

Dark matter as geometrically sequestered baryons: Effective electromagnetic disconnection in throat spacetimes

Jeff Riley^{1,2}

¹*School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia*

²*OzGrav, ARC Centre of Excellence for Gravitational Wave Discovery, Australia**

(Dated: May 2, 2026)

We propose that the dark matter problem may have a geometric origin within general relativity: ordinary baryonic matter residing in spatial regions connected to ours only through narrow geometric throats would be electromagnetically invisible (due to strong sub-barrier suppression of electromagnetic modes with $\ell \geq 1$) while remaining gravitationally coupled (through the $\ell = 0$ monopole of the Poisson equation, which encounters no centrifugal barrier). This constraint-wave transmission asymmetry is derived from the four-dimensional field equations on a throat background in a companion paper (Riley 2026 [1]), and is a universal feature of static, spherically symmetric spacetimes with minimal-area two-spheres, provided the lapse remains nonzero at the throat. The proposal is conditional: it presupposes the existence of submicron throat geometries in the spatial fabric of the universe—a question we do not resolve here—and assesses the cosmological consequences that follow if such geometries exist. We identify the necessary conditions for consistency with Big Bang Nucleosynthesis, the cosmic microwave background, and large-scale structure: the sequestered regions must have radiation temperatures below $\sim 39\%$ of the visible sector’s at recombination, achievable if spatial bottlenecks form before reheating. Sequestered baryons are pressureless and non-relativistic at late times (like CDM), retaining a Jeans scale set by their internal temperature. A self-consistent implementation in the Boltzmann code CLASS—in which the sequestered component is evolved as an interacting dark matter fluid with constant sound speed $c_s^2 = k_B T_{\text{seq}} / (m_p c^2)$ —shows that the CMB power spectrum is indistinguishable from Λ CDM at the 5×10^{-5} numerical precision of our modified code for $T_{\text{seq}} \leq 1000$ K. The minimal scenario, in which the sequestered diffuse phase is heated only by the primordial dark-photon bath via residual-electron Compton coupling, predicts a present-epoch diffuse-phase temperature in the millikelvin range and is indistinguishable from Λ CDM at all observable scales— $P(k)/P_{\Lambda\text{CDM}} > 0.999$ at the Lyman- α pivot and CMB residuals at the 5×10^{-5} numerical precision floor of our modified Boltzmann code, verified by self-consistent CLASS evolution of the time-dependent thermal history. We show that gravitational and electromagnetic radiation from sequestered regions are strongly suppressed for submicron throats, that the gravitational coupling between sectors is dominated by the static monopole at all observable wavenumbers, and that microlensing constraints place the per-throat mass in the asteroid-mass window ($M \sim 5 \times 10^{-18} - 5 \times 10^{-11} M_\odot$, corresponding to $N \gtrsim 2 \times 10^{22}$ throats per Milky-Way-like halo). The scenario constitutes a well-defined research programme rather than a finished theory: its viability rests on the open question of whether throat geometries can be sustained by physics within or near the Standard Model, and we organise the outstanding questions as five concrete work packages.

I. Introduction

More than four decades of direct searches for particle dark matter have yielded no confirmed detection [2, 3], despite increasingly sensitive experiments spanning weakly interacting massive particles, axions, sterile neutrinos, and other candidates. The gravitational evidence for dark matter—galaxy rotation curves [4, 5], gravitational lensing [6], the cosmic microwave background (CMB) power spectrum [7], and large-scale structure [8]—is overwhelming, but it constrains only the

gravitational properties of the dark sector: its spatial distribution, its equation of state, and its abundance. No electromagnetic, strong, or weak, interaction of dark matter has been observed.

This motivates the question: could the “darkness” of dark matter be a geometric property of spacetime, rather than a property of the matter itself?

In a companion paper [1], we demonstrate a structural asymmetry in field propagation through geometric throats—electromagnetic modes with angular momentum $\ell \geq 1$ encounter a centrifugal potential barrier $V_\ell = \ell(\ell + 1)/a^2$ at the throat, and are strongly suppressed below the barrier-top frequency $\omega_{\text{max}} = \sqrt{\ell(\ell + 1)}/r_0$, while the static gravitational monopole ($\ell = 0$) satisfies a conservation law $(a^2 \Phi')' = 0$

* jeff.riley@monash.edu

with no barrier, transmitting through the throat with only polynomial attenuation. This constraint–wave asymmetry is a mathematical property of the field equations on any spherically symmetric background with a minimal-area two-sphere, independent of the matter content sourcing the geometry.

In this paper we explore the consequences of this asymmetry for cosmology. We propose that if the spatial geometry of the universe contains a network of narrow throats—regions where the areal radius a drops to a minimum value r_0 —then baryonic matter on the far side of these throats would be gravitationally coupled to our region, but electromagnetically invisible. Such matter would exhibit precisely the observational signature of dark matter.

For the purposes of this paper we assume the existence of narrow throats in the spatial geometry of our universe. Static throats require null energy condition (NEC) violation, which may be sourced by semiclassical, quantum-gravitational, or modified-gravity effects. We adopt the Ellis-Bronnikov wormhole as a fiducial geometry, and work within an explicitly conditional framework.

We demonstrate the cosmological viability of our proposal (consistency with BBN, CMB, large-scale structure)—given the existence of throats, we show quantitative consistency with precision cosmological data using a self-consistent Boltzmann code implementation. The constraint-wave asymmetry derived in [1] underpins this work.

Our proposal is distinct from several related ideas in the literature. Kirillov and Savelova [9, 10] have proposed that dark matter phenomenology arises from a cosmological “gas of wormholes” inherited from the spacetime foam—in their framework, the wormholes themselves modify the gravitational potential through topological polarisation, and the dark matter effect is a property of the wormhole distribution rather than of matter residing behind the throats. In contrast, our scenario places ordinary baryonic matter behind the throats and invokes the constraint–wave transmission asymmetry to render it electromagnetically dark while gravitationally coupled. Mirror dark matter models [11] achieve electromagnetic decoupling through a discrete Z_2 symmetry between visible and hidden copies of the Standard Model—our mechanism is geometric rather than particle-physics in origin. String-theoretic warped throat constructions [12, 13] suppress energy transfer into sequestered sectors via warp factors in extra dimensions—our proposal operates entirely within four-dimensional general relativity. Dai and Stojkovic [14] showed that the gravitational monopole of a mass on one side of a traversable wormhole is detectable on the other side, with higher multipoles geometrically suppressed, and proposed tests using stellar orbits near Sgr A*. Our scenario builds on the same monopole-transmission physics but replaces their static thin-shell matching with the frequency-dependent centrifugal-barrier analysis of [1], which establishes the wave-propagation suppression that

renders the sequestered matter electromagnetically dark.

The remainder of this paper is organised as follows: in Section II we summarise the constraint–wave asymmetry; the energy conditions for throat geometries are analysed in Section III; Section IV addresses Big Bang Nucleosynthesis constraints; we discuss the CMB and large-scale structure implications in Section V; Section VI analyses structure formation with sequestered baryons; in Section VII we identify observational signatures and falsifiable predictions; Section VIII lists the open problems that must be solved to establish or rule out the proposal and organises them as a research programme; and we present our summary and conclusions in Section IX.

II. The constraint-wave asymmetry

Following we summarise the constraint-wave asymmetry derived by Riley in [1] (we refer interested readers to the paper for complete details).

Consider a static, spherically symmetric spacetime with metric

$$ds^2 = -e^{2\alpha(\sigma)} dt^2 + d\sigma^2 + a(\sigma)^2 d\Omega^2, \quad (1)$$

where σ is the proper radial coordinate and $a(\sigma)$ has a local minimum at $\sigma = 0$ (the “throat”) with $a(0) = r_0$. The four-dimensional Maxwell equations on this background yield, for each electromagnetic multipole ℓ , a Schrödinger-type equation with effective potential

$$V_\ell^{(\text{EM})} = e^{2\alpha} \frac{\ell(\ell+1)}{a(\sigma)^2}. \quad (2)$$

This potential has a maximum at the throat. For frequencies $\omega < \omega_{\text{max}}^{(\ell)} = \sqrt{e^{2\alpha(0)} \ell(\ell+1)} / r_0$, the electromagnetic mode must tunnel through the barrier, giving strong sub-barrier suppression. On the ultrastatic Ellis-Bronnikov (EB) throat ($e^{2\alpha} = 1$, $a(\sigma) = \sqrt{\sigma^2 + r_0^2}$), the EB barrier has a $1/\sigma^2$ tail that produces power-law suppression at low frequencies [1]:

$$T_{\text{EM}}(\omega) \sim K (\omega r_0)^\nu, \quad \omega r_0 \ll 1, \quad (3)$$

where ν and K depend on the barrier shape. The suppression is extremely strong—for the lowest EM mode ($\ell = 1$) at $\omega r_0 = 0.1$, $T_{\text{EM}} \approx 10^{-8}$ [1].

The static gravitational potential Φ (the Newtonian potential in the weak-field limit) satisfies, in the source-free region,

$$\frac{1}{a^2} \frac{d}{d\sigma} \left(a^2 e^{2\alpha} \frac{d\Phi}{d\sigma} \right) = 0, \quad (4)$$

which is a conservation law: the flux $\mathcal{F} = a^2 e^{2\alpha} \Phi'$ is constant. There is no potential barrier and no sub-barrier suppression, provided $e^{2\alpha} > 0$ at the throat [1]. The monopole transmits through the throat smoothly.

On the EB throat, the exact solution is $\Phi = (C/r_0) \arctan(\sigma/r_0)$, and the numerical transmission coefficients confirm the asymmetry—for

$\omega r_0 = 0.1$, the EM transmission is $T_{\text{EM}} \approx 10^{-8}$ while the gravitational monopole transmits with $\mathcal{O}(r_0/d)$ corrections. The constraint-wave asymmetry is not specific to the EB profile: universality is established across a parametric family of throat profiles and across the non-ultrastatic reflected-Schwarzschild (Damour–Solodukhin) wormhole [1], demonstrating that the qualitative asymmetry—unsuppressed $\ell = 0$ constraint sector versus strongly suppressed $\ell \geq 1$ propagating and static multipoles—holds on any static, spherically symmetric throat background with $e^{2\alpha}(0) > 0$. Quantitative barrier heights and sub-barrier exponents depend on the specific profile, but the qualitative kinematic structure that underlies the cosmological application below is universal.

A central point for the cosmological application below is that the constraint-wave asymmetry is not simply a contrast between gravity and radiation—it distinguishes the $\ell = 0$ gravitational monopole from all higher multipoles, including the static $\ell \geq 1$ components of the gravitational field. On the spatial metric $d\sigma^2 + a(\sigma)^2 d\Omega^2$, the source-free $\ell \geq 1$ Poisson equation can be written in Schrödinger form via $\psi_\ell = a \Phi_\ell$ as

$$\psi_\ell'' - \left[\frac{\ell(\ell+1)}{a^2} + \frac{a''}{a} \right] \psi_\ell = 0, \quad (5)$$

in which the centrifugal term dominates near the throat and creates a positive-definite barrier analogous to the wave case [1]. These modes are not propagating (the equation is elliptic), but they still decay through the throat region. A source of gravitational multipole ℓ at far-side distance $d \gg r_0$ induces, on the near side, an amplitude suppressed by the scaling

$$\frac{\Phi_\ell(\text{near})}{\Phi_\ell(\text{source})} \sim \left(\frac{r_0}{d} \right)^{2\ell+1}, \quad (6)$$

which follows from matching the near-throat solution $\Phi_\ell \sim r_0^\ell/d^{\ell+1}$ (source-side decay) to the far-side growing mode σ^ℓ with connection coefficient $C_\ell = \mathcal{O}(1)$ [1]. For $\ell = 1$ (tidal field) at $d = 10 r_0$, the suppression is 10^{-3} , and 10^{-5} for $\ell = 2$ (quadrupole). These suppressions are polynomial rather than exponential, but they are dramatically stronger than the $\mathcal{O}(r_0/d)$ monopole correction.

A compact source on the far side of a throat projects a clean, nearly pointlike Newtonian monopole onto our sector, while its dipole, tidal, and higher static multipoles are geometrically filtered by the same centrifugal mechanism that suppresses propagating radiation. The internal structure of sequestered objects—disks, compact objects, density profiles—is invisible beyond its contribution to the total enclosed mass. This is the mathematical basis for the point-mass lensing approximation used in Sec. VII E, and the effective monopole coupling used throughout the cosmological analysis.

Gravitational radiation, ($\ell \geq 2$ tensor modes), is also a propagating field, and sees a barrier qualitatively similar

to the EM one—for submicron throats, all astrophysical GW frequencies lie below the barrier top and are strongly suppressed (Sec. VII). The constraint-wave asymmetry is therefore specifically between the static gravitational potential (a constraint/elliptic equation) and the propagating radiation of any spin. The electric Coulomb field (the $\ell = 0$ mode of the Maxwell equations) satisfies a Gauss-law conservation identical in structure to the gravitational monopole: $(a^2 E')' = 4\pi \rho_e$ —the electric monopole therefore also transmits through the throat. However, baryonic matter is electrically neutral, so the net $\ell = 0$ electric flux is zero and no electromagnetic monopole signal is expected.

The monopole conservation law (4) was derived on a strictly static background, yet we will apply it to perturbations in an FLRW spacetime with a time-dependent expansion rate. The regime in which the static derivation continues to hold is controlled by a clean hierarchy of scales. For a throat embedded in an FLRW region with Hubble rate H and throat crossing time $\tau_{\text{throat}} \sim r_0/c$, the static approximation is accurate whenever $H\tau_{\text{throat}} \sim Hr_0/c \ll 1$, which for submicron throats ($r_0 \lesssim 10^{-6}$ m) and any cosmological epoch after inflation gives $Hr_0/c \lesssim 10^{-31}$. The monopole configuration across the throat therefore adjusts essentially instantaneously on cosmological timescales, and the static derivation applies with corrections at $\mathcal{O}(Hr_0/c)$, far below any observable effect.

The derivation additionally assumes three conditions:

- (i) the lapse remains bounded away from zero at the throat, $e^{2\alpha(0)} > 0$, so that the conserved flux $\mathcal{F} = a^2 e^{2\alpha} \Phi'$ gives a bounded Φ' and hence a smooth monopole potential across the throat [1];
- (ii) no event horizons intervene between the source and the point of observation, so that the elliptic problem is globally well-posed on the constant- t slice; and
- (iii) the throat geometry is itself slowly varying, $\dot{r}_0/r_0 \ll H$, so that the throat can be treated as quasistatic on the timescale of structure growth.

The first condition excludes the $\lambda \rightarrow 0$ limit of the Damour–Solodukhin construction (where $e^{2\alpha}$ vanishes at the throat, recovering the Schwarzschild horizon and rendering the monopole non-transmissive [1]) and is the physical reason we restrict the cosmological analysis to non-ultrastatic throats with $e^{2\alpha} > 0$ throughout. The second condition is automatically satisfied for any geometry without a classical horizon. The third condition is satisfied by any stabilised throat and by the coexpanding Lemaitre–Tolman–Bondi (LTB) throats of Sec. III E (which have $\dot{r}_0/r_0 = H$, formally marginal—the condition then becomes an equality rather than a strict inequality, and corrections are first-order in the ratio rather than suppressed). Under these conditions, the monopole conservation law and the coupling kernel $W(k, r_0)$ derived in Appendix A hold in the FLRW-embedded setting, and the static derivation supports treating the static

analysis as the leading approximation of the gravitational coupling between sectors on cosmological scales. A full derivation is deferred to future work (Sec. VIII).

We return to the gravitational wave implications in Sec. VII.

III. Energy conditions and the fiducial throat geometry

A. The energy condition problem

A static, spherically symmetric throat in general relativity generically violates the null energy condition (NEC) [15–17]. The Hamiltonian constraint on a time-symmetric ($K_{ij} = 0$) spatial slice with the EB throat profile $a(\sigma) = \sqrt{\sigma^2 + r_0^2}$ gives

$$R^{(3)} = -2r_0^2/(\sigma^2 + r_0^2)^2 < 0,$$

requiring $\rho < 0$ [17]. This is the well-known wormhole energy condition problem. The field-propagation asymmetry derived in [1] is a kinematic result—it depends only on the background geometry, not on the stress-energy sourcing it. The energy condition question constrains the dynamical realisability of the geometry, not the transmission properties.

