

Cosmological signals of a mirror twin Higgs

Nathaniel Craig,^{a,b} Seth Koren^a and Timothy Trott^a

^a*Department of Physics, University of California,
Santa Barbara, CA 93106, U.S.A.*

^b*Kavli Institute for Theoretical Physics,
Santa Barbara, CA 93106, U.S.A.*

E-mail: ncraig@physics.ucsb.edu, koren@physics.ucsb.edu,
ttrott@physics.ucsb.edu

ABSTRACT: We investigate the cosmology of the minimal model of neutral naturalness, the mirror Twin Higgs. The softly-broken mirror symmetry relating the Standard Model to its twin counterpart leads to significant dark radiation in tension with BBN and CMB observations. We quantify this tension and illustrate how it can be mitigated in several simple scenarios that alter the relative energy densities of the two sectors while respecting the softly-broken mirror symmetry. In particular, we consider both the out-of-equilibrium decay of a new scalar as well as reheating in a toy model of twinned inflation, Twinflation. In both cases the dilution of energy density in the twin sector does not merely reconcile the existence of a mirror Twin Higgs with cosmological constraints, but predicts contributions to cosmological observables that may be probed in current and future CMB experiments. This raises the prospect of discovering evidence of neutral naturalness through cosmology rather than colliders.

KEYWORDS: Beyond Standard Model, Cosmology of Theories beyond the SM

ARXIV EPRINT: [1611.07977](https://arxiv.org/abs/1611.07977)

Contents

| | | |
|----------|-------------------------------------------------|-----------|
| 1 | Introduction | 1 |
| 2 | The mirror Twin Higgs | 3 |
| 3 | Thermal history of the mirror Twin | 4 |
| 3.1 | Twin degrees of freedom | 5 |
| 3.2 | Decoupling | 5 |
| 3.3 | Cosmological constraints | 12 |
| 3.3.1 | Effective number of neutrinos | 13 |
| 3.3.2 | Neutrino masses | 16 |
| 3.3.3 | Bounds | 17 |
| 4 | Reheating by the decay of a scalar field | 19 |
| 4.1 | Asymmetric reheating | 20 |
| 4.2 | Imprints on the CMB | 26 |
| 4.2.1 | Analytic estimate of N_{eff} | 26 |
| 4.2.2 | Numerical calculation of N_{eff} | 28 |
| 4.2.3 | Neutrino masses | 32 |
| 4.3 | Thermal production | 32 |
| 5 | Twinflation | 35 |
| 5.1 | Toy model | 36 |
| 5.2 | Observability | 37 |
| 6 | Conclusion | 42 |

1 Introduction

The electroweak hierarchy problem is one of the primary motivators for accessible physics beyond the Standard Model and has led to an expansive set of searches at the LHC and beyond. Recent null results in searches for conventional approaches to the hierarchy problem motivate the exploration of alternative solutions. “Neutral naturalness” provides one such promising alternative, in which the lightest states responsible for protecting the weak scale are partly or wholly neutral under the Standard Model (SM). In these theories, discrete symmetries enforce cancellations between finite threshold corrections to the Higgs mass. The discrete symmetries may be approximate or exact, although solutions with approximate symmetries typically require a plethora of new particles near the TeV scale.

Perhaps the simplest avatar of neutral naturalness is the “mirror” Twin Higgs [1], in which the new physics near the weak scale consists of an identical copy of the Standard

Model related by an exact \mathbb{Z}_2 exchange symmetry. Higgs portal-type couplings between the Higgs doublets of the Standard Model and the twin sector lead to accidental global symmetries that protect the Higgs mass. The lightest partner particles are entirely neutral under the Standard Model, subject only to indirect bounds from precision Higgs coupling measurements. In conjunction with supersymmetry or compositeness at 5–10 TeV, this provides a complete solution to the “little” and “big” hierarchy problems consistent with current LHC limits. In this respect, the Twin Higgs naturally reconciles the observation of a light Higgs with the absence of evidence for new physics thus far at the LHC.

The primary challenge to the mirror Twin Higgs comes not from LHC data, but from cosmology. An exact \mathbb{Z}_2 exchange symmetry predicts mirror copies of light Standard Model states, which contribute to the energy density of the early universe. In particular, twin neutrinos and a twin photon provide a new source of dark radiation that is strongly constrained by CMB and BBN measurements [2, 3]. While these constraints could be avoided if the two sectors were at radically different temperatures, the Higgs portal couplings required by naturalness keep the two sectors in thermal equilibrium down to relatively low temperatures. Constraints on dark radiation in the mirror Twin Higgs have motivated models in which the \mathbb{Z}_2 symmetry is approximate (such as the orbifold [4, 5], holographic [6–8], fraternal [9], and vector-like [10] Twin Higgs), in which case the dark radiation component can be made naturally small. This problem was examined recently in [11], where the \mathbb{Z}_2 symmetry in the fermion Yukawa couplings was broken in order to find an arrangement that would reduce the residual dark radiation from the twin particles.¹ However, such cosmological fixes come at the cost of minimality, as models with approximate \mathbb{Z}_2 symmetries require a considerable amount of additional structure near the TeV scale.

In this work we take an alternative approach and investigate ways in which early universe cosmology can reconcile the mirror Twin Higgs with current CMB and BBN observations. In doing so, we find compelling scenarios that transfer the signatures of electroweak naturalness from high-energy colliders to cosmology. We consider several possibilities in which the energy density of the light particles in the twin sector is diluted by the out-of-equilibrium decay of a new particle after the two sectors have thermally decoupled. Crucially, the new physics in the early universe respects the exact (albeit spontaneously broken) \mathbb{Z}_2 exchange symmetry of the mirror Twin Higgs. This symmetry may be used to classify representations of the particle responsible for this dilution. We concentrate on two minimal cases: in the first, the long-lived particle is \mathbb{Z}_2 -even and the asymmetry is naturally induced by kinematics. In the second, there is a pair of particles which are exchanged by the \mathbb{Z}_2 symmetry and which may be responsible for inflation.² Moreover, in these cases the new physics does not merely reconcile the existence of a mirror twin sector with cosmological constraints, but predicts contributions to cosmological observables that may be probed in current and future CMB experiments. This raises the prospect of dis-

¹For recent related work on the cosmology and cosmological signatures of non-minimal Twin Higgs scenarios, see e.g. [12–18].

²A third case exists, in which the particle is \mathbb{Z}_2 -odd. This may additionally be related to the spontaneous \mathbb{Z}_2 -breaking in the Higgs potential, although we find that a realisation of such a scenario is dependent upon the UV completion of the model.

covering evidence of electroweak naturalness first through cosmology, rather than colliders, and provides natural targets for future cosmological constraints on minimal realizations of neutral naturalness.

This paper is organized as follows: we begin in section 2 by reviewing the salient features of the mirror Twin Higgs. In section 3 we discuss the thermal history of the mirror Twin Higgs, with a particular attention to the interactions keeping the Standard Model and twin sector in thermal equilibrium and the cosmological constraints on light degrees of freedom. In section 4 we present a simple model where the out-of-equilibrium decay of a particle with symmetric couplings to the Standard Model and twin sector leads to a temperature difference between the two sectors after they decouple. We turn to inflation in section 5, constructing a model of “twinflation” in which the softly broken \mathbb{Z}_2 -symmetry extends to the inflationary sector and leads to two periods of inflation. The first primarily reheats the twin sector, while the second primarily reheats the Standard Model sector. We conclude in section 6.

2 The mirror Twin Higgs

We begin by briefly reviewing the salient details of the mirror Twin Higgs. The reader is referred to any of the references listed in the previous section for further details. The theory consists of the Standard Model and an identical copy, related by a \mathbb{Z}_2 exchange symmetry at a scale $\Lambda \gg v$. The two sectors are connected only by Higgs portal-type interactions between the two SU(2) doublet scalars.³ Subject to conditions on the quartic coupling, the Higgs sector enjoys an approximate SU(4) global symmetry.⁴

The Higgs potential is best organized in terms of the accidental SU(4) symmetry involving the SU(2) Higgs doublets of the SM and twin sectors, H_A and H_B . The general tree-level Twin Higgs potential is given by (see e.g. [9])

$$V(H_A, H_B) = \lambda(|H_A|^2 + |H_B|^2 - f^2/2)^2 + \kappa(|H_A|^4 + |H_B|^4) + \sigma f^2 |H_A|^2. \quad (2.1)$$

The first term respects the accidental SU(4) global symmetry, the second breaks SU(4) but preserves the \mathbb{Z}_2 and the final term softly breaks the \mathbb{Z}_2 . Clearly, $\kappa, \sigma \ll \lambda$ are required for the SU(4) to be a good symmetry of the potential. The coupling κ should naturally be of order the expected SU(4)-breaking radiative corrections to the potential induced by Yukawa interactions with the top/twin top, $\kappa \sim 3y_t^4/(8\pi^2) \log(\Lambda/m_t) \sim 0.1$ for a cut-off $\Lambda \sim 10$ TeV (y_t being the top quark Yukawa coupling and m_t its mass). Requiring $\lambda \gg \kappa$ therefore implies $\lambda \gtrsim 1$. As the SM and twin isospin gauge groups are disjoint subgroups of the SU(4), the spontaneous breaking of the SU(4) coincides with the SM and twin electroweak symmetry breaking. Three Goldstone bosons are eaten by the broken gauge

³Here and in what follows we neglect possible kinetic mixing between the two $U(1)_Y$ gauge bosons; such mixing is not generated in the low-energy theory at three loops [1], and may be forbidden in UV completions where the mirror symmetry relates sectors with unified gauge groups.

⁴Properly speaking, the model must contain an SO(8) global symmetry in order to enjoy a residual custodial symmetry [7, 8], but in linear realizations the SU(4) is sufficient provided that higher-dimensional operators violating the custodial symmetry are adequately suppressed.

bosons in each sector, leaving one Goldstone remaining. This will acquire mass through the breaking of the $SU(4)$ that is naturally smaller than the twin scale f . For future reference, it is convenient to define the real scalar degrees of freedom in the gauge basis as $h_A = \frac{1}{\sqrt{2}}\Re(H_A^0) - v_A$ and $h_B = \frac{1}{\sqrt{2}}\Re(H_B^0) - v_B$, where $\langle H_A^0 \rangle = v_A$ and $\langle H_B^0 \rangle = v_B$.

The surviving Goldstone boson should be dominantly composed of the h_A gauge eigenstate in order to be SM-like. The soft \mathbb{Z}_2 -breaking coupling σ is required to tune the potential so that the vacuum expectation values (vevs) are asymmetric and that the Goldstone is mostly aligned with the h_A field direction. The (unique) minimum of the Twin Higgs potential (2.1) occurs at $v_A \approx \frac{f}{2}\sqrt{\frac{\lambda(\kappa-\sigma)-\kappa\sigma}{\lambda\kappa}}$ and $v_B \approx \frac{f}{2}\sqrt{\frac{\sigma+\kappa}{\kappa}}$. The required alignment of the vacuum in the H_B direction occurs if $\sigma \approx \kappa$, which has been assumed in these expressions for the minimum. The consequences of this are that $v_A \approx v/\sqrt{2}$ and $v_B \approx f/\sqrt{2} \gg v$ (where v is the vev of the SM Higgs, although $v_A \approx 174$ GeV is the vev that determines the SM particle masses and electroweak properties), so that the SM-like Higgs h is identified with the Goldstone mode and is naturally lighter than the other remaining real scalar, a radial mode H whose mass is set by the scale f . The component of h in the h_B gauge eigenstate is $\delta_{hB} \approx v/f$ (to lowest order in v/f). Measurements of the Higgs couplings restrict $f \gtrsim 3v$ [9], and the naive tuning of the weak scale associated with this inequality is of order $f^2/2v^2$.

The spectrum of states in the broken phase consists of a SM-like pseudo-Goldstone Higgs h of mass $m_h^2 \sim 8\kappa v^2$, a radial twin Higgs mode H of mass $m_H^2 \sim 2\lambda f^2$, a conventional Standard Model sector of gauge bosons and fermions and a corresponding mirror sector. The current masses of quarks, gauge bosons, and charged leptons in the twin sector are larger than their Standard Model counterparts by $\sim f/v$, while the twin QCD scale is larger by a factor $\sim (1 + \log(f/v))$ due to the impact of the higher mass scale of heavy twin quarks on the renormalisation group (RG) evolution of the twin strong coupling. The relative mass of twin neutrinos depends on the origin of neutrino masses, some possibilities being $\sim f/v$ for Dirac masses and $\sim f^2/v^2$ for Majorana masses from the Weinberg operator. Mixing in the scalar sector implies that the SM-like Higgs couples to twin sector matter with an $\mathcal{O}(v/f)$ mixing angle, as does the radial twin Higgs mode to Standard Model matter. These mixings provide the primary portal between the Standard Model and twin sectors.

The Goldstone Higgs is protected from radiative corrections from \mathbb{Z}_2 -symmetric physics above the scale f . While the mirror Twin Higgs addresses the little hierarchy problem, it does not address the big hierarchy problem, as nothing stabilizes the scale f against radiative corrections. However, the scale f can be stabilized by supersymmetry, compositeness, or perhaps additional copies of the twin mechanism without requiring new states beneath the TeV scale. Minimal supersymmetric UV completions can furthermore remain perturbative up to the GUT scale [19, 20].

3 Thermal history of the mirror Twin

The primary challenge to the mirror Twin Higgs comes from cosmology, rather than collider physics. The mirror Twin contains not only states responsible for protecting the Higgs against radiative corrections (such as the twin top), but also a plethora of extra states due to the \mathbb{Z}_2 symmetry that are irrelevant to naturalness. The lightest of these, namely

the twin photon and twin neutrinos, contribute significantly to the energy density of the early universe around the era of matter-radiation equality, since they have a temperature comparable to that of the Standard Model plasma at all times. This is because the same Higgs portal coupling that makes the Higgs natural also keeps the two sectors in thermal equilibrium down to $\mathcal{O}(\text{GeV})$ temperatures. Then the identical particle content in the twin and Standard Model sectors guarantees that they remain at comparable temperatures even after they decouple — for every massive Standard Model species that becomes non-relativistic and transfers its entropy to the rest of the plasma, its twin counterpart does the same within a factor of f/v in temperature.

In this section we undertake a detailed study of the decoupling between the Standard Model and twin sectors as well as the constraints from precision cosmology.

3.1 Twin degrees of freedom

In thermal equilibrium, each relativistic degree of freedom has roughly the same energy density. In general, we express the energy density of the universe ρ during the radiation-dominated era as $\rho \equiv g_\star \frac{\pi^2}{30} T^4$, where we define g_\star through this relation as the effective number of relativistic degrees of freedom and T the temperature of the SM photons. This then determines the evolution of the scale factor through the first Friedmann equation

$$H = \frac{1}{M_{\text{pl}}} \left[\frac{\pi^2}{90} g_\star T^4 \right]^{1/2} \tag{3.1}$$

(assuming spatial flatness), where M_{pl} is the reduced Planck mass. In general, the energy density of a particular species i may be computed from $\rho_i = g_i \int \frac{d^3p}{(2\pi)^3} f_i(p, T_i) E(p)$, where g_i are the number of internal degrees of freedom, $E(p)$ is the energy as a function of momentum p , while $f_i(p, T_i)$ is the phase-space number density and is a Bose-Einstein or Fermi-Dirac distribution if the species is in equilibrium at temperature T_i . The number of effective relativistic degrees of freedom may then be defined for each sector separately as $g_\star^{\text{SM}}(T)$ and $g_\star^{\text{t}}(T)$ satisfying $\rho_{\text{SM}}(T) = \frac{\pi^2}{30} g_\star^{\text{SM}}(T) T^4$ and $\rho_{\text{t}}(T) = \frac{\pi^2}{30} g_\star^{\text{t}}(T) T^4$, respectively, where $\rho_{\text{SM}}(T)$ and $\rho_{\text{t}}(T)$ are the total energy densities of SM and twin particles. The values of $g_\star(T)$ for the SM and twin sectors are shown in figure 1, where all species within each sector are in thermal equilibrium. These can then be used to calculate the total number g_\star as a function of temperature, by weighting twin sector energy density by its temperature: $g_\star(T) = g_\star^{\text{SM}}(T) + g_\star^{\text{t}}(\hat{T})(\hat{T}/T)^4$, where \hat{T} is the twin sector photon temperature when the SM photon temperature is T .

Likewise, entropy densities for each sector i are defined as $s_i(T) = \frac{2\pi^2}{45} g_\star^i(T) T^3$. We neglect the small differences between the number of relativistic degrees of freedom defined from energy and entropy densities, which are not significant over the range of temperatures of interest here.

3.2 Decoupling

In the early universe, the two sectors are thermally linked by interactions mediated by the Higgs, which, through mixing with both h_A and h_B components, allows for SM fermions

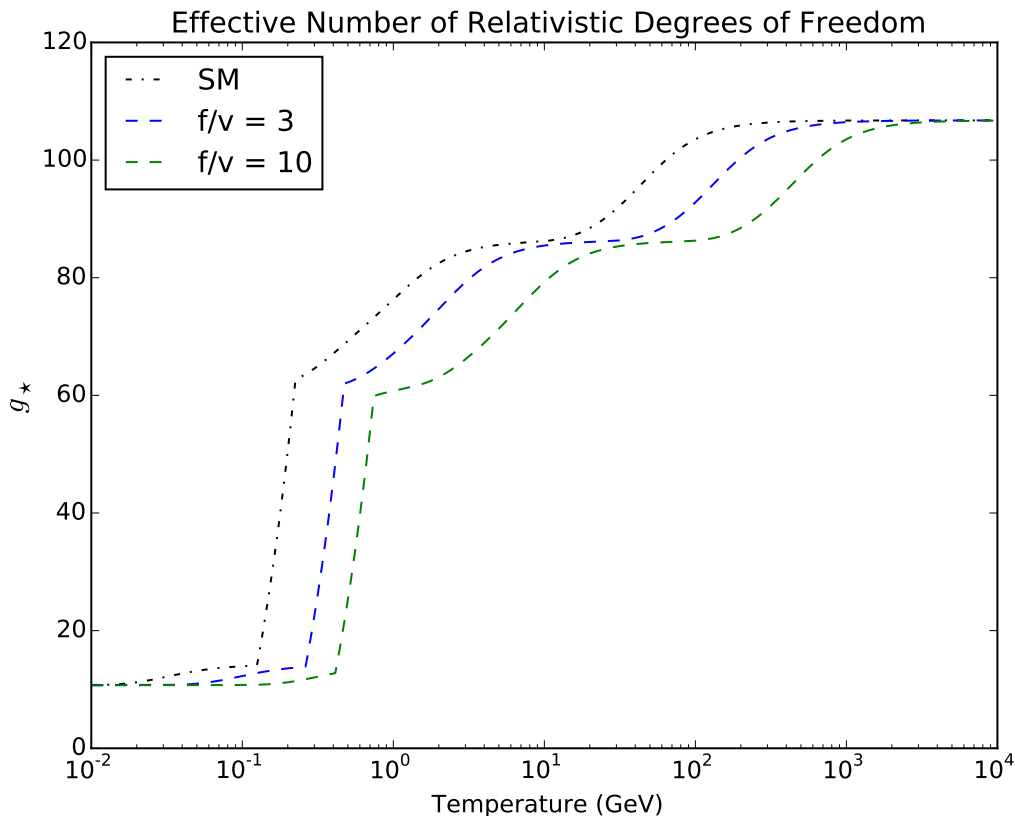


Figure 1. The effective number of relativistic degrees of freedom for mirror Twin Higgs models for different values of f/v . The dash-dotted line is the for the Standard Model $g_*^{\text{SM}}(T)$ and the dashed lines are the twin sector degrees of freedom $g_*^{\dagger}(T)$. The evolution of g_* during the QCD phase transition (QCDPT) is not well-understood, so we assign the SM QCDPT a central value of 175 MeV and a width of 50 MeV and interpolate linearly between the values of g_* at 225 MeV for free partons and at 125 MeV for pions. Further discussion may be found in [21]. For the twin sector we use a central value and width which are $(1 + \log(\frac{f}{v}))$ times larger than the SM values. Note that new mass thresholds, expected to appear at energies ~ 10 TeV in UV completions of the twin Higgs, have not been included.

and weak bosons to scatter off or annihilate into their twin counterparts. However, once the temperature drops sufficiently for this Higgs-mediated interaction to become rare on the expansion time-scale, the sectors decouple and thereafter thermally evolve independently. More precisely, thermal decoupling will occur once the rate at which energy can be exchanged between SM and twin particles (through the Higgs) falls below the Hubble rate.

