of Δa_e can indirectly probe $\Delta\kappa_e$ generally to $O(1)$ levels in these models, but that in some instances projected bounds extend to > 10 or < 1 . Instead, for the case where contributions to $\mathcal{C}_{e\gamma}$ are suppressed by two additional loops as compared to $\mathcal{C}_{eH}$, deviations in $\Delta\kappa_e$ could be very large, even if future measurements of Δa_e are in agreement with SM predictions.

In the following subsections, we investigate the two extremes of these possible loop suppressions in more detail. First, we explore the parameter space of the φ model, including additional future tests of its couplings at FCC-ee and FCC-hh. Then, we look into models of scalar leptoquarks, exploring to what extent the link between Δa_e and $\Delta \kappa_e$ can be altered in models with additional flavour structure.

4.3.2 A heavy second Higgs doublet

As argued above, the only tree level UV completion of $\mathcal{O}_{eH}$ which gives a two loop relative suppression in $\mathcal{O}_{e\gamma}$ consists of an extra scalar $SU(2)_L$ doublet with couplings to leptons and to Higgs bosons. This model hence represents the ‘best case scenario’ for new physics in $\Delta \kappa_e$, in the sense that we might expect that current and near-future indirect constraints from Δa_e could still allow very large deviations in $\Delta \kappa_e$. The possibility of a very large electron Yukawa within a 2HDM was studied previously in Ref. [102]. Here, we take a simplified approach, assuming that $m_\varphi \gg m_h$, and only studying the minimal couplings needed to enhance the electron Yukawa coupling, to determine the impact of correlated tests of this scenario upcoming at HL-LHC or FCC-ee. The relevant Lagrangian is⁶

$$-\mathcal{L}_\varphi \supset (y_\varphi^e)_{ij} \varphi^\dagger \bar{e}_R^i l_L^j + (y_\varphi^u)_{ij} \varphi^\dagger i \sigma_2 \bar{q}_{Li}^T u_{Rj} + \lambda_\varphi \left(\varphi^\dagger H \right) \left(H^\dagger H \right) + \text{h.c.} \quad (4.39)$$

Matching at tree level to the SMEFT gives [19]

$$[\mathcal{C}_{eH}]_{ij} = \frac{\lambda_\varphi (y_\varphi^e)_{ji}^*}{M_\varphi^2}, \quad [\mathcal{C}_{le}]_{ijkl} = -\frac{(y_\varphi^e)_{li}^* (y_\varphi^e)_{kj}}{2M_\varphi^2}, \quad \mathcal{C}_\phi = \frac{|\lambda_\varphi|^2}{M^2}, \quad [\mathcal{C}_{lequ}^{(1)}]_{ijkl} = -\frac{(y_\varphi^u)_{kl} (y_\varphi^e)_{ji}^*}{2M_\varphi^2}. \quad (4.40)$$

6. Additional scalar potential terms would also be present, however we don’t consider them here as they are not related directly to chirally enhanced κ_e corrections.

Within this model, the Yukawa modifying coefficient, $\mathcal{C}_{eH}$, is thus composed of two couplings which can be separately probed through other processes via contributions from $\mathcal{C}_\phi$ and $\mathcal{C}_{le}$. Current direct constraints on $\mathcal{C}_{le}$ are derived from Bhabha scattering measurements performed at LEP II [4, 103, 104]. These constraints are expected to be significantly improved by off- Z -pole measurements of the leptonic ratio R_e at FCC-ee [4]. Limits on $\mathcal{C}_\phi$ can be placed from searches for double Higgs production at the LHC [8, 105] and HL-LHC [7]. At FCC-ee the trilinear Higgs coupling can be probed at the Zh run [106] while at FCC-hh measurements of double Higgs production will provide strong sensitivity [6].

Figure 4.4 shows these future projections in the plane of the two couplings λ_φ and $(y_\varphi^e)_{11}$ for two different values of the mass, along with the current allowed region and future sensitivity of Δa_e . We assume here that only these two couplings are non-zero. Contours of different values of $\Delta\kappa_e$ are shown with dashed diagonal lines. It can be seen that current indirect constraints from Δa_e allow extremely large values of $\Delta\kappa_e \sim 1000$, while the combination of measurements at FCC-ee and FCC-hh will indirectly probe values of couplings corresponding to $|\Delta\kappa_e| \lesssim 500$ at 2σ . These are all larger than the current direct constraint on κ_e . Future measurements of Δa_e could bring the indirect constraint on $\Delta\kappa_e$ down to about 100.

As seen in (4.40), φ also matches to $\mathcal{C}_{lequ}^{(1)}$ at tree level if it couples to up type quarks. This coefficient generates contributions to $\Delta\kappa_e$ at the one loop level and to Δa_e at two-loop level, with the largest effects from the operator involving the third generation quarks. From Fig. 4.2, we see that this operator alone could induce $O(10)$ effects in $\Delta\kappa_e$ while remaining consistent with current constraints. To investigate this in the context of the φ model, in addition to λ_φ and $(y_\varphi^e)_{11}$, we now allow a non-zero value of $(y_\varphi^u)_{33}$. Fig. 4.5 shows the allowed regions for this scenario in the plane of $y_\varphi^e \equiv (y_\varphi^e)_{11}$ and $y_\varphi^t \equiv (y_\varphi^u)_{33}$, while holding the coupling to the Higgs boson fixed at $\lambda_\varphi = 1$.

We see that including $O(1)$ Yukawa-like couplings to top quarks significantly tightens




Figure 4.4: Parameter space for a φ extension coupled to electrons. Blue shaded regions are consistent with the indicated Δa_e measurement at 95% C.L. Dashed lines show the κ_e enhancement corresponding to the given parameter values. Solid (dotted) lines represent constraints on the couplings from current (projected) measurements at 95% C.L. The R_e constraints are taken from [4], $\kappa_\lambda^{\text{LHC}}$ and $\kappa_\lambda^{\text{FCCee}}$ are from [5] and $\kappa_\lambda^{\text{FCCChh}}$ assumes 6% uncertainty at 95% CL [6, 7].




Figure 4.5: Parameter space for a second Higgs doublet φ with Yukawa-like couplings to electrons and top quarks, and assuming $\lambda_\varphi = 1$. Shaded regions and line styles have the same meaning as in Fig. 4.4. The κ_t^{LHC} limit is from [8], and the κ_t^{FCCee} projection is from [9].

constraints from Δa_e with respect to Fig. 4.4. However, at smaller values of y_φ^t this effect quickly diminishes, until at $|y_\varphi^t| \sim 10^{-2}$, the one-loop contribution to Δa_e (from C_{lequ}^{1133}) is subdominant to the two-loop contribution (from C_{eH}^{11}), and the situation becomes equivalent to the $y_\varphi^t = 0$ case in Fig. 4.4. For negative $y_\varphi^t \sim -0.1$ a cancellation can occur between the $[\mathcal{C}_{lequ}^{(1)}]_{1133}$ and $[\mathcal{C}_{eH}]_{11}$ contributions to Δa_e therefore lifting constraints. Such a cancellation would require significant tuning of the model parameters (including also the finite two-loop matching contributions at the UV scale which we have neglected in this analysis). Constraints on y_φ^t can also be derived from measurements of the top quark Yukawa coupling modifier κ_t . However, the FCC-ee sensitivity to this coupling modifier does not improve on perturbativity bounds except in the negative quadrant of the $M_\varphi = 2$ TeV plot.

In summary, in a model with an extra Higgs doublet, there is open parameter space in which $|\kappa_e| \gg 1$ could arise, which could be tested by an FCC-ee Higgs pole run and which would not otherwise be probed in the near future by Δa_e or by other FCC measurements.

4.3.3 Leptoquarks

Scalar and vector leptoquarks seemingly provide another promising option for new physics in κ_e , given they are the only extensions which match at tree level onto $\mathcal{C}_{lequ}^{(1)}$ and $\mathcal{C}_{ledq}$ respectively, both of which could naively produce large effects in κ_e while remaining consistent with current constraints (see Fig. 4.2). Inspection of Tab. 4.5, however, reveals both the scalar leptoquarks (ω_1 and Π_7) and the vector leptoquarks ($\mathcal{U}_2$ and $\mathcal{Q}_5$) actually correspond to the most pessimistic scenario for new physics, in which $\Delta\kappa_e$ and Δa_e arise at the same loop order. For the scalar leptoquarks this is primarily due to a tree-level matching to $\mathcal{C}_{lequ}^{(3)}$ [19] which runs at one loop into $\mathcal{C}_{e\gamma}$, while for the vector leptoquarks the leading effect is a finite one-loop contribution to $\mathcal{C}_{e\gamma}$ (see e.g. [107, 108, 109]). The connection between muonic dipoles and Higgs couplings in scalar leptoquark models was previously explored in Refs. [110, 111, 112]. We explore the electronic leptoquark option here, to examine whether

the broad loop-counting conclusions of Sec. 4.3.1 can be affected by particular parameter choices or flavour structures in the model.

We focus on the scalar leptoquark extensions ω_1 and Π_7 where the relevant Lagrangian terms are

$$\begin{aligned}
-\mathcal{L}_{\text{LQs}} \supset & (y_{\omega_1}^{ql})_{ij} \omega_1^\dagger \bar{q}_{Li}^c i \sigma_2 l_{Lj} + (y_{\omega_1}^{eu})_{ij} \omega_1^\dagger \bar{e}_{Ri}^c u_{Rj} + \lambda_{\omega_1} (\omega_1^\dagger \omega_1) (H^\dagger H) \\
& + (y_{\Pi_7}^{lu})_{ij} \Pi_7^\dagger i \sigma_2 \bar{l}_{Li}^T u_{Rj} + (y_{\Pi_7}^{eq})_{ij} \Pi_7^\dagger \bar{e}_{Ri} q_{Lj} \\
& + \lambda_{\Pi_7} (\Pi_7^\dagger \Pi_7) (H^\dagger H) + \tilde{\lambda}_{\Pi_7} (\Pi_7^\dagger H) (H^\dagger \Pi_7) + \text{h.c.},
\end{aligned} \tag{4.41}$$

where i and j are flavour indices, and colour and weak isospin indices are left implicit. These extensions match at tree level to the dimension-six SMEFT coefficients [19]

$$\left(\mathcal{C}_{lequ}^{(1)} \right)_{ijkl} = \frac{(y_{\omega_1}^{eu})_{jl} (y_{\omega_1}^{ql})_{ki}^*}{2M_{\omega_1}^2} + \frac{(y_{\Pi_7}^{eq})_{jk}^* (y_{\Pi_7}^{lu})_{il}}{2M_{\Pi_7}^2}, \tag{4.42}$$

$$\left(\mathcal{C}_{lequ}^{(3)} \right)_{ijkl} = -\frac{(y_{\omega_1}^{eu})_{jl} (y_{\omega_1}^{ql})_{ki}^*}{8M_{\omega_1}^2} + \frac{(y_{\Pi_7}^{eq})_{jk}^* (y_{\Pi_7}^{lu})_{il}}{8M_{\Pi_7}^2}. \tag{4.43}$$

Since identical couplings appear in both matching expressions, the one-loop mixing of $\mathcal{C}_{lequ}^{(3)}$ into $\mathcal{C}_{eB}$ and $\mathcal{C}_{eW}$, and of $\mathcal{C}_{lequ}^{(1)}$ into $\mathcal{C}_{eH}$, provides a strong connection between $\Delta\kappa_e$ and Δa_e in these models.

We compute the one-loop matching of ω_1 and Π_7 states onto all SMEFT operators contributing to $\Delta\kappa_e$ or Δa_e , with **SOLD** [113, 114] and provide these expressions in App. G.1. We neglect effects related to two-loop anomalous dimensions and two-loop matching at the EW scale. The dependence on the up-type Yukawa, present identically in the RGEs of Eq. (4.22) and the matching relations of App. G.1, indicates that leptoquark electron-top couplings ($(y_{\omega_1}^{eu})_{13} (y_{\omega_1}^{ql})_{31}^*$ or $(y_{\Pi_7}^{eq})_{13} (y_{\Pi_7}^{lu})_{13}^*$) will give dominant contributions to both $\Delta\kappa_e$ and Δa_e . Contributions from leptoquark electron-charm couplings ($(y_{\omega_1}^{eu})_{12} (y_{\omega_1}^{ql})_{21}^*$ or $(y_{\Pi_7}^{eq})_{12} (y_{\Pi_7}^{lu})_{12}^*$) are naively suppressed by at least $\frac{m_c}{m_t}$, but can in fact still be important due to RGE effects




Figure 4.6: 95% C.L. constraints on $\Delta\kappa_e$ implied by current measurements of Δa_e , assuming a UV completion of either ω_1 or Π_7 with only electron-top couplings.

below the EW scale. Within the matching relations, finite contributions proportional to quartic leptoquark couplings to the SM Higgs (λ_{ω_1} or λ_{Π_7}) are present but subdominant with respect to log-enhanced pieces. Nevertheless, these couplings can be phenomenologically relevant for Higgs-lepton coupling deviations [109], and we therefore include them in our analysis.

The ATLAS and CMS collaborations have searched for pair-produced scalar leptoquarks decaying to electron-top [115, 116, 117] and electron-charm [118] final states, excluding masses up to 1.8 TeV depending on coupling assumptions. We consider two benchmark scenarios, $M_{\omega_1/\Pi_7} = 2 \text{ TeV}$ and 10 TeV, noting that the lighter case may be excluded by improved direct searches before a Higgs-pole run at FCC-ee. We perform tree- and one-loop matching to the SMEFT at these scales. The analysis then follows closely that of Sec. 4.2.2; we compute the induced $\Delta\kappa_e$ and Δa_e using the numerical solution for the one-loop RGEs from `DsixTools` and the results of [90]. We provide expressions for these quantities in terms of the model parameters of Eq. 4.41, at the benchmark mass scales, in App. G.2.

Current Δa_e measurements already restrict $\Delta\kappa_e$ for ω_1 and Π_7 with electron-top couplings, as shown in Fig. 4.6. Under either $\Delta a_e^{\text{Rb+Cs}}$ or Δa_e^{Rb} assumptions, enhancements lie below projected FCC-ee sensitivity. These explicit bounds are slightly stronger than the model-independent estimates in Tab. 4.6. For leptoquarks with only electron-charm couplings (previously studied as a way to explain deviations in lepton magnetic moments while




Figure 4.7: Parameter space for ω_1 scalar leptoquarks with both electron-top and electron-charm couplings assuming $\lambda_{\omega_1} = 0$. Red and blue shaded regions are compatible with the stated $\Delta\kappa_e$ enhancement and measurement of Δa_e respectively. The green dashed line shows the 95% upper limit on the electron-charm couplings derived from current constraints on $\text{BR}(D_s^+ \rightarrow e^+ \nu)$. The darker blue region is not clearly visible in the left hand plot, but exists as a very narrow band at the top of the lighter blue band.

avoiding strong $\mu \rightarrow e\gamma$ limits [119, 120]), we find that achieving $\Delta\kappa_e \sim 0.6$ would require $(y_{\omega_1}^{eu})_{12}(y_{\omega_1}^{ql})_{21}^*$ or $(y_{\Pi_7}^{eq})_{12}(y_{\Pi_7}^{lu})_{12}^*$ to be above the perturbativity bounds on this combination of couplings (see Eqs. (F.1) and (F.2)). Hence, a single scalar leptoquark with only electron-charm couplings cannot generate large $\Delta\kappa_e$.

Allowing both electron-top and electron-charm couplings, for specific coupling choices, cancellations can relax the strong constraints on $\Delta\kappa_e$ derived under the electron-top only assumption.

We display the allowed parameter space for the ω_1 extension in Fig. 4.7. Here, the thinness of the blue bands demonstrates the current Δa_e^{Rb} and $\Delta a_e^{\text{future}}$ mandate a significant tuning of the ratio between the couplings to generate a large $\Delta\kappa_e$ ⁷. For $M_{\omega_1} = 2$ TeV, tuning can allow $\Delta\kappa_e \sim -70$, whilst not violating perturbativity bounds on electron-charm couplings, however constraints from $D_s \rightarrow e\nu$ (discussed in App. D) restrict this maximum enhancement to $\Delta\kappa_e \sim -6$. A heavier $M_{\omega_1} = 10$ TeV field will evade any meson decay

⁷ We note that the specific value of the required ratio depends weakly on the non-perturbative parameter $c_T^{(c)}$ [90]. We set this to zero here, but an $O(1)$ value for this parameter would change the required coupling ratio by $O(10\%)$.

or pair-production bounds but has a lower maximum enhancement magnitude $\Delta\kappa_e \sim -3$. Both cases represent parameter space which will be tested by an FCC-ee κ_e measurement at a dedicated Higgs pole run. Similar conclusions hold for the Π_7 state when both electron-top and electron-charm couplings are allowed.

Such a situation, with both top-electron and top-charm couplings, would imply that the leptoquarks would also mediate flavour changing neutral currents involving the top. Specifically, decays $t \rightarrow ce^+e^-$ are generated at tree level through some combinations of these couplings, while $t \rightarrow cZ$ would be induced at one loop by closing the electron loop and radiating a Z . Current constraints are only set on the latter process and do not effectively bound the model at all, see App. G.3. Projected $t \rightarrow cZ$ constraints at FCC-ee and FCC-hh are also very weak for this model. Given that it occurs at tree level, a promising search channel at future e^+e^- colliders is $e^+e^- \rightarrow t\bar{c}$. Corresponding projected constraints for a 240 GeV run at CEPC [121] are at the level of $[y^{ul}]_{21}[y^{qe}]_{31} < 0.6$ for a leptoquark of mass 2 TeV (see App. G.3), and could provide a test of the model, depending on the leptoquark masses that remain viable after HL-LHC. However we note that this is sensitive to a different combination of couplings to that appearing on the axes of Fig. 4.7, so this projection cannot be shown directly on this plane.

Finally, we examine whether cancellations between running contributions to Δa_e and finite effects from the quartic couplings λ_{ω_1} and λ_{Π_7} can allow for large $\Delta\kappa_e$. Neglecting electron-charm couplings now, Fig. 4.8 shows that for $M_{\omega_1/\Pi_7} = 2$ TeV neither state can achieve $\Delta\kappa_e > 0.6$ whilst respecting Δa_e and perturbativity constraints. For heavier masses, the allowed $\Delta\kappa_e$ enhancements are further diminished.

In summary, scalar leptoquarks are unlikely to contribute to a measurable deviation in κ_e at an FCC-ee Higgs pole run, since the corresponding couplings are already well tested by Δa_e . The only way that a significant deviation in κ_e can be achieved is by a leptoquark with a tuned ratio of charm-electron to top-electron couplings, and in which the charm-electron




Figure 4.8: Parameter space for ω_1 (a) and Π_7 (b) leptoquarks assuming $M_{\omega_1/\Pi_7} = 2$ TeV and only $(y_{\omega_1}^{eu})_{13}(y_{\omega_1}^{ql})_{31}^*$ and λ_{ω_1} or $(y_{\Pi_7}^{eq})_{13}(y_{\Pi_7}^{lu})_{13}^*$ and λ_{Π_7} are non-zero.

coupling is significantly (about 60 to 80 times) larger than the top-electron coupling, and of opposite sign. This is a rather unnatural flavour structure, which would appear to have no justification other than to achieve this effect.