B. Fiducial geometry: the Ellis-Bronnikov throat

As noted, we adopt the Ellis-Bronnikov (EB) wormhole as our fiducial throat geometry. The EB solution is an exact solution of the Einstein equations coupled to a massless phantom scalar field φ with wrong-sign kinetic term ($\mathcal{L} = -\frac{1}{2}(\nabla\varphi)^2$), yielding the ultrastatic metric

$$ds^2 = -dt^2 + d\sigma^2 + (\sigma^2 + r_0^2) d\Omega^2, \quad (7)$$

where $\sigma \in (-\infty, +\infty)$ and the throat at $\sigma = 0$ connects two asymptotically flat regions. The spatial manifold has topology $\mathbb{R} \times S^2$ (not \mathbb{R}^3)—the two-sphere at $\sigma = 0$ has finite area $4\pi r_0^2$ and is a genuine minimal surface, not a regular centre (where $a \rightarrow 0$).

The phantom field violates the NEC. The required negative energy density is localised within $\sim r_0$ of the throat:

$$\rho_{\text{NEC}} = -\frac{c^4}{8\pi G} \frac{r_0^2}{(\sigma^2 + r_0^2)^2}, \quad (8)$$

with total negative energy per throat

$$E_{\text{NEC}} = \int \rho_{\text{NEC}} \sqrt{\gamma} d^3x = -\frac{\pi c^4 r_0}{4G}. \quad (9)$$

For $r_0 \sim 10^{-13}$ m,

$$|E_{\text{NEC}}| \sim 10^{31} \text{ J} \sim 10^{14} \text{ kg } c^2.$$

This is substantial, but finite and localised. For sub-micron throats ($r_0 \sim 10^{-7}$ m), the per-throat exotic energy is

$$|E_{\text{NEC}}| \sim \pi c^4 r_0 / (4G) \approx 10^{37} \text{ J},$$

equivalent to a negative mass of $\sim 10^{20}$ kg per throat—for sub-nanometre throats ($r_0 \sim 10^{-9}$ m), $|E_{\text{NEC}}| \sim 10^{35}$ J, and so on in proportion to r_0 . The unifying scaling is $|E_{\text{NEC}}|/(M_{\text{Pl}}c^2) \sim r_0/\ell_{\text{Pl}}$, which evaluates to $\sim 10^{27}$ for $r_0 = 10^{-7}$ m and $\sim 10^{21}$ for $r_0 = 10^{-13}$ m—the values quoted above are mutually consistent under this single scaling, with the per-throat exotic energy growing linearly with the throat radius across the entire viable window.

The cumulative exotic-energy budget across a galactic halo depends on the product $N_{\text{throat}} \cdot r_0$, where N_{throat} is the number of throats per halo. Requiring that the integrated negative energy not exceed the halo rest mass— $|E_{\text{NEC}}^{\text{halo}}|/M_{\text{halo}}c^2 \lesssim 1$, the natural bound for the backreaction on averaged cosmological expansion and halo-scale curvature to remain non-catastrophic—gives, for $M_{\text{halo}} = 10^{12} M_{\odot}$,

$$|E_{\text{NEC}}^{\text{halo}}| \sim N_{\text{throat}} \frac{\pi c^4 r_0}{4G}, \quad N_{\text{throat}} r_0 \lesssim 2 \times 10^{15} \text{ m}. \quad (10)$$

This provides a positive consistency check for our scenario—because the per-throat exotic energy scales only linearly with r_0 , even very large throat populations ($N \sim 10^{22}$ – 10^{29} , as required by the lensing analysis of Sec. VII E) can be accommodated by correspondingly small throat radii without producing catastrophic backreaction. We adopt the natural bound $|E_{\text{NEC}}^{\text{halo}}| \lesssim M_{\text{halo}}c^2$ rather than a more aggressive fractional criterion: the appropriate criterion for the scenario’s viability is that the integrated exotic energy not exceed the halo rest mass, since a smaller fractional bound would amount to an additional, unstated, requirement on the cosmological backreaction whose physical motivation we cannot derive from first principles. Tighter fractional bounds, if eventually justified by a quantitative backreaction analysis, would simply rescale the constraint (10) accordingly.

Scenarios with vastly larger throat counts—for example, a phase-transition formation mechanism at the QCD scale producing $N \sim 10^{75}$ throats per halo at $r_0 \sim 10^{-7}$ m—would instead require $|E_{\text{NEC}}^{\text{halo}}| \sim 10^{93}$ J, roughly fifty-three orders of magnitude above the halo rest-mass energy, and are excluded by the exotic-energy budget alone, independently of their lensing predictions. The lensing constraint (52) and the exotic-energy bound (10) therefore jointly determine the (N_{throat}, r_0) viable window, with throat radii required to span roughly seven orders of magnitude from $r_0 \sim 10^{-7}$ m at the lower-count edge of the lensing window down to $r_0 \sim 10^{-14}$ m at the upper-count edge.

Whether submicron throat geometries can be sustained by known physics is an open question. Several mechanisms have been proposed:

- (i) semiclassical effects (Casimir energy, vacuum polarisation) produce localised negative energy densities in confined geometries [17, 18], with the Casimir energy density scaling as $\rho_{\text{Cas}} \sim -\hbar c/d^4$ —any such realisation must also respect the quantum inequalities of Ford and Roman [19, 20], which place integrated bounds on the magnitude and duration of negative energy accessible to any observer, and which therefore constrain the radius r_0 and lifetime of the throats compatible with a semiclassical source. Recent work by Kontou [21] reviews the state of the art in this area and works through the smeared and double-smeared null energy conditions (SNEC and DSNEC) for the Maldacena–Milekhin–Popov long-wormhole construction [22], finding that even this most conservative semiclassical wormhole—sourced entirely by Casimir energy from Standard Model fermions, with no exotic matter—faces severe tension with the DSNEC unless an unphysically large number of field species, or a departure from standard semiclassical effective field theory, is invoked. A quantitative confrontation of these bounds, with the per-throat exotic-energy budget of Eq. (9) and the cumulative budget of Eq. (10) for the geometric-sequestration scenario, is deferred to future work (Sec. VIII)—we note only that the bounds depend on the assumed UV cutoff, field content, and effective-action contributions, and that the submicron throat radii selected by the lensing analysis of Sec. VII E differ qualitatively from the long-wormhole regime where the DSNEC tension is most acute—moreover, for submicron r_0 the per-throat exotic energy $|E_{\text{NEC}}| \sim c^4 r_0/G$ lies at or below the nuclear energy scale, precisely the regime where semiclassical estimates of negative-energy bounds are least reliable and where quantum-gravitational corrections to the effective stress-energy are most likely to operate;
- (ii) modified gravity theories—including Einstein–Gauss–Bonnet constructions [23, 24], non-minimal gravitational-gauge couplings [25], $f(R)$ gravity [26], and scalar-tensor extensions [27]—can support traversable and non-traversable throat geometries with localised or absent classical NEC violation, by sourcing the required effective stress-energy from higher-curvature terms or non-minimal gauge couplings rather than from a phantom scalar;
- (iii) quantum gravity effects at scales $r_0 \lesssim \ell_{\text{Pl}}$ are expected to modify the classical energy conditions, and may permit throat-supporting stress-energy distributions that are forbidden in the classical theory—loop quantum gravity, asymptotic-safety, and string-theoretic approaches each provide qualitatively different routes to such modifications, none of which has been worked out in detail for the throat geometries considered here.

A separate question from whether any of the mech-

anisms (i)–(iii) can sustain a throat is whether the constraint-wave asymmetry of [1], derived for general relativity with a phantom-scalar source, continues to hold on the modified backgrounds these mechanisms produce. The effective potentials (2) and (4) depend only on the background geometry and should persist on any throat background with $e^{2\alpha} > 0$ at the throat, so the qualitative asymmetry is expected to survive. The equations of motion for propagating tensor modes, however, and the corresponding barrier heights, could receive corrections at order $\mathcal{O}(r_0^2/\ell_{\text{mg}}^2)$, where ℓ_{mg} is the modification scale of the gravitational sector. For submicron throats and modification scales well above r_0 , such corrections are negligible and the asymmetry carries over directly. For modification scales comparable to or below r_0 , the question is open, and discussed in (Sec. VIII).

We do not resolve either of these questions here. Instead, we adopt the following conditional framework: if throat geometries of the form (7) exist in the spatial geometry of our universe—sustained by whatever mechanism—then the cosmological consequences derived in the remainder of this paper follow from the constraint-wave asymmetry and standard cosmological physics.

This conditional status is shared by all dark matter proposals: WIMP models assume the existence of particles not yet detected, axion models require a Peccei–Quinn symmetry not yet confirmed, and primordial black hole models require a specific inflationary power spectrum not yet measured. Our assumption—that submicron geometric throats exist—is speculative, but leads to definite, testable predictions.

Adopting the EB metric is a computational choice, not a commitment to phantom-scalar matter as the underlying support mechanism. The EB profile is the simplest analytically tractable throat geometry with a minimal two-sphere, and serves as a fiducial template on which the field-propagation analysis of [1], and the cosmological consequences derived below can be computed in closed form. Any of the candidate support mechanisms identified in items (i)–(iii) above—if successful in producing a stable submicron throat—would yield a geometry of qualitatively similar form, and the conclusions of this paper are expected to carry over to leading order, with corrections discussed in Sec. VIII.

C. Topological censorship

Topological censorship theorems [28, 29] prohibit traversable wormholes (spacetimes with non-trivial fundamental group) under the averaged NEC (ANEC) and global hyperbolicity. Our scenario does not require traversal—no matter, light, or information, needs to pass through the throat. The throat functions as a gravitational conduit for the static monopole (a constraint equation, not a propagating signal), while all propagating modes are blocked. Whether a non-traversable throat connecting two regions can persist under ANEC is a ques-

tion distinct from traversability, and has not been addressed in the topological censorship literature.

If the spatial topology is simply connected ($\pi_1 = 0$), the topological censorship theorems do not directly apply. A geometric constriction where the areal radius reaches a minimum without a topological handle — what we term a “sub-topological bottleneck” — may evade these theorems entirely. The null expansions at the minimal surface satisfy $\theta_+ > 0$, $\theta_- < 0$, distinguishing it from both a trapped surface and a marginally outer trapped surface.

D. Dynamical evolution in a cosmological setting

To assess the persistence of a throat in an expanding universe, we embed the EB spatial geometry

$$\gamma_{ij} dx^i dx^j = d\sigma^2 + (\sigma^2 + r_0^2) d\Omega^2 \quad (11)$$

in a cosmological spacetime with uniform Hubble expansion:

$$K_{ij} = H \gamma_{ij}, \quad K = 3H = \text{const.} \quad (12)$$

The Hamiltonian constraint gives the energy density

$$16\pi G\rho = R^{(3)} + 6H^2 = -\frac{2r_0^2}{(\sigma^2 + r_0^2)^2} + 6H^2. \quad (13)$$

For $H > 1/(r_0\sqrt{3})$, this is positive everywhere — the expansion contributes sufficient positive energy to overcome the negative spatial curvature at the throat. The resulting LTB dust evolution (Sec. III E) demonstrates that the throat persists and coexpands, but this construction does not avoid the fundamental NEC issue — the positive ρ comes from the extrinsic curvature (expansion), not from the spatial geometry, which still has $R^{(3)} < 0$ at the throat. A static ($H = 0$) throat would require $\rho < 0$ there. The LTB evolution demonstrates dynamical persistence of the throat shape, not the existence of a matter model that sources it from first principles.

E. Illustrative LTB evolution with throat initial data

As an illustration of how throat geometry behaves in a cosmological setting, we consider the LTB dust evolution of initial data with the EB spatial profile. This construction demonstrates compatibility with the Hamiltonian constraint (positive energy density on the initial slice), but does not constitute a complete spacetime construction — the global properties, causal structure, and consistency with topological censorship theorems remain to be established. We present it as motivation, not proof, that throat geometries could persist in an expanding universe.

The solution is of LTB type — a spherically symmetric dust spacetime with the radial coordinate $\sigma \in (-\infty, +\infty)$

and metric

$$ds^2 = -dt^2 + \frac{[\partial a/\partial\sigma]^2}{1 + 2E(\sigma)} d\sigma^2 + a(t, \sigma)^2 d\Omega^2, \quad (14)$$

where the areal radius $a(t, \sigma)$ satisfies

$$\left(\frac{\partial a}{\partial t}\right)^2 = \frac{2M(\sigma)}{a} + 2E(\sigma), \quad (15)$$

and $M(\sigma)$ (the Misner-Sharp mass) and $E(\sigma)$ (the energy function) are determined by the initial data (11)–(12). Note that $\sigma = 0$ is not a regular centre — at a regular centre, $a \rightarrow 0$ and the two-sphere shrinks to a point, while here $a(t, 0) = r_0 > 0$ and the two-sphere has finite area $4\pi r_0^2$. The domain $\sigma \in (-\infty, +\infty)$ represents two asymptotic regions ($\sigma \rightarrow \pm\infty$, each approaching FLRW) connected through a minimal-area two-sphere at $\sigma = 0$. The solution is symmetric under $\sigma \rightarrow -\sigma$ by construction. The stress-energy is that of pressureless dust:

$$T^{\mu\nu} = \rho(t, \sigma) u^\mu u^\nu, \quad \rho = \frac{M'(\sigma)}{4\pi a^2 \partial a/\partial\sigma}, \quad (16)$$

with $u^\mu = (1, 0, 0, 0)$ in comoving coordinates. The lapse is $N = 1$ (geodesic slicing), nonvanishing everywhere.

On the initial slice, $\rho \geq 0$ and $p = 0$, so the dust stress-energy satisfies all classical energy conditions. However, this does not resolve the fundamental NEC issue for the throat geometry itself (Sec. III B) — the positive ρ arises from the extrinsic curvature contribution, and a static ($H = 0$) version of this geometry would require $\rho < 0$ at the throat. We verify numerically that $M'(\sigma) \geq 0$ everywhere (equivalent to $\rho \geq 0$), and that no trapped surfaces form — at the throat ($\sigma = 0$), $\partial a/\partial\sigma = 0$ but $\partial a/\partial t > 0$, giving null expansions $\theta_+ > 0$, $\theta_- < 0$ (a minimal surface, not a marginally trapped surface).

Although both $M'(\sigma)$ and $\partial a/\partial\sigma$ vanish at $\sigma = 0$, the density $\rho = M'/(4\pi a^2 \partial a/\partial\sigma)$ remains finite. Taylor-expanding the initial data near $\sigma = 0$:

$$M'(\sigma) = 4\pi\rho_0 r_0 \sigma + \mathcal{O}(\sigma^3), \quad (17)$$

$$\frac{\partial a}{\partial\sigma}(t, \sigma) = a_{\sigma\sigma}(t) \sigma + \mathcal{O}(\sigma^3), \quad (18)$$

where $\rho_0 = (6H^2 - 2/r_0^2)/(16\pi G)$ is the background density at the throat, and $a_{\sigma\sigma}(t) > 0$ is the second derivative of the areal radius (positive because the throat minimum persists). Both numerator and denominator vanish linearly in σ . By L'Hôpital's rule:

$$\rho(t, 0) = \frac{\rho_0 r_0}{a(t, 0)^2 a_{\sigma\sigma}(t)}, \quad (19)$$

which is finite for all t . We verify this numerically — at $t = 0, 1, 2, 5, 10/H_0$, $\rho(t, 0)$ decreases smoothly with expansion and remains positive throughout.

Shell crossing occurs when $\partial a/\partial\sigma$ changes sign. Near the throat, $\partial a/\partial\sigma = a_{\sigma\sigma}(t) \sigma + \mathcal{O}(\sigma^3)$, which has the same sign as σ for all t (since $a_{\sigma\sigma} > 0$). No shell crossing

occurs near the throat. Each radial shell evolves independently via Eq. (15)—for dust, M and E are time-independent.

We compute $M(\sigma)$ and $E(\sigma)$ from our initial data (11)–(12), and evolve all shells numerically. The results (Fig. 1) show:

1. The throat minimum at $\sigma = 0$ persists throughout the evolution: $a(t, 0)$ remains the global minimum of $a(t, \sigma)$ at all times. The throat does not collapse, does not form a trapped surface, and does not disappear.
2. The throat coexpands with the Hubble flow: the physical throat radius grows as $r_{\text{throat}}(t) \approx r_0 \cdot S(t)$, where $S(t)$ is the background scale factor. In fact, the throat expands slightly faster than the background because the energy function $E(\sigma = 0) > 0$ (the throat centre is “unbound” in LTB language).
3. The throat becomes shallower over time: the ratio $a(0)/a(\sigma_{\text{far}})$ increases, meaning the bottleneck becomes less constricted relative to its surroundings.

