

Nnaturalness

Nima Arkani-Hamed,¹ Timothy Cohen,² Raffaele Tito D’AgNolo,¹
 ANson Hook,³ HyuNg Do Kim,⁴ and David PiNner⁵

¹ *School of Natural Sciences, Institute for Advanced Study, Princeton, New Jersey 08540, USA*

² *Institute of Theoretical Science, University of Oregon, Eugene, OR 97403, USA*

³ *Stanford Institute for Theoretical Physics, Stanford University, Stanford, CA 94305, USA*

⁴ *Department of Physics and Astronomy and Center for Theoretical Physics,
 Seoul National University, Seoul 151-747, Korea*

⁵ *Princeton Center for Theoretical Science, Princeton University, Princeton, NJ 08544, USA*

Abstract

We present a new solution to the electroweak hierarchy problem. We introduce N copies of the Standard Model with varying values of the Higgs mass parameter. This generically yields a sector whose weak scale is parametrically removed from the cutoff by a factor of $1/\sqrt{N}$. Ensuring that reheating deposits a majority of the total energy density into this lightest sector requires a modification of the standard cosmological history, providing a powerful probe of the mechanism. Current and near-future experiments will explore much of the natural parameter space. Furthermore, supersymmetric completions which preserve grand unification predict superpartners with mass below $m_W \times M_{\text{pl}}/M_{\text{GUT}} \sim 10$ TeV.

I. MECHANISM

This letter describes a new mechanism, dubbed “Nnaturalness,” which solves the hierarchy problem. It predicts no new particles at the LHC, but does yield a variety of experimental signatures for the next generation of CMB and large scale structure experiments [1, 2]. Well-motivated supersymmetric incarnations of this model predict superpartners beneath the scale $m_W \times M_{\text{pl}}/M_{\text{GUT}} \sim 10$ TeV, accessible to a future 100 TeV collider [3, 4].

The first step is to introduce N sectors which are mutually non-interacting. The detailed particle content of these sectors is unimportant, with the exception that the Standard Model (SM) should not be atypical; many sectors should contain scalars, chiral fermions, unbroken gauge groups, etc. For simplicity, we imagine that they are exact copies of the SM, with the same gauge and Yukawa structure.

It is crucial that the Higgs mass parameters are allowed to take values distributed between $-\Lambda_H^2$ and Λ_H^2 , where Λ_H is the (common) scale that cuts off the quadratic divergences. Then for a wide range of distributions, the generic expectation is that some sectors are accidentally tuned at the $1/N$ level, $|m_H^2|_{\text{min}} \sim \Lambda_H^2/N$. We identify the sector with the smallest non-zero Higgs vacuum expectation value (vev), $\langle H \rangle = v$, as “our” SM. This picture is illustrated schematically in Fig. 1.

In order for small values of m_H^2 to be populated, the distribution of the mass parameters must pass through zero. For concreteness, we take a simple uniform distri-




FIG. 1: A sketch of the N naturalness setup. The sectors have been ordered so that they range from $m_H^2 \sim \Lambda_H^2$ to $-\Lambda_H^2$. The sector with the smallest vacuum expectation value contains our copy of the SM.

bution of mass squared parameters, indexed by an integer label i such that

$$(m_H^2)_i = -\frac{\Lambda_H^2}{N}(2i + r), \quad -\frac{N}{2} \leq i \leq \frac{N}{2}, \quad (1)$$

where $i = 0 =$ “us” is the lightest sector with a non-zero vev: $(m_H^2)_{\text{us}} = -r \times \Lambda_H^2/N \simeq -(88 \text{ GeV})^2$ is the Higgs mass parameter inferred from observations. The parameter r can be seen as a proxy for fine-tuning,¹ since

¹ There are a variety of other ways one might choose to implement a measure of fine-tuning in this model. For example, one could assume the distribution of Higgs mass squared parameters

it provides a way to explore how well the naive relation between the cutoff and the mass scale of our sector works in a detailed analysis. Specifically, $r = 1$ corresponds to uniform spacing, while $r < 1$ models to an accidentally larger splitting between our sector and the next one. A simple physical picture for this setup is that the new sectors are localized to branes which are displaced from one another in an extra dimension. In this scenario, the lack of direct coupling is clear, and the variation of the mass parameters can be explained geometrically: the m_H^2 parameters may be controlled by the profile of a quasi-localized field shining into the bulk.

As a consequence of the existence of a large number of degrees of freedom, the hierarchy between Λ_H and the scale Λ_G where gravity becomes strongly coupled is reduced. The renormalization of the Newton constant implies $\Lambda_G^2 \sim M_{\text{pl}}^2/N$. If perturbative gauge coupling unification is to be preserved $\Lambda_G \gtrsim M_{\text{GUT}}$, implying that $N \lesssim 10^4$. This gives a cutoff no greater than $\Lambda_H \sim 10$ TeV, thus predicting a little hierarchy that mirrors the GUT-Planck splitting in the UV. At the scale Λ_H , new dynamics (*e.g.*, SUSY) must appear to keep the Higgs from experiencing sensitivity to even higher scales. Alternatively, the full hierarchy problem can be solved with $N \sim 10^{16}$, so that $\Lambda_H \sim \Lambda_G \sim 10^{10}$ GeV. Note that this number of copies, while sufficient, is unnecessary for a complete solution. There may be two classes of new degrees of freedom: the N copies that participate directly in the N naturalness picture, and another completely sterile set of degrees of freedom that still impact the renormalization of Λ_G .

So far we have described a theory with a S_N permutation symmetry, broken softly by the m_H^2 parameters, such that each of the sectors is SM-like. Sectors for which $m_H^2 < 0$ are similar to our own, with the exception that particle masses scale with the Higgs vev, $v_i \sim v\sqrt{i}$. In addition, once $i \gtrsim 10^8$ the quarks are all heavier than their respective QCD scales. Those sectors do not exhibit chiral symmetry breaking, nor do they contain baryons. Sectors with $m_H^2 > 0$ are dramatically different from ours. In these sectors, electroweak symmetry is broken at low scales due to the QCD condensate Λ_{QCD} . Fermion masses are generated by the four-fermion interactions that are induced by integrating out the complete $SU(2)$ Higgs multiplet. Thus, $m_f \sim y_f y_t \Lambda_{\text{QCD}}^3 / (m_H^2)_i \lesssim 100$ eV, where y_t is the top Yukawa coupling. All fermionic and gauge degrees of freedom are extremely light relative to the ones in our

sector.

With so many additional degrees of freedom, the naive cosmological history is dramatically excluded. In particular, if all sectors have comparable temperatures in the early Universe, then one expects $\Delta N_{\text{eff}} \sim N$ (see Eq. (9)). Thus, the hierarchy problem gets transmuted into the question of how to predominantly reheat only those sectors with a tuned Higgs mass.

To accomplish this, we need to introduce a last ingredient into the story, the “reheaton” field, so named because it is responsible for reheating the Universe via its decays. We call this field S^c for models where the reheaton is a fermion, and ϕ if the reheaton is a scalar. The cosmological history of the model begins in a post-inflationary phase where the energy density of the Universe is dominated by the reheaton. As stated multiple times we can not be unique, therefore we assume that the reheaton couples universally to all sectors. Note that the scalars must be near their true minimum when reheating occurs. This can be accomplished by having either low scale inflation, or else a coupling of the Higgses to the Ricci scalar.

In the next section, we present a set of models in which the reheaton dynamically selects and populates only the lightest sectors, despite preserving the aforementioned softly broken S_N symmetry. Sec. III then provides constraints on these models, and Sec. IV contains our conclusions and highlights potential signals.

II. MODELS

We have argued that the hierarchy problem can be solved by invoking a large number of copies of the SM, along with some dynamical mechanism which dominantly populates the lightest sector with a non-zero Higgs vev. This section details some simple explicit models that realize a viable cosmological history.

As anticipated in the previous section, we imagine that at a post-inflationary stage the energy density of the Universe is dominated by a reheaton that couples universally to all the new sectors. Its decays populate the SM and its copies. The goal is to deposit as much energy as possible into the sector with the smallest Higgs vev. This may be accomplished by arranging the decays of the reheaton such that the branching fraction into the i^{th} sector scales as $\text{BR}_i \sim (m_H)_i^{-\alpha}$ for some positive exponent α . To this end, we construct models that share three features:

- (i) The reheaton is a gauge singlet;
- (ii) It is parametrically lighter than the naturalness cutoff, $m_{\text{reheaton}} \lesssim \Lambda_H/\sqrt{N}$;
- (iii) Its couplings are the most relevant ones possible that involve the Higgs boson of each sector.

is random with some (arbitrary) prior, and then ask statistical questions regarding how often the resulting theory is compatible with observations.

