

Neutrino mass and secluded dark matter puzzles solved together

Mattia Di Mauro^{*}*Istituto Nazionale di Fisica Nucleare, Sezione di Torino, Via P. Giuria 1, 10125 Torino, Italy*

(Received 28 November 2025; accepted 17 April 2026; published 26 May 2026)

We present a secluded dark-matter (DM) framework based on an extra $U(1)_X$ gauge symmetry. The model contains a Dirac DM particle χ , three heavy neutrinos N_I with masses $M_{N,I}$, and a singlet scalar R that mixes with the Standard Model Higgs doublet Φ by an angle α . A symmetry forbids the $\Phi - R$ portal at tree level; the leading portal then arises at one loop from the same Yukawa structures that generate active neutrino masses $m_{\nu,I}$, implying $\tan(2\alpha) \propto \sum_I m_{\nu,I} M_{N,I}^2 / (v_h m_H^2)$, where v_h and m_H are the Standard Model Higgs vacuum-expectation-value and mass. For heavy-neutrino masses in the multi-TeV range, this yields a naturally tiny mixing, $\tan(2\alpha) \sim 5 \times 10^{-11} (M_N/10 \text{ TeV})^2$, which strongly suppresses DM signals in direct, indirect, and collider searches. For PeV-scale heavy neutrinos the loop-induced portal is enhanced and the DM-nucleon cross section can instead enter the reach of direct-detection experiments. The visible and dark sectors thermalize at temperatures of order a few times the mass of the lightest heavy neutrino, then subsequently decouple, and typically evolve with a slightly hotter dark bath. In the secluded regime, with $\tan(2\alpha) \ll 1$ and $m_\chi > m_{H_p}$, the relic density is set by p -wave annihilation $\chi\bar{\chi} \rightarrow H_p H_p$ (with H_p the Higgs-like particle of the dark sector), and the dark-sector Yukawa couplings required to reproduce the observed abundance are $\mathcal{O}(0.1 - 1)$, as in the standard WIMP case. For heavy-neutrino masses $\gtrsim 10 \text{ TeV}$, the mediator decays before nucleosynthesis without spoiling big bang nucleosynthesis observables, while the tiny portal suppresses present-day signals below current and near-future sensitivities. This links two long-standing puzzles—the absence of DM signals and the smallness of neutrino masses—within a predictive thermal framework.

DOI: [10.1103/np7k-2fck](https://doi.org/10.1103/np7k-2fck)

I. INTRODUCTION

Dark matter (DM) has so far revealed itself only through its gravitational effects on astrophysical and cosmological scales, while no signal of DM interactions with Standard Model (SM) particles has been observed in the laboratory [1–3]. A viable DM candidate must be cosmologically long-lived, electrically neutral, weakly coupled to visible matter, nonrelativistic by matter-radiation equality, and have small self-interactions [3,4]. Therefore, if DM is made of particles, its existence points to physics beyond the Standard Model (BSM), since no known SM particle can account for it.

Weakly interacting massive particles (WIMPs) are a well-motivated class of candidates [5,6], naturally realized in BSM frameworks and with properties compatible with DM [7–9]. The WIMP thermal freeze-out process,

assuming electroweak-scale masses and couplings, automatically provides a relic density of order the observed one, $\Omega_{\text{DM}} h^2 \simeq 0.12$ [10]. This is the so-called “WIMP miracle” [11], which has motivated large experimental programs in direct detection [12], collider searches [13], and indirect detection [14], complemented by precise cosmological measurements [10].

However, present direct-detection experiments such as LZ and XENONnT exclude spin-independent (SI) WIMP–nucleon cross sections down to $\mathcal{O}(10^{-48}) \text{ cm}^2$ for weak-scale masses [15–17]. In models where the same portal controls both freeze-out and scattering, these bounds exclude large portions of parameter space [18–23]. One way to evade these limits is to work near the resonant regime, in which the DM mass is about half of the mediator mass, $m_{\text{DM}} \simeq m_\phi/2$. In this regime, DM annihilation is enhanced and the couplings required to match the observed $\Omega_{\text{DM}} h^2$ can be made small enough to evade direct-detection bounds [20,22], though this typically entails some mass tuning.

A more generic solution is to retain a thermal origin for DM while decoupling the annihilation process from present-day DM-SM interactions. This is precisely the *secluded* DM scenario, originally proposed in [24,25], in

^{*}Contact author: dimauro.mattia@gmail.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

which the relic density is set by annihilations into light mediators in the dark sector (a dark Higgs ϕ or a dark gauge boson Z'), while the portal connecting to the SM (e.g., Higgs mixing or kinetic mixing, collectively denoted by ϵ) is very small. Recently, Ref. [26] thoroughly reassessed this framework in three BSM realizations, verifying that the mediator remains in thermal contact and decays before big-bang nucleosynthesis (BBN), and highlighting promising collider tests that remain sensitive despite tiny portals. For portal couplings $\epsilon \ll 10^{-3}$, the relic density is dominantly controlled by DM annihilation into mediator pairs, while SI scattering lies well below current limits and can remain viable even for future experiments approaching the neutrino floor [27]. Such a tiny portal might at first appear fine-tuned and theoretically unmotivated.

Interestingly, the SM already contains another unusually small parameter: the neutrino mass. Recent constraints from cosmology, oscillations, and direct kinematic measurements indicate neutrino masses at (and upper limits in) the tens-of-meV range [28–31], i.e., $\sim 10^{-11}$ of the proton mass and $\sim 10^{-7}$ of the electron mass.

In this work, we propose that the required smallness of the SM–dark portal is *not* fine-tuned but can be naturally connected to the small neutrino masses. Concretely, we construct a setup in which a symmetry forbids the tree-level portal, and the leading SM–dark coupling is generated radiatively by the same Yukawa structures (or a lepton-number–violating spurion) that produce neutrino masses via a type-I or inverse seesaw [32–37]. The induced portal—either a Higgs mixing angle $\sin \alpha$ or a gauge kinetic mixing ϵ —then scales with m_ν , up to loop factors and mass ratios. Consequently, direct, indirect, and collider signals are naturally suppressed—because m_ν is tiny—while the relic abundance remains set by $\chi\chi \rightarrow$ mediator pairs and is largely decoupled from the portal.

The rest of the paper is organized as follows. In Sec. II we present the model symmetries and particle content. Section III describes the neutrino interactions in the model and details the type-I seesaw mechanism. In Sec. IV we present the main features of the secluded regime, while Sec. V evaluates the loop-induced portal coupling and shows its relation to m_ν . Section VI explains high- T thermal-contact interactions and decoupling, and how DM reaches the correct relic density. Section VII reports the constraints on the model from spin-independent direct-detection experiments. Finally, in Sec. VIII we draw our conclusions.

II. MODEL LAGRANGIAN AND PARTICLE CONTENT

We consider a UV-complete dark-sector setup that realizes the secluded DM mechanism [24–26]. The model is based on an extra Abelian gauge group $U(1)_X$ with gauge boson Z'_μ (and coupling constant g_X), a Dirac fermion DM candidate χ with $U(1)_X$ charge q_X^χ (and singlet under the

SM gauge group), three right-handed neutrinos N_i ($i = 1, 2, 3$), and a real scalar singlet R that mixes with the SM Higgs doublet Φ (see, e.g., Refs. [38,39]).

The relevant Lagrangian terms are

$$\begin{aligned} \mathcal{L} \supset & -\frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} - \frac{\epsilon}{2}F'_{\mu\nu}B^{\mu\nu} + \partial_\mu R \partial^\mu R \\ & - V(\Phi, R) + \bar{\chi}(i\not{D} - m_\chi)\chi \\ & - \left(y_p \bar{\chi} \chi R + Y_\nu^{ai} \bar{L}_a \tilde{\Phi} N_i + \frac{1}{2} Y_N^{ij} R \overline{N_i^c} N_j + \text{H.c.} \right), \end{aligned} \quad (1)$$

where $L_a = (\nu_{aL}, \ell_{aL})^T$ are the SM lepton doublets, $\tilde{\Phi} = i\sigma_2 \Phi^*$, $B_{\mu\nu}$ is the hypercharge field strength, and Y_ν (3×3 complex) and $Y_N = Y_N^T$ (complex symmetric) are the neutrino Yukawa matrices. The covariant derivative is $D_\mu = \partial_\mu + iq_X g_X Z'_\mu$ (with $q_X \rightarrow q_X^\chi$ on χ), and we define the fine-structure constant of the hidden sector as $\alpha_X \equiv g_X^2/(4\pi)$.

The Lagrangian in Eq. (1) is written to emphasize the interactions relevant for the radiative generation of the Higgs-singlet portal. For completeness, one may also include the renormalizable bare Majorana mass term for the singlet neutrinos,

$$\Delta\mathcal{L} \supset -\frac{1}{2} \overline{N_i^c} (M_0)^{ij} N_j + \text{H.c.}, \quad (2)$$

which explicitly violates lepton number, as in the standard type-I seesaw. In this work we adopt the technically natural limit in which this term is set to zero (or is negligible) at the UV scale, so that the heavy-neutrino masses are generated dominantly by the Yukawa interaction proportional to Y_N once R acquires a vacuum-expectation-value (VEV). Expanding $R = (v_r + \rho)/\sqrt{2}$ in Eq. (1) yields

$$-\frac{1}{2} Y_N^{ij} R \overline{N_i^c} N_j + \text{H.c.} \supset -\frac{1}{2} \overline{N_i^c} (M_N)^{ij} N_j + \text{H.c.}, \quad (3)$$

where $(M_N)^{ij} \equiv v_r/\sqrt{2} Y_N^{ij}$. Therefore, the effective heavy-neutrino mass matrix that enters the seesaw mechanism is $M_N = v_r Y_N/\sqrt{2}$ (in the basis where it is diagonal we denote its eigenvalues by $M_{N,i}$ throughout the paper). In this option the most general mass matrix is $M_{\text{heavy}} = M_0 + v_r Y_N/\sqrt{2}$, while our baseline assumption corresponds to $M_0 = 0$, so that the origin of the heavy-neutrino scale is tied to the RNN spurion that also controls the radiatively induced portal.

The scalar potential is

$$\begin{aligned} V(\Phi, R) = & \mu_H^2 \Phi^\dagger \Phi + \mu_R^2 R^2 + \lambda_H (\Phi^\dagger \Phi)^2 \\ & + \lambda_R R^4 + \kappa (\Phi^\dagger \Phi) R^2. \end{aligned} \quad (4)$$

After electroweak symmetry breaking the scalar fields Φ and R become

$$\Phi = \begin{pmatrix} G^+ \\ \frac{v_h + h + iG^0}{\sqrt{2}} \end{pmatrix}, \quad R = \frac{v_r + \rho}{\sqrt{2}}, \quad (5)$$

with $v_h \simeq 246$ GeV. We choose R to be neutral, $q_X^R = 0$; thus $U(1)_X$ is unbroken by R and no additional Goldstone boson appears beyond the SM $G^\pm, G^0$. In the CP -even basis (h, ρ) the mass matrix is

$$\mathcal{M}^2 = \begin{pmatrix} 2\lambda_H v_h^2 & \kappa v_h v_r \\ \kappa v_h v_r & 2\lambda_R v_r^2 \end{pmatrix}, \quad \tan 2\alpha = \frac{\kappa v_h v_r}{\lambda_R v_r^2 - \lambda_H v_h^2}, \quad (6)$$

which is diagonalized by a rotation of angle α ,

$$\begin{pmatrix} H \\ H_p \end{pmatrix} = \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} h \\ \rho \end{pmatrix}, \quad (7)$$

yielding the SM-like state H and the singlet-like state H_p with masses m_H and m_{H_p} ,

$$m_{H,H_p}^2 = \lambda_H v_h^2 + \lambda_R v_r^2 \mp \sqrt{(\lambda_R v_r^2 - \lambda_H v_h^2)^2 + (\kappa v_h v_r)^2}. \quad (8)$$

We take χ to be vectorelike under $U(1)_X$, i.e., $q_{\chi_L} = q_{\chi_R} = q_X^\chi$. Since R is neutral, a bare Dirac mass term $m_0 \bar{\chi} \chi$ is gauge invariant and the renormalizable Yukawa portal $-y_p \bar{\chi} \chi R$ is allowed. Writing $R = (v_r + \rho)/\sqrt{2}$, the physical DM mass is

$$m_\chi = m_0 + \frac{y_p v_r}{\sqrt{2}}, \quad (9)$$

so y_p and m_χ are independent parameters in general. Since $q_X^R = 0$, gauge invariance of the $R \bar{N}^c N$ term implies that the heavy neutrinos are also neutral under $U(1)_X$, and therefore do not couple to Z' at tree level.

SM fermions are neutral under $U(1)_X$ as well; hence they couple to Z' only via kinetic mixing, controlled by ϵ [40,41]. In the following we neglect this mixing (or assume it is sufficiently small). Because R does not break $U(1)_X$, the Z' mass can be generated via a Stückelberg mechanism [42,43]. In the present work Z' plays only a spectator role: all SM fields are neutral under $U(1)_X$ and, with negligible kinetic mixing, Z' has no appreciable couplings to the visible sector. Moreover, since the heavy neutrinos are also neutral under $U(1)_X$, the neutrino portal and the loop-induced Higgs–singlet mixing discussed below are entirely controlled by the scalar and neutrino Yukawa sectors, independently of the gauge boson. Finally, in the secluded regime the DM freeze-out and direct-detection phenomenology are dominated by the scalar channel $\chi \bar{\chi} \rightarrow H_p H_p$ and by Higgs–singlet mixing, so our results are largely insensitive to the value of $m_{Z'}$ provided Z' is not the lightest dark-sector state.

For later use, we collect the leading fermionic interactions in the mass basis. SM fermions couple as

$$\mathcal{L}_{H,H_p-f\bar{f}} = -\sum_f \frac{m_f}{v_h} (\cos \alpha H - \sin \alpha H_p) \bar{f} f, \quad (10)$$

while the DM portal yields

$$\mathcal{L}_{H,H_p-\chi\bar{\chi}} = -\frac{y_p}{\sqrt{2}} (\sin \alpha H + \cos \alpha H_p) \bar{\chi} \chi. \quad (11)$$

Therefore, when the mixing angle is very small, as it is in the secluded regime (see Sec. IV), SM fermions couple predominantly to the SM Higgs H , while DM couples predominantly to the singlet scalar H_p .

The singlet Yukawa term generates Majorana masses for the heavy neutrinos, $M_N = \frac{v_r}{\sqrt{2}} Y_N$, and in the heavy-neutrino mass basis $\hat{M}_N = \text{diag}(M_1, M_2, M_3)$ one obtains

$$\mathcal{L}_{H,H_p-NN} = -\sum_{I=1}^3 \frac{M_I}{2v_r} (\sin \alpha H + \cos \alpha H_p) \bar{N}_I^c N_I + \text{H.c.} \quad (12)$$

Finally, the Dirac Yukawa term Y_ν controls the light-heavy couplings,

$$\mathcal{L}_{H,H_p-\nu N} = -\frac{1}{\sqrt{2}} (\cos \alpha H - \sin \alpha H_p) \bar{\nu}_{L\alpha} Y_\nu^{\alpha i} N_i + \text{H.c.} \quad (13)$$

In the secluded regime, SM neutrinos interact mainly with the SM Higgs H , while heavy neutrinos N_I interact mainly with the singlet scalar H_p .

For very small angles α , this setup yields a secluded thermal history in which the relic abundance is dominantly controlled by annihilations into dark-sector states (notably $\chi \bar{\chi} \rightarrow H_p H_p$ when kinematically open), while the portal to the SM is governed by the mixing angle α (and, if present, by kinetic mixing ϵ).

III. TYPE-I SEESAW

We summarize here the main aspects of light and heavy neutrino masses, couplings, and the seesaw mechanism. The relevant Yukawa interactions are

$$\mathcal{L}_Y \supset -Y_\nu^{\alpha i} \bar{L}_\alpha \tilde{\Phi} N_i - \frac{1}{2} Y_N^{ij} R \bar{N}_i^c N_j + \text{H.c.}, \quad (14)$$

with Y_ν a generic complex 3×3 matrix and $Y_N = Y_N^T$ complex symmetric. After electroweak symmetry breaking, the neutral components generate the Dirac and Majorana mass matrices

$$m_D \equiv \frac{v_h}{\sqrt{2}} Y_\nu, \quad M_N \equiv \frac{v_r}{\sqrt{2}} Y_N, \quad (15)$$

and induce scalar couplings to $\nu_L N$ and NN proportional to Y_ν and Y_N , respectively.