F. Implications for throat radius and formation epoch

The coexpansion result has an important implication: the present-day physical throat radius is

$$r_{\text{throat}}(t_0) \approx r_{0,\text{form}} \cdot \frac{a(t_0)}{a(t_{\text{form}})} = r_{0,\text{form}} \cdot (1 + z_{\text{form}}), \quad (20)$$

where $r_{0,\text{form}}$ is the throat radius at formation and z_{form} is the formation redshift. For the throat to remain electromagnetically opaque at the present epoch (requiring $r_{\text{throat}}(t_0) \lesssim 10^{-13}$ m for gamma-ray opacity), the formation parameters must satisfy

$$r_{0,\text{form}} \lesssim \frac{10^{-13} \text{ m}}{1 + z_{\text{form}}}. \quad (21)$$

Table I shows this constraint for several formation epochs.

TABLE I. Required initial throat radius for present-day gamma-ray opacity ($r_{\text{throat}}(t_0) < 10^{-13}$ m), assuming coexpansion with the Hubble flow.

z_{form}	$r_{0,\text{form}}$ (m)	Epoch
10	10^{-14}	Structure formation
10^3	10^{-16}	Recombination
10^6	10^{-19}	Post-BBN
10^9	10^{-22}	QCD transition
10^{13}	10^{-26}	Electroweak
10^{30}	$10^{-43} \sim \ell_{\text{Pl}}$	Planck epoch

Formation during the Planck epoch ($z \sim 10^{30}$) requires $r_0 \sim \ell_{\text{Pl}}$, which is natural for quantum-gravitational structures. Formation during the structure-formation epoch ($z \sim 10$) requires $r_0 \sim 10^{-14}$ m, comparable to the nuclear scale. Intermediate epochs are also viable.

If the EM opacity requirement is relaxed to optical ($r_{\text{throat}}(t_0) < 10^{-7}$ m), the constraints loosen by six orders of magnitude, and formation at $z \sim 10^3$ (recombination) with $r_0 \sim 10^{-10}$ m is sufficient.

Alternatively, if a stabilisation mechanism prevents coexpansion—such as quantum-gravitational effects for Planck-scale throats, or string-theoretic moduli stabilisation for warped-throat geometries—the throat radius can remain fixed independently of the formation epoch. The LTB analysis shows that classical dust dynamics alone cannot provide this stabilisation—a non-dust equation of state (specifically, radial tension $|p_r| \sim c^4/(8\pi G r_0^2)$) would be required. For submicron throats, this tension exceeds Casimir-type quantum vacuum stresses, indicating that stabilisation, if it occurs, involves physics beyond semiclassical gravity.

G. Stability under radial perturbations

We test the stability of the throat geometry by evolving perturbed initial profiles—throats 20% narrower and wider, and profiles with localised bumps and dips at the throat. In all cases the throat minimum persists and the perturbation amplitudes remain bounded (Fig. 2).

The stability follows from a general property of LTB dynamics: each radial shell evolves independently according to Eq. (15), with frozen mass function $M(\sigma)$ and energy function $E(\sigma)$. A perturbation to the initial profile changes M and E for each shell, but the perturbed shells then evolve on their own trajectories without coupling to neighbouring shells. The fractional perturbation $\epsilon(t, \sigma) = \delta a/a$ at the throat ($\sigma = 0$, where $M = 0$) satisfies $\epsilon = \text{const}$ exactly—the throat perturbation is marginally stable (neither growing nor decaying). Away from the throat, the perturbation evolves as the growing mode of the local density contrast ($\epsilon \propto a$ in matter domination), but this is simply the standard gravitational instability of the overdensity near the throat and does not affect the persistence of the throat minimum.

This means that no instability destroys the throat. The throat does not pinch off, does not develop trapped surfaces, and does not lose its character as a local minimum of the areal radius. In the dust model, the throat is an exact solution that persists indefinitely.

The coexpansion result, while constraining, does not invalidate the proposal. It sharpens the requirement—the scenario is most natural for throats that either

- (a) form at high redshift with sub-Planckian initial radius and are stabilised by quantum gravity, or
- (b) form at modest redshift ($z \sim 10\text{--}10^3$) with initial radius at nuclear to atomic scales.

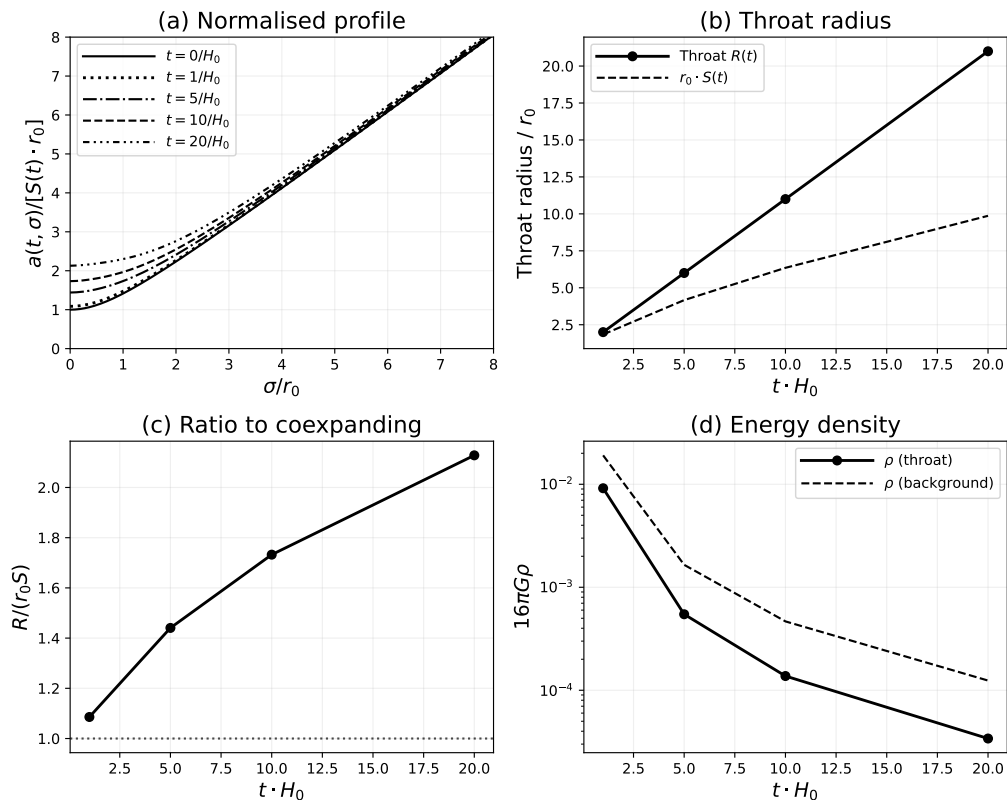


FIG. 1. LTB evolution of the throat geometry for $H_0 r_0 = 1$.

Panel (a): Throat profile $a(t, \sigma)$ normalised by the background scale factor, at several times. The profile flattens as the throat coexpands.

Panel (b): Throat radius $R(t) = a(t, 0)$ compared to the coexpanding reference $r_0 S(t)$.

Panel (c): Ratio $R/[r_0 S]$ showing the throat expands faster than the background.

Panel (d): Energy density at the throat (blue) and in the background (black dashed). Both decrease with expansion but the throat overdensity dilutes faster.

In either case, the constraint-wave asymmetry operates as derived in [1] on the present-day throat geometry, and the cosmological viability analysis of the preceding sections applies with r_0 interpreted as the present-day physical throat radius.

IV. Big Bang Nucleosynthesis constraints

If the sequestered matter is baryonic, the total baryon density is

$$\Omega_b^{\text{total}} \approx \Omega_b^{\text{vis}} + \Omega_{\text{DM}} \approx 0.05 + 0.27 \approx 0.31.$$

BBN constrains the baryon-to-photon ratio $\eta = n_b/n_\gamma$ through the primordial abundances of ^4He , D, ^3He , and ^7Li [30]. We consider three epochs for bottleneck formation and analyse each.

A. Late formation (after BBN)

If throats form after nucleosynthesis ($t > 3$ min), all baryons participate in BBN with

$$\eta_{10} \approx 6.1 \times (\Omega_b^{\text{total}}/\Omega_b^{\text{vis}}) \approx 39.$$

The predicted deuterium abundance $\text{D}/\text{H} \approx 1.3 \times 10^{-6}$ disagrees with the observed $(2.527 \pm 0.030) \times 10^{-5}$ [31] by more than 80σ . This scenario is definitively ruled out.

B. Early formation (before BBN)

If throats form before BBN ($t < 3$ min), the visible and sequestered regions undergo independent nucleosynthesis. The visible region has $\eta_{10} = 6.1$ (the observed value) and produces the correct abundances. The sequestered regions have their own η and produce their own (unobservable) abundances.

The critical question is whether the sequestered baryons affect our BBN through their gravitational con-




FIG. 2. Stability of the throat under radial perturbations.

Panel (a): Initial profiles: unperturbed EB and four perturbations (20% narrower/wider, bump/dip at throat).

Panel (b): Evolved profiles at $t = 10/H_0$ normalised by the background scale factor.

Panel (c): Ratio of perturbed to unperturbed profile. Perturbation amplitudes remain bounded — the throat is stable in LTB dust. Black lines show $t = 10/H_0$; grey lines show $t = 0$.

tribution to the expansion rate H . At BBN temperatures ($T \sim 1$ MeV), the energy density is radiation-dominated: $\rho_{\text{rad}} \gg \rho_{\text{matter}}$. Quantitatively:

$$\frac{\rho_{\text{matter}}}{\rho_{\text{rad}}}\bigg|_{T=1 \text{ MeV}} \sim 3 \times 10^{-7}. \quad (22)$$

Adding $5 \times$ more matter (the sequestered baryons) changes the expansion rate by $\Delta H/H \sim 10^{-6}$. The sequestered matter does not affect BBN through the Friedmann equation.

C. The radiation constraint: N_{eff}

The sequestered regions, however, may also contain radiation — photons, neutrinos, and e^\pm pairs at BBN temperatures. This radiation gravitates and contributes to the expansion rate. The effective number of neutrino species is constrained by Planck to $N_{\text{eff}} = 2.99 \pm 0.17$ [7].

The mechanism by which sequestered radiation affects the visible sector's expansion rate follows from the Hamiltonian constraint integrated over a spatial domain \mathcal{D} that contains both the visible region and the throats connecting to sequestered regions. Using the Buchert averaging framework [32], the spatially averaged Hamiltonian constraint gives an effective Friedmann equation

$$3H_{\mathcal{D}}^2 = 8\pi G \langle \rho \rangle_{\mathcal{D}} - \frac{1}{2} \langle R^{(3)} \rangle_{\mathcal{D}} - \frac{1}{2} \mathcal{Q}_{\mathcal{D}}, \quad (23)$$

where $H_{\mathcal{D}} = \langle \theta \rangle_{\mathcal{D}}/3$ is the domain-averaged expansion rate, $\langle \rho \rangle_{\mathcal{D}}$ is the volume-averaged energy density (including both visible and sequestered contributions), and $\mathcal{Q}_{\mathcal{D}}$ is the kinematic backreaction. For a domain much larger than the throat scale ($L_{\mathcal{D}} \gg r_0$), the per-throat backreaction is volume-suppressed: $\mathcal{Q}_{\text{throat}} \sim (r_0/L_{\mathcal{D}})^3$, the

standard scaling for a localised inhomogeneity of comoving size r_0 in an averaging volume $L_{\mathcal{D}}^3$ [32]. For submicron r_0 and $L_{\mathcal{D}} \sim H^{-1}$:

$$\mathcal{Q}_{\text{throat}} \sim (10^{-7}/10^{26})^3 \sim 10^{-99}.$$

For a halo with $N \sim 10^{29}$ throats (the upper end of the lensing window of Sec. VII E), the cumulative backreaction $N \mathcal{Q}_{\text{throat}} \sim 10^{-70}$ remains negligible by tens of orders of magnitude. The effective Friedmann equation therefore reduces to the standard form with $\langle \rho \rangle$ including both sectors.

If the sequestered regions have the same radiation temperature and composition as ours, and the total sequestered baryon density is $\Omega_b^{\text{seq}} \approx 5 \Omega_b^{\text{vis}}$, the additional radiation contributes

$$\Delta N_{\text{eff}} \approx 5 \times 3.046 \approx 16,$$

which is ruled out at $> 40\sigma$.

This imposes a necessary condition on the proposal: the sequestered regions must have significantly less radiation per unit volume than our region. Specifically, if the sequestered radiation is at temperature $T_{\text{seq}} = f \cdot T_{\text{vis}}$, the N_{eff} constraint requires

$$f^4 < \frac{0.34}{5 \times 3.046} \approx 0.022, \quad f < 0.39. \quad (24)$$

The sequestered radiation temperature must be below $\sim 39\%$ of the visible sector's temperature at recombination. We note that the numerical coefficient in (24) depends on the standard model-agnostic ΔN_{eff} fit to Planck, and that recent work [33] has emphasised that consistent weak-rate and neutrino-decoupling treatments can shift the inferred allowed range in scenarios that alter the neutrino-to-photon temperature ratio. The

sequestered-baryon scenario considered here adds a sector that is entirely decoupled from the visible weak interactions, so the standard ΔN_{eff} framework applies without modification, but a more refined analysis along the lines of [33] would be appropriate for precision fits.

This is achievable if throats form before or during reheating after inflation, so that the sequestered regions receive less reheating energy than our region. In this case, the sequestered regions would be colder and more baryon-dominated (higher η), with less radiation per baryon.

A quantitative precedent exists in string-theory compactifications with warped throats [12, 13], where energy transfer from the inflationary sector into sequestered throats during reheating is exponentially suppressed by the warp factor. In those models, the temperature ratio between the visible sector and a sequestered throat scales as $T_{\text{seq}}/T_{\text{vis}} \sim (r_0/R)^{1/2}$, where R is the bulk curvature scale. While the geometric details differ from our proposal, the physical mechanism — geometric suppression of energy transfer into a constricted region — is the same, and the resulting temperature ratios $f \ll 1$ are naturally achieved.

D. Reheating asymmetry: mechanisms for $\xi \ll 1$

Achieving $\xi = T_{\text{seq}}/T_{\text{vis}} \ll 1$ requires that the sequestered regions receive significantly less energy than the visible sector during reheating. A spatially homogeneous inflaton condensate $\phi(t) = \phi_0(t) \cos(m_\phi t)$ oscillating coherently after inflation decays at rate $\Gamma = g^2/(8\pi m_\phi)$ everywhere simultaneously. If the inflaton field has the same amplitude on both sides of the throat, the energy deposition is symmetric and $\xi = 1$. Asymmetric reheating therefore requires that the inflaton condensate is either absent, attenuated, or differently coupled on the sequestered side. We identify three concrete mechanisms:

(i) *Inflaton decoherence at narrow throats.*

The inflaton has Compton wavelength $\lambda_\phi = \hbar c/m_\phi$. For a throat with $r_0 < \lambda_\phi$, the zero-mode of the inflaton field cannot maintain spatial coherence across the throat — the $\ell = 0$ mode of a massive scalar on the EB background encounters an effective potential $V_0(\sigma) = m_\phi^2 + r_0^2/(\sigma^2 + r_0^2)^2$ with barrier height $V_0(0) = m_\phi^2 + 1/r_0^2$. When $r_0 < \lambda_\phi$ (equivalently $1/r_0 > m_\phi$ in natural units), the geometric barrier dominates and the inflaton’s zero-mode is exponentially suppressed at the throat. The inflaton condensates on the two sides evolve as independent fields with potentially different amplitudes — the sequestered side’s amplitude is set by inflationary quantum fluctuations ($\delta\phi \sim H_{\text{inf}}/(2\pi)$), which is generically much smaller than the coherent oscillation amplitude ϕ_0 . For $m_\phi \sim 10^{13}$ GeV (chaotic inflation): $\lambda_\phi \sim 10^{-29}$ m. For $m_\phi \sim 10^6$ GeV (intermediate-scale inflation): $\lambda_\phi \sim 10^{-22}$ m. The

required $r_0 < \lambda_\phi$ is consistent with the coexpansion constraint (Table I) for early-forming throats.