While the requirement of a light reheaton field may appear to require an additional coincidence, it can be easily accommodated in an extra-dimensional picture. In order to couple to all the sectors, the reheaton must be a bulk field. Then, before canonical normalization, its kinetic term carries a factor of N . If the reheaton enjoys a shift symmetry that is respected in the bulk, it will receive a Λ_H -sized mass from each brane on which the shift symmetry is violated. Here we assume that the dynamics above Λ_H respect the shift symmetry. As long as the shift symmetry is only violated on the boundaries, the reheaton mass will be parametrically the same as the weak scale after canonical normalization. In the case of a fermionic reheaton, this simple picture corresponds to the brane-localization of its Dirac partner.

The two simplest models, which we denote ℓ and ϕ , are

$$\mathcal{L}_\ell \supset -\lambda S^c \sum_i \ell_i H_i - m_S S S^c, \quad (2)$$

if the reheaton is a fermion S^c , and

$$\mathcal{L}_\phi \supset -a \phi \sum_i |H_i|^2 - \frac{1}{2} m_\phi^2 \phi^2, \quad (3)$$

if the reheaton is a scalar ϕ . For the theory to be perturbative, we need the coupling λ to obey a 't Hooft-like scaling $\lambda \sim 1/\sqrt{N}$. Naively we would expect the same scaling for a , but we find that a stronger condition needs to be imposed ($a \sim 1/N$) to insure that the loop induced mass for ϕ is not much larger than Λ_H/\sqrt{N} . Even with this scaling, the loop-induced tadpole for ϕ will be too large unless the sign of a is taken to be arbitrary for each sector. Note that a breaks a \mathbb{Z}_2 symmetry on ϕ , so that this choice is consistent with technical naturalness. Including the arbitrary sign, the sum over tadpole contributions only grows as \sqrt{N} , and so the natural range of ϕ is restricted to $\Lambda_H\sqrt{N}$. The Higgses will then receive a contribution to their m_H^2 parameters of order $a\langle\phi\rangle \sim \Lambda_H^2/\sqrt{N}$. While these contributions may be large compared to our weak scale, as long as they are smaller than $\mathcal{O}(\Lambda_H^2)$, they can be safely absorbed into the quadratically-divergent contributions to m_H^2 . Of course, these are upper bounds on the couplings; as we will discuss later in the section, they can be consistently taken smaller, so long as the reheat temperature is sufficiently high.

Before moving on to discuss the details of reheating, we remark on the existence of cross-quartics of the form $\kappa |H_i|^2 |H_j|^2$. Even if these are absent in the UV theory, they will be induced radiatively. After electroweak symmetry breaking in the various sectors, these can potentially affect the spectrum, and so it is critical to the N -naturalness mechanism that they be sufficiently

suppressed. Given an arbitrary, S_N symmetric cross-quartic, κ , the m_H^2 parameters will shift by approximately $-\kappa \Lambda_H^2 N/8 + \mathcal{O}(\kappa^2 N)$, while the mixing effects are subdominant. Thus, the general picture of hierarchical weak scales remains intact so long as $\kappa \lesssim 1/N$.

At a minimum, cross-quartics of this form will be induced gravitationally, regardless of the reheaton dynamics. These quartically-divergent gravitational couplings arise at three loops, giving $(16\pi^2)^3 \kappa_g \sim \lambda_h^2 (\Lambda_H/M_{\text{pl}})^4 \sim (\lambda_h/N)^2 (\Lambda_H/\Lambda_G)^4$, where λ_h is the SM-like Higgs self quartic. Here we have taken the scale that cuts off these divergences to be Λ_H , as would be appropriate for a supersymmetric UV completion (for which these quartics are absent). In either case, these gravitational couplings are parametrically safe, since they scale as $(1/N)^2$.

In addition, potentially dangerous cross-quartics can be generated by reheaton exchange. In the ℓ model, the cross-quartic is generated at one loop: $\kappa_\ell \sim \lambda^4/16\pi^2 \lesssim 1/N^2$, after enforcing the large- N scaling of λ . In the ϕ model, these quartics are generated at tree-level, $\kappa_\phi \sim a^2/m_\phi^2$. Naively this appears borderline problematic, since κ_ϕ scales as $1/N$. However, the arbitrary sign of a , which was necessary to mitigate the tadpole of ϕ , will once again soften the sum over sectors, so that $\sum a_i v_i^2 \sim a \Lambda_H^2 \sqrt{N}$. Combined with the large- N scaling of a , these quartics are rendered safely negligible.

A. Reheating

If the reheaton is sufficiently light, then we may analyze the leading reheaton decay operators using an effective Lagrangian computed by integrating out H_i . This immediately makes it clear why we want the reheaton to be coupled with the most relevant coupling possible, since these will suffer the fastest suppression as $|m_H| \rightarrow \infty$. Integrating out the Higgs and gauge bosons in the ℓ model, the leading decays of S^c are given by, *e.g.*

$$\begin{aligned} \mathcal{L}_\ell^{\langle H \rangle \neq 0} &\supset \mathcal{C}_1^\ell \lambda \frac{v}{m_Z^2 m_S} \nu^\dagger \bar{\sigma}^\mu S^c f^\dagger \bar{\sigma}_\mu f; \\ \mathcal{L}_\ell^{\langle H \rangle = 0} &\supset \mathcal{C}_2^\ell \lambda \frac{y_t}{m_H^2} S \ell Q_3^\dagger u_3^{c\dagger}, \end{aligned} \quad (4)$$

where m_Z is the relevant Z^0 -boson mass and the \mathcal{C}_i^ℓ are numerical coefficients. We have omitted decays through W and Higgs bosons for sectors with $\langle H \rangle \neq 0$ as they scale in the same way. We include them in all numerical computations.

From this low energy Lagrangian we can easily infer that a light reheaton dominantly populates the lightest negative Higgs mass sector. Denoting with m_{h_i} the physical Higgs mass in sectors with $\langle H \rangle \neq 0$, the reheaton decay widths scale as $\Gamma_{m_H^2 < 0} \sim 1/m_{h_i}^2$ and $\Gamma_{m_H^2 > 0} \sim$




FIG. 2: Feynman diagrams for the most important decays in the ϕ model. The left (right) column is for $\langle H \rangle \neq 0$ ($\langle H \rangle = 0$). The top (bottom) row is for $m_\phi \gg |m_H|$ ($m_\phi \ll |m_H|$).

$1/m_{H_i}^4$ in sectors with and without electroweak symmetry breaking, respectively. Thus the reheaton preferentially decays into sectors with light Higgs bosons and non-zero vevs. If, instead, the reheaton were heavy enough to decay directly to on-shell Higgs or gauge bosons, the branching fractions would be democratic into those sectors, and the energy density in our sector would not come to dominate the energy budget of the Universe.

In the scalar case the decays are different, but the scaling of the decay widths is exactly the same. This can be seen once more by integrating out the Higgs and gauge bosons in all the sectors:

$$\begin{aligned} \mathcal{L}_\phi^{\langle H \rangle \neq 0} &\supset \mathcal{C}_1^\phi a y_q \frac{v}{m_h^2} \phi q q^c; \\ \mathcal{L}_\phi^{\langle H \rangle = 0} &\supset \mathcal{C}_3^\phi a \frac{g^2}{16\pi^2} \frac{1}{m_H^2} \phi W_{\mu\nu} W^{\mu\nu}, \end{aligned} \quad (5)$$

where again the \mathcal{C}_i^ϕ are numerical coefficients, and $W_{\mu\nu}$ is the $SU(2)$ field strength. As in the fermionic case, this Lagrangian leads to decay widths that scale as $\Gamma_{m_H^2 < 0} \sim 1/m_{h_i}^2$ and $\Gamma_{m_H^2 > 0} \sim 1/m_{H_i}^4$ in sectors with and without electroweak symmetry breaking, respectively, through the diagrams shown in Fig. 2. We have not included the one-loop decay $\phi \rightarrow \gamma\gamma$ in Eq. (5) for sectors with $\langle H \rangle \neq 0$. This operator scales as $1/m_h^2$ and is important for sectors with $N \gtrsim 10^8$; we find that this is never the leading decay once the bounds on N discussed in Sec. III are taken into account.