Collecting the VEV-induced terms, the Majorana neutrino mass matrix in the (ν_L, N^c) basis is

$$\mathcal{M}_\nu = \begin{pmatrix} 0 & m_D \\ m_D^T & M_N \end{pmatrix}. \quad (16)$$

In the seesaw regime $\|m_D\| \ll \|M_N\|$, defining the active-sterile mixing matrix

$$\Theta \equiv m_D M_N^{-1}, \quad (17)$$

one can block diagonalize $\mathcal{M}_\nu$ perturbatively in Θ . To leading order this yields the standard type-I seesaw relation

$$M_\nu \simeq -m_D M_N^{-1} m_D^T = -\Theta M_N \Theta^T, \quad (18)$$

while the heavy-neutrino mass matrix receives only $\mathcal{O}(\Theta^2)$ corrections, so that $M_{\text{heavy}} \simeq M_N$.

In the one-generation case with real parameters m_D and M_N , the eigenvalues for the light and heavy neutrinos are

$$m_{\text{light}} \simeq -\frac{m_D^2}{M_N}, \quad m_{\text{heavy}} \simeq M_N \left(1 + \frac{m_D^2}{M_N^2}\right). \quad (19)$$

In terms of the Yukawa couplings, using $m_D = y_\nu v_h / \sqrt{2}$ and $M_N = y_N v_r / \sqrt{2}$, the seesaw relation for one generation can be written as

$$m_\nu \simeq -\frac{y_\nu^2 v_h^2}{\sqrt{2} y_N v_r}, \quad M_{\text{heavy}} \simeq \frac{y_N v_r}{\sqrt{2}}. \quad (20)$$

Since M_N is complex symmetric, it is diagonalized by a Takagi factorization,

$$U_N^T M_N U_N = \hat{M}_N \equiv \text{diag}(M_1, M_2, M_3), \quad M_I > 0, \quad (21)$$

which defines the heavy Majorana mass eigenstates N_I . Equivalently, in terms of diagonal Yukawas one has $\hat{Y}_N \equiv U_N^T Y_N U_N = \text{diag}(y_{1,N}, y_{2,N}, y_{3,N})$ with $M_I = (v_r / \sqrt{2}) y_{I,N}$. The light-neutrino mass matrix M_ν is diagonalized by the PMNS matrix U_ν ,

$$U_\nu^T M_\nu U_\nu = \hat{m}_\nu \equiv \text{diag}(m_{\nu,1}, m_{\nu,2}, m_{\nu,3}), \quad m_{\nu,i} \geq 0, \quad (22)$$

yielding the physical light eigenstates ν_i with masses $m_{\nu,i}$. The Majorana mass term breaks global lepton number by two units ($\Delta L = 2$), so the light neutrinos are Majorana.

The Yukawa couplings to scalars are most transparent in the light/heavy mass basis. After rotating $\nu_{L\alpha} = (U_\nu)_{\alpha i} \nu_{iL}$

and $N_i = (U_N)_{iI} N_I$, the ν - N couplings can be written as

$$\mathcal{L}_{H, H_p - \nu N} = -\frac{1}{\sqrt{2}} (\cos \alpha H - \sin \alpha H_p) \overline{\nu_{iL}} (Y_\nu^{\text{mass}})_{iI} N_I, \quad (23)$$

with $Y_\nu^{\text{mass}} \equiv U_\nu^\dagger Y_\nu U_N$. This shows that, in general, each heavy state N_I couples to all light states ν_i .

In this paper we work in the basis where $\hat{M}_N$ is diagonal, and we parametrize the Dirac mass matrix via the Casas-Ibarra form [44]

$$m_D = i U_\nu \sqrt{\hat{m}_\nu} O \sqrt{\hat{M}_N}, \quad O^T O = \mathbb{1}, \quad (24)$$

where O is a complex orthogonal matrix encoding the high-energy freedom not fixed by low-energy observables. In the *aligned limit* we take $O = \mathbb{1}$ and use the freedom to label the heavy eigenstates so that each heavy mass M_I is paired with the corresponding light eigenvalue $m_{\nu,I}$. In this case Y_ν^{mass} becomes diagonal and one finds

$$(Y_\nu^\dagger Y_\nu)_{II} = \sum_{\alpha=e,\mu,\tau} |y_{\alpha I}|^2 = \frac{2m_{\nu,I} M_I}{v_h^2}, \quad (25)$$

together with the effective seesaw mass parameter

$$\tilde{m}_I \equiv \frac{(m_D^\dagger m_D)_{II}}{M_I} = m_{\nu,I} = \frac{v_h^2 |y_{\nu,I}|^2}{2 M_I}, \quad (26)$$

where $y_{\nu,I} \equiv (Y_\nu^{\text{mass}})_{II}$. Outside this aligned limit ($O \neq \mathbb{1}$), Y_ν^{mass} is nondiagonal and each light mass $m_{\nu,i}$ arises from a coherent combination of several $(Y_\nu^{\text{mass}})_{iI}$ entries; in that case the full matrix structure in Eq. (18) must be retained.

IV. WIMP AND SECLUDED REGIMES

Two qualitatively different regimes arise depending on the hierarchy between the DM mass and the singlet-like scalar mass, m_{H_p} .

WIMP-like regime ($m_\chi < m_{H_p}$). The WIMP regime applies for $m_\chi < m_{H_p}$ with $y_p \sim \mathcal{O}(0.1 - 1)$ and a sizable scalar mixing $\sin \alpha \sim \mathcal{O}(0.1)$. In this case, annihilation into SM states is efficient because the dominant channels are $\chi \bar{\chi} \rightarrow f \bar{f}$ via s -channel exchange of H and H_p ,

$$\langle \sigma v \rangle_{f \bar{f}} \propto y_p^2 (\sin \alpha \cos \alpha)^2, \quad (27)$$

so the same portal that controls annihilation also controls interactions with nuclei. For weak-scale masses, values of $y_p \sim 0.1 - 1$ can readily reproduce the observed relic density through thermal freeze-out. However, because the SI scattering amplitude is also proportional to the Higgs-singlet mixing, this region is strongly constrained by direct detection. A well-known exception occurs near the resonant regime, $m_\chi \simeq m_{H_p}/2$ (or $m_\chi \simeq m_H/2$), where annihilation

is enhanced and the required couplings can be reduced [26], at the price of some mass tuning.

Secluded regime ($m_\chi > m_{H_p}$). In the secluded case, valid for $m_\chi > m_{H_p}$ with moderately small $y_p \sim \mathcal{O}(0.1)$ and $\sin \alpha \ll 1$, the relic abundance is set primarily by annihilations into dark-sector states, most importantly $\chi\bar{\chi} \rightarrow H_p H_p$ through t/u -channel χ exchange. For a Dirac fermion with a scalar Yukawa interaction, this channel is p -wave suppressed, so parametrically

$$\sigma v_{\text{rel}}(\chi\bar{\chi} \rightarrow H_p H_p) = \frac{y_p^4}{64\pi m_\chi^2} \frac{\sqrt{1-r}}{(1-\frac{r}{2})^4} \left(1 - r + \frac{r^2}{8}\right) v_{\text{rel}}^2, \quad (28)$$

where $r = m_{H_p}^2/m_\chi^2$ and v_{rel} is the DM relative velocity.

This has two important implications. First, the freeze-out process can still be efficient because v_{rel}^2 at freeze-out is $\mathcal{O}(0.1)$, whereas today in the Galactic halo $v_{\text{rel}}^2 \sim 10^{-6}$, so late-time annihilation signals are further suppressed and indirect-detection constraints are typically very weak. Second, since the dominant annihilation does not rely on the portal to SM fields, the relic density can be achieved even when the mixing angle is extremely small.

SI scattering on nucleons proceeds through t -channel exchange of H or H_p and scales as (see Sec. VII)

$$\sigma_{\chi N}^{\text{SI}} \propto y_p^2 \cos^2 \alpha \sin^2 \alpha \propto y_p^2 (\tan 2\alpha)^2, \quad (29)$$

so in the secluded regime the nuclear cross section is naturally suppressed by $\sin^2 \alpha$. Likewise, annihilation into SM fermions [Eq. (27)] becomes negligible for $\sin \alpha \ll 1$, strongly weakening indirect-detection signatures associated with SM final states. For $y_p \sim \mathcal{O}(0.1-1)$, the process $\chi\bar{\chi} \rightarrow H_p H_p$ can nonetheless yield $\Omega_{\text{DM}} h^2 \simeq 0.12$ (see Sec. VI).

A key requirement for a *thermal* secluded history is that, despite the small portal, the mediator H_p remains sufficiently coupled to the SM plasma around freeze-out so that the dark and visible sectors share a common temperature; equivalently, the interaction rate for processes such as $H_p \leftrightarrow \text{SM}$ (or inverse decays/scatterings mediated by mixing) must exceed the Hubble expansion rate at $T \sim m_\chi/20$, and H_p must decay before BBN.

Throughout the rest of the paper we focus on the secluded regime. For this reason, we use $\tan(2\alpha)$ as a convenient measure of the portal strength: it directly quantifies the Higgs-singlet mixing induced by the scalar potential and controls both the mixing-suppressed scattering rate and the subleading annihilation channels into SM states.

V. EVALUATION OF THE LOOP-INDUCED PORTAL

A. Zero tree-level portal

The scalar portal $(R^2)(\Phi^\dagger \Phi)$ (and similarly the Abelian kinetic mixing $\epsilon F_Y^{\mu\nu} F_{\mu\nu}^D$) is gauge invariant and is therefore

generically present in a 4D renormalizable effective field theory (EFT) unless it is forbidden by additional UV structure [40,45]. In our framework we *impose* the boundary condition

$$\kappa(\Lambda) = 0 \quad (30)$$

at a UV scale Λ , and arrange that the only interactions that communicate between the SM and the dark sector are those associated with the neutrino sector, i.e., the Yukawa matrices Y_ν and Y_N (equivalently, the spurion M_N) that also generate the light-neutrino masses. In this way any nonzero portal is necessarily *radiative* and aligned with the breaking that produces m_ν . In the 't Hooft sense, the setup is technically natural: in the limit where the neutrino-sector spurions responsible for lepton-number violation are switched off, lepton number is restored and both m_ν and the portals vanish simultaneously.

A useful consequence is that renormalization group (RG) running does not regenerate a portal if it is set to zero at Λ . Below the heavy-neutrino thresholds the low-energy spectrum contains no light fields that carry simultaneously SM and dark quantum numbers. As a result, the beta function for the mixed quartic is purely multiplicative,

$$\beta_\kappa(\mu) \equiv \frac{d\kappa}{d\ln\mu} = \kappa(\mu) \mathcal{F}(\lambda_H, \lambda_R, y_p, g_i, \dots), \quad (31)$$

so $\kappa=0$ is an RG fixed point and $\kappa(\Lambda)=0$ implies $\kappa(\mu)=0$ for all $\mu < M_N$ in the absence of threshold effects. The first nonzero contribution therefore arises precisely when the heavy neutrinos are integrated out, where the neutrino Yukawas provide the only allowed “bridge” between Φ and R . Therefore, fixing $\kappa(\Lambda)=0$ keeps it zero at lower scales, up to threshold corrections at M_N .

Two standard UV mechanisms can enforce $\kappa(\Lambda)=0$ (and/or $\epsilon(\Lambda)=0$) in a technically natural way. Since this assumption is central to the predictivity of the setup, we spell out more explicitly what is meant by “enforce” and why these mechanisms are stable under radiative corrections.

For a real gauge-singlet scalar R , the most general renormalizable potential allows, in addition to the even terms shown in Eq. (4), also odd operators,

$$V_{\text{odd}}(\Phi, R) = \frac{\mu_1}{1!} R + \frac{\mu_3}{3!} R^3 + \frac{\eta_1}{2} R(\Phi^\dagger \Phi), \quad (32)$$

where Φ is the SM Higgs doublet. A tadpole term $\mu_1 R$ can always be removed by shifting $R \rightarrow R + c$ and will be set to zero in what follows. In complete analogy with the sequestered boundary condition $\kappa(\Lambda)=0$ for the quartic portal, we impose

$$\eta_1(\Lambda) = 0, \quad \mu_3(\Lambda) = 0, \quad (33)$$

so that the leading Higgs-singlet communication arises from the radiatively induced quartic portal $\kappa_{\text{loop}}(\Phi^\dagger \Phi) R^2$

discussed in Sec. VB. We interpret Eq. (33) as part of the sequestered UV boundary conditions: in the technically natural limit in which the messenger sector couples to R only through even powers, these odd operators are not regenerated by renormalization-group running below the heavy thresholds. We report in Appendix D the one-loop diagrams that generate the effective odd couplings η_1 and μ_3 (as well as their parametric estimates) in the absence of an $R \rightarrow -R$ selection rule. The main conclusions are that the neutrino-induced linear portal is extremely suppressed for $M_N \sim \mathcal{O}(10 - 50)$ TeV and $m_\nu \sim 0.05$ eV, leading to a Higgs-singlet mixing angle $\sin \alpha_{\eta_1} \lesssim 10^{-9}$ and therefore negligible Higgs-mediated direct-detection rates, whereas the cubic coupling μ_3 is typically of order $\mu_3 \sim y_N^3 M_N / (4\pi^2)$ and can significantly affect the singlet vacuum structure and the scalar mass spectrum unless it is absorbed into the renormalized potential by matching conditions.

1. Sequestering of sectors

A simple way to forbid *all* renormalizable tree-level portals is to assume that, at the UV scale Λ , the visible and dark sectors are geometrically or dynamically separated so that the renormalizable Lagrangian *factorizes*,

$$\mathcal{L}_{\text{ren}}(\Lambda) = \mathcal{L}_{\text{SM}} + \mathcal{L}_{\text{dark}} + \mathcal{L}_{\text{mess}}, \quad (34)$$

where $\mathcal{L}_{\text{mess}}$ contains only fields that couple to both sectors. In this situation operators such as $(\Phi^\dagger \Phi)(R^2)$ and $F_Y^{\mu\nu} F_{\mu\nu}^D$ are absent because they *directly* couple fields localized in different sectors. Concretely, this happens in: (i) *extradimensional* constructions where the SM and dark fields live on different branes (or at separated points in the extra dimension), so local brane interactions cannot involve fields from the other brane; and (ii) *sequestered supersymmetry* scenarios (including anomaly mediation-type setups) where locality in the Kähler potential (or conformal dynamics) suppresses contact terms between hidden and visible chiral superfields. In both cases, any interaction that couples the two sectors must proceed through *messengers* that have access to both sides (bulk fields, heavy modes, or explicit connector multiplets) [46–50].

In our EFT, the only such connector fields are the heavy neutrinos N_I (and the lepton doublets in the SM), which carry the neutrino spurions responsible for m_ν . Therefore, the first nonvanishing contribution to κ arises only radiatively and is necessarily proportional to the same spurions that break lepton number and generate neutrino masses. Put differently, sequestering is not simply an *ad hoc* setting of κ to zero: it is a UV *selection rule* stating that there are no local renormalizable operators linking the two sectors at Λ , so $\kappa(\Lambda) = 0$ is stable and can only be violated by loops of the connector fields [46–48, 51].

A potential concern is whether Planck-suppressed effects (gravity) regenerate κ . Indeed, higher-dimensional operators such as $(\Phi^\dagger \Phi)(R^\dagger R)/M_{\text{Pl}}^2$ are generically expected

unless additional structure is present. However, even if present, they yield an effective quartic of order $\kappa_{\text{grav}} \sim v^2/M_{\text{Pl}}^2 \sim 10^{-32}$ after EWSB, far below the loop-induced values relevant here; and in sequestered/anomaly-mediated constructions such contact terms are further suppressed [48, 49]. Thus, the assumption $\kappa(\Lambda) = 0$ is technically natural in the sense that taking the messenger couplings to zero restores an enhanced factorization symmetry between the sectors.

2. Gauge-portal protection (kinetic mixing)

For the kinetic mixing $\epsilon F_Y^{\mu\nu} F_{\mu\nu}^D$, the key point is that it is only gauge invariant if *both* gauge factors are Abelian in the low-energy EFT. A robust way to forbid it in the UV is to embed the dark $U(1)_D$ into a non-Abelian group,

$$G_D \supset U(1)_D, \quad (35)$$

so that above the breaking scale the operator $F_Y F_D$ is not gauge invariant (because there is no gauge-invariant dimension-4 operator linear in a non-Abelian field strength). Then $\epsilon(\Lambda) = 0$ is enforced by gauge symmetry, and kinetic mixing can only appear *after* G_D breaks to $U(1)_D$, through loops of states that simultaneously carry hypercharge and the appropriate G_D quantum numbers [40, 45, 52–54]. In our setup there are no such light bicharged states, so any induced ϵ is again loop- and threshold suppressed, and can naturally be aligned with the same spurions that connect the sectors (e.g., the neutrino portal) [45, 52].