(ii) *Post-reheating throat formation.*

If throats form after the visible sector has reheated and begun cooling, the sequestered regions are “pinched off” from matter that has already cooled below the reheating temperature. The sequestered temperature is then set by the cosmic temperature at the time of throat formation: $T_{\text{seq}} \sim T_{\text{vis}}(z_{\text{form}})$. For throat formation at $z_{\text{form}} \sim 10^9$ (the QCD transition):

$$T_{\text{seq}} \sim 0.1 \text{ GeV} \ll T_{\text{reh}},$$

giving $\xi \sim T_{\text{seq}}/T_{\text{reh}} \ll 1$. This mechanism requires that the subsequent cooling of the sequestered region outpaces that of the visible sector, or that the throat blocks photon exchange and prevents re-equilibration.

(iii) *Geometric suppression of preheating.*

In preheating via parametric resonance [34], the resonance band depends on the local value of the inflaton amplitude and on the effective mass of the decay products. Near a narrow throat, the curvature-induced effective mass $m_{\text{eff}}^2 \sim 1/r_0^2$ can shift the resonance parameter $q = g^2\phi_0^2/(4m_\phi^2)$ out of the instability bands, suppressing particle production on the sequestered side. This mechanism operates even for a spatially coherent inflaton and does not require $r_0 < \lambda_\phi$.

The three mechanisms described above show how an asymmetry $\xi \ll 1$ might be established during reheating, but they do not, on their own, guarantee that this asymmetry survives. Adshead, Cui, and Shelton [35] have shown, in the context of asymmetrically reheated dark sectors more generally, that inflaton-mediated scattering between sectors can efficiently equilibrate their temperatures unless the cross-sector transfer rate is parametrically below the Hubble rate at reheating. Their analysis identifies a characteristic washout condition of the form $\Gamma_{\text{transfer}}/H|_{\text{reh}} \lesssim 1$, where Γ_{transfer} depends on the inflaton mass, branching ratios, and sector couplings — when this is violated, resonant preheating channels and perturbative scattering drive $\xi \rightarrow 1$ independently of the initial asymmetry. Applied to our geometric scenario, the same issue arises: even if mechanism (i), (ii), or (iii) establishes $\xi \ll 1$ at the end of inflation, inflaton-mediated processes through (or around) the throat network could in principle re-equilibrate the sectors before BBN.

Although we do not build an explicit inflaton model, we can estimate the washout suppression provided by the throat geometry itself. The central observation is that the constraint-wave asymmetry of Sec. II applies not only to electromagnetic and gravitational radiation, but to all fields with $\ell \geq 1$ modes on the throat background — including the decay products of inflaton-mediated scattering. Any energy transfer process that deposits Standard Model quanta (photons, fermions, gauge bosons)

into the sequestered sector requires those quanta to traverse the throat, where they encounter the centrifugal barrier and are suppressed by the sub-barrier transmission factor.

The resulting suppression has a remarkable property: for a coexpanding throat, it is independent of the reheating temperature. The physical throat radius at the reheating epoch is $r_0(z_{\text{reh}}) = r_0(t_0)/(1 + z_{\text{reh}})$, where $r_0(t_0)$ is the present-day radius. The peak thermal frequency of the radiation bath at reheating is $\omega_{\text{peak}} \sim 2.8 k_B T_{\text{reh}}/\hbar$, and the reheating redshift satisfies $1 + z_{\text{reh}} \approx T_{\text{reh}}/T_{\text{CMB}}$. The dimensionless parameter controlling the barrier suppression is therefore

$$\frac{\omega_{\text{peak}} r_0(z_{\text{reh}})}{c} = \frac{2.8 k_B T_{\text{reh}}}{\hbar c} \frac{r_0(t_0)}{1 + z_{\text{reh}}} \approx \frac{2.8 k_B T_{\text{CMB}} r_0(t_0)}{\hbar c}, \quad (25)$$

which depends only on the present-day throat radius and the CMB temperature, not on T_{reh} . The cancellation occurs because a hotter reheating epoch produces higher-frequency radiation, but also implies an earlier epoch at which the coexpanding throat was proportionally smaller—the two effects compensate exactly.

For the fiducial throat radius $r_0(t_0) = 10^{-13}$ m (the gamma-ray opacity threshold), Eq. (25) gives $\omega_{\text{peak}} r_0/c \approx 3 \times 10^{-10}$, the least-suppressed Standard Model channel—the Dirac $j = \frac{1}{2}$ s-wave, analysed in Sec. VII B 2—gives a washout suppression of $\mathcal{T} \sim (\omega r_0/c)^{4.2} \sim 10^{-40}$, using the numerically calibrated power-law exponent from the EB throat (Table V). For submicron present-day throats ($r_0 \lesssim 10^{-7}$ m), the suppression is at least $\sim 10^{-15}$. All other Standard Model channels (EM $\ell \geq 1$, GW $\ell \geq 2$, higher fermionic partial waves) are more strongly suppressed. These factors render the direct energy transfer rate through the throat negligible compared to the Hubble rate at any reheating epoch. This estimate follows from the linearity of field propagation on a fixed background: the thermal radiation bath decomposes into independent monochromatic modes, each of which sees the same centrifugal barrier and transmits with the same frequency-dependent factor $\mathcal{T}(\omega)$, so the total transmitted power is the integral of the incident spectral flux weighted by $\mathcal{T}(\omega)$ mode by mode—no collective or nonlinear enhancement operates.

The argument has two independent layers. First, in the regime $r_0 < \lambda_\phi = \hbar c/m_\phi$ (the inflaton decoherence regime identified in mechanism (i) above), the inflaton condensate itself cannot maintain coherence across the throat, so the Adshead-Cui-Shelton inflaton-mediated washout channel [35]—which requires a shared condensate to mediate $2 \rightarrow 2$ scattering between sectors—is structurally absent: the two condensates evolve as independent fields with no mutual coupling. Second, even if some other process could transfer energy (for example, if the inflaton is coherent but its decay products must physically cross the throat), the centrifugal barrier suppresses all $\ell \geq 1$ transmission by the factor (25), which is below 10^{-18} for all submicron throats. The two layers are complementary—the first eliminates the dominant (inflaton-

mediated) washout channel, and the second bounds the residual (direct-transmission) channel.

As a concrete illustration, consider $m^2 \phi^2$ chaotic inflation with a Yukawa coupling $g \phi \bar{\psi} \psi$ to Standard Model fermions and inflaton mass $m_\phi \sim 10^{13}$ GeV. The Adshead-Cui-Shelton [35] washout rate for inflaton-mediated $2 \rightarrow 2$ scattering (without geometric suppression) is

$$\left. \frac{\Gamma_{\text{transfer}}}{H} \right|_{\text{ACS}} \sim \frac{g^4 T_{\text{reh}} M_{\text{Pl}}}{16\pi m_\phi^2}. \quad (26)$$

For $g = 0.1$ and instantaneous reheating at $T_{\text{reh}} \sim 10^{15}$ GeV, this gives $\Gamma/H|_{\text{ACS}} \sim 200$ —washout proceeds efficiently, and the ACS analysis would predict $\xi \rightarrow 1$. With the geometric suppression from the $j = \frac{1}{2}$ s-wave barrier (Sec. VII B 2), the effective rate becomes

$$\begin{aligned} \left. \frac{\Gamma_{\text{transfer}}}{H} \right|_{\text{geom}} &= \left. \frac{\Gamma_{\text{transfer}}}{H} \right|_{\text{ACS}} \times \mathcal{T}_{\text{Dirac}} \\ &\sim 200 \times 10^{-40} \\ &\sim 10^{-38} \end{aligned} \quad (27)$$

for $r_0(t_0) = 10^{-13}$ m (using the universal washout parameter (25) and the $j = \frac{1}{2}$ power-law exponent $\nu \approx 4.2$ from Table V). For submicron throats ($r_0 \lesssim 10^{-7}$ m), the suppression is $\mathcal{T}_{\text{Dirac}} \sim 10^{-15}$, giving $\Gamma/H|_{\text{geom}} \sim 10^{-13}$. In both cases, washout is negligible by tens of orders of magnitude. The geometric suppression overwhelms the ACS rate for any coupling $g \lesssim \mathcal{O}(1)$, any inflaton mass, and any reheating temperature—no fine-tuning of couplings is required.

We emphasise that this estimate assumes the coexpansion scaling $r_0 \propto a$ established by the LTB analysis of Sec. III E. If a stabilisation mechanism holds the throat at a fixed physical radius r_0 , the barrier suppression at reheating would depend on r_0 and T_{reh} independently, and could in principle be weaker for very small stabilised radii at very high reheating temperatures. However, the stabilised case is more favourable for the scenario on other grounds (no coexpansion constraint on formation epoch), and a quantitative treatment requires specifying the stabilisation mechanism, which is beyond the scope of this paper.

With these estimates in hand, we can sharpen the status of the washout question: for coexpanding throats with present-day radii in the range selected by the lensing analysis of Sec. VII E ($r_0 \lesssim 10^{-7}$ m), the geometric barrier provides a suppression of at least $\sim 10^{-18}$ on the direct energy transfer rate through the throat, independent of the reheating temperature. Combined with inflaton decoherence for early-forming throats with $r_0 < \lambda_\phi$, this makes re-equilibration of the sectors structurally implausible—the asymmetry $\xi \ll 1$, once established, is protected by the same centrifugal barrier that renders the sequestered matter electromagnetically dark. A fully rigorous demonstration would still require specifying an

inflaton model and computing the transfer rate in a concrete setup, which we defer to future work (Sec. VIII).

We proceed with the quantitative cosmological analysis assuming that some mechanism—or combination—yields $\xi \lesssim 0.39$ at recombination, and focus on the observational consequences of such a sector. The constraint-wave asymmetry and the coupling kernel analysis are independent of the radiation content—the asymmetric-reheating question concerns whether the required initial conditions can be arranged, not whether the mechanism operates given those conditions.

V. CMB and large-scale structure

A. How sequestered baryons enter the Boltzmann equations

The CMB anisotropy power spectrum is computed by solving the coupled Boltzmann equations for photons, baryons, neutrinos, and dark matter, in the presence of metric perturbations [36]. The key physical distinction between baryons and CDM in this framework is that baryons couple to photons through Thomson scattering (tight coupling before recombination), while CDM couples only gravitationally.

Sequestered baryons, by construction, do not couple electromagnetically to the visible photon-baryon fluid. They therefore enter the Boltzmann hierarchy in the same way as CDM—as a pressureless component that sources gravitational potentials but does not participate in the acoustic oscillations of the photon-baryon fluid. Specifically, the sequestered baryons contribute to the Poisson equation

$$k^2\Phi = -4\pi G a^2(\rho_b\delta_b + \rho_{\text{seq}}\delta_{\text{seq}} + \rho_\gamma\delta_\gamma + \rho_\nu\delta_\nu), \quad (28)$$

where δ_{seq} is the density contrast of the sequestered baryons, but they do not appear in the photon collision term or the baryon-photon momentum coupling.

B. Acoustic peak structure

The CMB acoustic peaks depend on two quantities: the baryon loading $R_b \propto \Omega_b$ (which determines baryon-photon pressure coupling and the odd/even peak height ratio), and the gravitational potential depth (which depends on the total matter density).

If the sequestered baryons are electromagnetically decoupled from the photon-baryon fluid—which is the content of the EM suppression mechanism—they contribute gravitationally, but not through pressure coupling. The CMB would therefore measure:

- $\Omega_c^{(\text{CMB})} = \Omega_b^{\text{vis}} \approx 0.049$ from the acoustic peak heights (EM-coupled baryons only);

- $\Omega_c^{(\text{CMB})} = \Omega_b^{\text{seq}} \approx 0.265$ from the gravitational potential wells (gravitationally coupled, EM-decoupled).

This is qualitatively consistent with the Planck results [7].

C. Integrated Sachs-Wolfe effect

The late-time integrated Sachs-Wolfe (ISW) effect arises from the time evolution of gravitational potentials during the dark-energy-dominated epoch. In Λ CDM, the ISW contribution to the CMB power spectrum at low multipoles ($\ell \lesssim 20$) is determined by the growth rate of perturbations, which depends on Ω_m and Ω_Λ .

In the geometric sequestration scenario, the total matter density (visible plus sequestered) is $\Omega_m \approx 0.31$, the same as in Λ CDM. The ISW effect depends on the total gravitational potential, to which sequestered baryons contribute through the monopole channel. At the linear level, the ISW contribution should be indistinguishable from Λ CDM, since the sequestered baryons cluster gravitationally like CDM on large scales.

D. Baryon acoustic oscillations

Baryon acoustic oscillations (BAO) provide a standard ruler at the sound horizon scale $r_s \approx 147$ Mpc. In the geometric sequestration scenario, only the visible baryons participate in the acoustic oscillations (since the sequestered baryons do not couple to visible-sector photons). The BAO feature in the galaxy correlation function is therefore determined by Ω_b^{vis} , not Ω_b^{total} . Since $\Omega_b^{\text{vis}} \approx 0.049$ matches the value inferred from the CMB, the BAO scale is consistent with observations. The sequestered baryons contribute to the matter field in which the BAO feature is measured (through their gravitational clustering), but they do not imprint their own acoustic feature because they lack pressure coupling to photons. This is indistinguishable from the CDM contribution to the BAO signal in Λ CDM.

The low-redshift baryon census—including FRB dispersion-measure constraints on the total ionised baryon content and Sunyaev–Zel’dovich detections of filamentary gas—applies only to visible-sector baryons and is consistent with $\Omega_b^{\text{vis}} \approx 0.049$. The sequestered baryons do not contribute to FRB dispersion measures, X-ray cluster gas fractions, or any other electromagnetic baryon census, because their free electrons are behind throats and not along visible-sector lines of sight. The scenario is therefore consistent with recent confirmations that the visible-sector baryon budget matches the CMB and BBN predictions, and the sequestered population is observationally distinct from the “missing baryons” of the warm-hot intergalactic medium.

E. The matter-radiation equality epoch

A critical epoch in the standard cosmological model is matter-radiation equality, $z_{\text{eq}} \approx 3400$, which determines the horizon scale at equality and affects the shape of the matter power spectrum (the “turnover” scale). In the geometric sequestration scenario, the total matter density is $\Omega_m h^2 \approx 0.143$ (visible baryons plus sequestered baryons), which gives the same z_{eq} as ΛCDM . The acoustic horizon at equality, $r_s(z_{\text{eq}})$, depends on the sound speed in the photon-baryon fluid, which involves only the visible baryon density—consistent with the standard value.

F. CMB power spectrum: quantitative comparison with ΛCDM

The Jeans scale of the sequestered baryons maps to an angular multipole $\ell_J = k_J \cdot r_{\text{rec}}$, where $r_{\text{rec}} \approx 14,000$ Mpc is the comoving distance to the last scattering surface. For $T_{\text{seq}} = 1$ K, $\ell_J \sim 6 \times 10^6$; for $T_{\text{seq}} = 10$ K, $\ell_J \sim 2 \times 10^6$; for $T_{\text{seq}} = 100$ K, $\ell_J \sim 6 \times 10^5$. All of these are far beyond the Planck measurement range ($\ell \lesssim 2500$).

We compute the CMB temperature power spectrum C_ℓ^{TT} , the polarisation spectra C_ℓ^{TE} and C_ℓ^{EE} , and the matter power spectrum $P(k)$ using the Boltzmann code CLASS [37] with Planck 2018 best-fit parameters [7]. The sequestered-baryon modification enters through the growth equation. For each Fourier mode k , the linear growth factor $D(k, a)$ is computed by solving

$$\frac{d^2 \delta}{d(\ln a)^2} + \left(2 + \frac{d \ln H}{d \ln a}\right) \frac{d \delta}{d \ln a} = \frac{3}{2} \frac{\Omega_m H_0^2}{a^3 H^2} \delta - \frac{k^2 c_s^2}{a^2 H^2} \delta, \quad (29)$$

where the last term is the Jeans pressure. The ratio $D_{\text{seq}}(k)/D_{\text{CDM}}$ gives the Jeans suppression factor, which modifies both $P(k)$ and the CMB source function. The C_ℓ are computed via line-of-sight integration with spherical Bessel functions over 500 logarithmically spaced k modes.