Before moving to a more detailed discussion of signals and constraints it is worth pointing out two important differences between the ϕ and ℓ models that will lead us to modify the latter. Given the scaling of the widths we can approximately neglect the contributions to cosmological observables from the $\langle H \rangle = 0$ sectors. In the simple case that the vevs squared are equally spaced, $v_i^2 \sim 2i \times v_{\text{us}}^2$, as in Eq. (1) with $r = 1$, we find that the branching ratio

into the other sectors is $\sum 1/i \sim \log N$.

In the ϕ model, this logarithmic sensitivity to N is not realized. Since the reheaton decays into sectors with non-zero vevs via mixing with the Higgs, the decays become suppressed by smaller and smaller Yukawa couplings as h_i becomes heavy. After the charm threshold is crossed $m_\phi < 2m_{c_i}$ we can neglect the contribution of the new sectors to cosmological observables (with one exception that we discuss in the next section). This behavior is displayed in the left panel of Fig. 3, where we show the fraction of energy density deposited in each sector.

The second important difference is that in the ℓ model the reheaton couples directly to neutrinos and, in the sectors with electroweak symmetry breaking, it mixes with them. This leads to two effects. First, the physical reheaton mass grows with N , implying that the structure of the ℓ model forces the reheaton to be heavy at large N , and can be inconsistent depending on the value of λ . Additionally, this mixing can generate a freeze-in abundance [5] of neutrinos in the other sectors from the process $\nu_{\text{us}} \nu_{\text{us}} \rightarrow \nu_{\text{us}} \nu_i$ via an off-shell Z^0 . Tension with neutrino overclosure and overproduction of hot dark matter leads to an upper bound on the maximum number of sectors. In practice, it is hard to go beyond $N \simeq 10^3$.

However, there is a simple extension of the ℓ model that at once mitigates its UV, *i.e.*, large N , sensitivity and solves the problems arising from a direct coupling to neutrinos. If the reheaton couples to each sector only through a massive portal (whose mass grows with v_i), then the branching ratios will scale with a higher power of the Higgs vev after integrating out the portal states. As an example, consider introducing a 4th generation of vector-like leptons (L_4, L_4^c), (E_4, E_4^c), and (N_4, N_4^c) to each sector. Then relying on softly broken $U(1)$ symmetries, we can couple the reheaton to L_4 only via the Lagrangian

$$\begin{aligned} \mathcal{L}_{L_4} &\supset \mathcal{L}_{\text{mix}} + \mathcal{L}_Y + \mathcal{L}_M, \\ \mathcal{L}_{\text{mix}} &= -\lambda S^c \sum_i (L_4 H)_i - \mu_E \sum_i (e^c E_4)_i, \\ \mathcal{L}_Y &= -\sum_i \left[Y_E (H^\dagger L_4 E_4^c)_i + Y_E^c (H L_4^c E_4)_i \right. \\ &\quad \left. + Y_N (H L_4 N_4^c)_i + Y_N^c (H^\dagger L_4^c N_4)_i \right], \\ \mathcal{L}_M &= -\sum_i \left[M_E (E_4^c E_4)_i + M_L (L_4^c L_4)_i \right. \\ &\quad \left. + M_N (N_4^c N_4)_i \right] - m_S S S^c, \end{aligned} \quad (6)$$

where we have assumed universal masses and couplings across all the sectors for simplicity. We again need $\lambda \sim 1/\sqrt{N}$ for perturbativity. Note that we are assuming that




FIG. 3: Energy density deposited in each sector as a function of sector number, normalized to the energy density in our sector. The left panel is for the ϕ model with $a = 1$ MeV. The right panel is for the L_4 model with $\lambda \times \mu_E = 1$ MeV, $M_L = 400$ GeV, $M_{E,N} = 500$ GeV, $Y_E = Y_N = 0.2$, and $Y_E^c = Y_N^c = -0.5$. The solid lines are the result of a full numerical calculation. The dashed lines show the expected scalings. As discussed in the text, the steps in the ϕ model are proportional to Yukawa couplings due to the fact that ϕ decays via mixing with the Higgs. When $i \gtrsim 10^9$ in the L_4 model, the process $S^c \rightarrow 2e + \nu$ cannot proceed on-shell, which results in the deviation from the naive scaling as denoted by $m_S = 2m_e + m_\nu$. Both figures were made using the zero temperature branching ratios of the reheaton; thermal corrections are under control so long as T_{RH} is smaller than the weak scale in our sector, as discussed at the end of Sec. II.

the bilinear $\mu_E e^c E$ only couples a single flavor of right handed lepton to the new 4th generation fields, in order to avoid flavor violation bounds in the charged lepton sector. The predictions relevant to cosmology (see Fig. 5) are insensitive to the choice of flavor; we choose couplings involving the τ for the additional constraints discussed in Sec. III C below since this choice yields the strongest bounds.

To explore the differences between the L_4 and ℓ models let us again consider the limit in which the reheaton is light. If we integrate out the Higgs and gauge bosons along with the new vector-like leptons, the leading operators for the decays of S^c are given by

$$\begin{aligned} \mathcal{L}_{L_4}^{(H) \neq 0} &\supset C_1^{L_4} \lambda' \frac{g^2}{m_W^2} \left(e^{c\dagger} \bar{\sigma}^\mu S^c \right) \left(f^\dagger \bar{\sigma}_\mu f' \right); \\ \mathcal{L}_{L_4}^{(H) = 0} &\supset C_2^{L_4} \lambda \frac{y_t y_b Y_E M_E \mu_E}{16 \pi^2 m_H^4} \left(e^{c\dagger} \bar{\sigma}^\mu S^c \right) \left(u_3^{c\dagger} \bar{\sigma}_\mu d_3^c \right), \end{aligned} \quad (7)$$

where once more the $C_i^{L_4}$ are numerical coefficients, M_4 is used to represent the physical mass of the relevant heavy lepton, and for convenience we have defined $\lambda'_i \equiv (\lambda v_i^2 \mu_E / M_{4i}^4) f(Y, M)$. Here f is a function of dimension one that depends on the Yukawa couplings and vector-like masses in Eq. (7), but not on the Higgs vev. The M_{4i} masses receive a contribution from v_i that eventually dominates. When this happens S^c decays become suppressed by large powers of the Higgs vev. From the effective Lagrangian above, it is easy to conclude that the widths scale as $\Gamma_{m_H^2 < 0} \sim \text{const}$ for the first few sectors, since M_{4i} is approximately independent of v_i . When the

Yukawa contribution to the masses begins to dominate, such that $M_{4i} \sim v_i$, the scaling becomes $\Gamma_{m_H^2 < 0} \sim 1/v_i^8$. Contributions to observables from the sectors with positive Higgs mass squared are negligible: the decay is both three-body and loop-suppressed, and the width scales as $1/v_i^8$ in all the sectors.

The diagrams that lead to these decays are shown in Fig. 4, and the energy density deposited in each sector is depicted in the right panel of Fig. 3. It is obvious that in this model cosmological observables are sensitive only to the few sectors for which the vector-like masses dominate over the Higgs vev, making it insensitive to the UV. This comes at the price of introducing new degrees of freedom near the weak scale. As we will discuss in the following section, the vector-like masses cannot be arbitrarily decoupled, but they must be large enough to avoid tension with direct searches and the measured properties of our Higgs.

Finally, we end this section by briefly commenting on the presence of an upper bound for the reheating temperature T_{RH} such that the mechanism is preserved. Specifically, T_{RH} should be at most of order of the weak scale. If the temperature were larger, our Higgs mass would be dominated by thermal corrections resulting in a change in the scalings of the branching ratios. Our Higgs would obtain a large positive thermal mass and no longer be preferentially reheated over the other sectors. Noting that

$$T_{RH} \simeq 100 \text{ GeV} \sqrt{\frac{\langle \Gamma_{\text{reheaton}} \rangle_T}{10^{-14} \text{ GeV}}}, \quad (8)$$




FIG. 4: Feynman diagrams for the most important decays in the L_4 model. The left (right) column is for $\langle H \rangle \neq 0$ ($\langle H \rangle = 0$). The top (bottom) row is for $m_S \gg |m_H|$ ($m_S \ll |m_H|$).

where $\langle \Gamma_{\text{reheaton}} \rangle_T$ denotes a thermal average of the reheaton width that incorporates the effect of time dilation. Then Eq. (8) places an upper bound on the couplings of the reheaton. In the ϕ model, the $\phi - h$ mixing angle is bounded to be $\theta_{\phi h} \sim (a v / m_h^2)_{\text{us}} \lesssim 10^{-6}$ ($100 \text{ GeV} / m_\phi$) $^{1/2}$. In the L_4 model, most of the viable region of parameter space predicts on-shell decays to our W boson (see Fig. 5 below). Therefore, the width of S^c is dominated by this two-body decay and the constraint on T_{RH} translates into a rough bound of $\lambda'_{\text{us}} \lesssim 10^{-7}$ when $m_S \simeq 100 \text{ GeV}$. For the benchmark values used for the figures below, this in turn translates into a bound $\lambda \times \mu_E \lesssim 10^{-2} \text{ GeV}$.