Thermal decoupling is traditionally formulated from the Boltzmann equations describing the evolution of single-particle phase space number densities, wherein collisions induce instantaneous changes to the shape of these distributions. When the collisions occur faster than the expansion rate, the phase space probability density functions of the interacting species are expected to relax to an equilibrium distribution (Boltzmann, neglecting quantum statistics, will be applicable to our case). However, once the rate of collisions falls below the expansion rate, collisions become rare on cosmological time scales and the phase

space distributions depart from equilibrium. The decoupling temperature is determined as that at which the scattering rate of a participating particle, Γ , drops below the Hubble rate, assuming that this occurs instantaneously across the entire phase space where the number density is significant. This formulation can be used to determine the time at which a particular species of particle will cease to scatter off twin particles on cosmological time scales.

In the case of interest here, however, both sectors of particles remain thermalised within themselves while the interactions between sectors freeze-out. This implies that the phase space number densities are still Boltzmann distributions throughout decoupling, with a different temperature for each sector. As it is the twin sector temperature that ultimately determines the impact of the light twin degrees of freedom on the cosmological observables (discussed below in section 3.3), we wish to describe the thermal evolution of the two sectors by that of their entire energy or entropy content and the bulk heat flows between them. They may then be identified as thermally decoupled once the rate at which they exchange energy falls below the expansion rate.

If the SM and twin sector plasmas have temperatures T and \hat{T} respectively, then calling q the net heat flow density from the SM to the twin sector, the rate at which the twin entropy densities s_t and s_{SM} evolve is determined by

$$\frac{ds_t}{dt} + 3Hs_t = \frac{1}{\hat{T}} \frac{dq}{dt} = \frac{1}{\hat{T}} \left(\frac{dq_{in}}{dt} - \frac{dq_{out}}{dt} \right) \quad (3.2)$$

$$\frac{ds_{SM}}{dt} + 3Hs_{SM} = \frac{-1}{T} \frac{dq}{dt} = -\frac{1}{T} \left(\frac{dq_{in}}{dt} - \frac{dq_{out}}{dt} \right). \quad (3.3)$$

Here, H is the Hubble rate. The heat flow rate has been decomposed into the sum of the energy transferred into and out of the twin sector by collisions in the second equality in each line, where $\frac{dq_{in}}{dt}$ and $\frac{dq_{out}}{dt}$ are both positive.

The rate of heat flow q may be calculated by performing a phase space average of the rate that energy is transferred from the SM to the twin sector through particle interactions. Since the decay rates of top quarks or weak bosons are fast compared to their scattering rate and the Hubble rate, energy transferred to them is instantaneously transferred to the rest of the plasma. Similarly, the scattering rate of lighter fermions off other particles of the same sector (such as photons or gluons) is much faster than their interaction rate with twin fermions. Energy transferred to the lighter fermions therefore quickly diffuses throughout their respective plasmas. The rate of heat flow between sectors may therefore be well approximated by the rate at which energy is transferred from SM particles to twin particles in Higgs mediated interactions. This may occur through elastic scattering of SM particles off twin particles or annihilations of SM particle/antiparticle pairs into twin particles (or the reverse). The energy density transferred to twin particle i from SM particle j in scattering is given by

$$\frac{dq_{ij \rightarrow ij}}{dt} = \frac{g_i g_j}{(2\pi)^6} \int \int \frac{d^3k}{2E_i(k)} \frac{d^3h}{2E_j(h)} f_i(k, \hat{T}) f_j(h, T) \left(4E_i(k) E_j(h) \int v_{rel}(E_i(p) - E_i(k)) \frac{d\sigma_{ij \rightarrow ij}}{d\Omega} d\Omega \right), \quad (3.4)$$

where p is the outgoing 4-momentum of particle i . In the cosmic comoving frame, the phase space number densities f_i and f_j are just Boltzmann factors, although evaluated at the different temperatures of each sector. The factor g_i is the number of internal degrees

of freedom of particle i , which here includes colour (the cross section should not be colour averaged, as each colour of quark is present in the plasma in equal abundances and each mediates the exchange of energy, so have their contributions summed). Finally, $E_i(k)$ is the on-shell energy of particle i with momentum k , while $\frac{d\sigma_{ij \rightarrow ij}}{d\Omega}$ is the differential scattering cross section for species i scattering off j per solid angle Ω and v_{rel} is the usual relative speed of the incoming particles. As described in [22], the factor in the integrand giving the energy transferred per reaction is simply a component of a 4-vector,

$$X = 4E_i(k)E_j(h) \int (p - k)v_{\text{rel}} \frac{d\sigma_{ij \rightarrow ij}}{d\Omega} d\Omega. \tag{3.5}$$

This may be calculated in the centre-of-mass frame and then boosted back into the cosmic comoving frame where the integrals in (3.4) can be evaluated, similarly to the thermal averaging procedure described in [23].

The integral (3.4) may be decomposed into two terms giving the positive and negative energy changes of the twin particle, which respectively contribute to $\frac{dq_{\text{in}}}{dt}$ and $\frac{dq_{\text{out}}}{dt}$. When evaluated in the centre-of-mass frame, these terms correspond to the cases where the scattering angle of the twin particle is respectively less than and greater than the angle between its initial momentum and the total momentum of the system. However, when $T \neq \hat{T}$, we find the integrals involved in this decomposition substantially more arduous than when they are evaluated together.

Energy transferred through annihilations may be similarly calculated as

$$\begin{aligned} \frac{dq_{j\bar{j} \rightarrow i\bar{i}}}{dt} &= \frac{g_j^2}{(2\pi)^6} \int \int \frac{d^3k}{2E_j(k)} \frac{d^3h}{2E_j(h)} f_j(k)f_j(h) \left(4E_j(k)E_j(h) \int v_{\text{rel}}(E_j(h) + E_j(k)) \frac{d\sigma_{j\bar{j} \rightarrow i\bar{i}}}{d\Omega} d\Omega \right) \\ &\quad - \frac{g_i^2}{(2\pi)^6} \int \int \frac{d^3k}{2E_i(k)} \frac{d^3h}{2E_i(h)} f_i(k)f_i(h) \left(4E_i(k)E_i(h) \int v_{\text{rel}}(E_i(h) + E_i(k)) \frac{d\sigma_{i\bar{i} \rightarrow j\bar{j}}}{d\Omega} d\Omega \right), \end{aligned} \tag{3.6}$$

where $\frac{d\sigma_{j\bar{j} \rightarrow i\bar{i}}}{d\Omega}$ is now the differential annihilation cross section. This rate may be evaluated as described above and is more directly amenable to the factorisation of the integrals observed in [23]. See also [24] for further details of similar calculations. The first term of (3.6) is the energy transferred from the SM to the twin sector and contributes to $\frac{dq_{\text{in}}}{dt}$ in (3.2), while the second term is the energy transferred from the twin sector to the SM and contributes to $\frac{dq_{\text{out}}}{dt}$.

In thermal equilibrium, the rate of energy transferred through collisions into one sector will be balanced by that of energy transferred out of it so that there is negligible net heat flow. This state will be rapidly attained (compared to the age of the universe) if $\frac{dq_{\text{in,out}}}{dt} \gg 3H\hat{T}s_t$. However, as the universe expands and the plasma cools, the energy transfer rates fall faster than the Hubble rate. This is demonstrated in the figure 2 below. Once they drop below the Hubble rate, energy exchange ceases on cosmological time scales and the sectors thermally decouple, thereafter thermodynamically evolving independently.

To determine the decoupling temperature of the sectors, we calculate the rates of positive energy exchange for the twin particles interacting with the SM particles. The cross sections are calculated using a tree-level effective fermion-twin fermion contact interaction

that, in the full twin Higgs model, would be UV completed by a SM Higgs exchange (the heavier mass of the radial mode would make its exchange subdominant). The interaction strength is determined by the masses of the fermions through their Yukawa couplings, as well as the mixing angle of the SM-like mass state h with the gauge eigenstate h_B , giving a 4-fermion coupling of strength $\frac{m_f m_{\hat{f}}}{m_h^2 f^2}$ (here m_f and $m_{\hat{f}}$ are the masses of fermions f and \hat{f}). See [11, 19] for a more detailed discussion of the cross sections. This effective interaction is appropriate for the temperatures of interest here and helps to simplify the integrals of (3.4). In order to further simplify the integrations of (3.4) when it is to be decomposed into terms in which the energy exchange is positive and negative, we calculate $\frac{dq_{\text{in}}}{dt}$ under the assumption that the sectors have the same temperature (this ensures that the rate $\frac{dq_{\text{out}}}{dt}$ is identical). This is then combined with the rate of energy transferred from annihilation. A similar calculation of these rates was recently performed in [11], for cases where the Yukawa couplings do not respect the \mathbb{Z}_2 twin symmetry.

In figure 2 we compare the energy transfer rate to the Hubble rate in order to determine when decoupling occurs. As long as the energy exchange rate exceeds the expansion rate, the sectors will be thermalised and have the same temperature. Decoupling then occurs once this rate drops below the Hubble rate. From figure 2, this occurs at a temperature ~ 2 GeV. However, even after the energy exchange rate drops below the Hubble rate, the sectors will remain at the same temperature unless some event that either injects or redistributes entropy occurs within a sector (such as the temperature dropping below a mass threshold). As the heavy quark masses roughly coincide with the decoupling temperature, these do cause the twin sector to be mildly reheated with respect to the SM below decoupling. However, the resulting temperature difference is small and the energy exchange rates are expected to continue to be well-approximated by the rates presented in figure 2 beyond decoupling.

The lower plot of figure 2 illustrates the decomposition of the energy exchange rates into contributions from interactions involving different SM quarks. The interaction cross sections are proportional to the Yukawa couplings of the interacting fermions. The greatest heat exchange is therefore expected to be mediated by the most massive particles, provided that their abundances are not too Boltzmann suppressed. As expected, at temperatures ~ 1 GeV, the bottom quark is the best conduit of thermal equilibration, followed by the charm quark and then the τ (with colour factors enhancing the former two with respect to the latter). The rate of heat flow that the top quarks and weak bosons can mediate at these temperatures (or below) is negligible because of Boltzmann suppression. The bend in the curves at temperatures ~ 5 GeV in the lower plot corresponds to a transition from temperatures where the dominant energy exchange rate is through scatterings to those where it occurs through annihilations, as can be seen in the upper plot. The annihilation rate into twin bottom quarks is the dominant component at high enough energies (again because of the larger Yukawa coupling), but this becomes rapidly threshold suppressed as the temperature drops. As can also be inferred in the upper plot, the energy exchange rate through annihilations involving the twin charmed quarks and tau leptons overtakes that of twin bottom quarks at similar temperature, but are still subdominant to scatterings.

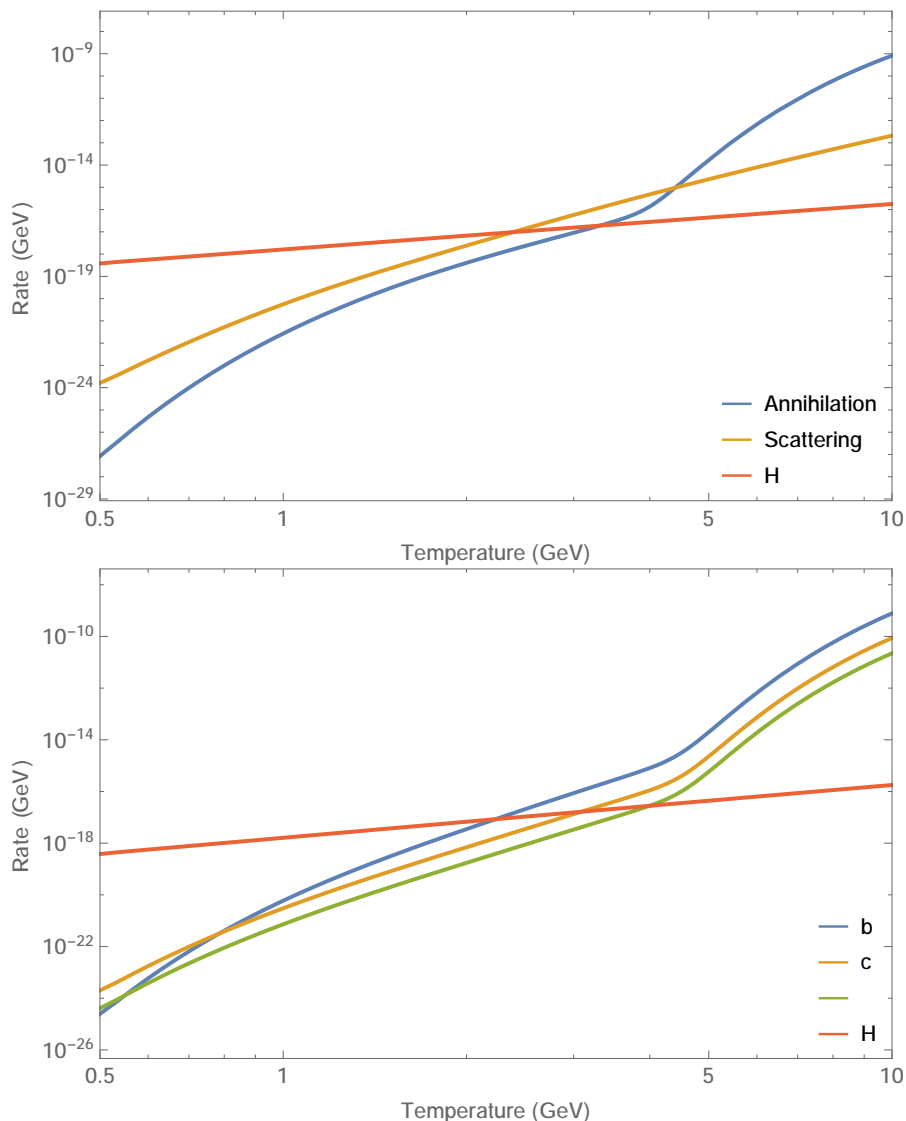


Figure 2. Rates of energy density exchange per twin entropy density ($\frac{1}{3s_t T} \frac{dq_{in}}{dt}$) decomposed into contributions from scattering and annihilation (top) and for interactions involving different species of SM fermions (bottom), along with the Hubble parameter, for $f/v = 4$. The decoupling temperature is that where the sum of the energy exchange rates equals the Hubble rate, which occurs at $T_{\text{decoup}} \approx 2$ GeV.

The decoupling temperature depends upon f/v , which sets both the mass scale of the twin sector and the strength of the Higgs-mediated coupling. As f/v is increased, decoupling occurs earlier because of the greater Boltzmann suppression, although this is only a relatively small effect that, for $f/v = 10$, increases the decoupling temperature by only 4 GeV.

When the twin sector is colder than the SM (which will be important for much of what follows) the heat flow is typically dominated by annihilations of SM into twin particles. However, the energy exchange from elastic scattering can be comparable to that from

annihilations, as illustrated in figure 2. Although the energy exchange in an annihilation will generally exceed that of a scattering because all of the energy involved in the process must be transferred, the annihilation rate also becomes more Boltzmann or threshold suppressed when the temperature drops below the mass of the heavier twin particles. It is therefore not always clear that energy transfer through annihilations dominates.

Decoupling is not exactly instantaneous and there is some range of temperatures over which the rate of heat flow freezes-out. The net heat flow rate $\frac{dq}{dt}$ is greater for larger temperature differences between sectors. The generation of a potentially large temperature difference within this brief epoch of sector decoupling, such as those discussed below in section 4, may be cut off when the heat flow rate becomes comparable to the Hubble rate. For a given SM temperature T , the minimum twin-sector temperature \hat{T}_{\min} during the decoupling period may be roughly estimated as that which satisfies

$$H \sim \frac{1}{3s_t \hat{T}} \left. \frac{dq}{dt} \right|_{\hat{T}=\hat{T}_{\min}}. \tag{3.7}$$

Twin temperatures colder than \hat{T}_{\min} will partially thermalise back to this value. As the participating fermions are not non-relativistic, instantaneous decoupling is not as accurate an approximation as it is, for example, for chemical decoupling of a WIMP, although it is still reliable.

In figure 3, we show the minimum temperature that the twin sector may have as a function of SM temperature for heat flow to freeze out, estimated using (3.7). Only annihilations have been included in the determination of the minimum temperature, although we have verified that, for these temperatures, the scatterings contribute only $\lesssim 10\%$ to the heat flow. Note that while the energy exchange rate, such as $\frac{1}{T} \frac{dq_{\text{in}}}{dt}$ in (3.2), in scattering processes may be faster, the net energy flow rate, or heat flow ($\frac{1}{T} \frac{dq}{dt}$ in (3.2)), which is the difference between energy exchange rates into and out of the sector, is actually dominated by annihilations. Generally, we find that decoupling begins at temperatures ~ 4 GeV. The temperature difference can reach an order of magnitude without relaxing once the SM temperature drops to ~ 1 GeV.

While the extent of thermal decoupling is temperature dependent, the maximum temperature difference that will not relax grows quickly as the SM temperature drops. Then we may describe the two sectors as being decoupled if, in a given cosmology, all events that raise the temperature of one sector relative to the other (such as the crossing of a mass threshold and the resulting entropy redistribution, the most significant of which is the confinement of colour) induce temperature differences that are too small to partially relax.

At energies $\lesssim 1$ GeV in figure 2, the reliability of the calculation of the heat flow rate diminishes because of the strengthening of the strong coupling and the eventual confinement of colour. Fortunately, for a cooler twin sector, which will be of interest in subsequent sections, annihilations from the SM dominate other processes over most of the parameter space. These are the least sensitive to higher order corrections and non-perturbative effects because of their higher temperature, and hence energy, compared to the potentially cooler twin sector. The range of temperatures illustrated in figures 2 and 3 have been selected to roughly illustrate the duration of decoupling, but may extend below the range where the

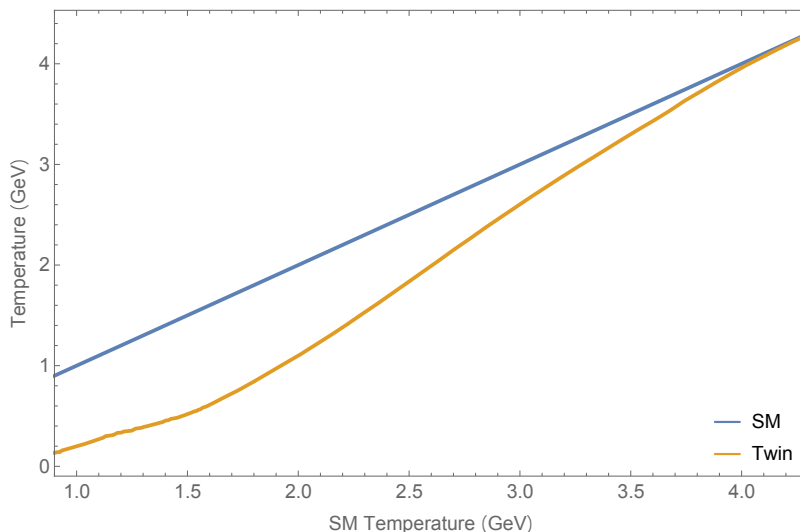


Figure 3. Minimum temperature of the twin sector that will not be heated by interactions with a hotter SM plasma, as a function of SM temperature, for $f/v = 4$. Also shown is the SM temperature, for reference.

perturbative calculation of the heat flow rate is valid. For example, at temperatures below the twin sector QCDPT, which occurs at $\sim \left(1 + \log\left(\frac{f}{v}\right)\right)$ higher temperatures than in the SM, the partonic calculation of twin quark/anti-quark pair production must be replaced by a hadronic one. Furthermore, the growth of the twin strong coupling necessitates that the quark-Higgs Yukawa couplings be RG evolved to the scale of the energy exchanged, which can induce an $\mathcal{O}(1)$ change to the cross section, although this has only a relatively small effect on the decoupling temperature. It is nevertheless clear that decoupling is mostly complete by then and that these uncertainties are not large enough to affect this conclusion.

