

In Search of an Invisible Z'

Joe Davighi

Theoretical Physics Department, CERN, 1211 Geneva 23, Switzerland

E-mail: joseph.davighi@cern.ch

ABSTRACT: We consider extending the Standard Model by an anomaly-free and possibly flavour non-universal $U(1)_X$ gauge symmetry, whose breaking gives a Z' boson that does not affect electroweak precision observables at tree-level or at 1-loop. Provided it does not also couple to electrons, such a Z' boson would be largely invisible to an electroweak precision machine like FCC-ee or CEPC. We show that, while this class of Z' models can also evade tests of quark flavour violation, the constraint of anomaly-cancellation implies that valence quarks, muons, and taus are all charged under $U(1)_X$, with the up quark charge being necessarily large. The conclusion holds even if one augments the SM by three right-handed neutrinos to try and absorb anomalies. This means such Z' bosons cannot simultaneously hide below the TeV scale from $pp \rightarrow \ell\ell$ Drell–Yan measurements at the LHC and, even if we entertain esoteric models in which the lepton charges are numerically very small, we cannot escape dijet searches at the LHC. For equitable quark and lepton charges, $pp \rightarrow \ell\ell$ already excludes such a Z' up to $M/g \gtrsim 10$ TeV, with a reach of 20 TeV expected by the end of the High-Luminosity LHC. The dijet bounds currently sit around 5 TeV, while sensitivity up to 10 TeV could be achieved at HL-LHC. We thus find an excellent complementarity between FCC-ee and HL-LHC in covering all anomaly-free Z' bosons up to several TeV.

Contents

1	Introduction	1
2	Electroweak neutrality	3
2.1	Constraints from 1-loop SMEFT running into the Z -pole	4
2.2	Electroweak neutrality beyond the Z -pole	5
2.3	Flavour Neutrality	5
3	Anomaly cancellation	5
3.1	Linear anomaly conditions	6
3.2	Corollary: zero 1-loop RGE with y_t and g_i	7
3.3	Non-linear anomaly conditions	10
4	Classifying Z' models with electroweak neutrality	10
4.1	Valence quarks are charged	10
4.2	At least two leptons are charged	11
4.3	Classifying solutions with one neutral lepton	12
4.4	Including right-handed neutrinos	12
4.5	Compatibility with semi-simple UV completions	14
4.6	Comments regarding quark and lepton flavour structures	15
5	Phenomenology off the Z pole	16
5.1	Benchmark Models	16
5.2	$4\ell: e^+e^- \rightarrow \ell\ell$	17
5.3	$2q2\ell: pp \rightarrow \ell\ell$	18
5.4	$4q: pp \rightarrow jj$	19
5.5	Comment on low-energy probes of LFUV	22
5.6	Comparison of benchmark models	22
6	Conclusion	26

1 Introduction

In Ref. [1] it was shown that very few single-particle extensions of the Standard Model (SM) can escape running into electroweak precision observables (EWPOs) at 1-loop level. The level of precision anticipated at FCC-ee, particularly on its Z -pole run, promises to probe all such states up to scales of several TeV, depending on the gauge and Lorentz representation of the particle.

The exceptions are: (i) a scalar diquark field that is coupled only to top quarks; (ii) a vector $SU(2)_L$ triplet, or W' , whose flavour-violating couplings to light quarks are

precisely tuned against its flavour-conserving couplings; (iii) a vector $SU(3)$ octet coupled only to top quarks; and (iv) a vector singlet, or Z' , with couplings to fermions that satisfy certain linear equations.¹ Of these options, it is hard to imagine the W' with such tuned couplings to the SM fermions emerging from any reasonable UV model, and moreover it is already ruled out to $M/g \gtrsim 1000$ TeV by its flavour-violating impact on neutral meson mixing, making its accidental invisibility on the Z -pole a moot point. The two coloured options, while tuned to escape the 1-loop RGEs, nonetheless exhibit a large-ish running into EWPOs at next-to-leading order that scales with the top Yukawa and the strong gauge coupling [4] and, more importantly, they will already be probed up to several TeV by the LHC thanks to their direct couplings to gluons. For the final case, namely the Z' , there is a wide family of possible solutions, which can arise from a reasonable UV model in which the SM is extended by some gauged $U(1)_X$ symmetry. Since the Z' is a gauge singlet, and its couplings can be flavour conserving, it is unclear *prima facie* whether any of these Z' models can simultaneously hide from other experiments, principally at the LHC.

In this paper, we look to fill the remaining gap, by asking if any of these Z' models that are invisible on the Z -pole can also escape constraints from current experiments and remain viable at a low-scale, let's say at $M/g \lesssim$ TeV, beyond FCC-ee. The answer is a resounding no. We find that the complementary constraints that already exclude such electroweakly-invisible Z' bosons come not from indirect flavour measurements, as anticipated in [1], but from direct searches at high-energy.

To show this, we will assume the Z' is the heavy gauge boson associated to a spontaneously broken $U(1)_X$ gauge symmetry. Then the $U(1)_X$ charges of SM fermions must be chosen such that gauge anomalies cancel. This imposes a system of linear and non-linear constraints which must be solved over the rational numbers. We find that these constraints become extremely restrictive when combined with those conditions that result from requiring Z -pole invisibility. In particular, we show that:²

- Right-handed up and charm quarks are charged under $U(1)_X$
- At least two generations of right-handed down-type quarks are charged under $U(1)_X$
- At least two generations of right-handed leptons are charged under $U(1)_X$
- The up-quark charge is moreover *large*, necessarily satisfying $2|u| \geq \max(|d_i|, |e_i|)$, meaning the Z' coupling to protons is unsuppressed.

¹Colourless scalar fields transforming in the 4-plet of $SU(2)_L$, with hypercharge 1/2 or 3/2, do not contribute to the RGEs for SMEFT operators constrained by the Z -pole, but they *do* however generate finite 1-loop contributions to the C_{HD} operator [2, 3], that shifts the W mass. As shown in [3], these effects can be pushed to higher order (in loops or the $1/\Lambda$ expansion) by considering the custodially symmetric combination of these two scalar fields, which could originate from a scalar transforming in the **16** of custodial $SO(4)$ symmetry. Such a model might leave its largest imprint on Higgs self-coupling measurements.

²In a similar spirit Ref. [5] found that anomaly cancellation when combined with constraints related to explaining the $(g-2)_\mu$ anomaly, such as it was in 2022, ends up being very predictive concerning the space of viable models.

All of these statements hold even if we consider augmenting the SM fermion content by arbitrarily many right-handed neutrinos charged under $U(1)_X$ to try and ‘soak up’ gauge anomalies.

The result is that such a Z' model is, inevitably, strongly constrained already at the LHC, with current constraints already excluding M/g up to 5 TeV. To reach this conclusion, we try to decouple as many further constraints as possible. Firstly, if right-handed electrons are charged, then there are also strong constraints coming from LEP-II $e^+e^- \rightarrow l^+l^-$ cross-section measurements. As highlighted in recent studies [6, 7], the power of the electroweak programme at FCC-ee is not limited to extreme precision on the Z -pole: we infer that the Z' models we consider, if coupled to e_R , would be probed up to 30 TeV at FCC-ee by measurements *off* the Z -peak. By assigning the two non-zero lepton charges to muon and tau, and demanding the electron charge be zero, these constraints from e^+e^- colliders such as FCC-ee disappear. In this case, we provide an analytic parametrization of all the viable $U(1)_X$ charge solutions, which ends up equivalent to solving for Pythagorean triples.

Continuing with our search for an invisible Z' , even if we decouple electrons in addition to the Z -pole running, there remain strong constraints coming from LHC $pp \rightarrow \ell\ell$ Drell–Yan measurements, which are very strong especially for muons in the final state, and given that up and down quarks must have unsuppressed couplings to the Z' in order to not run into the EWPOs. We find current LHC data from $pp \rightarrow \ell\ell$ at high- p_T requires $M/g \gtrsim 10$ TeV, which should improve to reach roughly 20 TeV after the High-Luminosity LHC (HL-LHC). Finally, one can try to weaken also these $pp \rightarrow \ell\ell$ constraints by finding viable charge assignments in which the lepton charges are numerically much smaller than the quark charges. But even in this case, we cannot help running into strong constraints coming from di-jet searches at the LHC, which give very strong constraints within the kinematical reach of the searches which extends up to 5 TeV. We give a crude estimate of the bounds from di-jet resonance searches, which apply only for narrow-width Z' s, using Ref. [8], and also estimate bounds from non-resonant searches performed using angular distributions of dijets [9].

The outline of the paper is as follows. In §2 we recap the conditions for a Z' to not run into the EWPOs from [1]. In §3 we review the conditions on a Z' to arise from an anomaly-free $U(1)_X$ extension of the SM, and highlight an interplay between these conditions and those from not running into the Z -pole. In §4 we systematically explore the space of Z' models satisfying these criteria, deriving the various statements sketched in the Introduction. Then in §5 we sketch the main phenomenological bounds on the models, coming from 4-lepton, semi-leptonic, and 4-quark probes. We compare various benchmark models, and those that *do* run into the EWPOs, in Fig. 1, before concluding in §6.

2 Electroweak neutrality

We assume the Z' comes from a gauged $U(1)_X$ extension of the SM, and we label the charges of the various SM fields by $\{q_i, l_i, u_i, d_i, e_i, H\}$ in what we hope is an obvious notation. For simplicity, we suppose the $U(1)_X$ symmetry is spontaneously broken by a SM singlet scalar

field ϕ , with unit charge under $U(1)_X$, that gets a vev $\langle\phi\rangle = v_X/\sqrt{2}$, resulting in the Z' getting a mass $M_X = g_X v_X$.

2.1 Constraints from 1-loop SMEFT running into the Z -pole

We begin by recapping the linear constraints on the $U(1)_X$ charges that should be satisfied to not run into EWPOs up to one-loop level, which was studied in [1]. The SMEFT operators that effect the EWPOs at tree-level are, in the Warsaw basis,

$$\{[O_{H\ell}^{(1,3)}]_{ii}, [O_{Hq}^{(1,3)}]_{ii}, [O_{He}]_{ii}, [O_{Hu}]_{ii} (i = 1, 2 \text{ only}), [O_{Hd}]_{ii}, O_{HD}, O_{HWB}, [O_{\ell\ell}]_{1221}\}. \quad (2.1)$$

For the Z' gauge boson to not shift EWPOs at tree-level therefore requires the Higgs is neutral, $H = 0$, and that the Z' has no flavour off-diagonal coupling to left-handed muon-electron bilinears. This can be achieved either via a universal coupling to both ($l_1 = l_2$) or by alignment between the electron and muon mass and flavour eigenstates. We see shortly that cancelling also the one-loop running into (2.1) implies the former scenario.

At one-loop and at leading order in the SMEFT Wilson coefficients, many additional SMEFT operators run into the set (2.1). Neglecting all Yukawa couplings other than y_t , and keeping the running with all three SM gauge couplings $g_{1,2,3}$, it is straightforward to identify the conditions for the Z' to not run into the EWPOs, in other words to be ‘electroweakly neutral’ at this order. Recapping from [1]:

- Running with the top Yukawa coupling y_t vanishes iff $q_3 = u_3$. This is the condition for $U(1)_X$ permitting a renormalisable top Yukawa coupling, which we want to enforce anyway;
- Running with the $SU(2)_L$ gauge coupling g_2 vanishes iff all left-handed fermion fields are neutral under $U(1)_X$, *i.e.* $l_i = q_i = 0 \forall i \in \{1, 2, 3\}$. Note this implies $[C_{\ell\ell}]_{1221}$ is not generated at tree-level. Together with the previous condition, this also implies $u_3 = 0$, so the top quark is neutral under $U(1)_X$;
- Lastly, running with the $U(1)_Y$ gauge coupling g_1 vanishes iff the further linear condition on the remaining non-zero charges is satisfied:

$$\sum_i (-e_i + 2u_i - d_i) = 0. \quad (2.2)$$

We discuss the implications of this final condition when we turn to anomaly cancellation.