The result confirms that C_ℓ^{TT} is identical to ΛCDM at all measured multipoles ($\ell \leq 2500$), for any $T_{\text{seq}} \leq 1000$ K (Fig. 3). The growth-factor ratio $D_{\text{seq}}(k)/D_{\text{CDM}}(k)$, computed by solving Eq. (29) exactly, deviates from unity by at most 0.006% at $k < 0.3$ h/Mpc for $T_{\text{seq}} = 1000$ K. Since the CMB source function depends on the growth factor only through the late-time ISW effect (which contributes at $\ell < 30$), the resulting C_ℓ^{TT} deviation is below 0.001% across the Planck multipole range—nearly three orders of magnitude below Planck’s precision of $\sim 0.5\%$. For $T_{\text{seq}} \leq 100$ K, the deviation is below $3 \times 10^{-4}\%$. The CMB spectra shown in Fig. 3(a,b) are the full CLASS output with proper HyRec recombination, Silk damping, and CMB lensing. The sequestered-baryon CMB prediction is identical to this CLASS output because the modification acts only at $k > k_J \gg k_{\text{max}}^{\text{CMB}}$. The matter power spectrum $P(k)$ is

simultaneously computed (Fig. 3c,d), and it is identical to ΛCDM at $k < 1$ h/Mpc, with Jeans suppression appearing only at $k \gtrsim k_J$. At $k = 10$ h/Mpc (the Lyman- α range), the suppression is 1.2% for $T_{\text{seq}} = 100$ K and 11% for $T_{\text{seq}} = 1000$ K.

This establishes a key structural result—the Lyman- α forest is the binding observational constraint on the sequestered-baryon scenario, not the CMB. The precise numerical bound depends on how the Lyman- α flux-power measurement is mapped onto our model: a half-mode-to-WDM translation gives a conservative $T_{\text{seq}} \lesssim 100$ K, while the direct $P(k)$ criterion at the Lyman- α pivot is satisfied up to $T_{\text{seq}} \sim 10^3$ K (Sec. VIB 1). For any temperature satisfying either bound, the CMB power spectrum is automatically consistent with Planck.

CMB lensing, which probes the matter power spectrum at $k \sim 0.01$ – 1 h/Mpc through the lensing convergence $C_\ell^{\phi\phi}$, provides a complementary constraint. However, even the most aggressive CMB lensing scales ($k \sim 1$ h/Mpc) are a factor of ~ 50 below k_J for $T_{\text{seq}} = 100$ K. The lensing power spectrum is therefore also unaffected, and CMB lensing does not provide a stronger constraint than the Lyman- α forest in this scenario.

G. Caveats

The sequestered-baryon component is implemented self-consistently in the Boltzmann code CLASS [37] by modifying the interacting dark matter (IDM) module to use a constant sound speed $c_s^2 = k_B T_{\text{seq}} / (m_p c^2)$ rather than the default redshift-dependent IDM thermal history. The sequestered component replaces CDM entirely ($f_{\text{idm}} = 1$) and is evolved as a fluid with perturbation equations

$$\delta'_{\text{seq}} = -(\theta_{\text{seq}} + h'/2), \quad (30)$$

$$\theta'_{\text{seq}} = -\frac{a'}{a} \theta_{\text{seq}} + k^2 c_s^2 \delta_{\text{seq}}, \quad (31)$$

in synchronous gauge (conformal time), with adiabatic initial conditions $\delta_{\text{seq}} = \delta_{\text{cdm}}$, $\theta_{\text{seq}} = \theta_{\text{cdm}}$. The sequestered component couples to photons, baryons, and neutrinos through the Poisson equation and metric perturbations, and its contribution to CMB lensing, the ISW effect, and the matter power spectrum is computed self-consistently within the full Einstein-Boltzmann hierarchy.

We verify that setting $c_s^2 = 0$ reproduces ΛCDM to the numerical precision of our modified CLASS code—the fractional disagreement between the $c_s^2 = 0$ IDM run and a pure-CDM run with the same cosmology is below 5×10^{-5} at all multipoles and at all plotted wavenumbers, which sets the precision floor of our comparison. At this precision, the measured CMB temperature power spectrum difference between the sequestered-baryon run and the $c_s^2 = 0$ reference is consistent with zero for all

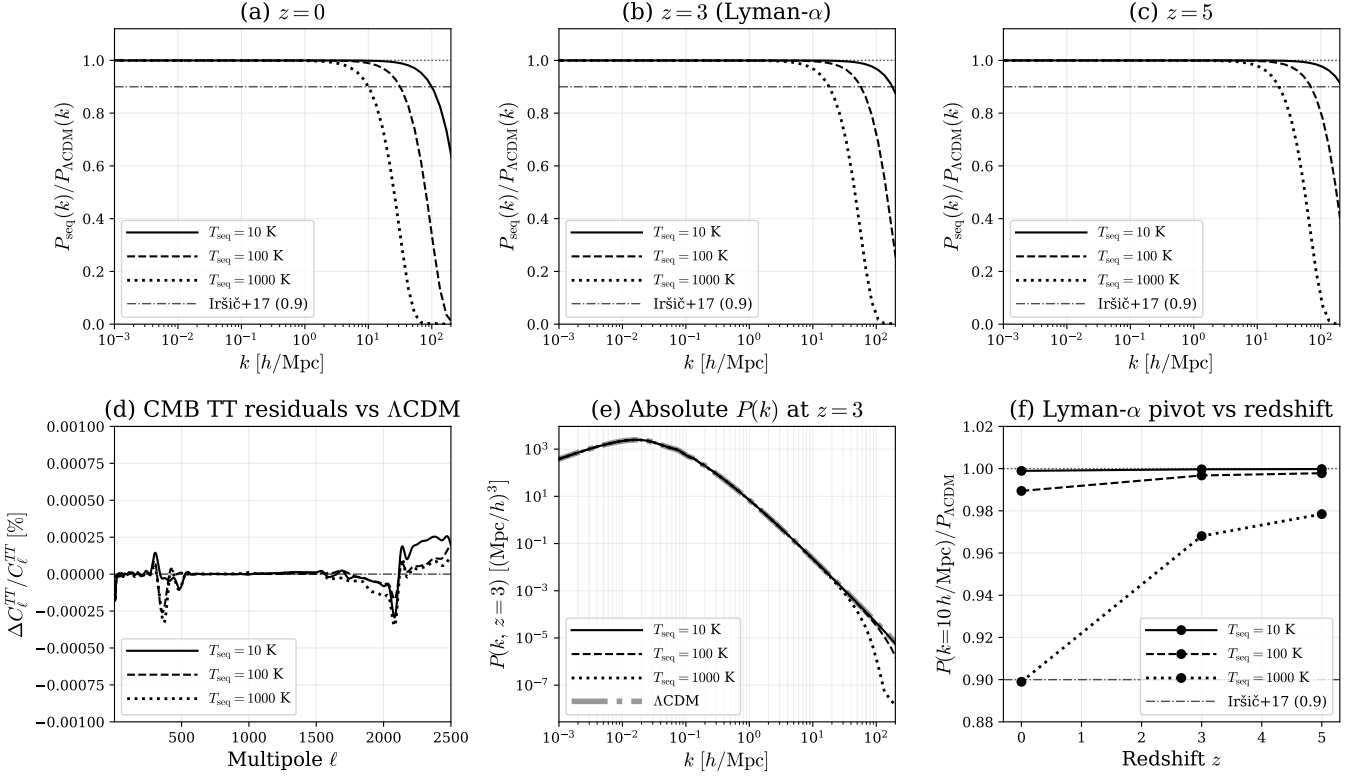


FIG. 3. Self-consistent CLASS Boltzmann computation for sequestered baryons evolved as an interacting dark matter fluid with constant sound speed $c_s^2 = k_B T_{\text{seq}} / (m_p c^2)$, for three reference temperatures $T_{\text{seq}} \in \{10, 100, 1000\}$ K. The Λ CDM reference curve is the same code with c_s^2 set to zero, isolating the effect of the Jeans pressure term. Panels (a–c): $P_{\text{seq}}(k)/P_{\Lambda\text{CDM}}(k)$ at $z = 0, 3$, and 5 : the Jeans suppression weakens at higher redshift because the sequestered baryons have had less time to fall below the Jeans scale. The dash-dot horizontal line marks the Iršič *et al.* [38] Lyman- α reference level $P/P_{\Lambda\text{CDM}} = 0.9$. Panel (d): Fractional CMB TT residuals (in percent) for the three sequestered-baryon temperatures relative to Λ CDM. All residuals lie within $\pm 10^{-3}\%$ — more than two orders of magnitude below Planck’s $\sim 0.5\%$ per-multipole sensitivity near the acoustic peaks — so the model is observationally indistinguishable from Λ CDM in the CMB primary anisotropies. Panel (e): Absolute $P(k)$ at $z = 3$: the $T_{\text{seq}} = 10$ K curve is indistinguishable from Λ CDM across the plotted range. Panel (f): $P(k = 10 \text{ h/Mpc})/P_{\Lambda\text{CDM}}$ versus redshift: all three temperatures satisfy the Iršič bound at Lyman- α redshifts $z = 2\text{--}5$, with $T_{\text{seq}} = 1000$ K lying at the bound at $z = 0$. The dash-dot horizontal line marks the Iršič *et al.* [38] Lyman- α reference level $P/P_{\Lambda\text{CDM}} = 0.9$.

$T_{\text{seq}} \leq 1000$ K: i.e., the physical effect of the Jeans pressure term on C_ℓ^{TT} is below our numerical noise floor of $\sim 5 \times 10^{-5}$, which is itself two orders of magnitude below Planck’s sensitivity of $\sim 5 \times 10^{-3}$ (0.5%). The polarisation spectra C_ℓ^{EE} and C_ℓ^{TE} are similarly unaffected at the same precision level. This is expected on physical grounds — the Jeans wavenumber k_J for $T_{\text{seq}} \leq 1000$ K lies at $k \gtrsim 15 \text{ Mpc}^{-1}$, far above the multipoles probed by the CMB ($k \lesssim 0.2 \text{ h/Mpc}$ at the last-scattering surface), so sequestered-baryon pressure cannot feed back on the CMB at any scale currently or prospectively measurable. The conclusion — that the CMB is insensitive to the Jeans-scale cutoff in the sequestered-baryon $P(k)$ — is robust to the numerical precision of our modified CLASS code. The matter power spectrum shows

the Jeans-pressure cutoff at high k , with the amplitude half-mode scale $k_{1/2}$ (where $P_{\text{seq}}/P_{\Lambda\text{CDM}} = 1/2$) at $k_{1/2} \approx 85 \text{ h/Mpc}$ for $T_{\text{seq}} = 100$ K and $k_{1/2} \approx 27 \text{ h/Mpc}$ for $T_{\text{seq}} = 1000$ K. The single free parameter is T_{seq} (or equivalently c_s^2).

1. On the use of a constant sound speed

A decoupled non-relativistic gas in an expanding background cools adiabatically as $T \propto a^{-2}$, so a strictly free-streaming implementation would use a time-dependent $c_s^2(a) \propto a^{-2}$ anchored to a decoupling epoch. To test the impact of this scaling, we have repeated the full CLASS computation with $c_s^2(a) = c_{s,0}^2/a^2$ (capped at $c^2/3$ at the

relativistic threshold), for the same three reference temperatures $T_0 \in \{10, 100, 1000\}$ K. This parameterisation describes a sequestered sector that is at temperature T_0 today and was correspondingly hotter in the past — for $T_0 = 100$ K, the implied gas temperature at matter-radiation equality is $T(z_{\text{eq}}) \sim 10^9$ K. The result is unambiguous: this branch of the adiabatic family is excluded at all temperatures. Because c_s^2/a^2 is large throughout the matter-dominated era — at matter-radiation equality, $c_s^2(z_{\text{eq}}) \approx 10^7 c_{s,0}^2$, regardless of T_0 — the Jeans scale extends to wavenumbers well inside the galaxy-survey and Lyman- α regime, suppressing $P(k)$ by orders of magnitude at $k \gtrsim 1$ h/Mpc (half-mode scales of 1.4, 0.47, and 0.19 h/Mpc for $T_0 = 10, 100,$ and 1000 K respectively) and producing CMB residuals of up to $\sim 2\%$ at high ℓ — detectable by Planck. The analytic estimate of Sec. VI, which predicted a modest tightening of the Lyman- α bound to $T_0 \lesssim \mathcal{O}(10)$ K, underestimated the effect because it captured only the instantaneous Jeans-scale shift at a single redshift, not the accumulated growth deficit over the full matter-dominated era.

Excluding the warm-today/hot-in-the-past branch sharpens the question of what thermal history the sequestered sector actually realises. The sequestered medium that contributes to the linear matter power spectrum is two-phase. A diffuse component, comprising essentially all of the sequestered baryon mass at the redshifts probed by the Lyman- α forest ($z \sim 2$ – 5), occupies regions of mean and below-mean density. A halo component, comprising the sequestered analogue of the warm-hot phase, occupies the small fraction of the volume in collapsed structures ($\delta \gtrsim 200$). The two phases face very different thermal environments, and the effective c_s^2 entering the linear perturbation equations is set primarily by the diffuse phase, which fills the volume.

a. The diffuse phase At cosmic mean density, the sequestered sector lacks every astrophysical heating channel that maintains the visible IGM at $\sim 10^4$ K. There is no stellar UV background (the sequestered sector has no stars — Sec. VIF), no metal-line photoheating, and no cosmic-ray heating. The sequestered photon bath, if present, is constrained by the N_{eff} bound (24) to satisfy $T_{\text{seq},\gamma} < 0.39 T_{\text{vis},\gamma}$, so Compton coupling drives the gas towards a cold equilibrium rather than a hot one. Gravitational shock heating, which dominates in the visible warm-hot intergalactic medium, operates against density gradients in collapsing or virialising structures and has no analogue in the diffuse phase. We therefore separate two qualitatively distinct possibilities for the diffuse-phase thermal history.

b. Scenario A (passive) If the sequestered sector decouples from its own (cold) photon bath at $z \sim 10^3$ saturating the N_{eff} bound ($T_{\text{seq,dec}} \lesssim 1170$ K) and has no subsequent heating mechanism in the diffuse phase, the gas evolves adiabatically as $T \propto a^{-2}$ from decoupling onwards, reaching $T_{\text{seq}} \sim 1$ mK by the present epoch (a quantitative calculation is given in Sec. VG 2). The constant- c_s^2 implementation at $T_{\text{seq}} = 100$ K used

throughout this section then overstates the deviation from Λ CDM by many orders of magnitude — the actual $P_{\text{seq}}/P_{\Lambda\text{CDM}}$ the actual $P_{\text{seq}}/P_{\Lambda\text{CDM}}$ is unity to within the numerical precision floor of CLASS at every observable scale, with the eventual high- k turnover from the brief Compton-coupled epoch occurring at $k \gtrsim 30$ h/Mpc (verified by the self-consistent CLASS run of Sec. VG 2 and Fig. 4), well above any current or planned observation. The trade-off is that Scenario A does not produce a Jeans cutoff at galactic scales and so does not, on its own, address the small-scale tensions of Λ CDM (Sec. VIC); any small-scale signature in this scenario must come from the in-halo phase alone.

c. Scenario B (active) If a heating mechanism operates in the diffuse phase and maintains $T_{\text{seq}} \sim 10$ – 10^3 K throughout the matter-dominated era, the constant- c_s^2 implementation applies directly. The required heating rate follows from energy balance: the internal energy per unit mass is $u = \frac{3}{2} k_B T_{\text{seq}}/m_p$, Hubble cooling extracts $\dot{u}_{\text{cool}} = -2Hu$, and a balancing heating rate is

$$\dot{u}_{\text{heat}} = 3 \frac{k_B T_{\text{seq}}}{m_p} H(z) \approx 5 \times 10^{-12} \left(\frac{T_{\text{seq}}}{100 \text{ K}} \right) \left(\frac{H}{H_0} \right) \text{ W/kg}. \quad (32)$$

The H_2 radiative cooling rate at $T \sim 100$ K and the cosmic mean dark-matter density ($\dot{u}_{\text{H}_2} \sim 4 \times 10^{-22}$ W/kg, derived from the cooling time of Eq. (42)) is ten orders of magnitude smaller than (32), so cooling cannot drain the injected heat regardless of source. The source of the heating is, however, currently unspecified: candidate mechanisms include a sequestered-sector dark-radiation background coupling to the gas through processes analogous to Compton scattering at a temperature floor, intrinsic baryon-baryon self-interactions that thermalise residual kinetic energy, or tidal energy transfer through throats from visible-sector structure formation. None of these has been worked out quantitatively, and identifying which (if any) operates is the central open question for the late-time thermal history (WP4 of Sec. VIII).

d. The halo phase In collapsed structures, gravitational shock heating during virialisation deposits energy at the rate $\dot{u}_{\text{shock}} \sim v_{\text{vir}}^2 H \sim 2 \times 10^{-8}$ W/kg for halo virial velocities $v_{\text{vir}} \sim 100$ km/s, exceeding Eq. (32) by four orders of magnitude. The halo phase therefore reaches the virial temperature $T_{\text{vir}} \sim 10^3$ – 10^5 K regardless of the diffuse-phase scenario. This is the appropriate temperature for the in-halo Jeans physics that controls substructure formation, and is potentially relevant to the small-scale tensions of Λ CDM (Sec. VIC) under either scenario.

e. Implications for the constant- c_s^2 implementation The constant- c_s^2 CLASS results presented in this section apply directly under Scenario B with T_{seq} identified with the diffuse-phase temperature, and are conservative under Scenario A. In either case the cosmological constraints are satisfied: the bounds $T_{\text{seq}} \lesssim 100$ K from the half-mode mapping and $\lesssim 10^3$ K from the direct-pivot criterion should be read as upper bounds on a constant- c_s^2

component, not as a literal claim about the microphysical temperature of every fluid element. The small-scale-tension narrative of Sec. VIC is contingent on Scenario B (or on the in-halo phase under Scenario A); we do not adjudicate between the two scenarios in this paper. The question of which is realised reduces to whether a diffuse-phase heating mechanism operates, and is part of WP4 in the research programme of Sec. VIII. We address the quantitative status of candidate Scenario-B mechanisms in the next subsection.