A complementary (and often orthogonal) protection is a discrete symmetry in the dark sector, such as a dark charge conjugation C_D under which the dark gauge field flips sign, $A_\mu^D \rightarrow -A_\mu^D$ (equivalently $F_{\mu\nu}^D \rightarrow -F_{\mu\nu}^D$). Since $F_Y^{\mu\nu}$ is even and $F_{\mu\nu}^D$ is odd, the operator $F_Y^{\mu\nu} F_{\mu\nu}^D$ is odd and therefore forbidden as long as C_D is exact. In this case ϵ can only be generated by interactions that *also* break C_D . If the only C_D -breaking couplings are the same spurions that ultimately generate neutrino masses (e.g., the Majorana-sector spurions M_N in our construction), then ϵ is forced to be radiative and proportional to those spurions, mirroring the logic for κ and ensuring that taking the spurions to zero simultaneously restores $m_\nu = 0$ and $\epsilon = 0$ [25, 45, 52, 55].

If additional UV messengers exist, they generically induce an extra threshold $\Delta\kappa_{\text{UV}}$; in that case κ_{loop} should be viewed as a guaranteed minimal contribution and the phenomenology interpolates toward the usual Higgs-portal scenario when $|\Delta\kappa_{\text{UV}}| \gg |\kappa_{\text{loop}}|$.

Why this is technically natural. In both mechanisms, setting $\kappa(\Lambda) = 0$ (and/or $\epsilon(\Lambda) = 0$) is not a tuning of unrelated parameters: it corresponds to an enhanced symmetry (or factorization) of the UV theory. Radiative corrections cannot generate the forbidden operators unless the symmetry-breaking messenger interactions are present.

This is precisely why, once the heavy neutrinos are integrated out, the RG evolution of κ is multiplicative and the origin is an RG fixed point. In other words, the portal is small because it is *radiatively tied* to the same spurions that make m_ν small, and any additional UV completion that respects the above structures will reproduce the same qualitative conclusion.

B. Loop-induced portal from the neutrino sector

Assuming the tree-level operator $\kappa_0(\Phi^\dagger\Phi)(R^2)$ is absent, the leading interaction between Φ and R is generated by the one-loop box diagram in Fig. 1, involving two ΦLN insertions and two RNN insertions. The matching between the interaction term $\kappa_0(\Phi^\dagger\Phi)(R^2)$ and the loop-induced one is most transparent in the limit of vanishing external momenta p_{ext} . Since in our setup the heavy-neutrino masses are $M_{N,I} \gtrsim \mathcal{O}(10)$ TeV, all phenomenological applications satisfy $p_{\text{ext}}^2 \ll M_{N,I}^2$, so the zero-momentum matching is an excellent approximation: higher orders in $p_{\text{ext}}^2/M_{N,I}^2$ correspond to derivative operators and do not affect the quartic portal.

For a single generation, dimensional regularization and $\overline{\text{MS}}$ renormalization yield the Wilson coefficient

$$\kappa_{\text{loop}}(\mu) = \frac{|Y_N|^2 |Y_\nu|^2}{16\pi^2} \left[\ln \frac{M_N^2}{\mu^2} - 1 \right], \quad (36)$$

so choosing $\mu \simeq M_N$ minimizes the logarithm and gives the finite threshold correction

$$\kappa_{\text{loop}}(M_N) = -\frac{|Y_N|^2 |Y_\nu|^2}{16\pi^2}. \quad (37)$$

Using the one-generation seesaw relation in the aligned limit, $(Y_\nu^\dagger Y_\nu) = 2m_\nu M_N / v_h^2$ with $M_N = y_N v_r / \sqrt{2}$, this can be rewritten as

$$\kappa_{\text{loop}}(M_N) = -\frac{y_N^2 M_N}{8\pi^2 v_h^2} m_\nu, \quad (38)$$

making explicit the parametric scaling $\kappa_{\text{loop}} \propto m_\nu$. The overall sign is scheme dependent, but physical quantities




FIG. 1. Box diagram generating the mixed quartic $(R^2)(\Phi^\dagger\Phi)$ at one loop from neutrino Yukawa interactions.

such as the mixing angle and scalar masses are RG- and scheme independent.

For three generations, in the heavy-neutrino mass basis and in the aligned Casas-Ibarra limit one finds stepwise threshold matching at each $\mu \simeq M_{N,I}$,

$$\kappa_{\text{loop}} = -\sum_{I=1}^3 \frac{y_{N,I}^2 M_{N,I}}{8\pi^2 v_h^2} m_{\nu,I}, \quad (39)$$

where each term is understood as the finite threshold left after integrating out N_I at its mass. If the $M_{N,I}$ are nondegenerate, a single choice of μ cannot remove all logarithms simultaneously; the correct multiscale treatment is stepwise EFT matching. Below the thresholds, the RG evolution remains purely multiplicative and changes κ by at most a factor of a few between M_N and the electroweak scale, so Eq. (39) provides the relevant effective portal for freeze-out, direct detection, and collider observables.

C. Relation between the mixing angle and m_ν

At tree level the scalar mixing angle α between the SM-like Higgs H and the singlet-like state H_p is controlled by κ [Eq. (6)]. Using $\kappa = \kappa_{\text{loop}}$ and $M_N = y_N v_r / \sqrt{2}$, one obtains for a single generation

$$\begin{aligned} \tan 2\alpha &= \frac{\kappa v_h v_r}{\lambda_R v_r^2 - \lambda_H v_h^2} = \frac{2\kappa v_h v_r}{m_{H_p}^2 - m_H^2} \\ &= -\frac{|m_\nu| y_N M_N^2}{2\sqrt{2}\pi^2 v_h (m_{H_p}^2 - m_H^2)} \\ &\simeq -\frac{4.66 \times 10^{-11}}{1 - (m_{H_p}/125 \text{ GeV})^2} y_N \left(\frac{|m_\nu|}{0.05 \text{ eV}} \right) \left(\frac{M_N}{10 \text{ TeV}} \right)^2. \end{aligned} \quad (40)$$

For three generations (e.g., in the aligned Casas-Ibarra limit) the result generalizes via $y_N M_N^2 |m_\nu| \rightarrow \sum_I y_{N,I} M_{N,I}^2 |m_{\nu,I}|$.

Therefore, the Higgs–singlet mixing that controls the SM–dark portal is directly proportional to the light-neutrino masses. In particular, for representative seesaw parameters the induced mixing is of order $\tan 2\alpha \sim 10^{-11}$, rendering direct, indirect, and collider signals essentially undetectable with current and near-future experiments, while keeping a consistent thermal secluded history (see also Ref. [26]).

VI. DARK MATTER RELIC DENSITY

Here we detail how DM attains the observed relic density, following the thermal history of the visible and hidden sectors from very high temperatures, when the heavy neutrinos are ultrarelativistic ($T \gg M_N$), down to the BBN epoch at $t_{\text{BBN}} \sim 1$ s.

We adopt the secluded regime $m_\chi > m_{H_p}$ with $\tan(2\alpha) \ll 1$, so that the relic density is controlled by

$\chi\bar{\chi} \rightarrow H_p H_p$ evaluated at the *dark* temperature T' , while the Higgs portal is too feeble to maintain (or reestablish) equilibrium with the SM once the heavy-neutrino bridge has disappeared. In this regime, the portal affects the relic density mainly *indirectly*: it sets the H_p lifetime and determines whether the dark bath can safely dump its entropy into the SM before BBN.

A. BBN requirement: Prompt H_p decays

We start from the end of the thermal history. After DM freezes out in the dark sector, the dominant annihilation products are long-lived dark scalars H_p that remain part of the dark bath and subsequently decay into SM states through their Higgs admixture. If these decays occur after BBN, the injected electromagnetic and hadronic energy can spoil light-element abundances. We therefore require $\tau(H_p) \lesssim t_{\text{BBN}} \sim 1$ s.

For a Dirac fermion f with $\mathcal{L} \supset -g_{H_p f \bar{f}} H_p \bar{f} f$, the decay width of H_p into $f \bar{f}$ is given by

$$\Gamma_{\text{ferm}}(H_p) = \sum_f \frac{N_c^f m_{H_p} m_f^2}{8\pi v_h^2} \sin^2 \alpha \left(1 - \frac{4m_f^2}{m_{H_p}^2}\right)^{3/2}, \quad (41)$$

where N_c^f is the color factor and the sum is performed over all kinematically allowed f with $m_{H_p} > 2m_f$. Neglecting threshold effects and summing over open fermionic channels for $m_{H_p} \in [20, 100]$ GeV, a convenient parametrization is

$$\tau_{\text{ferm}}(H_p) = \frac{1}{\Gamma_{\text{ferm}}(H_p)} \approx \frac{4.3 \times 10^{-21} \text{ s}}{(\tan 2\alpha)^2} \left(\frac{20 \text{ GeV}}{m_{H_p}}\right), \quad (42)$$

up to $\mathcal{O}(1)$ factors from channel thresholds. Using the loop-induced mixing and the $\tan 2\alpha$ relation of Eq. (41) one obtains for one generation

$$\tau_{\text{ferm}}(H_p) \simeq \left(\frac{1.20 \text{ s}}{y_N^2}\right) \left(\frac{20 \text{ GeV}}{m_{H_p}}\right) \left(\frac{16 \text{ TeV}}{M_N}\right)^4 \left(\frac{0.05 \text{ eV}}{m_\nu}\right)^2 \times \left[1 - \left(\frac{m_{H_p}}{125 \text{ GeV}}\right)^2\right]^2. \quad (43)$$

This makes explicit the strong sensitivity to the heavy neutrinos $\tau_{\text{ferm}}(H_p) \propto M_N^{-4}$, which ultimately traces back to the radiative origin of the portal and the seesaw suppression: $\tan(2\alpha) \propto \kappa_{\text{loop}} \propto m_\nu M_N^3$ (for fixed y_N and v_r). For three generations the same scaling holds with the replacement $y_N M_N^2 m_\nu \rightarrow \sum_I y_{N,I} M_{N,I}^2 m_{\nu,I}$.

Figure 2 confirms this scaling by showing the (m_{H_p}, M_N) region where $H_p \rightarrow f \bar{f}$ occurs before BBN ($\tau_{\text{ferm}} < \tau_{\text{BBN}}$), assuming $y_N = 1$ and $m_\nu = 0.05$ eV. For $m_{H_p} \lesssim 100$ GeV, larger m_{H_p} typically requires smaller M_N to keep $\tau_{\text{ferm}} < 1$ s, while a pronounced dip appears as




FIG. 2. Values of the heavy-neutrino mass M_N and dark scalar mass m_{H_p} for which the decay of H_p occurs before BBN. We fix $y_N = 1$ and $m_\nu = 0.05$ eV.

$m_{H_p} \rightarrow m_H$ due to the $1/(m_{H_p}^2 - m_H^2)^2$ enhancement in $\tan 2\alpha$. For $m_{H_p} > m_H$, the trend reverses and the allowed M_N increase with m_{H_p} . Away from the near-degeneracy region, one typically needs $M_N \gtrsim 10$ TeV to ensure $\tau_{\text{ferm}} \lesssim 1$ s.

Since we work with WIMP-range m_χ and require $m_\chi > m_{H_p}$ (secluded regime), and because BBN safety favors multi-tens-TeV M_N to guarantee prompt H_p decays into SM fermions, in what follows we assume the following mass hierarchy

$$M_N > m_\chi > m_{H_p}. \quad (44)$$

B. Dark matter thermal history

When several heavy Majorana eigenstates N_I are present, the slowest (last) to lose equilibrium is the *lightest* one. It is therefore sufficient (and conservative) to track a single state N with mass $M_N \equiv \min_I(M_{N,I})$. In the aligned limit we denote by m_ν the light eigenvalue paired with this lightest N .

1. Era I ($T \gg M_N$): Full thermal contact

For temperatures above the heavy-neutrino threshold the SM and the hidden sector are connected by the neutrino portal. The relevant processes include:

- (i) decays/inverse decays $N \leftrightarrow L_\alpha H, \bar{L}_\alpha H^\dagger$;
- (ii) $\Delta L = 1$ scatterings with one neutrino Yukawa and one SM coupling, such as $N L_\alpha \leftrightarrow Q_3 t$ and $L_\alpha A \leftrightarrow N H$;
- (iii) $\Delta L = 2$ scatterings induced by virtual N ; and
- (iv) portal reactions involving the dark scalar R (and via R the DM), including $N R \leftrightarrow L_\alpha H$ and, in principle, $R R \leftrightarrow H H$.

A complete discussion of these channels, including reduced cross sections and reaction densities, is given in Appendixes A and B.

For the lightest state N (mass M_N), the thermally averaged decay rate in the unbroken phase and aligned limit is

$$\langle \Gamma_D \rangle_T = \Gamma_D \frac{K_1(\xi)}{K_2(\xi)}, \quad \Gamma_D = \frac{m_\nu M_N^2}{4\pi v_h^2}, \quad \xi \equiv \frac{M_N}{T}, \quad (45)$$

to be compared with $H(T) = 1.66\sqrt{g_*}T^2/M_{\text{Pl}}$. Solving $\langle \Gamma_D \rangle_T = H(T)$ in the relativistic regime gives

$$\frac{T_*}{M_N} \simeq \left[\frac{m_\nu M_{\text{Pl}}}{13.28\pi\sqrt{g_*}v_h^2} \right]^{1/3} \simeq 2.8 \left(\frac{m_\nu}{0.05 \text{ eV}} \right)^{1/3} \left(\frac{106.75}{g_*} \right)^{1/6}, \quad (46)$$

so N is in thermal equilibrium with the SM bath for $T \gtrsim \text{few} \times M_N$. At $T = M_N$ one finds

$$\left. \frac{\langle \Gamma_D \rangle_T}{H} \right|_{T=M_N} \sim \mathcal{O}(20-30) \quad \text{for } m_\nu \sim 0.05 \text{ eV}, \quad (47)$$

showing that decays/inverse decays alone efficiently maintain contact near and above $T \sim M_N$ for seesaw-motivated parameters.

The leading $\Delta L = 1$ scatterings provide an additional contribution. Summing the dominant top- and gauge-assisted channels one finds

$$\frac{\Gamma_{\text{scatt}}^{(\Delta L=1)}(T)}{H(T)} \simeq \frac{2c_1}{1.66\sqrt{g_*}} \frac{m_\nu M_{\text{Pl}}}{v_h^2} (y_t^2 + g^2) \frac{M_N}{T} \sim \mathcal{O}(1) \frac{M_N}{T}, \quad (48)$$

with $c_1 \simeq (3-7) \times 10^{-4}$. Thus $\Delta L = 1$ scattering rates are comparable to H around $T \sim M_N$ (and subleading at $T \gg M_N$), reinforcing equilibration that is dominated by decays/inverse decays. In contrast, $\Delta L = 2$ rates become relevant only at very high temperatures ($T \gtrsim 10^{12}$ GeV) and are negligible here.

Among the portal processes, $RR \leftrightarrow HH$ is controlled by the loop-induced quartic and is completely negligible: since $\sigma(RR \rightarrow HH) \sim \kappa_{\text{loop}}^2/(16\pi s)$, one has $\Gamma_{RR \rightarrow HH} \sim \kappa_{\text{loop}}^2 T \ll H$ for the extremely small κ_{loop} implied by the seesaw. Conversely, the mixed Yukawa process $NR \leftrightarrow L_\alpha H$, controlled by y_N and y_ν , can be non-negligible. For $y_N = \mathcal{O}(1)$ it satisfies $\Gamma/H \sim 0.1-0.2$ near $T \sim M_N$ (Appendix B) and provides an additional energy-transfer channel between the two sectors. In practice, once N is thermalized by decays/inverse decays, even a subleading $NR \leftrightarrow LH$ rate is sufficient to keep R (and via y_p also χ) in equilibrium with the SM at early times.

Overall, in Era I the neutrino portal is efficient: decays/inverse decays and $\Delta L = 1$ scatterings ensure N tracks the




FIG. 3. Per-particle rates vs expansion rate for the neutrino-portal processes as a function of $\xi = M_N/T$: decay/inverse decay $N \leftrightarrow LH$ (red solid), $\Delta L = 1$ scatterings (blue dashed), $NR \leftrightarrow L_\alpha H$ (black dotted), and $\Delta L = 2$ scatterings (orange dot dashed). We fix $M_N = 20$ TeV, $m_\nu = 0.05$ eV, and $y_N = 1$. The green band shows $H(T)$; the cyan band reflects the uncertainty on the coefficient c_1 in the $\Delta L = 1$ rate. The vertical yellow band marks the transition to the nonrelativistic regime for N where its abundance becomes Boltzmann suppressed.

SM bath, while the $NR \leftrightarrow LH$ channel (for sizeable y_N) communicates this equilibrium to the dark sector.