In the standard mirror Twin Higgs cosmology, knowing the decoupling temperature tells us how the temperatures of the two sectors will be related at subsequent times. The sectors separately evolve adiabatically after decoupling, though they redshift in the same way and differences in temperature only arise from events that redistribute entropy. Non-minimal cosmological events that could potentially cause the temperatures of each sector to diverge can therefore only be effective if they leave each sector colder than this approximate decoupling temperature.

3.3 Cosmological constraints

Given that the twin and Standard Model sectors remain in thermal equilibrium to $\mathcal{O}(\text{GeV})$ temperatures, the simplest mirror Twin Higgs scenario is cosmologically inviable due to the presence of light twin species (photons and neutrinos) with abundances comparable to those of the SM. The cosmological observables through which evidence of light species may be inferred are typically represented by N_{eff} , the “effective number of neutrino species” in the early universe; their individual masses, which determine their free-streaming distances; and the “effective mass” m_ν^{eff} , which parameterises their contribution to the present-day energy density of non-relativistic matter. These observables are probed by both the CMB and large scale structure (LSS).

3.3.1 Effective number of neutrinos

The parameter N_{eff} describes the amount of radiation-like energy density during the evolution of the CMB anisotropies before photon decoupling. It is defined as the effective number of massless neutrinos with temperature as predicted in the standard cosmology that would give equivalent energy density in radiation:

$$\rho_r = \rho_\gamma + \frac{7}{8} \left(\frac{4}{11} \right)^{4/3} N_{\text{eff}} \rho_\gamma, \tag{3.8}$$

where ρ_r is the energy density of radiation and ρ_γ is the energy density of photons (the factor of $\left(\frac{4}{11}\right)^{4/3}$ arises from the relative reheating of the photons from electron/positron annihilation, which occurs after most of the neutrinos have decoupled, and the factor of $7/8$ is from the opposite spin statistics). A deviation from the Standard Model prediction of 3.046 [25] is denoted by $\Delta N_{\text{eff}} = N_{\text{eff}} - 3.046$. This definition of radiation, or equivalently, relativistic degrees of freedom, becomes less clear if the new fields have a non-negligible mass, as we discuss further below.

We here review the CMB physics of dark radiation, summarising the discussion in [26]. See also [2] for further review. The angular size and scale of the first acoustic peak is well-measured and this approximately fixes the scale factor at matter-radiation equality a_{eq} . If we imagine fixing all other Λ CDM parameters, extra radiation would delay the epoch of matter-radiation equality. This would have a pronounced effect on the power spectrum in the vicinity of the first acoustic peak through the early Integrated Sachs-Wolfe (eISW) effect. The modes corresponding to this feature entered the horizon close to matter-radiation equality and the evolution of their potentials is highly sensitive to the radiation energy density. However, the impact of a $\Delta N_{\text{eff}} \sim \mathcal{O}(1)$ deviation on the peak height can be simultaneously balanced by increasing the amount of non-relativistic matter, to the extent to which other observations providing independent constraints upon Ω_c permit (for Λ CDM+ N_{eff} , a variation of $\sim 10\%$ in $\Omega_c h^2$ is consistent with present CMB+BAO measurements [2], although these variations must be consistent with other observables). This degeneracy is not expected to be broken by CMB-S4 [27].

Given that a_{eq} is approximately fixed, the utility of N_{eff} arises because, in simple extensions of the Λ CDM model, it approximately corresponds to the suppression of power in the small scale CMB anisotropies that arises from Silk damping. The reason for this is roughly that, although the greater expansion rate induced by the extra radiation reduces the time that CMB photons have to diffuse before decoupling, it also reduces the sound horizon size more severely. As the angular size of the sound horizon is determined by the location of the acoustic peaks and is also well measured, the reduction in the sound horizon must be compensated for by a reduction in the angular diameter distance to the CMB. This effectively raises the angular distance over which photon diffusion proceeds and results in a prediction of smoother temperature anisotropies at small scales. This correspondence with the Silk damping allows N_{eff} to be approximately factorised from other parameters and constrained independently, providing a direct observational avenue for detecting the presence of new, massless fields [26] (see [28] for further implications for model building).

This relationship arises because the fixing of a_{eq} implies that N_{eff} effectively determines the energy density of the universe, and hence the Hubble rate, during CMB decoupling. Note, however, that further extensions of Λ CDM may complicate this correspondence, in particular deviations from the standard Big Bang Nucleosynthesis prediction of the primordial helium abundance.

The contribution to N_{eff} (or ΔN_{eff}) in the mirror Twin Higgs arises from two sources: the twin photons, which can be treated as massless dark radiation with an appropriate twin temperature T_{eq}^t at the time of matter-radiation equality, and the twin neutrinos, whose non-zero masses may need to be accounted for. For the twin photons, the contribution to N_{eff} is simple; their equation of state is always $w = 1/3$ and their energy density is given by $g \frac{\pi^2}{30} (T_{\text{eq}}^t)^4$, where $g = 2$. The twin temperature at matter-radiation equality is found from the SM temperature using comoving entropy conservation,

$$\frac{T_{\text{eq}}^t}{T_{\text{eq}}^{\text{SM}}} = \left(\frac{g_{\star}^t(T_{\text{decoup}})}{g_{\star}^{\text{SM}}(T_{\text{decoup}})} \right)^{1/3} \left(\frac{g_{\star}^{\text{SM}}(T_{\text{eq}}^{\text{SM}})}{g_{\star}^t(T_{\text{eq}}^t)} \right)^{1/3}, \tag{3.9}$$

where the two sectors have the same number of thermalized degrees of freedom by this time. Here, $T_{\text{eq}}^{\text{SM}}$ is the SM photon temperature at matter-radiation equality and T_{decoup} is the sector decoupling temperature.

Since neutrinos are massive, their behavior is more complicated. Their equation of state parameter takes on a scale factor dependence which is controlled by their mass. In the Standard Model, this sensitivity is negligible because present CMB bounds imply that neutrinos are ultra-relativistic at a_{eq} to good approximation [2]. However, the factor by which the twin neutrino masses are enhanced may raise them to order T_{eq}^t or greater (see section 2 for discussion of the scaling of the masses with f/v).

To better describe the impact of the extra twin (semi-)relativistic degrees of freedom on the CMB, we choose to define N_{eff} through the effects of neutrinos at matter-radiation equality, when the impact on the expansion rate of the universe for most of the period relevant for the evolution of the CMB is greatest. Note that, in their presentation of joint exclusion bounds on N_{eff} and $\sum m_{\nu}$ (the sum of SM neutrino masses) or m_{ν}^{eff} (effective mass contributing to the present-day non-relativistic matter density of an extra sterile neutrino), the Planck collaboration define N_{eff} as the value in (3.8) at temperatures sufficiently high that the neutrinos are fully relativistic. Our values cannot be directly compared with their analysis, although we consider ours to be a reasonable rough estimate that is more representative of the CMB constraints. The ensuing correction from the finite neutrino masses is, in the cases considered in this work, a small effect anyway.

To determine this correction and provide a definition of N_{eff} that better describes the impact of quasi-relativistic particles on the CMB, we first define the epoch of matter-radiation equality as the time at which the average equation of state parameter of the universe is $\bar{w} = 1/6$ (the equation of state is defined as $\rho = \bar{w}P$, where ρ is energy density and P is pressure). We can express this condition as

$$\left. \frac{d \ln H}{d \ln a} \right|_{a_{\text{eq}}} = -\frac{7}{4}, \tag{3.10}$$

as in [29]. Call the quasi-relativistic neutrino energy density $\tilde{\rho}(a)$ with time-evolving equation of state parameter $w(a)$, which is to be balanced against some extra non-relativistic energy density $\Delta\rho_{\text{CDM}}(a) \propto a^{-3}$ to keep a_{eq} the same. This amount of non-relativistic energy density $\Delta\rho_{\text{CDM}}$ is

$$\Delta\rho_{\text{CDM}}(a_{\text{eq}}) = \rho_r(a_{\text{eq}}) - \rho_m(a_{\text{eq}}) - 2a_{\text{eq}} \left. \frac{d\tilde{\rho}}{da} \right|_{a_{\text{eq}}} - 7\tilde{\rho}(a_{\text{eq}}), \quad (3.11)$$

where ρ_r and ρ_m are the energy densities of the radiation and non-relativistic matter. For a perfect fluid, $\frac{d\tilde{\rho}}{da} = -3(1+w(a))\tilde{\rho}/a$ (neglecting the anisotropic stress that is expected only to contribute to a weak phase shift in the CMB [30]), this results in a Hubble parameter of

$$H^2(a_{\text{eq}}) = \frac{2}{3M_{\text{pl}}^2} [\rho_r(a_{\text{eq}}) + 3w(a_{\text{eq}})\tilde{\rho}(a_{\text{eq}})]. \quad (3.12)$$

This suggests a definition of the effective number of neutrinos, N_{eff} , via

$$H^2(a_{\text{eq}}) = \frac{2}{3M_{\text{pl}}^2} \left(\rho_\gamma + N_{\text{eff}} \rho_{\nu, m=0}^{\text{th}} \right) \Big|_{a_{\text{eq}}} \quad (3.13)$$

$$N_{\text{eff}} \equiv \sum_i \frac{w_i}{1/3} \frac{\rho_i}{\rho_{\nu, m=0}^{\text{th}}}, \quad (3.14)$$

where ρ_i is the contribution to the energy density from some species i with equation of state parameter w_i and $\rho_{\nu, m=0}^{\text{th}}$ is the energy density of a massless neutrino with a thermal distribution in the standard cosmology. Then $3w$ gives the ‘relativistic fraction’ of the energy density. Note that this is simply a ratio of the pressure exerted by the new fields to that of a massless neutrino. The effectiveness of this approximation was discussed in [31] in the context of thermal axions (while effective at keeping a_{eq} fixed, changes to odd peak heights subsequent to the first are imperfectly cancelled and require further changes to H_0 to compensate — see section 3.3.2 below).

Calling T_ν^i the temperature at which the neutrinos in sector i freeze-out and a_ν^i the corresponding scale factor, then assuming instantaneous decoupling, the phase space number density for scale factor a is given by a redshifted Fermi-Dirac distribution [32]

$$f_\alpha^i(p) \approx \left[1 + e^{pa/(a_\nu^i T_\nu^i)} \right]^{-1} \quad (3.15)$$

for the α neutrino mass eigenstate in the i sector ($m_\alpha^i \ll T_\nu^i$, so has been dropped). The energy density and pressure are

$$\rho_{\nu_\alpha}^i = \frac{g_\alpha}{2\pi^2} \int_0^\infty dp p^2 \sqrt{p^2 + (m_\alpha^i)^2} f_\alpha^i(p) \quad (3.16)$$

$$P_{\nu_\alpha}^i = \frac{g_\alpha}{2\pi^2} \int_0^\infty dp \frac{p^4}{3\sqrt{p^2 + (m_\alpha^i)^2}} f_\alpha^i(p), \quad (3.17)$$

where $g_\alpha = 2$ is the number of degrees of freedom for a neutrino species.

Since the neutrino decoupling temperature depends on the strength of the weak interaction as $T_\nu \propto G_F^{-2/3}$, while $G_F \propto v^2$, then the twin neutrino decoupling temperature T_ν^t is related to the SM neutrino decoupling temperature T_ν^{SM} by

$$T_\nu^t = (f/v)^{4/3} T_\nu^{\text{SM}}. \tag{3.18}$$

We can then simply use (3.16) and (3.17) at matter-radiation equality to find ΔN_{eff} (assuming instantaneous decoupling). We thus obtain

$$H^2(a_{\text{eq}}) = \frac{2}{3M_{\text{pl}}^2} \left(\rho_\gamma^{\text{SM}} + 3.046\rho_{\nu,m=0}^{\text{th}} + \rho_\gamma^t + \sum_\alpha 3w_{\nu_\alpha} \rho_{\nu_\alpha}^t \right) \Big|_{a_{\text{eq}}} \tag{3.19}$$

and

$$\Delta N_{\text{eff}} = \left(\frac{11}{4} \right)^{4/3} \frac{120}{7\pi^2 (T^{\text{SM}})^4} \left(\rho_\gamma^t + \sum_\alpha 3w_{\nu_\alpha}^t \rho_{\nu_\alpha}^t \right), \tag{3.20}$$

where we now have equation of state parameters w_{ν_α} for each neutrino, while ρ_γ^{SM} and ρ_γ^t are the SM and twin photon energy densities, $\rho_{\nu,m=0}^{\text{th}}$ and $\rho_{\nu_\alpha}^t$ are the neutrino energy densities.

3.3.2 Neutrino masses

Because they are so weakly interacting, the neutrinos have a long free-streaming scale given by the distance travelled in a Hubble time v_ν/H , with $v_\nu \propto m_\nu^{-1}$ the speed of the neutrino once it becomes non-relativistic. This defines a free-streaming momentum scale $k_{fs} = \sqrt{\frac{3}{2}} \frac{aH}{v_\nu} \propto m_\nu$, above which neutrinos do not cluster. Below this scale, perturbations in the matter density consist coherently of neutrinos and other matter, but well above it only non-neutrino matter contributes to density perturbations. This results in a suppression of the matter power spectrum on large scales which is proportional to the fraction of energy density in the free-streaming matter. Since this occurs at late times when neutrinos are non-relativistic, the energy density is simply $\rho_{\nu_\alpha} = n_{\nu_\alpha} m_{\nu_\alpha}$ for each neutrino species α , where n_{ν_α} is the number density. Constraints on the sum of neutrino masses then come from the observations of power on small scales, which is suppressed relative to that expected for massless neutrinos by a factor $\lesssim 1-8f_\nu$, where $f_\nu = \Omega_\nu/\Omega_m$ is the fraction of non-relativistic energy in neutrinos at late times [33].

More generally, inferences of the matter power spectrum constrain the present-day energy density fraction of free-streaming species that do not cluster on small scales and have since become non-relativistic, $\Omega_\nu = (\sum m_\nu + m_\nu^{\text{eff}})/(94.1 \text{ eV})$, where $\sum m_\nu$ is the sum of SM neutrino masses and m_ν^{eff} is the sum of twin neutrino masses weighted by their number density

$$m_\nu^{\text{eff}} = \frac{n_\nu^t}{n_\nu^{\text{SM}}} \sum_\alpha m_{\nu_\alpha}^t. \tag{3.21}$$

Here n_ν^t is the number density of a relic twin neutrino flavour and n_ν^{SM} is that for a SM neutrino. It is assumed that the neutrinos have been thermally produced as hot relics.

The relic abundance of a neutrino species is given by its number density when it decoupled, diluted by the factor by which the universe has since expanded. The scale factors

at which neutrino decoupling occurs in the two sectors, a_ν^{SM} and a_ν^{t} can be determined from (3.18), the relative temperatures in the two sectors and comoving entropy conservation, to obtain

$$a_\nu^{\text{t}} = a_\nu^{\text{SM}} \left(\frac{v}{f}\right)^{4/3} \left(\frac{g_\star^{\text{t}}(T_{\text{decoup}})}{g_\star^{\text{SM}}(T_{\text{decoup}})}\right)^{1/3} \quad (3.22)$$

where the same mass thresholds have been assumed in each sector below their neutrino decoupling temperatures, so that $g_\star^{\text{SM}}(T_\nu^{\text{SM}}) = g_\star^{\text{t}}(T_\nu^{\text{t}})$. The neutrino number densities are then

$$\frac{n_\nu^{\text{t}}}{n_\nu^{\text{SM}}} = \left(\frac{T_\nu^{\text{t}} a_\nu^{\text{t}}}{T_\nu^{\text{SM}} a_\nu^{\text{SM}}}\right)^3 = \frac{g_\star^{\text{t}}(T_{\text{decoup}})}{g_\star^{\text{SM}}(T_{\text{decoup}})}. \quad (3.23)$$

For f/v from 3 to 10 and using $T_{\text{decoup}} \sim 2\text{--}6\text{ GeV}$ from section 3.2, we find $g_\star^{\text{t}}(T_{\text{decoup}})/g_\star^{\text{SM}}(T_{\text{decoup}}) \sim 0.8$ and thus arrive at

$$m_\nu^{\text{eff}} \approx 0.8 \left(\frac{f}{v}\right)^n \sum_\alpha m_{\nu_\alpha}^{\text{SM}}, \quad (3.24)$$

where $n = 1$ for Dirac masses and $n = 2$ for Majorana masses.

If they are sufficiently light and hot, the twin neutrinos only affect the CMB as dark radiation and their masses may then only be inferred from tests of the matter power spectrum. However, if heavier and colder, they are better described as a hot dark matter component. Their impact on the CMB is discussed in [34], where the shape of the power spectrum can depend upon the individual neutrino kinetic energies through their characteristic free-streaming lengths. The early Integrated Sachs-Wolfe effect (eISW) is also sensitive to the masses if the neutrinos become non-relativistic during decoupling (thereby affecting the radiation energy density and the growth of inhomogeneities) [33].

There is a significant degeneracy in cosmological fits to the CMB between Ω_m and H_0 (the Hubble constant) [35], where raising the non-relativistic matter fraction, such as with nonrelativistic neutrinos, can be accommodated by a decrease in H_0 (or equivalently, the dark energy density), which keeps the angular diameter distance to the CMB approximately fixed. This degeneracy can be broken by measurements of the baryon acoustic oscillations (BAOs), which are sensitive to the expansion rate of the late universe and provide an independent measurement of Ω_m and H_0 . It is through combination with these results that bounds from Planck on neutrino masses are strongest [2].

3.3.3 Bounds

The authors are unaware of any specialised analysis of the present and projected future cosmological constraints on scenarios with both massless dark radiation and additional light, semi-relativistic sterile neutrinos. In the absence of this, we use bounds from [2] as a rough indication of the present level of sensitivity to these parameters, which we nevertheless expect to be a reliable indication of the (in)viability of this model. The 95% confidence limits on these parameters are $N_{\text{eff}} = 3.2 \pm 0.5$ and $\sum m_\nu < 0.32\text{ eV}$ when each are constrained separately with the other fixed. This, of course, overlooks correlations between the impacts of masses and ΔN_{eff} on the CMB and LSS. Bounds on an additional

sterile neutrino as the only source of dark radiation are also presented with number density, or equivalently, contribution to ΔN_{eff} , left to float. These are similar to the limit on $\sum m_\nu$. It was found in [36] that, allowing $\sum m_\nu$ and m_ν^{eff} to float independently for a single extra sterile neutrino, the bound mildly relaxes to $m_\nu^{\text{eff}} \lesssim 1 \text{ eV}$, although the bound may be stronger depending on the combination of data sets chosen (the lensing power spectrum presently prefers higher neutrino masses and raises the combined bounds if included). Other bounds from LSS on $\sum m_\nu$ exist and are potentially stronger than those placed from the CMB, possibly as low as $m_\nu^{\text{eff}} \lesssim 0.05 \text{ eV}$, again depending on data sets combined (see [37, 38]), although these are subject to greater uncertainties in the inference of the power spectra of dark matter halos from galaxies surveys and the Ly α forest.

It must also be noted that the shape of the CMB temperature anisotropies depends upon both the mass of individual neutrino components (through their free-streaming distance) and their contribution to the energy density of the nonrelativistic matter that does not cluster on small scales. However, it is not expected that improvements in bounds on the former will be made from improved measurements of the primary CMB itself, but rather from weak lensing of the CMB, in conjunction with future measurements from DESI of the BAOs to break degeneracy with Ω_m . The lensing spectrum, like inferences of the matter power spectrum made in galaxy surveys, is expected to measure the suppression of small scale power and therefore to strengthen constraints upon m_ν^{eff} , rather than the individual neutrino masses. One of the goals of CMB-S4 will be the detection of neutrino masses, given the present lower bound $\sum m_\nu \gtrsim 0.06 \text{ eV}$ from oscillations. Projected bounds are as low as $\sim 0.02 \text{ eV}$ [27], although this assumes no extra dark radiation or sterile neutrinos. A projection of the joint bound on N_{eff} (from extra massless dark radiation) and m_ν^{eff} combining improved measurements CMB temperature measurements, lensing and BAOs indicates a limit of $m_\nu^{\text{eff}} \lesssim 0.1 \text{ eV}$ at 1σ [27]. Any contribution from additional states to m_ν^{eff} may therefore be testable and bounded by the excess of the neutrino mass inference over the minimum neutrino mass, although laboratory measurements or measurements of ΔN_{eff} will be required to further ascertain the contribution from the new particles.