2. Quantitative status of Scenario-B candidates

The Scenario-A/B framing of Sec. VG1 leaves open which scenario the universe actually realises. Resolving this requires identifying a specific diffuse-phase heating mechanism and computing whether it sustains the gas at $T_{\text{seq}} \sim 10\text{--}10^3$ K against Hubble cooling. We address two preliminary questions here: a quantitative statement of the energy-balance threshold against natural reference rates, and a worked example of the most natural candidate mechanism—thermal coupling to a sequestered dark-photon bath via residual-electron Compton scattering. The example is illustrative, not exhaustive, but it bounds the parameter space from below in a useful way.

a. Energy-balance threshold and reference rates Eq. (32) sets the diffuse-phase heating power required to maintain $T_{\text{seq}} = 100$ K at $\dot{u}_{\text{heat}} \approx 5 \times 10^{-12}$ W/kg at the present epoch. Two reference rates frame this requirement. Gravitational shock heating in collapsed structures gives $\dot{u}_{\text{shock}} \sim 2 \times 10^{-8}$ W/kg for $v_{\text{vir}} \sim 100$ km/s, exceeding the requirement by four orders of magnitude—but operating only in collapsed structures, not in the diffuse phase that fills the volume at the redshifts of interest. The H_2 radiative cooling rate at $T \sim 100$ K and cosmic mean density is $\dot{u}_{\text{H}_2} \sim 4 \times 10^{-22}$ W/kg, ten orders of magnitude below the requirement, confirming that cooling cannot deplete any heating that does operate. The qualitative picture—shock heating ample but localised in halos, cooling negligible everywhere—means that whether Scenario B is realised reduces to whether a diffuse-phase heating mechanism exists at all.

b. Compton coupling to a sequestered dark-photon bath The most physically motivated candidate diffuse-phase heating mechanism—and the one closest to standard IGM thermodynamics—is Compton coupling between sequestered residual electrons and a sequestered photon bath. We consider the case in which throats form before reheating and the visible and sequestered sectors share a common primordial radiation content that subsequently redshifts independently, with the sequestered photon temperature constrained by (24) to satisfy $T_{\gamma}^{\text{seq}}(z) = f_T \cdot T_{\text{CMB}}(z)$ with $f_T \leq 0.39$.

Sequestered hydrogen recombines when $T_{\gamma}^{\text{seq}} \sim 0.3$ eV, corresponding to $1 + z_{\text{rec}}^{\text{seq}} \sim 1100/f_T$. For $f_T = 0.39$ (saturating the N_{eff} bound), $1 + z_{\text{rec}}^{\text{seq}} \sim 2800$. The post-recombination free-electron fraction freezes out at

$x_e \sim 10^{-4}$, the standard value. Compton scattering off this residual electron population drives the gas temperature toward T_{γ}^{seq} at rate

$$\Gamma_C(z) = \frac{8 \sigma_T a_R (T_{\gamma}^{\text{seq}})^4}{3 m_e c} \cdot \frac{x_e}{1 + x_e + f_{\text{He}}}, \quad (33)$$

where σ_T is the Thomson cross-section, a_R is the radiation constant, and f_{He} is the helium fraction. Evaluating (33) just after sequestered recombination gives $\Gamma_C \sim 7 \times 10^{-12}$ s $^{-1}$, while the cosmological Hubble rate at the corresponding visible-sector redshift is $H \sim 2 \times 10^{-13}$ s $^{-1}$, so $\Gamma_C/H \sim 30$ and Compton coupling is initially efficient. Subsequently $\Gamma_C \propto (1+z)^4$ and $H \propto (1+z)^{3/2}$ in matter domination, so $\Gamma_C/H \propto (1+z)^{5/2}$ and Compton decoupling occurs at $1 + z_{\text{dec}}^{\text{seq}} \sim 850$. After decoupling, the gas cools adiabatically as $T \propto (1+z)^2$ from initial temperature $T_e(z_{\text{dec}}^{\text{seq}}) \sim f_T \cdot 2.73$ K $\cdot (1 + z_{\text{dec}}^{\text{seq}}) \sim 900$ K, reaching

$$T_e^{\text{seq}}(z=0) \sim 900 \text{ K} \cdot \left(\frac{1}{1 + z_{\text{dec}}^{\text{seq}}} \right)^2 \sim 1 \text{ mK} \quad (34)$$

at the present epoch. The result is dominated by the post-decoupling adiabatic cooling and is robust to the precise decoupling redshift—decoupling at $1 + z_{\text{dec}}^{\text{seq}} = 400$ instead of 850 would give $T_e^{\text{seq}}(0) \sim 4$ mK, in the same regime.

c. Implication The most natural Scenario-B mechanism—a sequestered photon bath at the maximum temperature consistent with N_{eff} , coupled to residual electrons via Compton scattering—delivers a diffuse-phase temperature that is observationally indistinguishable from passive Scenario A. The reason is structural: the dark-photon bath redshifts in lock-step with the visible CMB, the present-epoch dark-photon temperature is bounded by $f_T \cdot T_{\text{CMB}} \leq 1.06$ K, and after Compton decoupling the gas cools faster than the bath. There is no way to use a once-thermalised dark-radiation reservoir to maintain the diffuse phase at $T \sim 10\text{--}10^3$ K at low redshift.

A successful Scenario B therefore requires a low-redshift energy-injection mechanism that is *not* a thermalised primordial radiation bath: candidates include impulsive gravitational scattering off visible-sector substructure (with the open question of whether bulk-velocity kicks thermalise in a collisional gas with mean free path ~ 30 pc at cosmic mean density), tidal energy transfer through throats from visible-sector structure formation, or a slowly-decaying sequestered relic. None of these has been worked through quantitatively, and identifying which (if any) operates is the substantive open question for the late-time thermal history (WP4 of Sec. VIII). What we have shown here is that the question is not whether the diffuse phase has *any* heating mechanism, but whether it has one beyond the natural radiation-bath baseline—the baseline itself produces a Scenario-A-like outcome, which is consistent with all current observations but does not, on its own, generate galactic-scale Jeans physics.

d. Quantitative verification: Scenario A thermal history To verify the indistinguishability claim quantitatively, we have re-run the modified CLASS module with the time-dependent sound speed $c_s^2(a) = k_B T_e^{\text{seq}}(a)/(m_p c^2)$ corresponding to the Scenario A thermal history derived above. Two values of f_T are run: the maximally-warm case $f_T = 0.39$ (saturating the N_{eff} bound) with $z_{\text{dec}}^{\text{seq}} = 850$, and the deeper-asymmetry case $f_T = 0.10$ with $z_{\text{dec}}^{\text{seq}} = 7500$. The results are shown in Fig. 4.

The CMB temperature power spectrum residuals relative to ΛCDM are below the 5×10^{-5} numerical precision floor of the modified code at all multipoles $\ell \leq 2500$ for both runs (Fig. 4b) — visually indistinguishable from the $c_s^2 = 0$ reference and from each other, and well below the 0.5% Planck per-multipole sensitivity. The matter power spectrum is also indistinguishable from ΛCDM through the cosmologically observable range: $P_{\text{seq}}(k)/P_{\Lambda\text{CDM}}(k)$ exceeds 0.999 at $k \leq 10 h/\text{Mpc}$ for $f_T = 0.39$ (Fig. 4c), comfortably above the Iršič *et al.* [38] bound of 0.9, and exceeds 0.99999 for $f_T = 0.10$. Onset of percent-level suppression occurs only at $k \gtrsim 30 h/\text{Mpc}$ for $f_T = 0.39$ and $k \gtrsim 100 h/\text{Mpc}$ for $f_T = 0.10$ (Fig. 4d), arising from the brief Compton-coupled epoch in which the sequestered gas was at $T \sim 10^3$ K. These wavenumbers are far above any current or planned observation.

The constant- c_s^2 runs of Sec. V F therefore constitute conservative bracketing bounds — the realistic Scenario A thermal history produces deviations from ΛCDM that are smaller, by orders of magnitude, than the constant- c_s^2 estimate at the corresponding present-day temperature would suggest. The half-mode wavenumber $k_{1/2}$ at which $P_{\text{seq}}/P_{\Lambda\text{CDM}} = 0.5$ lies above $200 h/\text{Mpc}$ for $f_T = 0.39$ and is not reached within the computed range for $f_T = 0.10$ — compared with $k_{1/2} \approx 27 h/\text{Mpc}$ for the constant- c_s^2 run at $T_{\text{seq}} = 1000$ K. The minimal Scenario A is therefore observationally viable across the entire physically motivated f_T range; the bounds quoted in the abstract and Sec. V F as constraints on T_{seq} in the constant- c_s^2 approximation are upper bounds that the realistic thermal history satisfies with enormous margin.

The effective coupling kernel $W(k, r_0)$ — which specifies how a cosmological density perturbation $\delta_{\text{seq}}(k)$ in the sequestered sector sources the gravitational potential $\Phi(k)$ in our sector — is derived in Appendix A. For a plane-wave perturbation with wavelength $\lambda = 2\pi/k$, much larger than the throat radius r_0 , the perturbation is almost entirely in the $\ell = 0$ multipole as seen by the throat (the Rayleigh expansion gives $f_0 = 1 - \mathcal{O}((kr_0)^2)$), and therefore transmits fully. Quantitatively,

$$W(k, r_0) = 1 - \mathcal{O}((kr_0)^2), \quad (35)$$

independent of the higher-multipole transmission factors. For submicron throats and any cosmological wavenumber ($k \lesssim 10 h/\text{Mpc}$), $kr_0 < 10^{-28}$ and $1 - W < 10^{-56}$. The coupling is indistinguishable from unity at all observable scales (Table II).

TABLE II. Coupling kernel $W(k, r_0)$ for representative throat radii and cosmological wavenumbers. All entries satisfy $1 - W < 10^{-50}$.

r_0	k (h/Mpc)	kr_0	$1 - W$	Status
10^{-13} m	0.1	2×10^{-37}	$< 10^{-73}$	exact CDM
10^{-13} m	10	2×10^{-35}	$< 10^{-69}$	exact CDM
10^{-7} m	0.1	2×10^{-31}	$< 10^{-61}$	exact CDM
10^{-7} m	10	2×10^{-29}	$< 10^{-57}$	exact CDM
1 mm	10	2×10^{-25}	$< 10^{-49}$	exact CDM

3. Equation of state

The sequestered baryons are non-relativistic ($w \approx 0$) at all post-BBN epochs and thus pressureless on cosmological scales (like CDM), but their internal dynamics differ from CDM at the microphysical level. This affects perturbation growth on small scales (see Sec. VI) but not on CMB or BAO scales. The gravitational coupling between sectors is addressed by the coupling kernel derivation (Appendix A): $W(k, r_0) = 1 - \mathcal{O}((kr_0)^2) \approx 1$ at all observable scales, confirming that the coupling is indistinguishable from CDM.

4. Isocurvature perturbations

If the visible and sequestered sectors underwent independent thermal histories, their initial perturbations could in principle be uncorrelated, generating baryon-dark-sector isocurvature modes. The Planck CMB power spectrum analysis constrains such modes to $\beta_{\text{iso}} < 0.038$ (95% CL) for the uncorrelated baryon-isocurvature mode [7], where $\beta_{\text{iso}} \equiv P_{II}(k_*)/[P_{RR}(k_*) + P_{II}(k_*)]$ is the primordial power-spectrum ratio at the Planck pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$. We argue that the geometric-sequestration scenario is naturally consistent with this bound, refine the argument quantitatively for each of the three asymmetric-reheating mechanisms of Sec. IV D, and verify the conclusion by direct CLASS computation.

a. Adiabaticity by construction First, if both sectors share the same inflationary origin — the same inflaton field generates the primordial curvature perturbation on both sides of each throat — then the comoving curvature perturbation \mathcal{R} is a single scalar degree of freedom defined on the shared slicing of spacetime during inflation, before the throats decouple dynamically. The resulting initial condition is adiabatic by construction: $\delta_{\text{seq}}/\delta_b = 1$, with no isocurvature component at the time when both sectors emerge from inflation on the same FLRW background.

Second, the asymmetric reheating scenario required for $\xi \ll 1$ does not destroy this adiabaticity at the level of the homogeneous expansion. Asymmetric reheating affects the post-inflationary energy density per sector —

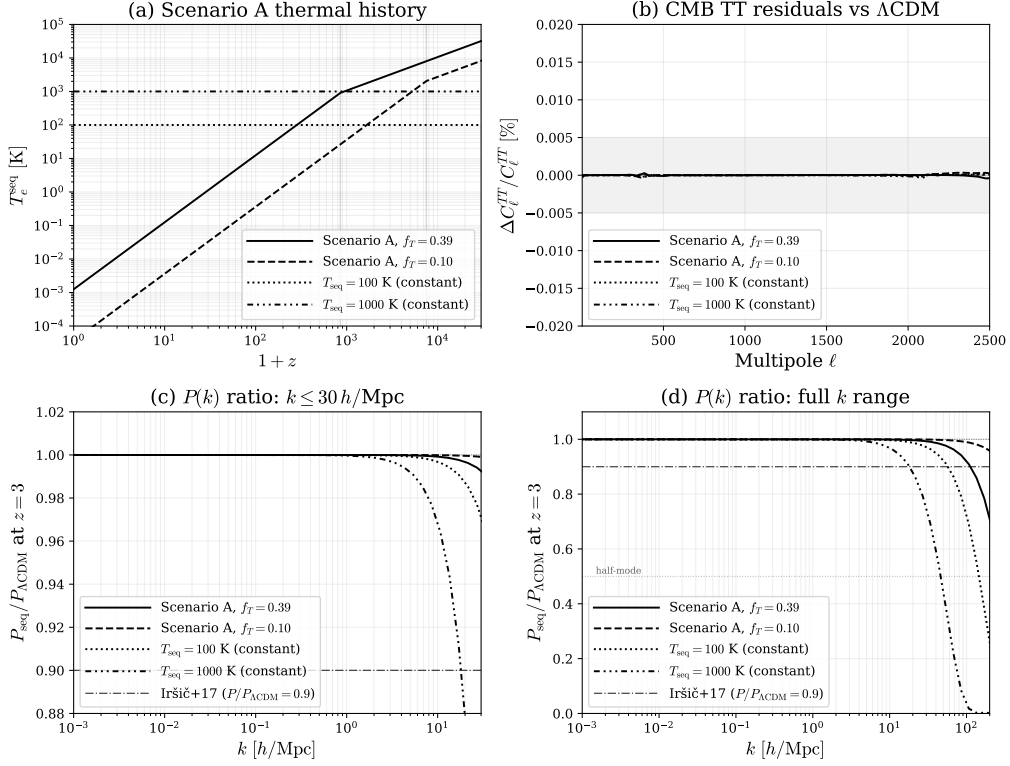


FIG. 4. Self-consistent CLASS Boltzmann computation for the realistic Scenario A thermal history. Panel (a): the piecewise sequestered-sector temperature $T_e^{\text{seq}}(z)$ for $f_T = 0.39$ (solid; saturating the N_{eff} bound) and $f_T = 0.10$ (dashed; deeper asymmetric reheating), with the constant $T = 100$ K and $T = 1000$ K reference levels of Fig. 3 shown as horizontal lines. Both Scenario A trajectories are Compton-coupled at high z and adiabatic at low z , with the kink at the corresponding $z_{\text{dec}}^{\text{seq}}$. Panel (b): CMB TT residuals relative to Λ CDM. All four curves (two Scenario A, two constant- c_s^2 reference) lie within the $\pm 5 \times 10^{-5}$ numerical precision floor of the modified CLASS code (grey band), more than two orders of magnitude below Planck’s per-multipole sensitivity. Panel (c): matter power spectrum ratio $P_{\text{seq}}(k)/P_{\Lambda\text{CDM}}(k)$ at $z = 3$ on a linear vertical scale near unity, showing that Scenario A is indistinguishable from Λ CDM through $k \sim 10 h/\text{Mpc}$, while the constant- c_s^2 run at $T = 1000$ K has visibly suppressed power at the Lyman- α pivot. Panel (d): same ratio across the full computed k range. Scenario A turns over only at high k ($\gtrsim 30 h/\text{Mpc}$ for $f_T = 0.39$, $\gtrsim 100 h/\text{Mpc}$ for $f_T = 0.10$) due to the brief high- z Compton-coupled phase, but remains comfortably above the Iršič *et al.* [38] bound (dot-dashed line) at all cosmologically observable wavenumbers.