In Fig. 3 we compare the relevant per-particle rates to $H(T)$ ¹ as a function of ξ for the benchmark $M_N = 20$ TeV, $m_\nu = 0.05$ eV, and $y_N = 1$. The thermally averaged decay/inverse-decay rate $\langle \Gamma_D \rangle$ exceeds $H(T)$ around and above $T \sim M_N$, ensuring efficient equilibration between N and the SM bath. The $\Delta L = 1$ scatterings, shown as a band from the uncertainty in c_1 , contribute at the level of H near $T \sim M_N$ and remain comparable to H for $T \gtrsim M_N$, reinforcing thermal contact. The portal scattering $NR \leftrightarrow L_\alpha H$ is also efficient for the chosen parameters [$y_N = \mathcal{O}(1)$], with $\Gamma/H \sim \mathcal{O}(0.1-1)$ throughout the range where N is still abundant; once N is in equilibrium, the dark scalar and χ are therefore brought into equilibrium with the SM as well. In contrast, the $\Delta L = 2$ rates remain well below $H(T)$ at these scales and do not affect equilibration. For $T \lesssim M_N$ ($\xi \gtrsim 1$), the number density of N becomes Boltzmann suppressed; all N -mediated processes effectively decouple, and the neutrino portal ceases to maintain thermal contact.

2. Era II ($T \lesssim M_N$): Portal decoupling and entropy bookkeeping

As T drops below the (lightest) heavy-neutrino mass, the equilibrium number density of N becomes Boltzmann

¹The figure shows *per-particle* rates Γ , which by construction do not exhibit Boltzmann suppression for $T \lesssim M_N$. That suppression enters instead through the *reaction density* $\gamma(T) \equiv n_N^{\text{eq}}(T)\langle \Gamma \rangle$, with $n_N^{\text{eq}}(T) \propto (M_N T)^{3/2} e^{-M_N/T}$.

suppressed. Even though the *per-particle* rates in Fig. 3 remain finite, the *reaction densities* $\gamma(T) = n_N^{\text{eq}}(T)\langle\Gamma\rangle$ inherit the $e^{-M_N/T}$ suppression, and the total energy exchange between the SM and dark bath rapidly shuts off. We define the decoupling temperature T_{dec} as the epoch when the *energy-transfer* rate falls below H ; numerically this typically occurs at

$$T_{\text{dec}} \sim (0.2 - 0.3)M_N, \quad (49)$$

with mild dependence on y_N and on the precise criterion (see Appendix C).

After visible-dark decoupling, the *comoving entropies* of the two sectors are separately conserved,

$$s_{\text{vis}}(T)a^3 = \text{const}, \quad s_{\text{dark}}(T', \{\mu_i\})a^3 = \text{const}, \quad (50)$$

and the temperature ratio $\zeta \equiv T'/T$ is fixed by entropy conservation. If (i) each bath can be treated as radiation with $\mu_i = 0$ and (ii) number-changing processes inside the dark bath are efficient enough to enforce $\mu_i = 0$, one obtains the standard relation [56]

$$\zeta(T) = \left[\frac{g_{*S}^{\text{vis}}(T)}{g_{*S}^{\text{vis}}(T_{\text{dec}})} \right]^{1/3} \left[\frac{g_{*S}^{\text{dark}}(T'_{\text{dec}})}{g_{*S}^{\text{dark}}(T')} \right]^{1/3}. \quad (51)$$

This simplified expression holds only while the assumptions above remain valid. In particular, when H_p (or χ) becomes semi-/nonrelativistic or after dark chemical freeze-out (when $\mu_i \neq 0$), one must use the general entropy expression (see Appendix C). In typical GeV–TeV DM benchmarks with multi-TeV M_N , one finds $\zeta \simeq 1$ –2 at the epoch relevant for DM freeze-out, depending on the dark spectrum.

We stress that after T_{dec} the Higgs portal does not reequilibrate the two sectors. This follows from the radiative nature of the portal: $\kappa(M_N) = \kappa_{\text{loop}} \sim 10^{-13} - 10^{-12}$ for seesaw-motivated parameters, and its subsequent RG evolution is purely multiplicative and changes κ by at most an $\mathcal{O}(1$ –10) factor down to the electroweak scale. Consequently, all Higgs-portal processes remain negligible for thermal contact at late times.

3. Era III: Freeze-out in a decoupled dark bath with $T' = \zeta T$

For $m_\chi > m_{H_p}$, the annihilation $\chi\bar{\chi} \rightarrow H_p H_p$ keeps the dark bath internally thermalized at temperature $T' = \zeta T$, while the portal is far too small to reequilibrate with the SM. Working with $x \equiv m_\chi/T$ and $x' \equiv m_\chi/T' = x/\zeta$, the Boltzmann equation for the visible-entropy yield $Y_\chi \equiv n_\chi/s_{\text{vis}}$ reads

$$\frac{dY_\chi}{dx} = -\frac{s_{\text{vis}}(T)}{xH(T)} \langle\sigma v\rangle_{\chi\bar{\chi} \rightarrow H_p H_p}(T') [Y_\chi^2 - (Y_\chi^{\text{eq}}(T'))^2], \quad (52)$$

with $Y_\chi^{\text{eq}}(T') = n_\chi^{\text{eq}}(T')/s_{\text{vis}}(T) \propto x'^2 K_2(x')$ the equilibrium abundance evaluated at the *dark* temperature. During radiation domination, the Hubble rate depends on the *total* radiation energy density from both sectors,

$$H(T) = \frac{\pi}{\sqrt{90} M_{\text{Pl}}} T^2 [g_{*}^{\text{vis}}(T) + g_{*}^{\text{dark}}(T') \zeta^4]^{1/2} = \quad (53)$$

$$= \frac{1.66 \sqrt{g_{*}^H(T, T')}}{M_{\text{Pl}}} T^2, \quad (54)$$

where $g_{*}^H(T, T') \equiv g_{*}^{\text{vis}}(T) + \zeta^4 g_{*}^{\text{dark}}(T')$. In our minimal setup g_{*}^{dark} is small, so the correction to H is typically subleading, but we keep it for completeness.

The annihilation $\chi\bar{\chi} \rightarrow H_p H_p$ is p wave in the limit of a pure scalar coupling $y_p \bar{\chi}\chi H_p$, so its thermal average scales as $\langle\sigma v\rangle \propto 1/x'$ (see Appendix C). The freeze-out condition is $n_\chi^{\text{eq}}(T') \langle\sigma v\rangle(T') \simeq H(T)$, which yields the usual logarithmic solution $x'_f \simeq 20$ –30 for m_χ in the GeV–TeV range and $y_p \sim \mathcal{O}(0.1$ –1), similar to standard WIMPs but evaluated at the dark temperature.

The relic density follows from the asymptotic yield $Y_{\chi, \infty}$,

$$\Omega_\chi h^2 \simeq (1.05 \times 10^9 \text{ GeV}^{-1}) \frac{\sqrt{g_{*}^H(T_f, T'_f)} \zeta x_f'^2}{g_{*S}^{\text{vis}}(T_f)} \frac{64\pi m_\chi^2}{M_{\text{Pl}} 3y_p^4}, \quad (55)$$

which makes explicit two key effects of the two-temperature setup: (i) the explicit ζ factor from the annihilation integral, $J_{2T} \propto 1/\zeta$ for p -wave freeze-out, and (ii) the mild dependence on the total expansion rate through g_{*}^H . Inverting Eq. (55), one finds the coupling required to reproduce $\Omega_\chi h^2 \simeq 0.12$,

$$y_p \simeq 0.43 \left(\frac{x'_f}{25} \right)^{1/2} \left(\frac{m_\chi}{100 \text{ GeV}} \right)^{1/2} \left(\frac{g_{*S}^{\text{vis}}}{86.25} \right)^{-1/4} \left(\frac{g_{*}^H}{100} \right)^{1/8}, \quad (56)$$

i.e., $y_p \sim \mathcal{O}(0.1$ –1) over the mass range of interest.

Equivalently, a useful heuristic for the *required* thermal rate is

$$\langle\sigma v\rangle_{\text{req}} \approx \frac{2 \times 10^{-26} \text{ cm}^3 \text{ s}^{-1}}{\zeta} \left[1 + \mathcal{O}\left(\frac{g_{*}^{\text{dark}}}{g_{*}^{\text{vis}}} \zeta^4 \right) \right], \quad (57)$$

so a hotter dark bath ($\zeta > 1$) requires a *smaller* annihilation cross section than the canonical WIMP value. Physically, for fixed x'_f the dark bath has a larger equilibrium density $n_\chi^{\text{eq}}(T')$ than the visible bath would have at the same T , and freeze-out therefore occurs for a slightly reduced annihilation strength.

Finally, we emphasize the internal consistency of the thermal history: (i) at early times the neutrino portal ensures full equilibrium, so the initial condition for Y_χ is the equilibrium abundance at $T' = T$; (ii) once N becomes Boltzmann suppressed, the sectors decouple and a temperature ratio ζ is generated by subsequent entropy dumps; (iii) the radiatively induced Higgs portal remains far too small to undo this decoupling, both because κ_{loop} is tiny and because its RG running is only multiplicative; and (iv) DM freezes out in the dark bath through $\chi\bar{\chi} \rightarrow H_p H_p$ and the resulting long-lived H_p decays before BBN in the parameter region of Fig. 2.

VII. NUCLEAR CROSS-SECTION UPPER LIMITS

Direct-detection experiments constrain the SI elastic scattering of DM off nuclei. In our setup, SI scattering proceeds at tree level via t -channel exchange of the two CP -even scalars H and H_p . The rate is therefore controlled by the Higgs-singlet mixing angle α . Crucially, in the neutrino-aligned scenario considered here the mixing (equivalently the portal quartic κ) is radiatively generated and proportional to the light neutrino masses, which implies a dramatic suppression of the nuclear cross section.

The LZ experiment sets currently the strongest constraints on SI WIMP-nucleon scattering in the WIMP-mass range: its most stringent limit is $2.2 \times 10^{-48} \text{ cm}^2$ at $m_\chi \simeq 40 \text{ GeV}$ (90% CL), with comparable bounds $\sim \mathcal{O}(10^{-48}) \text{ cm}^2$ across $m_\chi \sim 10\text{--}10^3 \text{ GeV}$ [17]. The SI cross section on a nucleon is

$$\sigma_{\chi N}^{\text{SI}} = \frac{\mu_N^2 m_N^2 f_N^2 y_p^2}{2\pi v_h^2} \sin^2 \alpha \cos^2 \alpha \left(\frac{1}{m_H^2} - \frac{1}{m_{H_p}^2} \right)^2, \quad (58)$$

where $\mu_N = m_\chi m_N / (m_\chi + m_N)$ is the χ - N reduced mass and $f_N \simeq 0.30$ is the nucleon scalar form factor. Equation (58) exhibits a “blind spot” for $m_{H_p} \simeq m_H$, where the H and H_p contributions cancel.

For the secluded benchmarks of interest, $|\alpha| \ll 1$ and typically $m_{H_p} \ll m_H$. In this regime one has $\sin \alpha \cos \alpha \simeq \frac{1}{2} \sin 2\alpha \simeq \frac{1}{2} \tan 2\alpha$ and $(1/m_H^2 - 1/m_{H_p}^2)^2 \simeq 1/m_{H_p}^4$, so that

$$\sigma_{\chi N}^{\text{SI}} \propto y_p^2 (\sin \alpha \cos \alpha)^2 \frac{1}{m_{H_p}^4} \simeq \frac{y_p^2}{4} \tan^2(2\alpha) \frac{1}{m_{H_p}^4}. \quad (59)$$

This makes the parametric suppression transparent: the same tiny mixing that sequesters the hidden sector from the SM suppresses the nuclear cross section quadratically, and the neutrino-aligned origin of α enforces this smallness.

A more useful expression is obtained by trading the mixing angle for the portal quartic κ through Eq. (38), and inserting the loop-induced κ_{loop} generated at the heavy-neutrino thresholds. The result can be written entirely in terms of physical parameters, with *no* residual dependence

on the singlet VEV v_ν ,

$$\sigma_{\chi N}^{\text{SI}} = \frac{\mu_N^2 m_N^2 f_N^2 y_p^2}{64\pi^5 m_H^4 m_{H_p}^4 v_h^4} \left(\sum_{I=1}^3 y_{N,I}^2 m_{\nu,I}^2 M_{N,I}^4 \right). \quad (60)$$

Therefore, direct detection constrains the specific neutrino-spurion combination $\sum_I y_{N,I}^2 m_{\nu,I}^2 M_{N,I}^4$ that controls the radiative portal. The strong $M_{N,I}^4$ scaling in Eq. (60) arises after rewriting the result in terms of the *physical* seesaw parameters $m_{\nu,I}$ rather than the Dirac Yukawas Y_ν .

Numerically, for $m_\chi \gg m_N$ (so $\mu_N \simeq m_N$),

$$\sigma_{\chi N}^{\text{SI}} \simeq 3.9 \times 10^{-57} \text{ cm}^2 \left(\frac{y_p}{1} \right)^2 \left(\frac{10 \text{ GeV}}{m_{H_p}} \right)^4 \cdot \sum_{I=1}^3 \left(\frac{y_{N,I}}{1} \right)^2 \left(\frac{m_{\nu,I}}{0.05 \text{ eV}} \right)^2 \left(\frac{M_{N,I}}{10 \text{ TeV}} \right)^4. \quad (61)$$

Even for $y_p \sim \mathcal{O}(1)$ (as required by the relic density in Sec. VI), $y_{N,I} \sim \mathcal{O}(1)$, and $m_{H_p} \sim 10\text{--}100 \text{ GeV}$, this is many orders of magnitude below current SI limits, because the rate is suppressed by $m_{\nu,I}^2$ and by the loop factor inherent to the radiative portal.

Fixing y_p to the value that reproduces $\Omega_\chi h^2$ [Eq. (56)], the LZ bound can be translated into a conservative upper limit on the heavy neutrino scale. For one generation and $m_\chi = 100 \text{ GeV}$ one finds

$$M_N \lesssim 1.7 \times 10^3 \text{ TeV} \left(\frac{1}{y_p y_N} \right)^{1/2} \left(\frac{0.05 \text{ eV}}{m_\nu} \right)^{1/2} \left(\frac{m_{H_p}}{10 \text{ GeV}} \right). \quad (62)$$

In Fig. 4 we show the bounds on M_N as a function of m_χ , obtained by applying the SI limits from Ref. [17] to our




FIG. 4. Upper limit on the heavy-neutrino mass M_N as a function of the DM mass m_χ , obtained by applying the spin-independent limits from Ref. [17] to our model [see Eq. (61)].

model. This bound lies far above the multi-TeV range relevant for our thermal history (and the BBN requirement on the H_p lifetime in Sec. VI). Therefore, throughout the parameter space of interest [$M_N \sim \mathcal{O}(10\text{--}100)$ TeV, $y_N \sim 1$, $m_{H_p} \sim \text{few} - 100$ GeV], the loop-induced portal is safely consistent with current direct-detection limits and remains well below the neutrino floor.

In summary, direct-detection experiments do constrain the radiative portal, but the suppression by the neutrino masses implies that current LZ limits only exclude extremely large heavy-neutrino masses (up to the few-PeV scale for $y_p, y_N \sim 1$ and $m_{H_p} \sim 10$ GeV). In practice, for the seesaw-motivated multi-TeV regime relevant for our thermal history, the SI cross section is many orders of magnitude below the neutrino floor and well beyond the reach of next-generation xenon-based detectors.

VIII. CONCLUSIONS

We have presented a predictive realization of thermal *secluded* DM in which the leading renormalizable communication between the SM and the hidden sector is absent at tree level and is generated *radiatively* by the same spurions that control light-neutrino masses. The model extends the SM by an Abelian gauge group $U(1)_X$ and includes a Dirac fermion χ (the DM candidate), a Higgs-like singlet scalar R , and three heavy singlet neutrinos N_I implementing a type-I seesaw. The DM particle χ is charged under $U(1)_X$, while R is neutral, so that its VEV does not break the new gauge symmetry.

A central assumption of the setup is that the usual renormalizable portals between the two sectors—the Higgs portal $(R^2)(\Phi^\dagger\Phi)$ and, optionally, the Abelian kinetic mixing $\epsilon F_Y^{\mu\nu} F_{\mu\nu}^X$ —are set to zero at a UV scale Λ . This can be technically natural in standard UV realizations (e.g., sequestering or symmetry protection), and implies that the leading visible-dark interaction is induced at one loop by the neutrino sector. Integrating out the heavy singlets generates a mixed quartic κ_{loop} and hence a Higgs-singlet mixing angle α that is directly correlated with the physical seesaw parameters. Parametrically,

$$\tan(2\alpha) \sim \sum_I 5 \times 10^{-11} \left(\frac{m_{\nu,I}}{0.05 \text{ eV}} \right) \left(\frac{M_{N,I}}{10 \text{ TeV}} \right)^2 \times \mathcal{O}(1), \quad (63)$$

so that the portal strength is ultimately controlled by the smallness of the SM active neutrino masses m_ν . In this sense, the null results of WIMP searches are not the consequence of tuning a generic Higgs portal to a tiny value: once the tree-level portal is absent, the radiatively generated portal is automatically minuscule because it is proportional to the same neutrino-sector parameters that generate m_ν .