Constraints on ΔN_{eff} from improved measurements of the damping tail as part of CMB-S4 are projected to be $\sim 0.02\text{--}0.05$ at 1σ [27]. In the following sections, we use an optimistic estimate of 0.02 for its reach in order to identify as much of the potentially testable parameter space as possible.

To estimate the impact of current and projected CMB limits on the mirror Twin Higgs, we consider two scenarios: the minimal Standard Model neutrino mass spectrum of $m_\nu = [0.0, 0.009 \text{ eV}, 0.06 \text{ eV}]$ and a degenerate spectrum of $m_\nu = [0.1 \text{ eV}, 0.1 \text{ eV}, 0.1 \text{ eV}] / 3$ from [2]. In figure 4 we plot the predictions of the mirror Twin Higgs for ΔN_{eff} and m_ν^{eff} for both types of spectra, as well as for both Dirac and Majorana masses (which scale differently with f/v). As is plainly evident, the mirror Twin Higgs is ruled out cosmologically, no matter the choices of neutrino masses one makes, if only for the presence of the twin photon. In the standard cosmology, the twin sector will have roughly the same temperature as the SM, giving $4.6 \lesssim \Delta N_{\text{eff}} \lesssim 6.3$ for $f/v < 10$, according to the definition of (3.20). This range depends upon f/v through the twin neutrino decoupling temperature (3.18), which determines the extent to which the twin photons are reheated

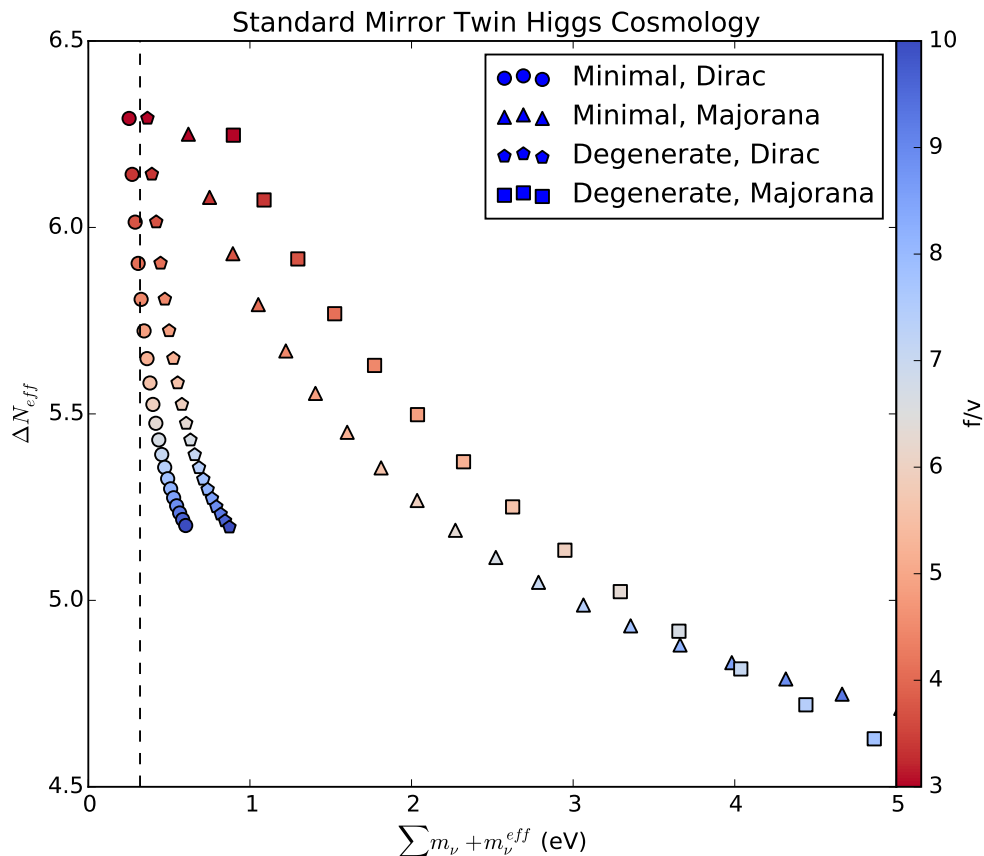


Figure 4. Predicted values of ΔN_{eff} and $\sum m_\nu + m_\nu^{\text{eff}}$ for minimal and degenerate neutrino mass spectra with both Dirac and Majorana masses for f/v from 3 to 10. The Planck 2015 constraint [2] is the dashed line; the corresponding N_{eff} upper bound is well below the bottom of the plot. All points are excluded by the combination of bounds on ΔN_{eff} and $\sum m_\nu + m_\nu^{\text{eff}}$.

relative to the twin neutrinos after twin electron/positron annihilations. This is sufficiently large that even the cold dark matter fraction cannot be adjusted to keep matter-radiation equality fixed, resulting inevitably in changes to the height and shape of the first acoustic peak. The energy density in neutrinos is predicted to be above the present observational upper bounds for most neutrino mass configurations, with the exception of the minimal values permitted by neutrino oscillation measurements with $f/v \lesssim 6$. We therefore discuss cosmological mechanisms in which the twin radiation is diluted to levels compatible with these observational bounds in the subsequent sections of this paper.

4 Reheating by the decay of a scalar field

We now turn to simple scenarios that reconcile the mirror Twin Higgs with cosmological bounds, while taking care to respect the softly-broken \mathbb{Z}_2 symmetry. We begin with the out-of-equilibrium decay of a particle with symmetric couplings to the Standard Model and twin sectors, in which the desired asymmetry is generated kinematically. That is to say, the dimensionless couplings between the decaying particle and the two sectors are equal, and

asymmetric energy deposition into the two sectors is a direct consequence of the asymmetric mass scales. In this respect, the scenario is philosophically similar to N naturalness [39], albeit with a parsimonious $N = 2$ sectors. See also [40, 41] and [24] for other recent related ideas of using long-lived particles for the dilution of dark sectors.

For simplicity, here we will focus on the case of a real scalar X coupled symmetrically to the A and B sector Higgs doublets. Due to the difference in masses between the sectors after electroweak symmetry breaking, simple kinematic effects give X a larger branching ratio into the Standard Model. This occurs over a range of X masses within a few decades of the weak scale. If X decays out-of-equilibrium below the decoupling temperature of the two sectors, this injects different amounts of energy into the two sectors, effectively suppressing the temperature of the twin sector relative to the Standard Model. This relative cooling suppresses the contribution of the light degrees of freedom of the mirror Twin Higgs to below cosmological bounds. Insofar as the asymmetry is driven entirely by kinematic effects arising from $v \ll f$, the resulting temperature inequality between the two sectors is proportional to powers of v/f .

The requisite suppression of the twin sector temperature relative to the Standard Model temperature necessitates that the X dominate the cosmology before it decays. Our main discussion will follow the simplest case of an X which dominates absolutely before it decays, comprising all of the energy density of the universe and effectively acting as a ‘reheaton’. Afterwards, we will discuss the possibility of a ‘thermal history’ for X — a scenario where X is in thermal equilibrium with the two sectors, then chemically decouples at some high temperature and grows to dominate the cosmology before it decays. This scheme will result in additional stringent constraints on the viable parameter space.

4.1 Asymmetric reheating

A \mathbb{Z}_2 -even scalar X which is a total singlet under the SM and twin gauge groups admits the renormalisable interactions

$$V \supset \lambda_x X(X + x) \left(|H_A|^2 + |H_B|^2 \right) + \frac{1}{2} m_X^2 X^2, \quad (4.1)$$

where m_X is the mass of X (neglecting corrections from mixing that will be shown below to be tiny), λ_x is a dimensionless coupling and x is a dimensionful parameter, which one may imagine identifying as a vacuum expectation value (vev) of X in an UV theory. Note that these interactions preserve the accidental $SU(4)$ symmetry of the Twin Higgs. The X field may additionally possess self-interactions, which we omit here as they do not play a significant role in what follows.

The interactions in (4.1) allow X to decay into light states in the Standard Model and twin sectors. If X reheats the universe through out-of-equilibrium decays, the reheating temperatures of the two sectors will be determined by its partial decay widths, assuming that the decay products do not equilibrate. In the instantaneous decay approximation, X decays when the Hubble parameter falls to its decay rate $\Gamma_X \sim H$. As we will show in section 4.2, in order to evade cosmological constraints we need the X to decay mostly into the SM, so we may estimate $\Gamma_X \sim \Gamma(X \rightarrow \text{SM})$. Then the energy that was contained in

the X is transferred into radiation energy density, with the resulting temperature of the radiation given by (see [42])

$$T \sim 1.2 \sqrt{\frac{\Gamma_X M_{\text{pl}}}{\sqrt{g_\star}}} \quad (4.2)$$

where g_\star is the effective number of relativistic degrees of freedom, as defined in section 3, of the particles that are being reheated. Our numerical calculation of the reheating temperature, which will be presented in section 4.2.2, indicates that the approximation $T \sim 0.1 \sqrt{\Gamma_X M_{\text{pl}}}$ reliably reproduces the reheating temperature over the range of interest.

As shown in section 3.2, the two sectors thermally decouple when the temperature falls below $T_{\text{decoup}} \sim 1 \text{ GeV}$, so reheating must take place to below this temperature. At even lower temperatures, big bang nucleosynthesis (BBN) places strong constraints on energy injected into the SM at temperatures below $\mathcal{O}(1-10) \text{ MeV}$ [43]. Requiring that the SM reheating temperature is above $\sim 10 \text{ MeV}$, these constraints on the SM reheating temperature become constraints on the decay rate of the X into the SM, which in the above approximation becomes

$$5 \times 10^{-21} \text{ GeV} \lesssim \Gamma_X \lesssim 3 \times 10^{-16} \text{ GeV}. \quad (4.3)$$

This then constrains the couplings λ_x and x of the X to the Higgs sector. Importantly, it means that X must couple very weakly, in order to be long-lived enough to reheat to a low temperature, as will be shown below.

The asymmetry in partial widths arises from different effects depending upon the mass of X . For masses below the SM Higgs threshold, it is predominantly differences in mass mixing with the two Higgs doublets that produces the asymmetry, where the size of the mixing angles determines the effective coupling of X to the SM and twin particles and therefore its branching fractions. For masses below the twin scale, the relative size of the mixing scales inversely with the vevs in each sector. Thus the hierarchy $v \ll f$ already present in the Higgs sector can automatically give rise to a hierarchy in partial widths. Note that additional threshold effects can enhance the asymmetry further, in particular when X has mass above threshold for a significant decay channel in the SM, but below the corresponding mass threshold in the twin sector. Decays into on-shell Higgses complicate this picture further. In what follows, we first give an analytic calculation of the mass mixing effect, then present a more precise calculation of the decay widths into each sector.

To lowest order, X decays via its interactions with the SM and twin Higgs, and only to other fermions and gauge bosons through its mass mixing with the Higgs scalars. Expanding the X potential after the $SU(4)$ is spontaneously broken, the mixing term between X and h_A in the scalar mass matrix is $\sqrt{2}\lambda_x x v_A$, while that between X and h_B is $\sqrt{2}\lambda_x x v_B$. The h_A and h_B components of the X mass eigenstate, which we denote respectively as δ_{XA} and δ_{XB} , can then be determined. The expressions for the mixing angles are in general complicated, but they simplify in limits $m_X < f$ and $m_X \gg f$:

$$(\delta_{XA}, \delta_{XB}) \approx \begin{cases} \frac{4\lambda_x x v_A}{m_X^2 - m_h^2} \left(\frac{1}{\sqrt{2}}, \frac{v_A}{f} \right) & m_X < f \\ \frac{\lambda_x x f}{m_X^2} \left(\frac{\sqrt{2}v_A}{f}, 1 \right) & m_X \gg f \end{cases} \quad (4.4)$$

to lowest order in $(v/f)^2$ and κ/λ . The partial width for the decay of X into SM states (excluding the Higgs) is

$$\Gamma(X \rightarrow \text{SM}) \approx |\delta_{XA}|^2 \Gamma_h(m_h = m_X), \quad (4.5)$$

where $\Gamma_h(m_h = m_X)$ denotes the decay width of a SM Higgs if it were to have mass m_X . Note that the Higgs partial width must be computed using the vev $v_A \approx v/\sqrt{2}$ to determine the masses and couplings of the SM particles. The partial width of the X into twin states is computed the same way using δ_{XB} and the vev $v_B \approx f/\sqrt{2}$.

From the mixing angles (4.4), it is already apparent over what mass range asymmetric reheating from X decays will work. These give

$$\frac{\Gamma(X \rightarrow \text{SM})}{\Gamma(X \rightarrow \text{Twin})} \sim \begin{cases} f^2/v_A^2 \gg 1 & m_X < f \\ v_A^2/f^2 \ll 1 & m_X \gg f. \end{cases} \quad (4.6)$$

Thus when the mass of X is less than the twin scale, the Standard Model will be reheated to a higher temperature than the twin sector, but in the large mass limit this mechanism works in the opposite direction and would appear to lead to preferential reheating of the twin sector.

More precise statements about the relative branching ratios and resulting temperatures require additional care. In addition to decaying through mass mixing, X can decay into the Higgs mass eigenstates themselves if above threshold. As the energy is ultimately transferred to the SM and twin sectors, we then need to consider how these states decay and account for the further mixing of the Higgs mass eigenstates into Higgs gauge eigenstates.

For $m_X > 2m_h$, decay can occur into the lighter (SM-like) Higgs mass eigenstate h with partial width

$$\Gamma(X \rightarrow hh) \approx \frac{\lambda_x^2 x^2}{16\pi m_X} \sqrt{1 - \left(\frac{2m_h}{m_X}\right)^2}. \quad (4.7)$$

Similarly, for $m_X > 2m_H$, decays can proceed into HH with a similar partial width, but with the h mass replaced with that of the H . Above the intermediate threshold $m_X > m_h + m_H$, there is also the mixed decay

$$\Gamma(X \rightarrow hH) \approx \frac{\lambda_x^2}{2\pi m_X} \sqrt{1 - \left(\frac{m_H + m_h}{m_X}\right)^2} (f\delta_{AX} + 2v_A\delta_{BX})^2. \quad (4.8)$$

Here, $\delta_{AX} \approx -\delta_{hA}\delta_{XA} - \delta_{hB}\delta_{XB}$ is the component of the h_A gauge eigenstate in the X mass eigenstate and $\delta_{BX} \approx \delta_{hB}\delta_{XA} - \delta_{hA}\delta_{XB}$ is the corresponding component of the h_B gauge eigenstate, where δ_{hA} and δ_{hB} are, respectively, the components of the SM Higgs in the h_A and h_B gauge eigenstates to zeroth order in λ_x . Combining all ingredients, this decay width is of order $\lambda_x^4 x^2$. Since it is only the total decay width that is constrained to be small by the demand that the SM reheating temperature lie in the required window, this fixes only a product of λ_x and x . If $x \sim v$, then the mixed decay to hH is effectively second order in the small coupling λ_x^2 and can be neglected relative to the other partial widths. Conversely if $x \ll v$, then λ_x is much larger and this decay cannot be neglected.

In what follows we will work in the region of parameter space where mixed decays to hH are negligible.

The rate of heat flow into each sector may be well approximated by adding the decay rates of X into each channel and weighting these by the fraction of energy transferred into the particular sector. Of course, when X decays into Higgs particles, these in turn decay out of equilibrium into both the Standard Model and twin sectors. As the Higgs decays are almost instantaneous, the fraction of energy transferred into each sector is simply that carried by the Higgs decay products multiplied by their branching fractions for each sector. The total rate at which X particles are transferred into the SM plasma is

$$W(X \rightarrow SM) \approx \Gamma(X \rightarrow SM) + \Gamma(X \rightarrow hh)Br(h \rightarrow SM) + \Gamma(X \rightarrow HH)(Br(H \rightarrow SM) + Br(H \rightarrow hh)Br(h \rightarrow SM)). \quad (4.9)$$

The corresponding rate for energy deposition into the twin sectors is simply given by the replacement of $SM \mapsto \text{Twin}$. The first term is the rate at which X decays directly into the SM through mass mixing with the Higgs. The second is the fraction of X energy that is transferred into lighter Higgs states that subsequently decay into the SM. The third is the analogous term for decays into the heavy Higgs, where cascade decays of the H into the h and subsequently other SM particles must be included. Note that decays of the heavy Higgs into the light Higgs make up a majority of decay width, because of the large quartic coupling required for the twin Higgs potential.

Below the hh threshold, it is possible for X to decay via one on-shell and one off-shell Higgs boson. The partial width for off-shell Higgs production was calculated for $X \rightarrow hh^* \rightarrow hb\bar{b}$ and found to be negligible compared to two-body decays through mass mixing and so we omit three-body decay widths in what follows.

Ultimately, the complete partial widths for the decay of X into the Standard Model and twin sectors includes the sum of decays into Higgs bosons h and H and direct decays into the fermions and gauge bosons of the two sectors. We compute the latter to an intended level of accuracy of $\sim 10\%$ (including, e.g., NLO QCD corrections to decays into light-flavor quarks), mostly following [44]. The resulting partial widths into the Standard Model and twin sectors are shown as a function of m_X in figure 5 with the ratio of branching fractions displayed in figure 6.

Over much of the space below the Higgs mass, the branching ratio exhibits the expected $(f/v)^2$ scaling from the mass mixing. Below ~ 40 GeV, suppression of the twin partial width arises because the twin bottom quark pair production threshold is crossed. As m_X nears m_h , the SM branching fraction grows by ~ 4 orders of magnitude as the WW^* , ZZ^* , and then WW and ZZ decays go above threshold. Since the analogous thresholds are at much higher energies in the twin sector, the enhancement is not paralleled by decays into the twin sector until m_X is close to the twin scale. There is therefore a large range of masses $m_h \lesssim m_X \lesssim m_H$ over which the SM branching fraction dominates by several orders of magnitude.

Above the $X \rightarrow hh$ threshold, the ratio of decay widths is roughly constant in mass up to the HH threshold. The twin sector decay rate is dominated by decays of on-shell light

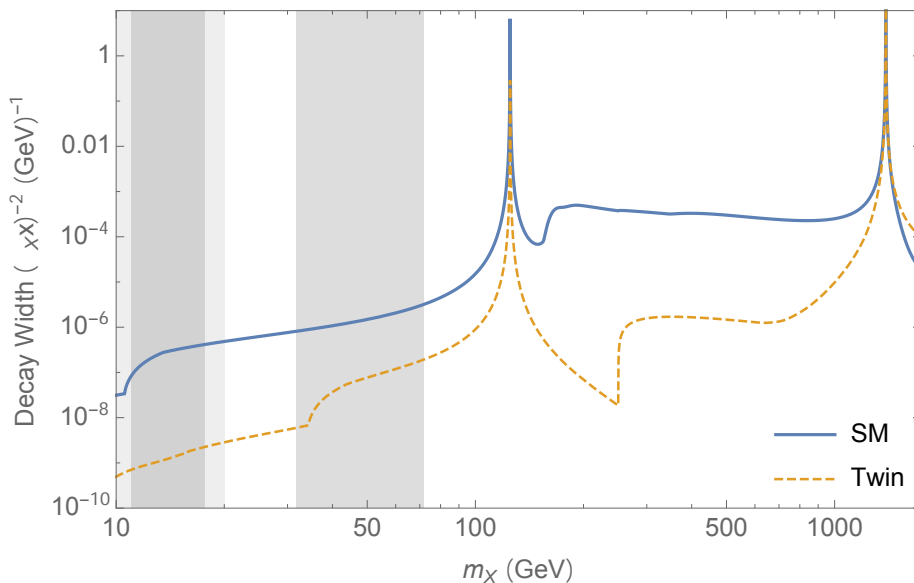


Figure 5. The partial widths of the X into the SM (solid blue line) and twin sector (dashed orange) for $f/v = 3$ in units of $(\lambda_{xx})^2$. The light gray bands indicate regions of QCD-related uncertainty in the SM calculation, while the darker gray bands indicate the corresponding regions of uncertainty for the twin calculation.