Finally, we note that at large N there is a more stringent upper bound on the reheating temperature determined by the perturbativity of λ . Requiring $\lambda \lesssim 4\pi/\sqrt{N}$ and $m_S \sim 100 \text{ GeV}$, we find that it is still possible to reheat to a few GeV even with $N \sim 10^{16}$, where this estimate has been done using the complete numerical implementation of the mixings.

In principle, we must also ensure that other sectors are not overly heated by scattering from our own plasma after reheating. However, the aforementioned constraints on the reheaton couplings sufficiently suppress this contribution to their energy density.

B. Baryogenesis

A viable mechanism for baryogenesis is an even more crucial part of our mechanism for solving the hierarchy problem than in typical natural theories for new physics, where it can be treated in a modular way. One challenge is that our reheating temperature should be near or below the electroweak phase transition. Additionally,

baryogenesis cannot occur in all of the copies of the SM, or there would be too much matter in the Universe.

One simple approach, which makes use of features intrinsic to the model, is to imagine that the reheaton S^c carries a lepton number asymmetry. This asymmetry is distributed to the various sectors through the decays of S^c . Only in the sectors nearest ours is this lepton asymmetry converted into a baryon asymmetry. The small number abundance of baryons results from the low reheat temperature. At temperatures just below the electroweak phase transition, the sphaleron rate is exponentially suppressed, and only a small fraction of the lepton asymmetry is converted into a baryon asymmetry. The baryon asymmetry in sectors with $m_H^2 > 0$ is even further suppressed; since $m_W \lesssim \Lambda_{\text{QCD}}$, the sphalerons remain active at temperatures below the baryon masses. Any asymmetry in these sectors will eventually be redistributed back into the leptons. We have now laid out the necessary ingredients of our mechanism and we are ready to explore their phenomenology in more detail.

III. SIGNALS AND CONSTRAINTS

The signals and experimental constraints for N naturalness come from two sources: mixing between the sectors and energy density deposited in the new sectors by the reheaton decays. The cosmological observables sensitive to the energy density in each sector can be further divided into two categories.

First we discuss measurements that can detect new light particles. These signatures are dominated by the sectors closest to us and can not be avoided by changing the UV scalings of the model. They provide the most characteristic signatures of the theory. Then we study the impact of stable massive particles from the new sectors. This last set of constraints is dominated by sectors with the largest Higgs masses and can be ignored in the L_4 model, where the large i physics is decoupled. In the last two subsections we discuss the bounds arising from mixing between the sectors, followed by possible collider signatures.

A. Massless degrees of freedom

As discussed previously, our models have a large number of massless or nearly massless degrees of freedom. For example, all additional sectors contain photons and neutrinos. There are several kinds of cosmological observations that are sensitive to new relativistic particles. For instance the measurement of the Hubble parameter during either Big Bang Nucleosynthesis or at the epoch

of photon decoupling, and bounds on hot dark matter from the matter power spectrum.

The sensitivity of the expansion of the Universe to new relativistic degrees of freedom is usually phrased in terms of the number of effective neutrinos

$$\Delta N_{\text{eff}} = \frac{1}{\rho_{\nu}^{\text{us}}} \sum_{i \neq \text{us}} \rho_i. \quad (9)$$

Current bounds are $\Delta N_{\text{eff}} \lesssim 1$ during BBN [6] and $\Delta N_{\text{eff}} \lesssim 0.6$ at photon decoupling [7]. In both cases we quote an approximate 95% C.L. constraint. The CMB bound applies to free-streaming radiation [8]. However, the photons in some of the new sectors are still in equilibrium with or have just decoupled from electrons at that time and might be more similar to a perfect fluid. Until recently it was impossible to distinguish between the two types of radiation, as they affect the CMB damping tail in the same way [9]. The detection of a phase shift in the CMB anisotropies [10] has broken this degeneracy, and it is now possible to set a 95% C.L. bound: $N_{\text{fluid}} \lesssim 1$

for $\Delta N_{\text{eff}} = 0$ [11]. Here we have defined N_{fluid} in the same way as ΔN_{eff} , normalizing the energy density of non-free-streaming radiation to that of a neutrino in our sector.

In the following, we do not distinguish between the two types of radiation. We use ΔN_{eff} to denote the sum of the two components. Given the bounds discussed above and the two dimensional exclusions in [11], this is sufficient to show that the model has large areas of parameter space consistent with current data. In the future, it would be interesting to explore CMB observations in more detail, as it is a generic prediction of this type of theories to have roughly comparable amounts of free-streaming and non-free-streaming extra radiation.

Having set N_{fluid} to zero, it is straightforward to estimate the contribution to ΔN_{eff} from our new sectors, since the ratio of energy densities ρ_i/ρ_{us} is determined by the decay widths of the reheaton: $\rho_i/\rho_{\text{us}} \simeq \Gamma_i/\Gamma_{\text{us}}$. For example, assume that the reheaton is lighter than the lightest Higgs across all the sectors; then we have

$$\begin{aligned} \Delta N_{\text{eff}}^{\phi} &\sim \sum_{i=1}^{N_b} \frac{1}{2i+1} + \frac{y_c^2}{y_b^2} \sum_{i=N_b+1}^{N_c} \frac{1}{2i+1} \simeq \frac{1}{2} \left(\log 2N_b + \frac{y_c^2}{y_b^2} \log \frac{N_c}{N_b} \right), & N_{b,c} &= \left(\frac{m_{\phi}^2}{8m_{b,c}^2} - \frac{1}{2} \right), \\ \Delta N_{\text{eff}}^{L_4} &\sim \sum_{i=1}^{N_V} i^0 + \sum_{i=N_V+1}^{N/2} \frac{1}{(2i+1)^4} \simeq N_V, & N_V &\simeq \left(\frac{M^2}{Y^2 v^2} - \frac{1}{2} \right), \end{aligned} \quad (10)$$

where M represents one of the vector-like masses in the L_4 model and Y one of the new Yukawas. In this estimate we have neglected the contribution from $m_H^2 > 0$ sectors and the effect of g_* in each sector, to highlight the scaling of ΔN_{eff} . From this simple exercise we see that ΔN_{eff} is dominated by the bottom of the spectrum. The sectors past $i = N_{b,c}$ or $i = N_V$ receive a negligible fraction of the total energy density and do not contribute to ΔN_{eff} . Using Eq. (10) to go beyond a simple parametric estimate gives results that are in tension with current bounds. For example $m_{\phi} \simeq 50$ GeV implies $N_b \simeq 17$ and $\Delta N_{\text{eff}} \simeq 2$. However, these estimates are only qualitative, and break down in a large fraction of the parameter space of the models. The results from a full numerical computation are shown in Fig. 5.

There are two main messages that can be extracted from this calculation. First, we can satisfy current constraints for a range of reheaton masses up to a few hundred GeV. Second, the models predict values of ΔN_{eff} within the range of sensitivity relevant for CMB stage 4

experiments [1]. These next generation detectors, which should start taking data within the next five years, will probe $\Delta N_{\text{eff}} \gtrsim 0.03$. If no beyond the SM discovery is made, then the only way to suppress this signal is to introduce ‘‘fine tuning,’’ which in the context of these models is the limit $r \lesssim 0.1$. Alternatively, we could imagine alleviating this tension by taking the vector-like masses in the L_4 model far below the weak scale, in potential conflict with electroweak/Higgs measurements.

