Solving the Hierarchy Problem at Reheating with a Large Number of Degrees of Freedom

Nima Arkani-Hamed, Timothy Cohen, Raffaele Tito D'Agnolo, Anson Hook, Hyung Do Kim, and David Pinner

Phys. Rev. Lett. **117**, 251801 — Published 15 December 2016

DOI: 10.1103/PhysRevLett.117.251801

Solving the Hierarchy Problem at Reheating

Nima Arkani-Hamed,[1] Timothy Cohen,[2] Raffaele Tito D'Agnolo,[1] Anson Hook,[3] Hyung Do Kim,[4] and David Pinner[5]

[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 mechanism to stabilize the electroweak hierarchy. 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 $m_W \times M_{\rm pl}/M_{\rm GUT} \sim 10$ TeV.

I. MECHANISM

This letter describes a new mechanism, dubbed "N-naturalness," 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 $m_W \times M_{\rm pl}/M_{\rm GUT} \sim 10$ 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, $\left|m_H^2\right|_{\rm min} \sim \Lambda_H^2/N$. 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 distribution of mass squared parameters, indexed by an integer label i such that

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

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

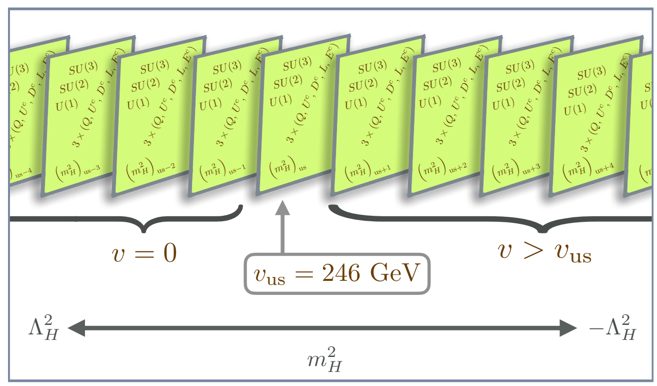

FIG. 1: A sketch of the N-naturalness setup. The sectors have been ordered so that they range from $m_H^2 \sim \Lambda_H^2$ to $-\Lambda_H^2$. The sector with the smallest vacuum expectation value contains our copy of the SM.

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

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 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 $\Lambda_G^2 \sim M_{\rm pl}^2/N$. If perturbative gauge coupling unification is to be preserved $\Lambda_G \gtrsim M_{\rm GUT}$, 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 $\Lambda_{\rm QCD}$. Fermion masses are generated by the four-fermion interactions that are induced by integrating out the complete $SU(2)$ Higgs multiplet. Thus, $m_f \sim y_f\, y_t\, \Lambda_{\rm QCD}^3/\left(m_H^2\right)_i \lesssim 100$ 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 $\Delta N_{\rm eff} \sim N$ (see Eq. (**??**) in the supplemental material). 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. Constraints on these models are provided as Supplemental Materials [5], and Sec. III 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 $i^{\rm th}$ sector scales as ${\rm BR}_i \sim (m_H)_i^{-\alpha}$ 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, $m_{\rm reheaton} \lesssim \Lambda_H/\sqrt{N}$;

(iii) Its couplings are the most relevant ones possible that involve the Higgs boson of each sector.

---

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

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

if the reheaton is a fermion S^c, and

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

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 $\Lambda_H/\sqrt{N}$. 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 $(16\,\pi^2)^3 \kappa_g \sim \lambda_h^2 (\Lambda_H/M_{\rm pl})^4 \sim (\lambda_h/N)^2 (\Lambda_H/\Lambda_G)^4$, 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 we want the reheaton to be coupled with the most relevant coupling possible, since these will suffer the fastest suppression as $|m_H| \to \infty$. Integrating out the Higgs and gauge bosons in the ℓ model, the leading decays of S^c are given by, *e.g.*

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

where m_Z is the relevant Z^0-boson mass and the $\mathcal{C}_i^\ell$ 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 $\Gamma_{m_H^2>0} \sim$

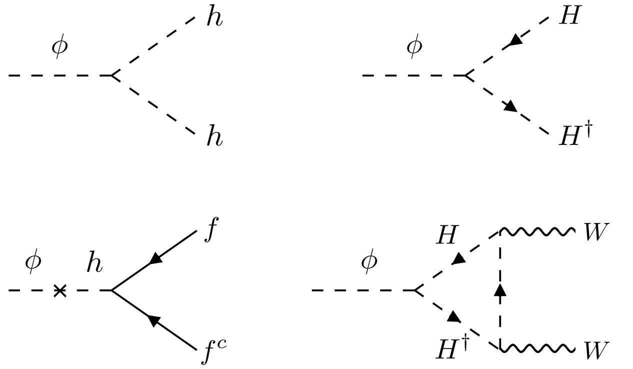

FIG. 2: Feynman diagrams for the most important decays in the ϕ model. The left (right) column is for $\langle H\rangle \neq 0$ $\big(\langle H\rangle = 0\big)$. The top (bottom) row is for $m_\phi \gg |m_H|$ $\big(m_\phi \ll |m_H|\big)$.