We have focused on the phenomenologically relevant *secluded* regime with $m_\chi > m_{H_p}$ and $\tan(2\alpha) \ll 1$, where H_p denotes the singlet CP -even scalar. In this regime the relic density is set dominantly by hidden-sector annihilations, $\chi\bar{\chi} \rightarrow H_p H_p$, while annihilations into SM final states are suppressed by the tiny mixing. The dominant channel is p wave, so the late-time annihilation rate is further reduced compared to standard s -wave WIMPs. After the neutrino portal decouples around $T \sim M_N$, the dark bath can remain at a temperature $T' = \zeta T$ (typically $\zeta \simeq 1\text{--}2$), and DM freezes out in the dark sector. Reproducing $\Omega_{\text{DM}} h^2 \simeq 0.12$ requires a Yukawa coupling $y_p = \mathcal{O}(0.1\text{--}1)$, i.e., WIMP-like values, but with the crucial difference that present-day signals are not tied to the thermal annihilation rate because the SM portal is radiative and extremely small. Requiring that the mediator subsequently decays sufficiently early, $H_p \rightarrow f\bar{f}$, before BBN, selects heavy-neutrino masses in the multi-TeV range for $\mathcal{O}(1)$ neutrino-sector couplings in the aligned limit.

A key phenomenological consequence is the predicted size of the SI WIMP-nucleon cross section. In our model SI scattering arises from t -channel exchange of H and H_p and is entirely controlled by the mixing angle α . Using the radiative relation between α and the neutrino spurions, the SI cross section can be expressed directly in terms of physical neutrino-sector parameters, schematically $\sigma_{\chi N}^{\text{SI}} \propto y_p^2 m_{H_p}^{-4} \sum_I y_{N,I}^2 m_{\nu,I}^2 M_{N,I}^4$, with no dependence on the singlet VEV. Numerically, for $M_{N,I} \sim 10\text{--}100$ TeV (the range favored by our thermal history and BBN), the resulting cross section lies many orders of magnitude below current direct-detection bounds and, in practice, well beneath the neutrino floor. Only for much larger heavy-neutrino masses does the loop-induced portal become sizable: for $M_{N,I} \sim \mathcal{O}(\text{PeV})$ the predicted $\sigma_{\chi N}^{\text{SI}}$ can approach the reach of next-generation xenon searches. Interestingly, this higher-scale regime can also overlap with the parameter space where standard thermal or resonant leptogenesis is viable, potentially linking the same heavy singlet sector to the origin of the baryon asymmetry [57–59].

Overall, our results show that a WIMP-mass thermal relic with $\mathcal{O}(0.1\text{--}1)$ hidden-sector couplings can be fully consistent with the null results of direct, indirect, and collider searches when the leading portal is radiative and aligned with neutrino-mass breaking. In this class of models the smallness of m_ν does not merely coexist with the absence of DM signals, but quantitatively enforces it: the portal is predictably tiny, the relic density is controlled by secluded annihilations, and direct-detection bounds only become relevant for extremely large heavy-neutrino scales. This provides a simple and UV-motivated paradigm in which the lack of standard WIMP signatures is a generic expectation rather than an unexplained coincidence.

ACKNOWLEDGMENTS

M. D. M. thanks N. Fornengo and S. Gariazzo for reading the paper and providing us with very helpful comments and suggestions. M. D. M. acknowledges support from the research grant Theoretical Astroparticle Physics (TAsP) funded by Istituto Nazionale di Fisica Nucleare (INFN) and from the Italian Ministry of University and Research (MUR), PRIN 2022 “EXSKALIBUR—Euclid-Cross-SKA: Likelihood Inference Building for Universe’s Research,” Grants No. 20222BBYB9, No. CUP I53D23000610 0006, and from the European Union—Next Generation EU.

DATA AVAILABILITY

The data are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: THERMAL AVERAGING AND EFFECTIVE RATES

In a radiation-dominated expanding Universe, departures of the number density of a species X from equilibrium relax back with an *effective* rate that sums all number-changing channels that create or destroy X , normalized to the equilibrium abundance. Linearizing the Boltzmann equation around n_X^{eq} ,

$$\begin{aligned} \dot{n}_X + 3Hn_X &= -(n_X - n_X^{\text{eq}})\Gamma_X^{\text{eff}}(T), \\ \Gamma_X^{\text{eff}}(T) &= \frac{\sum_i \gamma_i(T)}{n_X^{\text{eq}}(T)}, \end{aligned} \quad (\text{A1})$$

where the sum runs over all relevant processes i and γ_i are the corresponding reaction densities (rates per unit volume). Chemical equilibrium is maintained when $\Gamma_X^{\text{eff}} \gg H$. A practical tracking criterion that accounts for the explicit temperature dependence of n_X^{eq} is

$$\frac{\Gamma_X^{\text{eff}}(T)}{H(T)} \gg \frac{3}{2} + \xi_X, \quad \xi_X \equiv \frac{m_X}{T}. \quad (\text{A2})$$

This follows from linearizing the Boltzmann equation and demanding that the relaxation time is much shorter than the Hubble time, including the dilution term. However, in the paper we have just used the simplified condition $\frac{\Gamma_X^{\text{eff}}(T)}{H(T)} \gg 1$.

Reaction densities and reduced cross sections. For a generic $2 \rightarrow 2$ process $ab \rightarrow cd$, the reaction density is

$$\begin{aligned} \gamma_{ab \rightarrow cd}(T) &= \int d\Pi_a d\Pi_b d\Pi_c d\Pi_d (2\pi)^4 \\ &\times \delta^{(4)}(p_a + p_b - p_c - p_d) \\ &\times f_a^{\text{eq}} f_b^{\text{eq}} (1 \pm f_c^{\text{eq}})(1 \pm f_d^{\text{eq}}) |\mathcal{M}_{ab \rightarrow cd}|^2, \end{aligned} \quad (\text{A3})$$

with $d\Pi_i \equiv d^3 p_i / [(2\pi)^3 2E_i]$ and $f_{F,B}^{\text{eq}}(E) = 1/(e^{E/T} \pm 1)$

for fermions/bosons. Using standard manipulations and introducing the reduced cross section $\hat{\sigma} \equiv 2s\sigma$ (see, e.g., Refs. [59–62]), this can be written for (effectively) massless initial states as

$$\gamma_{ab \rightarrow cd}(T) = \frac{T}{64\pi^4} \int_{s_{\min}}^{\infty} ds \hat{\sigma}_{ab \rightarrow cd}(s) \sqrt{s} K_1\left(\frac{\sqrt{s}}{T}\right), \quad (\text{A4})$$

where K_1 is a modified Bessel function and $s_{\min}$ includes thermal masses when relevant. Quantum-statistical factors and thermal corrections are encoded in $\hat{\sigma}$ and in the effective threshold $s_{\min}$.

When Maxwell–Boltzmann factorization holds one may also estimate $\gamma_{ab \rightarrow cd} \simeq n_a^{\text{eq}} n_b^{\text{eq}} \langle \sigma v \rangle$; Eqs. (A3) and (A4) provide the exact thermal result used in our calculations.

1. Useful templates

(i) Pair annihilation $XX \leftrightarrow \text{SM}$. Near equilibrium,

$$\dot{n}_X + 3Hn_X = -\langle \sigma v \rangle (n_X^2 - n_X^{\text{eq}2}) \Rightarrow \Gamma_X^{\text{eff}} = 2\langle \sigma v \rangle n_X^{\text{eq}}. \quad (\text{A5})$$

(The factor 2 counts the two X destroyed per reaction.)

(ii) Semiannihilation $XX \leftrightarrow XY$.

$$\Gamma_X^{\text{eff}} = \langle \sigma v \rangle n_X^{\text{eq}}. \quad (\text{A6})$$

(iii) Coannihilation $Xa \leftrightarrow \text{SM}$ (partner a thermal).

$$\Gamma_X^{\text{eff}} \simeq \langle \sigma v \rangle_{Xa} n_a^{\text{eq}}. \quad (\text{A7})$$

(iv) Decay/inverse decay $X \leftrightarrow \text{SM}$.

$$\Gamma_X^{\text{eff}} = \Gamma_D \frac{K_1(\xi_X)}{K_2(\xi_X)}, \quad \xi_X \equiv \frac{m_X}{T}, \quad (\text{A8})$$

with $K_1/K_2 \simeq \xi_X/2$ for $\xi_X \ll 1$ and $K_1/K_2 \rightarrow 1$ for $\xi_X \gg 1$.

(v) $2 \rightarrow 2$ production of a single X : $ab \rightarrow X \dots$.

$$\Gamma_X^{\text{eff}} = \frac{\gamma_{ab \rightarrow X \dots}(T)}{n_X^{\text{eq}}(T)}, \quad (\text{A9})$$

with γ from Eq. (A4). This is the precise counterpart of the heuristic $\Gamma \sim \sum n^{\text{eq}} \langle \sigma v \rangle$.

(vi) Number-changing $3 \rightarrow 2$.

$$\begin{aligned} \dot{n}_X + 3Hn_X &= -\langle \sigma v^2 \rangle (n_X^3 - n_X^{\text{eq}3}) \\ &\Rightarrow \Gamma_X^{\text{eff}} = 3\langle \sigma v^2 \rangle (n_X^{\text{eq}})^2. \end{aligned} \quad (\text{A10})$$

(vii) Elastic (kinetic) scattering with bath species b .

$$\Gamma_{\text{kin}} \simeq n_b^{\text{eq}} \langle \sigma_T v \rangle, \quad (\text{A11})$$

relevant for temperature (rather than number) equilibration.

In natural units, $n^{\text{eq}} \sim \text{GeV}^3$ and $\gamma \sim \text{GeV}^4$, hence $\Gamma^{\text{eff}} = \gamma/n^{\text{eq}} \sim \text{GeV}$. Boltzmann and threshold effects are automatically encoded in Eqs. (A3) and (A4) and in $K_1(\xi_X)/K_2(\xi_X)$; as a result, Γ^{eff}/H drops rapidly once $\xi_X \gtrsim \mathcal{O}(1-3)$.

APPENDIX B: DETAILS OF $\Delta L = 1$ SCATTERINGS AT HIGH TEMPERATURE

As discussed in the main text, for temperatures larger than the lightest heavy-neutrino mass, $T \gtrsim M_N$, the SM bath is in thermal contact with the hidden sector through the neutrino portal. In this section, H denotes the SM Higgs doublet in the unbroken phase, not the CP -even mass eigenstate H of Appendix II.

The relevant interaction terms are

$$-\mathcal{L}_Y \supset Y_\nu^{aI} \bar{L}_a \tilde{H} N_I + \frac{1}{2} Y_N^{IJ} R \bar{N}_I^c N_J + y_p \bar{\chi} \chi R + y_t \bar{Q}_3 \tilde{H} t_R + \text{H.c.}, \quad (\text{B1})$$

$$\mathcal{L}_{\text{gauge}} \supset \sum_f \bar{f} \gamma^\mu (g_2 T^a W_\mu^a + g_Y Y B_\mu) f + (D_\mu H)^\dagger (D^\mu H) + \bar{\chi} i \gamma^\mu (\partial_\mu + i g_X q_\chi Z'_\mu) \chi, \quad (\text{B2})$$

with $\tilde{H} = i\sigma_2 H^*$, f running over SM fermions, T^a the $SU(2)_L$ generators, Y hypercharge, and D_μ the SM covariant derivative acting on H . The process Feynman diagrams are reported in Fig. 5. We list below the main processes keeping the hidden sector in thermal contact with the SM:

- (i) *Decays and inverse decays* (with violation of lepton number $\Delta L = 1$):

$$N_I \leftrightarrow L_\alpha H, \quad N_I \leftrightarrow \bar{L}_\alpha H^\dagger. \quad (\text{B3})$$

These control the production and destruction of heavy neutrinos N_I and are the leading neutrino-portal processes in the relativistic and mildly non-relativistic regimes.

- (ii) $\Delta L = 1$ *Yukawa-mediated scatterings* with one Y_ν and one SM coupling, e.g.,

$$N_I L_\alpha \leftrightarrow Q_3 t, \quad L_\alpha A \leftrightarrow N_I H, \quad (\text{B4})$$

and their crossed channels, mediated by Higgs or lepton exchange, where A is a SM electroweak gauge boson. These are $\mathcal{O}(Y_\nu^2 y_t^2)$ and $\mathcal{O}(Y_\nu^2 g^2)$ processes and efficiently contribute to keeping N_I in kinetic and chemical contact with the SM plasma for $T \gtrsim M_I$.

- (iii) $\Delta L = 2$ *scatterings* induced by virtual N_I , such as

$$L_\alpha H \leftrightarrow \bar{L}_\beta H^\dagger, \quad L_\alpha L_\beta \leftrightarrow HH, \quad (\text{B5})$$

which encode lepton-number violation and are relevant for washout and for the low-energy Weinberg operator once N_I are integrated out. These processes are relevant only at very high temperatures ($T \gtrsim 10^{12}$ GeV) and are negligible near (and below) $T \sim M_I$. We will not consider them further.

- (iv) *Portal processes involving the dark scalar and DM*, mediated by N_I and R , e.g.,

$$N_I N_J \leftrightarrow R, \quad N_I R \leftrightarrow L_\alpha H, \quad RR \leftrightarrow HH, \quad (\text{B6})$$

and, via the R - χ coupling, reactions that transfer energy and number between χ and the SM through the N_I - R chain. Once N_I are in equilibrium with the SM, these interactions help ensure that the dark sector is also thermally linked.

In the following we quantify the rates of decays/inverse decays, the leading $\Delta L = 1$ scatterings, and the portal processes involving the dark scalar and DM. We work in the unbroken phase and for $T \gg m_{\text{EW}}$, taking all external masses (except $M_{N,I}$) negligible. When several heavy Majorana eigenstates N_I are present, the slowest (last) to lose equilibrium is the *lightest* one. It is therefore sufficient (and conservative) to track a single state N with mass $M_N \equiv \min_I(M_{N,I})$. In the aligned limit, we denote by m_ν the light eigenvalue paired with this lightest N .

1. Tree-level width $N_I \rightarrow L_\alpha H$

We start from the neutrino Yukawa interaction for a given mass eigenstate N_I ,

$$\mathcal{L}_{\text{int}} \supset -y_{aI} \bar{L}_a \tilde{H} N_I + \text{H.c.}$$

Expanding in components, this contains decays of the form $N_I \rightarrow \ell_\alpha H$. Summing over lepton flavors yields the total width for decays and inverse decays

$$\Gamma_{D,I} \equiv \Gamma(N_I \rightarrow LH) + \Gamma(N_I \rightarrow \bar{L}H^\dagger) = \frac{M_{N,I}}{8\pi} \sum_\alpha |y_{aI}|^2 \mathcal{F}(r_H, r_{\ell_\alpha}), \quad (\text{B7})$$

with

$$\mathcal{F}(r_H, r_{\ell_\alpha}) = \sqrt{(1 - r_H - r_{\ell_\alpha})^2 - 4r_H r_{\ell_\alpha}} (1 + r_{\ell_\alpha} - r_H),$$

$$r_H \equiv \frac{m_H^2}{M_{N,I}^2}, \quad r_{\ell_\alpha} \equiv \frac{m_{\ell_\alpha}^2}{M_{N,I}^2}. \quad (\text{B8})$$

In the unbroken phase, or whenever $M_{N,I} \gg m_H, m_{\ell_\alpha}$, $\mathcal{F}(r_H, r_{\ell_\alpha}) \rightarrow 1$, so that




FIG. 5. Representative neutrino-portal processes maintaining thermal contact between the SM, the heavy neutrinos N_I , and the dark sector at $T \gg M_I$. (i) Decays and inverse decays $N_I \leftrightarrow L_\alpha H, \bar{L}_\alpha H^\dagger$, (ii) $\Delta L = 1$ scatterings: $N_I L_\alpha \leftrightarrow Q_3 t$ and $L_\alpha A \leftrightarrow N_I H$ (t-channel), (iii) $\Delta L = 2$ scatterings: $LH \leftrightarrow \bar{L}H^\dagger, LL \leftrightarrow HH$ via s-channel N_I , and (iv) Portal processes with R (and χ via R): $N_I N_J \leftrightarrow R, N_I R \leftrightarrow L_\alpha H, RR \leftrightarrow HH$.

$$\Gamma_{D,I} = \frac{M_{N,I}}{8\pi} (y^\dagger y)_{II}. \quad (\text{B9})$$

averaged decay rate per particle (see Appendix A) for the lightest of the heavy neutrinos N ,

This is the standard expression used in leptogenesis analyses [57–59].