4.3.4 Comments on fine-tuning

If the leading correction arises at dimension-6, then to realise a large ($\gg 1$) Yukawa modification requires fine-tuning the dimension-4 Yukawa coupling such that the electron mass is small despite having a large coupling to the Higgs. The greater the fine-tuning, the greater the coupling enhancement. Indeed, one may quantify this for the dimension-6 contribution in the usual way, through a fine-tuning estimate

$$\Delta = \left| \frac{\partial \ln \kappa_e}{\partial \ln y_e} \right| = \left| \frac{(1 - \kappa_e)(3 - \kappa_e)}{2\kappa_e} \right|, \quad (4.44)$$

which behaves as

$$\lim_{\kappa \gg 1} \Delta = \frac{\kappa_e}{2}, \quad (4.45)$$

and exhibits vanishing fine-tuning for $\kappa_e = 1$ or $\kappa_e = 3$, reflecting the cases where the dim-4 or dim-6 contributions dominate.

As a result, we see that in the limit where a new physics scenario gives rise to a Yukawa coupling which is significantly greater in magnitude than the SM coupling and not an integer (which corresponds to a single operator dominating), the parameters within the UV theory will have to be fine-tuned by an amount directly proportional to the magnitude of the enhancement. Pragmatically speaking, a theory which enhances the electron Yukawa by a factor of 10 will generically need to be fine-tuned at around the 20% level. Realistic UV-completions give rise to a tower of operators, thus this dimension-6 based analysis will be quantitatively modified. Nonetheless, the qualitative picture will remain valid as having a large Higgs coupling but small mass will generically require fine-tuning to realise the latter.

One should also consider radiative corrections from an IR perspective. Naïvely there is no symmetry nor dimensional analysis of Eq. (4.1) which can, from an IR perspective, forbid corrections at the matching scale or within the full UV theory of the form

$$\delta\mathcal{L}_{\text{Yuk}} = \frac{1}{(4\pi)^2} M^2 C_{eH} H \bar{l} e + \dots \quad , \quad (4.46)$$

where M is a UV mass scale. That this does not overcome the ultimate Yukawa correction originating from the dimension-6 contribution requires that

$$v^2 \gtrsim \frac{M^2}{(4\pi)^2} \quad , \quad (4.47)$$

naïvely requiring that any new states ought not to be heavier than ~ 3 TeV.⁸ This has little impact on the models considered here as it stands, however it does suggest that if a significant Yukawa modification were observed then natural new physics shouldn't lie too

8. This aspect is alluded to in [122] and has independently been emphasised to MM by Neal Weiner and Riccardo Rattazzi.

far above the TeV scale, despite that the electron Yukawa coupling is so small. However, as the previous discussion reveals, if $> \mathcal{O}(1)$ Yukawa coupling modifications are realised then some fine-tuning between higher and lower dimension contributions is necessary in any case, so naturalness arguments are weakened and it could be that the new physics responsible lies at higher mass scales.

We now turn to the question of RGE effects within the UV theory. One expects, on symmetry grounds, that within the UV theory there will be RGE contributions to y_e proportional to C_{eH} . Being model-dependent, it is not a theorem that such contributions must exist, however this is an important consideration since it impacts the question of radiative stability and fine-tuning of parameters within the UV theory. The reason is that we have already established that the value of y_e at the matching scale must be fine-tuned in order to realise significantly enhanced couplings. This fine-tuning is exacerbated if the precise value of y_e is not radiatively stable within the UV theory, which in the absence of additional symmetries is not generically expected. If the UV theory does appear radiatively stable, one should scrutinise it to ascertain if some fine-tuning, including RG effects, has been employed elsewhere in order to render small the couplings which contribute to the running of y_e .

4.4 Discussion

The electron Yukawa is presently very poorly constrained. Higgs boson decay measurements are projected, at HL-LHC, to be sensitive to up to 120 times the SM value. This could be improved on significantly during a proposed FCC-ee run on the Higgs resonance, in which it can be accessed through s -channel production. Although it is not part of the baseline running scenarios as outlined in the recently published FCC feasibility study [6], if it goes ahead this run would test electron Yukawa couplings down to 1.6 times the SM value. As a fundamental Standard Model parameter, confirming the value of this coupling would be an important milestone, but it is worth evaluating what any deviation in this coupling would

tell us about BSM physics. In this work we have explored the broad theoretical implications of an $\gtrsim \mathcal{O}(1)$ enhancement of the electron Yukawa, correlating it with other measurements in the landscape of current and future planned experiments.

CHAPTER 5

NEUTRINO EXPERIMENTS AND LONG-BASELINE SIGNATURES OF HIDDEN SECTORS

Next, we turn to studies of new physics at low energy scales. Primarily, we are interested in inverting the precision experimental paradigm—how can we utilize precision experiments to conduct discovery studies? In fact, there are many cases where experiments build to study very specific features of the SM can in fact be used to search for new particles, at little to no additional cost to the experimental program. In this chapter, and the next, we examine two cases where this approach is fruitful. This work is based on the published article [123].

MiniBooNE and MicroBooNE are nominally neutrino precision experiments: they are intended to study the oscillation patterns of neutrinos as they propagate through space. This alone is a feature that cannot be explained by the SM alone - a consistent explanation requires new particle species.¹ However, the geometry of these setups also mimic the “beam dump” style experiments. In beam dump experiments, a detector is placed behind a large section of opaque material, like cement or rock. The only particles capable of traversing such material are either neutrinos or new particles beyond the SM. However, contrary to most beam dump experiments, MiniBooNE and MicroBooNE are also capable of reconstructing kinematics, in particular the invariant mass of a scalar decaying to two photons. In fact, these collaborations already collect this exact dataset, searching for neutral pions that are produced by neutrino scattering and subsequently decay. Consequently, one can bound the coupling of a scalar produced in the MiniBooNE focusing horn, which decays to two photons. The model which is viable is called *muonphilic scalar*.

1. For a review of neutrino physics, see e.g. [124, 125].




Figure 5.1: Schematic cartoon of our MiniBooNE and MicroBooNE signal processes. **Top:** the 8 GeV proton beam impinges on a fixed Be target and produces mesons including charged kaons, whose decays yield a flux of muons and associated scalars S through $\pi^\pm, K^\pm \rightarrow \mu^\pm \nu_\mu S$ decays. **Bottom:** S may also be created via a spray of pions/kaons whose decays yield a secondary muon beam in the decay pipe. The muons then interact with the beam dump, where they produce long-lived muon-philic scalar particles S through bremsstrahlung-like processes. Note that the meson decay mode is independent of the beam dump, while the muon bremsstrahlung mode requires scattering in the beam dump. In both cases, for $m_S < 2m_\mu$, the scalars eventually decay to diphotons in the downstream detector. For details of the experimental setup and detector performance, see Ref. [10].

5.1 Model Overview

We consider a muon-philic scalar S , a SM singlet of mass m_S with the Yukawa interaction

$$\mathcal{L}_{\text{int}} \supset y S \bar{\mu} \mu, \quad (5.1)$$

which we have written in four-component Dirac fermion notation. This interaction can arise from a gauge-invariant dimension-5 operator

$$\mathcal{L}_{\text{eff}} \supset \left(\frac{y}{v}\right) S H^\dagger L_2 \mu^c + \text{h.c.} \quad (5.2)$$

where H is the SM Higgs doublet, $v = 246$ GeV is its vacuum expectation value, L_2 is the second generation lepton $SU(2)_L$ doublet, and μ^c is the $SU(2)_L$ singlet muon field.




Figure 5.2: **Left:** Feynman diagram for scalar production from $K, \pi \rightarrow \mu\nu S$ decays. **Middle:** Scalar production via μ - N scattering in the beam dump. **Right:** Scalar decay through a muon loop using only the coupling in Eq. (5.1).

This non-renormalizable interaction may be UV-completed with a heavy vector-like fourth generation of leptons, which mixes with the muon and is integrated out in the low-energy effective theory [126, 127, 128].

If the operator in Eq. (5.1) is the scalar's only coupling to SM fermions, then for $m_S < 2m_\mu$, S will decay to diphotons through muon loops. The width for this process is [15]

$$\Gamma_S = \frac{\alpha^2 y^2 m_S^3}{64\pi^3 m_\mu^2} x^2 |1 + (1-x)f(x)|^2, \quad (5.3)$$

where $x \equiv 4m_\mu^2/m_S^2$ and $f(x) = \arcsin^2(x^{-1/2})$ for $x > 1$. Note that for nearly all viable parameter values in this low mass range, S is a long-lived particle with a decay length

$$\frac{c\tau_S}{\gamma} \approx 300 \text{ m} \left(\frac{10^{-4}}{y} \right)^2 \left(\frac{40 \text{ MeV}}{m_S} \right)^3, \quad (5.4)$$

where $\tau_S = 1/\Gamma_S$ is the S lifetime and γ is the boost factor. This estimate uses Eq. (5.3) in the $x \gg 1$ limit.

In our parameter space of interest (see Fig. 5.6), the Yukawa coupling in Eq. (5.1) always satisfies $y \gtrsim 10^{-5}$, which is sufficient to bring S into chemical equilibrium with SM particles in the early universe. Therefore, in order to avoid cosmological limits on extra light degrees of freedom during Big Bang nucleosynthesis, we generically require $m_S \gtrsim \text{few MeV}$ [14, 129].




Figure 5.3: Heatmaps (arbitrary units) of the processes in the beam dump that produce muon-philic scalars, for $r_{\perp} < 0.9\text{ m}$, where $r_{\perp}$ is the cylindrical radius in the $x - y$ plane. This displays the scalar production inside the decay pipe, however the analysis considers production events that occur outside the pipe as well. The top (middle) diagram shows the vertices where the $\pi^{\pm}$ ($K^{\pm}$) decays, which may contain scalars in their decay products. The lower diagram shows the spatial distribution of scalar emission from muons scattering in the beam dump, specifically for $m_S = 100\text{ MeV}$, although this distribution does not vary significantly over the scalar mass. The plot has several features: the target can be seen on the far left, with the electromagnet adjacent, and the beam dump is visible on the right.

5.2 Analysis

Our search concept leverages secondary muon production at MiniBooNE and MicroBooNE following a sequence of steps, depicted schematically in Fig. 5.1 with relevant Feynman diagrams in Fig. 5.2:

1. The 8 GeV proton beam from the Fermilab Booster Neutrino Beamline (BNB) strikes the Be target, producing a forward shower of pions and kaons;
2. The charged pions and kaons are focused down the beamline via an electromagnet, called the *horn*, and then typically decay in flight in the decay pipe;
3. New scalar singlets S can be produced either directly through radiative 3-body $\pi^{\pm}, K^{\pm} \rightarrow \mu^{\pm} \nu_{\mu} S$ decays (Fig. 5.1, top), or through $\mu N \rightarrow \mu N S$ scattering from the secondary interactions in the beam dump (Fig. 5.1, bottom);




Figure 5.4: Energy distribution of charged pions (dark green), charged kaons (dark blue), and secondary muons (dark red) produced from proton-Be collisions at 8 GeV. Also shown are energy distributions for scalars produced from pion decays (light green), kaon decays (light blue) and from secondary muon scattering in the beam dump (orange). The scalar distributions are normalized to $y = 1$ to enhance visibility on the plot, but in our analysis we explore the region $y \in [10^{-5}, 10^{-2}]$. These curves were generated with a **GEANT** simulation utilizing the setup described in Refs. [10, 11].

4. The S particles propagate on-shell and decay visibly to $\gamma\gamma$ in the downstream detector.

To perform our numerical analysis, we utilize samples of pions, kaons, and muons from 10^7 protons-on-target (POT) created from a custom **GEANT4** [130] simulation of the BNB in neutrino mode. The detailed simulation includes all relevant beamline elements – the target, horn electromagnet and magnetic field, collimator, decay pipe, shielding, and beam dump – and is based on the description in Refs. [10, 11]. The 10^7 POT produce $\sim 10^6$ charged kaons and $\sim 2 \times 10^7$ muons (primarily via decays of pions produced at the target).

For the decay production channel, we generate a Monte Carlo (MC) sample of the scalars produced in K/π decays by sampling the available phase space of the decays, and weighting each event by the squared matrix element. The amplitude and branching ratio formulas are given in Ref. [131], and our sampling method agrees with their reported $K \rightarrow \mu\nu S$ branching




Figure 5.5: Example signal distributions generated by scalar decays in the MiniBooNE and MicroBooNE detectors. The data points are measured $\gamma\gamma$ -like events, and the background prediction arises primarily from neutrino-induced $\pi^0 \rightarrow \gamma\gamma$ events, and digitized from [12, 13]. The MiniBooNE error bars include statistical and systematic uncertainties. The MicroBooNE error bars are statistical only. These particular signal distributions correspond to 20σ (top panel) and 6σ (bottom panel) discrepancies with the data.

ratios. The pion MC is generated using the same amplitude, with the general substitution $K \rightarrow \pi$. The kinetic energy distribution of scalars produced via these modes, for $m_S = 24$ MeV, is shown in Fig. 5.4. The locations of the decays are shown in Fig. 5.3 (top and middle panels).

For the scalars produced by muons scattering in the beam dump, we use the **FeynRules** [132] Mathematica package to construct the Feynman rules for muon-philic interactions in Eq. (5.1). This model is then used as input to the **CalcHEP** [133] event generator to generate scalar emission $\mu N \rightarrow \mu NS$ cross sections (Fig. 5.2, center) and MC events. These events are then re-weighted according to the electromagnetic form factor of the iron nucleus (see Ref. [134]

for a more detailed treatment). The probability of emitting a scalar in a muon-nucleus scattering event, as a function of muon energy incident on the beam dump, is

$$P_{\text{emit}}(E_\mu) \simeq \frac{n_N}{\langle dE/dx \rangle} \int_0^{E_\mu} \sigma(E) dE, \quad (5.5)$$

where $n_N = 8.6 \times 10^{22}/\text{cm}^3$ is the number density of nuclei in the iron beam dump, $\sigma(E)$ is the energy-dependent $\mu N \rightarrow \mu N S$ cross section, and $\langle dE/dx \rangle$ is the average energy loss of the muon in the medium, conservatively taken to be 16 MeV/cm [135]. The typical path length in the beam dump is thus $l \sim 50$ cm [136]. For each muon produced by the GEANT4 simulation, the differential scalar emission probability in Eq. (5.5) is used to generate a random scattering event. A CalcHEP MC event of matching incident muon energy is randomly selected to generate the scalar 4-momentum. The production locations for scalars with $m_S = 100$ MeV are shown in Fig. 5.3 (bottom panel).

Example energy distributions of the mesons and scalars for $m_S = 24$ MeV are shown in Fig. 5.4. Using the 4-momenta of the scalars generated via the meson decay and scattering channels, we then evaluate the probability of S decay within the fiducial volume of MiniBooNE/MicroBooNE, via the loop diagram in Fig. 5.2, right. This is computed using the distance from the emission event to the detector, its line-of-sight width along the direction of the scalar momentum, and the decay length. The approximate decay length is given in Eq. 5.4, but in our analysis we use the full expression (5.3) for each value of m_S . Diphoton invariant mass distributions for two representative scalar masses are shown alongside data and background expectations in Fig. 5.5. We find that the largest signal rate given the MiniBooNE and MicroBooNE target and detector geometry comes from boosted $K^\pm \rightarrow \mu^\pm \nu_\mu S$ decays. Pion decays $\pi^\pm \rightarrow \mu^\pm \nu_\mu S$ are only kinematically allowed for $m_S < m_\pi - m_\mu = 34$ MeV, but even at $m_S = 24$ MeV (Fig. 5.5, bottom) the kaon decays still dominate. The contribution to the total $\gamma\gamma$ rate from muon scattering is largest at $m_S \sim 30$ MeV, but is generally subdominant over the whole relevant mass range; unlike the meson decay channels, the scattering

rate depends on the beam dump thickness and density, and thus this contribution may be more important in other experimental setups.

To evaluate our sensitivity to the S decay diphoton signal, we use the fact that the same final state was used in MiniBooNE and MicroBooNE’s measurements of neutrino-induced neutral current (NC) π^0 production, $\nu_\mu N \rightarrow \nu_\mu N \pi^0 (\pi^0 \rightarrow \gamma\gamma)$. In both analyses, the diphoton invariant mass $m_{\gamma\gamma}$ is reported, and we recast these analyses as searches for resonances with $m_{\gamma\gamma} = m_S$. Notably, the NC π^0 signal definition was slightly different for the two experiments. In MiniBooNE, the measurement included events with ≥ 0 visible protons since MiniBooNE was insensitive to proton(s) at the interaction vertex [12]. MicroBooNE was able to separate their measurement into NC π^0 events with either 0 or 1 associated visible protons [13]. We consider the 0 proton case in MicroBooNE here since no protons are expected in an S decay.

The MiniBooNE (MicroBooNE) measurement was performed in neutrino mode with 6.46×10^{20} (5.89×10^{20}) POT and utilized a fiducial mass of 450 tons (77 tons). MiniBooNE’s measurement in antineutrino mode is not considered here for simplicity. For MiniBooNE, we impose a “track mass” cut, which is used to separate diphoton and single-electron events, providing a natural lower bound on sensitivity,

$$m_{\gamma\gamma} \equiv \sqrt{2E_1 E_2 (1 - \cos \theta_{12})} > f(E_1 + E_2), \quad (5.6)$$

where $E_{1,2}$ are the energies of the two photons, θ_{12} is their separation angle in the lab frame, and the specific function $f(E_1 + E_2) \approx 32 \text{ MeV}$ is given in Ref. [137]. For MicroBooNE, we employ the minimum photon energies used in their analysis ($E_{\text{high}} > 30 \text{ MeV}$ and $E_{\text{low}} > 20 \text{ MeV}$), but allow arbitrary separation angles. The smeared signal distribution of $m_{\gamma\gamma}$ is approximated as a Gaussian centered at the scalar mass with 15% (33%) resolution, consistent with the measured MiniBooNE (MicroBooNE) π^0 invariant mass in this region [138, 12]. We assume that the 2γ detection efficiency is a flat 40% (6%) across the

invariant mass range used in this analysis, 32–128 MeV (10–120 MeV) in MiniBooNE (MicroBooNE), which is reasonably consistent with Refs. [12, 13]. There is minimal sensitivity above these mass ranges primarily because the scalar decays before reaching the detectors.

The lower threshold at MicroBooNE is noteworthy. Since MicroBooNE does not expect any true neutrino-induced π^0 (2γ , 0p) events in the 0–10 MeV bin [13] and it is difficult to estimate the detection efficiency behavior within this region, we have conservatively set the lower threshold to be 10 MeV. However, in principle, LArTPCs are sensitive to invariant masses below this value and we therefore encourage these experiments to explore sensitivity to $m_{\gamma\gamma} < 10$ MeV, where pion decay could dominate and for which there is a large swath of allowed parameter space that could be used to explain the $g - 2$ anomaly.

The primary background for our analysis is π^0 production in neutrino-induced $\text{NC}\pi^0$ scattering, and in setting our limits we compare our predicted signal yield against the observed diphoton distributions from Refs. [12, 13], noting that no significant excess is apparent in the datasets. For MiniBooNE, we calculate the significance of an injected signal (s_i) with a standard χ^2 test,

$$\chi^2 = \sum_{i,j=\text{bins}} (d_i - s_i - b_i) M_{ij}^{-1} (d_j - s_j - b_j) , \quad (5.7)$$

where d_i is the observed data, b_i is the background prediction, and M_{ij} is the uncertainty matrix. We construct M_{ij} using the statistical uncertainties on the data and assume fully correlated bin-to-bin systematics associated with flux (12.4%), cross section (8.4%), and the detector (5.1%), consistent with Ref. [12]. We note that these systematics are nominally applicable to MiniBooNE’s $\text{NC}\pi^0$ analysis, rather than a generic 2γ search, and that the detailed correlations between bins may be important. For MicroBooNE, we use a Poisson extended maximum likelihood χ^2 definition [135, 139] since the number of reported events

is < 20 across many of the relevant bins:

$$\chi^2_{\text{poiss}} = 2 \sum_{i=\text{bins}} \left[N_{\text{pred},i} - d_i + d_i \ln(d_i/N_{\text{pred},i}) \right], \quad (5.8)$$

where $N_{\text{pred},i} = s_i + b_i$. This MicroBooNE estimate is a “statistical uncertainty only” approximation, which ignores systematic uncertainties. However, the relevant reported invariant mass bins (10–120 MeV) are dominated by statistical errors and a more complete treatment is therefore not expected to markedly change our reported limit. A more detailed study by the collaboration, in particular to search for generic 2γ (rather than $\text{NC}\pi^0$) events, validate the χ^2 thresholds with fake data studies (given the low statistics), and include systematics, would be valuable. We define a 5σ exclusion limit in the region where the p -value associated with the χ^2 falls below 2.8×10^{-7} .