At this point, we already wish to emphasize that the condition from g_2 running means these Z' models are intrinsically chiral, coupling only to right-handed fermion fields. While anomaly-cancellation is still possible (see §3), this does however immediately imply that full Yukawa coupling matrices are not permitted at the renormalisable level. Rather, only certain elements of the Yukawa matrices are permitted at dimension-4; encouragingly, the top Yukawa coupling $\overline{q_3} H^c u_3$ is always permitted thanks to the top quark being neutral, which was implied by requiring no 1-loop running into the Z -pole with y_t or g_3 . But otherwise, some additional dynamics is needed to populate the full Yukawa matrices, and

one can push this to a higher scale without destabilising the lowest layer of new physics that we assume is associated to the Z' .³

2.2 Electroweak neutrality beyond the Z -pole

As mentioned in the Introduction, the electroweak program of FCC-ee extends beyond the virtues of its tera- Z run, as emphasized for instance in the studies [6, 7] (see also [10–12]). Of particular relevance to our Z' models, FCC-ee measurements above the Z -peak will put strong constraints on 4-fermion operators involving electrons.

The constraints on such operators coming from LEP-II, in particular from $e^+e^- \rightarrow \ell^+\ell^-$ cross-section and forward-backward asymmetry measurements, are already strong and probe a model in which the electron is charged up to scales of several TeV (as we show for our class of Z' models in §5.2). At FCC-ee, the bounds on dimension-6 4-lepton operators containing electrons will jump to $20 \div 30$ TeV [7]. Therefore, in our hunt for a light Z' model that could escape detection at FCC-ee, it had better not couple to electrons either, and so we should restrict to solutions in which $e_1 = 0$ also. We do so in §4.3.

2.3 Flavour Neutrality

The condition that all left-handed quarks and leptons are neutral under $U(1)_X$ means that such models can evade most constraints coming from flavour violation probes. This is because, while for left-handed quarks we know there is non-zero misalignment between the flavour and gauge eigenstates in order to reproduce the CKM matrix, the right-handed quarks can be aligned. If only right-handed fields are charged under the Z' , this means that models can be consistently constructed in which the Z' does not induce any flavour violation. More realistically, the quark flavour violation will be suppressed by the size of mixing angles amongst the right-handed quark fields, which must be sufficiently small.⁴ The same goes for leptons. So, this class of Z' models that are invisible to electroweak precision have a consistent limit in which they are invisible also to tests of flavour violation. This means that flavour-conserving constraints on 4-fermion operators will dominate the phenomenology, as we will see in §5.

3 Anomaly cancellation

The $U(1)_X$ extension of the SM should also be anomaly-free. In a $U(1)_X$ extension of the SM there are six possible triangle anomalies involving $U(1)_X$ gauge bosons and SM gauge bosons, including the mixed anomaly between $U(1)_X$ and gravity. Requiring these

³Viewed from a different perspective, the chiral and flavour non-universal nature opens up the possibility that $U(1)_X$ can explain aspects of the flavour puzzle: some Yukawa couplings must originate from higher-dimensional operators which are therefore suppressed with respect to the renormalisable top quark Yukawa coupling. But our focus here is on the phenomenology of the Z' and not on solving the flavour puzzle.

⁴Small right-handed mixing angles are a face-value expectation from the gauged $U(1)_X$ symmetries that we end up considering. This is because right-handed fields acquire non-universal charges under $U(1)_X$ and so in general we expect different EFT suppression factors appearing in each column of the Yukawa matrices, resulting in hierarchically suppressed right-handed mixing. See §4.6.

perturbative anomalies cancel⁵ results in six homogeneous polynomial equations that must be solved over the rationals: four are linear, one is quadratic, and one is cubic in the charges. One of the linear conditions, namely the mixed anomaly with $SU(2)_L$, is automatically satisfied because we already learnt that all $SU(2)_L$ doublets must be $U(1)_X$ neutral to avoid the running into EWPOs with g_2 .

Before we get started solving these equations, the reader might object that, since the $U(1)_X$ gauge symmetry is necessarily spontaneously broken, it is possible for gauge anomalies to be cancelled in the low-energy EFT by Wess–Zumino–Witten (WZW) terms [15, 16] à la Preskill [17], that we suppose arise from integrating out some heavy chiral fermions that ultimately restore anomaly cancellation in the underlying renormalisable theory. But it is reasonable to discard this class of EFTs in our pursuit of a Z' that is invisible on the Z -pole. Firstly, if the heavy fermions that cancel anomalies are SM singlets, they can at most cancel the pure 't Hooft anomaly associated with gauging $U(1)_X$ and the mixed anomaly between $U(1)_X$ and gravity. Such states are simply heavy right-handed neutrinos, and we discuss this extension of the SM field content explicitly in §4.4, finding that our main conclusions go through unchanged. Secondly, if we entertain heavy chiral fermions that cancel mixed anomalies between $U(1)_X$ and the SM, then they must be charged (although not necessarily chirally) under the SM gauge group. If they are electroweakly-charged, as in *e.g.* Ref. [18], the extra chiral fermions couple directly to the electroweak gauge bosons and so give 1-loop contributions to the EWPOs anyway, so we should strike out this option in our pursuit of an invisible Z' . If they are coloured, they radically alter the LHC phenomenology of such a model, with the heavy quarks likely providing the leading signature to target rather than the Z' .⁶

We therefore stick to the simpler scenario in which the $U(1)_X$ charges are chosen so that all gauge anomalies cancel given the SM chiral fermion contributions only, with the exception of possible right-handed neutrino contributions that we treat in §4.4.

3.1 Linear anomaly conditions

Given the mixed anomaly with $SU(2)_L$ is automatically satisfied, there are three linear anomaly cancellation conditions (ACCs) to solve, plus the two non-linear ones. Allowing a right-handed electron charge for now, recall there are eight rational charges left to fix, namely $\{u_1, u_2, d_i, e_i\}$ where $i \in \{1, 2, 3\}$. To proceed, it is convenient to define the sums of charges:

$$D = \sum_i d_i, \quad E = \sum_i e_i, \quad (3.1)$$

since it is these sums that appear in all the linear anomaly conditions, given their invariance under permuting family indices. Note the ‘ g_1 running’ condition (2.2) is then $2(u_1 + u_2) =$

⁵In a $U(1)_X$ extension of the SM it is sufficient to consider only the cancellation of perturbative gauge anomalies. As long as we do not add an odd number of chiral fermions with $SU(2)_L$ isospin $j \in 2\mathbb{Z} + 1/2$, there are no non-perturbative gauge anomalies to further consider [13, 14].

⁶Note that, precisely because the heavy quark is posited to cancel a mixed anomaly between colour and $U(1)_X$, we cannot make it arbitrarily heavier than the Z' , without simultaneously making the Z' arbitrarily weakly coupled to the SM fermions.

$E+D$. The mixed anomalies with $SU(3)_c$ and $U(1)_Y$ (the one with a single X leg) vanish iff $u_1 + u_2 = -D$ and $8(u_1 + u_2) = -2D - 6E$ respectively, which implies $E = D = -(u_1 + u_2)$. The g_1 running condition is then only satisfied if

$$u_1 + u_2 = 0 \quad \Rightarrow \quad E = D = 0. \quad (3.2)$$

All the linear constraints are now satisfied. Note that the mixed anomaly with gravity automatically vanishes without needing to enforce it as a constraint, because the other conditions already imply that the sums of charges for every fermion type equal zero.

One can also flip the logic, and consider first imposing all the linear ACCs including the mixed anomaly with gravity. It is known that this forces the ‘sums of charges’ D , E *etc* to be a linear combination of hypercharge and $B - L$, and so knowing also that all the left-handed fields are neutral immediately imply all sums of charges are zero. Then the 1-loop running with g_1 into Z -pole observables cancels exactly upon summing over flavours of each type of fermion. For instance, the contributions of 4-lepton SMEFT operators O_{ee}^{ijj} that arise from integrating out the Z' into the Higgs-bilepton current are

$$\dot{C}_{He}^{ii} \sim \sum_{j=1}^3 \begin{array}{c} e_i \\ \diagdown \\ \bullet \\ \diagup \\ e_i \end{array} \begin{array}{c} \text{---} \\ \diagdown \\ \bullet \\ \diagup \\ \text{---} \end{array} \begin{array}{c} e_j \\ \diagdown \\ \bullet \\ \diagup \\ e_j \end{array} \begin{array}{c} \text{---} \\ \diagdown \\ \bullet \\ \diagup \\ \text{---} \end{array} Z_\mu \propto g_1^2 \sum_{j=1}^3 e_j = 0 \quad \text{by ACCs}, \quad (3.3)$$

where a square dot denotes a dimension-6 operator obtained from integrating out the Z' at tree-level, here a 4-fermion operator with $(\bar{R}R)(\bar{R}R)$ chiral structure, while a circular dot denotes a SM coupling, in this case a gauge coupling. Separate cancellations occur when summing over down-type and up-type quarks in the loop. The same goes for the O_{Hu} and O_{Hd} Higgs-fermion bilinears at 1-loop. Thus the modifications of Z couplings to right-handed fermions, sometimes denoted δg_{Ri}^{Zf} for $f \in \{u, d, e\}$, vanish at this order.

This diagrammatic argument also makes it clear that the cancellations are exact only because we are neglecting the masses of the leptons and quarks running inside the loop. This assumption corresponds to dropping the dependence of the SMEFT RGEs on all Yukawa couplings other than y_t in Ref. [1]. We infer that the leading corrections to such Z couplings should scale with $m_{b,c,\tau}^2/M_X^2$, noting that there are no contributions with tops running in the loop because $u_3 = q_3 = 0$. Comparable contributions would likely come from including dimension-8 operators for a TeV scale Z' , given $\frac{m_b m_X}{m_Z^2} \sim \left(\frac{m_X}{2 \text{ TeV}}\right)^2$.

3.2 Corollary: zero 1-loop RGE with y_t and g_i

This non-renormalisation at 1-loop of the Z boson current is an example of a more general feature of the class of Z' models we have arrived at. The 1-loop RGEs, keeping only effects in y_t and the gauge couplings g_i , in fact vanish completely for *all* dimension-6 SMEFT operators, not just the subset (2.1) that enter the EWPOs at tree-level. It is again straightforward to see this by considering 1-loop diagrams – or by inspecting the complete formulae for 1-loop RGEs derived in Refs. [19–21].

Yukawa dependence. Firstly, let us consider running of any SMEFT operators with y_t in these models. Because the charges of u_3 , q_3 , and H are all required to be zero by our electroweak neutrality conditions, we do not generate any dimension-6 SMEFT operators at tree-level that can connect to a top Yukawa insertion, meaning all the 1-loop diagrams involving one SMEFT coefficient and insertions of the top Yukawa vanish. For example,

$$= 0 \quad (3.4)$$

because the 4-fermion operator vanishes at tree-level, and likewise for any dimension-6 SMEFT operator involving Higgses. Thus, if we drop all Yukawa couplings other than y_t , there is no 1-loop running with Yukawas for these models from the matching scale (at which the Z' is integrated out) down to the electroweak scale.