In the early Universe the heavy neutrinos N_I are immersed in a thermal bath and are produced with a momentum distribution rather than at rest. The relevant quantity for equilibration is the *time-dilated*, thermally

$$\langle \Gamma_D \rangle_T = \Gamma_D \frac{K_1(\xi)}{K_2(\xi)}, \quad \xi \equiv \frac{M_N}{T}, \quad (\text{B10})$$

where $K_{1,2}$ are modified Bessel functions of the second kind and we assume Maxwell-Boltzmann statistics for N_I . To assess thermal contact we compare $\langle \Gamma_D \rangle_T$ with the

Hubble rate during radiation domination,

$$H(T) = 1.66\sqrt{g_*(T)} \frac{T^2}{M_{\text{Pl}}} = 1.66\sqrt{g_*} \frac{M_N^2}{M_{\text{Pl}} \xi^2}, \quad (\text{B11})$$

where g_* is the effective number of relativistic degrees of freedom (d.o.f.). Using Eqs. (B9) and (B10) and in the aligned Casas-Ibarra limit $((y^\dagger y)_{II} = 2m_{\nu,I} M_I / v_h^2)$ the ratio $\langle \Gamma_{D,I} \rangle_T / H(T)$ simplifies to

$$\frac{\langle \Gamma_D \rangle_T}{H(T)} = \frac{m_\nu M_{\text{Pl}}}{4\pi 1.66\sqrt{g_*} v_h^2} \xi^2 \frac{K_1(\xi)}{K_2(\xi)}. \quad (\text{B12})$$

We define the equilibration (or decoupling) temperature T_* by

$$\left. \frac{\langle \Gamma_D \rangle_T}{H(T)} \right|_{T=T_*} = 1. \quad (\text{B13})$$

For $T_* \gg M_N$ (i.e., $\xi_* \ll 1$), we can use $K_1(\xi)/K_2(\xi) \simeq \xi/2$, so Eq. (B12) yields

$$\frac{T_*}{M_N} \simeq 2.8 \left(\frac{m_\nu}{0.05 \text{ eV}} \right)^{1/3} \left(\frac{106.75}{g_*} \right)^{1/6}, \quad (\text{B14})$$

so that N enter equilibrium as the Universe cools to $T \sim \text{few} \times M_N$, and are well in equilibrium for $T \lesssim \text{few} \times M_N$. Evaluating Eq. (B12) at $T = M_N$ ($\xi = 1$) gives

$$\left. \frac{\langle \Gamma_D \rangle_T}{H(T)} \right|_{T=M_N} = \frac{m_\nu M_{\text{Pl}}}{4\pi 1.66\sqrt{g_*} v_h^2} \frac{K_1(1)}{K_2(1)} \sim \mathcal{O}(20 - 30) \quad \text{for } m_\nu \sim 0.05 \text{ eV}, \quad (\text{B15})$$

confirming that decays and inverse decays alone are fast enough to maintain thermal contact around and below $T \sim \text{few} \times M_N$ for seesaw-motivated parameters.

2. Top-assisted scattering $N_I L_\alpha \rightarrow Q_3 t$

We now consider the process $N_I(p_1) + L_\alpha(p_2) \rightarrow Q_3(p_3) + t(p_4)$ mediated by t -channel Higgs exchange of the interaction term $-\mathcal{L} \supset y_{\alpha I} \bar{L}_\alpha \tilde{H} N_I + y_t \bar{Q}_3 H t_R + \text{H.c.}$ In the relativistic regime $s \gg M_I^2, m_H^2(T)$, after spin/color sums and angular integration one finds a reduced cross section of the form

$$\begin{aligned} \hat{\sigma}_{N_I L \rightarrow Q_3 t}(s) &\equiv 2s \sigma_{N_I L \rightarrow Q_3 t}(s) \\ &= \frac{3}{8\pi} (Y_\nu^\dagger Y_\nu)_{II} |y_t|^2 \mathcal{F}_t \left(\frac{s}{m_H^2(T)} \right), \end{aligned} \quad (\text{B16})$$

with

$$\mathcal{F}_t(x) \equiv \left(1 + \frac{2}{x} \right) \ln(1+x) - 2, \quad x \equiv \frac{s}{m_H^2(T)}. \quad (\text{B17})$$

With hard thermal loop (HTL) thermal masses² $m_H^2(T) \sim c_H g^2 T^2$ and typical $s \sim (2T)^2$ one has $x = \mathcal{O}(5 - 20)$, so numerically $\mathcal{F}_t(x) \sim 0.3 - 1.0$. A fully resummed derivation (including HTL propagators and exact $\mathcal{F}_t$) can be found in [60–62].

3. Gauge-assisted scattering $L_\alpha A \rightarrow N_I H$

Now we focus on the process $L_\alpha(p_1) + A(p_2) \rightarrow N_I(p_3) + H(p_4)$ where A is an electroweak gauge boson (W_α or B). These channels involve one neutrino Yukawa and one gauge vertex and are representative of $\mathcal{O}(Y_\nu^2 g^2)$ scatterings. In the unbroken phase the relevant interactions are

$$-\mathcal{L}_Y \supset Y_\nu^{aI} \bar{L}_\alpha \tilde{H} N_I + \text{H.c.}, \quad (\text{B18})$$

$$\mathcal{L}_{\text{gauge}} \supset g \bar{L}_\alpha \gamma^\mu T^a L_\alpha A_\mu + (D_\mu H)^\dagger (D^\mu H), \quad (\text{B19})$$

with analogous hypercharge couplings. A representative contribution is the t -channel lepton exchange; s - and u -channel diagrams complete the gauge-invariant set and yield the same parametric dependence once summed.

Summing over spins, polarizations and gauge indices, and working in the relativistic regime $s \gg M_I^2, m_{L,H,A}^2(T)$, one obtains the reduced cross section

$$\begin{aligned} \hat{\sigma}_{LA \rightarrow N_I H}(s) &\equiv 2s \sigma_{LA \rightarrow N_I H}(s) \\ &= \frac{C_A}{8\pi} |Y_\nu^{aI}|^2 g^2 \mathcal{F}_g \left(\frac{s}{m_{\text{th}}^2(T)} \right), \end{aligned} \quad (\text{B20})$$

where: $C_A = \mathcal{O}(1)$ is a group-theory factor, $m_{\text{th}}^2(T)$ is a representative thermal scale built from $m_L^2(T)$, $m_H^2(T)$, $m_A^2(T)$, and $\mathcal{F}_g$ is a dimensionless function encoding angular structure and IR regularization. A convenient parametrization with the correct limits is

$$\mathcal{F}_g(x) \equiv \left(1 + \frac{1}{x} \right) \ln(1+x) - 1, \quad x \equiv \frac{s}{m_{\text{th}}^2(T)}. \quad (\text{B21})$$

For $m_{\text{th}}^2(T) \sim c g^2 T^2$ and $s \sim (2T)^2$ one typically has $x = \mathcal{O}(5 - 30)$ and hence $\mathcal{F}_g(x) \sim 0.5 - 2$. Summing over flavors,

$$\hat{\sigma}_{LA \rightarrow N_I H}(s) = \frac{C_A}{8\pi} (Y_\nu^\dagger Y_\nu)_{II} g^2 \mathcal{F}_g \left(\frac{s}{m_{\text{th}}^2(T)} \right). \quad (\text{B22})$$

²HTL masses are temperature-dependent effective masses that particles acquire from interactions in a hot plasma, computed within the hard thermal loop approximation [63].

Detailed HTL-resummed expressions and precise values of $\mathcal{F}_g$ for the individual gauge channels can be found in [60–62].

4. Per-particle rates and comparison with Hubble rate for the top- and gauge-assisted scatterings

For equilibration we require the *per-particle* rate, normalized to one heavy neutrino N_I ,

$$\Gamma_{\text{scatt}}^{(\Delta L=1)}(T) \equiv \frac{\gamma_{\Delta L=1}(T)}{n_{N_I}^{\text{eq}}(T)}, \quad (\text{B23})$$

where $\gamma_{\Delta L=1}$ is the sum of the relevant $\Delta L = 1$ reaction densities involving N_I . In the relativistic limit ($T \gg M_I$), the equilibrium number density of a Majorana fermion ($g_N = 2$) is

$$n_{N_I}^{\text{eq}}(T) = \frac{3\zeta(3)}{2\pi^2} T^3. \quad (\text{B24})$$

For a slowly varying reduced cross section $\hat{\sigma}(s) \approx \hat{\sigma}_0$, Eq. (A4) gives

$$\gamma(T) \simeq \frac{\hat{\sigma}_0 T^4}{16\pi^4}, \quad (\text{B25})$$

using $\int_0^\infty dx x^2 K_1(x) = 2$. Thus

$$\Gamma_{\text{scatt}}(T) = \frac{\gamma(T)}{n_{N_I}^{\text{eq}}(T)} \simeq \frac{\hat{\sigma}_0}{24\zeta(3)\pi^2} T. \quad (\text{B26})$$

The scattering rate for top-assisted scatterings can be determined using Eq. (B16),

$$\begin{aligned} \Gamma_{\text{top},I}^{(\Delta L=1)}(T) &\simeq \frac{3\bar{\mathcal{F}}_t}{192\zeta(3)\pi^3} (Y_\nu^\dagger Y_\nu)_{II} y_t^2 T \\ &\equiv c_t (Y_\nu^\dagger Y_\nu)_{II} y_t^2 T, \end{aligned} \quad (\text{B27})$$

with $c_t = 3\bar{\mathcal{F}}_t/(192\zeta(3)\pi^3) \approx (2-4) \times 10^{-4}$, which parametrizes the theoretical uncertainties associated with thermal masses, IR regularization, and phase-space approximations. Instead, for gauge-assisted processes and using Eq. (B22), we find

$$\begin{aligned} \Gamma_{\text{gauge},I}^{(\Delta L=1)}(T) &\simeq \frac{C_A \bar{\mathcal{F}}_g}{192\zeta(3)\pi^3} (Y_\nu^\dagger Y_\nu)_{II} g^2 T \\ &\equiv c_g (Y_\nu^\dagger Y_\nu)_{II} g^2 T, \end{aligned} \quad (\text{B28})$$

with $c_g = C_A \bar{\mathcal{F}}_g/(192\zeta(3)\pi^3) \approx (1-3) \times 10^{-4}$.

Summing the dominant channels,

$$\begin{aligned} \Gamma_{\text{scatt}}^{(\Delta L=1)}(T) &\equiv \Gamma_{\text{top},I}^{(\Delta L=1)} + \Gamma_{\text{gauge},I}^{(\Delta L=1)} \\ &\simeq c_1 (Y_\nu^\dagger Y_\nu)_{II} (y_t^2 + g^2) T, \end{aligned} \quad (\text{B29})$$

with an effective coefficient

$$c_1 \approx (3-7) \times 10^{-4}, \quad (\text{B30})$$

once thermal masses, angular dependence and channel multiplicities are taken into account, in agreement with detailed numerical studies [60–62]. The variation in the parameters c_t , c_g and consequently c_1 parametrizes the theoretical uncertainty from thermal masses, IR regularization, and phase-space approximations in the $\Delta L = 1$ scatterings.

In the aligned limit $\Gamma_{\text{scatt}}^{(\Delta L=1)}(T)$ can be rewritten as

$$\Gamma_{\text{scatt}}^{(\Delta L=1)}(T) \simeq 2c_1 \frac{m_{\nu,I} M_I}{v_h^2} (y_t^2 + g^2) T. \quad (\text{B31})$$

Therefore, the ratio between the scattering and the Hubble rate is

$$\frac{\Gamma_{\text{scatt}}^{(\Delta L=1)}(T)}{H(T)} \simeq \frac{2c_1}{1.66\sqrt{g_*}} \frac{m_{\nu,I} M_{\text{Pl}}}{v_h^2} (y_t^2 + g^2) \frac{M_I}{T}. \quad (\text{B32})$$

For $m_{\nu,I} \sim 0.05$ eV, $g_* \sim 10^2$, $y_t^2 + g^2 = \mathcal{O}(1)$ and $c_1 \sim 5 \times 10^{-4}$ one obtains for $T \sim M_I$

$$\frac{\Gamma_{\text{scatt}}^{(\Delta L=1)}}{H} \gtrsim \mathcal{O}(1). \quad (\text{B33})$$

5. Portal processes involving R and χ

(a) $RR \leftrightarrow HH$. The cross section of the process $RR \leftrightarrow HH$ depends on κ_{loop} . In fact, for $T \gg m_H, m_R$ it scales as

$$\sigma(RR \rightarrow HH) \sim \frac{\kappa_{\text{loop}}^2}{16\pi s}, \quad s \sim T^2. \quad (\text{B34})$$

The per-particle rate is

$$\Gamma_{RR \rightarrow HH} \sim n_R^{\text{eq}} \sigma \sim \kappa_{\text{loop}}^2 T, \quad (\text{B35})$$

up to $\mathcal{O}(10^{-1})$ factors. For the values relevant in our model, $\kappa_{\text{loop}} \sim 10^{-11} - 10^{-10}$, hence $\Gamma_{RR \rightarrow HH}/T \sim 10^{-20}$ and $\Gamma_{RR \rightarrow HH} \ll H(T)$ at all temperatures of interest. Thus $RR \leftrightarrow HH$ is completely negligible for thermal contact.

(b) $N_I R \rightarrow LH$. The process $N_I R \rightarrow L_\alpha H$ proceeds through the combination of the Majorana coupling $R N_I N_I$ and the Dirac Yukawa $L_\alpha H N_I$. In the mass

basis (and in the limit where R couples diagonally to N_I), the relevant interactions are

$$-\mathcal{L} \supset \frac{1}{2} y_{N,I} R \overline{N_I^c} N_I + y_{aI} \overline{L_a} \tilde{H} N_I + \text{H.c.}, \quad (\text{B36})$$

so that $N_I R \rightarrow L_a H$ proceeds dominantly via s -channel N_I exchange. In the unbroken phase and for $T \gg m_{\text{EW}}$, taking all external masses (except M_I) negligible, the reduced cross section $\hat{\sigma}(s) \equiv 2s\sigma(s)$ is

$$\hat{\sigma}_{N_I R \rightarrow LH}(s) = \frac{1}{8\pi} |y_{N,I}|^2 (Y_\nu^\dagger Y_\nu)_{II} \mathcal{F}_{NR} \left(\frac{s}{M_I^2} \right), \quad (\text{B37})$$

where $\mathcal{F}_{NR}$ is a dimensionless kinematic function that encodes the exact angular dependence and the mild s/M_I^2 dependence. For $s \gtrsim (2T)^2 \gg M_I^2$ one has

$$\mathcal{F}_{NR}(x) \xrightarrow{x \gg 1} 1 \quad \text{up to } \mathcal{O}\left(\frac{M_I^2}{s}\right) \text{ corrections.} \quad (\text{B38})$$

We therefore take an effective thermal average $\overline{\mathcal{F}}_{NR} \sim 0.5\text{--}2$ for $T \sim \text{few} \times M_I$.

Using Eq. (A4) the per-particle rate is

$$\begin{aligned} \Gamma_{N_I R \rightarrow LH}(T) &\equiv \frac{\gamma_{N_I R \rightarrow LH}}{n_{N_I}^{\text{eq}}} \simeq \frac{\hat{\sigma}_0}{24\zeta(3)\pi^2} T \\ &= c_{NR} |y_{N,I}|^2 (Y_\nu^\dagger Y_\nu)_{II} T, \end{aligned} \quad (\text{B39})$$

with

$$c_{NR} = \frac{\overline{\mathcal{F}}_{NR}}{192\zeta(3)\pi^3} \approx (0.7 - 2.8) \times 10^{-4}. \quad (\text{B40})$$

This is directly analogous to the c_t and c_g factors obtained for the top- and gauge-assisted $\Delta L = 1$ scatterings.