A few additional features of the ΔN_{eff} calculation are worth discussing. In the L_4 case the plot is valid for a large range of N , namely $30 \lesssim N \lesssim 10^9$. The upper bound is determined by requiring $\lambda \lesssim 4\pi/\sqrt{N}$ and mixing between e^c and the vector-like leptons less than 1%. It is trivial to go beyond $N = 10^9$, and even possible to reach $N = 10^{16}$, by lowering the reheaton coupling – this comes at the expense of an overall decrease in reheating temperature, even though the result for ΔN_{eff} would not change. For $N < 30$, ΔN_{eff} is smaller than shown in the figure. In the ϕ case, the results are more sensitive to N ,




FIG. 5: ΔN_{eff} contours as a function of reheaton mass and the r parameter defined in Eq. (1). $\Delta N_{\text{eff}} \simeq 0.03$ corresponds to the sensitivity of CMB stage 4 experiments. The current upper bound at the CMB epoch is around 0.6. The left panel is for the ϕ model with $a = 1$ MeV. The right panel is for the L_4 model with $\lambda \times \mu_E = 1$ MeV, $M_L = 400$ GeV, $M_{E,N} = 500$ GeV, $Y_E = Y_N = 0.2$, and $Y_E^c = Y_N^c = -0.5$. As discussed in the text, the L_4 result is valid for a large range of N , namely $30 \lesssim N \lesssim 10^9$. Both figures were made using the zero temperature branching ratios of the reheaton; see the end of Sec. II for a discussion.

as shown in Eq. (10). We chose the largest N that is both compatible with overclosure (see the next subsection) and also interesting from a model building perspective, given the relation to the Planck/GUT hierarchy ($N = 10^4$).

The shapes of the ΔN_{eff} contours are easy to explain in terms of kinematics. In L_4 the allowed region corresponds to the reheaton decaying to our sector via a two-body channel, versus a three-body decay into all the other $m_H^2 < 0$ sectors. This is highlighted by the $m_S = m_{W_2}$ line in the plot. In the ϕ model the situation is different. The mixing with the Higgs natu-

rally introduces a number of mass thresholds that reduce ΔN_{eff} . At very low ϕ masses, decays to a pair of b -quarks are kinematically allowed only in our sector. As the ϕ mass increases, the reheaton can mix resonantly with our Higgs and subsequently decay to a pair of W or Z bosons. The last aspect of these results that is not captured by the simple estimate in Eq. (10) is the fact that $(\Delta N_{\text{eff}})_{\text{CMB}} > (\Delta N_{\text{eff}})_{\text{BBN}}$. It is easy to show that this must be the case by appealing to conservation of entropy in each of the sectors. If we compute the ratio of ΔN_{eff} in sector i at the two different epochs, we obtain

$$\frac{(\Delta N_{\text{eff}}^i)_{\text{CMB}}}{(\Delta N_{\text{eff}}^i)_{\text{BBN}}} = \frac{g_*^i(T_{\text{CMB}})}{g_*^i(T_{\text{BBN}})} \left(\frac{g_{*S}^{\text{us}}(T_{\text{BBN}})}{g_{*S}^{\text{us}}(T_{\text{CMB}})} \right)^{4/3} \left(\frac{g_{*S}^i(T_{\text{BBN}})}{g_{*S}^i(T_{\text{CMB}})} \frac{g_{*S}^{\text{us}}(T_{\text{CMB}})}{g_{*S}^{\text{us}}(T_{\text{BBN}})} \right)^{4/3} \simeq \left(\frac{g_{*S}^i(T_{\text{BBN}})}{g_{*S}^i(T_{\text{CMB}})} \right)^{1/3} \geq 1. \quad (11)$$

The first term in the first equality counts the number of relativistic degrees of freedom in sector i at the two different temperatures. The second factor accounts for the fact that neutrinos in our sector are decoupled after BBN, so their temperature during the CMB epoch is

lower than that of photons. The last term comes from entropy conservation in our sector and sector i . In the last equality we have used $g_* \simeq g_{*S}$.

To conclude the discussion of ΔN_{eff} , recall that the result depends almost exclusively on the reheaton branch-

ing ratios and that it is largely insensitive to the value of its overall coupling. A single choice of λ and a is sufficient to understand the complete parameter space. In contrast, the precise value of the vector masses and Yukawa couplings in the L_4 model can change the results considerably, as it is already clear from Eq. (10). When the vector-like masses are around the TeV scale or above, the models are excluded, while $M \simeq 500$ GeV yields predictions that are consistent with current data, as shown in Fig. 5. We leave a more detailed exploration of the parameter space and a discussion of possible collider signatures to future work.

The second class of light particles that can impact our cosmological history are those that are non-relativistic at matter radiation equality, but might have free-streamed enough to suppress the matter power spectrum. Particles that become non-relativistic at a time $t_{\text{NR}} < t_{\text{EQ}}$ suppress structure up to scales $\lambda_{\text{FS}} = c\sqrt{t_{\text{EQ}} t_{\text{NR}}}(2 + \log t_{\text{EQ}}/t_{\text{NR}})/a(t_{\text{EQ}})$. The neutrinos from many of the new sectors would have λ_{FS} larger than one Mpc. At these scales the matter power spectrum can be computed reliably in the linear regime and can be used to infer another upper bound on their energy density. To roughly estimate current constraints we compute the energy density in particles that can suppress structure at one Mpc or above. We find that for Dirac neutrinos the energy density is well below 1% of the total dark matter energy density in all the plane of Fig. 5 for both the ϕ and L_4 models, while for Majorana neutrinos this is true within the $(\Delta N_{\text{eff}})_{\text{CMB}} = 0.5$ contours.

The hot dark matter population may provide another signal. The tower of sterile neutrinos results in a characteristic impact on the matter power spectrum. Furthermore, the hot dark matter signal is primarily determined by the reheaton branching ratios (and hence the spacing between the lightest sectors), so once a value of $\Delta N_{\text{eff}} \neq 0$ is measured it is possible to make predictions for the distortion of the matter power spectrum and vice versa. In general our theories produce non-trivially related modifications in several CMB observables and we leave to future work a more detailed study. Our generic expectation is that neutrino cosmology is modified at the $\mathcal{O}(1)$ level due to slightly heavier albeit less abundant neutrinos in the closest sectors with electroweak symmetry breaking.

B. Massive stable particles

Relic neutrinos account for a fraction $\Omega_\nu^{\text{us}} h^2 \simeq \sum m_\nu(\text{eV})/91.5 \gtrsim 10^{-3}$ of the energy density in the Universe. It is natural to ask if the heavier neutrinos in the sectors with $\langle H \rangle \neq 0$ can lead to overclosure problems. Furthermore, electrons and protons can be sim-

ilarly problematic. This is perhaps surprising, since in the standard picture their symmetric component is completely negligible today. However, in the other sectors their masses are \sqrt{i} larger and subsequently their annihilation cross-sections decrease as $1/i$.²

In all cases, the relic density of the new stable particles comes from two different sources. There is a contribution that grows with i from the sectors where the stable particles reach thermal equilibrium (including a possible freeze-in abundance from our sector) and a second contribution that decreases with i from sectors where the particles never thermalize. Let us focus on this first contribution (for the moment we will neglect the freeze-in abundance from our sector):

$$\Omega h^2 = \frac{s_0}{\rho_c^0} \sum_{i=-1}^{-N_d} m_i Y_i^{\text{fo}} + \dots = a (N_d)^p + \dots \quad (12)$$

Here we use Ωh^2 to indicate the relic density of either neutrinos, electrons or protons; ρ_c^0 is the critical energy density today; s_0 is the entropy density; m_i is the mass of the stable particle; Y_i^{fo} is its yield at freeze-out; N_d is the sector after which the stable particles are not ever in thermal equilibrium with the other particles in their sector; and a and p are positive numbers. In general $a \sim \Omega^{\text{us}} h^2$ and $p > 1$. The reason for $p > 1$ is that $m_i \sim \sqrt{i}$ (or $\sim i$ for Majorana neutrinos) and up to a certain sector number Y_i^{fo} also grows with i , since neutrinos, electrons and protons all freeze-out earlier and earlier.