The leading RG running with Yukawa couplings will, for generic models in our class, be due to the bottom quark Yukawa $y_b \sim 0.01$ (using the value of y_b at a scale $\mu = 1$ TeV [22], assuming SM RGE up to that scale). This generates y_b^2 -suppressed running into 4-fermion operators with $(\bar{L}L)(\bar{R}R)$ chirality structure, from the diagrams

$$\Rightarrow \dot{C}_{qf}^{33ii} \sim y_b^2 \frac{g_X^2}{M_X^2} X_{d_3} X_{f_i} \quad f \in \{u, d, e\}, \quad (3.5)$$

This would give, for instance, a very small contribution to $e^+e^- \rightarrow b\bar{b}$ but only if e_1 were charged. The more important contribution, phenomenologically, would likely be y_b -suppressed running into Z -pole observables, due to

$$\Rightarrow \dot{C}_{Hf}^{ii} \sim y_b^2 \frac{g_X^2}{M_X^2} X_{d_3} X_{f_i} \quad f \in \{u, d, e\}, \quad (3.6)$$

leading to very small RG-induced modifications to the Z -boson couplings, including $Z \rightarrow b_R b_R$, $Z \rightarrow \mu_R \mu_R$, and $Z \rightarrow \tau_R \tau_R$ measured at the Z -pole. Given our Z' couples at tree-level only to right-handed fermions, there are no other diagrams to consider for running into any operators due to Yukawa dependence.

Small as these contributions already are, we shall see in §4.3 that the class of solutions that decouple from FCC-ee both on *and off* the Z -pole necessarily has one d_i vanishing, and it is attractive to take this to be $d_3 = 0$ due to the flavour structure implied by $U(1)_X$ (§4.6). In those models, the leading Yukawa running is now quadratic in $y_c \sim 3 \times 10^{-3}$, and so is a truly tiny RG effect compared to the leading ‘tree-level’ phenomenology that we will study in §5.

Gauge coupling dependence. The running with the gauge couplings can be captured via diagrams like that we already discussed in (3.3), which gives the modification to the gauge boson currents from which one can obtain the running into 4-fermion operators and Higgs-bifermion operators. Again these contributions cancel exactly upon setting all Yukawas other than y_t to zero, although the cancellation is qualitatively different in the two cases. For the Yukawa running, we saw the largest contributions (the y_t running) were simply set to zero by choosing $U(1)_X$ charge assignments that decouple from the Higgs and top; for the gauge running, on the other hand, the $U(1)_X$ charge assignment leads to a precise cancellation between non-vanishing one-loop diagrams like that in (3.3). This cancellation follows from flavour universality of the SM gauge interactions, combined with the strict flavour *non*-universality of the $U(1)_X$ force, in particular the ‘tracelessness’ conditions $D = E = U = 0$ that followed from anomaly cancellation plus our Z -pole running criteria. This general argument shows that it is not only the g_i -dependent running into the Higgs-bifermions operators constrained on the Z -pole that vanishes at one-loop, but also the g_i -dependent running into 4-fermion operators. This applies even for those 4-fermion operators that are generated at tree-level, *viz.*

$$C_{ff'}^{ijj} \sim \sum_{f'' \in \{u,d,e\}} \sum_{k=1}^3 \begin{array}{c} f_i \\ \bar{f}_i \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} f_k'' \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} f_j' \\ \bar{f}_j' \end{array} = 0. \quad (3.7)$$

As above, the cancellations are only perfect because we neglect the mass of all the SM fermions running in the loop, and so there will be small corrections to the g_i running that scale with the light Yukawas.

Finite parts. Finally, we comment on the finite parts of the 1-loop diagrams that should enter into the computation of low-energy observables, going beyond the leading log RGE approximation to the running. Such finite contributions have been found to be quantitatively important effects to include in phenomenological studies of $\text{tera-}Z$, for instance in Ref. [23]. In our case, the cancellation of the 1-loop diagrams will hold not only for the divergent parts of the loop integral that are captured via the SMEFT RGEs, but also for the finite parts. This is clear because the relevant $U(1)_X$ charges enter simply as overall pre-factors in the computation of each diagram.

3.3 Non-linear anomaly conditions

Continuing with our discussion of anomaly cancellation, there remain two non-linear conditions on the charges that should be satisfied for gauge anomalies to cancel. The non-linearity here means that the conclusions we can deduce from enforcing these constraints, which requires some elementary number theory, are somewhat less obvious, but no less powerful.

First, the mixed anomaly with hypercharge corresponding to a triangle diagram with two X legs, that is quadratic in charges, vanishes iff

$$4u_1^2 = \sum_i d_i^2 + e_i^2, \quad (3.8)$$

where we used $u_2 = -u_1$. Secondly, the cubic anomaly, corresponding to a triangle diagram with three X legs, vanishes iff

$$0 = \sum_i 3d_i^3 + e_i^3, \quad (3.9)$$

where we have used the fact that the linear constraints already imply $u_1^3 + u_2^3 = u_1^3 + (-u_1)^3 = 0$. In summary, we need to solve the non-linear equations (3.8) and (3.9) subject to the linear constraints $\sum_i d_i = \sum_i e_i = 0$.

4 Classifying Z' models with electroweak neutrality

4.1 Valence quarks are charged

We can deduce some immediate consequences of the non-linear ACCs. The right-hand-side of eq. (3.8) is a sum of squares which implies the left-hand-side is non-negative. Hence, a non-sterile⁷ solution requires

$$u_1 \neq 0, \quad u_2 = -u_1. \quad (4.1)$$

Thus, both the right-handed up and charm quarks are unavoidably charged under $U(1)_X$. The coupling to up quarks in particular means that production of the Z' boson in pp collisions has no flavour suppression, and so there will be strong constraints from the LHC.

Also notice that (3.8) does not only imply that the up quark has non-zero charge, but also that its charge is *large*, because:

$$2|u| \geq \max(|d_i|, |e_i|), \quad (4.2)$$

⁷We define a ‘sterile solution’ to be one in which all $U(1)_X$ charges of SM fields are zero, in which case the Z' boson would decouple entirely, modulo a possible kinetic mixing with the hypercharge gauge boson. If the kinetic mixing parameter $\epsilon X_{\mu\nu} B^{\mu\nu}$ is turned on, the Z' mimics a heavy copy of the SM hypercharge gauge boson, picking up couplings to the Higgs and all SM fermions after diagonalising the gauge field kinetic terms. The effective scale in that case would be $\Lambda = M/(\epsilon g_X)$, with the combination ϵg_X being the effective coupling, and this Z' would be probed up to scales of order 80 TeV [1] by the FCC-ee Z -pole run due to the direct couplings to all SM fields. Of course, such a Z' could also be very light ($M \ll M_Z$) and weakly-coupled, in which case it would serve as a dark photon – a phenomenologically different beast to the heavy, strongly coupled extensions of the SM that we are concerned with here.

i.e. the up charge (multiplied by two) is at least as big as the largest lepton or down-type quark charge. It follows immediately that the Z' has large coupling to protons and so is abundantly produced in a proton-proton collider.

Knowing this, the next most important thing to know phenomenologically is whether the Z' also couples to charged leptons, and how many flavours thereof. If it does, the Z' will contribute to Drell–Yan production of di-lepton pairs at the LHC (§5.3). And even if it does not, or if the lepton charges can be numerically suppressed, condition (4.1) implies the Z' will contribute to di-jet production at the LHC, which gives strong constraints for a Z' coupled to up quarks [8] as is the case here (§5.4).

We also emphasize that the up-type quarks necessarily have all different charges under $U(1)_X$, *viz.* $(u_1, u_2 = -u_1, u_3 = 0)$. This means there are no Z' models with $U(2)_u$ quark flavour symmetry, which would require $u_1 = u_2$, that satisfy our criteria for electroweak neutrality. But, as discussed in §2, this does not necessarily mean there are strong bounds from *e.g.* kaon mixing or other precision flavour probes, because the strong $U(2)$ -breaking is only present for right-handed quarks and there the mixing angles can be consistently tuned to be arbitrarily small.

4.2 At least two leptons are charged

It is quick to show that at least one right-handed lepton is necessarily charged, or else the $U(1)_X$ is sterile *i.e.* all charges are zero. To see this, let us try setting all three RH lepton charges e_i to zero. The cubic anomaly constraint becomes

$$0 = d_1^3 + d_2^3 + d_3^3. \quad (4.3)$$

Rather famously, this equation has no rational solutions in which all d_i are non-zero, but rather the only solutions satisfy $d_1 d_2 d_3 = 0$, of the form $d_1 = -d_2$, $d_3 = 0$ (up to permuting flavour indices). Substituting this into the quadratic, together with $e_i = 0$, we get

$$2u_1^2 = d_1^2, \quad (4.4)$$

which has no rational solutions. Thus, by contradiction, we establish that at least one e_i charge must be non-zero for there to exist rational solutions. But then the linear condition $E = \sum_i e_i = 0$ can only be satisfied if at least *two* e_i be non-zero. Thus, in addition to coupling to up and charm quarks, an anomaly-free Z' that does not run into the Z -pole observables necessarily couples at tree-level to (at least) two flavours of leptons. This combination of couplings means there are strong constraints coming from LHC Drell–Yan data at high- p_T (§5.3), including at least one light flavour lepton channel.

At least two down-type quarks are charged. The same mathematical argument used in §4.2 applies to the down quarks: that is, (3.8) and (3.9) together imply the d_i cannot all be zero, or else there would be no rational solution for the e_i charges, and then the linear constraint $D = \sum_i d_i = 0$ implies at least two of the d_i are non-zero.

4.3 Classifying solutions with one neutral lepton

As advertized in §2.2, to decouple the Z' from FCC-ee, indeed even from old LEP-II data, means it should not couple to the electron at tree-level. It is therefore of interest to set $e_1 = 0$ and study the space of solutions satisfying this condition.⁸ The space of all such Z' models, satisfying our anomaly-free and electroweak-neutrality criteria, is equivalent to the set of Pythagorean triples and so admits a straightforward analytic parametrization, as follows.

We know $\sum e_i = 0$ from the linear equations, therefore if we suppose one lepton is neutral, taking $e_1 = 0$ wlog, then all solutions are of the form $e_2 = -e_3$. The cubic then becomes $\sum_i d_i^3 = 0$, which implies by Fermat's theorem that $d_i \propto \{0, 1, -1\}$ also up to permutation. Therefore, all solutions in which one right-handed lepton is neutral are of the ‘parity-symmetric’ form, namely for each right-handed field of type $f \in \{u, d, e\}$ the charge assignment is $(f_1, f_2, f_3) = X_f(1, -1, 0)$ up to permutation. The X_f values, which we can define to all be non-negative without loss of generality, must then be chosen to satisfy the quadratic, which becomes

$$2X_u^2 = X_d^2 + X_e^2. \quad (4.5)$$

For example, one solution is $X_u = X_d = X_e = 1$, which we discuss further below when we turn to phenomenology. Equation (4.5) can be mapped to Pythagoras' equation, by rewriting it as

$$X_u^2 = \left(\frac{X_d - X_e}{2}\right)^2 + \left(\frac{X_d + X_e}{2}\right)^2; \quad (4.6)$$

since X_d and X_e are either both even or both odd to satisfy (4.5), we know $(X_d \pm X_e)/2$ are themselves both integers. The complete solution (over the integers) to Pythagoras' equation has been known for over two millenia. The set of all Z' models of interest can thence be parametrized in terms of two integers $p, q \in \mathbb{Z}^2$,

$$X_u = p^2 + q^2, \quad X_d = p^2 + 2pq - q^2, \quad X_e = q^2 + 2pq - p^2, \quad (4.7)$$

from Euclid's formula. We pick out some benchmarks to study in §5.

4.4 Including right-handed neutrinos

We have found that combining the requirements of electroweak neutrality (on and off the Z peak) with anomaly cancellation is very predictive for Z' extensions of the SM: to sum up, it implies the Z' couples at tree-level to up and charm quarks, at least two down-type quarks, muons, and taus, and that all inequivalent solutions up to family permutations are parametrized by a pair of integers.