5.3 Results

In Fig. 5.6 we show our sensitivity in the coupling y versus mass m_S parameter space. Our main results are the red shaded region which represents the MiniBooNE exclusion at 5σ confidence, and the purple shaded region with the analogous interpretation for MicroBooNE. We also include a projection for the SBND LArTPC-based experiment at Fermilab [140], set to take first data in 2024². We consider a realistic detector geometry (77 ton fiducial mass) at a distance of 60 m from the beam dump, 10^{21} POT, and use the same invariant mass resolution and efficiency assumptions as MicroBooNE. The excluded regions all have similar characteristic shapes. At large coupling and high mass, the decay length is much shorter than the detector baseline, and the signal flux is exponentially suppressed. In the low-coupling region, scalar production becomes increasingly unlikely. In the low-mass limit, the analyses are no longer suited to detect the diphoton signal. It is noticeable that SBND has a much

2. This projection is not to be considered as coming from the SBND collaboration.

shorter baseline - this allows it to exclude further into the shorter lifetime region. It also has a higher geometric acceptance, which allows it to extend further into the small coupling regime. The SBND $\text{NC}\pi^0$ background estimate comes from scaling MicroBooNE to the SBND location using the ray tracing simulation, which estimates an increase in background event rates of a factor of $N_{\text{bkg,SBND}} \approx 90N_{\text{bkg,MicroB}}$. This reasonably agrees with the simple estimate $N_{\text{bkg}} \propto N_{\text{POT}}V/r^2$, with V the detector volume and r the baseline.

The green band in Fig. 5.6 represents the favored region for reconciling the anomalous observed experimental value of $g - 2$ to within 2σ agreement [141]. The combination of MiniBooNE and MicroBooNE constraints exclude scalars of mass between 10 and 60 MeV in this band. The gray shaded region directly above the green band is excluded by muon $g - 2$ measurements, as loop contributions from S would contribute beyond the observed value and reintroduce a discrepancy between theory and experiment at the $> 5\sigma$ level. For $m_S > 2m_\mu$, Fig. 5.6 also shows constraints from $S \rightarrow \mu^+\mu^-$ decays assuming a 100% branching fraction into this channel. We show limits from the BaBar $e^+e^- \rightarrow 4\mu$ search dashed projections for improving B -factory limits with future Belle-II analyses [14]. For $m_S > \text{few GeV}$, there is also a CMS constraint based on reinterpreting a $Z \rightarrow 4\mu$ decay search [14].

Below the dimuon mass threshold, Fig. 5.6 also shows several projections for future searches. The blue dashed region is the projection for a dedicated muon beam-dump experiment proposed in Ref. [15], which could cover much of the remaining $g - 2$ band. The orange dashed curve is based on a proposed muon missing momentum (M^3) search strategy [16]. Although M^3 is nominally sensitive to invisibly-decaying particles produced in muon interactions, the scalar S can be sufficiently long-lived on the length scales of the proposed experiment such that it is effectively invisible, and M^3 would probe the same parameter space as the model considered here. As mentioned earlier, access to the 0 – 10 MeV bin in $m_{\gamma\gamma}$ at MicroBooNE/SBND would allow us to place strong limits on $m_S < 10$ MeV from S production in charged pion decays, which would likely be competitive with the dedicated




Figure 5.6: Parameter space for the singlet scalar in Eq. (5.1), decaying to diphotons via Eq. (5.3). The red and purple shaded regions are excluded at 5σ by existing MiniBooNE and MicroBooNE data, respectively, based on the analysis in this paper and the magenta dashed curve represents a future 2σ projection for SBND sensitivity. The green band represents the parameter space where this scenario resolves the muon $g - 2$ anomaly. For $S > 2m_\mu$ we show constraints from BaBar and CMS based on the $S \rightarrow \mu^+\mu^-$ decay channel along with future projections for Belle-II [14]. For $S < 2m_\mu$ we also show projections for a dedicated muon beam-dump experiment [15] and for the M^3 muon missing momentum proposal [16]. Also shown are cosmological bounds for $m_S < 1$ MeV where S particles thermalize with the SM as a relativistic species during Big Bang nucleosynthesis [17].

muon beam dump projections.

5.4 Discussion and Outlook

The results of this work should be taken as a proof-of-concept validation of the potential of reinterpreting data collected at neutrino short-baseline experiments in the context of muon interactions beyond the SM. There are some experimental subtleties that were not taken into account in this analysis, and thus we expect the results to be correct to within an $\mathcal{O}(1)$ factor. Nonetheless, given the amount of parameter space which is accessible given current data, especially the highly-motivated $g - 2$ region, we strongly urge the MiniBooNE and

MicroBooNE collaborations to undertake a generic 2γ search (including systematics and detailed bin-to-bin correlations not present in publicly-available data) in order to achieve more accurate results.

Not only is a muon program at neutrino sources economical, but it is also uniquely sensitive to muon-philic new physics scenarios, which is an under-explored region of parameter space due to the lack of dedicated muon sources. Finally, we emphasize the practical importance of these secondary muon programs: the costs are small and the required timescales are short. An investment in exploring the implications from muon physics at these experiments can not only inform and shape the future construction of other large-scale experiments, but also could provide the first hints towards understanding the persistent puzzles of the SM.

CHAPTER 6

NEW PHYSICS PROBES IN PRECISION MUON EXPERIMENTS

In this chapter we present another, very different example of using precision experiments to conduct discovery searches. In the Standard Model, there is perhaps one important experimental anomaly where theoretical predictions do not match observation. This is $(g - 2)_\mu$, the anomalous magnetic moment of the muon. This was first measured to be in tension with theory predictions at Brookhaven National Laboratory [142], and recently that measurement was confirmed at Fermilab [141]. The value of $(g - 2)_\mu$ determines the effective coupling between photons and the muon, and can be very precisely measured by storing muons in a ring with a background magnetic field and measuring the frequency at which they precess. This effect can be computed in the SM, and the effects from the Electroweak sector have been computed to extremely high precision. However, there is considerable uncertainty as to the contributions from QCD matter. These effect typically must be computed with lattice QCD, which offers an important foothold in making theory predictions. However, there is another data-driven approach: to determine the same amplitude in a controlled environment, then use the information gained to make a precise prediction for the hadronic contribution to muon $(g - 2)_\mu$. This is the purpose of the MUonE experiment.

MUonE was originally proposed to measure elastic μe scattering and extract the hadronic vacuum polarization contribution to the running of the electromagnetic fine-structure constant α . To make this measurement, MUonE has developed a tracking system capable of reconstructing charged tracks with precise angular resolution [143, 144]. This state-of-the-art tracking system also enables MUonE to detect displaced vertices, making it a promising platform for a long-lived particle search. Ref. [145] studied dark photon production through $\mu e \rightarrow \mu e A'$ with $A' \rightarrow e^+ e^-$, considering both a 4-track search including all final state leptons and a 3-track search where the primary electron need not be observed. Using a

simulation-based analysis, they illustrated that MUonE can probe regions in the “gap” for dark photon masses up to 70 MeV.

Interestingly, an experimental “gap” exists in parameter space where the hidden force coupling is too small for the mediator to be produced at colliders, but it also is not long-lived enough to provide a signal in existing beam dump experiments. For characteristic beam energies of $\mathcal{O}(100)$ GeV, a hidden boson in the “gap” would produce a displaced vertex a few centimeters away from the interaction point, and thus requires a compact detector geometry to be identified.

At the time of writing this thesis, MUonE finished its test run and is currently analyzing the data collected. The next stage is an official run of the experiment, projected to begin in early 2027 [146]. This recent progress calls for a detailed assessment of the sensitivity of MUonE to hidden force models. In this work, we consider three well-motivated vector boson models and generalize the displaced vertex search at MUonE to include both μN and μe scattering production channels. We conduct a simulation-based analysis and implement a search strategy following Ref. [145, 18], aimed at making reliable projections. Our main result is a characterization of the sensitivity of MUonE to the existence of a dark photon, $U(1)_{B-L}$ gauge boson, and $U(1)_{L_e-L_\mu}$ gauge boson.

This chapter is based on the work published in [147].

6.1 Simple Models of Vector Boson Mediators

We consider hidden force models that include a new massive $U(1)$ vector boson V_μ which couples to some (or all) of the SM fermions. The interaction term in the low energy limit is

$$\mathcal{L} \supset g_V \sum_{f \in \text{SM}} c_f \bar{f} \gamma^\mu V_\mu f, \quad (6.1)$$

where the sum is over all SM fermions. Here, $g_V c_f$ is the coupling between the massive vector boson and SM fermion f . Additional dark fermions or scalars may be present and charged under the new hidden force, thus forming a “dark sector.” We assume that any such fields are more massive than the mediator, so that the force carrier V decays only into SM products.

Notably, the list of viable $U(1)$ gauge symmetries that can be consistently added to the SM is severely restricted by anomaly cancellation. The models we consider can all be consistently added without requiring additional fields to cancel gauge anomalies, making them well-motivated and flexible extensions of the SM.

- **Dark photon:** this scenario posits a $U(1)_D$ under which no SM fields are charged. However, heavy particles charged under both $U(1)_D$ and $U(1)_{EM}$ can induce a small gauge-invariant kinetic mixing term in the IR. Therefore, the mediator A'_μ of the $U(1)_D$ inherits the same structure of charge assignments as the SM photon [148]. It couples to all charged particles in SM with strength $\epsilon e Q$, where ϵ is the mixing parameter, e is the QED coupling, and Q is the electric charge of the charged particle.
- **Gauged $U(1)_{B-L}$:** this scenario gauges the existing global $U(1)_{B-L}$ symmetry in the SM [149, 150]. This avoids a gauge anomaly if one assumes right-handed neutrinos exist, and we take them to have heavy Majorana masses. The $Z'_{\mu, B-L}$ gauge boson couples to all quarks (leptons) with coupling $g_{B-L}/3$ (g_{B-L}).
- **Gauged $U(1)_{L_e-L_\mu}$:** this scenario gauges a difference of L_e and L_μ [151]. The $Z'_{e\mu}$ gauge boson couples to the muon (muon neutrino) and electron (electron neutrino) with coupling $\pm g_{e\mu}$, respectively. In principle, any of the three combinations can be chosen, but $U(1)_{L_e-L_\mu}$ is the model for which MUonE is especially sensitive.

Their properties are summarized in Table. 6.1.

	g_V	c_u	c_d	c_ν	c_ℓ
Dark photon	ϵe	$2/3$	$-1/3$	0	-1
Gauged $L_e - L_\mu$	$g_{e\mu}$	0	0	$c_{\nu_e} = 1, c_{\nu_\mu} = -1, c_{\nu_\tau} = 0$	$c_e = 1, c_\mu = -1, c_\tau = 0$
Gauged $B - L$	g_{B-L}	$1/3$	$1/3$	-1	-1

Table 6.1: The vector boson mediator considered in this work. Columns show the notation of the coupling g_V and the charge of each SM fermions, respectively.




Figure 6.1: Feynman diagrams of vector boson production at MUonE.

Fig. 6.1 shows the Feynman diagrams for vector boson production at MUonE. All mediators can be produced by “dark bremsstrahlung” in μe or μN scattering. For μN scattering, bremsstrahlung off of the nucleus itself is severely suppressed, and thus for a given vector boson mass the cross section can be straightforwardly rescaled between models, i.e.

$$\frac{\sigma(\mu N \rightarrow \mu N V)}{\sigma(\mu N \rightarrow \mu N A')} = \frac{(g_V c_\mu)^2}{(\epsilon e)^2}, \quad (6.2)$$

where V stands for a general vector mediator. ϵe can be any reference value in principle, and is chosen to be $\epsilon = 10^{-3}$ during the production simulation. On the contrary, vector bosons from μe scattering can be emitted from both the muon and electron legs. The interference behavior changes depending on the relative sign of the Vee and $V\mu\mu$ vertices, an effect which we find numerically to be an $\mathcal{O}(10)\%$ change in cross section. This effect can be seen in Fig. 6.4.

The decay rate of the vector boson V_μ into a pair of SM fermions in the CoM frame is

$$\Gamma_{\bar{f}f} = \frac{N_f(g_V c_f)^2}{12\pi} m_V \left(1 + \frac{2m_f^2}{m_V^2} \right) \sqrt{1 - \frac{4m_f^2}{m_V^2}}. \quad (6.3)$$

In the region excluded by MUonE, the mass of the vector boson is below $2m_\mu$ and $2m_\pi$, so the only relevant decays are to electrons, neutrinos, or photons. For e and μ , $N_f = 1$. For ν , $N_f = 1/2$. Assuming only one decay channel into a pair of Dirac fermions ($N_f = 1$) with $m_f \ll m_V$, the decay length in the lab frame is given by

$$d_V = \frac{|\vec{p}_V|}{m_V \Gamma_V} \simeq 74 \text{ mm} \left(\frac{10^{-5}}{g_V c_f} \right)^2 \left(\frac{p_V}{10 \text{ GeV}} \right) \left(\frac{100 \text{ MeV}}{m_V} \right)^2, \quad (6.4)$$

where $\vec{p}_V$ is the momentum of the vector boson in the lab frame, so that $|\vec{p}_V|/m_V$ represents the boost factor times velocity. The decay length of other vector bosons can be computed in a similar way, with the width summed over all possible final states.

6.2 Vector Boson Signatures at the *MUonE* Experiment

In this section, we introduce the signal of displaced decays of BSM vector bosons at the MUonE experiment, following Ref. [145, 18].

6.2.1 *The Apparatus*

The MUonE experiment aims to measure the hadronic vacuum polarization (HVP) contribution to the running of α [143, 144]. To do so, MUonE will follow the strategy put forth by the NA7 experiment [152, 153]. This approach uses precise measurements of the μe elastic scattering cross section to extract the HVP contribution to α , providing an important cross-check with the result from lattice QCD [154]. This quantity is crucial for making a precise theoretical prediction of the muon anomalous magnetic moment [155, 156] and investigating




Figure 6.2: A schematic diagram for the MUonE apparatus with displaced vertex signature, borrowed from Fig. [18]. Orange line: incoming muon beam and outgoing scattered muon. Blue dashed line: displaced decaying vector boson. Green lines: decay products of the vector boson. Black squares: tracking stations. Gray squares: Beryllium targets. Blue and dark blue squares: the ECAL and the muon filter.

the current tension with experimental measurements [142, 157, 141].

To achieve this measurement, MUonE delivers a 160 GeV muon beam from the CERN M2 beamline onto 1.5 cm thick Beryllium (Be) targets. The targets are put into 40 identical modules aligned with the beamline, each separated by 1 m. Each module contains one Be target and three $10\text{ cm} \times 10\text{ cm}$ tracking stations, with the first station located 15 cm away from the target. The location of the 2nd and 3rd tracking stations has not been precisely determined yet. We assume the 3rd one to be located 1 m away from the target, while the 2nd one stands in the middle between the 1st one and the 3rd one. More tracking stations may be inserted in the middle to ensure achievement of the proposed angular and vertex reconstruction resolution. At the very end of the tracking stations, an electromagnetic calorimeter (ECAL) and a muon filter with size $1\text{ m} \times 1\text{ m}$ are placed for the purpose of particle identification (PID). With 2 years of data taking, the proposed total luminosity is $1.5 \times 10^4\text{ pb}^{-1}$ for μe scattering, corresponding to a total of $\sim 10^{16}$ muons on target.

The elastic $\mu e \rightarrow \mu e$ scattering can be distinguished from other events since the outgoing

muon angle θ_μ and electron angle θ_e strictly follow a one-on-one correspondence, known as the elastic curve [158]. To utilize the elastic curve for background rejection, MUonE employs a CMS-based high-precision tracking system for charged track reconstruction [144, 159, 160], which achieves an angular resolution within 0.02 mrad of the outgoing particles. No magnetic field is applied, so charge discrimination and momentum measurement are not possible. The energy of the outgoing muon is not measured, and the energy resolution of the ECAL has not been finally determined yet. For this reason, a missing energy search is not possible in MUonE.

Fig. 6.2 shows a schematic picture of the MUonE apparatus together with our proposed displaced vertex signature. The signal event topology is discussed in the next subsection.

6.2.2 Search Strategy

If a vector bosons from a hidden force are produced in MUonE, signal events would include a pair of charged tracks originating from a displaced vertex, corresponding to the decay of the vector boson. The final state thus contains three or four charged tracks: one coming from the primary interaction vertex, identified as the deflected muon, two reconstructed to be coming from a common displaced vertex and identified as a pair of electrons, and the potentially fourth one coming from the atomic electron that may also be involved in the production mechanism.

Fig. 6.3 shows the signal event process in a given module. For a successful charged track reconstruction, the tracks must pass through all three tracking stations in the module where they are produced. This requires the displaced decay to happen before the first tracking station. As shown in Ref. [145], the fourth charged track from the primary electron does not need to be observed to distinguish signal events from backgrounds, so we inclusively search for the final state with or without this electron being captured. For this reason, we do not show its track in Fig. 6.3. To avoid backgrounds from mis-reconstructing a prompt vertex as




Figure 6.3: Signal event geometry in the corresponding volume, borrowed from Fig. [18]. We show the ECAL and muon filter after the module, but note that this can be a few meters away from the module that we consider.

Variable	Selection Criteria
Decay z coordinate	$25 \text{ mm} < z < 140 \text{ mm}$
Decay daughter energy	$> 1 \text{ GeV}$
Decay daughter opening angle	$> 1 \text{ mrad}$
Charged track geometry	Hit all 3 trackers
Modules	Last 5

Table 6.2: Event selection criteria for our proposed search. In our numerical results, these requirements are imposed on our signal and background MC events. These criteria ensure that SM backgrounds are negligible for our signal of interest (see Supplemental Material for a discussion).

a displaced vertex, the displaced vertex should be at least $10\delta z$ away from both the target and the first tracking station, with $\delta z \simeq 1 \text{ mm}$ being the reconstruction resolution along the z direction (i.e. the streamline direction) [144, 145, 18, 146]. The decay volume is therefore defined as

$$25 \text{ mm} \leq z_{V \text{ decay}} \leq 140 \text{ mm}, \quad (6.5)$$

where $z_{V \text{ decay}}$ stands for the decay location of the vector boson along z coordinate. We further require the decay products to have an opening angle larger than 1 mrad to ensure accurate reconstruction.

For the purpose of background rejection, we require that all three charged tracks should enter the ECAL/muon filter for a successful PID. Energy loss for charged particles penetrating materials may significantly impact the PID efficiency. For example, an electron with energy 1 GeV will lose approximately 5% of its energy when passing through one single Be target, see Ref. [161, 53] for more details. For this reason, we only use the last 5 modules of the streamline as a conservative choice. Charged tracks that pass through all three tracking stations in these modules are guaranteed to be within the angular acceptance of the ECAL. An energy threshold of 1 GeV is applied to the muon and the electron-positron pair from displaced decay as a mimic of the detector response to avoid potential reconstruction errors.