In the ϕ model this thermal abundance turns out to be the only relevant one. Specifically, electrons and positrons provide the dominant constraint. Once the bound on the reheating temperature is taken into account, the freeze-in abundance from our sector is negligible. Furthermore the overclosure bound on N kicks in before including heavy enough sectors where electrons would not thermalize. Therefore the bound arises only from thermal freeze-out ($n_e^i \sim 1/\langle \sigma_e v \rangle_i$) and it is straightforward to estimate:

$$\begin{aligned} \Omega_e^\phi h^2 &= \sum_{i=1}^{N_{\text{th}}} \frac{m_e^i n_e^i}{\rho_c^0} \simeq \frac{(m_e^{\text{us}} T_0^{\text{us}})^3}{\rho_c^0} \frac{N_{\text{th}}^{5/2}}{M_{\text{pl}} v_{\text{us}} \alpha^2} \\ &\lesssim 0.1 \times \Omega_{\text{DM}} h^2 \implies N_\phi \lesssim 10^5, \end{aligned} \quad (13)$$

where the sum runs up to the heaviest sector where the electrons have thermalized as denoted by N_{th} , T_0^{us} is the photon temperature in our sector today, m_e^{us} is our electron mass, M_{pl} is the Planck mass, and all other quan-

² For protons, this scaling is only valid once the quark masses exceed Λ_{QCD} ; the scaling is slower for the nearer sectors.

tities were defined previously. For this estimate we have assumed that our sector dominates the energy density of the Universe when electron-positron annihilations freeze-out, *i.e.*, the ΔN_{eff} constraint is satisfied. Furthermore we have conservatively assumed freeze-out happens just after reheating (at $T_{\text{RH}}^{\text{us}} \simeq v_{\text{us}}$) in all the sectors. Finally, note that we have required that electrons and positrons make up only 10% of dark matter, the rough bound for particles that behave very differently from cold, collisionless dark matter. To be more conservative we could require them to make up only 1% of dark matter, which would reduce the maximum allowed value of N_ϕ by 60%, still leaving open $N_\phi = 10^4$.

To conclude this section we note that the rapid scaling of the energy density with sector number in L_4 protects the model from overclosure. This implies that N can be taken all the way to 10^{16} and still be consistent with data, at the expense of a low reheat temperature.

C. Mixing Between Sectors

Upon integrating out the reheaton, the low energy theory will contain cross-couplings between the sectors. Stringent bounds from stellar and supernova cooling place limits on the size of these mixings.

a. $\epsilon_i \mathbf{F}_{\mu\nu} \mathbf{F}_i^{\mu\nu}$ In the presence of kinetic mixing, the electrically charged particles of other sectors will have milli-charge couplings to our photon. The most stringent bound on this coupling is derived from energy loss in stars [12, 13]. In sectors with $m_H^2 > 0$, the charged particles are all extremely light – much lighter than stellar temperatures – so that democratic kinetic mixing leads to $\mathcal{O}(N)$ accessible final states for plasmon decay. Thus we require $\sqrt{\sum_i \epsilon_i^2} \lesssim 10^{-14}$, in which ϵ_i is the coefficient of kinetic mixing between our photon and that of sector i , and i runs over all sectors with $m_H^2 > 0$.

Accordingly, there must be no bi-fundamental matter in the UV, the inclusion of which would generate kinetic mixing at one loop. Even in the absence of those states, kinetic mixing may be generated in the IR through the coupling to the reheaton.³ In this case the bounds may be easily avoided by the smallness of the coupling. As described in Sec. II, the portal couplings must decrease with increasing N in order to have a consistent large- N limit.

³ Kinetic mixing is not generated at any order if the coupling between sectors is mediated by a single, real scalar field, since there can be no effective coupling of such a field to the electromagnetic field-strength tensor. Therefore, this effect can be safely neglected in the ϕ model.

For example, in the L_4 model, we must have $\lambda \sim \lambda_0/\sqrt{N}$. The kinetic mixing parameter is generated only at three loops with four powers of the portal coupling:

$$\epsilon_i \sim \frac{\alpha}{4\pi} \left(\frac{\lambda_0}{4\pi} \right)^4 \frac{1}{N^2 i}. \quad (14)$$

Note that ϵ_i decreases with i due to the scaling of $(m_H^2)_i$, and so kinetic mixing is dominated by the sectors nearest to our own. Then the stellar bounds may be avoided as long as $\lambda_0/4\pi \lesssim 10^{-3}\sqrt{N}$, and no suppression is required for $N \gtrsim 10^6$ beyond the natural large- N scaling.

b. $\epsilon_i^n \nu_i^\dagger \bar{\sigma}^\mu \mathbf{D}_\mu \nu$ At one loop in the L_4 model, the reheaton mass-mixes with neutrinos. After integrating out the reheaton, this induces kinetic mixing between neutrinos of different sectors. However, because the vector-like leptons only couple to the charged leptons, the effective coupling to the neutrinos is Yukawa-suppressed:

$$\epsilon_i^n \sim (m_\ell)_{\text{us}} (m_\ell)_i \left(\frac{\lambda Y_E^c \mu_E}{16 \pi^2 (M_A)_i m_S} \right)^2. \quad (15)$$

For sectors with $(M_A)_i \sim M_L$, $\epsilon_i \sim \sqrt{i}$. Once $Y_E v_i \gtrsim M_L$, the kinetic mixing decreases as $\epsilon_i \sim 1/\sqrt{i}$.

Energy loss in SN1987a [14] limits the size of the kinetic mixing. The neutrino production rate from neutral-current bremsstrahlung requires $\sqrt{\sum_i (\epsilon_i^n)^2} \lesssim 10^{-4}$. Due to the growth of ϵ_i^n with i for small values of i , the sum is dominated by those sectors for which the vector-like lepton masses are larger than their chiral masses. For typical parameters, such as those shown in Fig. 5, this is the case for only $\mathcal{O}(10)$ sectors. Taking into account the bound on T_{RH} from Sec. II, this gives $\sqrt{\sum_i (\epsilon_i^n)^2} \lesssim 10^{-13} (M_L/4\pi v)^4$, so that there is no constraint as long as the vector-like masses are taken to be sufficiently close to the weak scale.

c. $\epsilon_i^c \mathbf{G}_F (\nu_i^\dagger \bar{\sigma}^\mu e^c) (p^\dagger \bar{\sigma}_\mu n)$ There is a somewhat more powerful constraint from SN1987a due to charge-current neutrino production. The mass-mixing of the reheaton with neutrinos leads to an effective four-Fermi operator, with

$$\epsilon_i^c \sim (m_\tau m_\nu)_i \left(\frac{\lambda Y_E^c v_{\text{us}} \mu_E M_L}{4\pi m_S (M_A^2)_{\text{us}} (M_A)_i} \right)^2. \quad (16)$$

In the case of Majorana neutrino masses, ϵ_i^c grows like \sqrt{i} , so that $\sqrt{\sum_i (\epsilon_i^c)^2} \sim N$. Once again taking into account the limit on T_{RH} , we have $\sqrt{\sum_i (\epsilon_i^c)^2} \lesssim 10^{-24} N (M_L/\sqrt{4\pi Y_E} v)^4$. The supernova bound is only $\sim 10^{-5}$, so that even for $N \sim 10^{16}$, the coupling is unconstrained for M_L near the weak scale.

Finally, limits on active-sterile neutrino oscillations can also bound ϵ_i^c , both from cosmological measurements as

well as active neutrino disappearance [15]. However, due to the Yukawa suppression of the neutrino mixing, the most relevant limits are those involving the tau neutrino, which are comparatively weak. Absent resonant mixing due to accidental degeneracies, which we expect to be atypical in our parameter space, the bounds from neutrino oscillations are negligible.

D. Colliders

Models of N naturalness can provide collider signatures through both direct production of the reheaton as well as rare decays of SM particles. However, the smallness of the reheaton couplings, due to both large- N suppression and T_{RH} constraints, precludes these signatures from being a generic feature of our models. In the ϕ model, for example, rare Higgs decays proceed through $\phi - h$ mixing. The dominant signature in this case is invisible decays of the SM-like Higgs boson, with $\text{BR}_{\text{inv}} \sim \theta_{\phi h}^2 \Delta N_{\text{eff}} / (1 + \Delta N_{\text{eff}}) \lesssim 10^{-12}$ after requiring sufficiently low T_{RH} . Even using optimistic estimates, future colliders such as TLEP or a 100 TeV machine (with 10 ab^{-1} of luminosity) will only produce $\sim 10^6$ [16] or 10^{10} [3, 4] Higgses, respectively, rendering such decays unobservable.