One might wonder whether the above constraints can be relaxed by including up to three right-handed neutrinos, with non-universal charges $\{n_i\}$ with which one can try to ‘absorb’ gauge anomalies. Given the linear conditions excluding the mixed anomaly with

⁸To investigate solutions with three non-zero e_i charges, one could apply the geometric tools pioneered in the most general $\text{SM} \times U(1)_X$ case in Refs. [24–26].

gravity, none of which notice the inclusion of SM singlets, imply that $U = D = E = 0$, in order to not re-introduce the mixed anomaly with gravity we require also $N := \sum_{i=1}^3 n_i = 0$. Then, the only modification is that the cubic becomes

$$0 = \sum_i 3d_i^3 + e_i^3 + n_i^3. \quad (4.8)$$

Since the quadratic equation is not modified, the conclusion that right-handed up and charm quarks are charged is unchanged. Phenomenologically, there is no way to decouple the di-jet constraints.

At first glance it appears slightly more encouraging to try to find solutions in which, say, all the charged leptons decouple, if we are hoping to hide from LHC Drell–Yan constraints. Setting all $e_i = 0$, we now open up non-trivial solutions to the cubic in which all three d_i are non-zero thanks to the inclusion of $\{n_i\}$. The quadratic constraint would then be

$$4u_1^2 = d_1^2 + d_2^2 + d_3^2. \quad (4.9)$$

Considered in isolation, this equation has infinitely many solutions that can be straightforwardly parametrized in terms of four rational variables (by a generalisation of Euclid’s formula), with all d_i necessarily even. But it turns out that none of these solutions intersect the linear constraints imposed above. To show this, our strategy is to substitute the linear constraint $D = \sum_i d_i = 0$ into this quadratic to reduce the number of variables by one, and then appeal to Legendre’s theorem on the ternary quadratic form.

It is convenient to perform a change of variables on the $\{d_i\}$. Following [27], for instance, we define $d_{32} := d_3 - d_2$ and $\bar{d} := -2d_1 + d_2 + d_3$. Note that the linear map from $\{d_1, d_2, d_3\}$ to $\{D, d_{32}, \bar{d}\}$ and its inverse are rational. In the new variables, the quadratic becomes $3d_{32}^2 + \bar{d}^2 + 2D^2 = 24u_1^2$. In this form, it is easy to impose the linear condition, which is just $D = 0$, giving

$$d_{32}^2 + 3\left(\frac{\bar{d}}{3}\right)^2 = 2(2u_1)^2. \quad (4.10)$$

We have cast the quadratic equation in the form $ax^2 + by^2 = cz^2$, where the coefficients a , b and c are coprime and square-free positive integers. By Legendre’s theorem, this does not admit any rational solutions for $\{x, y, z\}$ because the Legendre symbol $(ac|b) = (2|3) = -1$.⁹ Accordingly, there are no rational solutions for our original variables $\{u_1, d_i\}$ to which $\{x, y, z\}$ are related by a rational linear transformation.

We thus prove by contradiction that there exist no anomaly-free non-sterile solutions that satisfy our Z -pole invisibility criteria in which the charged leptons are all neutral, even allowing for arbitrarily many right-handed neutrinos with any rational charges. As before, the linear constraint $E = \sum e_i = 0$ then implies at least two e_i are non-zero,¹⁰ which implies the Z' must couple at tree-level to either electrons or muons. As argued

⁹Recall that, in number theory, the Legendre symbol $(a|p)$ is, given an odd prime p , defined to equal 1 if a is congruent to a (non-zero) perfect square modulo p , equal to 0 if a is a multiple of p , and equal to -1 otherwise. In the case at hand, we have that $(2|3) = -1$ because every square is congruent to 0 or 1 modulo 3.

¹⁰The same proof can be given to show that at least two down-type quarks are also charged.

above, in order to hide the Z' from e^+e^- colliders, *i.e.* to make the Z' invisible both on and off the Z -peak at FCC-ee, we learn unambiguously that such a Z' couples at tree-level to muons as well as quarks, making it ideal to search for at the LHC.

4.5 Compatibility with semi-simple UV completions

While Z' extensions of the SM might be motivated by many phenomenological reasons, for instance to explain various ‘anomalies’ in flavour physics or to explain aspects of the flavour puzzle, one theoretical argument against such extensions is that they typically run counter to our desire for unification into a semi-simple gauge group. Unification remains attractive for curing the SM of the Landau pole associated with hypercharge while also explaining the observed quantization of hypercharge. Extending the SM gauge group by a second $U(1)_X$ factor near the TeV scale doubles the challenge of unifying into a semi-simple gauge group; to get an idea, note that if there were only one generation of SM fermions (or equivalently if we impose flavour universality), then only the choice $B - L$ is compatible with a semi-simple UV gauge embedding.

With three generations, there are significantly more options for gauged $U(1)_X$ extensions of the SM that have semi-simple embeddings: the options were completely classified in Ref. [28], assuming up to three RH neutrinos, by considering all possible Lie algebra embeddings of $\mathfrak{su}(3) \oplus \mathfrak{su}(2) \oplus \mathfrak{a} \hookrightarrow \mathfrak{g} \hookrightarrow \mathfrak{su}(48)$ consistent with the SM fermion representation (where \mathfrak{a} denotes an abelian Lie algebra containing hypercharge, \mathfrak{g} is the semi-simple Lie algebra we wish to embed in, and $\mathfrak{su}(48)$ is the maximal SM flavour symmetry with all interactions turned off), and then checking for (local and global) anomaly cancellation.

We now ask which of these extensions satisfy our additional criteria for not running at 1-loop into the EWPOs. The key observation that makes this in fact very simple is that electroweak neutrality required all LH fermions be neutral, as explained in §2. Given that, we can infer that the only anomaly-free $U(1)_X$ choices with semi-simple embeddings are of the form [28]

$$u_i = d_i = e_i = n_i = (a, b, -a - b) \quad (4.11)$$

up to permutations within families. The condition from y_t and g_3 running, which imposed $u_3 = 0$, then restricts us to a unique solution up to an overall normalisation

$$u_i = d_i = e_i = n_i = (1, -1, 0), \quad (4.12)$$

where again we are free to permute the family indices for d , e , and n , which leads to different phenomenology (see §5). So, strikingly, there is a *unique* anomaly-free $U(1)_X$ extension of the SM that doesn’t run into the Z -pole observables *and* which has a semi-simple UV completion. We further emphasize that including (at least two) right-handed neutrinos is now essential to reveal this possibility.

This $U(1)_X$ charge assignment embeds inside an $Sp(4)_R$ ‘electroweak flavour unification’ symmetry, studied in [29–31], in which right-handed up- and down-type fields (both quarks and leptons) of two generations are unified into 4-dimensional multiplets which must

be implemented alongside quark-lepton unification via the $SU(4)$ of Pati and Salam [32].¹¹ The $Sp(4)_R$ symmetry could be further embedded inside an $Sp(6)_R$; a candidate semi-simple embedding is then given by the group $SU(4) \times Sp(6)_L \times Sp(6)_R$, which can break down to the SM via our $SM \times U(1)_X$ extension. An alternative embedding of the right-handed $U(1)_X$ is inside the group $SO(6)_R$. All of these options are free of local and global gauge anomalies.

This unique charge assignment forms one of our benchmarks when we come to study the phenomenology in §5; from the phenomenological perspective it is an instructive benchmark because of the equal couplings to quarks and leptons.

4.6 Comments regarding quark and lepton flavour structures

All the chiral extensions of the SM considered above do not permit the full Yukawa coupling matrices at the renormalisable level. As mentioned in §2, the conditions $H = q_3 = u_3$ do, at least, guarantee the top quark Yukawa coupling is consistent with the $SM \times U(1)_X$ gauge symmetry for all models we consider. Beyond that, the gauge non-invariance of many elements of the Yukawa matrices means they must be populated by higher-dimension operators, which requires further states in the UV beyond just the $U(1)_X$ gauge field and the scalar field that breaks it, such as vector-like fermions. The Yukawa matrices in these models are therefore expected to possess hierarchical structure, with (at least) an order-1 top quark Yukawa coupling. Since our goal in this paper is not to address the flavour puzzle, but rather to explore the phenomenological complementarity of different experiments in probing anomaly-free Z' models, we are content to make only a few comments concerning flavour textures.

The parity-symmetric solutions in §4.3, for instance, admit a full column of each Yukawa matrix $Y_{u,d,e}$ at the renormalisable level. For up-type quarks this must be the third column, and for down-type quarks it is natural to assume the same since then y_b is also unsuppressed. But for leptons, we already took $e_1 = 0$ for phenomenological reasons. Such $U(1)_X$ symmetries therefore admit the first column of the lepton Yukawa but not the second or third. A natural implementation of these model presumably entails a different mechanism deeper in the UV (such as an additional gauge symmetry) that results in suppression of y_e .

Finally, it is worth remarking that the quark Yukawa textures naïvely predicted by gauging $U(1)_X$ are consistent with small mixing angles amongst the right-handed quarks, since, if we take $d_3 = 0$ to be the neutral right-handed down quark flavour, then we generally expect $Y_{u,d,e} \sim \begin{pmatrix} \delta & \epsilon & 1 \\ \delta & \epsilon & 1 \\ \delta & \epsilon & 1 \end{pmatrix}$, where ϵ and δ are small parameters related to the breaking of $U(1)_X$.¹² With such a texture, right-handed rotations would have the form $V_R \sim \begin{pmatrix} 1 & \delta & \delta\epsilon \\ \cdot & 1 & \epsilon \\ \cdot & \cdot & 1 \end{pmatrix}$, which further justifies our expectation that models in this class can have small flavour-violation and thus be more-or-less flavour neutral.

¹¹The results of [28], for instance via the `TestYourOwnCharges.nb` Mathematica notebook provided [here](#), can be used to deduce other semi-simple UV completions for this Z' model, all of which feature either some electroweak flavour unification for right-handed fields, or feature $SO(10)$ unification for all three families.

¹²For instance, in a Froggatt-Nielsen-like setup [33], these take the form $(v_X/M)^n$ where M is the mass of a heavy fermion and n is some integer.

5 Phenomenology off the Z pole

The set of Z' models that we have identified in the previous Sections, which are anomaly-free and do not run at 1-loop into Z -pole observables, all feature couplings to right-handed quarks and leptons, but nothing else. When the Z' is integrated out at tree-level, these models give rise to a set of flavour-conserving 4-fermion SMEFT operators with ‘RR’ chirality structure, which we can categorise into 4ℓ , $2q2\ell$, and $4q$ operators. The former are strongly constrained if there is a coupling to electrons, by cross-section measurements at electron-positron colliders (§5.2). Seeking invisibility at lepton colliders, this motivates coupling the Z' only to muon and tau leptons. The semi-leptonic operators, being flavour-conserving, are then best constrained by Drell–Yan data from the LHC (§5.3), which are reasonably strong for all lepton flavours in the final state. To hide from Drell–Yan $pp \rightarrow \ell\ell$ also, we can try to further restrict to Z' models in which the lepton charges are numerically very small. But even in this most invisible case, the $4q$ operators can be constrained from di-jet resonance searches at the LHC (§5.4).

We neglect the effects of renormalisation group (RG) running of the SMEFT operators from the matching scale Λ down to the relevant experimental energy scale. As explained in §3.2, this enters only at higher-order in the Wilson coefficients and/or is suppressed by light Yukawa couplings. For the models we focus on with $e_1 = 0$ and $d_3 = 0$, the leading RG effects are suppressed by $y_c^2 \sim 10^{-5}$ in addition to the loop factor and heavy mass scale, and are thus totally negligible.