We summarize the proposed search strategy in Table 6.2.

6.2.3 Signal Event Generation

To evaluate the sensitivity of MUonE to hidden vector bosons, we generate 20000 $\mu^- N \rightarrow \mu^- NV$ and $\mu^- e^- \rightarrow \mu^- e^- V$ events with a beam energy of 160 GeV using **MadGraph 3.2** [162, 163, 164, 165], with the model files constructed by **FeynRules** [132]. These events are weighted by the nuclear form factor, which dresses the electromagnetic vertex of the nucleus with $G(t)_2^{1/2} = \sqrt{G_2^{\text{el}} + G_2^{\text{in}}}$, where

$$G_2^{\text{el}}(t) = \left(\frac{a^2 t}{1 + a^2 t} \right)^2 Z^2 \left(\frac{1}{1 + t/d} \right)^2, \quad (6.6)$$

$$G_2^{\text{in}}(t) = \left(\frac{a'^2 t}{1 + a'^2 t} \right)^2 Z \left(\frac{1 + t(\mu_p^2 - 1)/(4m_p^2)}{(1 + t/(0.71 \text{ GeV}^2))^4} \right)^2. \quad (6.7)$$

For Beryllium, $a = 106Z^{-1/3}/m_e$, $d = 0.164 \text{ GeV}^2 A^{-2/3}$ for the elastic portion and $a' = 571Z^{-2/3}/m_e$ and $\mu_p = 2.79$ for the inelastic portion [15, 166]. This form factor is valid in the limit $t \ll m_p^2$ [167]. The **MadGraph** output contains very few events that violate this condition, but nonetheless, we discard those that approach it ($t > m_p^2/2$) by setting their




Figure 6.4: Production cross sections per unit charge, normalized to $g_V^2 = 1$ (Table. 6.1). At low masses, production from scattering off electrons is much more likely, as the electron is more likely to radiate out a gauge boson than the heavy nucleus, leading to a larger cross section. At large vector boson masses, production from nuclear scattering dominates, extending the experimental reach of MUonE to ~ 100 - 150 MeV.

weight to zero. We find that this has a negligible effect on the resulting bound.

Fig. 6.4 shows the production cross section for the vector gauge boson in the μe and μN channel as a function of the mass. The scaling behavior differs between the two production channels due to the inversion of the mass hierarchy between the target and beam. As shown in the plot, μe scattering dominates the cross section when the vector boson mass is below ~ 60 MeV because the target electron contributes significantly to the bremsstrahlung process, an effect which is heavily suppressed for the nuclear scattering case. On the other hand, μN scattering dominates in the large-mass regime because of the larger CoM energy for the nuclear target. With a beam energy of 160 GeV, the CoM energy in the nuclear scattering case is $\sqrt{s_{\mu N}} = 54$ GeV, compared to the electron scattering at just $\sqrt{s_{\mu e}} = 400$ MeV. The scattering off the nucleus thus enjoys a much wider available phase space, important for the production of a heavy vector boson. An important consequence of this is that adding the μN scattering significantly increases the maximum vector boson mass that MUonE is sensitive to.

In addition, as noted in Section 6.1, the relative sign of the electron and muon couplings to the dark vector boson in the $U(1)_{L_e-L_\mu}$ model can change the interference between diagrams in $\mu e \rightarrow \mu e V$. Numerically, we find this effect to give an $\mathcal{O}(10)\%$ correction.

We comment on the analytical calculation of the cross section with Weizsaecker-Williams (WW) approximation [168, 169]. Using the results of Ref. [170] with the WW approximation, we find a factor of 1-2 difference in cross section with the results from **MadGraph**. However, we find a larger discrepancy of a factor of ~ 8 between the **MadGraph** result and the “Improved WW” technique used by Ref. [171] to construct analytic MUonE bounds. This discrepancy could be explained by the assumption that $m_{beam} \ll m_{A'}$ that is necessary for the derivation of the cross section in Ref. [170] (see e.g. Ref [167]), but invalid for the muon-flavored beam used by MUonE. The difference in predicted cross sections partly explains the difference in projected sensitivities in this work and in Ref. [171], with the other differences discussed in Sec. 6.3.

Figs. 6.5 and 6.6 display the kinematic distributions of the products in $\mu e \rightarrow \mu e V$ and $\mu N \rightarrow \mu N V$ scatterings for $m_V = 12$ MeV and $m_V = 44$ MeV. As seen in the figures, scattering off electrons produces much softer vector gauge bosons than those produced from scattering off nuclei. The energy cut of $E_{\text{prod}} > 1$ GeV implies that vector bosons of $E_V < 2$ GeV are completely invisible. Above this threshold, the probability of passing the E_{prod} cut for both resulting electrons increases and smoothly approaches unity at high E_V . For $m_V = 12$ MeV, a significant portion of the produced vector bosons lie well below 2 GeV, leading to a low detection efficiency. This offsets the much greater production cross section of the electron channel, so both channels yield a similar number of signal events, as shown later in Sec. 6.3. For $m_V = 44$ MeV, the energies of the vector bosons produced in μe scattering are above 2 GeV, which improves the efficiency. However, at $m_V = 44$ MeV, the production cross sections for both channels are comparable, again resulting in a similar contribution to signal events from the nuclear and electron scattering channels.




Figure 6.5: Energy and angle distribution for the produced vector boson and deflected muon for $m_{A'} = 12$ MeV.

After production via dark bremsstrahlung, the vector boson decays, with decay length given in Eq. 6.3. The kinematics of the decay are computed manually by boosting to the rest frame of the vector boson, where the decay occurs isotropically, and then inverse boosting. In the regions of parameter space excluded by MUonE, the $m_V < 2m_\mu$, so the only available decay modes at tree-level are to electrons or neutrinos. Electrons are the only visible channel, and the visible branching ratios (for $m_A \gg 2m_e$) are $1, 2/5, 1/2$ for the dark photon, $U(1)_{B-L}$, and $U(1)_{L_e-L_\mu}$ models, respectively. The lower branching ratio to visible products can slightly reduce the reach of MUonE, see Sec. 6.3 for a discussion. Below $m_{A'} < 2m_e$, the dark vector decays to neutrinos or is very nearly stable, and both scenarios are invisible to MUonE. After decay, the resulting visible leptons are required to pass the




Figure 6.6: Energy and angle distribution for the produced vector boson and deflected muon for $m_{A'} = 44$ MeV.

cuts in Table. 6.2, and we compile those that pass into an expected number of signal events.

6.2.4 Backgrounds

The background falls into two categories: (a) SM processes that happen within the target or the first tracking station and create a prompt vertex, which may be mis-reconstructed as a displaced vertex due to limited resolution. (b) SM processes that can naturally produce a displaced vertex. The first category can be safely rejected by requiring a $10\delta z$ displacement away from the target and the first tracking station, while the second category requires a simulation to examine its impact.

Ref. [145] performed a systematic, simulation-based analysis for the background process

with an inclusive fourth track. It was concluded that the backgrounds can all be safely rejected. Most hadronic processes come with many extra charged particles that leave hits on the tracking stations or the ECAL, leading to rejection. The remaining small portion of events is relatively clean, and the PID system is required to reject them. The vast majority of background events arise from $K^0 \rightarrow \pi^+\pi^-$ and $\Lambda \rightarrow p\pi^-$. The simulation finds that about 4000 such background events are expected to fake the signal topology, and can be safely rejected by ECAL as long as the per-particle-fake rate is less than 1.6%. A small number of events have $K^0 \rightarrow \pi e \nu$ or $\Lambda \rightarrow p e \nu$, where one real electron/positron is faking the displaced decay and the other is from a charged pion or proton. Simulation finds that about 8 events are expected, and can be rejected by the ECAL. Finally, about 0.5 events are expected to have two real electrons faking the displaced decay, such as $K^0 \rightarrow \pi^0 \pi^0$ with one π^0 decaying into $e^+e^-\gamma$. Fortunately, all these events have at least one non-electron particle hitting the ECAL (usually a charged pion from the primary vertex, or an energetic photon), leading these events to be rejected. As a conclusion, all backgrounds can be safely rejected by our search strategy, and we assume the backgrounds are zero in our statistics.

6.3 Analysis and Results

Based on the discussion of Sec. 6.2, we perform a parameter scan over the coupling-mass plane. We generate signal event samples for 16 mass points logarithmically spaced between 2 MeV and 300 MeV. For each mass point, we apply the selection criteria on an event-by-event basis for couplings ranging from 10^{-6} –1, covering 30 logarithmically spaced points. For each parameter point, the total number of signal events is calculated by

$$N_{\text{sig}} = \epsilon_{\text{sig},N} \sigma_N \mathcal{L}_N + \epsilon_{\text{sig},e} \sigma_e \mathcal{L}_e . \quad (6.8)$$

Here $\epsilon_{\text{sig},N(e)}$ is the fraction of signal events from μN (μe) scattering that pass the cuts in Table. 6.2, $\sigma_{N,e}$ is the production cross section calculated using **MadGraph**, and $\mathcal{L}_{N,e}$ is the effective luminosity. For μN scattering, $\mathcal{L}_N = 1.5 \times 10^4/8/4 \text{ pb}^{-1}$, while for μe scattering we have $\mathcal{L}_e = 1.5 \times 10^4/8 \text{ pb}^{-1}$. The factor of 8 comes from the fact that only last 5 modules can be used (Sec. 6.2.2), while the factor of 4 in $\mathcal{L}_N$ comes from the Be atomic number.

Figs. 6.7–6.9 show our main results for the dark photon, $U(1)_{B-L}$ gauge boson, and $U(1)_{L_e-L_\mu}$ gauge boson. For each plot, we show the projected sensitivity at 95% confidence level assuming zero background from μe and μN scattering with red dotted and dashed curves, respectively. The combined result is shown with the red solid curve. In all cases, we find that the μe scattering can reach a lower coupling than the μN scattering with vector boson mass below $\sim 30 \text{ MeV}$. This is due to the larger energy acquired by the vector boson from the Be nucleus as seen in Figs. 6.5 and 6.6, which leads to a longer decay length that should be compensated by a larger coupling. The μN scattering, on the other hand, can reach a much higher vector boson mass because of the larger CoM energy.

The existing limits are shown in the gray shaded area, while future sensitivities within the next 5 years are shown by gray dashed curves. Our recasting of these limits relies on DARKCAST [172]. We also refer the readers to Ref. [173] for a comprehensive study on $U(1)$ vector bosons. Note that all of our searches are based on the assumption that the vector bosons only decay back to SM particles. Invisible searches, assuming decays into stable dark sector particles, do not apply.

In the following subsections, we comment on our results and emphasize the differences among these models.

6.3.1 Dark photon model

Fig. 6.7 shows the MUonE projection for dark photon searches. We found that the MUonE reach extends to dark photon masses as heavy as $\sim 150 \text{ MeV}$ for couplings that fall into




Figure 6.7: The result for dark photon search. MUonE sensitivity is shown with the red dotted curve (for μe scattering), the red dashed curve (for μN scattering), and the red solid curve (combined). Existing limits are shown in gray shaded regions. Selected future limits that are planned for a similar timeline as MUonE are also shown. Specifically, we add an arrow for the SeaQuest projection to indicate that it is expected to probe the parameter space below the curve.

the “gap”. Note that the sensitivity is significantly smaller than that in Ref. [171]. Besides the discussion of the WW approximation in Sec. 6.2.3, the restriction to only the last 5 modules suppresses the luminosity. We point out that this is essential to have a reliable PID accuracy for background rejection. Existing limits are from $(g - 2)_e$ [174, 175], beam dumps [176, 177, 178, 179, 180, 167, 181, 182, 183, 184, 185, 186, 187, 188, 189, 190, 191, 192]), electron fixed-target [193, 194], collider experiments [195, 196, 197, 198, 199, 200, 201, 202], and rare meson decays [189, 203, 200, 190, 204, 205, 206, 207, 208, 201]. Future projections shown here include HPS [209, 210], NA64e with 10^{12} electrons on target [211], D^* decay search at LHCb upgrade [212, 213], SeaQuest [214], and Mu3e [215].




Figure 6.8: The result for gauged $U(1)_{B-L}$ search. MUonE sensitivity is shown with the red dotted curve (for μe scattering), the red dashed curve (for μN scattering), and the red solid curve (combined). Existing limits are shown in gray shaded regions. Selected future limits that are planned for a similar timeline as MUonE are also shown. Specifically, we add an arrow for the SeaQuest projection to indicate that it is expected to probe the parameter space below the curve.

6.3.2 The gauged $U(1)_{B-L}$ model

Fig. 6.8 shows the MUonE sensitivity to the $U(1)_{B-L}$ gauge boson. The existing dark photon limits and projections discussed in the previous subsection are re-interpreted for the $U(1)_{B-L}$ gauge boson using DARKCAST [172].

The main difference arises from the coupling to the neutrinos. This opens an additional decay channel into a neutrino pair, shortening the decay length by a factor of $2/5$ compared to the dark photon model. To maintain the same decay length, the coupling must be reduced slightly, which in turn lowers the cross section. Moreover, since neutrinos are invisible, the observable signal rate is further suppressed by the branching ratio into e^+e^- , which carries the same factor of $2/5$. As a result, the largest coupling that is accessible for MUonE for a




Figure 6.9: The result for gauged $U(1)_{L_e-L_\mu}$ search. MUonE sensitivity is shown with the red dotted curve (for μe scattering), the red dashed curve (for μN scattering), and the red solid curve (combined). Existing limits are shown in gray shaded regions. Selected future limits that are planned for a similar timeline as MUonE are also shown. Specifically, we add an arrow for the SeaQuest projection to indicate that it is expected to probe the parameter space below the curve.

given mass drops by an $\mathcal{O}(1)$ factor, and the maximum mass also reduces to about 100 MeV, slightly weaker than in the dark photon case. In addition, the $U(1)_{B-L}$ gauge theory is further constrained by neutrino experiments such as Borexino [216, 217], Texono [218], CHARM II, SPEAR, DORIS, and PETRA [219, 220], resulting in the extra gray-shaded region for masses around a few tens of MeV. These additional limits cover the projection of Mu3e, so it is not shown in Fig. 6.8. We note that the HPS projection is absent due to the lack of sufficient information to translate the result into this model.

6.3.3 The gauged $U(1)_{L_e-L_\mu}$ model

Fig. 6.9 shows the MUonE sensitivity for the gauged $U(1)_{L_\mu-L_e}$ model, with existing dark photon limits re-interpreted by running DARKCAST [172]. The discussion on the impact of neutrino coupling for the $U(1)_{B-L}$ model also applies here, though this gauge boson couples to one fewer neutrino flavor. This slightly enhances the sensitivity over the $U(1)_{B-L}$ case.

Unlike the previous two models, the $U(1)_{L_e-L_\mu}$ gauge boson does not couple to quarks at tree-level (although it can do so at loop-level). Consequently, the production of $U(1)_{L_\mu-L_e}$ gauge boson in hadronic experiments such as CHARM, NuCAL, and NOMAD is strongly suppressed. Nevertheless, electron beam-dump experiments have achieved similar sensitivities, so the existing bounds are not strongly affected.

The absence of a τ neutrino coupling suppresses the sensitivity of many neutrino experiments, but raises it in Super-K due to modified neutrino oscillations [221]. Among the neutrino experiments mentioned in the $U(1)_{B-L}$ case, the Texono [218] and Borexino [216, 217] remain relevant as their primary search channel is through scattering between electrons and electron neutrinos. As the $U(1)_{B-L}$ case, the Mu3e and HPS projections are absent for the reasons explained above.

6.4 Discussion

In this paper, we demonstrate that the MUonE experiment provides a unique opportunity to probe hidden vector bosons through displaced decay signatures. In exploring the dark photon, $U(1)_{B-L}$, and $U(1)_{L_e-L_\mu}$ models as benchmark examples, we extended the prior work of Ref. [145] to include the nuclear scattering, $\mu N \rightarrow \mu NV$, in addition to the $\mu e \rightarrow \mu e V$ channel. We perform a dedicated simulation of signal events and apply the event selection criteria. Our analysis illustrates that μN scattering, owing to its higher CoM energy, substantially enhances the reach of MUonE for heavier mediators, extending the sensitivity up to $m_V \sim 150$ MeV, a factor of two improvement over the case with μe scattering

alone. Together with the performance of μe scattering at low mass, the combined production channels show a promising sensitivity.

CHAPTER 7

CONCLUSION

There are many reasons to expect physics beyond the Standard Model, yet no new particles have been discovered at the LHC. That being said, it's not known whether the next discovery in physics will be above, at, or below the Electroweak scale, and each regime has viable models that serve to answer outstanding questions in physics. Between the LHC and the FCC-hh, the focus of experimental physics will be towards precision experiments, with important questions like: can we find any deviations from SM predictions in $\mu \rightarrow e\gamma$, the electron's electric dipole moment, or in $(g - 2)_\mu$? Ultimately, even if all of these experiments return null results, these efforts will help strengthen constraints on the scale of new physics, and inform search strategies and detector design at the FCC-hh.

As we have seen throughout this thesis, this period should not really be viewed as a staging ground, but a period when many well-motivated models will be probed indirectly. The HL-LHC offers the possibility of not only discovering new heavy states, but providing the first hint of new physics through flavor-violating decays of the Higgs. In particular, lepton precision experiments and LHC heavy Higgs searches complement each other, together exploring regions of parameter space where loop-induced lepton flavor violation is a prominent signal. The FCC-ee could discover signs of new physics in any of its proposed experimental programs, and we have characterized its capabilities particularly in the case of the electron Yukawa. To this end, electron $g - 2$ experiments offer competing bounds, in particular when they are suppressed at one-loop with respect to electron Yukawa enhancements. Neutrino experiments, which provide concrete evidence that the SM is incomplete, also offer insights into lower mass scale particles, in a region that also contains many candidates for new physics and possible connections to dark matter. And lastly, the $(g - 2)_\mu$ tension remains to be resolved, and we showed that the same experiments can be used to probe new gauge forces.

This thesis highlights the interplay between the categories of precision and discovery, which are usually considered as two distinct experimental philosophies. The line between them has blurred, and often the utility of such experiments can be substantially broadened by considering the same apparatus through the opposite lens. Whether or not new physics reveals itself through lepton precision experiments remains to be seen, but in this pursuit, it's vital that we consider each experiment through the lenses of both precision and discovery.

APPENDIX A

GENERATING YUKAWAS AT LOOP LEVEL IN THE MSSM

In the MSSM, the holomorphicity of the superpotential imposes a $\mathbb{Z}_2$ symmetry on the dimension 4 terms in the lagrangian, resulting in a type-II 2HDM where the leptons couple only to Φ_1 . This $\mathbb{Z}_2$ is softly broken by the so-called μ -term.¹ At the same time, SUSY is broken softly by the masses of the superpartners, which for the sleptons are parameterized by a scalar mass matrices $m_{\ell_L}^2, m_{e_R}^2$. These mass matrices are additional spurions for flavor symmetry.