Direct production of ϕ is similarly suppressed, since it must proceed through the same mixing angle. Even for $m_\phi < m_h$, in which case the production cross sections are somewhat larger, direct production of ϕ will be unobservable. For example, a SM-like Higgs with a mass of 10 GeV gives a cross section only 2.5 times larger at TLEP and approximately 14 times larger at a 100 TeV $p-p$ machine than a Higgs at 125 GeV. Nevertheless, if ϕ is sufficiently light, there may be a variety of interesting signatures; see, e.g., [17] for a study of current constraints which probe mixing angles down to $\theta_{\phi h} \sim 10^{-5}$. We leave a detailed study of detection prospects for $m_\phi \lesssim 10$ GeV to future work, see e.g., [18] for an analysis of a similar scenario.

In the L_4 model, the new sectors and the reheaton are similarly difficult to observe; however, the vector-like leptons of our own sector may be accessible. As discussed above, the vector-like mass parameters should all be of order the weak scale. This implies they would likely be observable, both directly and through $h \rightarrow \gamma\gamma$ and precision electroweak measurements. For a recent study of these bounds, see [19]. In particular, attempting to evade these constraints by raising the mass of the vector-like leptons will re-introduce tension with ΔN_{eff} , as described in the Section III A.

E. A Heavy Axion

An outstanding puzzle within the SM is the strong CP problem. If we have N copies of the SM, then naively they all have their own theta angles θ_i . However, if the S_N symmetry is only softly broken by Higgs mass terms, then all of these angles would be equal. A shared axion would be able to set to zero all the θ_i 's at the same time. The only difference between the version proposed here and the single sector story is that here there are N contributions to its mass from each Λ_{QCD} . The potential for the axion has three contributions

$$V(a) \ni \begin{cases} \frac{\Lambda_{\text{QCD},i}^6}{m_{H,i}^2} \left(\frac{a}{f_a} - \theta_i\right)^2 & \text{for } i < 0 \\ m_{\pi,i}^2 f_{\pi,i}^2 \left(\frac{a}{f_a} - \theta_i\right)^2 & \text{for } 0 \leq i < N_u \\ \Lambda_{\text{QCD},i}^4 \left(\frac{a}{f_a} - \theta_i\right)^2 & \text{for } i \geq N_u \end{cases} \quad (17)$$

where $\Lambda_{\text{QCD},i}$ is the QCD scale for the i^{th} sector, $i < 0$ corresponds to sectors with $\langle H_i \rangle = 0$, and $N_u \sim 10^5$ is the sector with the smallest vev for which $m_u > \Lambda_{\text{QCD}}$. The contribution from the sectors with $\langle H_i \rangle = 0$ are due to higher dimensional operators from integrating out the Higgs doublet, the sectors with $m_u < \Lambda_{\text{QCD}}$ yield the familiar contribution to the axion potential, and the final term is the result for pure QCD with no light quarks. Numerically, the first term can always be neglected, the second term dominates as long as $N < N_u$, and only the third term is relevant for $N \gg N_u$.

In order to estimate how much heavier this state will be as compared to the standard case, we have calculated Eq. (17) numerically including the one-loop running of Λ_{QCD} . For the first two sums in Eq. (17), we used chiral perturbation theory to calculate the contribution to the axion mass. For the last sum, we normalized it such that it is equal to the chiral perturbation theory result when $m_u = \Lambda_{\text{QCD}}$. A numerical fit to the axion mass gives approximately

$$\frac{m_a(N)}{m_a(1)} \simeq \begin{cases} 4 \times 10^3 \left(\frac{N}{10^4}\right)^1 & \text{for } N < N_u \\ 2 \times 10^{14} \left(\frac{N}{10^{16}}\right)^{0.9} & \text{for } N \gg N_u \end{cases} \quad (18)$$

It is critical that the soft-breaking of the S_N symmetry by the different Higgs vevs does not lead to any issues via higher dimensional operators. For example, one class of operators that leads to a change in θ_i between the sectors are

$$\mathcal{O}_{\Delta\theta_{\text{QCD}}} \sim Y_u H_i Q_i u_i^c \frac{|H^i|^2}{\Lambda_G^2}. \quad (19)$$

Because the different sectors have different Higgs vevs, a

chiral rotation shows that the theta angles all differ by $\sim |H_i|^2/\Lambda_G^2$. Plugging this into Eq. (17), solving for the axion vev, and requiring that our theta angle is smaller than 10^{-10} , we find that $N < 10^{10}$ if a shared axion is the solution to the strong CP problem. This approach requires the important assumption that whatever resolves the hierarchy problem between Λ_H and Λ_G does not introduce these operators or any other Higgs dependent phases.

IV. DISCUSSION

In this paper we have proposed a new solution to the hierarchy problem. The need for a huge integer N is obviously the least appealing feature of our setup. It is perhaps not entirely unreasonable to have the mild $N \sim 10^4$ compatible with the existence of a supersymmetric GUT scale, but this seems outlandish in the $N \sim 10^{16}$ limit. At the moment it is difficult to see how such a large integer can be explained dynamically, in the same way as we usually explain hierarchies by, *e.g.* dimensional transmutation. On the other hand, this is simply another large set of degrees of freedom, and we do not deeply understand where the even vaster number of degrees of freedom in a macroscopic expanding universe comes from, so perhaps the large N may eventually find a different sort of natural explanation. The theoretical consistency of the proposal also makes a number of demands on the UV theory, such as the absence of sizable cross-couplings between the sectors, which may be technically natural but may again strain credulity. However, we find it fascinating that huge values of N are experimentally viable. This is highly non-trivial, and indeed in the simplest models we did find significant constraints on N . While we have examined all the zeroth-order phenomenological constraints we know of, it is important to continue to look for constraints on (and signals of!) the scenarios with high values of N ($\gg 10^4$).

It is also interesting to compare N naturalness with other approaches. It bears a superficial resemblance to large extra dimensions, which add 10^{32} degrees of freedom in the form of KK gravitons, as well as the scenario of Dvali [20] which invokes 10^{32} copies of the SM. In each of these cases, M_{pl} is renormalized down to the TeV scale. Of course this predicts (as yet unseen) new particles accessible to the LHC [21]. By contrast, N naturalness solves the hierarchy problem with cosmological dynamics; the weak scale is parametrically removed from the cutoff, and so it does not demand new physics to be accessible at colliders.

N naturalness has some features in common with low-energy SUSY as well. Both models invoke a softly bro-

ken symmetry: SUSY is broken by soft terms, and the S_N symmetry is broken by varying Higgs masses. Also in both cases, the most obvious implementations of the idea are experimentally excluded. If SUSY is directly broken in the MSSM sector, we have the famous difficulties with charge and color breaking; in the case of N naturalness, direct reheating of all N sectors is grossly excluded by N_{eff} . Thus in both cases we need to have “mediators.” SUSY must be dominantly broken in another sector and have its effects mediated to the MSSM. Similarly, reheating must be dominantly communicated to the reheaton, which subsequently dumps its energy density into the other sectors. Finally, both models have additional scales that are not, on the face of it, tied to the physics responsible for naturalness. In SUSY there is a “ μ problem” in that the vector-like Higgsino mass must be comparable to the soft scalar masses, while in N naturalness the reheaton mass must be close to the bottom of the spectrum of Higgs masses. While in both cases there are simple pictures for how this can come about, these coincidences do not emerge automatically.

Moving beyond purely field theoretic mechanisms, there is the recent proposal of the relaxion [22], which invokes an extremely long period of inflation coupled with axionic dynamics to relax to a low weak scale. While both the relaxion and N naturalness mechanisms are cosmological, the physical mechanism of the relaxation, associated with the huge number of e -foldings of inflation, is *in principle* unobservable given our current accelerating Universe, much like the vast regions of the multiverse outside our cosmological horizon are imperceptible. By contrast, the cosmological dynamics associated with reheaton decay in N naturalness are sharply imprinted on the particle number abundance in all the sectors. They are not only in principle observable but, as we have stressed (at least for a small number of sectors “close” to ours), are detectable in practice within our Universe.

It is also interesting to contrast N naturalness with the picture of an eternally inflating multiverse, with environmental selection explaining the smallness of the cosmological constant, as well as potentially at least part of the hierarchy problem. This picture is, after all, the first cosmological approach to fine-tuning puzzles. While it is very far from well-understood and has yet to make internal theoretical sense, it is the only cartoon we have for understanding the cosmological constant problem and does not involve any model-building gymnastics. Furthermore, fine-tuning for the Higgs mass also has a plausible environmental explanation. Especially in the context of minimal split SUSY [23], these ideas give us a picture which simultaneously accounts for the apparent fine-tuning of the cosmological constant and the Higgs mass, while maintaining the striking quantitative successes of

natural SUSY theories in the form of gauge coupling unification and dark matter. Nonetheless, it is important to continue to look for alternatives, minimally as a foil to the landscape paradigm. N naturalness is a concrete example of an entirely different cosmological approach to tuning puzzles, and in particular relies on the existence of only a single vacuum.