$1/m^4_{H_i}$ in sectors with and without electroweak symmetry breaking, respectively. Thus the reheaton preferentially decays into sectors with light Higgs bosons and nonzero 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:

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

where again the $\mathcal{C}_i^\phi$ 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 $\Gamma_{m_H^2<0} \sim 1/m^2_{h_i}$ and $\Gamma_{m_H^2>0} \sim 1/m^4_{H_i}$ in sectors with and without electroweak symmetry breaking, respectively, through the diagrams shown in Fig. 2. We have not included the one-loop decay $\phi \to \gamma\,\gamma$ 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 the Supplemental Materials [5] 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, $v_i^2 \sim 2\,i \times v_{\rm us}^2$, 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 nonzero 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 $m_\phi < 2\, m_{c_i}$ 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 [6] of neutrinos in the other sectors from the process $\nu_{\rm us}\,\nu_{\rm us} \to \nu_{\rm us}\,\nu_i$ 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 $4^{\rm th}$ 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

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

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 as-

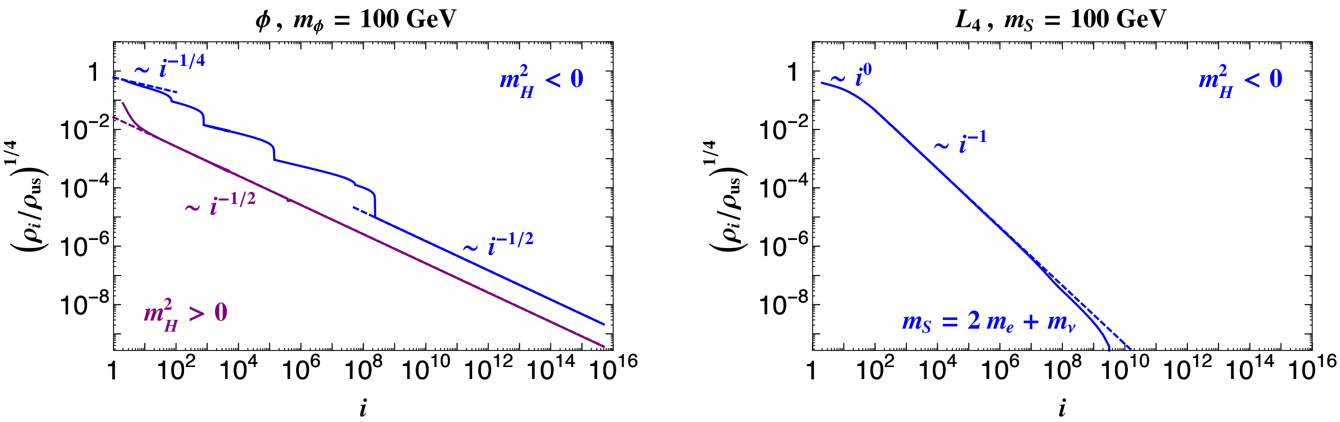

FIG. 3: Energy density deposited in each sector as a function of sector number, normalized to the energy density in our sector. The left panel is for the ϕ model with $a = 1$ MeV. The right panel is for the L_4 model with $\lambda \times \mu_E = 1$ MeV, $M_L = 400$ GeV, $M_{E,N} = 500$ GeV, $Y_E = Y_N = 0.2$, 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 $S^c \to 2\,e + \nu$ cannot proceed on-shell, which results in the deviation from the naive scaling as denoted by $m_S = 2\,m_e + m_\nu$. Both figures were made using the zero temperature branching ratios of the reheaton; thermal corrections are under control so long as $T_{\rm RH}$ is smaller than the weak scale in our sector, as discussed at the end of Sec. II.

suming that the bilinear $\mu_E\, e^c\, E$ only couples a single flavor of right handed lepton to the new $4^{\rm th}$ generation fields, in order to avoid flavor violation bounds in the charged lepton sector. The predictions relevant to cosmology (see Fig. **??** in the supplemental material) are insensitive to the choice of flavor; we choose couplings involving the τ for the additional constraints discussed in Sec. **??** in the supplemental material 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

$$\begin{aligned}
&\mathcal{L}_{L_4}^{\langle H\rangle \neq 0} \supset \mathcal{C}_1^{L_4} \lambda' \, \tfrac{g^2}{m_W^2} \left(e^{c\dagger} \bar{\sigma}^\mu S^c\right)\left(f^\dagger \bar{\sigma}_\mu f'\right); \\
&\mathcal{L}_{L_4}^{\langle H\rangle = 0} \supset \mathcal{C}_2^{L_4} \lambda \, \tfrac{y_t\, y_b}{16\,\pi^2} \, \tfrac{Y_E\, M_E\, \mu_E}{m_H^4} \left(e^{c\dagger} \bar{\sigma}^\mu S^c\right)\left(u_3^{c\dagger} \bar{\sigma}_\mu d_3^c\right),
\end{aligned} \tag{7}$$

where once more the $\mathcal{C}_i^{L_4}$ are numerical coefficients, M_4 is used to represent the physical mass of the relevant heavy lepton, and for convenience we have defined $\lambda'_i \equiv \left(\lambda\, v_i^2\, \mu_E / M_{4i}^4\right) f(Y, M)$. 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 $\Gamma_{m_H^2<0} \sim$ 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 $T_{\rm RH}$ such that the mechanism is preserved. Specifically, $T_{\rm RH}$ 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