Notably, since the sleptons inherit the same quantum numbers as the leptons, they also transform the same under the global symmetries of the MSSM, including the flavor $U(3)^5$. This dictates how the slepton mass matrices can be combined with factors of the lepton Yukawas to generate corrections. The charge assignments under the global $U(3)^5$ for the flavor spurions in the MSSM are

$$m_{\ell_L}^2 \sim \mathbf{3}_L \times \bar{\mathbf{3}}_L \quad , \quad m_{e_R}^2 \sim \mathbf{3}_e \times \bar{\mathbf{3}}_e \quad , \quad y_1 \sim \mathbf{3}_L \times \bar{\mathbf{3}}_e \quad (\text{A.1})$$

All combinations of parameters that have the opposite flavor charge as $\bar{\ell}_L e_R$ can appear through loop effects in the 2HDM. With μ as a spurion for $\mathbb{Z}_2$ and $m_{\ell_{L/R}}^2$ as spurions for flavor symmetry, we can construct a loop correction to the Yukawas that transmit these spurions to the 2HDM subsector of the MSSM. For some regular functions g, h , the loop contribution take the form, schematically,

$$\mathcal{L}_{2HDM} \supset \sum_{\text{diagrams}} \frac{k}{16\pi^2} \mu/M \times \left[g(m_{\ell_L}^2/M^2) y_1 h(m_{e_R}^2/M^2) \right]_{ij} \left(\bar{\ell}_L^i \Phi_2 e_R^j \right) . \quad (\text{A.2})$$

Here, k is combination of couplings appearing in the loops, and M is at the scale of the

1. $\mathbb{Z}_2$ is usually also softly broken by a $\Phi_1^\dagger \Phi_2$ mass term, in order to avoid any domain wall problem.




Figure A.1: Loop diagram that leads to flavor violation the MSSM.

superpartners. One important corresponding loop diagram in the MSSM, which indeed has this form, is depicted in Fig. A.1, and is computed in more detail in Ref. [222]. In our analysis, we would assume that the functions g and h are sufficient to randomize the flavor coefficients of the loop correction, so that the y_2 that is generated is statistically independent from y_1 .

APPENDIX B

LOOP FUNCTIONS FOR THE DIPOLE OPERATORS

We list the loop functions used in this analysis. The complete list of diagrams and their loop functions is given in Ref. [1]. The loop function often use the special functions $\Phi(x)$ and $\Phi(x, y)$, defined in Ref. [223] to be

$$\Phi(z) = \frac{1}{\lambda} \left[-4\text{Li}_2\left(\frac{1}{2}(1-\lambda)\right) + 2\log^2\left(\frac{1}{2}(1-\lambda)\right) - \log^2(4z) + \frac{1}{3}\pi^2 \right], \quad (\text{B.1})$$

where in this expression, $\lambda = \sqrt{1 - 1/z}$, and

$$\begin{aligned} \Phi(x, y) = \frac{1}{\lambda} & \left(2\log\left(\frac{1}{2}(1+x-y-\lambda)\right) \log\left(\frac{1}{2}(1-x+y-\lambda)\right) - \log(x)\log(y) \right. \\ & \left. - 2\text{Li}_2\left(\frac{1}{2}(1+x-y-\lambda)\right) - 2\text{Li}_2\left(\frac{1}{2}(1-x+y-\lambda)\right) + \frac{\pi^2}{3} \right), \quad (\text{B.2}) \end{aligned}$$

where here $\lambda = \sqrt{(1-x-y)^2 - 4xy}$. The loop functions for the fermion loops are

$$f_1(x) = \Phi(1/4x) \qquad f_2(x) = (1-2x)f_1(x) + 2(2+\log(x)) \quad (\text{B.3})$$

The loop functions involving the charged higgs are

$$\tilde{f}_5(x) = 3x - \frac{1}{2}(13-6x)\log x + (2-x)(2-3x) \left[\text{Li}_2(1-1/x) - \frac{\pi^2}{6} \right] \quad (\text{B.4})$$

$$f_5(x, y) = \frac{y}{x-y} \left[\tilde{f}_5(x) - \tilde{f}_5(y) \right] \quad (\text{B.5})$$

$$\tilde{f}_6(x,) = -3x + \frac{1}{2}(1-6x)\log x + x(2-3x) \left[\text{Li}_2(1-1/x) - \frac{\pi^2}{6} \right] \quad (\text{B.6})$$

$$f_6(x, y) = \frac{y}{x-y} \left[\tilde{f}_6(x) - \tilde{f}_6(y) \right] \quad (\text{B.7})$$

And lastly the loop integral from the Higgs kinetic term,

$$\begin{aligned}
F_{\text{kin}}(\{m_{h^k}\}, s_w) = & (s_w^2 - 4) \left(f_{8a}(x_{Zh_k}, c_w^2) - f_{8a}(x_{Zh_1}, c_w^2) + \right. \\
& \left. \frac{1}{s_w^2} (f_{8b}(x_{Zh_k}, c_w^2) - f_{8b}(x_{Zh_1}, c_w^2)) \right) + \\
& f_{8c}(x_{Wh_k}) - f_{8c}(x_{Wh_1}) + \\
& \frac{1}{2s_w^2} \left(f_{8d}(x_{Wh_k}) - f_{8d}(x_{Wh_1}) + \right. \\
& \left. f_{8e}(x_{Wh_k}, x_{WH^\pm}) - f_{8e}(x_{Wh_1}, x_{WH^\pm}) \right)
\end{aligned} \tag{B.8}$$

Where

$$f_{8a}(x, y) = \log(x) \frac{1 + 2xy}{y(x - 1)}, \quad (\text{B.9})$$

$$f_{8b}(x, y) = \frac{12y^2 - 16y + 3}{1 - x} x \Phi(1/4y) + \log(x) \frac{10xy + 1}{1 - x} + \frac{12x^2y^2 - 16xy + 3}{x - 1} \Phi(1/4xy), \quad (\text{B.10})$$

$$f_{8c}(x) = (12x^2 - 16x + 3) \Phi_1(1/4x) - 2[\log(x) + 2](6x + 1) \quad (\text{B.11})$$

$$f_{8d}(x) = \frac{4x^2 + 3x}{9} \pi^2 + \log^2(x) \frac{8x^3 + 6x^2 - 2}{6x} - \log(x) \frac{8(x + 1)}{3} + \frac{8x^4 + 6x^3 - 2x - 3}{3x^2} \text{Li}_2(1 - x) + \frac{2x^3 - 19x^2 - 4x + 3}{6x^2} \Phi(1/4x) - \frac{8x - 5}{3}, \quad (\text{B.12})$$

$$f_{8e}(x, y) = \log^2(x) \frac{x^3y(3 - y) + 3x^2y^2 - x(3y^2 + 4y + 1) + y^2 + y}{2x^3} + \log(x) \log(y) \frac{x^3(y^3 - 4y^2 + 3y) - 3x^2y^2(y - 1) + xy^2(3y + 1) - y^3}{2x^3(y - 1)} + \log(x) \frac{y - xy - x}{x^2} + \log(y) \frac{x^2(1 - 2y^2 + 7y) + xy(4y + 1) - 2y^2}{2x^2(y - 1)} + \frac{x^4(3y - y^4 + 5y^3 - 7y^2) + x^3(4y^4 - 8y^3)}{2x^4(y - 1)} - \frac{x^2(6y^4 - y^3 + y^2) - x(4y^4 + 2y^3) + y^4}{2x^4(y - 1)} \times \Phi(y, y/x) + \frac{x^3(3y - y^2) + 3x^2y^2 - x(3y^2 + 4y) + y^2 + y}{x^3} \text{Li}_2(1 - x) + \frac{4x^3 + 6x^2 - 6x + 1}{2x^3(y - 1)} y \Phi(1/4x) + \frac{2xy - x^2y + x - y}{x^2}. \quad (\text{B.13})$$

APPENDIX C

LOOP CORRECTIONS TO THE DARK-SIDE

In the type-II 2HDM, we must ensure that the loop corrections to κ_τ do not split κ_τ and κ_b beyond the measured difference. Using Table. 3.1 we can write

$$\kappa_\tau - \kappa_b = \frac{\cos(\beta - \alpha) \tan \beta}{1 + \varepsilon^2 \tan^2 \beta} \left(\varepsilon^2 \tan^2 \beta + \frac{a_{\tau\tau} \varepsilon \tan \beta}{\sin^2 \beta} \right) \quad (\text{C.1})$$

In the intermediate region $1 \ll \tan \beta \ll \varepsilon^{-1}$ and using the tree-level cancellation condition $\cos(\beta - \alpha) = 2 \cot \beta$, the experimental constraints on κ_b and κ_τ from ATLAS (Table 3.3) imply that the $\mathcal{O}(1)$ prefactor must satisfy

$$|a_{\tau\tau}| \lesssim \frac{0.2}{\varepsilon \tan \beta} \quad (\text{C.2})$$

For $\varepsilon = 0.1/(4\pi)^2$, the relevant limits are all stronger than about $\tan \beta \lesssim 30$, implying the constraint on $a_{\tau\tau}$ is no stronger than roughly $|a_{\tau\tau}| \lesssim 10$. In conclusion, loop corrections will provide an irrelevant correction to the tree-level dark side condition in Eq. 3.68.

APPENDIX D

THEORETICAL EXPRESSIONS FOR MESON DECAYS

The EFT operator coefficients $\mathcal{C}_{ledq}$ and $\mathcal{C}_{lequ}^{(1,3)}$ can contribute to electronic meson decays with large chiral enhancements over the SM.

D.1 $B_s \rightarrow ee$

The low-energy effective Lagrangian describing $b \rightarrow s\ell\ell$ transitions is given by

$$\mathcal{L}_{bs\ell\ell} \supset \frac{4G_F}{\sqrt{2}} V_{ts}^* V_{tb} \sum_k C_k O_k, \quad (\text{D.1})$$

with operators

$$O_9 = \frac{\alpha}{4\pi} (\bar{s}_L \gamma_\mu b_L) (\bar{\ell}_i \gamma^\mu \ell_i), \quad O_{10} = \frac{\alpha}{4\pi} (\bar{s}_L \gamma_\mu b_L) (\bar{\ell}_i \gamma^\mu \gamma^5 \ell_i), \quad (\text{D.2})$$

$$O_S = \frac{\alpha}{4\pi} (\bar{s}_L b_R) (\bar{\ell}_i \ell_i), \quad O_P = \frac{\alpha}{4\pi} (\bar{s}_L b_R) (\bar{\ell}_i \gamma^5 \ell_i). \quad (\text{D.3})$$

The SM values of the corresponding Wilson coefficients are [224]:

$$C_9^{\text{SM}} = 4.114, \quad C_{10}^{\text{SM}} = -4.193, \quad C_S^{\text{SM}} = 0, \quad C_P^{\text{SM}} = 0. \quad (\text{D.4})$$

The SMEFT Wilson coefficient $\mathcal{C}_{ledq}$ with third generation quark indices matches at the electroweak scale onto the coefficients C_S and C_P as:

$$C_S = -C_P = \frac{\pi v^2}{\alpha V_{tb}} [\mathcal{C}_{ledq}]_{ii33}, \quad (\text{D.5})$$

where the quark doublet flavour index is defined in the up quark mass basis, i.e. $q_3 \equiv (t, V_{tj} d_j)^T$. Note that if instead this coefficient were aligned in the down quark mass basis, there would be no down-type flavour change.

The coefficients C_S and C_P receive a multiplicative factor from the running from the electroweak scale to the b mass scale: $C_{S,P}(m_b) = 1.38 C_{S,P}(m_Z)$ [70]. Then the branching ratio for $B_s \rightarrow ee$ is given by [224]:

$$\begin{aligned} \text{BR}(B_s \rightarrow ee) = \text{BR}(B_s \rightarrow ee)_{\text{SM}} & \left(\left| 1 + 1.38 \frac{C_P}{C_{10}^{\text{SM}}} \frac{m_{B_s}^2}{2m_e(m_b + m_s)} \right|^2 \right. \\ & \left. + \left| 1.38 \frac{C_S}{C_{10}^{\text{SM}}} \frac{m_{B_s}^2}{2m_e(m_b + m_s)} \right|^2 \right). \end{aligned} \quad (\text{D.6})$$

The factor of m_e in the denominator of the terms dependent on C_S and C_P gives an enormous chiral enhancement to decays mediated by these scalar and pseudoscalar operators. The SM prediction for this decay is very small [225, 226]:

$$\text{BR}(B_s \rightarrow ee)_{\text{SM}} = (8.30 \pm 0.36) \times 10^{-14}, \quad (\text{D.7})$$

and experimentally only an upper limit has been set, several orders of magnitude larger than the SM expectation [92]:

$$\text{BR}(B_s \rightarrow ee) < 11.2 \times 10^{-9} \quad (95\% \text{ CL}). \quad (\text{D.8})$$

This translates to a bound on $[\mathcal{C}_{ledq}]_{1133}$ of

$$|[\mathcal{C}_{ledq}]_{1133}| < 4.1 \times 10^{-3} \text{ TeV}^{-2} \quad (95\% \text{ CL}). \quad (\text{D.9})$$

We emphasise that the flavour index on the quark doublet is assumed to be in the up-aligned basis here. If instead the index is down-aligned, then this coefficient does not mediate a $b \rightarrow s$ flavour change and there is no bound on the coefficient from this process.

D.2 $D_s \rightarrow e\nu$

The low-energy effective Lagrangian describing $c \rightarrow s\ell\nu$ decays is written

$$\begin{aligned} \mathcal{L}_{cs\ell\nu} = & -\frac{4G_F V_{cs}}{\sqrt{2}} \left[(1 + \epsilon_L)(\bar{c}_L \gamma^\mu s_L)(\bar{\ell}_L \gamma_\mu \nu_L) + \epsilon_R(\bar{c}_R \gamma^\mu s_R)(\bar{\ell}_L \gamma_\mu \nu_L) \right. \\ & \left. + \epsilon_T(\bar{c}_L \sigma^{\mu\nu} s_L)(\bar{\ell}_L \sigma_{\mu\nu} \nu_L) + \epsilon_{S_L}(\bar{c}_R s_L)(\bar{\ell}_R \nu_L) + \epsilon_{S_R}(\bar{c}_L s_R)(\bar{\ell}_R \nu_L) \right] + h.c.. \end{aligned} \quad (\text{D.10})$$

In the SM, all the ϵ coefficients are zero. The branching ratio of $D_s^+ \rightarrow e^+ \nu$ is given by [227]

$$\text{BR}(D_s^+ \rightarrow e^+ \nu) = \frac{G_F^2}{8\pi} \tau_{D_s} f_{D_s}^2 |V_{cb}|^2 m_{D_s} m_e^2 \left| 1 + \epsilon_L + \frac{m_{D_s}}{m_e(m_c + m_s)}(\epsilon_{S_R} - \epsilon_{S_L}) \right|^2. \quad (\text{D.11})$$

where the ϵ coefficients are assumed to be defined at the c mass scale. The decay constant is $f_{D_s} = 249.9 \pm 0.5$ MeV [228, 229, 230]. The SMEFT coefficient $[\mathcal{C}_{lequ}^{(1)}]_{1122}$ matches to ϵ_{S_L} as:¹

$$\epsilon_{S_L}(m_c) = -1.78 \frac{v^2}{2V_{cs}} [\mathcal{C}_{lequ}^{(1)*}]_{1122}(m_W), \quad (\text{D.12})$$

where the factor of 1.78 arises from the running from the electroweak scale to the c mass scale [70]. We assume here that the quark doublet flavour index is defined in the up mass basis; if instead it is defined in the down mass basis, the factor of V_{cs} would not appear in the denominator of the matching expression.

The measured upper bound on the branching ratio is three orders of magnitude larger than the SM expectation [93]

$$\text{BR}(D_s^+ \rightarrow e^+ \nu) < 1.0 \times 10^{-4} \quad (95\% \text{ CL}), \quad (\text{D.13})$$

1. Note that in general, there are other SMEFT matching relations that could in principle play a role; for example $[\mathcal{C}_{lequ}^{(3)}]_{1122}$ matches to ϵ_T and $[\mathcal{C}_{ledq}]_{1122}$ matches to ϵ_{S_R} . However the former does not contribute to the $\text{BR}(D_s^- \rightarrow e^- \bar{\nu})$ (except in very small amounts through RG mixing of ϵ_T into ϵ_{S_L}), and the latter does not feature in our main analysis since it cannot contribute measurably to the electron Yukawa.

leading to a bound on the SMEFT Wilson coefficient of

$$|[\mathcal{C}_{lequ}^{(1)*}]_{1122}| < 0.18 \text{ TeV}^{-2} \quad (95\% \text{ CL}). \quad (\text{D.14})$$

APPENDIX E

$\mu \rightarrow e$ CONVERSION

Assuming all of the quark Yukawas take their SM values, the leading contributions to the rate for $\mu \rightarrow e$ conversion can be written following Appendix 2 in [44],

$$\Gamma(\mu \rightarrow e) = \left| -\frac{D}{2m_\mu} [Le\gamma(m_h)]_{e\mu} + \tilde{g}_{LS}^{(p)} S^{(p)} + \tilde{g}_{LS}^{(n)} S^{(n)} + \tilde{g}_{LV}^{(p)} V^{(p)} \right|^2 + \quad (\text{E.1})$$

$$\left| -\frac{D}{2m_\mu} [Le\gamma(m_h)]_{\mu e}^* + \tilde{g}_{RS}^{(p)} S^{(p)} + \tilde{g}_{RS}^{(n)} S^{(n)} + \tilde{g}_{RV}^{(p)} V^{(p)} \right|^2, \quad (\text{E.2})$$

where the terms containing $Le\mu$ follow from photon-exchange t-channel diagrams, the S terms follow from tree-level Higgs exchange, and the V terms from an evaluation of the 1-loop diagram. For further detail on the V term, we refer to ref. [44], however assuming $\mathcal{C}_{eH}$ is real, we find approximately that

$$\tilde{g}_{LV/RV} \approx (1.1 \times 10^{-9}) \times \frac{\alpha v^2}{2\sqrt{2}\pi m_\mu^2} [\mathcal{C}_{eH}(m_h)]_{e\mu/\mu e}. \quad (\text{E.3})$$

The S terms contain quark Yukawas and factors of the nuclear matrix elements $f^{(q,N)} = \langle N | m_q \bar{q}q | N \rangle / m_p$ where $N \in \{n, p\}$. In our scenario they are given by

$$g_{LS}^{(N)} = -\sqrt{2} \frac{m_p v}{m_h^2} [\mathcal{C}_{eH}(m_h)]_{e\mu} \sum_q f^{(q,N)} \quad (\text{E.4})$$

$$g_{RS}^{(N)} = -\sqrt{2} \frac{m_p v}{m_h^2} [\mathcal{C}_{eH}(m_h)]_{\mu e}^* \sum_q f^{(q,N)}. \quad (\text{E.5})$$

Numerically, the matrix elements $f^{(q,N)}$ are given as well in ref. [44], taking the values

$$f^{(u,p)} = f^{(d,n)} = 0.024 \quad , \quad f^{(d,p)} = f^{(u,n)} = 0.033 \quad (\text{E.6})$$

$$f^{(s,p)} = f^{(s,n)} = 0.25 \quad , \quad f^{(other,p)} = f^{(other,n)} = 0.051. \quad (\text{E.7})$$

Finally, in units of $m_\mu^{5/2}$, nuclear overlap integrals and muon capture rates are given by ref. [231], summarized in the following table.

	D	$S^{(p)}$	$S^{(n)}$	$V^{(p)}$	$\Gamma_{\text{capture}}/(10^6 s^{-1})$
Gold	0.189	0.0614	0.0918	0.0974	13.07
Aluminum	0.0362	0.0155	0.0167	0.0161	0.7054

Table E.1: Wavefunction overlap integrals for muons captured in the two relevant nuclei.