5.1 Benchmark Models

To assess the strength of the various bounds, it is instructive to define two benchmark (BM) solutions, both within the class identified in §4.3 for which electrons are neutral, as follows:

$$\text{Benchmark 1: } (X_e = X_{u,d}) \quad (X_e, X_u, X_d) \propto (1, 1, 1), \quad (5.1)$$

$$\text{Benchmark 2: } (X_e \ll X_{u,d}) \quad (X_e, X_u, X_d) \propto (1, 29, 41). \quad (5.2)$$

The first BM model has equal charges for quarks and leptons, and gives a suitable BM for assessing the strength of the Drell–Yan constraints. As discussed in §4.5, this benchmark is special from the theoretical point of view because it is the only possible Z' extension in our class that has a semi-simple UV completion. The second BM has a numerically very small lepton charge,¹³ for which the Drell–Yan constraints also become subleading, leaving us with LHC dijet searches as the leading constraints.

We point out that a benchmark in which $X_e \gg X_{u,d}$, which we might hope to decouple from the LHC entirely and make truly invisible by coupling only to RH taus and muons, is *not* possible: this is because it is X_u that appears on the RHS of the quadratic ACC, which recall was $X_d^2 + X_e^2 = 2X_u^2$, and so X_u cannot be made small. In particular, we have the strict inequality

$$X_e \leq \sqrt{2}X_u. \quad (5.3)$$

¹³To obtain BM2, the Pythagorean triple solving (4.6) corresponds to a right-angled triangle whose perpendicular sides are of near-equal length, in this case the (20, 21, 29) triangle.

Thus, the benchmarks in (5.1) cover the two main limiting cases available to us, in which leptons either have comparable charge to quarks, or they have far smaller charge.

5.2 $4\ell: e^+e^- \rightarrow \ell\ell$

If the Z' couples at tree-level to electrons, there are bounds coming from measurements of the $e^+e^- \rightarrow \ell_i\ell_i$ cross-section (and forward-back asymmetry measurements) at electron-positron colliders. There are strong constraints already coming from the LEP-II measurements, which we calculate here for the Z' models of interest. The reach of FCC-ee will far exceed the LEP-II bounds, thanks to the proposed set of FCC-ee runs above the Z -pole.

The LEP-II constraints have been computed recently for Z' models in *e.g.* Refs. [34, 35], based on the χ^2 likelihood computed in [36]. The observables that go into computing the likelihood are the cross-sections $\sigma_{\ell\ell} := \sigma(e^+e^- \rightarrow \ell^+\ell^-)$ for $\ell \in \{e, \mu, \tau\}$, and the forward-backward asymmetry observables (defined as the forward cross-section minus the backward cross-section), for $\mu^+\mu^-$ and $\tau^+\tau^-$ final states. The relevant Wilson coefficients turned on in our Z' models are O_{ee}^{11jj} , where j can be any lepton flavour; without turning on the electron coupling, there are no bounds at leading order. It is therefore instructive to hone in on the ‘model-independent’ constraint that is unavoidable if we turn on only the electron coupling, which is the bound from σ_{ee} on $C_{ee}^{1111} = -g_X^2 X_e^2 / (2M_X^2)$, where X_e is the e_1 charge under $U(1)_X$. The 95% C.L. constraint, obtained from the $\chi^2 - \chi_{\min}^2 = 3.84$ contour (given we have a single independent model-parameter),¹⁴ is

$$\frac{M_X}{X_e g_X} \gtrsim 5.2 \text{ TeV}. \quad (5.4)$$

We remark that, if the Wilson coefficient C_{ee}^{1111} had the opposite sign, we find the bound would be $2 \div 3$ times weaker than this (*i.e.* 2 TeV), which reflects an underlying mild tension in the $e^+e^- \rightarrow e^+e^-$ data from LEP-II. This is in agreement with the results shown in [36, Fig. 1].

Of course, we already learnt that at least the tau or muon charge must also be turned on to satisfy the anomaly cancellation criteria, which only strengthens the bounds. For example, consider a ‘parity symmetric’ solution of the kind discussed in §4.3, with electron and muon charged equal and oppositely (for which the lepton Yukawa would mimic those for the quarks, *i.e.* with y_τ being renormalisable). The SMEFT Lagrangian is $\mathcal{L} = -\frac{g_X^2}{M_X^2} X_e^2 (\bar{e}_R \gamma_\mu e_R - \bar{\mu}_R \gamma_\mu \mu_R)^2$. Including also $\sigma_{\mu\mu}$ and the muon forward-backward asymmetry LEP-II measurement in the χ^2 likelihood, the bound increases a little to

$$\frac{M_X}{X_e g_X} \gtrsim 5.9 \text{ TeV}. \quad (5.5)$$

If one considered the option with e_R and τ_R charged, and μ_R neutral, including this time the $\sigma_{\tau\tau}$ and τ FB asymmetry observables, then the bound increases with respect to (5.4) by only 40 GeV or so.

¹⁴We are very grateful to Lukas Allwicher for sharing code for the LEP-II χ^2 likelihood that we used in this Section, and for related helpful discussions.

All these bounds on $M_X/X_e g_X$ would leap forward in reach with FCC-ee, thanks to its planned sequence of high-statistics runs above the Z -peak. As recently shown in Ref. [7], the expected increase in reach, for example for the C_{ee}^{1111} Wilson coefficient that is most important here, is by nearly an order of magnitude. In fact, given we know that the Z' must couple to at least tau or muon in addition to electrons, the strongest relevant bounds from [7] are on the (necessarily RR chirality) operators $O_{ee}^{11,22}$ and $O_{ee}^{11,33}$, which reach up to¹⁵

$$\frac{M_X}{g_X} \gtrsim 30 \text{ TeV} \quad (\text{FCC-ee above } Z\text{-peak}), \quad (5.7)$$

which actually well exceeds the scale of 7 TeV (in the case of flavour-universal couplings to SM fermions) that is probed due to 1-loop running into the Z -pole observables [1]. Therefore, as anticipated, we should discard this option in our hunt for a Z' that is invisible to FCC-ee.

5.3 $2q2\ell$: $pp \rightarrow \ell\ell$

Even if we decouple the Z' from electrons, it must then couple to both μ_R and τ_R , which leads to important bounds from Drell–Yan $pp \rightarrow \mu\mu$ and $pp \rightarrow \tau\tau$ measurements at the LHC. Here we use the HighPT package [37] to estimate the current strength of these constraints. The underlying ATLAS and CMS searches [38, 39] that go into our likelihood use 139 fb^{-1} of LHC data. A similar implementation of the bounds from Drell–Yan can be performed using the `smelli` software [34, 40, 41].

Here we focus on the solutions found in §4.3 in which the electron is neutral. Even though all charges are specified by two integer parameters p and q via (4.7), it is transparent to write things in terms of $X_{u,d,e}$ for now. Integrating out the Z' gives the following tree-level dimension-6 semi-leptonic SMEFT operators

$$\mathcal{L} = \mp \frac{g_X^2}{M_X^2} X_e X_u (\bar{\tau}_R \gamma_\mu \tau_R - \bar{\mu}_R \gamma_\mu \mu_R) (\bar{u}_R \gamma^\mu u_R - \bar{c}_R \gamma^\mu c_R) \quad (5.8)$$

$$\mp \frac{g_X^2}{M_X^2} X_e X_d (\bar{\tau}_R \gamma_\mu \tau_R - \bar{\mu}_R \gamma_\mu \mu_R) (\bar{d}_R \gamma^\mu d_R - \bar{s}_R \gamma^\mu s_R) \quad (5.9)$$

The sign ambiguities out the front correspond to permuting the signs of u_i or d_i charges, keeping the e_i permutation fixed.

To get a feel for the bounds, it is helpful to consider benchmark 1 (BM1) defined in (5.1). Given present data, there is a significant dependence of the bound that we extract on the possible sign choices in (5.8), which trace back to fluctuations in the data giving

¹⁵We also point out that lepton flavour universality violation (LFUV) is generically predicted by these models; to recap, it is forced upon us by anomaly cancellation, which implies both that $\sum_i e_i = 0$ and that at least one e_i is non-zero. FCC-ee provides new opportunities for LFUV tests, such as measuring the ratio

$$R_{\tau/\mu} = \frac{\sigma(e^+e^- \rightarrow \tau^+\tau^-)}{\sigma(e^+e^- \rightarrow \mu^+\mu^-)}, \quad (5.6)$$

as discussed in [7]. Such LFUV ratios would provide competitive probes of these Z' scenarios, although they are not independent from the more-typical ‘ R_i ’ ratios that were measured at LEP.

preference for non-zero values of the Wilson coefficients turned on by our Z' model. We obtain

$$\frac{M_X}{g_X} \gtrsim \begin{cases} 6.7 \text{ TeV}, & C_{eu}^{2211} > 0, C_{ed}^{2211} > 0 \\ 13.8 \text{ TeV}, & C_{eu}^{2211} < 0, C_{ed}^{2211} > 0 \\ 9.2 \text{ TeV}, & C_{eu}^{2211} > 0, C_{ed}^{2211} < 0 \\ 11.3 \text{ TeV}, & C_{eu}^{2211} < 0, C_{ed}^{2211} < 0 \end{cases} \quad (5.10)$$

We thus learn that, for equitable quark and lepton charges, the bounds from LHC Drell–Yan $pp \rightarrow \ell\ell$ data are strong, roughly excluding the Z' up to scales of 10 TeV or so.

These bounds, which are already probing very high scales due to the order-1 coupling of the Z' to both valence quarks and to charged leptons, will continue to strengthen as the LHC analyses more data, albeit slowly since the statistical uncertainty on the measured event rates goes down with the square root of the gain in luminosity. By the end of the High Luminosity LHC upgrade, assuming an integrated luminosity of 3 ab^{-1} of pp collision have been accrued, and that the predicted Standard Model background rate is measured in all bins of the high- p_T Drell–Yan distributions, we project the bound to be

$$\frac{M_X}{g_X} \gtrsim 22 \text{ TeV} \quad (\text{High-Luminosity LHC}), \quad (5.11)$$

which we obtained again using the `HighPT` package [37], with the bound being largely driven by the di-muon search, and where we average over the bounds one would obtain with the four different choices for the relative signs of charges, as in (5.10).

Even though we discarded the options in which electrons are charged in our hunt for an invisible Z' due to the very strong constraints one would obtain at FCC-ee above the Z -peak, we can give the corresponding HL-LHC projections in the other two scenarios for completeness, that is, in which electron and tau are charged and in which electron and muon are charged. We find

$$\frac{M_X}{g_X} \gtrsim 19 \text{ TeV} \quad \text{for BM1 } (e\tau), \quad (5.12)$$

$$\frac{M_X}{g_X} \gtrsim 24 \text{ TeV} \quad \text{for BM1 } (e\mu), \quad (5.13)$$

where again we averaged over the various sign choices. Unsurprisingly, the option in which both light leptons are charged is expected to get the strongest constraint from Drell–Yan.

5.4 $4q: pp \rightarrow jj$

The closest we get to an invisible Z' is one in which the lepton charge, while necessarily non-zero, is very small compared to the quark charges, as provided for instance by BM2 in (5.1). If we normalise the gauge coupling so that the largest charge is unity, then the lepton charge is a few percent, and the semi-leptonic signal in Drell–Yan is much suppressed. Of course, the physical, normalisation-independent quantities to compare between different Z' models are not charges or gauge couplings, but rather the ratios between different Wilson coefficients generated by the Z' . In particular, for BM2 we have the scaling

$$|C_{4q}| \sim \frac{X_{u,d}}{X_e} |C_{2q2l}| \sim \left(\frac{X_{u,d}}{X_e} \right)^2 |C_{4l}|, \quad (5.14)$$

with $X_u/X_e = 29$ and $X_d/X_e = 41$, meaning that 4-quark operator coefficients are over an order-of-magnitude larger than the semi-leptonic operator coefficients studied in §5.3. This Z' could therefore hide only in its contribution to $q\bar{q} \rightarrow q\bar{q}$ Drell–Yan and thus in di-jet searches at the LHC, the results of which were helpfully recast for narrow-width BSM mediators in generic SM gauge representations in Ref. [8] based on ATLAS [42, 43] and CMS [44, 45] resonant searches. EFT-based non-resonant analyses for di-jets were performed in [9, 46].