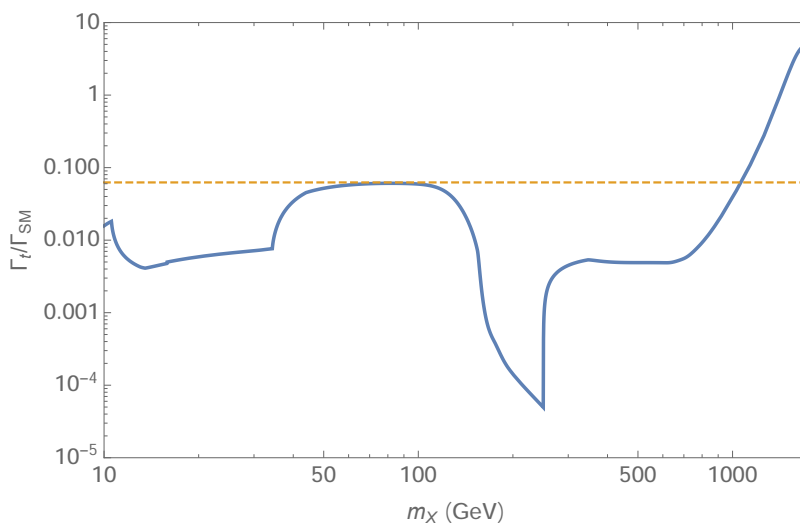


Figure 6. The ratio of branching fractions of the X into the SM and twin sectors at $f/v = 3$. The dashed line gives the expected $(v/f)^2$ scaling from the mass mixing; deviations are due to various mass threshold effects.

Higgs into twin states, $\Gamma(X \rightarrow \text{Twin}) \approx \Gamma(X \rightarrow hh)\text{Br}(h \rightarrow \text{Twin}) \propto 1/m_X$ as in (4.7). If the SM were also predominantly reheated through this channel, then the ratio of branching fractions would again be approximately $\delta_{hA}^2/\delta_{hB}^2 \approx (f/v)^2$. However, the SM decay width also receives a larger contribution from decays through mass mixing between the X and the Higgs gauge eigenstates.

For masses $m_X > 2m_h$, decays through mass mixing are dominated by the SM WW and ZZ channels. In this mass region, the decay rate of a Higgs into longitudinally polarized vector bosons scales as $\Gamma(h \rightarrow WW, ZZ) \sim m_X^3$, but the mixing angle scales as $\delta_{AX}^2 \sim 1/m_X^4$ (as in (4.4)), resulting in the same $\sim 1/m_X$ scaling and thus a roughly constant ratio in this range of masses. Near $m_X \sim 1$ TeV, decays into twin vector bosons through mass mixing begin to dominate, and there is no favourable asymmetry in the branching fractions, as discussed in this section. Even at higher masses, the effects of heavy Higgs decays into light Higgs do not compensate sufficiently, as this partial width scales with m_X in the same way as the partial width for longitudinally polarised weak bosons.

The constraint on the decay width from the required reheating temperature (4.3) translates into a constraint on the size of the coupling $\lambda_x x$. For $m_X \gtrsim m_h$, this gives $10^{-8.5} \text{ GeV} \lesssim \lambda_x x \lesssim 10^{-6} \text{ GeV}$, while for lower masses, this range increases to $10^{-7} \text{ GeV} \lesssim \lambda_x x \lesssim 10^{-5.5} \text{ GeV}$ at $m_X \sim 20 \text{ GeV}$.

The gray bands in figure 5 highlight regions where our analytic estimates of the partial widths encounter enhanced uncertainties arising from the bottom and charm thresholds in both sectors. Over most of these ranges, we estimate the size of these uncertainties to be either $\sim 10\%$ or confined to very small subregions. The thicknesses of these bands have been chosen conservatively, and ultimately the branching ratios should be accurate to within a factor of $\pm \Lambda_{QCD}$ of the bottom and charm mass thresholds. In particular, the prescription of [45] has been followed for approximating the bottom partial width close to the open flavour threshold. Resonant decay into gluons from bottomonia mixing has been neglected, although these resonant mass ranges are expected to be only $\sim \text{MeV}$ wide at the CP-even, spin-0 bottomonia masses $m_X = m_{\chi_{bi}}$ (see [45] and [46]). It should be noted, however, that at temperatures above that of the QCD phase transition, the quark decay products behave differently compared to that expected in a low temperature environment. In particular, for hot enough temperatures, the b or c quarks may not hadronise and the partonic partial widths may more reliable. The applicability of the treatment of the flavour thresholds used here may therefore not be valid if the decay occurs in the hot early universe. However, it is only very close to the threshold itself (within several GeV) that this uncertainty becomes significant. Finally, quark masses have been neglected in the gluon partial width. For m_X close to the flavour thresholds, this approximation breaks down, but the gluon branching fraction is only $\sim 10\%$ and so the error does not contribute to the uncertainty of the total width by more than this order (it is this uncertainty that is responsible for most of the extension of the length of the gray bands about the flavour threshold).

Close to the charm threshold, the analogous uncertainties are even more poorly understood. Below the charm threshold, hadronic decays of a light scalar are highly uncertain (see [47] for discussion). We avoid these regions altogether by restricting our considerations to m_X roughly above the twin charm threshold. Note that below the SM charm threshold, the smaller decay rate of a Higgs-like scalar necessitates larger couplings $\lambda_X x$ for X to have a lifetime within the required reheating window. The larger couplings then imply potentially stronger constraints from invisible mesonic decays. See [46–48] for further discussion and recent analysis of the pertinent experimental constraints.

Taken together, the results in figures 5 and 6 bear out the expectation that a scalar X with symmetric couplings to the Standard Model and twin sectors may nonetheless inherit a large asymmetry in partial widths from the hierarchy between the scales v and f . Across a wide range of masses m_X , the asymmetry is proportional to (or greater than) v^2/f^2 , tying the reheating of the two sectors to the hierarchy of scales.

Before proceeding to our computation of cosmological observables, we comment on an alternative variation on the reheating mechanism presented here that involves having X odd under the twin parity. This permits two renormalisable interactions with the Higgses to give a Higgs potential of the form:

$$\mathcal{V} \supset m_0^2 (|H_A|^2 + |H_B|^2) + \lambda_0 (|H_A|^4 + |H_B|^4) + \epsilon X^2 (|H_A|^2 + |H_B|^2) + \tilde{\epsilon} X (|H_A|^2 - |H_B|^2). \tag{4.10}$$

If X then acquires a vev at some scale, it may be possible to arrange for the resulting spontaneous breaking of the \mathbb{Z}_2 to give that required in the Higgs potential. However, we find that, in order for X to be long-lived and reheat the universe, its couplings to the Higgs must be highly suppressed and therefore that the resulting vev of X required to explain the soft \mathbb{Z}_2 -breaking in the Higgs potential must be many orders of magnitude above the twin scale. If this is to be identified with the characteristic mass scale of X , then a UV-completion of the twin Higgs is required for anything further to be said of the prospects of this possibility. However, if such a UV completion has similar structure to the couplings in (4.10), then asymmetric reheating may require a cancellation between the odd and even couplings of X to the Higgs potential in order to suppress its twin-sector branching fraction (because the odd coupling appears with opposite signs in the coupling between X and the h_A and h_B states). We do not consider this possibility further.

4.2 Imprints on the CMB

For appropriate values of m_X , the out-of-equilibrium decay of X reheats the two sectors to different temperatures and effectively dilutes the energy density in the twin sector. We obtain an analytic estimate of the effects of the X decay on the number of light degrees of freedom observed from the CMB by approximating both the decay of X and the decoupling of species as instantaneous in section 4.2.1. We then demonstrate that this estimate is reliable over most of the parameter space of interest with a numerical calculation in section 4.2.2. In section 4.2.3 we consider neutrino masses and their joint constraints with N_{eff} .

4.2.1 Analytic estimate of N_{eff}

If X dominates the energy density of the universe and then decays, depositing energy ρ_{SM} and ρ_t into the SM and twin sectors respectively, then the temperature ratio is determined by

$$\frac{\rho_t}{\rho_{\text{SM}}} = \frac{g_\star^t(T_{\text{reheat}}^t)}{g_\star^{\text{SM}}(T_{\text{reheat}}^{\text{SM}})} \left(\frac{T_{\text{reheat}}^t}{T_{\text{reheat}}^{\text{SM}}} \right)^4 \approx \frac{\Gamma(X \rightarrow \text{Twin})}{\Gamma(X \rightarrow \text{SM})}, \tag{4.11}$$

where $T_{\text{reheat}}^{\text{SM}}$ and T_{reheat}^t are the reheating temperatures for each sector, while g_\star^{SM} and g_\star^t are the SM and twin effective number of relativistic degrees of freedom, respectively. We

have assumed that the two sectors are cool enough that they have already decoupled. We point out that not only does the number of effective degrees of freedom in each sector need to be evaluated at the temperature of that sector, but that g_\star^t and g_\star^{SM} differ as functions of temperature due to the differences in the spectra of the sectors, as seen in figure 1. As is well-known [42], reheating is a protracted process that occurs over a time-scale given by the lifetime of the reheaton. During this time, the temperature of the plasma cools slowly because, while the energy is being replenished by the decay of the reheaton, it is simultaneously diluted and redshifted with the expansion of the universe. It is assumed in (4.11) that any primordial energy density in either sector is subdominant.

The temperatures of both sectors then redshift in the same way, so the only additional differences between their temperatures arise from changes to the effective number of degrees of freedom in each sector. By conservation of comoving entropy within each sector, each evolves as $T_{\text{eq}}^i/T_{\text{reheat}}^i = (g_\star^i(T_{\text{reheat}}^i)/g_\star^i(T_{\text{eq}}^i))^{1/3} a(T_{\text{reheat}})/a(T_{\text{eq}})$ where T_{eq}^i is the temperature of the sector at matter-radiation equality, which the CMB probes as explained in section 3.3, and $a(T)$ is the scale factor as a function of temperature. In the mirror Twin Higgs model, the two sectors have the same number of light degrees of freedom at recombination (three neutrinos and a photon, assuming that the neutrinos are still relativistic), so

$$\left(\frac{T_{\text{eq}}^t}{T_{\text{eq}}^{\text{SM}}}\right)^4 = \left(\frac{T_{\text{reheat}}^t}{T_{\text{reheat}}^{\text{SM}}}\right)^4 \left(\frac{g_\star^t(T_{\text{reheat}}^t)}{g_\star^{\text{SM}}(T_{\text{reheat}}^{\text{SM}})}\right)^{4/3} = \frac{\Gamma(X \rightarrow \text{Twin})}{\Gamma(X \rightarrow \text{SM})} \left(\frac{g_\star^t(T_{\text{reheat}}^t)}{g_\star^{\text{SM}}(T_{\text{reheat}}^{\text{SM}})}\right)^{1/3}. \quad (4.12)$$

As our range of reheat temperatures encompasses the QCD phase transitions of both sectors, the factors of g_\star can be important.

Given the temperatures of the two sectors after X decays, we can obtain a simple estimate of the contribution to N_{eff} that neglects the impact of masses of the twin neutrinos discussed in section (3.3.1),

$$(\Delta N_{\text{eff}})_{m_\nu=0} = \frac{4}{7} \left(\frac{11}{4}\right)^{4/3} g_\star^{\text{SM}}(T_{\text{eq}}^{\text{SM}}) \frac{\rho_t(T_{\text{eq}}^t)}{\rho_{\text{SM}}(T_{\text{eq}}^{\text{SM}})} \quad (4.13)$$

$$\approx 7.4 \times \frac{\text{Br}(X \rightarrow \text{Twin})}{\text{Br}(X \rightarrow \text{SM})} \left(\frac{g_\star^t(T_{\text{reheat}}^t)}{g_\star^{\text{SM}}(T_{\text{reheat}}^{\text{SM}})}\right)^{1/3}. \quad (4.14)$$

In this limit the most recent Planck data give a 2σ bound of $\Delta N_{\text{eff}} \lesssim 0.40$ assuming pure $\Lambda\text{CDM}+N_{\text{eff}}$ [2]. This translates into the requirement $\frac{\rho_t(T_{\text{eq}}^t)}{\rho_{\text{SM}}(T_{\text{eq}}^{\text{SM}})} \approx \frac{\Gamma(X \rightarrow \text{Twin})}{\Gamma(X \rightarrow \text{SM})} \lesssim 0.05$, ignoring possible differences in g_\star .

Of course, as discussed in section 3, the twin neutrino masses are relevant at the temperature of matter-radiation equality, so we can obtain a more meaningful estimate of ΔN_{eff} using the results of section 3.3.1 evaluated at the twin temperature determined above:

$$\Delta N_{\text{eff}} = \left(\frac{11}{4}\right)^{4/3} \frac{120}{7\pi^2 (T_{\text{eq}}^{\text{SM}})^4} \left(\rho_\gamma^t(T_{\text{eq}}^t) + \sum_\alpha 3w_{\nu_\alpha}^t(T_{\text{eq}}^t) \rho_{\nu_\alpha}^t(T_{\text{eq}}^t) \right) \quad (4.15)$$

$$T_{\text{eq}}^t = T_{\text{eq}}^{\text{SM}} \left(\frac{\Gamma(X \rightarrow \text{Twin})}{\Gamma(X \rightarrow \text{SM})}\right)^{1/4} \left(\frac{g_\star^t(T_{\text{reheat}}^t)}{g_\star^{\text{SM}}(T_{\text{reheat}}^{\text{SM}})}\right)^{1/12} \quad (4.16)$$

with $T_{\text{eq}}^{\text{SM}} \approx 0.77 \text{ eV}$ [2] the photon temperature. While the right-hand side of this equality has implicit dependence on T_{eq}^{t} through g_{\star}^{t} , this is only important if the reheating occurs between the SM and twin QCDPTs and the neglecting of the factors of g_{\star} is otherwise reliable. With the further inclusion of Standard Model neutrino masses or an extra sterile neutrino, the bound described above weakens to $\Delta N_{\text{eff}} \lesssim 0.7$. As discussed in section 3.3.3, we are not aware of any analyses specific to our model involving both pure dark radiation and three sterile neutrinos with masses of order the photon decoupling temperature of the CMB and possibly cooler temperatures. In the absence of such an analysis, we use the inequality $\Delta N_{\text{eff}} \lesssim 0.7$ to indicate where the present CMB measurements are likely to constrain the light degrees of freedom of this model, leaving a more detailed analysis of the CMB constraints as future work. In this case, the bound on the decay width ratio is $\frac{\Gamma(X \rightarrow \text{Twin})}{\Gamma(X \rightarrow \text{SM})} \lesssim 0.09$. The next generation of CMB experiments are projected to strengthen this constraint to $\Delta N_{\text{eff}} \lesssim 0.02$ at the 1σ level [49].

4.2.2 Numerical calculation of N_{eff}

A more precise study of the effect of X decay on the number of effective neutrino species at recombination may be performed by numerically solving a system of differential equations for the entropy in X and the two sectors as a function of time. Following the analysis of chapter 5.3 of [42] we have

$$H = \frac{1}{a} \frac{da}{dt} = \sqrt{\frac{1}{3M_{\text{Pl}}^2} (\rho_X + \rho_{\text{SM}} + \rho_t)} \quad (4.17)$$

$$\frac{d\rho_X}{dt} + 3H\rho_X = -\Gamma_X \rho_X \quad (4.18)$$

$$\rho_i = \frac{3}{4} \left(\frac{45}{2\pi^2 g_{\star}^i} \right)^{1/3} S_i^{4/3} a^{-4} \quad (4.19)$$

$$S_i^{1/3} \frac{dS_i}{dt} = \left(\frac{2\pi^2 g_{\star}^i}{45} \right)^{1/3} a^4 \left(\rho_X \Gamma_{X \rightarrow i} + \frac{dq_{j \rightarrow i}}{dt} \right), \quad (4.20)$$

where S_i are comoving entropy densities and it has been assumed that X is cold by the time it decays so that $\rho_X = m_X n_X$ with number density n_X (this is reliable as we only consider $m_X > 10 \text{ GeV}$, which is above the decoupling temperature of $\sim 1 \text{ GeV}$). The rate of heat flow from sector j to i per proper volume, $\frac{dq_{j \rightarrow i}}{dt}$, is defined in (3.6). To account for the temperature-dependence of the effective number of relativistic degrees of freedom in each sector, these equations are solved iteratively in the profiles of $g_{\star}^i(T^i)$.

The equations are solved in three stages: before, during and after the decoupling of the SM and twin sectors. The ratio f/v is fixed to 4 for this analysis. Initial conditions were chosen with $\rho = 10^{-12} \rho_X$, for combined SM and twin energy densities ρ . However, it is only the requirement that the initial energy density of X dominates over that of the SM and twin sectors that is important for simulating the cosmology over the times of interest here, as the entirety of the latter is then generated by the subsequent decay. The results close to the decoupling and reheating epochs are otherwise insensitive to the initial conditions and ultimately match onto the standard outcome [42] expected by equating the Hubble

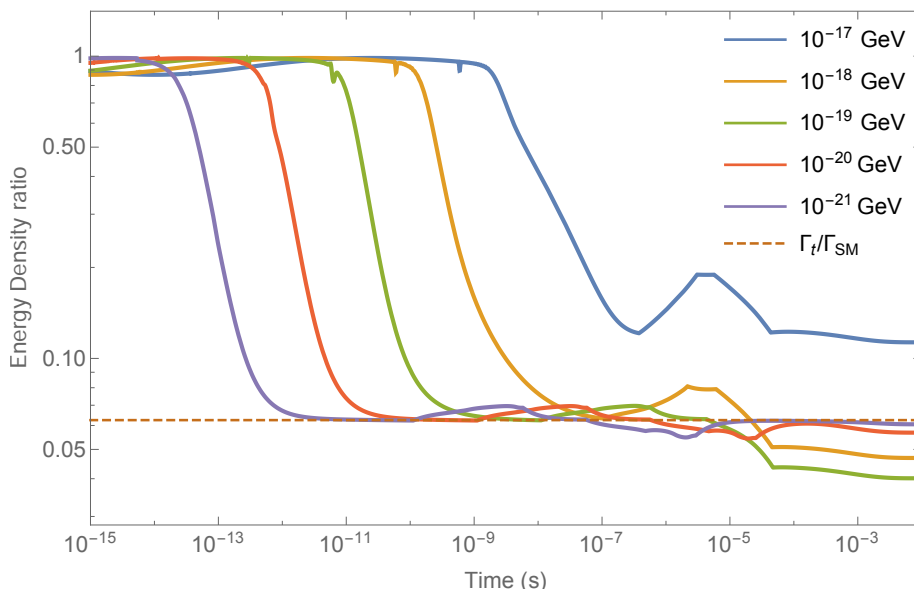


Figure 7. Ratio of twin to SM energy densities throughout decoupling and reheating, for different decay rates Γ_X . The dashed line corresponds to the prediction of from the ratio of decay widths, here selected to be $1/16$.

rate with the decay rate of X . The sectors are assumed to be in thermal equilibrium and sharing entropy until a temperature of 10 GeV , below which they are evolved separately with the heat flows $\frac{dq_{i \rightarrow j}}{dt}$ switched on. Elastic scatterings were neglected from the heat flow rate to accelerate the computation. It was verified for the results found below that their contribution to the heat flow was always $\lesssim 10\%$ while the heat flow was itself not dominated by the Hubble rate. Heat flow was switched off again once the twin temperature reaches 0.1 GeV , by which time thermal decoupling is long-since complete, and the sectors are subsequently evolved separately. Again, although the strengthening of the colour force and the QCDPT make the perturbative tree-level computation of the scattering rates unreliable at temperatures below $\sim 1\text{ GeV}$, as found in section 3.2 and also in the results below, the sectors decouple above these temperatures. Notably, the impact of X on the expansion rate causes decoupling to occur at slightly hotter temperatures than expected from the analysis of section 3.2 for the decoupling in the standard cosmology.