APPENDIX F

PERTURBATIVITY BOUNDS

Following the general discussion in [232], we report here the perturbativity limits on relevant coupling combinations for the models we study in Sec. 4.3. For the Π_7 leptoquark, we follow the discussion in [233], which gives

$$\text{Re}[y_{\Pi_7}^{ul} y_{\Pi_7}^{qe*}] \lesssim \frac{8\pi}{\sqrt{3}} \simeq 14.5 \quad \lambda_{\Pi_7} \lesssim \frac{4\pi}{\sqrt{3}}. \quad (\text{F.1})$$

Similarly, the case of the ω_1 leptoquark gives

$$\text{Re}[y_{\omega_1}^{eu} y_{\omega_1}^{ql*}] \lesssim \frac{8\pi}{\sqrt{3}} \quad \lambda_{\omega_1} \lesssim \frac{8\pi}{\sqrt{6}}. \quad (\text{F.2})$$

For the φ couplings, partial-wave unitarity yields

$$\lambda_\varphi \lesssim 4\pi \quad y_e^\varphi \lesssim \sqrt{8\pi}. \quad (\text{F.3})$$

APPENDIX G

MORE DETAILS ON SCALAR LEPTOQUARKS

G.1 Matching expressions for scalar leptoquarks

We provide here the expressions for the one loop matching of ω_1 and Π_7 onto relevant SMEFT coefficients,

$$\begin{aligned}
 [\mathcal{C}_{eH}]_{ij} = & \frac{3}{16\pi^2} \left[\frac{(y_{\omega_1}^{eu})_{jp}(y_{\omega_1}^{ql})_{ri}^*}{M_{\omega_1}^2} [Y_u^\dagger Y_u Y_u^\dagger]_{rp} \left(1 + \ln \left(\frac{\mu^2}{M_{\omega_1}^2} \right) \right) \right. \\
 & - \frac{(y_{\omega_1}^{eu})_{jp}(y_{\omega_1}^{ql})_{ri}^*}{2M_{\omega_1}^2} [Y_u^\dagger]_{rp} \left(2\lambda_{\omega_1} + \lambda_\phi \left(3 + 2 \ln \left(\frac{\mu^2}{M_{\omega_1}^2} \right) \right) \right) \\
 & + \frac{(y_{\Pi_7}^{eq})_{jp}^*(y_{\Pi_7}^{lu})_{ir}}{M_{\Pi_7}^2} [Y_u^\dagger Y_u Y_u^\dagger]_{pr} \left(1 + \ln \left(\frac{\mu^2}{M_{\Pi_7}^2} \right) \right) \\
 & \left. - \frac{(y_{\Pi_7}^{eq})_{jp}^*(y_{\Pi_7}^{lu})_{ir}}{2M_{\Pi_7}^2} [Y_u^\dagger]_{pr} \left(2\lambda_{\Pi_7} + \lambda_\phi \left(3 + 2 \ln \left(\frac{\mu^2}{M_{\Pi_7}^2} \right) \right) \right) \right], \quad (G.1)
 \end{aligned}$$

$$\begin{aligned}
 [\mathcal{C}_{eW}]_{ij} = & -\frac{3g_2}{16\pi^2} \left(\frac{(y_{\omega_1}^{eu})_{jp}(y_{\omega_1}^{ql})_{ri}^*}{16M_{\omega_1}^2} [Y_u^\dagger]_{rp} \left[3 + 2 \ln \left(\frac{\mu^2}{M_{\omega_1}^2} \right) \right] \right. \\
 & \left. - \frac{(y_{\Pi_7}^{eq})_{jp}^*(y_{\Pi_7}^{lu})_{ir}}{16M_{\Pi_7}^2} [Y_u^\dagger]_{pr} \left[1 + 2 \ln \left(\frac{\mu^2}{M_{\Pi_7}^2} \right) \right] \right), \quad (G.2)
 \end{aligned}$$

$$\begin{aligned}
 [\mathcal{C}_{eB}]_{ij} = & \frac{g_1}{16\pi^2} \left(\frac{(y_{\omega_1}^{eu})_{jp}(y_{\omega_1}^{ql})_{ri}^*}{16M_{\omega_1}^2} [Y_u^\dagger]_{rp} \left[19 + 10 \ln \left(\frac{\mu^2}{M_{\omega_1}^2} \right) \right] \right. \\
 & \left. - \frac{(y_{\Pi_7}^{eq})_{jp}^*(y_{\Pi_7}^{lu})_{ir}}{16M_{\Pi_7}^2} [Y_u^\dagger]_{pr} \left[1 + 10 \ln \left(\frac{\mu^2}{M_{\Pi_7}^2} \right) \right] \right), \quad (G.3)
 \end{aligned}$$

$$\mathcal{C}_{HD} = \frac{-g_1^4}{2880\pi^2 M_{\omega_1}^2} - \frac{49g_1^4 + 90(\sqrt{2}\lambda_{\Pi_7} - \eta_{\Pi_7})^2}{5760\pi^2 M_{\Pi_7}^2} \quad (G.4)$$

$$\mathcal{C}_{H\Box} = \frac{(-g_1^4 - 180\lambda_{\omega_1}^2)}{2880\pi^2 M_{\omega_1}^2} - \frac{49g_1^4 + 90(\sqrt{2}\lambda_{\Pi_7} - \eta_{\Pi_7})^2}{5760\pi^2 M_{\Pi_7}^2}, \quad (G.5)$$

where we have neglected pieces proportional to SM lepton Yukawas, and leptoquark couplings are defined as in Eq. (4.41). These expressions were calculated using the **SOLD** [113, 114]

package and are consistent with those in Refs. [234, 90, 235].

G.2 Expressions for $\Delta\kappa_e$ and Δa_e

We give here the expressions for Δa_e and $\Delta\kappa_e$ in terms of the ω_1 and Π_7 model parameters, as defined in Eq. (4.41), at the mass scales $M_{\omega_1/\Pi_7} = 2$ TeV and $M_{\omega_1/\Pi_7} = 10$ TeV. We find

$$\Delta a_e = 10^{-11} \left[180(y_{\omega_1}^{eu})_{13}(y_{\omega_1}^{ql})_{31}^* + (3.0 + 0.35c_T^{(c)})(y_{\omega_1}^{eu})_{12}(y_{\omega_1}^{ql})_{21}^* - 190(y_{\Pi_7}^{eq})_{13}^*(y_{\Pi_7}^{lu})_{13} - (2.9 + 0.34c_T^{(c)})(y_{\Pi_7}^{eq})_{12}^*(y_{\Pi_7}^{lu})_{12} \right] \Big|_{2 \text{ TeV}}, \quad (\text{G.6})$$

$$\Delta a_e = 10^{-11} \left[12(y_{\omega_1}^{eu})_{13}(y_{\omega_1}^{ql})_{31}^* + (0.14 + 0.014c_T^{(c)})(y_{\omega_1}^{eu})_{12}(y_{\omega_1}^{ql})_{21}^* - 12(y_{\Pi_7}^{eq})_{13}^*(y_{\Pi_7}^{lu})_{13} - (0.13 + 0.013c_T^{(c)})(y_{\Pi_7}^{eq})_{12}^*(y_{\Pi_7}^{lu})_{12} \right] \Big|_{10 \text{ TeV}}, \quad (\text{G.7})$$

where $c_T^{(c)}$ is a non-perturbative parameter defined in [90], expected to be of order 1. And

$$\Delta\kappa_e = \left[(290 + 73\lambda_{\omega_1})(y_{\omega_1}^{eu})_{13}(y_{\omega_1}^{ql})_{31}^* - (0.23 - 0.27\lambda_{\omega_1})(y_{\omega_1}^{eu})_{12}(y_{\omega_1}^{ql})_{21}^* + (270 + 73\lambda_{\Pi_7})(y_{\Pi_7}^{eq})_{13}^*(y_{\Pi_7}^{lu})_{13} + (0.036 + 0.27\lambda_{\Pi_7})(y_{\Pi_7}^{eq})_{12}^*(y_{\Pi_7}^{lu})_{12} \right] \Big|_{2 \text{ TeV}}, \quad (\text{G.8})$$

$$\Delta\kappa_e = \left[(18 + 2.6\lambda_{\omega_1})(y_{\omega_1}^{eu})_{13}(y_{\omega_1}^{ql})_{31}^* - (16 - 9.5\lambda_{\omega_1}) \times 10^{-3}(y_{\omega_1}^{eu})_{12}(y_{\omega_1}^{ql})_{21}^* + (16 + 2.6\lambda_{\Pi_7})(y_{\Pi_7}^{eq})_{13}^*(y_{\Pi_7}^{lu})_{13} + (10 + 9.5\lambda_{\Pi_7}) \times 10^{-3}(y_{\Pi_7}^{eq})_{12}^*(y_{\Pi_7}^{lu})_{12} \right] \Big|_{10 \text{ TeV}}. \quad (\text{G.9})$$

G.3 Top FCNCs

Given the coupling structure for both leptoquarks and the sizeable couplings needed in order to accommodate a large deviation of the electron Yukawa, one expects also a contribution to FCNCs involving the top quark. At lepton colliders, top FCNC processes can be probed both in single top production, i.e. $e^+e^- \rightarrow tq$, or in the decays of the tops produced in the $t\bar{t}$ run. In the following, we briefly analyse the prospects coming from both processes, in the cases relevant for us, as well as existing bounds.

$$e^+e^- \rightarrow tq$$

Future prospects for this process have been analysed in [121], where all bounds on the relevant effective operators are reported for a 240 GeV run of a circular e^+e^- collider. The generic constraints one expects on four-fermion operators of the form $[\mathcal{C}]_{11i3}$, $i = 1, 2$, are of the order of 0.05 TeV^{-2} . Setting a mass for the leptoquarks of $M = 10 \text{ TeV}$, one finds that, in both cases

$$[y_{\Pi_7}^{ul}]_{21}[y_{\Pi_7}^{qe*}]_{31} < 14, \quad [y_{\omega_1}^{eu}]_{12}[y_{\omega_1}^{ql*}]_{31} < 14. \quad (\text{G.10})$$

For $M = 2 \text{ TeV}$ instead one finds

$$[y_{\Pi_7}^{ul}]_{21}[y_{\Pi_7}^{qe*}]_{31} < 0.6, \quad [y_{\omega_1}^{eu}]_{12}[y_{\omega_1}^{ql*}]_{31} < 0.6. \quad (\text{G.11})$$

$$t \rightarrow cX$$

The decays to a charm quark and a Higgs or photon necessarily go through chirally-flipping structures ($\mathcal{C}_{uH}$ or dipole operators), which are suppressed by the electron Yukawa in our setup. The decay $t \rightarrow cZ$, on the other hand, can be mediated by vector structures proportional to the same coupling squared (with different flavour indices), e.g. $[y_{\omega_1}^{eu}]_{13}[y_{\omega_1}^{eu}]_{12}^*$. In SMEFT, the branching ratio can be written as [236]:

$$\mathcal{B}(t \rightarrow cZ) = \frac{1}{\Gamma_t} \frac{m_t^3 v^2}{32\pi} \left(1 - \frac{m_Z^2}{m_t^2}\right)^2 \left(1 + \frac{2m_Z^2}{m_t^2}\right) \left(|[\mathcal{C}_{Hq}^{(1-3)}]_{23}|^2 + |[\mathcal{C}_{Hu}]_{23}|^2\right), \quad (\text{G.12})$$

where $\mathcal{C}_{Hq}^{(1-3)} \equiv \mathcal{C}_{Hq}^{(1)} - \mathcal{C}_{Hq}^{(3)}$. Keeping only the log term in the coefficients, one finds

$$[\mathcal{C}_{Hq}^{(1-3)}]_{23}(m_t) = \frac{1}{16\pi^2 M_{\omega_1}^2} \log \frac{m_t}{M_{\omega_1}} \frac{g_2^2 - g_1^2}{6} [y_{\omega_1}^{ql}]_{21} [y_{\omega_1}^{ql}]_{31}^* \quad (\text{G.13})$$

$$[\mathcal{C}_{Hu}]_{23}(m_t) = -\frac{1}{16\pi^2 M_{\omega_1}^2} \log \frac{m_t}{M_{\omega_1}} \frac{g_1^2}{3} [y_{\omega_1}^{eu}]_{13} [y_{\omega_1}^{eu}]_{12}^* \quad (\text{G.14})$$

for the case of the ω_1 , and

$$[\mathcal{C}_{Hq}^{(1-3)}]_{23}(m_t) = \frac{1}{16\pi^2 M_{\Pi_7}^2} \log \frac{m_t}{M_{\Pi_7}} \frac{g_1^2}{3} [y_{\Pi_7}^{qe}]_{21} [y_{\Pi_7}^{qe}]_{31}^* \quad (\text{G.15})$$

$$[\mathcal{C}_{Hu}]_{23}(m_t) = \frac{1}{16\pi^2 M_{\Pi_7}^2} \log \frac{m_t}{M_{\Pi_7}} \frac{g_1^2}{3} [y_{\Pi_7}^{ul}]_{31} [y_{\Pi_7}^{ul}]_{21}^* \quad (\text{G.16})$$

for Π_7 . Given the current experimental constraint of $\mathcal{B}(t \rightarrow cZ) < 1.2 \times 10^{-4}$ (95% C.L.) [237], this doesn't give any meaningful constraint on the couplings. The FCC-ee projection for $\mathcal{B}(t \rightarrow qZ)$ is at 8×10^{-6} (95% C.L.) [238], yielding the still very weak bound

$$y_{12}y_{13} \lesssim 6 \times 10^4, \quad (\text{G.17})$$

for a mass $M = 10$ TeV.

REFERENCES

- [1] Wolfgang Altmannshofer, Benoît Assi, Joachim Brod, Nick Hamer, J. Julio, Patipan Uttayarat, and Daniil Volkov. Electron EDM and $\Gamma(\mu \rightarrow e\gamma)$ in the 2HDM. 10 2024.
- [2] Aram Hayrapetyan et al. Search for physics beyond the standard model in top quark production with additional leptons in the context of effective field theory. *JHEP*, 12:068, 2023.
- [3] Lukas Allwicher, Darius A. Faroughy, Florentin Jaffredo, Olcyr Sumensari, and Felix Wilsch. Drell-Yan tails beyond the Standard Model. *JHEP*, 03:064, 2023.
- [4] Admir Greljo, Hector Tiblom, and Alessandro Valenti. New physics through flavor tagging at FCC-ee. *SciPost Phys.*, 18(5):152, 2025.
- [5] Jaco ter Hoeve, Luca Mantani, Juan Rojo, Alejo N. Rossia, and Eleni Vryonidou. Higgs trilinear coupling in the standard model effective field theory at the high luminosity LHC and the FCC-ee. *Phys. Rev. D*, 112(1):013008, 2025.
- [6] M. Benedikt et al. Future Circular Collider Feasibility Study Report: Volume 1, Physics, Experiments, Detectors. 4 2025.
- [7] de Blas et al. Physics Briefing Book: Input for the 2026 update of the European Strategy for Particle Physics. 2025.
- [8] Georges Aad et al. A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery. *Nature*, 607(7917):52–59, 2022. [Erratum: *Nature* 612, E24 (2022)].
- [9] J. de Blas et al. Higgs Boson Studies at Future Particle Colliders. *JHEP*, 01:139, 2020.
- [10] A. A. Aguilar-Arevalo et al. The MiniBooNE Detector. *Nucl. Instrum. Meth. A*, 599:28–46, 2009.
- [11] A. A. Aguilar-Arevalo et al. The Neutrino Flux Prediction at MiniBooNE. *Phys. Rev. D*, 79:072002, 2009.
- [12] Alexis A. Aguilar-Arevalo et al. Measurement of ν_μ and $\bar{\nu}_\mu$ induced neutral current single π^0 production cross sections on mineral oil at $E_\nu \sim \mathcal{O}(1\text{GeV})$. *Phys. Rev. D*, 81:013005, 2010.
- [13] P. Abratenko et al. Measurement of neutral current single π^0 production on argon with the MicroBooNE detector. *Phys. Rev. D*, 107(1):012004, 2023.
- [14] Rodolfo Capdevilla, David Curtin, Yonatan Kahn, and Gordan Krnjaic. Systematically testing singlet models for $g - 2$. *JHEP*, 04:129, 2022.
- [15] Chien-Yi Chen, Maxim Pospelov, and Yi-Ming Zhong. Muon Beam Experiments to Probe the Dark Sector. *Phys. Rev. D*, 95(11):115005, 2017.
- [16] Yonatan Kahn, Gordan Krnjaic, Nhan Tran, and Andrew Whitbeck. M^3 : a new muon missing momentum experiment to probe $g-2$ and dark matter at Fermilab. *JHEP*, 09:153, 2018.
- [17] Gordan Krnjaic and Samuel D. McDermott. Implications of BBN Bounds for Cosmic Ray Upscattered Dark Matter. *Phys. Rev. D*, 101(12):123022, 2020.
- [18] Gordan Krnjaic, Duncan Rocha, and Isaac R. Wang. Discovering Dark Matter with the MUonE Experiment. *Phys. Rev. Lett.*, 134(16):161801, 2025.
- [19] J. de Blas, J. C. Criado, M. Perez-Victoria, and J. Santiago. Effective description of

- general extensions of the Standard Model: the complete tree-level dictionary. JHEP, 03:109, 2018.
- [20] Nathaniel Craig, Minyuan Jiang, Ying-Ying Li, and Dave Sutherland. Loops and Trees in Generic EFTs. JHEP, 08:086, 2020.
 - [21] Christophe Grojean, Guilherme Guedes, Jasper Roosmale Nepveu, and Gabriel M. Salla. A log story short: running contributions to radiative Higgs decays in the SMEFT. JHEP, 12:065, 2024.
 - [22] B. Grzadkowski, M. Iskrzynski, M. Misiak, and J. Rosiek. Dimension-Six Terms in the Standard Model Lagrangian. JHEP, 10:085, 2010.
 - [23] Antonio Pich. Effective field theory: Course. pages 949–1049, 6 1998.
 - [24] G. C. Branco, P. M. Ferreira, L. Lavoura, M. N. Rebelo, Marc Sher, and Joao P. Silva. Theory and phenomenology of two-Higgs-doublet models. Phys. Rept., 516:1–102, 2012.
 - [25] John F. Gunion and Howard E. Haber. The CP conserving two Higgs doublet model: The Approach to the decoupling limit. Phys. Rev. D, 67:075019, 2003.
 - [26] Howard Georgi and D.V. Nanopoulos. Suppression of flavor changing effects from neutral spinless meson exchange in gauge theories. Physics Letters B, 82(1):95–96, 1979.
 - [27] John F. Donoghue and Ling-Fong Li. Properties of charged higgs bosons. Phys. Rev. D, 19:945–955, Feb 1979.
 - [28] L. Lavoura and Joao P. Silva. Fundamental CP violating quantities in a $SU(2) \times U(1)$ model with many Higgs doublets. Phys. Rev. D, 50:4619–4624, 1994.
 - [29] L. Lavoura. Signatures of discrete symmetries in the scalar sector. Phys. Rev. D, 50:7089–7092, 1994.
 - [30] F. J. Botella and Joao P. Silva. Jarlskog - like invariants for theories with scalars and fermions. Phys. Rev. D, 51:3870–3875, 1995.
 - [31] Ralf Hempfling. Yukawa coupling unification with supersymmetric threshold corrections. Phys. Rev. D, 49:6168–6172, 1994.
 - [32] Lawrence J. Hall, Riccardo Rattazzi, and Uri Sarid. The Top quark mass in supersymmetric $SO(10)$ unification. Phys. Rev. D, 50:7048–7065, 1994.
 - [33] Marcela Carena, M. Olechowski, S. Pokorski, and C. E. M. Wagner. Electroweak symmetry breaking and bottom - top Yukawa unification. Nucl. Phys. B, 426:269–300, 1994.
 - [34] Andrzej J. Buras, Piotr H. Chankowski, Janusz Rosiek, and Lucja Slawianowska. $\Delta M_{d,s}, B^0 d, s \rightarrow \mu^+ \mu^-$ and $B \rightarrow X_s \gamma$ in supersymmetry at large $\tan \beta$. Nucl. Phys. B, 659:3, 2003.
 - [35] Athanasios Dedes and Apostolos Pilaftsis. Resummed Effective Lagrangian for Higgs Mediated FCNC Interactions in the CP Violating MSSM. Phys. Rev. D, 67:015012, 2003.
 - [36] Marcela Carena, A. Menon, R. Noriega-Papaqui, A. Szyrkman, and C. E. M. Wagner. Constraints on B and Higgs physics in minimal low energy supersymmetric models. Phys. Rev. D, 74:015009, 2006.
 - [37] Apostolos Pilaftsis and Carlos E. M. Wagner. Higgs bosons in the minimal supersym-