In the aligned Casas-Ibarra limit one finds

$$\Gamma_{N_I R \rightarrow LH}(T) \simeq 2c_{NR} \frac{y_{N,I}^2 m_{\nu,I} M_I}{v_h^2} T. \quad (\text{B41})$$

For illustration, take $m_{\nu,I} = 0.05$ eV, $v_h = 246$ GeV, $g_* = 106.75$, and $c_{NR} \sim 10^{-4}$. Using Eq. (B41), the rate at $T \simeq M_I$ is

$$\frac{\Gamma_{N_I R \rightarrow LH}}{H} \simeq 0.12 \left(\frac{c_{NR}}{10^{-4}} \right) \left(\frac{y_{N,I}}{1} \right)^2 \left(\frac{m_{\nu,I}}{0.05 \text{ eV}} \right) \left(\frac{106.75}{g_*} \right)^{1/2}, \quad (\text{B42})$$

independent of M_I at $T = M_I$. Thus, for $y_{N,I} \sim 1$ this channel contributes at the level $\Gamma/H \sim 0.1$ near $T \simeq M_I$, while larger Yukawas $y_{N,I} \gtrsim 3$ would make it fully efficient ($\Gamma/H \gtrsim 1$). For fixed M_I , the ratio scales as $\Gamma/H \propto M_I/T$

at higher temperatures, decreasing slowly for $T \gg M_I$, analogously to the standard $\Delta L = 1$ scatterings.

a. Summary

Process hierarchy and temperature trends. For seesaw-motivated parameters [$m_{\nu,I} \sim 0.05$ eV, $y_{N,I} = \mathcal{O}(1)$, $M_I \gtrsim \text{few TeV}$, $g_* \simeq 10^2$]:

(i) Decays/inverse decays $N_I \leftrightarrow LH$:

$$\frac{\langle \Gamma_{D,I} \rangle_T}{H} = \frac{m_{\nu,I} M_{\text{Pl}}}{4\pi 1.66 \sqrt{g_*} v_h^2} \xi^2 \frac{K_1(\xi)}{K_2(\xi)}.$$

For $\xi \ll 1$ (very high T), this scales as $\propto \xi^3$ and is small; as the Universe cools to $\xi \sim \mathcal{O}(1)$ ($T \sim M_I$) the ratio grows to $\mathcal{O}(10\text{--}30)$ and *dominates* equilibration.

(ii) $\Delta L = 1$ scatterings (top- and gauge-assisted):

$$\frac{\Gamma_{\text{scatt}}^{(\Delta L=1)}}{H} \simeq \frac{2c_1}{1.66 \sqrt{g_*}} \frac{m_{\nu,I} M_{\text{Pl}}}{v_h^2} (y_t^2 + g^2) \frac{M_I}{T} \propto \frac{M_I}{T}.$$

They are typically $\mathcal{O}(1)$ at $T \sim M_I$ and provide efficient additional channels; for $T \gg M_I$ the ratio decreases slowly as M_I/T .

(iii) Portal scattering $N_I R \leftrightarrow L_a H$: at $T = M_I$

$$\begin{aligned} \left. \frac{\Gamma_{N_I R \rightarrow LH}}{H} \right|_{T=M_I} &\simeq 0.1 - 0.2 \left(\frac{c_{NR}}{10^{-4}} \right) \left(\frac{m_{\nu,I}}{0.05 \text{ eV}} \right) \\ &\times \left(\frac{y_{N,I}}{1} \right)^2, \end{aligned} \quad (\text{B43})$$

where the numerical range reflects the mild dependence on g_* . For $y_{N,I} = \mathcal{O}(1)$ and $m_{\nu,I}$ in the experimentally favored range, the portal scattering is *subdominant* compared to decays and the standard $\Delta L = 1$ scatterings, but still non-negligible. For larger Yukawas $y_{N,I} \gtrsim 3$ it can become fully efficient ($\Gamma/H \gtrsim 1$), further assisting SM–dark-sector thermal contact.

(iv) Loop-induced scalar portal $RR \leftrightarrow HH$: $\Gamma/H \ll 1$ throughout the parameter space of interest and can be safely neglected for thermal contact.

Overall: As the Universe cools down to $T \sim \text{few} \times M_I$, N_I decays/inverse decays become the *dominant* equilibration channel. $\Delta L = 1$ scatterings are efficient and comparable to decays around $T \sim M_I$, while $N_I R \leftrightarrow LH$ provides an additional but subleading portal for $y_{N,I} \sim 1$, becoming important only for larger $y_{N,I}$. At $T \ll M_I$ all N_I -mediated reactions decouple due to Boltzmann suppression of N_I .

APPENDIX C: DECOUPLING AND FREEZE-OUT WITH TWO TEMPERATURES ($T' \neq T$)

This appendix collects the ingredients used in the main text to describe the decoupling of the neutrino portal and the subsequent dark-sector freeze-out in the regime where the SM bath and the hidden sector evolve with different temperatures $T' \neq T$. We adopt the notation of Appendix B: M_I denotes the heavy-neutrino masses, $\zeta \equiv T'/T$ is the dark-to-visible temperature ratio, $g_{*S}^{\text{vis}}(T)$ and $g_{*S}^{\text{vis}}(T)$ are the usual visible-sector energy and entropy relativistic d.o.f., with analogous definitions for the dark sector.

1. Decoupling around $T \sim M_I$ and entropy bookkeeping

As the temperature drops below the mass of the lightest heavy neutrino $M_{N,I}$, the heavy-neutrino abundance becomes Boltzmann suppressed and portal-mediated energy exchange shuts off. From that moment on, the comoving entropies of the two sectors are separately conserved,

$$s_{\text{vis}}(T)a^3 = \text{const}, \quad s_{\text{dark}}(T', \{\mu_i(T')\})a^3 = \text{const}, \quad (\text{C1})$$

where μ_i is the chemical potential for the particle species i . The temperature ratio at later times is fixed by separate entropy conservation after decoupling. Assuming full chemical equilibrium in the dark sector at the decoupling temperature $T'_{\text{dec}} [\mu_i(T'_{\text{dec}}) = 0]$, the general expression is (see, e.g., Ref. [56])

$$\zeta(T) \equiv \frac{T'}{T} = \left[\frac{s_{\text{vis}}(T)}{s_{\text{vis}}(T_{\text{dec}})} \right]^{1/3} \left[\frac{s_{\text{dark}}(T'_{\text{dec}}, \{\mu_i = 0\})}{s_{\text{dark}}(T' = \zeta T, \{\mu_i(T')\})} \right]^{1/3}. \quad (\text{C2})$$

If the following two physical conditions apply: (i) the relevant species in each sector behave as a radiation bath so that $s = \frac{2\pi^2}{45} g_{*S} T^3$, and (ii) *number-changing* reactions are fast enough to enforce $\mu_i = 0$, Eq. (C2) reduces to the familiar result (see, e.g., Ref. [56])

$$\zeta(T) = \left[\frac{g_{*S}^{\text{vis}}(T)}{g_{*S}^{\text{vis}}(T_{\text{dec}})} \right]^{1/3} \left[\frac{g_{*S}^{\text{dark}}(T'_{\text{dec}})}{g_{*S}^{\text{dark}}(T')} \right]^{1/3}. \quad (\text{C3})$$

In typical benchmarks with $m_\chi \sim \mathcal{O}(\text{GeV})$ and $M_I \sim \mathcal{O}(\text{TeV})$, the ensuing ζ lies in the ballpark $\zeta \simeq 1$ –2 (the precise value follows from the g_{*S} evolution in the two sectors).

2. Temperature ratio $\zeta \equiv T'/T$: General expression and benchmarks

In our secluded setup the late-time dark bath is dominated by the real scalar H_p (mass m_{H_p}) and, depending on the benchmark, possibly by a light dark gauge boson Z' . The Dirac DM χ (mass m_χ) is Boltzmann suppressed at

TABLE I. Representative values of $\zeta \equiv T'/T$ at $T \simeq m_\chi/20$, obtained by solving Eq. (C3) self-consistently with $T' = \zeta T$, $g_{*S}^{\text{vis}}(T)/g_{*S}^{\text{vis}}(T_{\text{dec}}) = 86.25/106.75$, $g_{*S}^{\text{dark}}(T'_{\text{dec}}) = 4.5$ (Dirac χ + real scalar H_p), and exact scalar entropy for H_p (no light Z' at late times, $\mu_i = 0$).

m_{H_p} (GeV)	ζ for m_χ (GeV)				
	10	30	100	300	1000
1	1.296	1.238	1.236	1.236	1.236
3	1.601	1.296	1.238	1.236	1.236
10	1.872	1.639	1.296	1.238	1.236
30	1.894	1.872	1.601	1.296	1.238
100	1.895	1.894	1.872	1.639	1.296

$T' \sim m_\chi/20$ and can be neglected in $g_{*S}^{\text{dark}}(T')$ unless explicitly stated.

In order to estimate the value of ζ at the epoch relevant for DM freeze-out, $T \simeq m_\chi/20$, we use the following inputs:

- (i) Visible sector: $g_{*S}^{\text{vis}}(T_{\text{dec}}) = 106.75$, $g_{*S}^{\text{vis}}(T \simeq m_\chi/20) = 86.25$.
- (ii) Dark sector at decoupling: $g_{*S}^{\text{dark}}(T'_{\text{dec}}) = 4.5$ (Dirac χ + real scalar H_p). If a dark gauge boson Z' is relativistic at decoupling, use $g_{*S}^{\text{dark}}(T'_{\text{dec}}) = 7.5$.
- (iii) Chemical potentials: we assume efficient number changing in the dark bath ($\mu_i = 0$) so that Eq. (C3) applies directly. After dark chemical freeze-out ($\mu_i > 0$) one must use the general form (C2), which typically yields a *smaller* ζ at fixed parameters.

Table I summarizes representative values of ζ as a function of m_{H_p} and m_χ in the minimal case (no light Z' at late times), assuming $\mu_i = 0$. The qualitative trends are: (i) for fixed m_χ , the heavier m_{H_p} , the larger ζ ; (ii) for fixed m_{H_p} , the larger m_χ , the smaller ζ . Across the broad range relevant here, ζ typically lies in the 1.1–1.9 band.

3. Dark freeze-out with $T' = \zeta T$

After visible-dark decoupling, the portal remains too feeble to reequilibrate the sectors. The dark bath stays internally thermalized at temperature $T' = \zeta T$ through fast secluded reactions. For the Dirac dark matter particle χ (mass m_χ) annihilating into a lighter real scalar H_p ($m_{H_p} < m_\chi$) via the Yukawa interaction $y_p \bar{\chi} \chi H_p$, the dominant number-changing channel is

$$\chi \bar{\chi} \rightarrow H_p H_p. \quad (\text{C4})$$

We work with the visible-entropy yield $Y_\chi \equiv n_\chi/s_{\text{vis}}$, using $s_{\text{vis}}(T) = \frac{2\pi^2}{45} g_{*S}^{\text{vis}}(T) T^3$, and define

$$x \equiv \frac{m_\chi}{T}, \quad x' \equiv \frac{m_\chi}{T'} = \frac{x}{\zeta}. \quad (\text{C5})$$

During radiation domination, the Hubble rate receives contributions from the radiation in both sectors,

$$\begin{aligned} H(T) &= \sqrt{\frac{8\pi G}{3}(\rho_{\text{vis}}(T) + \rho_{\text{dark}}(T'))} \\ &= \frac{\pi}{\sqrt{90} M_{\text{Pl}}} T^2 [g_*^{\text{vis}}(T) + g_*^{\text{dark}}(T') \zeta^4]^{1/2} \\ &\equiv \frac{1.66 \sqrt{g_*^H(T, T')}}{M_{\text{Pl}}} T^2, \end{aligned} \quad (\text{C6})$$

where

$$g_*^H(T, T') \equiv g_*^{\text{vis}}(T) + \zeta^4 g_*^{\text{dark}}(T') \quad (\text{C7})$$

is the energy d.o.f. combination entering H .

The Boltzmann equation for Y_χ reads

$$\frac{dY_\chi}{dx} = -\frac{s_{\text{vis}}(T)}{xH(T)} \langle \sigma v \rangle_{\chi\bar{\chi} \rightarrow H_p H_p}(x') [Y_\chi^2 - (Y_\chi^{\text{eq}}(x'))^2], \quad (\text{C8})$$

with $Y_\chi^{\text{eq}}(x') = n_\chi^{\text{eq}}(T')/s_{\text{vis}}(T)$ and where the annihilation rate must be evaluated at the *dark* temperature T' .

The thermal average cross section for the process $\chi\bar{\chi} \rightarrow H_p H_p$ that dominates the relic density is

$$\langle \sigma v \rangle_{\chi\bar{\chi} \rightarrow H_p H_p}(x') = \frac{6b(r)}{x'} = \frac{y_p^4}{64\pi m_\chi^2} \mathcal{S}_p(r) \frac{6}{x'}. \quad (\text{C9})$$

For a Majorana χ the result is multiplied by 1/2 due to identical initial states. The kinematic limits are $\mathcal{S}_p(0) = 1$ (light mediator) and $\mathcal{S}_p(r) \propto (1-r)^{1/2}$ near threshold ($r \rightarrow 1$).

The freeze-out condition equates the *dark* annihilation rate to the *visible* Hubble rate,

$$n_\chi^{\text{eq}}(T') \langle \sigma v \rangle(T') \simeq H(T), \quad (\text{C10})$$

with $n_\chi^{\text{eq}}(T') \simeq g_\chi \left(\frac{m_\chi T'}{2\pi}\right)^{3/2} e^{-x'}$ in the nonrelativistic regime. Using Eq. (C9) one obtains the standard logarithmic solution

$$x'_f \simeq \ln(\mathcal{A} y_p^4) - \frac{1}{2} \ln[\ln(\mathcal{A} y_p^4)], \quad (\text{C11})$$

with

$$\begin{aligned} \mathcal{A} &\equiv \frac{c_f g_\chi m_\chi M_{\text{Pl}}}{1.66(2\pi)^{3/2}} \frac{\zeta^2}{\sqrt{g_*^H(T_f, T'_f)}} \frac{6\mathcal{S}_p(r)}{64\pi m_\chi^2}, \\ T_f &\equiv \frac{m_\chi}{x_f}, \quad T'_f \equiv \frac{m_\chi}{x'_f}. \end{aligned} \quad (\text{C12})$$

Here $c_f = \mathcal{O}(1)$ encodes the usual matching convention. For $m_\chi \sim \text{GeV}$ and $y_p \sim 0.1-1$ one finds $x'_f \in [20, 30]$, as in the standard p -wave freeze-out.

The solution of Eq. (C8) can be written in terms of the annihilation integral in the two-temperature case,

$$J_{2T} \equiv \int_{x_f}^{\infty} \frac{\langle \sigma v \rangle(x')}{x'^2} dx = \frac{1}{\zeta} \left[\frac{a}{x'_f} + \frac{3b}{x'^2_f} \right]. \quad (\text{C13})$$

Using for the annihilation cross section Eq. (C9) (pure p -wave, $a = 0$) we find

$$J_{2T} = \frac{1}{\zeta} \frac{3y_p^4}{64\pi m_\chi^2} \frac{\mathcal{S}_p(r)}{x'^2_f}. \quad (\text{C14})$$

Defining

$$\lambda_{\text{eff}} \equiv 0.264 \frac{g_{*S}^{\text{vis}}(T_f)}{\sqrt{g_*^H(T_f, T'_f)}} M_{\text{Pl}} m_\chi, \quad (\text{C15})$$

the asymptotic yield and relic density are

$$\begin{aligned} Y_\infty &= \frac{1}{\lambda_{\text{eff}} J_{2T}}, \\ \Omega_\chi h^2 &\simeq (1.05 \times 10^9 \text{ GeV}^{-1}) \frac{\sqrt{g_*^H(T_f, T'_f)}}{g_{*S}^{\text{vis}}(T_f)} \frac{1}{M_{\text{Pl}} J_{2T}}. \end{aligned} \quad (\text{C16})$$

Specializing to p -wave only and inserting Eq. (C14),

$$\Omega_\chi h^2 \simeq (1.05 \times 10^9 \text{ GeV}^{-1}) \frac{\sqrt{g_*^H(T_f, T'_f)}}{g_{*S}^{\text{vis}}(T_f)} \frac{\zeta x'^2_f}{M_{\text{Pl}} 3y_p^4 \mathcal{S}_p(r)} \frac{64\pi m_\chi^2}{\mathcal{S}_p(r)}. \quad (\text{C17})$$

Inverting Eq. (C17) for the coupling,

$$y_p = \left[(1.05 \times 10^9 \text{ GeV}^{-1}) \frac{\sqrt{g_*^H(T_f, T'_f)}}{g_{*S}^{\text{vis}}(T_f)} \frac{\zeta x'^2_f}{M_{\text{Pl}} 3\Omega_\chi h^2 \mathcal{S}_p(r)} \frac{64\pi m_\chi^2}{\mathcal{S}_p(r)} \right]^{1/4}. \quad (\text{C18})$$

A useful numerical estimate is

$$y_p \simeq 0.43 \left(\frac{x'_f}{25} \right)^{1/2} \left(\frac{m_\chi}{100 \text{ GeV}} \right)^{1/2} \left(\frac{g_{*S}^{\text{vis}}}{86.25} \right)^{-1/4} \left(\frac{g_*^H}{100} \right)^{1/8}. \quad (\text{C19})$$

Therefore, a value of $y_p \sim \mathcal{O}(0.1-1)$ is needed to reach the observed DM relic density. We show in Fig. 6 the values of y_p for which DM achieves the observed relic density. Equivalently, a handy heuristic for the *required* thermal rate is




FIG. 6. Values of y_p for which DM achieves the observed relic density.

$$\langle\sigma v\rangle_{\text{req}} \approx \frac{2 \times 10^{-26} \text{ cm}^3 \text{ s}^{-1}}{\zeta} \left[1 + \mathcal{O}\left(\frac{g_*^{\text{dark}}}{g_*^{\text{vis}}} \zeta^4\right) \right], \quad (\text{C20})$$

i.e., a hotter (colder) dark bath with $\zeta > 1$ ($\zeta < 1$) requires a *smaller (larger)* annihilation cross section to reproduce the observed $\Omega_{\text{DM}} h^2$.