Note that for this type of benchmark model, *i.e.* featuring $|X_e| \ll |X_{u,d}|$, the anomaly cancellation condition $2X_u^2 = X_d^2 + X_e^2 \approx X_d^2$ directly implies $|X_d| \approx \sqrt{2}|X_u|$ which, in light of the results of [8], makes the dd and uu channels roughly equally strong. We can thus characterize BM2 via the approximate set of charges

$$X_d = 1, \quad X_u \approx \frac{1}{\sqrt{2}}, \quad X_e \ll 1, \quad (5.15)$$

where we choose a normalisation for the gauge coupling g_X in which the largest charge, that of down-type quarks, is unity.

Our goal in this Subsection is to use the results of [8, 9] to crudely estimate conservative bounds on such a Z' model coming from dijet resonant and non-resonant searches. The four resonant searches cover different overlapping regions of the kinematical range $m_{jj} \in (0.05, 5)$ TeV. We emphasize that the bound we estimate from the resonant searches, following [8], is qualitatively different from those for $4l$ and $2q2l$ that we studied in previous Subsections, which were obtained by integrating out the heavy Z' and matching onto the 4-fermion SMEFT operators to bound the effective scale $\Lambda_X = M_X/g_X$. For the narrow resonance di-jet search, the result is not a mass-independent bound on the scale as in §§5.2 and 5.3, but a bound that only holds within the kinematical range $M_X \leq 5$ TeV – indeed, as we next show, we cannot reliably get a bound even for this full range.

Constraining such a BSM non-resonant contribution that lifts the tails of the invariant mass distribution is a qualitatively different challenge for dijet distributions than for dileptons (§5.3), both experimentally and theoretically (see *e.g.* [9, 46, 47], where angular distributions of the dijet system are leveraged). While it is beyond the scope of this work to recast the experimental analyses to our model case, we nonetheless compare the resonant search bounds with an EFT-derived bound from [9] in the case that most closely matches our model, which gives a similar ballpark bound.

Narrow dijet resonance search. To reliably interpret the resonance search to put a bound on our Z' , the latter should be a narrow resonance with $\Gamma_X/M_X \lesssim 0.1$. We can compute the width by summing over the partial widths to $q\bar{q}$ bilinears:

$$\frac{\Gamma_X}{M_X} = \frac{g_X^2}{8\pi} (2|X_u|^2 + 2|X_d|^2 + \dots) \approx \frac{3g_X^2}{8\pi}, \quad (5.16)$$

where the \dots indicates the partial width to leptons which, for this BM2, is negligible. The narrow width condition thus puts an upper bound on the gauge coupling:

$$g_X < \sqrt{\frac{4\pi \times 0.1}{|X_u|^2 + |X_d|^2}} \approx 0.9 \quad (\text{BM2}), \quad (5.17)$$

given the choice of charge normalisation in (5.15). The dijet resonance search constraints from [8] can only be reliably used when this condition on g_X is satisfied. Because we are interpreting a resonance search, we do not integrate out the Z' and match to SMEFT. Rather, the relevant Lagrangian is described by the following couplings of the Z' to quarks,

$$L = g_X X_u (\bar{u} \gamma^\mu P_R u - \bar{c} \gamma^\mu P_R c) X_\mu + g_X X_d (\bar{d} \gamma^\mu P_R d - \bar{s} \gamma^\mu P_R s) X_\mu. \quad (5.18)$$

Given the equal-size first and second generation quarks, and the fact that the parton distribution functions (PDFs) are much suppressed for the second generation with respect to the first, we can drop the charm and strange quark contributions from now on.

From digitising the plot in [8], the perturbativity limit $g_X \lesssim 0.9$ (from $\Gamma_X/M_X < 1$) implies that we can use the constraint on the coupling to $d\bar{d}$, which is $X_d g_X \approx g_X$, out to masses $M_X \leq 4.2$ TeV. Likewise in the $u\bar{u}$ channel, the bound can be used up to slightly higher Z' masses $M_X \leq 4.4$ TeV, where we use the fact that $X_u g_X \approx g_X/\sqrt{2}$ from (5.15). Then at this end-point of validity, we can exclude this value of the mass and coupling; for that mass, all couplings $g_X > 0.9$ are then robustly excluded. To compare with our other EFT-derived constraints, which all depend on a scale $\sim M/(gX)$ rather than a mass and coupling individually, we can phrase the bound as

$$\frac{M_X}{g_X} \geq 4.9 \text{ TeV}, \quad \text{valid for } M_X \leq 4.4 \text{ TeV}. \quad (5.19)$$

We admit that this is a very crude estimate of the bound on our Z' coming from di-jet, but it is equally very conservative: we use only the resonance search, and we use it only in the region where the width over mass is < 0.1 . This is already sufficient to demonstrate that the bound is strong, at the level on 5 TeV on effective scale, meaning this Z' is far from being invisible at the LHC. The bound is driven by the large coupling to up quarks, which we showed is necessary for all Z' models that do not run into the Z -pole, which clearly emphasizes the perfect complementarity between the LHC and a future FCC-ee machine in probing *any* Z' extension of the SM that is not sterile.

4-quark operator bound from dijet angular distributions. In Ref. [9] the CMS collaboration perform an analysis of the angular distributions in dijet final states to obtain bounds on 4-quark effective operators using 35.9 fb^{-1} of 13 TeV data, for which the interpretation is not restricted by a kinematical upper limit on the mass. The scenario they consider that is closest to our family of Z' models is one with flavour-universal couplings to only right-handed quarks (corresponding to the Wilson coefficient η_{RR} in the experimental analysis [9]). Given BM2 has a similar strength coupling to u_R and d_R , we are not too far from this flavour-universality limit. Translating to our SMEFT Lagrangian, and taking both X_u and X_d to be approximated by the smaller of the two in our benchmark in order to get a conservative estimate of the bound, we infer the following EFT-derived bound on the Z' scale

$$\frac{M_X}{g_X} \geq 3.5 \text{ TeV} \quad \text{valid for all } M_X \text{ (up to perturbativity)}. \quad (5.20)$$

We reiterate that this bound would correspond to a Z' with couplings $X_u = X_d = 1/\sqrt{2}$, whereas for our ‘leptophobic’ anomaly-free class of models (realised by BM2 in (5.15)) we have $X_u \approx 1/\sqrt{2}$ and a larger $X_d = 1$, so if we performed a proper recast for BM2 one would expect the EFT bound to be somewhat stronger than (5.20) thanks to the enhanced down-quark coupling.

5.5 Comment on low-energy probes of LFUV

For completeness, we make some comments about low-energy flavour probes before moving to our summary. Recall from footnote 15 that our Z' models necessarily feature LFUV, with all three RH lepton charges being different. We there mentioned the possibility of probing this in e^+e^- colliders, from measuring the LFUV ratio $R_{\tau/\mu}$ (indeed, there would already be bounds coming from measurements of R_τ and R_μ at LEP). One might further wonder whether there are *low-energy* probes of this LFUV, for instance in B -meson decays and the class of $R_H := BR(B \rightarrow H\mu\mu)/BR(B \rightarrow H\ell\ell)$ measurements, where B is a b -hadron and H is some other hadron. These observables are extremely sensitive probes of new physics due to the change of quark flavour, typically $b \rightarrow s$ (as in $R_{K^{(*)}}$ [48], R_{K_s} and $R_{K^{*+}}$ [49], the very recent R_ϕ measurement which tests LFUV in $B_s^0 \rightarrow \phi\ell^+\ell^-$ [50]), for which the SM contribution is extremely suppressed.

The heavy Z' bosons under consideration can significantly impact these observables if it connects a flavour-violating bs vertex to a LFUV leptonic current. For left-handed bs currents, the Z' has no direct coupling to left-handed fields and so the BSM contribution goes through the same penguin loop as for the SM process, and is a negligible correction. There is, however, a right-handed contribution at tree-level that violates LFU. This will be linear in the right-handed mixing angles between b_R and s_R which, as we discussed in §2, are expected to be small in these models and can be consistently taken to zero to decouple the corresponding bounds.

It is perhaps worth noting that there are large, tree-level contributions to quark flavour *conserving* LFU tests. While high-precision measurements can here be made, for instance in J/ψ decays (with $\Gamma_{J/\psi \rightarrow e^+e^-}/\Gamma_{J/\psi \rightarrow \mu^+\mu^-}$ measured to be unity up to 1% precision [51] in the KEDR experiment [52]), these decays occur at tree-level in the SM and via photon exchange, and so the Z' will give a relative shift of order $g_X^2 M_{J/\psi}^2/e^2 M_X^2$ which is completely undetectable for the multi-TeV scale Z' mass that is forced upon us by the LHC dijet constraints alone.

Lastly, there are no LFUV effects in tau decays, because we only have neutral current LFUV but no concomitant LFV, which means that any process necessarily features an even number of taus and so cannot affect τ decays.

5.6 Comparison of benchmark models

We now put together our findings by comparing the dominant bounds for the benchmark Z' models considered in the main text, and also comparing against three benchmark scenarios studied in [1] that *do* run into the EWPOs, in Fig. 1. As we have argued, the conditions of electroweak neutrality plus anomaly cancellation require an order-1 coupling to up quarks,

which unavoidably results in strong bounds from the LHC experiments. The main phenomenological differences are then driven by which leptons are also charged (notably if electrons are charged, the Z' will be extremely well-probed at FCC *off* the Z -peak) and by whether the lepton charge is comparable to the quark charge or is numerically very small.

The expected sensitivities that we plot in Fig. 1 are for FCC-ee projections both on and off the Z -pole, where for the former we use [1] and for the latter we use [7]. For Drell–Yan we plot High-Luminosity LHC projections, which are straightforwardly extrapolated from current bounds assuming HL-LHC will accumulate an integrated luminosity of 3 ab^{-1} . For the dijet constraints we perform a similar naïve rescaling of the bound to an integrated luminosity of 3 ab^{-1} , using in this case the non-resonant search as quoted in (5.20) for the case of BM2, for which a very slow improvement of Λ scaling with the eighth root of the luminosity ratio is expected (similar to the non-resonant $pp \rightarrow \ell\ell$ case). An extrapolation of the resonant search result is much harder to estimate, because the improvement of the dijet reach depends crucially on the kinematic endpoints of the search. For the first three benchmarks described below, which correspond to the scenarios in [1], the coupling to quarks is vector-like (and flavour-universal in two cases) rather than right-handed as for BMs 1 and 2, and so we use the bounds for a different signal model in the non-resonant dijet analysis of CMS [9] to perform our HL-LHC extrapolation.