We note that there is no obstacle to augmenting N naturalness with an anthropic solution to the cosmological constant problem. The presence of extra sectors exponentially increases the number of available vacua. For example we could add to the SM a sector with m vacua and end up with m^N . Already $N \simeq 10^4$ with two vacua per sector is more than enough to scan the cosmological constant without relying on string theory landscapes. When solving the entire hierarchy problem with $N \simeq 10^{16}$, the vacua utilized to scan the cosmological constant can even be the two minima of the Higgs potential; this requires a high cutoff so that the second minimum is below Λ_H and the difference in the potential energy of the two minima is $\mathcal{O}(\Lambda_G)$.

To conclude, we would like to comment on the nature of the signals that we have discussed in this paper. For concreteness, three models that make N naturalness cosmologically viable were presented. However, it is easy to imagine a broader class of theories that realizes the same mechanism. We can relax the assumption that the Higgs masses are uniformly spaced (or even pulled from a uniform distribution) or that all the new sectors are exact copies of the SM. It is also possible to construct different models of reheating, with new physics near the weak scale to modify the UV behavior of the theory.

Nonetheless our sector can not be special in any way. There will always be a large number of other sectors with massless particles and with matter and gauge contents similar to ours, leading to the following signatures:

- We expect extra radiation to be observable at future CMB experiments.

- The neutrinos in the closest $m_H^2 < 0$ sectors are slightly heavier and slightly less abundant than ours. This implies $\mathcal{O}(1)$ changes in neutrino cosmology, which will start to be probed at this level in the next generation of CMB experiments [2].
- If the strong CP problem is solved by an axion, its mass will be much larger than the standard prediction.
- If $N \lesssim 10^4$ as motivated by grand unification, supersymmetry or new natural dynamics should appear beneath 10 TeV.

The natural parameter space is being probed now, and soon we may know if the N naturalness paradigm explains how the hierarchy problem has been solved by nature.

Acknowledgments

We are grateful to Bob Holdom, Jared Kaplan, Marilena LoVerde and Kris Sigurdson for useful discussions. We also want to thank Arka Banerjee and Neal Dalal for a preliminary investigation of the neutrino cosmology. NAH is supported in part by U.S. Department of Energy grant de-sc0009988. TC is supported by an LHC Theory Initiative Postdoctoral Fellowship, under the National Science Foundation grant PHY-0969510. RTD acknowledges support from the Institute for Advanced Study Marvin L. Goldberger Membership and DOE grant de-sc0009988. AH is supported by the Department of Energy grant DE-SC0012012 and National Science Foundation grant 1316699. HDK is supported by National Research Foundation of Korea(NRF) No. 0426-20140009 and No. 0409-20150110. TC and AH acknowledge the Aspen Center for Physics where this work was completed, which is supported by National Science Foundation grant PHY-1066293.

-
- [1] S. Dodelson *et al.*, “Working Group Report: Dark Energy and CMB,” in *Community Summer Study 2013: Snowmass on the Mississippi (CSS2013) Minneapolis, MN, USA, July 29-August 6, 2013*. 2013. [arXiv:1309.5386 \[astro-ph.CO\]](#).
- [2] R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, “Towards a cosmological neutrino mass detection,” *Phys. Rev.* **D92** (2015) no. 12, 123535, [arXiv:1509.07471 \[astro-ph.CO\]](#).
- [3] N. Arkani-Hamed, T. Han, M. Mangano, and L.-T. Wang, “Physics Opportunities of a 100 TeV Proton-Proton Collider,” [arXiv:1511.06495 \[hep-ph\]](#).
- [4] T. Golling *et al.*, “Physics at a 100 TeV pp collider: beyond the Standard Model phenomena,” [arXiv:1606.00947 \[hep-ph\]](#).
- [5] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, “Freeze-In Production of FIMP Dark Matter,” *JHEP* **03** (2010) 080, [arXiv:0911.1120 \[hep-ph\]](#).
- [6] R. Cooke, M. Pettini, R. A. Jorgenson, M. T. Murphy, and C. C. Steidel, “Precision measures of the primordial abundance of deuterium,” *Astrophys. J.* **781** (2014) no. 1, 31, [arXiv:1308.3240 \[astro-ph.CO\]](#).
- [7] **Planck** Collaboration, P. A. R. Ade *et al.*, “Planck 2015 results. XIII. Cosmological parameters,”

- arXiv:1502.01589 [astro-ph.CO].
- [8] N. F. Bell, E. Pierpaoli, and K. Sigurdson, “Cosmological signatures of interacting neutrinos,” *Phys. Rev.* **D73** (2006) 063523, arXiv:astro-ph/0511410 [astro-ph].
- [9] Z. Hou, R. Keisler, L. Knox, M. Millea, and C. Reichardt, “How Massless Neutrinos Affect the Cosmic Microwave Background Damping Tail,” *Phys. Rev.* **D87** (2013) 083008, arXiv:1104.2333 [astro-ph.CO].
- [10] B. Follin, L. Knox, M. Millea, and Z. Pan, “First Detection of the Acoustic Oscillation Phase Shift Expected from the Cosmic Neutrino Background,” *Phys. Rev. Lett.* **115** (2015) no. 9, 091301, arXiv:1503.07863 [astro-ph.CO].
- [11] D. Baumann, D. Green, J. Meyers, and B. Wallisch, “Phases of New Physics in the CMB,” *JCAP* **1601** (2016) 007, arXiv:1508.06342 [astro-ph.CO].
- [12] S. Davidson, B. Campbell, and D. C. Bailey, “Limits on particles of small electric charge,” *Phys. Rev.* **D43** (1991) 2314–2321.
- [13] S. Davidson, S. Hannestad, and G. Raffelt, “Updated bounds on millicharged particles,” *JHEP* **05** (2000) 003, arXiv:hep-ph/0001179 [hep-ph].
- [14] G. Raffelt and D. Seckel, “Bounds on Exotic Particle Interactions from SN 1987a,” *Phys. Rev. Lett.* **60** (1988) 1793.
- [15] M. Cirelli, G. Marandella, A. Strumia, and F. Vissani, “Probing oscillations into sterile neutrinos with cosmology, astrophysics and experiments,” *Nucl. Phys.* **B708** (2005) 215–267, arXiv:hep-ph/0403158 [hep-ph].
- [16] **TLEP Design Study Working Group** Collaboration, M. Bicer *et al.*, “First Look at the Physics Case of TLEP,” *JHEP* **01** (2014) 164, arXiv:1308.6176 [hep-ex].
- [17] J. D. Clarke, R. Foot, and R. R. Volkas, “Phenomenology of a very light scalar (100 MeV $\leq m_h \leq$ 10 GeV) mixing with the SM Higgs,” *JHEP* **02** (2014) 123, arXiv:1310.8042 [hep-ph].
- [18] G. Krnjaic, “Probing Light Thermal Dark-Matter With a Higgs Portal Mediator,” arXiv:1512.04119 [hep-ph].
- [19] W. Altmannshofer, M. Bauer, and M. Carena, “Exotic Leptons: Higgs, Flavor and Collider Phenomenology,” *JHEP* **01** (2014) 060, arXiv:1308.1987 [hep-ph].
- [20] G. Dvali, “Black Holes and Large N Species Solution to the Hierarchy Problem,” *Fortsch. Phys.* **58** (2010) 528–536, arXiv:0706.2050 [hep-th].
- [21] G. Dvali and M. Redi, “Phenomenology of 10^{32} Dark Sectors,” *Phys. Rev.* **D80** (2009) 055001, arXiv:0905.1709 [hep-ph].
- [22] P. W. Graham, D. E. Kaplan, and S. Rajendran, “Cosmological Relaxation of the Electroweak Scale,” *Phys. Rev. Lett.* **115** (2015) no. 22, 221801, arXiv:1504.07551 [hep-ph].
- [23] N. Arkani-Hamed and S. Dimopoulos, “Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC,” *JHEP* **06** (2005) 073, arXiv:hep-th/0405159 [hep-th].