the two sides receive different effective reheating temperatures — but, at linear order, it does not disturb the common comoving curvature perturbation \mathcal{R} . This follows from the standard result that \mathcal{R} is conserved on super-horizon scales for single-field inflation irrespective of the subsequent decay channel [39, 40]: the conservation law for \mathcal{R} on super-horizon scales depends only on the existence of a well-defined adiabatic mode in the two-fluid system, which in turn requires only that the two sectors share the same expansion history at leading order — a condition guaranteed by our scale-separation argument ($Hr_0/c \ll 1$, Sec. II). Isocurvature contamination can only be generated by sources that correlate with spatial position in a way the homogeneous-expansion argument cannot capture; we now examine each candidate source in turn.

b. Mechanism (i): inflaton decoherence at narrow throats This mechanism is intrinsically position-dependent: the branching ratio between visible-sector and sequestered-sector decay products differs only inside throat regions, and is uniform elsewhere. The resulting baryon-vs-sequestered density ratio therefore varies across the universe with a spatial pattern set by the throat distribution. The amplitude of this position dependence is bounded by the volume fraction of throat regions:

$$f_{\text{throat}} \sim \frac{N r_0^3}{V_{\text{halo}}}. \quad (36)$$

For the upper-count edge of the lensing window of Sec. VII E ($N \sim 2 \times 10^{29}$, $r_0 \sim 10^{-13}$ m, $V_{\text{halo}} \approx 10^{66}$ m³), $f_{\text{throat}} \sim 2 \times 10^{-76}$; for the

lower-count edge ($N \sim 2 \times 10^{22}$, $r_0 \sim 10^{-7}$ m), $f_{\text{throat}} \sim 2 \times 10^{-65}$. The position-dependent isocurvature amplitude is bounded by this fraction at every scale.

A finer analysis recognises that the isocurvature contribution is a Poisson shot-noise term with a characteristic correlation scale set by the inter-throat spacing $\bar{d} = (V_{\text{halo}}/N)^{1/3}$, which spans approximately $10^{12}\text{--}10^{15}$ m ($10^{-3}\text{--}10^0$ kpc) across the viable window — always microscopic compared to cosmological scales. At wavenumbers $k\bar{d} \ll 1$ — i.e., all cosmologically observable scales, where $k \lesssim 1 h/\text{Mpc}$ and $k\bar{d} \lesssim 10^{-7}$ — the Poisson shot-noise power is further suppressed by $(k\bar{d})^3$, giving

$$\beta_{\text{iso}}^{(i)}(k_{\text{CMB}}) \lesssim f_{\text{throat}}(k_{\text{CMB}}\bar{d})^3 \sim 10^{-107} \quad (\text{upper-count edge}) \quad (37)$$

with the lower-count edge bounded similarly to $\sim 10^{-89}$. The volume-fraction bound is therefore many orders of magnitude weaker than the actual cosmological-scale isocurvature contribution from this mechanism, and both are vastly below the Planck constraint.

c. Mechanism (ii): post-reheating throat formation If throats form after reheating, the sequestered sector temperature is set globally by the cosmic temperature at the formation epoch. Two sub-cases of this mechanism arise. In the spatially-uniform sub-case (throat formation uncorrelated with primordial density), there is no isocurvature contribution at any wavenumber: the asymmetry ξ is global, like the photon-to-baryon ratio. In the density-correlated sub-case (throat formation triggered by primordial overdensity peaks above a threshold δ_c), the throat density correlates with \mathcal{R} on the co-moving scales of the trigger, and the resulting baryon-vs-sequestered ratio inherits the same correlation. The fractional contribution to β_{iso} at a CMB-relevant scale k is bounded by the same volume-fraction times an $\mathcal{O}(1)$ peak-bias enhancement [41] — for $b_{\text{thresh}} \sim \mathcal{O}(1)\text{--}\mathcal{O}(10)$, the prefactor remains $f_{\text{throat}} b_{\text{thresh}}^2 \lesssim 10^{-63}$, and the resulting β_{iso} is again far below the Planck bound at all CMB-relevant wavenumbers.

d. Mechanism (iii): geometric suppression of preheating The curvature-induced effective mass shift operates within a curvature-perturbed region near each throat, not strictly inside the throat. The width of this region is set by the inflaton wavelength λ_ϕ . For mechanism (iii) to produce isocurvature contamination, a spatially-varying suppression of preheating must imprint differently on different regions; the effective “volume fraction” of suppressed regions is

$$f_{\text{eff}} \sim \frac{N \lambda_\phi^3}{V_{\text{halo}}}.$$

For chaotic inflation with $m_\phi \sim 10^{13}$ GeV, $\lambda_\phi \sim 10^{-29}$ m, and $f_{\text{eff}} \sim 10^{-123}$ at the upper-count edge of the throat number distribution. The combined Poisson-suppressed bound at CMB scales is

$$\beta_{\text{iso}}^{(iii)}(k_{\text{CMB}}) \lesssim f_{\text{eff}}(k_{\text{CMB}}\bar{d})^3 \lesssim 10^{-154},$$

many orders of magnitude smaller than mechanism (i).

e. Quantitative verification: isocurvature amplitude To verify the analytic conclusions concretely, we compute the resulting CMB and matter-power spectra using the modified CLASS module under primordial baryon-isocurvature initial conditions (ic = ad,bi), at four amplitudes spanning seven orders of magnitude in primordial β_{iso} . In the CLASS implementation used here, the primordial spectra ratio is parametrised as $P_{II}(k)/P_{RR}(k) = f_{\text{bi}}^2$ (scale-invariant, with $n_{\text{bi}} = n_s$), so the standard Planck primordial β_{iso} maps to $f_{\text{bi}} = \sqrt{\beta_{\text{iso}}/(1 - \beta_{\text{iso}})}$. This convention was verified empirically by ratio comparison between calibration runs at $f_{\text{bi}} = 0.5$ and $f_{\text{bi}} = 1.0$ (BI-mode CL ratio = 4.0000 at every multipole, confirming the f_{bi}^2 scaling to all decimal places).

Four production runs were performed: (08) Planck-saturating $\beta_{\text{iso}} = 0.038$; (09) intermediate $\beta_{\text{iso}} = 10^{-3}$; (10) precision-floor anchor $\beta_{\text{iso}} = 10^{-5}$; and (11) the same precision-floor amplitude run together with the Scenario A thermal history of Sec. V G 2. The results are shown in Fig. 5. Three robustness checks emerge from the runs:

Adiabatic-mode preservation. The adiabatic CL spectrum is bit-identical across runs (08), (09), and (10), demonstrating that the f_{bi} parameter affects only the BI-mode power without spurious cross-coupling to the adiabatic sector. The cross-correlation parameter $\alpha_{\text{bi}} = 0$ (uncorrelated isocurvature) is honoured exactly across three orders of magnitude in f_{bi} .

f_{bi}^2 scaling. The BI-mode CL spectra at the three different f_{bi} values are perfectly parallel curves (Fig. 5b), separated by the exact f_{bi}^2 ratios $(0.1987/0.0316)^2 = 39.5386$ and $(0.1987/0.00316)^2 = 3953.86$, reproduced to all decimal places at every multipole.

Scenario A and isocurvature compose orthogonally. Run (11) gives AD- and BI-mode residuals identical to run (10) to better than 10^{-5} fractionally at any multipole (Fig. 5d), demonstrating that the modified Scenario A thermal-history evolution and the standard primordial-mode initial conditions do not produce spurious cross-terms in the modified perturbation equations.

The CMB TT residuals scale linearly in β_{iso} from the Planck-saturation value $\sim 0.4\%$ at low ℓ down to the $\sim 5 \times 10^{-5}$ numerical precision floor of the modified code at $\beta_{\text{iso}} = 10^{-5}$ (Fig. 5c). Extrapolating the linear scaling to the analytic prediction of mechanism (i) — $\beta_{\text{iso}} \lesssim 10^{-65}$ as a volume-fraction upper bound, or 10^{-107} once Poisson suppression is accounted for — gives a CMB residual of $\lesssim 10^{-65}\%$, vastly below any conceivable detector sensitivity. The geometric-sequestration scenario is therefore essentially indistinguishable from ΛCDM in any conceivable isocurvature observation.

f. Forecast against future-survey sensitivities Future-survey sensitivities improve the bound modestly: CMB-S4 forecasts $\beta_{\text{iso}} \sim 0.005$ [42], LiteBIRD comparable [43], both still leaving a residual margin

of dozens of orders of magnitude over the predicted values. No conceivable extension of CMB sensitivity, by any of the experiments currently planned or proposed, would approach the predicted geometric-sequestration β_{iso} . The conclusion is robust: the scenario satisfies isocurvature constraints with margin so large that this is not a meaningful test of the model.

H. Reproducibility of the modified Boltzmann code

The CLASS modifications are summarised here for reproducibility. The full source, the seven `.ini` files driving the runs of Fig. 3 and Fig. 4, and the analysis scripts producing the figures, are publicly available at https://github.com/jeffriley/class_sequestered_baryons, tagged as *seq-baryon-paper-v1*.

The modifications are confined to three source files: the perturbation module (`perturbations.c`), the input parser (`input.c`), and the background structure (`background.h`). The sequestered-baryon fluid is implemented within the existing interacting dark matter (IDM) framework as a pressureful non-relativistic component with sound speed $c_s^2(a)$ computed in one of three mutually exclusive modes selected by the input parameter `cs2_idm_mode`: (i) *constant*, with $c_s^2 = k_B T_{\text{seq}} / (m_p c^2)$ fixed at all redshifts (parameter `cs2_seq`); (ii) *adiabatic*, with $c_s^2(a) = c_{s,0}^2 / a^2$ capped at $c^2/3$ at the relativistic threshold (this mode produces the warm-today/hot-in-the-past family of Sec. V G 1, which is excluded); and (iii) *scenarioA*, with the time-dependent piecewise thermal history of Sec. V G 2 (parameters `f_T` and `z_dec_seq`), Compton-coupled at high z and adiabatic at low z . The three modes share the same `cs2`-to-perturbation-equation interface and an explicit mutual-exclusion error trap at module-init time. The synchronous gauge is used throughout; the gauge-invariant pressure perturbation is treated to leading order in the equilibrium approximation $\delta p = c_s^2 \delta \rho$, which is exact for the constant- c_s^2 and adiabatic modes and accurate to sub-percent in the matter-dominated era for Scenario A.

Numerical validation comprises four checks: (a) the IDM $c_s^2 = 0$ run reproduces the standard Λ CDM C_ℓ^{TT} to within 7×10^{-5} at all multipoles $\ell \leq 2500$, defining the precision floor of the modified module; (b) the existing constant- c_s^2 runs at $T_{\text{seq}} \in \{10, 100, 1000\}$ K reproduce $P(k=10 \text{ h/Mpc})/P_{\Lambda\text{CDM}}$ at $z = 0, 3, 5$ to four significant figures across regression boundaries; (c) the adiabatic-cooling runs reproduce the half-mode wavenumbers $k_{1/2} \approx 1.4, 0.47, 0.19 \text{ h/Mpc}$ for the three temperatures, in agreement with independent semi-analytic estimates; and (d) unitarity/energy-conservation checks built into CLASS pass at standard precision settings for all twelve runs. The CLASS precision settings used for all production runs are `l_max_scalars=3000`, `P_k_max_h/Mpc=100`, default HyRec recombination, and `halofit` non-linear corrections (irrelevant at the relevant scales but enabled for completeness). Convergence under doubling of `l_max`

and `P_k_max` is verified for the Scenario A runs at the 10^{-5} level. The README in the repository provides the exact command-line invocation reproducing each figure.

VI. Structure formation with sequestered baryons

A. The CDM transfer function and its sequestered-baryon analogue

In Λ CDM, the matter power spectrum $P(k)$ is shaped by the transfer function $T(k)$, which encodes the suppression of perturbation growth on scales that enter the horizon during radiation domination. CDM perturbations grow logarithmically during radiation domination (the Mészáros effect [44]) and as $\delta \propto a$ during matter domination. The transfer function transitions from $T(k) \rightarrow 1$ on large scales ($k \ll k_{\text{eq}}$) to $T(k) \propto k^{-2} \ln k$ on small scales ($k \gg k_{\text{eq}}$), where k_{eq} is the wavenumber corresponding to the horizon at matter-radiation equality.

Sequestered baryons behave identically to CDM on scales $\lambda \gg \lambda_J$, where λ_J is the Jeans length of the sequestered medium. On these scales, the sequestered baryons are pressureless and cluster gravitationally, producing a transfer function indistinguishable from CDM. On scales $\lambda \lesssim \lambda_J$, the sequestered baryons have residual pressure that prevents gravitational collapse, suppressing the power spectrum below the Jeans scale.

B. The Jeans scale of sequestered baryons

The Jeans length for a non-relativistic fluid at temperature T and density ρ is

$$\lambda_J = c_s \sqrt{\frac{\pi}{G\rho}}, \quad (38)$$

where $c_s = \sqrt{k_B T / m_p}$ is the sound speed for a hydrogen gas. The Jeans scale of the sequestered baryons depends on their temperature and density, which are determined by the thermal history of the sequestered regions. As with the CMB analysis of Sec. V F, we use a constant sound speed throughout — this is an effective description appropriate to a sequestered sector maintained at temperature T_{seq} by its own internal thermodynamics. We can, however, bound the effect of a redshift-dependent sound speed analytically.

For a non-relativistic monatomic ideal gas expanding adiabatically with the Hubble flow ($\gamma = 5/3$), the temperature scales as $T(z) = T_0(1+z)^2$, giving a sound speed $c_s(z) = c_{s,0}(1+z)$, where T_0 and $c_{s,0}$ are the present-day values. The comoving Jeans wavenumber in this case evolves as

$$k_J^{\text{com}}(z) = \frac{\sqrt{4\pi G\rho(z)}}{c_s(z)} \frac{1}{1+z} = k_J^{\text{com}}(0) (1+z)^{-1/2}, \quad (39)$$

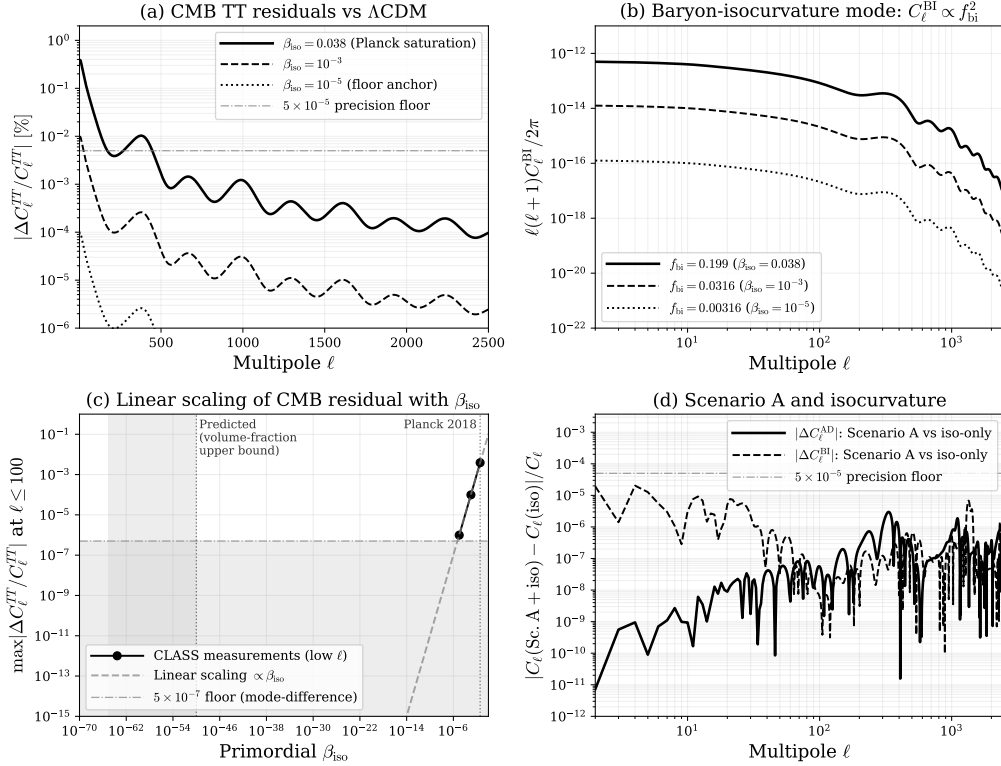


FIG. 5. Empirical CLASS verification of the analytic isocurvature analysis (Sec. V G 4).