The benchmarks we include in Fig. 1, and the corresponding bounds, are:

- ‘Higgs’: the Z' couples at tree-level to the Higgs and to all SM fermion generations, resulting in tree-level corrections to EWPOs. For instance, this could arise from gauging a linear combination of $B - L$ and hypercharge [53]. This scenario will be probed up to 80 TeV or so by FCC-ee Z -pole run. Similarly strong Z -pole bounds would apply to less-minimal $U(1)_X$ extensions as long as the Higgs is charged, such as those associated with gauging a component of flavour non-universal hypercharge, *e.g.* in [54–58].
- ‘Universal’: this toy Z' model does not couple to the Higgs, but couples universally to all SM fermion generations. An anomaly-free realisation of this scenario would be to gauge $U(1)_{B-L}$. It was shown that 1-loop running into EWPOs means this scenario will be probed to 8 TeV by the FCC-ee Z -pole run alone [1]. But in fact, we observe that for this charge assignment showcased in [1] the running into EWPOs with y_t in fact cancels exactly, due to $u_3 = q_3$, so this 8 TeV bound comes primarily from the running with g_2 . We therefore point out that, for a chiral Z' model coupled differently to u_3 and q_3 , the reach of the FCC-ee *tera-Z* run would likely far exceed 8 TeV thanks to strong running with y_t . We supplement this with the bound recently obtained in [7] by projecting FCC-ee *off*-peak constraints; to plot the ballpark bound of 30 TeV in Fig. 1, we take the bounds on 4-lepton operators that are flavour-conserving and probe the highest scales, which are $O_{ll,ee}^{1122}$, and $O_{ll,ee}^{1133}$, coming from $e^+e^- \rightarrow \mu^+\mu^-(\tau^+\tau^-)$ cross-section measurements and related observables. We see that the off Z -peak measurements actually probe higher scales than due to RGE running into the EWPOs, but this is in large part due to the cancellation of the y_t running for this benchmark model as already noted. Since this model also has order-1

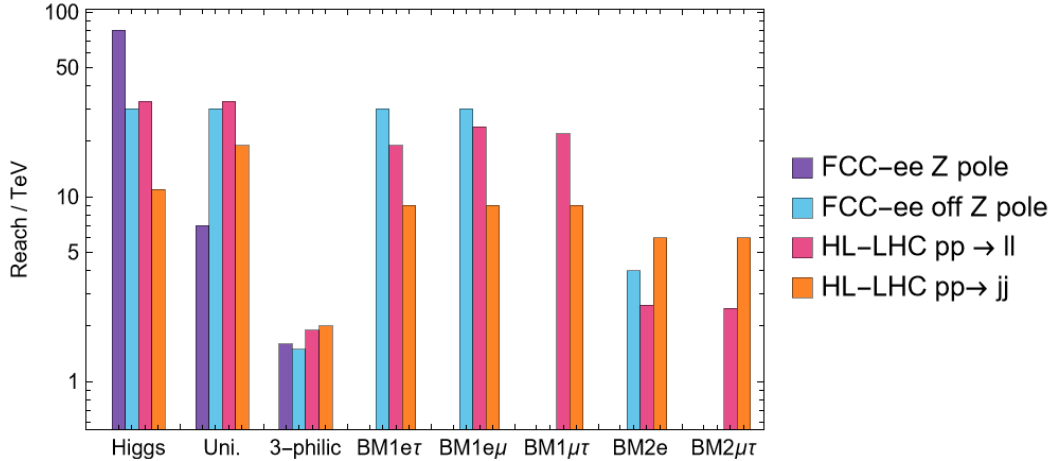


Figure 1: A comparison of the expected reach in M/g of FCC-ee *vs.* HL-LHC for a set of benchmark Z' models which are designed to cover, qualitatively speaking, all anomaly-free $U(1)_X$ -extensions of the SM. The first three scenarios are well-probed by EWPOs measured on the Z -pole at FCC-ee, as shown in [1]. The other four scenarios correspond to Z' models that do not run into the EWPOs at 1-loop, as classified in this paper. All such models necessarily have order-1 coupling to valence quarks, and so strong LHC bounds cannot be avoided. BM1 has equal couplings to RH quarks and RH leptons; if electrons are charged the model will be probed up to around 30 TeV by FCC-ee runs above the Z -pole, while even if electrons are not charged, $pp \rightarrow \ell\ell$ measurements at the LHC already exclude all BM1 variants up to 10 TeV, with HL-LHC projections reaching above 20 TeV. The final option BM2 has numerically tiny lepton charges, and so escapes all leptonic probes. But, as we show in the text, nothing that escapes the Z -pole at 1-loop can also escape the bounds from dijet searches at the LHC. We here estimate the sensitivity of HL-LHC in dijet searches by performing a naïve luminosity rescaling of the non-resonant search performed by CMS in [9].

couplings to quarks, we perform a projection for the expected HL-LHC sensitivity in Drell–Yan $pp \rightarrow \ell\ell$ using HighPT [37], including all three lepton flavour channels, which would already surpass 30 TeV even before the FCC-ee begins running. Lastly, for dijets we extrapolate the non-resonant bound from [9] relevant for a vector-like 4-quark operator, to obtain a strong projected sensitivity of 11 TeV.

- ‘3rd Gen’: this toy Z' model does not couple to the Higgs or to light fermions, but couples to all third generation SM fermions. An anomaly-free realisation of this scenario would be to gauge $U(1)_{B_3-L_3}$. The expected reach from EWPOs at FCC-ee was here shown to be 1.6 TeV from the 1-loop RGE running in [1], which is again subject to the huge caveat that the running with y_t accidentally cancels exactly due to assigning $u_3 = q_3$ for the benchmark. Studies at the level of single SMEFT operator bounds show that pure third generation 4-fermion operators would get bounds of order $10 \div 15$ TeV from the FCC-ee Z -pole run – see *e.g.* [7, 59]. For this third-family aligned Z' there are no tree-level bounds from the off Z -peak data because we do not couple to electrons, but at 1-loop the third generation 4-fermion operators will run into operators involving electrons, and so are still bounded by off- Z -peak measurements. In [7, Table VII] it was shown the bounds are of order 1.5 TeV, which

we include in Fig. 1. We also compute the sensitivity to this benchmark Z' model from $pp \rightarrow \tau\tau$ Drell–Yan measurements at the LHC, again using HighPT [37]. For this benchmark it is harder to estimate the HL sensitivity of the LHC dijet searches, because there is no EFT scenario in the non-resonant search that is third-family specific. In this case, we extract a very crude current bound of $M_X/g_X \gtrsim 1$ TeV or so from the resonant search via the $b\bar{b}$ production mode using [8], in a similar manner to §5.4. For the other models, we saw the bounds from non-resonant and resonant searches were similar in the region where both apply. If we assume similar holds for this third-generation specific scenario, we can naïvely extrapolate with the gain in luminosity to obtain a crude projection of 1.5 TeV or thereabouts.

- BM1 refers to our benchmark model 1 as introduced in the main text, (5.1). We include bounds for three variants of this benchmark corresponding to permutations of the lepton flavour assignment; the key difference is that when $e_1 \neq 0$ there is a bound reaching 30 TeV from $e^+e^- \rightarrow \ell^+\ell^-$ off the Z peak. The HL-LHC Drell–Yan estimates are also slightly different for the three variants (§5.3), but all are roughly 20 TeV. We estimate the HL-LHC dijet non-resonant bound directly from the η_{RR} signal model used by CMS in [9], which is here appropriate because $X_u = X_d = 1$, so we are able to set a slightly higher projection (than for BM2) of 9 TeV.
- BM2 refers to the second benchmark introduced in (5.1) in which the lepton charge is very small compared to the quark charge. The most important bound in this case will likely come from dijet measurements, where we rescale the non-resonant bound (5.20) with the luminosity gain to the eighth root to reach 6 TeV. Even though the lepton charge is very small, we find the expected sensitivity of HL-LHC Drell–Yan searches reaches nearly 2 TeV, driven by the dimuon channel. We also include a variant of BM2 (labelled BM2e in Fig. 1) in which electrons are charged for the sake of comparison, showing (from [7]) that even with a numerically very small charge $|e_1| = |d_1|/41$ the off Z -peak measurements can reach above 4 TeV in sensitivity, largely thanks to the expected performance of bottom and charm tagging. The case in which $e_1 = 0$, labelled BM2 $\mu\tau$, is totally decoupled from the FCC-ee electroweak program, like BM1 $\mu\tau$.

From Fig. 1 it is clear that there is excellent complementarity between e^+e^- and hh colliders for probing Z' models: in essence, the former puts strong constraints on the $U(1)_X$ coupling to leptons and to the Higgs, both through precision measurements at the Z -pole but also $e^+e^- \rightarrow f\bar{f}$ measurements off the Z -pole, while the latter constrains also the $U(1)_X$ couplings to quarks. Moreover, the conditions from requiring anomaly cancellation serve to tie the various charges together in an intricate way, so that there is no way to put the charges just into, say, heavy generation right-handed leptons or just into heavy flavour right-handed quarks, which would be the hardest scenario to detect.

6 Conclusion

In this paper, we sought to ‘fill a gap’ left by [1], which demonstrated that almost all single particle extensions of the SM will be probed at least to the TeV scale purely through measuring EWPOs to very high precision, thanks to the mixing of operators under RG running of the SMEFT. Essentially the only exception that could be viable at low effective scales, considering also flavour probes and direct LHC bounds on coloured states, was found to be a family of Z' extensions of the SM.

We here close this gap by showing that, by combining the requirements from not running into the EWPOs with the requirements of gauge anomaly cancellation, any Z' that can escape the Z -pole cannot also escape the LHC, especially moving into the High-Luminosity phase. This is primarily because all such Z' s must have order-1 couplings to up quarks. Anomaly cancellation also implies that at least two flavours of leptons must be charged.¹⁶ We emphasize that options in which electrons are directly charged would be probed up to 30 TeV by FCC-ee data *off* the Z -peak, surpassing the scales probed by the Z -pole due to RGE running; this reinforces the findings of the recent study [7] that measurements above the Z -pole should form an integral part of the electroweak program at FCC-ee. Requiring the electron is neutral, we are able to classify all remaining anomaly-free Z' models analytically in terms of a pair of integers, up to certain permissible family index permutations.

Continuing in pursuit of an invisible Z' , the two charged leptons that couple to the Z' ought to be muons and taus. In this case, Drell–Yan $pp \rightarrow \mu\mu, \tau\tau$ measurements from ATLAS and CMS put very strong bounds up to 10 TeV, which are projected to reach $20 \div 25$ TeV after HL-LHC. In a last ditch attempt to evade also these bounds from $pp \rightarrow \ell\ell$, we are led to consider rather bizarre models in which the lepton charge is numerically very small compared to the quark charge (which, incidentally, we prove cannot embed in any semi-simple UV completion). Even in this esoteric corner of model space, there remain strong constraints from LHC di-jet measurements that cannot be avoided, because the coupling to up quarks must be order-1. A conservative estimate of the current bound sits at 5 TeV. A naïve rescaling of the non-resonant bound on 4-quark operators coming from angular dijet analyses to the HL-LHC target integrated luminosity would give a projected sensitivity in the $6 \div 10$ TeV range.

Reversing the argument: any anomaly-free Z' model that is *not* already constrained above the 10 TeV mark after the LHC, for instance because it does not couple to light generation quarks, will necessarily run with y_t and/or the SM gauge couplings into the Z -pole observables at 1-loop and so will be probed at FCC-ee. If the Higgs is charged under $U(1)_X$, it is strikingly clear from [1] that the sensitivity of FCC-ee will be unrivalled. Thus, in addition to closing the gap left by [1], we can make a broader conclusion, which is that there ought to be a perfect complementarity between a high-precision e^+e^- collider (making the most of measurements both on and off the Z -pole) and the HL-LHC for probing generic anomaly-free $U(1)_X$ extensions of the SM gauge group.

¹⁶We showed that none of these results are modified by the inclusion of an arbitrary number of right-handed neutrino fields with which we might try to ‘soak up’ gauge anomalies.

Acknowledgement

I am grateful to Lukas Allwicher, Marzia Bordone, Admir Greljo, Sophie Renner, and Peter Stangl for discussions, and to Lukas Allwicher, Admir Greljo, and Sophie Renner for their helpful comments on the manuscript. I am especially grateful to Lukas Allwicher for sharing code used to compute the bounds on $4l$ operators from LEP-II data in §5.2.