The ratio of energy densities in each sector determines N_{eff} , from (4.15). A plot of this ratio over time is shown in figure 7, with the expectation under the approximations of the previous section shown as well. This approximation is reliable as long as the lifetime of X is much longer than the temperature at which decoupling concludes, here $\sim 1\text{ GeV}$. The larger asymptotic value of the ratio of the blue line arises because the lifetime lies close to the decoupling period, so that a significant fraction of the energy is transferred while the sectors are thermalised or partially thermalised and does not contribute toward asymmetric reheating. Equivalently, as will be discussed below, insufficient time elapses between decoupling and reheating for the twin energy density to dilute and be repopulated by the decays to the level predicted by (4.11). The subsequent bump represents the period

between the reheating of the twin sector by its QCD phase transition followed by that of the SM. The green and orange lines correspond to reheating temperatures that lie between SM and twin QCD phase transitions. In these cases, the reheating of the SM from the subsequent SM QCD phase transition raises its energy density relative to the twin sector above that expected from the ratio of branching fractions. As this occurs after the lifetime of the reheaton, the estimate of the reheating temperatures presented in (4.12) is still good as subsequent changes in the ratio due to the evolution of g_* are accounted for in our analysis of the reheating scenarios.

The steep drop in the energy density ratio corresponds to the brief period during which the energy density of the twin sector present at decoupling dilutes and redshifts, which continues until it reaches a comparable size to the energy density that is being replenished by reheating. If the twin-sector branching fraction is highly suppressed, as can occur in the “valley” region in figure 6 with $m_h \lesssim m_X \lesssim 2m_h$, then a longer time is required for this to happen, especially close to the decay epoch where the diminishing of the X population also contributes to a reduced reheating rate. These effects can prolong the time required for the energy density ratio to converge to the asymptotic prediction of (4.11).

Contour plots of ΔN_{eff} as a function of m_X and f/v appear in figure 8, along with current and predicted bounds using the analytic results of section 4.2.1. The minimum neutrino mass configuration with Dirac masses has also been assumed, although the results are relatively insensitive to this provided that the twin neutrino masses are not well above the eV scale. A SM reheating temperature of 0.7 GeV has been assumed. At this temperature, we have verified using the numerical calculation of section 4.2.2 that the twin sector reheating temperature is always roughly above the twin neutrino decoupling temperature over the parameter space of the figure, ensuring that the neutrinos thermalise once produced in the decays and hence that the predictions of section 4.2.1 are valid. A treatment of the case in which the twin neutrinos are produced below their decoupling temperature is beyond the scope of this analysis, but would involve the computation of the phase space spectrum of the neutrino decay products of the X .

Also, as discussed in section 3.2, a large temperature difference may partially relax back if reheating occurs close to sector decoupling. However, a reliable calculation of the heat flow at the temperatures of interest here must incorporate non-perturbative effects. We do not perform such a computation, but note that, at a slightly higher SM reheating temperature of 2 GeV where this computation is more reliable, ΔN_{eff} in figure 8 can be raised by up to an order of magnitude in the region with $f/v \lesssim 4$ and $150 \text{ GeV} \lesssim m_X \lesssim 200 \text{ GeV}$, notably where the twin sector partial width is suppressed relative to the SM by several orders of magnitude. The resulting ΔN_{eff} prediction is, nevertheless, still out of observable reach. At the lower SM reheating temperature assumed in figure 8, it is expected that decoupling will be further advanced and the enhancement in ΔN_{eff} would be weaker.

We emphasize that, if the lifetime of X is sufficiently close to the time of decoupling, or equivalently, that the reheating temperature is sufficiently close to the decoupling temperature, then the residual twin energy density left-over may be comparable to or greater than that regenerated by reheating. Consequently, the suppression in ΔN_{eff} would be less than

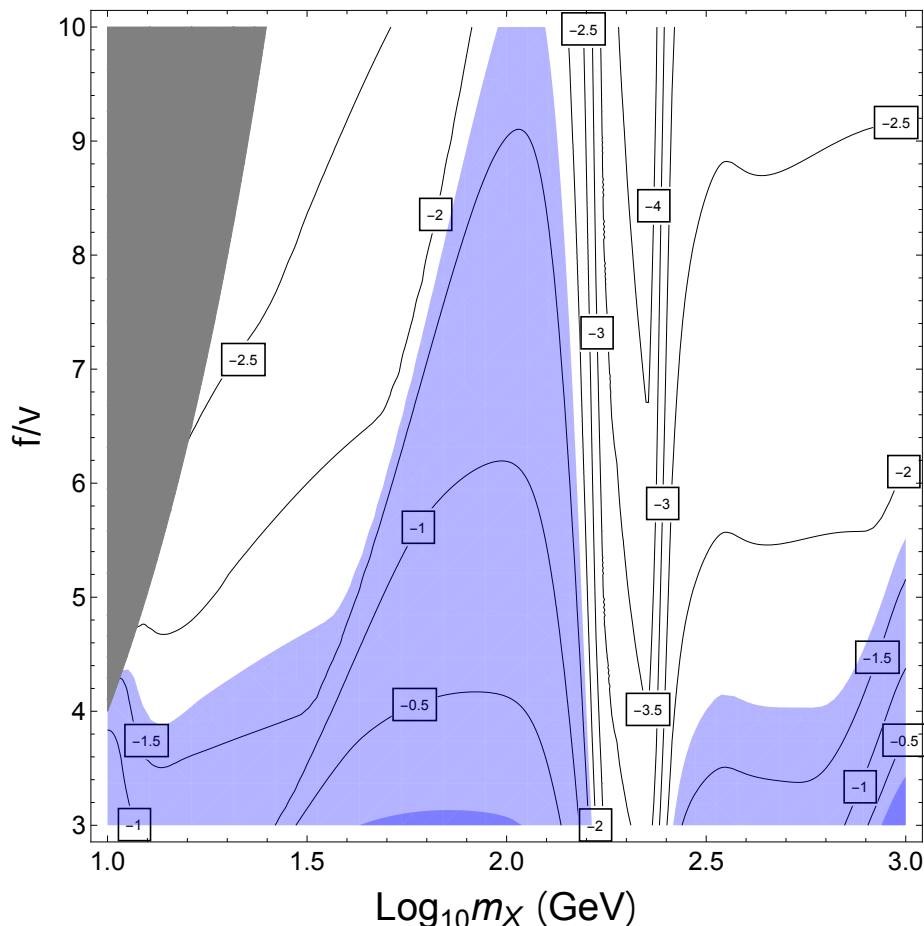


Figure 8. Contours of $\log_{10} \Delta N_{\text{eff}}$ as a function of m_X and f/v , for $T_{\text{reheat}}^{\text{SM}} = 0.7 \text{ GeV}$. The dark blue region is in tension with Planck, while the light blue region will be tested by CMB-S4. Gray regions are where the X mass is below the twin charm threshold and our calculation of the twin sector partial width is unreliable.

that predicted in (4.12). In this respect, the projection of figure 8 should be regarded as a lower bound on ΔN_{eff} . In the regions of high suppression, such as the “valley” region, the full asymmetry may not be generated before the complete decay of X when the reheating temperature is of similar order as the decoupling temperature. In particular, for the reheating temperature chosen here of 0.7 GeV and branching fraction $\text{Br}(X \rightarrow \text{Twin}) \sim 10^{-5}$, the numerical calculation of the energy density ratio saturates at $\sim 4 \times 10^{-5}$. We do not include this effect in figure 8 as its only impact is to mildly shift the unobservably small $\Delta N_{\text{eff}} = 10^{-4}$ contour. Lower reheating temperatures would agree with the prediction of (4.11) were it not for the caveat that the twin neutrinos may be produced out of equilibrium. However, this minimum value at which ΔN_{eff} is saturated can grow significantly with hotter reheating temperatures upon which it is highly dependent.

CMB-S4 observations will be able to probe a large portion of the most natural parameter space, save the region $m_h \lesssim m_X \lesssim 2m_h$ where decays into the Standard Model dominate well beyond the ratio f^2/v^2 , as previously discussed. Significantly, precision

Higgs coupling measurements at the LHC are unlikely to probe the mirror Twin Higgs model beyond $f \sim 4v$, so that the observation of additional dark radiation may be the *first* signature of a mirror Twin Higgs.

4.2.3 Neutrino masses

In addition to the bounds on N_{eff} , we must also respect the bounds on neutrino masses. The analysis remains nearly the same as in section 3.3.2, but now with the twin neutrinos at a lower temperature, as determined above. As mentioned above, for large enough f/v and SM reheating temperature sufficiently close to the lower bound, the reheating temperature of the twin sector may be below the twin neutrino decoupling temperature and the resulting energy density would be more difficult to compute. For simplicity, we choose $\lambda_x x$ large enough such that the twin reheating temperature is always above the twin neutrino decoupling temperature.

As before, we compute m_ν^{eff} as

$$m_\nu^{\text{eff}} = \frac{n_\nu^t}{n_\nu^{\text{SM}}} \sum_\alpha m_{\nu_\alpha}^t. \tag{4.21}$$

In relating the scale factors at neutrino decoupling in each sector, we now have to use the above temperature ratio to find, analogously to section 3.3.2, that

$$m_\nu^{\text{eff}} = \left(\frac{\Gamma_t}{\Gamma_{\text{SM}}} \right)^{3/4} \left(\frac{g_\star^t(T_{\text{reheat}}^t)}{g_\star^{\text{SM}}(T_{\text{reheat}}^{\text{SM}})} \right)^{1/4} \left(\frac{f}{v} \right)^n \sum_\alpha m_{\nu_\alpha}^{\text{SM}}, \tag{4.22}$$

where, again, $n = 1$ for Dirac masses and $n = 2$ for Majorana masses. Interestingly, if the branching ratios scale as $\Gamma_t/\Gamma_{\text{SM}} = (v/f)^2$, then we have $m_\nu^{\text{eff}} \propto (f/v)^{-3/2+n}$, so the contribution grows with f/v for Majorana masses, but is suppressed for Dirac masses.

As before, we consider the minimal mass spectrum of $m_\nu = [0.0, 0.009, 0.06 \text{ eV}]$ and a degenerate spectrum of $m_\nu = [0.1 \text{ eV}, 0.1 \text{ eV}, 0.1 \text{ eV}] / 3$. In figure 9 we plot the predictions of the X reheating for ΔN_{eff} and m_ν^{eff} for both spectra and both Dirac and Majorana masses using the approximations of section 3.3, for f/v from 3 to 10 and assuming the $\frac{\Gamma_t}{\Gamma_{\text{SM}}} \sim (v/f)^2$ scaling; there are regions in the space of m_X where the suppression of m_ν^{eff} would be much higher.

Dashed lines indicate the rough locations of present experimental limits from Planck 2015, and projected bounds from CMB-S4. As mentioned in section 3.3.2, we are unaware of any study of bounds on both m_ν^{eff} and ΔN_{eff} treated jointly. In the absence of this, we show present and projected constraints on N_{eff} and $\sum m_\nu$ from [50] and [27], ignoring correlations, as described in section 3.3.3.

4.3 Thermal production

In our discussion up to this point, we have been agnostic about the origin of the cosmic abundance of X and have operated under the assumption that it absolutely dominates the cosmology before it decays. Here, we consider the possibility that X was thermally produced through freeze-out and subsequently dominates the universe as a relic before

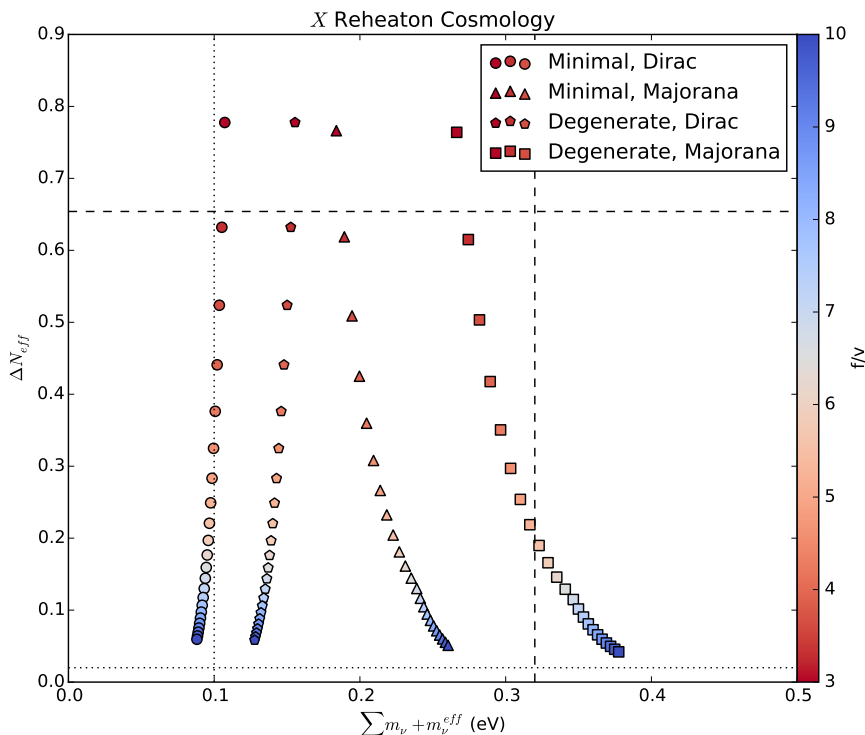


Figure 9. Predicted values of ΔN_{eff} and $\sum m_\nu + m_\nu^{\text{eff}}$ for minimal and degenerate neutrino mass spectra with both Dirac and Majorana masses for f/v from 3 to 10. The Planck 2015 [2] bounds on $\sum m_\nu$ and N_{eff} , as discussed in section 3.3.3, are represented by the dashed lines, and the projected CMB-S4 constraints are given by the dotted lines. It has been assumed that $\frac{\Gamma_t}{\Gamma_{\text{SM}}} \sim (v/f)^2$. Note however, that, from figure 8, this scaling of the partial widths holds only for the mass range $50 \text{ GeV} \lesssim m_X \lesssim 120 \text{ GeV}$, outside of which the twin partial width is more suppressed and the model is only testable through ΔN_{eff} over a smaller range in f/v .

decaying. This thermal history is viable, but places strong constraints on the mass and couplings of the X .

The energy density of relativistic species redshifts as $\rho_r \propto a^{-4} \propto T^4$, while the energy density of non-relativistic, chemically decoupled matter scales as $\rho_m \propto a^{-3}$. The energy density contained in the X can therefore only grow relative to the energy density in the thermal bath once it becomes non-relativistic. We found in section 4.2.1 that by recombination, $\rho_t/\rho_{\text{SM}} \lesssim 0.09$ is needed to evade current bounds on ΔN_{eff} . Thus we need to have the energy density in the X dominate over the SM and twin plasmas by more than this factor when it decays. If X becomes non-relativistic instantaneously at the moment that its temperature reaches some fraction $c \sim \mathcal{O}(0.1)$ of its mass, then, as $T \propto 1/a$ and ρ_X is $\sim 1/g_\star$ of the total energy density, the mass is required to satisfy $m_X \gtrsim 10/c \times g_\star(T = m_X) T_{X\text{reheat}}^{\text{SM}}$. Since the SM reheating temperature is strongly constrained to be above BBN, this effectively puts a lower limit on the mass of the X . Importantly, X must freeze-out when relativistic or its energy density will be further Boltzmann suppressed. The lower limit on the mass of the X becomes an upper limit on the X 's couplings — if it couples too strongly to the thermal bath, then it won't freeze out early enough to be hot.

In fact the situation is somewhat less favorable than the above analysis suggests, because it is relevant operators that must keep X in thermal equilibrium. For an X with the interactions introduced in section 4.1, the annihilations have rates that scale with temperature as $\Gamma \sim n_X \langle \sigma v \rangle \sim T$ for $T \gtrsim m_X, m_h$ (where n_X is the number density of X and $\langle \sigma v \rangle$ is its thermally averaged annihilation cross section). However, in a radiation-dominated universe, $H \sim T^2$. Thus, at high enough temperatures, X is not in thermal equilibrium with the plasma and it is only once the universe cools enough that it may thermalise. Then, as the temperature drops, $XX \rightarrow q\bar{q}$ annihilations become suppressed by the Higgs mass and subsequent Boltzmann suppression causes X to freeze-out. Note that the rates of these annihilation processes are controlled by the coupling λ_x , independently of x , which is unconstrained by itself (other processes mediated by $\lambda_x x$ are found to be subdominant in the ensuing analysis, for the range of λ_x over which thermal production is successful). If the coupling is too weak to begin with, then the X never thermalises and thermal production cannot happen. Thermal production therefore requires a careful balancing of parameters — small coupling λ_x is preferred for X to freeze-out hot and as early as possible, but the coupling is bounded from below by the requirement that X reach thermal equilibrium. This combination of constraints severely restricts the size of the parameter space over which thermal production is viable to cases in which the coupling is selected so that X enters and departs from thermal equilibrium at close to the same temperature.

To obtain numerical predictions for this scenario, the calculation of section 4.2.2 was modified to account for the time after the freeze-out of X before it becomes non-relativistic. During this period we use (3.15) and (3.16) for the energy density of the X , approximating decays as being negligible, before switching over to (4.18) when the temperature drops below the mass of the X . The approximation that the X does not decay appreciably while it is relativistic must be good if there is to be sufficient time for it to grow to dominate between becoming non-relativistic and decaying. The decay width of X was fixed to 5×10^{-21} GeV, corresponding to a reheating temperature close to the ~ 10 MeV lower limit, in order to maximise the amount of time over which the energy density of X may grow relative to the SM plasma, thereby providing the greatest possible reheating.

The predictions for ΔN_{eff} from a thermally produced X are shown in figure 10 for the small regions of parameter space where this is viable, with $f/v = 4$. We find that the dominant annihilation channels over this region are $XX \rightarrow t\bar{t}$ and $XX \rightarrow b\bar{b}$, mediated by the light Higgs, as well as their twin analogues, mediated by the heavy Higgs. As expected, the primordial energy density in the twin sector is too large compared to that generated by the X for the asymmetric reheating to be effective when m_X is too light ($\lesssim 100$ GeV in this case). Similarly, when the coupling is too strong, the X is held in equilibrium for longer and freezes-out underabundant compared to the twin energy density. However, when the coupling is too weak (the gray region), X never thermalises to begin with (close to the boundary with this region, X freezes-out almost immediately after thermalising). The peak in the contours occurs because of the “ H -funnel” in which the twin Higgs resonantly enhances annihilations into twin quarks. All of this region will be testable by CMB-S4.

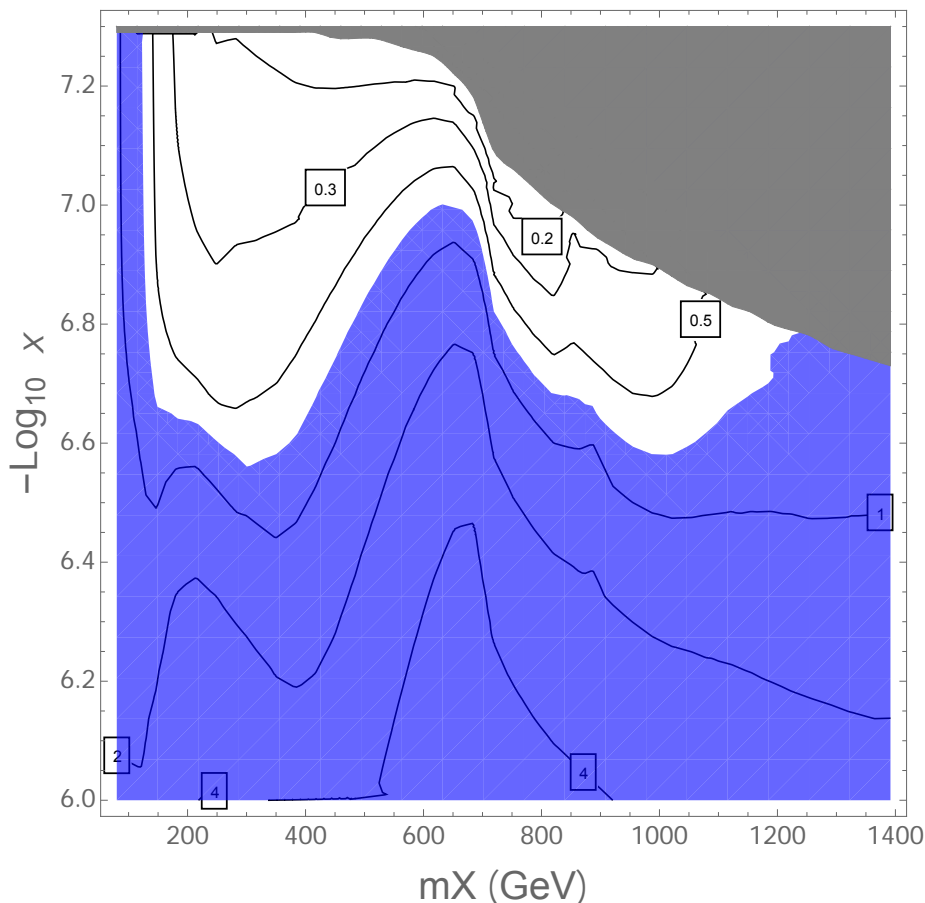


Figure 10. Parameter space where thermal production of X gives a large enough relic abundance to dilute the twin sector, for $f/v = 4$. In the gray region, the coupling is too weak for X to ever reach thermal equilibrium. The blue region is in tension with recent Planck measurements of ΔN_{eff} , whereas all of the white region will be tested by CMB-S4. Predictions presented here for ΔN_{eff} close to the gray boundary are more uncertain because of the high sensitivity of the freeze-out temperatures to the coupling.