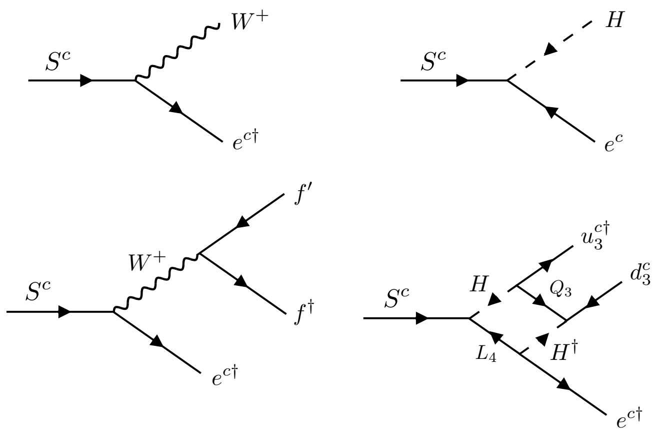

FIG. 4: Feynman diagrams for the most important decays in the L_4 model. The left (right) column is for $\langle H\rangle \neq 0$ $\big(\langle H\rangle = 0\big)$. The top (bottom) row is for $m_S \gg |m_H|$ $\big(m_S \ll |m_H|\big)$.

that

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

where $\langle \Gamma_{\rm reheaton}\rangle_T$ 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 $\theta_{\phi h} \sim \big(a\, v/m_h^2\big)_{\rm us} \lesssim 10^{-6}$ $(100 \text{ GeV}/m_\phi)^{1/2}$. In the L_4 model, most of the viable region of parameter space predicts on-shell decays to our W boson (see Fig. **??** in the supplemental material). Therefore, the width of S^c is dominated by this two-body decay and the constraint on $T_{\rm RH}$ translates into a rough bound of $\lambda'_{\rm us} \lesssim 10^{-7}$ when $m_S \simeq 100$ GeV. For the benchmark values used for the figures below, this in turn translates into a bound $\lambda \times \mu_E \lesssim 10^{-2}$ 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 $m_S \sim 100$ 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 $m_W \ll \Lambda_{\rm QCD}$, the sphalerons remain active at temperatures far 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. 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 [7] which invokes 10^{32} copies of the SM. In each of these cases, $M_{\rm pl}$ is renormalized down to the TeV scale. Of course this predicts (as yet unseen) new particles accessible to the LHC [8]. 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 broken 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 $N_{\rm eff}$. 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 [9], 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 [10], 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 uniformly 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] See Supplemental Material [url], which includes Refs. [11–23].
[6] 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]`.
[7] G. Dvali, "Black Holes and Large N Species Solution to the Hierarchy Problem," *Fortsch. Phys.* **58** (2010) 528–536, `arXiv:0706.2050 [hep-th]`.
[8] G. Dvali and M. Redi, "Phenomenology of 10^{32} Dark Sectors," *Phys. Rev.* **D80** (2009) 055001, `arXiv:0905.1709 [hep-ph]`.
[9] 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]`.
[10] 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]`.
[11] 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]`.
[12] N. F. Bell, E. Pierpaoli, and K. Sigurdson, "Cosmological signatures of interacting neutrinos," *Phys. Rev.* **D73** (2006) 063523, `arXiv:astro-ph/0511410 [astro-ph]`.
[13] 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]`.
[14] 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]`.
[15] 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]`.
[16] S. Davidson, B. Campbell, and D. C. Bailey, "Limits on particles of small electric charge," *Phys. Rev.* **D43** (1991) 2314–2321.
[17] S. Davidson, S. Hannestad, and G. Raffelt, "Updated bounds on millicharged particles," *JHEP* **05** (2000) 003, `arXiv:hep-ph/0001179 [hep-ph]`.
[18] G. Raffelt and D. Seckel, "Bounds on Exotic Particle Interactions from SN 1987a," *Phys. Rev. Lett.* **60** (1988) 1793.
[19] 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].

[20] **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].

[21] J. D. Clarke, R. Foot, and R. R. Volkas, "Phenomenology of a very light scalar (100 MeV ¡ m_h ¡ 10 GeV) mixing with the SM Higgs," *JHEP* **02** (2014) 123, arXiv:1310.8042 [hep-ph].

[22] G. Krnjaic, "Probing Light Thermal Dark-Matter With a Higgs Portal Mediator," arXiv:1512.04119 [hep-ph].

[23] W. Altmannshofer, M. Bauer, and M. Carena, "Exotic Leptons: Higgs, Flavor and Collider Phenomenology," *JHEP* **01** (2014) 060, arXiv:1308.1987 [hep-ph].