APPENDIX D: LOOP-INDUCED ODD OPERATORS: $\eta_1 R(\Phi^\dagger \Phi)$ AND $\mu_3 R^3$

For a real gauge-singlet scalar R , the most general renormalizable scalar potential contains, in addition to the even terms, the odd operators

$$V_{\text{odd}}(\Phi, R) = \mu_1 R + \frac{\mu_3}{3!} R^3 + \frac{\eta_1}{2} R(\Phi^\dagger \Phi), \quad (\text{D1})$$

where Φ is the SM Higgs doublet. In this section we show that, in the absence of an $R \rightarrow -R$ selection rule, the heavy-neutrino sector generically induces one-loop threshold corrections to η_1 and μ_3 . We present (i) the relevant one-loop topologies, (ii) compact expressions for the induced couplings, and (iii) parametric/numerical estimates in the regime $M_N \sim \mathcal{O}(10 - 50)$ TeV, $m_\nu \sim 0.05$ eV, and $y_N \sim \mathcal{O}(1)$.

1. One-loop induced effective couplings η_1 and μ_3

In the absence of an $R \rightarrow -R$ selection rule, the heavy-neutrino sector generically induces at one loop the odd operators in Eq. (D1). The corresponding topologies are shown in Figs. 7 and 8: (i) a fermion triangle with three RNN insertions generating the cubic $\mu_3 R^3$, and (ii) a mixed loop with two Higgs insertions from the Y_ν vertices and one RNN insertion from Y_N generating the linear portal $R(\Phi^\dagger \Phi)$.

The relevant interactions (schematically, suppressing flavor indices) are




FIG. 7. One-loop topology generating the linear portal $R(\Phi^\dagger \Phi)$, corresponding to a threshold correction to η_1 . Two Yukawa insertions Y_ν and one insertion Y_N are required.




FIG. 8. One-loop heavy-neutrino triangle generating the cubic R^3 , i.e., a threshold correction to μ_3 (three RNN insertions).

$$L \supset -\bar{L} \tilde{\Phi} Y_\nu N - \frac{1}{2} \bar{N}^c Y_N N R + \text{H.c.} \quad (\text{D2})$$

A one-loop tadpole term $\mu_1 R$ is also induced by a single RNN insertion; however, since μ_1 is renormalization-scheme dependent and can be removed by an appropriate counterterm (equivalently, by imposing that the renormalized one-point function vanishes in the chosen vacuum), we do not treat it as an independent parameter and we neglect it in the following.

For a single heavy Majorana neutrino with mass M_N and couplings as in Eq. (D2), the one-loop threshold corrections at a renormalization scale μ can be written as

$$\mu_3^{\text{1L}} = -\frac{y_N^3 M_N}{8\pi^2} \left[3 \ln \frac{M_N^2}{\mu^2} + 2 \right], \quad (\text{D3})$$

$$\eta_1^{\text{1L}} \simeq -\frac{y_N M_N}{4\pi^2} y_\nu^2 \left(\ln \frac{M_N^2}{\mu^2} - 1 \right), \quad (\text{D4})$$

where y_N and y_ν denote the corresponding one-generation limits of Y_N and Y_ν , and y_ν is defined such that $m_D^2(\Phi) = y_\nu^2(\Phi^\dagger \Phi)$ in the seesaw limit. We stress that the overall coefficients are subject to scheme-dependent finite terms, while the parametric dependence on y_N , M_N , and y_ν^2 is robust.

In the basis where the heavy-neutrino mass matrix is diagonal, $M_N = \text{diag}(M_{N_i})$, the above expressions generalize to the schematic structures

$$\mu_3^{\text{1L}} \sim -\frac{1}{8\pi^2} \sum_i y_{N_i}^3 M_{N_i} \left[3 \ln \frac{M_{N_i}^2}{\mu^2} + \mathcal{O}(1) \right], \quad (\text{D5})$$

$$\eta_1^{\text{1L}} \propto -\frac{1}{16\pi^2} \text{Re Tr}[Y_N M_N Y_\nu^\dagger Y_\nu] \times \mathcal{O}(1), \quad (\text{D6})$$

up to scheme-dependent finite terms and mild logarithmic factors. The parametric dependence on Y_N , M_N , and $Y_\nu^\dagger Y_\nu$ is fixed by the diagrammatic structure.

2. Numerical estimates for $M_N \sim \mathcal{O}(10 - 50)$ TeV

It is useful to rewrite Eq. (D4) in terms of light-neutrino masses using the seesaw scaling $y_\nu^2 \sim m_\nu M_N / v_h^2$, with $v_h = 246$ GeV. Taking $\mu \simeq M_N$ [so that $|\ln(M_N^2/\mu^2) - 1| \sim \mathcal{O}(1)$], we obtain the parametric estimate

$$|\eta_1^{\text{1L}}| \sim \frac{y_N m_\nu M_N^2}{4\pi^2 v_h^2} \simeq 2 \times 10^{-8} \text{ GeV} \left(\frac{y_N}{1} \right) \left(\frac{m_\nu}{0.05 \text{ eV}} \right) \left(\frac{M_N}{30 \text{ TeV}} \right)^2, \quad (\text{D7})$$

i.e., η_1^{1L} is typically at the level of $\mathcal{O}(10 \text{ eV})$ for multi-10 TeV heavy neutrinos and standard active-neutrino masses.

In contrast, the cubic induced by Eq. (D3) is not neutrino-mass suppressed,

$$|\mu_3^{\text{1L}}| \simeq \frac{y_N^3 M_N}{4\pi^2} \simeq 7.6 \times 10^2 \text{ GeV} \left(\frac{y_N}{1} \right)^3 \left(\frac{M_N}{30 \text{ TeV}} \right), \quad (\text{D8})$$

for $\mu \simeq M_N$. Therefore, without an $R \rightarrow -R$ selection rule, μ_3 is generically expected to be sizable and should be treated as an important parameter of the scalar potential (or, equivalently, one must specify a renormalization condition that removes it at a given scale).

3. Phenomenological impact

The phenomenologically most relevant odd operator is typically the linear portal $R(\Phi^\dagger \Phi)$, because after electroweak symmetry breaking it induces a Higgs-singlet mixing term $(\eta_1 v_h/2)Rh$. For $m_R \neq m_h$, the induced mixing angle scales as

$$\sin \alpha_{\eta_1} \simeq \frac{\eta_1 v_h}{2(m_h^2 - m_R^2)}, \quad (\text{D9})$$

where the factor 1/2 follows from the mixing term $(\eta_1 v_h/2)Rh$. Using Eq. (D7) and $m_R \ll m_h$ gives $\sin \alpha_{\eta_1} \sim 10^{-10} - 10^{-9}$ for $M_N \sim 20\text{--}50$ TeV and $y_N \sim 1$, i.e., extremely small. Consequently, Higgs-mediated spin-independent direct detection remains negligible for the neutrino-induced η_1^{1L} alone. For instance, for a scalar mediator H_p of mass m_{H_p} coupled to DM with Yukawa y_χ , one finds parametrically

$$\sigma_{\chi N}^{\text{SI}} \sim \frac{\mu_{\chi N}^2}{\pi} \left(\frac{f_{NMNY_\chi} \sin \alpha_{\eta_1}}{v_h m_{H_p}^2} \right)^2, \quad (\text{D10})$$

which is far below experimental sensitivity for $\sin \alpha_{\eta_1} \lesssim 10^{-9}$ unless additional UV contributions enhance η_1 .

The cubic $\mu_3 R^3$ instead primarily affects the scalar potential and the singlet vacuum structure. In particular, together with m_R^2 and λ_R , it controls whether $R = 0$ is a local minimum and how the singlet vacuum expectation value (if present) is shifted. Since Eq. (D8) indicates that μ_3^{1L} can be $\mathcal{O}(\text{TeV})$ for $M_N \sim \mathcal{O}(10 - 50)$ TeV, it can have a non-negligible impact on the scalar spectrum and self-interactions unless absorbed into the renormalized parameters by appropriate matching conditions. A complete treatment therefore requires specifying the renormalization conditions for the odd operators (tadpole and cubic) at the matching scale, or equivalently treating μ_3 (and η_1) as UV-sensitive parameters in addition to the irreducible even portal computed in the main text.

-
- [1] J. Silk *et al.*, *Particle Dark Matter: Observations, Models and Searches*, edited by G. Bertone (Cambridge University Press, Cambridge, England, 2010).
 - [2] G. Bertone and D. Hooper, *Rev. Mod. Phys.* **90**, 045002 (2018).
 - [3] M. Cirelli, A. Strumia, and J. Zupan, [arXiv:2406.01705](#).
 - [4] G. Bertone, D. Hooper, and J. Silk, *Phys. Rep.* **405**, 279 (2005).
 - [5] B. W. Lee and S. Weinberg, *Phys. Rev. Lett.* **39**, 165 (1977).
 - [6] J. E. Gunn, B. W. Lee, I. Lerche, D. N. Schramm, and G. Steigman, *Astrophys. J.* **223**, 1015 (1978).

- [7] J. Wess and B. Zumino, *Nucl. Phys.* **B70**, 39 (1974).
- [8] H. Goldberg, *Phys. Rev. Lett.* **50**, 1419 (1983).
- [9] J. R. Ellis, J. S. Hagelin, D. V. Nanopoulos, K. A. Olive, and M. Srednicki, *Nucl. Phys.* **B238**, 453 (1984).
- [10] N. Aghanim *et al.* (Planck Collaboration), [arXiv:1807.06209](#).
- [11] G. Steigman and M. S. Turner, *Nucl. Phys.* **B253**, 375 (1985).
- [12] M. Schumann, *J. Phys. G* **46**, 103003 (2019).
- [13] A. Boveia and C. Doglioni, *Annu. Rev. Nucl. Part. Sci.* **68**, 429 (2018).

- [14] J. M. Gaskins, *Contemp. Phys.* **57**, 496 (2016).
- [15] J. Aalbers *et al.*, *Phys. Rev. Lett.* **131**, 041002 (2023).
- [16] E. Aprile *et al.*, *Phys. Rev. Lett.* **131**, 041003 (2023).
- [17] J. Aalbers *et al.* (LZ Collaboration), *Phys. Rev. Lett.* **135**, 011802 (2025).
- [18] G. Arcadi, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, M. Pierre, S. Profumo, and F. S. Queiroz, *Eur. Phys. J. C* **78**, 203 (2018).
- [19] G. Arcadi, A. Djouadi, and M. Raidal, *Phys. Rep.* **842**, 1 (2020).
- [20] M. Di Mauro, C. Arina, N. Fornengo, J. Heisig, and D. Massaro, *Phys. Rev. D* **108**, 095008 (2023).
- [21] G. Arcadi, D. Cabo-Almeida, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, J. P. Neto, M. Pierre, S. Profumo, and F. S. Queiroz, *Eur. Phys. J. C* **85**, 152 (2025).
- [22] M. Di Mauro and B. Xie, *Phys. Rev. D* **113**, 015034 (2026).
- [23] C. Kong and M. Di Mauro, *Phys. Rev. D* **113**, 043031 (2026).
- [24] M. Pospelov, A. Ritz, and M. B. Voloshin, *Phys. Lett. B* **662**, 53 (2008).
- [25] M. Pospelov and A. Ritz, *Phys. Lett. B* **671**, 391 (2009).
- [26] M. Di Mauro and Y. Wang, *Phys. Rev. D* **113**, 075003 (2026).
- [27] J. Aalbers *et al.* (DARWIN Collaboration), *J. Cosmol. Astropart. Phys.* **11** (2016) 017.
- [28] N. Aghanim *et al.*, *Astron. Astrophys.* **641**, A6 (2020).
- [29] A. G. Adame *et al.* (DESI Collaboration), *J. Cosmol. Astropart. Phys.* **02** (2025) 021.
- [30] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro, and T. Schwetz, *J. High Energy Phys.* **12** (2024) 216.
- [31] M. Aker *et al.* (KATRIN Collaboration), *Science* **388**, adq9592 (2025).
- [32] P. Minkowski, *Phys. Lett.* **67B**, 421 (1977).
- [33] T. Yanagida, in *Proceedings of the Workshop on the Unified Theory and the Baryon Number in the Universe* (1979), <https://inspirehep.net/literature/19833>.
- [34] M. Gell-Mann, P. Ramond, and R. Slansky, *Conf. Proc. C* **790927**, 315 (1979), <https://inspirehep.net/literature/9686>.
- [35] R. N. Mohapatra and G. Senjanovic, *Phys. Rev. Lett.* **44**, 912 (1980).
- [36] D. Wyler and L. Wolfenstein, *Nucl. Phys.* **B218**, 205 (1983).
- [37] R. N. Mohapatra and J. W. F. Valle, *Phys. Rev. D* **34**, 1642 (1986).
- [38] S. Baek, P. Ko, and W.-I. Park, *J. High Energy Phys.* **02** (2011) 047.
- [39] S. Baek, P. Ko, and W.-I. Park, *J. High Energy Phys.* **07** (2013) 013.
- [40] B. Holdom, *Phys. Lett.* **166B**, 196 (1986).
- [41] R. Essig *et al.*, [arXiv:1311.0029](https://arxiv.org/abs/1311.0029).
- [42] H. Ruegg and M. Ruiz-Altaba, *Int. J. Mod. Phys. A* **19**, 3265 (2004).
- [43] B. Kors and P. Nath, *Phys. Lett. B* **586**, 366 (2004).
- [44] J. A. Casas and A. Ibarra, *Nucl. Phys.* **B618**, 171 (2001).
- [45] K. R. Dienes, C. F. Kolda, and J. March-Russell, *Nucl. Phys.* **B492**, 104 (1997).
- [46] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, *Phys. Lett. B* **429**, 263 (1998).
- [47] L. Randall and R. Sundrum, *Phys. Rev. Lett.* **83**, 3370 (1999).
- [48] M. A. Luty and R. Sundrum, *Phys. Rev. D* **67**, 045007 (2003).
- [49] M. Dine, P. J. Fox, E. Gorbatov, Y. Shadmi, Y. Shirman, and S. D. Thomas, *Phys. Rev. D* **70**, 045023 (2004).
- [50] H. Murayama, Y. Nomura, and D. Poland, *Phys. Rev. D* **77**, 015005 (2008).
- [51] L. Randall and R. Sundrum, *Nucl. Phys.* **B557**, 79 (1999).
- [52] T. Gherghetta, J. Kersten, K. Olive, and M. Pospelov, *Phys. Rev. D* **100**, 095001 (2019).
- [53] T. G. Rizzo, *Phys. Rev. D* **107**, 095014 (2023).
- [54] T. G. Rizzo, *Phys. Rev. D* **110**, 075037 (2024).
- [55] D. Feldman, B. Kors, and P. Nath, *Phys. Rev. Lett.* **99**, 251802 (2007).
- [56] T. Bringmann, P. F. Depta, M. Hufnagel, and K. Schmidt-Hoberg, *Phys. Lett. B* **817**, 136341 (2021).
- [57] M. Fukugita and T. Yanagida, *Phys. Lett. B* **174**, 45 (1986).
- [58] W. Buchmuller, P. Di Bari, and M. Plumacher, *Ann. Phys. (Amsterdam)* **315**, 305 (2005).
- [59] S. Davidson, E. Nardi, and Y. Nir, *Phys. Rep.* **466**, 105 (2008).
- [60] G. F. Giudice, A. Notari, M. Raidal, A. Riotto, and A. Strumia, *Nucl. Phys.* **B685**, 89 (2004).
- [61] C. S. Fong, E. Nardi, and A. Riotto, *Adv. High Energy Phys.* **2012**, 158303 (2012).
- [62] F. Hahn-Woernle, M. Plumacher, and Y. Y. Y. Wong, *J. Cosmol. Astropart. Phys.* **09** (2009) 010.
- [63] E. Braaten and R. D. Pisarski, *Nucl. Phys.* **B339**, 310 (1990).