5 Twinflation

As an alternative to the model presented above of late, out-of-equilibrium decays of a \mathbb{Z}_2 -symmetric scalar, one may imagine that the field driving primordial inflation reheats only the Standard Model to below the decoupling temperature of the two sectors. Production of the twin particles then ceases at some time after the temperature drops below the decoupling temperature during reheating.

To make this consistent with a softly-broken \mathbb{Z}_2 symmetry, we extend the inflationary sector and introduce a ‘twinflaton’ that couples solely to the twin sector. The combined inflationary and twinflationary sectors respect the \mathbb{Z}_2 symmetry. However, if the two sectors are entirely symmetric then one generally expects both inflationary dynamics to happen coincidentally, which would result in identical reheating. We therefore rely on soft \mathbb{Z}_2 -breaking to give an asymmetry between the two sectors that causes the twinflationary

sector to dominate the universe first. With the right arrangement we can end up with two distinct periods of inflation — a first caused by the “twinflaton” and a second that then reheats the Standard Model to below the decoupling temperature, having diluted the sources of twin-sector reheating from the first period.

One simple mechanism for \mathbb{Z}_2 -breaking which is well-suited for introducing asymmetry to inflationary sectors is to introduce an additional \mathbb{Z}_2 -odd scalar field η (as was done in [51]). This admits linear and quadratic interactions to antisymmetric and symmetric combinations of the inflationary sector fields, respectively. When η acquires a vev, this introduces an asymmetry in the fields to which it was coupled, dependent on the combination of its vev and its couplings. If η is coupled to both the inflationary sectors and the Higgs sectors, it could be the sole source of \mathbb{Z}_2 -breaking in a twinflationary theory. One may generally imagine that, in some UV completion, the mechanism that softly breaks the symmetry in the Higgs potential could also be the origin of the soft breaking of the inflationary sector.

Cosmologically, this possibility may have similar observational signatures as the model discussed in section 4, where the amount of twin-sector dark radiation is determined by the partial widths of the inflaton of the second inflationary epoch. If this dominantly couples to the SM, then ΔN_{eff} will be suppressed which, while successfully resolving the cosmological problems of the Mirror Twin Higgs, may also be observationally inaccessible. However, additional, distinctly inflationary signatures may make this potentially testable by other cosmological observations.

The mechanism of twinflation completes a catalog of models of asymmetric reheating by late decays, which may be indexed by representations of the twin parity: the case of a \mathbb{Z}_2 -even particle, in which a kinematic asymmetry in the partial widths provides the reheating asymmetry, the case of a \mathbb{Z}_2 -odd particle, which can also provide the spontaneous \mathbb{Z}_2 -breaking required in the Higgs potential, and the case where two distinct, long-lived particles couple to each sector, which may also be related to inflation.

5.1 Toy model

As a toy model we here consider ‘twinning’ the simple φ^2 chaotic inflation scenario. The inflationary dynamics in this case are easy to understand and we have the additional benefit that this inflationary model has been considered in the literature before as ‘Double Inflation’ (see [52, 53] and [54]). We furthermore specialize to ‘double inflation with a break’, where there are two distinct periods of inflation which produces a step in the power spectrum, and we consider the constraints that this places on our model. In this case, it is assumed that each inflaton field couples and therefore decays dominantly into the sector to which it belongs. We will comment briefly on the case without a break and the additional signals one could look for in that case.

The potential of the inflationary sector for inflaton φ_A and twinflaton φ_B is

$$V = \frac{1}{2}m_A^2\varphi_A^2 + \frac{1}{2}m_B^2\varphi_B^2, \quad (5.1)$$

where $m_A \neq m_B$ may arise from soft \mathbb{Z}_2 -breaking, perhaps related to the soft \mathbb{Z}_2 -breaking in the Higgs potential. In order for the ‘twinflation’ to occur first, we require that the energy

of the B field initially dominates the energy density of the universe. We take the initial positions of the fields to be the same and $m_B^2 \gg m_A^2$.⁵ Call $\varphi_A(0) = \varphi_B(0) = n\sqrt{2}M_{\text{pl}} = n\varphi_c$, where φ_c is the critical value at which inflation stops and $m_B = rm_A = rm$ with $n, r > 1$. The inflationary dynamics are then those of slowly-rolling scalar fields. At some point in the early universe we imagine that the slow-roll approximation holds for both fields and the inflationary sector dominates the universe. The dominating field then slow-rolls down its potential for $\frac{n^2-1}{2}$ e -folds, while the lighter field's velocity is suppressed by approximately $\frac{\varphi_A}{r^2\varphi_B}$. Solving the system numerically reveals that the motion of φ_A during this period can be neglected entirely.

After φ_B reaches the critical value $\sqrt{2}M_{\text{pl}}$, it stops slow-rolling and begins oscillating around the minimum of its potential. For there to be two distinct periods of inflation, there must be a period where these oscillations dominate the universe, which requires that the energy densities of each inflaton ρ_A and ρ_B satisfy $\rho_B(\varphi_c) = r^2m^2M_{\text{pl}}^2 > \rho_A(\varphi(0)) = n^2m^2M_{\text{pl}}^2$ and therefore $r > n$. For a φ_B^2 potential, the energy in these oscillations redshifts as $\rho_B \sim a^{-3}$. Eventually, the energy density in φ_B drops below that of φ_A and a new epoch of inflation, driven by φ_A , begins. This provides a further $\frac{n^2-1}{2}$ e -folds of inflation to give $n^2 - 1$ in total, while the B -sector energy density is diluted away.

Note that in order for our toy model to reheat below the decoupling temperature of the two sectors, reheating must occur well after the end of inflation. If, during the coherent oscillation of an inflaton, it becomes the case that the inflaton decay width $\Gamma \sim H$, then reheating will occur and result in temperature $T_{\text{reheat}} \sim 0.1\sqrt{\Gamma M_{\text{pl}}}$. However, if $\Gamma \gg H$ when inflation ends, then all of the energy in the inflaton is immediately transferred and we instead have reheating temperature $T_{\text{reheat}} \sim 0.1\sqrt{m_\alpha M_{\text{pl}}}$ for an inflaton of mass m_α . But in order for $T_{\text{reheat}} \lesssim 1 \text{ GeV}$, it is required that $m_\alpha \lesssim 10^{-7} \text{ eV}$, so this possibility that the inflaton is short lived is not viable. The procedure of twinning inflationary potentials may be generalised to other, more realistic models, provided that this constraint upon the reheating temperature can be satisfied.

5.2 Observability

One could always make a twinflationary scenario consistent with observational constraints by letting the second inflationary period of inflation last long enough. In our toy model, this would correspond to setting n high enough that the momentum modes which left the horizon during the first inflation have not yet re-entered the horizon — such a scenario would look exactly like single-field chaotic inflation.

Alternatively, we may also allow for n small enough that all the momentum modes that left the horizon during the second inflation are currently sub-horizon. In this case, fluctuations at large enough wavenumbers (equivalently, small enough length scales) are ‘processed’ (cross the horizon) at a different inflationary energy scale than those that were processed earlier, giving a step in the power spectrum. While Planck has measured the primordial power spectrum for modes with $10^{-4} \text{ Mpc}^{-1} \lesssim k \lesssim 0.3 \text{ Mpc}^{-1}$ (where the lower

⁵Note that merely giving the twin field a much larger initial condition does not instigate twinflation. The dynamics of the subdominant field in this case are such that it will track the dominant field and both will reach the critical value at the same time. This is easily confirmed numerically.

bound is set by the fact that smaller modes have not yet re-entered the horizon), proposed CMB-S4 experiments will increase this range [27] somewhat, as will be discussed further below. We wish to show that the power spectrum of our toy model is not ruled out and, furthermore, may be observed in the coming decades.

The height of the step in the primordial power spectrum is determined by the energy scale of each period of inflation, so modes crossing the horizon in the second inflationary period should be suppressed by a factor of $r^2 > n^2 \gtrsim 25$ compared to those exiting in the first period. This degree of suppression is ruled out by Planck for the range of modes over which it has reconstructed the power spectrum [50]. A computation of the primordial power spectrum for double inflation was given in [53]. It was found that significant damping does not occur for modes which cross outside the horizon during the first inflationary period, re-enter during the inter-inflationary period and again cross the horizon during the second inflationary period. It is only those scales which first cross the horizon during the second inflationary period that are significantly damped (although other features in the shape, such as oscillations, may be present for modes that are subhorizon during the intermediate period).

The relation of this characteristic scale to present-day observables is easily done using the framework given in [55]. Let the subscripts a, b, c, d, e respectively correspond to the beginning of the first inflationary period, the end of that period, the beginning of the second inflationary period, the end of that period, and the beginning of radiation domination. During the coherent oscillation periods, the inflaton acts as matter and the energy density falls as $\rho \propto a^{-3}$. Let k_i be the momentum whose mode is horizon-size at the i epoch; $k_i = a_i H_i$. The scales k_i can be related using the number of e -folds in each period, which are themselves determined from the first Friedmann equation. Denoting $N_{ij} = \ln \frac{a_j}{a_i}$, we have $k_a = e^{-N_{ab}} n k_b$, $k_b = e^{\frac{1}{2} N_{bc}} k_c$ and similarly for the other characteristic modes, where, in particular, slow-roll inflation predicts that $N_{ab} = N_{cd} = \frac{n^2 - 1}{2}$. The evolution of the characteristic momentum scales is shown schematically in figure 11. Finally, k_e can be determined using the conservation of comoving entropy:

$$k_e = \frac{\pi g_\star^{1/3}(T_0) g_\star^{1/6}(T_{\text{reheat}}) T_0 T_{\text{reheat}}}{3\sqrt{10} M_{\text{pl}}}, \tag{5.2}$$

where T_0 and a_0 are the temperature and scale factor today and T_{reheat} is the reheating temperature (which is sufficiently low that only SM particles are produced). We work explicitly with the convention $a_0 = 1$. The characteristic modes associated with the break can then be determined.

As mentioned above, [53] shows that damping occurs for modes that exit the horizon only during the second inflationary period, so we should take the characteristic damping scale to be the smallest such scale, which here corresponds roughly to k_b . This can be determined as

$$\begin{aligned} k_b &= n e^{\frac{1}{2} N_{bc} - N_{cd} + \frac{1}{2} N_{de}} k_e \\ &= n \left(\frac{r}{n}\right)^{1/3} \exp\left(-\frac{n^2 - 1}{2}\right) \left[\frac{\frac{1}{2} m^2 M_{\text{pl}}^2}{\frac{\pi^2}{30} g_\star(T_{\text{reheat}}) T_{\text{reheat}}^4} \right]^{1/6} \frac{\pi g_\star^{1/3}(T_0) g_\star^{1/6}(T_{\text{reheat}}) T_0 T_{\text{reheat}}}{3\sqrt{10} M_{\text{pl}}} \end{aligned} \tag{5.3}$$

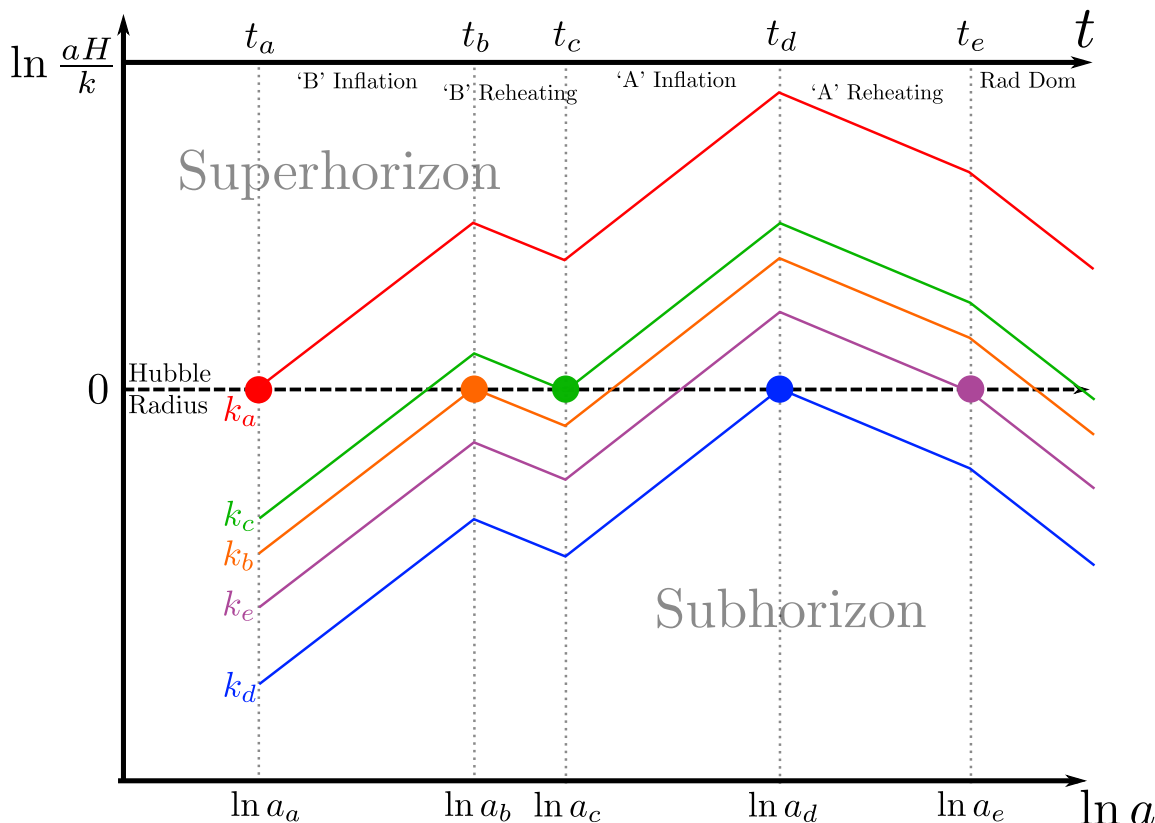


Figure 11. Schematic evolution of the characteristic scales in Twinflation, as seen by comparing wavenumbers to the Hubble radius over time. Note that the time axis is not a linear scale.

where k_c only differs by the factor of $(r/n)^{1/3}$ (which is roughly close to unity). Once again, between k_b and k_c are oscillatory features, so k_b should merely be taken as the rough characteristic scale of the damping.

Now the characteristic damping scale is determined by m , n , r , and T_{reheat} . Our observational bound on k_b is that Planck has not seen this suppression on momentum scales at which it has been able to reconstruct the primordial power spectrum from the angular temperature anisotropy power spectrum, which is roughly $k \lesssim 0.3 \text{ Mpc}^{-1}$. We have constraints on the reheating temperature from rethermalization of the twin sector or interrupted big bang nucleosynthesis $10 \text{ MeV} \lesssim T_{\text{reheat}} \lesssim 1 \text{ GeV}$, on having a period of intermediate matter domination between the two inflations $r > n$ and on the total number of e -folds $n^2 - 1 \gtrsim 25$ to solve cosmological problems. Note that we require fewer e -folds of inflation than is typically assumed in the standard cosmology. Since the low reheating temperature gives fewer e -folds from reheating up to today, less inflation is needed to explain the large causal horizon and flatness.

The normalization of the spectrum provides a further constraint, the most recent measurement of which come from Planck [50]. The scalar power spectrum at $k_* = 0.05 \text{ Mpc}^{-1}$ is measured to be $\mathcal{P}_{\mathcal{R}}(k_*) = e^{3.094 \pm 0.034} \times 10^{-10}$. Then for $k_* < k_c$ (i.e. k_* having left the horizon during the first period of inflation and not re-entered before the second, so no

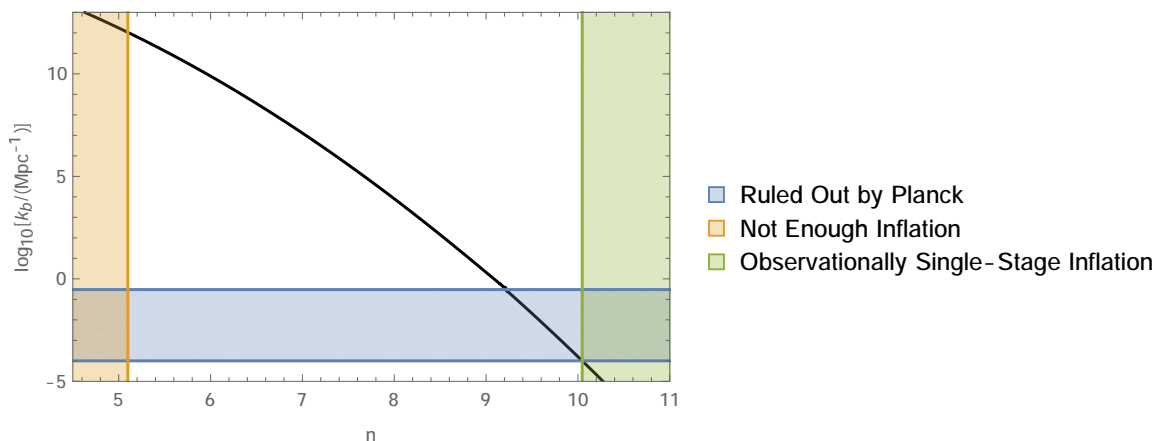


Figure 12. The prediction for the characteristic suppression scale as a function of the initial values of the fields. The mapped regions should be interpreted not as having hard boundaries, but rather fuzzy endpoints where they break down. Here we have used $T_{\text{reheat}} = 10 \text{ MeV}$ and $r = 2n$.

deviation from single-field inflation would be seen at this scale), the spectrum of [53] yields the constraint

$$2.03 \times 10^{-6} = \frac{r^2 m^2}{M_{\text{pl}}^2} \ln\left(\frac{k_b}{k_\star}\right) \left(\ln\frac{k_b}{k_\star} + \frac{n^2}{2}\right). \quad (5.4)$$

The characteristic scale (5.3) depends much more strongly on n than it does on any of the other parameters. In figure 12, we give a rough idea of the scale as a function of n , having set $T_{\text{reheat}} = 10 \text{ MeV}$ and $r = 2n$, while m is chosen to satisfy the normalization condition. We also show the constraint on k_b set by Planck. Note again that the region described as “observationally single-stage inflation” *does* still provide a solution to the problem of reconciling cosmology with the mirror Twin Higgs.

CMB-S4 will improve the constraint on k_b through its improved measurement of polarization anisotropies [27]. With only precision measurements of temperature anisotropies, the un-lensed power spectrum cannot be so easily reconstructed from the lensed spectrum. The effects of gravitational lensing of CMB place an upper limit on the size of primordial temperature anisotropies that can be measured [56], which Planck has saturated. However, the polarization anisotropy power spectrum allows the removal of lensing noise from the temperature spectrum so that higher primordial modes can be detected. The polarization power spectrum itself also gives us another window into the high- ℓ modes of the primordial power spectrum, as the signal does not become dominated by polarized foreground sources until higher scales near $\ell \sim 5000$. CMB-S4 is projected to make cosmic variance limited measurements of both the temperature and polarization anisotropy power spectra up to the modes where they become foreground-contaminated and so provide additional information on the shape of the primordial power spectrum [27]. The map from measurements of angular modes ℓ to constraints on spatial modes k depends on the evolution of the power spectrum between inflation and the CMB, so forecasting constraints requires careful study. However, these improvements will not test most of the parameter space presented in figure 12, where the step is predicted on extremely small distance scales.