References

- [1] L. Allwicher, M. McCullough and S. Renner, *New Physics at Tera-Z: Precision Renormalised*, [2408.03992](#).
- [2] Anisha, S. Das Bakshi, S. Banerjee, A. Biekötter, J. Chakraborty, S. Kumar Patra et al., *Effective limits on single scalar extensions in the light of recent LHC data*, *Phys. Rev. D* **107** (2023) 055028, [[2111.05876](#)].
- [3] G. Durieux, M. McCullough and E. Salvioni, *Charting the Higgs self-coupling boundaries*, *JHEP* **12** (2022) 148, [[2209.00666](#)].
- [4] L. Allwicher, G. Isidori, J. M. Lizana, N. Selimovic and B. A. Stefanek, *Third-family quark-lepton Unification and electroweak precision tests*, *JHEP* **05** (2023) 179, [[2302.11584](#)].
- [5] A. Greljo, P. Stangl, A. E. Thomsen and J. Zupan, *On $(g-2)_\mu$ from gauged $U(1)_X$* , *JHEP* **07** (2022) 098, [[2203.13731](#)].
- [6] S.-F. Ge, Z. Qian, M. J. Ramsey-Musolf and J. Zhou, *New Physics Off the Z-Pole: $e^+e^- \rightarrow f\bar{f}$ at Future Lepton Colliders*, [2410.17605](#).
- [7] A. Greljo, H. Tiblom and A. Valenti, *New Physics Through Flavor Tagging at FCC-ee*, [2411.02485](#).
- [8] M. Bordone, A. Greljo and D. Marzocca, *Exploiting dijet resonance searches for flavor physics*, *JHEP* **08** (2021) 036, [[2103.10332](#)].
- [9] CMS collaboration, A. M. Sirunyan et al., *Search for new physics in dijet angular distributions using proton–proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter and other models*, *Eur. Phys. J. C* **78** (2018) 789, [[1803.08030](#)].
- [10] J. de Blas, Y. Du, C. Grojean, J. Gu, V. Miralles, M. E. Peskin et al., *Global SMEFT Fits at Future Colliders*, in *Snowmass 2021*, 6, 2022, [2206.08326](#).
- [11] E. Celada, T. Giani, J. ter Hoeve, L. Mantani, J. Rojo, A. N. Rossia et al., *Mapping the SMEFT at high-energy colliders: from LEP and the (HL-)LHC to the FCC-ee*, *JHEP* **09** (2024) 091, [[2404.12809](#)].
- [12] F. Maltoni, C. Severi, S. Tentori and E. Vryonidou, *Quantum tops at circular lepton colliders*, *JHEP* **09** (2024) 001, [[2404.08049](#)].
- [13] J. Davighi, B. Gripaios and N. Lohitsiri, *Global anomalies in the Standard Model(s) and Beyond*, *JHEP* **07** (2020) 232, [[1910.11277](#)].
- [14] Z. Wan and J. Wang, *Beyond Standard Models and Grand Unifications: Anomalies, Topological Terms, and Dynamical Constraints via Cobordisms*, *JHEP* **07** (2020) 062, [[1910.14668](#)].
- [15] J. Wess and B. Zumino, *Consequences of anomalous Ward identities*, *Phys. Lett. B* **37** (1971) 95–97.

- [16] E. Witten, *Global Aspects of Current Algebra*, *Nucl. Phys. B* **223** (1983) 422–432.
- [17] J. Preskill, *Gauge anomalies in an effective field theory*, *Annals Phys.* **210** (1991) 323–379.
- [18] J. Davighi, *Anomalous Z' bosons for anomalous B decays*, *JHEP* **08** (2021) 101, [[2105.06918](#)].
- [19] E. E. Jenkins, A. V. Manohar and M. Trott, *Renormalization Group Evolution of the Standard Model Dimension Six Operators I: Formalism and λ Dependence*, *JHEP* **10** (2013) 087, [[1308.2627](#)].
- [20] E. E. Jenkins, A. V. Manohar and M. Trott, *Renormalization Group Evolution of the Standard Model Dimension Six Operators II: Yukawa Dependence*, *JHEP* **01** (2014) 035, [[1310.4838](#)].
- [21] R. Alonso, E. E. Jenkins, A. V. Manohar and M. Trott, *Renormalization Group Evolution of the Standard Model Dimension Six Operators III: Gauge Coupling Dependence and Phenomenology*, *JHEP* **04** (2014) 159, [[1312.2014](#)].
- [22] Z.-z. Xing, H. Zhang and S. Zhou, *Updated Values of Running Quark and Lepton Masses*, *Phys. Rev. D* **77** (2008) 113016, [[0712.1419](#)].
- [23] J. Gargalionis, J. Quevillon, P. N. H. Vuong and T. You, *Linear Standard Model extensions in the SMEFT at one loop and Tera-Z*, [2412.01759](#).
- [24] B. C. Allanach, B. Gripaios and J. Tooby-Smith, *Geometric General Solution to the $U(1)$ Anomaly Equations*, *JHEP* **05** (2020) 065, [[1912.04804](#)].
- [25] B. C. Allanach, B. Gripaios and J. Tooby-Smith, *Solving local anomaly equations in gauge-rank extensions of the Standard Model*, *Phys. Rev. D* **101** (2020) 075015, [[1912.10022](#)].
- [26] B. C. Allanach, B. Gripaios and J. Tooby-Smith, *Anomaly cancellation with an extra gauge boson*, *Phys. Rev. Lett.* **125** (2020) 161601, [[2006.03588](#)].
- [27] B. C. Allanach, J. Davighi and S. Melville, *An Anomaly-free Atlas: charting the space of flavour-dependent gauged $U(1)$ extensions of the Standard Model*, *JHEP* **02** (2019) 082, [[1812.04602](#)].
- [28] J. Davighi and J. Tooby-Smith, *Flatland: abelian extensions of the Standard Model with semi-simple completions*, *JHEP* **09** (2022) 159, [[2206.11271](#)].
- [29] J. Davighi and J. Tooby-Smith, *Electroweak flavour unification*, *JHEP* **09** (2022) 193, [[2201.07245](#)].
- [30] J. Davighi, *Gauge flavour unification: From the flavour puzzle to stable protons*, *Nuovo Cim. C* **46** (2022) 23, [[2206.04482](#)].
- [31] J. Davighi, G. Isidori and M. Pesut, *Electroweak-flavour and quark-lepton unification: a family non-universal path*, *JHEP* **04** (2023) 030, [[2212.06163](#)].
- [32] J. C. Pati and A. Salam, *Lepton Number as the Fourth Color*, *Phys. Rev. D* **10** (1974) 275–289.
- [33] C. D. Froggatt and H. B. Nielsen, *Hierarchy of Quark Masses, Cabibbo Angles and CP Violation*, *Nucl. Phys. B* **147** (1979) 277–298.
- [34] A. Greljo, J. Salko, A. Smolkovič and P. Stangl, *Rare b decays meet high-mass Drell-Yan*, *JHEP* **05** (2023) 087, [[2212.10497](#)].
- [35] B. Allanach and A. Mullin, *Plan B: new Z' models for $b \rightarrow sll$ anomalies*, *JHEP* **09** (2023)

- 173, [2306.08669].
- [36] A. Falkowski and K. Mimouni, *Model independent constraints on four-lepton operators*, *JHEP* **02** (2016) 086, [1511.07434].
- [37] L. Allwicher, D. A. Faroughy, F. Jaffredo, O. Sumensari and F. Wilsch, *HighPT: A tool for high- p_T Drell-Yan tails beyond the standard model*, *Comput. Phys. Commun.* **289** (2023) 108749, [2207.10756].
- [38] ATLAS collaboration, G. Aad et al., *Search for heavy Higgs bosons decaying into two tau leptons with the ATLAS detector using pp collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **125** (2020) 051801, [2002.12223].
- [39] CMS collaboration, A. M. Sirunyan et al., *Search for resonant and nonresonant new phenomena in high-mass dilepton final states at $\sqrt{s} = 13$ TeV*, *JHEP* **07** (2021) 208, [2103.02708].
- [40] J. Aebischer, J. Kumar, P. Stangl and D. M. Straub, *A Global Likelihood for Precision Constraints and Flavour Anomalies*, *Eur. Phys. J. C* **79** (2019) 509, [1810.07698].
- [41] P. Stangl, *smelli – the SMEFT Likelihood*, *PoS TOOLS2020* (2021) 035, [2012.12211].
- [42] ATLAS collaboration, M. Aaboud et al., *Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **121** (2018) 081801, [1804.03496].
- [43] ATLAS collaboration, G. Aad et al., *Search for new resonances in mass distributions of jet pairs using 139 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **03** (2020) 145, [1910.08447].
- [44] CMS collaboration, A. M. Sirunyan et al., *Search for narrow and broad dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter mediators and other new particles*, *JHEP* **08** (2018) 130, [1806.00843].
- [45] CMS collaboration, A. M. Sirunyan et al., *Search for high mass dijet resonances with a new background prediction method in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **05** (2020) 033, [1911.03947].
- [46] ATLAS collaboration, M. Aaboud et al., *Search for new phenomena in dijet events using 37 fb^{-1} of pp collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **96** (2017) 052004, [1703.09127].
- [47] S. Alioli, M. Farina, D. Pappadopulo and J. T. Ruderman, *Precision Probes of QCD at High Energies*, *JHEP* **07** (2017) 097, [1706.03068].
- [48] LHCb collaboration, R. Aaij et al., *Measurement of lepton universality parameters in $B^+ \rightarrow K^+\ell^+\ell^-$ and $B^0 \rightarrow K^{*0}\ell^+\ell^-$ decays*, *Phys. Rev. D* **108** (2023) 032002, [2212.09153].
- [49] LHCb collaboration, R. Aaij et al., *Tests of lepton universality using $B^0 \rightarrow K_S^0\ell^+\ell^-$ and $B^+ \rightarrow K^{*+}\ell^+\ell^-$ decays*, *Phys. Rev. Lett.* **128** (2022) 191802, [2110.09501].
- [50] LHCb collaboration, R. Aaij et al., *Test of lepton flavour universality with $B_s^0 \rightarrow \phi\ell^+\ell^-$ decays*, [2410.13748](#).
- [51] KEDR collaboration, V. M. Aulchenko et al., *Measurement of the ratio of the leptonic widths $\Gamma_{ee}/\Gamma_{\mu\mu}$ for the J/ψ meson*, *Phys. Lett. B* **731** (2014) 227–231, [1311.5005].
- [52] V. V. Anashin et al., *The KEDR detector*, *Phys. Part. Nucl.* **44** (2013) 657–702.

- [53] W. Altmannshofer, J. Davighi and M. Nardecchia, *Gauging the accidental symmetries of the standard model, and implications for the flavor anomalies*, *Phys. Rev. D* **101** (2020) 015004, [[1909.02021](#)].
- [54] B. C. Allanach and J. Davighi, *Third family hypercharge model for $R_{K^{(*)}}$ and aspects of the fermion mass problem*, *JHEP* **12** (2018) 075, [[1809.01158](#)].
- [55] B. C. Allanach and J. Davighi, *Naturalising the third family hypercharge model for neutral current B -anomalies*, *Eur. Phys. J. C* **79** (2019) 908, [[1905.10327](#)].
- [56] J. Davighi and B. A. Stefanek, *Deconstructed hypercharge: a natural model of flavour*, *JHEP* **11** (2023) 100, [[2305.16280](#)].
- [57] M. Fernández Navarro and S. F. King, *Tri-hypercharge: a separate gauged weak hypercharge for each fermion family as the origin of flavour*, *JHEP* **08** (2023) 020, [[2305.07690](#)].
- [58] M. Fernández Navarro, S. F. King and A. Vicente, *Minimal complete tri-hypercharge theories of flavour*, *JHEP* **07** (2024) 147, [[2404.12442](#)].
- [59] L. Allwicher, C. Cornella, G. Isidori and B. A. Stefanek, *New physics in the third generation. A comprehensive SMEFT analysis and future prospects*, *JHEP* **03** (2024) 049, [[2311.00020](#)].