- metric standard model with explicit CP violation. *Nucl. Phys. B*, 553:3–42, 1999.
- [38] John F. Gunion and Howard E. Haber. Conditions for CP-violation in the general two-Higgs-doublet model. *Phys. Rev. D*, 72:095002, 2005.
 - [39] Henning Bahl, Marcela Carena, Nina M. Coyle, Aurora Ireland, and Carlos E. M. Wagner. New tools for dissecting the general 2HDM. *JHEP*, 03:165, 2023.
 - [40] I. F. Ginzburg and I. P. Ivanov. Tree-level unitarity constraints in the most general 2HDM. *Phys. Rev. D*, 72:115010, 2005.
 - [41] H. Fritzsch. Weak-interaction mixing in the six-quark theory. *Physics Letters B*, 73(3):317–322, 1978.
 - [42] T. P. Cheng and Marc Sher. Mass Matrix Ansatz and Flavor Nonconservation in Models with Multiple Higgs Doublets. *Phys. Rev. D*, 35:3484, 1987.
 - [43] Wick Haxton, Kenneth McElvain, Tony Menzo, Evan Rule, and Jure Zupan. Effective theory tower for $\mu \rightarrow e$ conversion. *JHEP*, 11:076, 2024.
 - [44] Roni Harnik, Joachim Kopp, and Jure Zupan. Flavor Violating Higgs Decays. *JHEP*, 03:026, 2013.
 - [45] Tanya S. Roussy et al. An improved bound on the electron’s electric dipole moment. *Science*, 381(6653):adg4084, 2023.
 - [46] Wilhelm H. Bertl et al. A Search for muon to electron conversion in muonic gold. *Eur. Phys. J. C*, 47:337–346, 2006.
 - [47] K. Afanaciev et al. New limit on the $\mu^+ \rightarrow e^+ \gamma$ decay with the MEG II experiment. *Eur. Phys. J. C*, 85(10):1177, 2025. [Erratum: *Eur.Phys.J.C* 85, 1317 (2025)].
 - [48] F. Abdi et al. Mu2e Run I Sensitivity Projections for the Neutrinoless Conversion Search in Aluminum. *Universe*, 9(1):54, 2023.
 - [49] K. Afanaciev et al. A search for $\mu^+ \rightarrow e^+ \gamma$ with the first dataset of the MEG II experiment. *Eur. Phys. J. C*, 84(3):216, 2024. [Erratum: *Eur.Phys.J.C* 84, 1042 (2024)].
 - [50] Georges Aad et al. Searches for lepton-flavour-violating decays of the Higgs boson into $e\tau$ and $\mu\tau$ in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. *JHEP*, 07:166, 2023.
 - [51] Serguei Chatrchyan et al. Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC. *Phys. Lett. B*, 716:30–61, 2012.
 - [52] Georges Aad et al. Search for the Higgs boson decays $H \rightarrow ee$ and $H \rightarrow e\mu$ in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Lett. B*, 801:135148, 2020.
 - [53] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110(3):030001, 2024.
 - [54] Georges Aad et al. Highlights of the HL-LHC physics projections by ATLAS and CMS. 4 2025.
 - [55] Georges Aad et al. Search for heavy neutral Higgs bosons decaying into a top quark pair in 140 fb $^{-1}$ of proton-proton collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector. *JHEP*, 08:013, 2024.
 - [56] David Eriksson, Johan Rathsmann, and Oscar Stal. 2HDMC: Two-Higgs-Doublet Model Calculator Physics and Manual. *Comput. Phys. Commun.*, 181:189–205, 2010.
 - [57] Robert V. Harlander, Stefan Liebler, and Hendrik Mantler. SusHi Bento: Beyond NNLO and the heavy-top limit. *Comput. Phys. Commun.*, 212:239–257, 2017.

- [58] Robert V. Harlander, Stefan Liebler, and Hendrik Mantler. SusHi: A program for the calculation of Higgs production in gluon fusion and bottom-quark annihilation in the Standard Model and the MSSM. *Comput. Phys. Commun.*, 184:1605–1617, 2013.
- [59] Tao Han, Shuailong Li, Shufang Su, Wei Su, and Yongcheng Wu. Heavy Higgs bosons in 2HDM at a muon collider. *Phys. Rev. D*, 104(5):055029, 2021.
- [60] Hanfei Ye. Search for heavy Higgs bosons decaying into two tau leptons with the ATLAS detector using pp collisions at $\sqrt{s} = 13$ TeV. *PoS, EPS-HEP2021:629*, 2022.
- [61] Georges Aad et al. Search for heavy resonances decaying into a Z or W boson and a Higgs boson in final states with leptons and b -jets in 139 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *JHEP*, 06:016, 2023.
- [62] Georges Aad et al. Search for heavy resonances decaying into a pair of Z bosons in the $\ell^+ \ell^- \ell'^+ \ell'^-$ and $\ell^+ \ell^- \nu \bar{\nu}$ final states using 139 fb^{-1} of proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Eur. Phys. J. C*, 81(4):332, 2021.
- [63] Marcela Carena, Ian Low, Nausheen R. Shah, and Carlos E. M. Wagner. Impersonating the Standard Model Higgs Boson: Alignment without Decoupling. *JHEP*, 04:015, 2014.
- [64] Lukas Allwicher, Matthew McCullough, Sophie Renner, Duncan Rocha, and Benjamin Smith. The price of a large electron Yukawa modification. *JHEP*, 03:201, 2026.
- [65] Joachim Brod, Jonathan M. Cornell, Dimitrios Skodras, and Emmanuel Stamou. Global constraints on Yukawa operators in the standard model effective theory. *JHEP*, 08:294, 2022.
- [66] Giuliano Panico, Alex Pomarol, and Marc Riembau. EFT approach to the electron Electric Dipole Moment at the two-loop level. *JHEP*, 04:090, 2019.
- [67] Clifford Cheung and Chia-Hsien Shen. Nonrenormalization Theorems without Supersymmetry. *Phys. Rev. Lett.*, 115(7):071601, 2015.
- [68] Joan Elias Miró, James Ingoldby, and Marc Riembau. EFT anomalous dimensions from the S-matrix. *JHEP*, 09:163, 2020.
- [69] Juan Manuel Dávila, Anirban Karan, Emilie Passemar, Antonio Pich, and Luiz Vale Silva. The Electric Dipole Moment of the electron in the decoupling limit of the aligned Two-Higgs Doublet Model. 4 2025.
- [70] Elizabeth E. Jenkins, Aneesh V. Manohar, and Peter Stoffer. Low-Energy Effective Field Theory below the Electroweak Scale: Operators and Matching. *JHEP*, 03:016, 2018. [Erratum: *JHEP* 12, 043 (2023)].
- [71] J. D. Bjorken and Steven Weinberg. A Mechanism for Nonconservation of Muon Number. *Phys. Rev. Lett.*, 38:622, 1977.
- [72] Stephen M. Barr and A. Zee. Electric Dipole Moment of the Electron and of the Neutron. *Phys. Rev. Lett.*, 65:21–24, 1990. [Erratum: *Phys.Rev.Lett.* 65, 2920 (1990)].
- [73] S. Heinemeyer, D. Stockinger, and G. Weiglein. Two loop SUSY corrections to the anomalous magnetic moment of the muon. *Nucl. Phys. B*, 690:62–80, 2004.
- [74] Sven Heinemeyer, Dominik Stockinger, and Georg Weiglein. Electroweak and supersymmetric two-loop corrections to $(g-2)(\mu)$. *Nucl. Phys. B*, 699:103–123, 2004.
- [75] Dominik Stockinger. The Muon Magnetic Moment and Supersymmetry. *J. Phys. G*, 34:R45–R92, 2007.
- [76] King-man Cheung, Chung-Hsien Chou, and Otto C. W. Kong. Muon anomalous mag-

- netic moment, two Higgs doublet model, and supersymmetry. *Phys. Rev. D*, 64:111301, 2001.
- [77] Javier Fuentes-Martin, Pedro Ruiz-Femenia, Avelino Vicente, and Javier Virto. DsixTools 2.0: The Effective Field Theory Toolkit. *Eur. Phys. J. C*, 81(2):167, 2021.
 - [78] Richard H. Parker, Chenghui Yu, Weicheng Zhong, Brian Estey, and Holger Müller. Measurement of the fine-structure constant as a test of the Standard Model. *Science*, 360:191, 2018.
 - [79] Léo Morel, Zhibin Yao, Pierre Cladé, and Saïda Guellati-Khélifa. Determination of the fine-structure constant with an accuracy of 81 parts per trillion. *Nature*, 588(7836):61–65, 2020.
 - [80] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse. Measurement of the Electron Magnetic Moment. *Phys. Rev. Lett.*, 130(7):071801, 2023.
 - [81] Luca Di Luzio, Alexander Keshavarzi, Antonio Masiero, and Paride Paradisi. Model-Independent Tests of the Hadronic Vacuum Polarization Contribution to the Muon $g-2$. *Phys. Rev. Lett.*, 134(1):011902, 2025.
 - [82] Tatsumi Aoyama, Toichiro Kinoshita, and Makiko Nio. Revised and Improved Value of the QED Tenth-Order Electron Anomalous Magnetic Moment. *Phys. Rev. D*, 97(3):036001, 2018.
 - [83] Sergey Volkov. Calculating the five-loop QED contribution to the electron anomalous magnetic moment: Graphs without lepton loops. *Phys. Rev. D*, 100(9):096004, 2019.
 - [84] Sergey Volkov. Calculation of the total 10th order QED contribution to the electron magnetic moment. *Phys. Rev. D*, 110(3):036001, 2024.
 - [85] Barbara Anna Erdelyi, Ramona Gröber, and Nudzeim Selimovic. Probing new physics with the electron Yukawa coupling. *JHEP*, 05:135, 1 2025.
 - [86] Elizabeth E. Jenkins, Aneesh V. Manohar, and Michael Trott. Renormalization Group Evolution of the Standard Model Dimension Six Operators I: Formalism and lambda Dependence. *JHEP*, 10:087, 2013.
 - [87] Elizabeth E. Jenkins, Aneesh V. Manohar, and Michael Trott. Renormalization Group Evolution of the Standard Model Dimension Six Operators II: Yukawa Dependence. *JHEP*, 01:035, 2014.
 - [88] Rodrigo Alonso, Elizabeth E. Jenkins, Aneesh V. Manohar, and Michael Trott. Renormalization Group Evolution of the Standard Model Dimension Six Operators III: Gauge Coupling Dependence and Phenomenology. *JHEP*, 04:159, 2014.
 - [89] Alejandro Celis, Javier Fuentes-Martin, Avelino Vicente, and Javier Virto. DsixTools: The Standard Model Effective Field Theory Toolkit. *Eur. Phys. J. C*, 77(6):405, 2017.
 - [90] Jason Aebischer, Wouter Dekens, Elizabeth E. Jenkins, Aneesh V. Manohar, Dipan Sengupta, and Peter Stoffer. Effective field theory interpretation of lepton magnetic and electric dipole moments. *JHEP*, 07:107, 2021.
 - [91] Luca Naterop and Peter Stoffer. Renormalization-group equations of the LEFT at two loops: dimension-six operators. 7 2025.
 - [92] Roel Aaij et al. Search for the Rare Decays $B_s^0 \rightarrow e^+e^-$ and $B^0 \rightarrow e^+e^-$. *Phys. Rev. Lett.*, 124(21):211802, 2020.
 - [93] A. Zupanc et al. Measurements of branching fractions of leptonic and hadronic D_s^+

- meson decays and extraction of the D_s^+ meson decay constant. *JHEP*, 09:139, 2013.
- [94] Searches of lepton-flavour-violating decays of the Higgs bosons with the ATLAS detector at the HL-LHC. 2022. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2022-054>.
 - [95] Qin Qin, Qiang Li, Cai-Dian Lü, Fu-Sheng Yu, and Si-Hong Zhou. Charged lepton flavor violating Higgs decays at future e^+e^- colliders. *Eur. Phys. J. C*, 78(10):835, 2018.
 - [96] Niklaus Berger. The Mu3e Experiment. *Nucl. Phys. B Proc. Suppl.*, 248-250:35–40, 2014. 1st Conference on Charged Lepton Flavor Violation.
 - [97] Manabu Moritsu. Search for Muon-to-Electron Conversion with the COMET Experiment †. *Universe*, 8(4):196, 2022.
 - [98] Daniel Egana-Ugrinovic, Samuel Homiller, and Patrick Meade. Aligned and Spontaneous Flavor Violation. *Phys. Rev. Lett.*, 123(3):031802, 2019.
 - [99] Armen Tumasyan et al. Search for the Higgs boson decay to a pair of electrons in proton-proton collisions at $\sqrt{s}=13\text{TeV}$. *Phys. Lett. B*, 846:137783, 2023.
 - [100] M. Cepeda et al. Report from Working Group 2: Higgs Physics at the HL-LHC and HE-LHC. *CERN Yellow Rep. Monogr.*, 7:221–584, 2019.
 - [101] Hooman Davoudiasl and Pier Paolo Giardino. Electron $g-2$ foreshadowing discoveries at FCC-ee. *Phys. Rev. D*, 109(7):075037, 2024.
 - [102] Avital Dery, Claudia Frugiuele, and Yosef Nir. Large Higgs-electron Yukawa coupling in 2HDM. *JHEP*, 04:044, 2018.
 - [103] S. Schael et al. Electroweak Measurements in Electron-Positron Collisions at W-Boson-Pair Energies at LEP. *Phys. Rept.*, 532:119–244, 2013.
 - [104] Precision Electroweak Measurements and Constraints on the Standard Model. 12 2010.
 - [105] Armen Tumasyan et al. A portrait of the Higgs boson by the CMS experiment ten years after the discovery. *Nature*, 607(7917):60–68, 2022. [Erratum: *Nature* 623, (2023)].
 - [106] Matthew McCullough. An Indirect Model-Dependent Probe of the Higgs Self-Coupling. *Phys. Rev. D*, 90(1):015001, 2014. [Erratum: *Phys.Rev.D* 92, 039903 (2015)].
 - [107] Carla Biggio, Marzia Bordone, Luca Di Luzio, and Giovanni Ridolfi. Massive vectors and loop observables: the $g - 2$ case. *JHEP*, 10:002, 2016.
 - [108] Farinaldo S. Queiroz and William Shepherd. New Physics Contributions to the Muon Anomalous Magnetic Moment: A Numerical Code. *Phys. Rev. D*, 89(9):095024, 2014.
 - [109] Wolfgang Altmannshofer, Stefania Gori, Hiren H. Patel, Stefano Profumo, and Douglas Tuckler. Electric dipole moments in a leptoquark scenario for the B -physics anomalies. *JHEP*, 05:069, 2020.
 - [110] Andreas Crivellin, Dario Mueller, and Francesco Saturnino. Correlating $h \rightarrow \mu + \mu^-$ to the Anomalous Magnetic Moment of the Muon via Leptoquarks. *Phys. Rev. Lett.*, 127(2):021801, 2021.
 - [111] Andreas Crivellin, Christoph Greub, Dario Müller, and Francesco Saturnino. Scalar Leptoquarks in Leptonic Processes. *JHEP*, 02:182, 2021.
 - [112] Svjetlana Fajfer, Jernej F. Kamenik, and Michele Tamaro. Interplay of New Physics effects in $(g - 2)_\ell$ and $h \rightarrow \ell^+ \ell^-$ — lessons from SMEFT. *JHEP*, 06:099, 2021.

- [113] Guilherme Guedes, Pablo Olgoso, and José Santiago. Towards the one loop IR/UV dictionary in the SMEFT: One loop generated operators from new scalars and fermions. *SciPost Phys.*, 15(4):143, 2023.
- [114] Guilherme Guedes and Pablo Olgoso. From the EFT to the UV: the complete SMEFT one-loop dictionary. 12 2024.
- [115] Georges Aad et al. Combination of searches for pair-produced leptoquarks at $\sqrt{s}=13$ TeV with the ATLAS detector. *Phys. Lett. B*, 854:138736, 2024.
- [116] Georges Aad et al. Search for leptoquark pair production decaying into $te^-\bar{t}e^+$ or $t\mu^-\bar{t}\mu^+$ in multi-lepton final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Eur. Phys. J. C*, 84(8):818, 2024.
- [117] Armen Tumasyan et al. Inclusive nonresonant multilepton probes of new phenomena at $\sqrt{s}=13$ TeV. *Phys. Rev. D*, 105(11):112007, 2022.
- [118] Georges Aad et al. Search for pairs of scalar leptoquarks decaying into quarks and electrons or muons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. *JHEP*, 10:112, 2020.
- [119] Innes Bigaran and Raymond R. Volkas. Getting chirality right: Single scalar leptoquark solutions to the $(g-2)_{e,\mu}$ puzzle. *Phys. Rev. D*, 102(7):075037, 2020.
- [120] Ilja Doršner, Svjetlana Fajfer, and Shaikh Saad. $\mu \rightarrow e\gamma$ selecting scalar leptoquark solutions for the $(g-2)_{e,\mu}$ puzzles. *Phys. Rev. D*, 102(7):075007, 2020.
- [121] Liaoshan Shi and Cen Zhang. Probing the top quark flavor-changing couplings at CEPC. *Chin. Phys. C*, 43(11):113104, 2019.
- [122] Gian F. Giudice and Oleg Lebedev. Higgs-dependent Yukawa couplings. *Phys. Lett. B*, 665:79–85, 2008.
- [123] Cari Cesarotti, Yonatan Kahn, Gordan Krnjaic, Duncan Rocha, and Joshua Spitz. New μ forces from $\nu\mu$ sources. *Phys. Rev. D*, 110(5):055032, 2024.
- [124] Mohammad Sajjad Athar et al. Status and perspectives of neutrino physics. *Prog. Part. Nucl. Phys.*, 124:103947, 2022.
- [125] Nina Mireille Coyle. *Hunting for New Physics in the Higgs and Neutrino Sectors*. PhD thesis, Chicago U., 2023.
- [126] Brian Batell, Ayres Freitas, Ahmed Ismail, and David McKeen. Flavor-specific scalar mediators. *Phys. Rev. D*, 98(5):055026, 2018.
- [127] Brian Batell, Ayres Freitas, Ahmed Ismail, David McKeen, and Mudit Rai. Renormalizable models of flavor-specific scalars. *Phys. Rev. D*, 104(11):115032, 2021.
- [128] Daniel Egana-Ugrinovic, Samuel Homiller, and Patrick Meade. Light Scalars and the Koto Anomaly. *Phys. Rev. Lett.*, 124(19):191801, 2020.
- [129] Miguel Escudero, Dan Hooper, Gordan Krnjaic, and Mathias Pierre. Cosmology with A Very Light $L_\mu - L_\tau$ Gauge Boson. *JHEP*, 03:071, 2019.
- [130] S. Agostinelli et al. GEANT4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003.
- [131] Gordan Krnjaic, Gustavo Marques-Tavares, Diego Redigolo, and Kohsaku Tobioka. Probing Muonphilic Force Carriers and Dark Matter at Kaon Factories. *Phys. Rev. Lett.*, 124(4):041802, 2020.
- [132] Adam Alloul, Neil D. Christensen, Céline Degrande, Claude Duhr, and Benjamin