We have discussed a twinflationary model of double inflation with a break for simplicity, but there is a parametric regime where double inflation without a break gives the required amount of asymmetric reheating into the Standard Model. With two periods of inflation, the second period dilutes the energy density of the heavier field sufficiently that there is no observable signal of it produced in reheating. However, even with only one period, inflation can continue for long enough after the inflaton turns the corner in field space such that, at late times, the fraction of the inflaton in the B state relative to the A state is small enough that the expected energy densities that are transferred into each sector satisfy $\rho_B/\rho_A < 0.1$. This occurs as long as $r \gtrsim 1.2$, assuming that the mixing angle of the slow-rolling field with the φ_A and φ_B fields entirely determines the fraction of its energy that reheats each sector. There is thus a much larger range of r where this toy model of inflation passes N_{eff} bounds than our above analysis shows. The resulting imprint on the CMB could resemble that of the long-lived decay model of section 4, with ΔN_{eff} again being related to the ratio of branching fractions, although this is dependent upon the UV completion of the Twin Higgs.

When there is only one period of inflation, the step is smoothed out and less pronounced and it is necessary to locate the feature numerically. Furthermore, having multiple degrees of freedom available allows for non-trivial evolution of momentum modes after they become super-horizon, which does not occur in single-field inflation but may be calculated from the full solution to the field equations [54]. While a twinned potential leading to two periods of inflation generally predicts a step in the power spectrum, when there is no break the predictions, and thus constraints, this prediction become more model-dependent. Therefore we leave detailed predictions in that case for future study using realistic models and merely state that the range of $r = 1$ to n interpolates between the single field spectrum and that with a step, as one would expect.

There are also at least two other detectable effects one might expect in double inflation without a break and in general realistic twinflationary models. Interactions between inflaton fields may produce primordial non-Gaussianities, while the presence of additional oscillating degrees of freedom may produce isocurvature perturbations. These do not appear in our toy model because the heavy field is exponentially damped during the second inflation. CMB-S4 is projected to improve Planck's bounds on non-Gaussianities by a factor of ~ 2 and on isocurvature perturbations by perhaps an order of magnitude (though model-independent projections have not been made), so may be able to detect or place useful constraints on realistic twinflationary models [27].

We have introduced twinflation as a mirror Twin Higgs model which suppresses the cosmological effects of twin light degrees of freedom. It extends the mirror symmetry to the inflationary sector. The soft Z_2 symmetry-breaking of the Higgs sector may be used in the inflationary sector to cause distinct periods of inflation. There exists a parametric region where this is cosmologically indistinct from single-stage inflation, but also another in which it may be observable. As the direct product of inflation and the Mirror Twin Higgs, this is in some sense a minimal solution.

6 Conclusion

In this work we have considered scenarios in which cosmology provides meaningful insight on solutions to the electroweak hierarchy problem. In particular, we have demonstrated several simple mechanisms in which the cosmological history of a mirror Twin Higgs model is reconciled with current CMB constraints and provides signatures accessible in future CMB experiments. In the case of out-of-equilibrium decays, we have found that decays of \mathbb{Z}_2 -even scalars sufficiently dilute the energy density in the twin sector without the addition of any new sources of \mathbb{Z}_2 -breaking. In much of the parameter space, the residual contribution to ΔN_{eff} is directly proportional to the ratio of vacuum expectation values v^2/f^2 parameterizing the mixing between Standard Model and twin sectors (as well as the tuning of the electroweak scale), and may be within reach of CMB-S4 experiments. In the case of twinflation, we have found that a (broken) \mathbb{Z}_2 -symmetric inflationary sector may successfully dilute the energy density in the twin sector, as well as potentially leave signatures in the form of a step in the primordial power spectrum or in departures of primordial perturbations from adiabaticity and Gaussianity. In both cases, these models raise the tantalizing possibility that signatures of electroweak naturalness may first emerge in the CMB, rather than the LHC.

There are a variety of possible directions for future work. Here we have focused on the cosmological consequences of late-decaying scalars and twinned inflationary sectors without specifying their origin in a microscopic model. It would be interesting to construct complete models (where, e.g., supersymmetry or compositeness protect the scale f from UV contributions) in which the existence and couplings of late-decaying scalars arise as intrinsic ingredients of the UV completion. Likewise, we have considered only a toy model of twin chaotic inflation; it would be interesting to see if twinflation may be realized in complete inflationary models that match the observed spectral index and constraints on the tensor-to-scalar ratio.

While we have taken care to ensure that our scenarios respect the well-measured cosmological history beneath $T \sim 1 \text{ MeV}$, we have not addressed the origin of the observed baryon asymmetry. In the case of out-of equilibrium decays, there are a number of possibilities. It is plausible that a somewhat larger baryon asymmetry is generated through various conventional mechanisms and diluted by late decays. Alternatively, the decay mechanism itself may possibly be expanded to generate a baryon asymmetry or some other late decay may generate the baryon asymmetry below $\sim 1 \text{ GeV}$. In the case of twinflation, inflationary dilution of pre-existing baryon asymmetry requires that baryogenesis occur in association with reheating or via another mechanism at temperatures below $\sim 1 \text{ GeV}$. It would be worthwhile to study models for the baryon asymmetry consistent with these scenarios. Steps in this direction have been taken in [17], which attempted to relate this to asymmetric dark matter in the twin sector.

Likewise, any investigation of dark matter, be it related directly to the twin mechanism or otherwise, must also address implications of the dilution. Previous work attempting to construct dark matter candidates in the twin sector [11–18]) has relied upon explicit \mathbb{Z}_2 -breaking that is not present in the mirror model. Dark matter may alternatively be

unrelated to the Twin Higgs mechanism, such as a WIMP in some minimal extension of the electroweak sector that freezes-out as an overabundant thermal relic and is then diluted to the observed density during reheating. Alternatively, it may be that the dark matter abundance is produced directly during reheating. It would be interesting to study extensions of our scenarios that incorporate dark matter candidates directly related to the mechanism of dilution.

Finally, we have only approximately parameterized Planck constraints and the reach of CMB-S4 on twin neutrinos and twin photons. Ultimately, more precise constraints and forecasts may be obtained via numerical CMB codes. This strongly motivates the future study of CMB constraints on scenarios with three sterile neutrinos and additional dark radiation whose temperatures differ from the Standard Model thermal bath.

Acknowledgments

We thank Zackaria Chacko, Yanou Cui, Paddy Fox, Daniel Green, and Roni Harnik for useful discussions, and Nima Arkani-Hamed for enlightening remarks on N naturalness. The work of NC, SK, and TT is supported in part by the US Department of Energy under the grant DE-SC0014129. NC acknowledges the hospitality of the Kavli Institute for Theoretical Physics and the corresponding support of the National Science Foundation under Grant No. NSF PHY11-25915.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] Z. Chacko, H.-S. Goh and R. Harnik, *The Twin Higgs: Natural electroweak breaking from mirror symmetry*, *Phys. Rev. Lett.* **96** (2006) 231802 [[hep-ph/0506256](https://arxiv.org/abs/hep-ph/0506256)] [[INSPIRE](#)].
- [2] PLANCK collaboration, P.A.R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [[arXiv:1502.01589](https://arxiv.org/abs/1502.01589)] [[INSPIRE](#)].
- [3] R.H. Cyburt, B.D. Fields, K.A. Olive and T.-H. Yeh, *Big Bang Nucleosynthesis: 2015*, *Rev. Mod. Phys.* **88** (2016) 015004 [[arXiv:1505.01076](https://arxiv.org/abs/1505.01076)] [[INSPIRE](#)].
- [4] N. Craig, S. Knapen and P. Longhi, *Neutral Naturalness from Orbifold Higgs Models*, *Phys. Rev. Lett.* **114** (2015) 061803 [[arXiv:1410.6808](https://arxiv.org/abs/1410.6808)] [[INSPIRE](#)].
- [5] N. Craig, S. Knapen and P. Longhi, *The Orbifold Higgs*, *JHEP* **03** (2015) 106 [[arXiv:1411.7393](https://arxiv.org/abs/1411.7393)] [[INSPIRE](#)].
- [6] M. Geller and O. Telem, *Holographic Twin Higgs Model*, *Phys. Rev. Lett.* **114** (2015) 191801 [[arXiv:1411.2974](https://arxiv.org/abs/1411.2974)] [[INSPIRE](#)].
- [7] R. Barbieri, D. Greco, R. Rattazzi and A. Wulzer, *The Composite Twin Higgs scenario*, *JHEP* **08** (2015) 161 [[arXiv:1501.07803](https://arxiv.org/abs/1501.07803)] [[INSPIRE](#)].
- [8] M. Low, A. Tesi and L.-T. Wang, *Twin Higgs mechanism and a composite Higgs boson*, *Phys. Rev.* **D 91** (2015) 095012 [[arXiv:1501.07890](https://arxiv.org/abs/1501.07890)] [[INSPIRE](#)].

- [9] N. Craig, A. Katz, M. Strassler and R. Sundrum, *Naturalness in the Dark at the LHC*, *JHEP* **07** (2015) 105 [[arXiv:1501.05310](#)] [[INSPIRE](#)].
- [10] N. Craig, S. Knapen, P. Longhi and M. Strassler, *The Vector-like Twin Higgs*, *JHEP* **07** (2016) 002 [[arXiv:1601.07181](#)] [[INSPIRE](#)].
- [11] R. Barbieri, L.J. Hall and K. Harigaya, *Minimal Mirror Twin Higgs*, *JHEP* **11** (2016) 172 [[arXiv:1609.05589](#)] [[INSPIRE](#)].
- [12] I. Garcia Garcia, R. Lasenby and J. March-Russell, *Twin Higgs WIMP Dark Matter*, *Phys. Rev. D* **92** (2015) 055034 [[arXiv:1505.07109](#)] [[INSPIRE](#)].
- [13] N. Craig and A. Katz, *The Fraternal WIMP Miracle*, *JCAP* **10** (2015) 054 [[arXiv:1505.07113](#)] [[INSPIRE](#)].
- [14] I. Garcia Garcia, R. Lasenby and J. March-Russell, *Twin Higgs Asymmetric Dark Matter*, *Phys. Rev. Lett.* **115** (2015) 121801 [[arXiv:1505.07410](#)] [[INSPIRE](#)].
- [15] M. Farina, *Asymmetric Twin Dark Matter*, *JCAP* **11** (2015) 017 [[arXiv:1506.03520](#)] [[INSPIRE](#)].
- [16] M. Freytsis, S. Knapen, D.J. Robinson and Y. Tsai, *Gamma-rays from Dark Showers with Twin Higgs Models*, *JHEP* **05** (2016) 018 [[arXiv:1601.07556](#)] [[INSPIRE](#)].
- [17] M. Farina, A. Monteux and C.S. Shin, *Twin mechanism for baryon and dark matter asymmetries*, *Phys. Rev. D* **94** (2016) 035017 [[arXiv:1604.08211](#)] [[INSPIRE](#)].
- [18] V. Prilepina and Y. Tsai, *Reconciling Large And Small-Scale Structure In Twin Higgs Models*, [arXiv:1611.05879](#) [[INSPIRE](#)].
- [19] S. Chang, L.J. Hall and N. Weiner, *A supersymmetric twin Higgs*, *Phys. Rev. D* **75** (2007) 035009 [[hep-ph/0604076](#)] [[INSPIRE](#)].
- [20] N. Craig and K. Howe, *Doubling down on naturalness with a supersymmetric twin Higgs*, *JHEP* **03** (2014) 140 [[arXiv:1312.1341](#)] [[INSPIRE](#)].
- [21] M. Drees, F. Hajkarim and E.R. Schmitz, *The Effects of QCD Equation of State on the Relic Density of WIMP Dark Matter*, *JCAP* **06** (2015) 025 [[arXiv:1503.03513](#)] [[INSPIRE](#)].
- [22] D.C. Kelly, *Collision Operator for a Relativistic Lorentz Gas*, *Phys. Fluids* **12** (1969) 799.
- [23] J. Edsjo and P. Gondolo, *Neutralino relic density including coannihilations*, *Phys. Rev. D* **56** (1997) 1879 [[hep-ph/9704361](#)] [[INSPIRE](#)].
- [24] P. Adshead, Y. Cui and J. Shelton, *Chilly Dark Sectors and Asymmetric Reheating*, *JHEP* **06** (2016) 016 [[arXiv:1604.02458](#)] [[INSPIRE](#)].
- [25] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti and P.D. Serpico, *Relic neutrino decoupling including flavor oscillations*, *Nucl. Phys. B* **729** (2005) 221 [[hep-ph/0506164](#)] [[INSPIRE](#)].
- [26] Z. Hou, R. Keisler, L. Knox, M. Millea and C. Reichardt, *How Massless Neutrinos Affect the Cosmic Microwave Background Damping Tail*, *Phys. Rev. D* **87** (2013) 083008 [[arXiv:1104.2333](#)] [[INSPIRE](#)].
- [27] CMB-S4 collaboration, K.N. Abazajian et al., *CMB-S4 Science Book, First Edition*, [arXiv:1610.02743](#) [[INSPIRE](#)].
- [28] C. Brust, D.E. Kaplan and M.T. Walters, *New Light Species and the CMB*, *JHEP* **12** (2013) 058 [[arXiv:1303.5379](#)] [[INSPIRE](#)].

- [29] S. Dodelson, A. Melchiorri and A. Slosar, *Is cosmology compatible with sterile neutrinos?*, *Phys. Rev. Lett.* **97** (2006) 041301 [[astro-ph/0511500](#)] [[INSPIRE](#)].
- [30] D. Baumann, D. Green, J. Meyers and B. Wallisch, *Phases of New Physics in the CMB*, *JCAP* **01** (2016) 007 [[arXiv:1508.06342](#)] [[INSPIRE](#)].
- [31] M. Archidiacono, S. Hannestad, A. Mirizzi, G. Raffelt and Y.Y.Y. Wong, *Axion hot dark matter bounds after Planck*, *JCAP* **10** (2013) 020 [[arXiv:1307.0615](#)] [[INSPIRE](#)].
- [32] T.D. Jacques, L.M. Krauss and C. Lunardini, *Additional Light Sterile Neutrinos and Cosmology*, *Phys. Rev. D* **87** (2013) 083515 [[arXiv:1301.3119](#)] [[INSPIRE](#)].
- [33] J. Lesgourgues and S. Pastor, *Neutrino mass from Cosmology*, *Adv. High Energy Phys.* **2012** (2012) 608515 [[arXiv:1212.6154](#)] [[INSPIRE](#)].
- [34] S. Dodelson, E. Gates and A. Stebbins, *Cold + hot dark matter and the cosmic microwave background*, *Astrophys. J.* **467** (1996) 10 [[astro-ph/9509147](#)] [[INSPIRE](#)].
- [35] C. Howlett, A. Lewis, A. Hall and A. Challinor, *CMB power spectrum parameter degeneracies in the era of precision cosmology*, *JCAP* **04** (2012) 027 [[arXiv:1201.3654](#)] [[INSPIRE](#)].
- [36] E. Di Valentino, S. Gariazzo, M. Gerbino, E. Giusarma and O. Mena, *Dark Radiation and Inflationary Freedom after Planck 2015*, *Phys. Rev. D* **93** (2016) 083523 [[arXiv:1601.07557](#)] [[INSPIRE](#)].
- [37] M. Costanzi, B. Sartoris, M. Viel and S. Borgani, *Neutrino constraints: what large-scale structure and CMB data are telling us?*, *JCAP* **10** (2014) 081 [[arXiv:1407.8338](#)] [[INSPIRE](#)].
- [38] S. Gariazzo, C. Giunti, M. Laveder, Y.F. Li and E.M. Zavatin, *Light sterile neutrinos*, *J. Phys. G* **43** (2016) 033001 [[arXiv:1507.08204](#)] [[INSPIRE](#)].
- [39] N. Arkani-Hamed, T. Cohen, R.T. D’Agnolo, A. Hook, H.D. Kim and D. Pinner, *Solving the Hierarchy Problem at Reheating with a Large Number of Degrees of Freedom*, *Phys. Rev. Lett.* **117** (2016) 251801 [[arXiv:1607.06821](#)] [[INSPIRE](#)].
- [40] M. Reece and T. Roxlo, *Nonthermal production of dark radiation and dark matter*, *JHEP* **09** (2016) 096 [[arXiv:1511.06768](#)] [[INSPIRE](#)].
- [41] L. Randall, J. Scholtz and J. Unwin, *Flooded Dark Matter and S Level Rise*, *JHEP* **03** (2016) 011 [[arXiv:1509.08477](#)] [[INSPIRE](#)].
- [42] E.W. Kolb and M.S. Turner, *The Early Universe*, *Front. Phys.* **69** (1990) 1 [[INSPIRE](#)].
- [43] P.F. de Salas, M. Lattanzi, G. Mangano, G. Miele, S. Pastor and O. Pisanti, *Bounds on very low reheating scenarios after Planck*, *Phys. Rev. D* **92** (2015) 123534 [[arXiv:1511.00672](#)] [[INSPIRE](#)].
- [44] A. Djouadi, *The Anatomy of electro-weak symmetry breaking. I: The Higgs boson in the standard model*, *Phys. Rept.* **457** (2008) 1 [[hep-ph/0503172](#)] [[INSPIRE](#)].
- [45] M. Drees and K.-i. Hikasa, *Heavy Quark Thresholds in Higgs Physics*, *Phys. Rev. D* **41** (1990) 1547 [[INSPIRE](#)].
- [46] M. Baumgart and A. Katz, *Implications of a New Light Scalar Near the Bottomonium Regime*, *JHEP* **08** (2012) 133 [[arXiv:1204.6032](#)] [[INSPIRE](#)].
- [47] J.D. Clarke, R. Foot and R.R. Volkas, *Phenomenology of a very light scalar ($100\text{ MeV} < m_h < 10\text{ GeV}$) mixing with the SM Higgs*, *JHEP* **02** (2014) 123 [[arXiv:1310.8042](#)] [[INSPIRE](#)].

- [48] U. Haisch and J.F. Kamenik, *Searching for new spin-0 resonances at LHCb*, *Phys. Rev. D* **93** (2016) 055047 [[arXiv:1601.05110](#)] [[INSPIRE](#)].
- [49] K.N. Abazajian et al., *Neutrino Physics from the Cosmic Microwave Background and Large Scale Structure*, *Astropart. Phys.* **63** (2015) 66 [[arXiv:1309.5383](#)] [[INSPIRE](#)].
- [50] PLANCK collaboration, P.A.R. Ade et al., *Planck 2015 results. XX. Constraints on inflation*, *Astron. Astrophys.* **594** (2016) A20 [[arXiv:1502.02114](#)] [[INSPIRE](#)].
- [51] Z.G. Berezhiani, A.D. Dolgov and R.N. Mohapatra, *Asymmetric inflationary reheating and the nature of mirror universe*, *Phys. Lett. B* **375** (1996) 26 [[hep-ph/9511221](#)] [[INSPIRE](#)].
- [52] J. Silk and M.S. Turner, *Double Inflation*, *Phys. Rev. D* **35** (1987) 419 [[INSPIRE](#)].
- [53] D. Polarski and A.A. Starobinsky, *Spectra of perturbations produced by double inflation with an intermediate matter dominated stage*, *Nucl. Phys. B* **385** (1992) 623 [[INSPIRE](#)].
- [54] D. Wands, N. Bartolo, S. Matarrese and A. Riotto, *An observational test of two-field inflation*, *Phys. Rev. D* **66** (2002) 043520 [[astro-ph/0205253](#)] [[INSPIRE](#)].
- [55] S.E. Hong, H.-J. Lee, Y.J. Lee, E.D. Stewart and H. Zoe, *Effects of thermal inflation on small scale density perturbations*, *JCAP* **06** (2015) 002 [[arXiv:1503.08938](#)] [[INSPIRE](#)].
- [56] W. Hu and T. Okamoto, *Principal power of the CMB*, *Phys. Rev. D* **69** (2004) 043004 [[astro-ph/0308049](#)] [[INSPIRE](#)].