- Fuks. FeynRules 2.0 - A complete toolbox for tree-level phenomenology. Comput. Phys. Commun., 185:2250–2300, 2014.
- [133] Alexander Belyaev, Neil D. Christensen, and Alexander Pukhov. CalcHEP 3.4 for collider physics within and beyond the Standard Model. Comput. Phys. Commun., 184:1729–1769, 2013.
 - [134] Diana Forbes, Christian Herwig, Yonatan Kahn, Gordan Krnjaic, Cristina Mantilla Suarez, Nhan Tran, and Andrew Whitbeck. New searches for muonphilic particles at proton beam dump spectrometers. Phys. Rev. D, 107(11):116026, 2023.
 - [135] P. A. Zyla et al. Review of Particle Physics. PTEP, 2020(8):083C01, 2020.
 - [136] Donald E. Groom, Nikolai V. Mokhov, and Sergei I. Striganov. Muon stopping power and range tables 10-MeV to 100-TeV. Atom. Data Nucl. Data Tabl., 78:183–356, 2001.
 - [137] Georgia S. Karagiorgi. Searches for New Physics at MiniBooNE: Sterile Neutrinos and Mixing Freedom. 2010.
 - [138] R. B. Patterson, E. M. Laird, Y. Liu, P. D. Meyers, I. Stancu, and H. A. Tanaka. The Extended-track reconstruction for MiniBooNE. Nucl. Instrum. Meth. A, 608:206–224, 2009.
 - [139] G. Cowan. Statistical Data Analysis. Clarendon Press, 1998.
 - [140] R. Acciarri et al. A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam. 3 2015.
 - [141] D. P. Aguillard et al. Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm. Phys. Rev. Lett., 131(16):161802, 8 2023.
 - [142] G. W. Bennett et al. Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL. Phys. Rev. D, 73:072003, 2006.
 - [143] G. Abbiendi et al. Measuring the leading hadronic contribution to the muon $g-2$ via μe scattering. Eur. Phys. J. C, 77(3):139, 2017.
 - [144] Letter of Intent: the MUonE project. 6 2019.
 - [145] Iftah Galon, David Shih, and Isaac R. Wang. Dark photons and displaced vertices at the MUonE experiment. Phys. Rev. D, 107(9):095003, 2023.
 - [146] Umberto Marconi and Clara Matteuzzi. Private communication.
 - [147] Duncan Rocha and Isaac R. Wang. Filling the gap: hunting for vector bosons at the MUonE experiment with displaced decay signature. JHEP, 03:190, 2026.
 - [148] Bob Holdom. Searching for ϵ Charges and a New $U(1)$. Phys. Lett. B, 178:65–70, 1986.
 - [149] R. E. Marshak and Rabindra N. Mohapatra. Quark - Lepton Symmetry and B-L as the $U(1)$ Generator of the Electroweak Symmetry Group. Phys. Lett. B, 91:222–224, 1980.
 - [150] Rabindra N. Mohapatra and R. E. Marshak. Local B-L Symmetry of Electroweak Interactions, Majorana Neutrinos and Neutron Oscillations. Phys. Rev. Lett., 44:1316–1319, 1980. [Erratum: Phys.Rev.Lett. 44, 1643 (1980)].
 - [151] Robert Foot. New Physics From Electric Charge Quantization? Mod. Phys. Lett. A, 6:527–530, 1991.
 - [152] S. R. Amendolia et al. A Measurement of the Pion Charge Radius. Phys. Lett. B, 146:116–120, 1984.
 - [153] S. R. Amendolia et al. A Measurement of the Space - Like Pion Electromagnetic

- Form-Factor. Nucl. Phys. B, 277:168, 1986.
- [154] Sz. Borsanyi et al. Leading hadronic contribution to the muon magnetic moment from lattice QCD. Nature, 593(7857):51–55, 2021.
 - [155] T. Aoyama et al. The anomalous magnetic moment of the muon in the Standard Model. Phys. Rept., 887:1–166, 2020.
 - [156] A. Boccaletti et al. High precision calculation of the hadronic vacuum polarisation contribution to the muon anomaly. 7 2024.
 - [157] B. Abi et al. Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm. Phys. Rev. Lett., 126(14):141801, 2021.
 - [158] Pulak Banerjee et al. Theory for muon-electron scattering @ 10 ppm: A report of the MUonE theory initiative. Eur. Phys. J. C, 80(6):591, 2020.
 - [159] The Phase-2 Upgrade of the CMS Tracker. Technical Report CERN-LHCC-2017-009, CMS-TDR-014, CERN, Geneva, 2017.
 - [160] Ernesto Migliore. Status of the upgrade project of the CMS Tracker for HL-LHC. Technical Report 10, CERN, Geneva, 2022.
 - [161] R. Walters. Stopping-power & range tables for electrons, protons, and helium ions. NIST Standard Reference Database 124, NIST, 2017.
 - [162] Johan Alwall, Michel Herquet, Fabio Maltoni, Olivier Mattelaer, and Tim Stelzer. MadGraph 5 : Going Beyond. JHEP, 06:128, 2011.
 - [163] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. JHEP, 07:079, 2014.
 - [164] Stefano Frixione, Olivier Mattelaer, Marco Zaro, and Xiaoran Zhao. Lepton collisions in MadGraph5_aMC@NLO. 8 2021.
 - [165] O. Mattelaer and K. Ostrolenk. Speeding up MadGraph5_aMC@NLO. Eur. Phys. J. C, 81(5):435, 2021.
 - [166] Yung-Su Tsai. Pair Production and Bremsstrahlung of Charged Leptons. Rev. Mod. Phys., 46:815, 1974. [Erratum: Rev. Mod. Phys. 49, 421–423 (1977)].
 - [167] James D. Bjorken, Rouven Essig, Philip Schuster, and Natalia Toro. New fixed-target experiments to search for dark gauge forces. Phys. Rev., D80:075018, 2009.
 - [168] C. F. von Weizsacker. Radiation emitted in collisions of very fast electrons. Z. Phys., 88:612–625, 1934.
 - [169] E. J. Williams. Nature of the high-energy particles of penetrating radiation and status of ionization and radiation formulae. Phys. Rev., 45:729–730, 1934.
 - [170] Yu-Sheng Liu and Gerald A. Miller. Validity of the Weizsäcker-Williams approximation and the analysis of beam dump experiments: Production of an axion, a dark photon, or a new axial-vector boson. Phys. Rev. D, 96(1):016004, 2017.
 - [171] Giovanni Grilli di Cortona and Enrico Nardi. Probing light mediators at the MUonE experiment. Phys. Rev. D, 105(11):L111701, 2022.
 - [172] Philip Ilten, Yotam Soreq, Mike Williams, and Wei Xue. Serendipity in dark photon searches. JHEP, 06:004, 2018.
 - [173] Martin Bauer, Patrick Foldenauer, and Joerg Jaeckel. Hunting All the Hidden Photons.

- JHEP, 07:094, 2018.
- [174] Maxim Pospelov. Secluded U(1) below the weak scale. Phys. Rev. D, 80:095002, 2009.
 - [175] Motoi Endo, Koichi Hamaguchi, and Go Mishima. Constraints on Hidden Photon Models from Electron g-2 and Hydrogen Spectroscopy. Phys. Rev., D86:095029, 2012.
 - [176] E. M. Riordan et al. Search for short-lived axions in an electron-beam-dump experiment. Phys. Rev. Lett., 59:755, 1987.
 - [177] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann. Search for neutral metastable penetrating particles produced in the SLAC beam dump. Phys. Rev., D38:3375, 1988.
 - [178] A. Bross, M. Crisler, Stephen H. Pordes, J. Volk, S. Errede, and J. Wrbanek. A search for short-lived particles produced in an electron beam dump. Phys. Rev. Lett., 67:2942–2945, 1991.
 - [179] A. Konaka et al. Search for neutral particles in electron-beam-dump experiment. Phys. Rev. Lett., 57:659, 1986.
 - [180] M. Davier and H. Nguyen Ngoc. An unambiguous search for a light Higgs boson. Phys. Lett., B229:150, 1989.
 - [181] Sarah Andreas, Carsten Niebuhr, and Andreas Ringwald. New Limits on Hidden Photons from Past Electron Beam Dumps. Phys. Rev., D86:095019, 2012.
 - [182] J. Blümlein et al. Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment. Z. Phys., C51:341–350, 1991.
 - [183] J. Blümlein et al. Limits on the mass of light (pseudo)scalar particles from Bethe-Heitler e^+e^- and $\mu^+\mu^-$ pair production in a proton-iron beam dump experiment. Int. J. Mod. Phys., A7:3835–3850, 1992.
 - [184] Johannes Blümlein and Jurgen Brunner. New exclusion limits for dark gauge forces from beam-dump data. Phys. Lett., B701:155–159, 2011.
 - [185] Johannes Blümlein and Jürgen Brunner. New exclusion limits on dark gauge forces from proton bremsstrahlung in beam-dump data. Phys. Lett., B731:320–326, 2014.
 - [186] F. Bergsma et al. Search for Axion Like Particle Production in 400-GeV Proton - Copper Interactions. Phys. Lett., 157B:458–462, 1985.
 - [187] S.N. Gninenko. Constraints on sub-GeV hidden sector gauge bosons from a search for heavy neutrino decays. Phys. Lett., B713:244–248, 2012.
 - [188] P. Astier et al. Search for heavy neutrinos mixing with tau neutrinos. Phys. Lett., B506:27–38, 2001.
 - [189] G. Bernardi et al. Search for neutrino decay. Phys. Lett., B166:479, 1986.
 - [190] S.N. Gninenko. Stringent limits on the $\pi^0 \rightarrow \gamma X, X \rightarrow e^+e^-$ decay from neutrino experiments and constraints on new light gauge bosons. Phys. Rev., D85:055027, 2012.
 - [191] D. Banerjee et al. Search for a Hypothetical 16.7 MeV Gauge Boson and Dark Photons in the NA64 Experiment at CERN. Phys. Rev. Lett., 120(23):231802, 2018.
 - [192] D. Banerjee et al. Improved limits on a hypothetical X(16.7) boson and a dark photon decaying into e^+e^- pairs. Phys. Rev. D, 101(7):071101, 2020.
 - [193] H. Merkel et al. Search at the Mainz Microtron for light massive gauge bosons relevant for the muon g-2 anomaly. Phys. Rev. Lett., 112:221802, 2014.

- [194] S. Abrahamyan et al. Search for a new gauge boson in electron-nucleus fixed-target scattering by the APEX experiment. *Phys. Rev. Lett.*, 107:191804, 2011.
- [195] Bernard Aubert et al. Search for dimuon decays of a light scalar boson in radiative transitions $\Upsilon \rightarrow \gamma A^0$. *Phys. Rev. Lett.*, 103:081803, 2009.
- [196] David Curtin et al. Exotic decays of the 125 GeV Higgs boson. *Phys. Rev.*, D90:075004, 2014.
- [197] J. P. Lees et al. Search for a dark photon in e^+e^- collisions at BaBar. *Phys. Rev. Lett.*, 113:201801, 2014.
- [198] M. Ablikim et al. Dark photon search in the mass range between 1.5 and 3.4 GeV/ c^2 . 2017.
- [199] Roel Aaij et al. Search for dark photons produced in 13 TeV pp collisions. 2017.
- [200] F. Archilli et al. Search for a vector gauge boson in ϕ meson decays with the KLOE detector. *Phys. Lett.*, B706:251–255, 2012.
- [201] A. Anastasi et al. Limit on the production of a new vector boson in $e^+e^- \rightarrow U\gamma$, $U \rightarrow \pi^+\pi^-$ with the KLOE experiment. *Phys. Lett.*, B757:356–361, 2016.
- [202] Henso Abreu et al. Search for dark photons with the FASER detector at the LHC. *Phys. Lett. B*, 848:138378, 2024.
- [203] R. Meijer Drees et al. Search for weakly interacting neutral bosons produced in π^-p interactions at rest and decaying into e^+e^- pairs. *Phys. Rev. Lett.*, 68:3845–3848, 1992.
- [204] D. Babusci et al. Limit on the production of a light vector gauge boson in ϕ meson decays with the KLOE detector. *Phys. Lett.*, B720:111–115, 2013.
- [205] P. Adlarson et al. Search for a dark photon in the $\pi^0 \rightarrow e^+e^-\gamma$ decay. *Phys. Lett.*, B726:187–193, 2013.
- [206] G. Agakishiev et al. Searching for a dark photon with HADES. *Phys. Lett.*, B731:265–271, 2014.
- [207] A. Adare et al. Search for dark photons from neutral meson decays in pp and dAu collisions at $\sqrt{s_{NN}} = 200$ GeV. *Phys. Rev.*, C91:031901, 2015.
- [208] J. R. Batley et al. Search for the dark photon in π^0 decays. *Phys. Lett.*, B746:178–185, 2015.
- [209] P. H. Adrian et al. Search for a dark photon in electroproduced e^+e^- pairs with the Heavy Photon Search experiment at JLab. *Phys. Rev. D*, 98(9):091101, 2018.
- [210] Nathan Baltzell et al. The Heavy Photon Search Experiment. 3 2022.
- [211] Sergei Gninenko. Addendum to the NA64 Proposal: Search for the $A' \rightarrow invisible$ and $X \rightarrow e^+e^-$ decays in 2021. 2018.
- [212] Philip Ilten, Jesse Thaler, Mike Williams, and Wei Xue. Dark photons from charm mesons at LHCb. *Phys. Rev.*, D92:115017, 2015.
- [213] Philip Ilten, Yotam Soreq, Jesse Thaler, Mike Williams, and Wei Xue. Proposed inclusive dark photon search at LHCb. *Phys. Rev. Lett.*, 116:251803, 2016.
- [214] S. Gardner, R. J. Holt, and A. S. Tadepalli. New prospects in fixed target searches for dark forces with the SeaQuest experiment at Fermilab. *Phys. Rev.*, D93:115015, 2016.
- [215] Bertrand Echenard, Rouven Essig, and Yi-Ming Zhong. Projections for dark photon searches at Mu3e. *JHEP*, 01:113, 2015.

- [216] G. Bellini et al. Precision measurement of the 7Be solar neutrino interaction rate in Borexino. *Phys. Rev. Lett.*, 107:141302, 2011.
- [217] Roni Harnik, Joachim Kopp, and Pedro A. N. Machado. Exploring ν Signals in Dark Matter Detectors. *JCAP*, 1207:026, 2012.
- [218] S. Bilmis, I. Turan, T. M. Aliev, M. Deniz, L. Singh, and H. T. Wong. Constraints on Dark Photon from Neutrino-Electron Scattering Experiments. *Phys. Rev. D*, 92(3):033009, 2015.
- [219] Eric D. Carlson. Limits on a new $U(1)$ coupling. *Nucl. Phys.*, B286:378–398, 1987.
- [220] Claudia Frugiuele, Elina Fuchs, Gilad Perez, and Matthias Schlaffer. Constraining New Physics Models with Isotope Shift Spectroscopy. *Phys. Rev.*, D96(1):015011, 2017.
- [221] Mark B. Wise and Yue Zhang. Lepton Flavorful Fifth Force and Depth-dependent Neutrino Matter Interactions. *JHEP*, 06:053, 2018.
- [222] Stefan Antusch and Martin Spinrath. Quark and lepton masses at the GUT scale including SUSY threshold corrections. *Phys. Rev. D*, 78:075020, 2008.
- [223] Andrei I. Davydychev and J. B. Tausk. Two loop selfenergy diagrams with different masses and the momentum expansion. *Nucl. Phys. B*, 397:123–142, 1993.
- [224] Lukas Allwicher, Gino Isidori, and Marko Pesut. Flavored circular collider: cornering New Physics at FCC-ee via flavor-changing processes. *Eur. Phys. J. C*, 85(6):631, 2025.
- [225] Christoph Bobeth, Martin Gorbahn, Thomas Hermann, Mikolaj Misiak, Emmanuel Stamou, and Matthias Steinhauser. $B_{s,d} \rightarrow l^+ l^-$ in the Standard Model with Reduced Theoretical Uncertainty. *Phys. Rev. Lett.*, 112:101801, 2014.
- [226] Martin Beneke, Christoph Bobeth, and Robert Szafron. Power-enhanced leading-logarithmic QED corrections to $B_q \rightarrow \mu^+ \mu^-$. *JHEP*, 10:232, 2019. [Erratum: *JHEP* 11, 099 (2022)].
- [227] Martín González-Alonso and Jorge Martin Camalich. Global Effective-Field-Theory analysis of New-Physics effects in (semi)leptonic kaon decays. *JHEP*, 12:052, 2016.
- [228] N. Carrasco et al. Leptonic decay constants f_K, f_D , and f_{D_s} with $N_f = 2 + 1 + 1$ twisted-mass lattice QCD. *Phys. Rev. D*, 91(5):054507, 2015.
- [229] A. Bazavov et al. B - and D -meson leptonic decay constants from four-flavor lattice QCD. *Phys. Rev. D*, 98(7):074512, 2018.
- [230] Y. Aoki et al. FLAG Review 2024. 11 2024.
- [231] Ryuichiro Kitano, Masafumi Koike, and Yasuhiro Okada. Detailed calculation of lepton flavor violating muon electron conversion rate for various nuclei. *Phys. Rev. D*, 66:096002, 2002. [Erratum: *Phys.Rev.D* 76, 059902 (2007)].
- [232] Lukas Allwicher, Pere Arnau, Daniele Barducci, and Marco Nardecchia. Perturbative unitarity constraints on generic Yukawa interactions. *JHEP*, 10:129, 2021.
- [233] Lukas Allwicher, Luca Di Luzio, Marco Fedele, Federico Mescia, and Marco Nardecchia. What is the scale of new physics behind the muon $g-2$? *Phys. Rev. D*, 104(5):055035, 2021.
- [234] Valerio Gherardi, David Marzocca, and Elena Venturini. Matching scalar leptoquarks to the SMEFT at one loop. *JHEP*, 07:225, 2020. [Erratum: *JHEP* 01, 006 (2021)].
- [235] John Gargalionis, Jeremie Quevillon, Pham Ngoc Hoa Vuong, and Tevong You. Linear Standard Model extensions in the SMEFT at one loop and Tera-Z. *JHEP*, 07:136, 2025.

- [236] Jure Drobnak, Svjetlana Fajfer, and Jernej F. Kamenik. QCD Corrections to Flavor Changing Neutral Coupling Mediated Rare Top Quark Decays. Phys. Rev. D, 82:073016, 2010.
- [237] Georges Aad et al. Search for flavor-changing neutral-current couplings between the top quark and the Z boson with proton-proton collisions at $\sqrt{s}=13$ TeV with the ATLAS detector. Phys. Rev. D, 108(3):032019, 2023.
- [238] Hamzeh Khanpour, Mojtaba Mohammadi Najafabadi, and Seddigheh Tizchang. Top quark FCNC anomalous couplings at the future collider FCC-ee , October 2023.