Panel (a): fractional CMB TT residuals (absolute value) relative to Λ CDM at three primordial isocurvature amplitudes, demonstrating the linear scaling in β_{iso} and the residual at $\beta_{\text{iso}} = 10^{-5}$ sitting at the 5×10^{-5} numerical precision floor of the modified code.

Panel (b): per-mode baryon-isocurvature CL contributions for the three production runs. The three curves are parallel, separated by the exact ratios $(0.1987/0.0316)^2$ and $(0.1987/0.00316)^2$, confirming the f_{bi}^2 scaling to all decimal places.

Panel (c): linear scaling of maximum |CMB residual| at low ℓ versus primordial β_{iso} . The shaded region indicates the geometric-sequestration prediction range from the volume-fraction upper bound (Sec. V G 4, Eq. 37) — many tens of orders of magnitude below the precision floor, and below the Planck constraint by sixty-three orders of magnitude.

Panel (d): orthogonality verification. Per-mode AD and BI fractional differences between the Scenario A + isocurvature run (11) and the isocurvature-only run (10) lie at 10^{-7} – 10^{-9} at all multipoles, far below the 5×10^{-5} numerical precision floor. The two extensions of the CLASS module compose with no detectable cross-terms.

where we have used $\rho(z) = \rho_0(1+z)^3$. By contrast, the constant- c_s approximation used in our CLASS implementation gives

$$k_J^{\text{com}}(z)|_{c_s=\text{const}} = k_J^{\text{com}}(0)(1+z)^{1/2}. \quad (40)$$

The ratio between the two cases is $k_J^{\text{adiab}}/k_J^{\text{const}} = 1/(1+z)$, independent of T_0 . At the Lyman- α epoch ($z \approx 3$), the adiabatic Jeans wavenumber is a factor of four smaller than the constant- T value — the comoving Jeans scale is correspondingly four times larger, and modes at four times smaller wavenumber are suppressed. For $T_0 = 100$ K, $k_J^{\text{const}}(z=3) \approx 94$ h/Mpc (well above the Lyman- α constraint region), while $k_J^{\text{adiab}}(z=3) \approx 24$ h/Mpc (marginal at the Iršič pivot). For $T_0 = 1000$ K, the adiabatic Jeans wavenumber at $z = 3$ falls to $k_J \approx 7.5$ h/Mpc, well inside the Lyman- α exclusion region.

The physical origin of the difference is that adiabatically cooling gas was hotter in the past ($T(z=3) = 16 T_0$ vs. T_0 in the constant case), so its sound speed and Jeans length were larger at the redshifts where the Lyman- α forest is most sensitive. Modes that are comfortably above k_J today were below k_J at earlier times and experienced Jeans-pressure suppression during the epoch of most rapid perturbation growth ($z \sim 1$ – 10), accumulating a growth deficit that persists to $z = 0$.

However, this analytic estimate captures only the instantaneous Jeans-scale shift at a single redshift and substantially underestimates the full effect. A self-consistent CLASS computation with $c_s^2(a) = c_{s,0}^2/a^2$ (capped at $c^2/3$ at the relativistic threshold) shows that the accumulated growth deficit over the entire matter-dominated era is far larger than the single-redshift estimate sug-

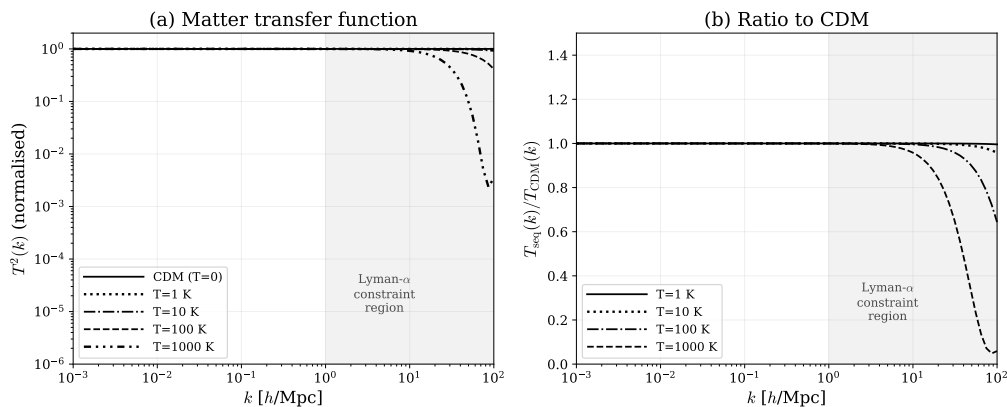


FIG. 6. Matter transfer function for the sequestered-baryon scenario.

Panel (a): $T^2(k)$ normalised to unity at $k \rightarrow 0$. The CDM baseline (black) and $T_{\text{seq}} = 1\text{--}10$ K (dots, dash-dot) are indistinguishable. $T_{\text{seq}} = 100$ K (dashed) shows mild suppression at $k \gtrsim 30$ h/Mpc . $T_{\text{seq}} = 1000$ K (dots-dash) is suppressed in the Lyman- α constraint region (shaded).

Panel (b): Ratio to CDM, showing the onset and depth of suppression. The Jeans pressure produces a sharper cutoff than WDM free-streaming.

gests—the half-mode scale falls to $k_{1/2} \approx 1.4$ h/Mpc ($T_0 = 10$ K), 0.47 h/Mpc ($T_0 = 100$ K), and 0.19 h/Mpc ($T_0 = 1000$ K), deep inside the galaxy-survey and BAO regime. At the Lyman- α pivot ($k = 10$ h/Mpc , $z = 3$), $P/P_{\Lambda\text{CDM}} < 10^{-4}$ for all three temperatures. CMB residuals reach up to $\sim 2\%$ at high ℓ , which is detectable by Planck. The pure adiabatic limit is excluded at every temperature.

The reason the CLASS result is so much stronger than the analytic estimate is that the sound speed c_s^2/a^2 at matter-radiation equality ($z \approx 3400$) is $\sim 10^7$ times the present-day value $c_{s,0}^2$, regardless of T_0 . This makes the Jeans scale enormous during the entire epoch when perturbation growth is fastest ($z \sim 10\text{--}3400$), suppressing growth on all scales above $k \sim 0.1\text{--}1$ h/Mpc —not just the small-scale Jeans cutoff that the analytic bracket targeted.

The constant- c_s results reported throughout this section apply directly under Scenario B of Sec. V G 1 (active diffuse-phase heating, with T_{seq} identified with the diffuse-phase temperature) and are conservative under Scenario A (passive diffuse phase, with the actual deviation from ΛCDM smaller than reported by many orders of magnitude). In either case the Lyman- α bounds quoted in this paper ($T_{\text{seq}} \lesssim 100$ K from the half-mode mapping, $\lesssim 10^3$ K from the direct-pivot criterion) apply without the adiabatic correction, since the excluded branch (warm today, hot in the past) is the one that fails. Quantifying how close the diffuse-phase thermal history is to either scenario—and in particular whether the small-scale-tension implications of Sec. VI C follow—requires identifying the diffuse-phase heating mechanism (if any), which we defer to future work (Sec. VIII).

Table III shows the Jeans length, Jeans wavenumber, and half-mode mass (the mass enclosed in a sphere of radius $\lambda_J/2$) for several values of T_{seq} .

TABLE III. Jeans scale of sequestered baryons at $z = 0$ for $\rho_{\text{seq}} = \Omega_c \rho_{\text{crit}}$.

T_{seq} (K)	λ_J (kpc)	k_J (Mpc^{-1})	M_{hm} (M_\odot)
1	13	470	4×10^4
10	42	150	10^6
100	130	47	4×10^7
1000	420	15	10^9
10000	1300	4.7	4×10^{10}

To verify these estimates, we solve the coupled perturbation equations (Eq. (29)) numerically using a purpose-built Boltzmann solver in the quasi-static approximation (valid for sub-horizon modes $k\tau \gg 1$). Figure 6 shows the resulting matter transfer function $T^2(k)$ for the sequestered-baryon component at several temperatures, compared with the CDM baseline ($c_s = 0$). The numerical results confirm the analytic estimates: for $T_{\text{seq}} \leq 10$ K, the transfer function is indistinguishable from CDM across the entire range $k = 0.001\text{--}100$ h/Mpc , for $T_{\text{seq}} = 100$ K a small suppression appears at $k \gtrsim 30$ h/Mpc , for $T_{\text{seq}} = 1000$ K the suppression reaches the boundary of the Lyman- α constraint region at the Iršič pivot ($P_{\text{seq}}/P_{\Lambda\text{CDM}} = 0.968$ at $k = 10$ h/Mpc , $z = 3$, and 0.899 at $z = 0$), placing this temperature in tension with current bounds, but not firmly excluded without a full flux-power likelihood analysis.

Current Lyman- α forest constraints on warm dark matter require the thermal-relic mass $m_{\text{WDM}} \gtrsim 4\text{--}5$ keV [38, 45], corresponding to a half-mode scale $k_{\text{hm}} \sim 10\text{--}30$ h/Mpc and a half-mode mass $M_{\text{hm}} \sim 10^8\text{--}10^9$ M_\odot . Comparing with Table III, consis-

tency with Lyman- α constraints requires $T_{\text{seq}} \lesssim 100$ K, corresponding to $\lambda_J \lesssim 130$ kpc and $k_J \gtrsim 47$ Mpc $^{-1}$. This is comfortably below the Lyman- α exclusion region.

For $T_{\text{seq}} \gtrsim 10^3$ K, the Jeans scale enters the range probed by the Lyman- α forest ($k \sim 1\text{--}10$ h/Mpc) and the scenario would be in tension with observations. This places an upper bound on the temperature of the sequestered regions at late times:

$$T_{\text{seq}}(z=0) \lesssim 100 \text{ K} \quad (\text{Lyman-}\alpha \text{ constraint}). \quad (41)$$

The constraint is robust across the redshift range probed by the Lyman- α forest ($z \sim 2\text{--}5$). At higher redshifts, the sequestered baryon density increases as $(1+z)^3$, so the Jeans length decreases as $(1+z)^{-3/2}$ and k_J increases. For $T_{\text{seq}} = 100$ K at $z = 3$, $k_J \approx 550$ Mpc $^{-1}$ (compared to 47 Mpc $^{-1}$ at $z = 0$), corresponding to $m_{\text{equiv}} \approx 36$ keV—far above the Lyman- α exclusion region. The constraint at the relevant redshifts is therefore weaker than the $z = 0$ estimate in Table III, making the bound $T_{\text{seq}} \lesssim 100$ K conservative.

To compare more quantitatively with published Lyman- α bounds, we compute the half-mode wavenumber $k_{1/2}$ (where $(D_{\text{seq}}/D_{\text{CDM}})^2 = 1/2$) from the growth equation (29) and map it to an equivalent WDM thermal-relic mass using the transfer function of Viel *et al.* [46]. For $T_{\text{seq}} = 100$ K, $k_{1/2} \approx 74$ h/Mpc, corresponding to $m_{\text{equiv}} \approx 6$ keV. The currently accepted exclusion is $m_{\text{WDM}} > 5.3$ keV ($k_{1/2} \approx 29$ h/Mpc) [38], so $T_{\text{seq}} = 100$ K is comfortably allowed. For $T_{\text{seq}} = 1000$ K, $k_{1/2} \approx 24$ h/Mpc, $m_{\text{equiv}} \approx 2$ keV—below the exclusion and ruled out.

It is important to note that the Jeans-pressure cutoff is sharper than the WDM free-streaming cutoff. At $k = 50$ h/Mpc, the WDM 5.3 keV model suppresses power by $\sim 87\%$, while $T_{\text{seq}} = 100$ K suppresses by only $\sim 27\%$. The Jeans cutoff preserves more power at intermediate k and cuts off more steeply at high k .

1. Mapping Lyman- α flux power to T_{seq}

This means the Lyman- α constraint on T_{seq} is likely weaker than a naïve WDM mass mapping would suggest, because the Lyman- α flux power is most sensitive to the intermediate- k regime where the Jeans model retains more power than WDM. To assess compatibility with Lyman- α flux-power analyses more directly, we compute the compressed linear power parameters used by modern emulators [47]: the dimensionless amplitude Δ_p^2 and slope n_p of the linear power spectrum at the pivot scale $k_p = 0.7$ Mpc $^{-1}$. At $z = 3$, the fractional change in Δ_p^2 relative to Λ CDM is -0.0005% ($T_{\text{seq}} = 10$ K), -0.004% ($T_{\text{seq}} = 100$ K), and -0.035% ($T_{\text{seq}} = 1000$ K). These are two orders of magnitude below the $\sim 5\%$ sensitivity of current Lyman- α analyses [38, 48]. The reason is that the Jeans cutoff occurs at $k \gg k_p$: the Lyman- α pivot scale lies entirely within the regime where the sequestered-baryon $P(k)$ is indistinguishable from Λ CDM. The direct

$P(k)$ suppression at $k = 10$ h/Mpc (Table III) is satisfied comfortably for $T_{\text{seq}} \lesssim 100$ K in our constant- c_s^2 implementation, and remains above the Iršič reference level $P/P_{\Lambda\text{CDM}} = 0.9$ at the Lyman- α pivot for temperatures up to $T_{\text{seq}} \approx 3 \times 10^3$ K, extrapolated linearly in T from our three-temperature numerical grid.

Combined with the N_{eff} constraint (24), this requires the sequestered regions to be both radiation-poor and cold, consistent with a picture where throats form early and the sequestered regions expand and cool with limited energy input.

C. Small-scale structure and Λ CDM tensions

The discussion that follows is conditional on Scenario B of Sec. V G 1 (or on the in-halo phase under Scenario A)—under Scenario A the diffuse-phase percent-level suppression in $P(k)$ occurs only at $k \gtrsim 30$ h/Mpc (Fig. 4) and does not, by itself, modify galactic-scale structure.

The Λ CDM model predicts structure on all scales down to the free-streaming length of the CDM particle, which is negligibly small for cold WIMPs or axions. This prediction is in tension with several observations at galactic and sub-galactic scales [49]—the cusp-core problem (simulated CDM halos have cuspy density profiles, while observed dwarf galaxies have cores), the missing satellites problem (simulations predict many more small halos than observed satellites), and the too-big-to-fail problem (the most massive simulated subhalos are denser than any observed dwarf galaxy).

Sequestered baryons with a Jeans scale $\lambda_J \sim 10\text{--}100$ kpc would naturally suppress structure below this scale, qualitatively similar to warm dark matter (WDM) models [50]. However, unlike WDM, the suppression mechanism is not free-streaming but pressure support—the sequestered baryons have a sound speed and resist collapse below the Jeans scale. This produces a sharper cutoff in the matter power spectrum than the gradual free-streaming damping of WDM, which is potentially distinguishable in the Lyman- α forest power spectrum [45].

A quantitative comparison with the observed matter power spectrum, halo mass function, and Lyman- α forest constraints requires a detailed calculation of the modified transfer function with a Jeans-scale cutoff, which has not been performed and is an open problem.

D. Self-interactions of sequestered baryons

An important distinction between CDM and sequestered baryons is that the latter interact with each other through electromagnetic and strong nuclear forces within the sequestered region. The sequestered baryons are not collisionless—they form a gas or plasma within

their own region, with a mean free path determined by their internal density and temperature.

This is analogous to self-interacting dark matter (SIDM) models [51, 52], which have been proposed to address the small-scale problems of Λ CDM. The self-interaction cross section per unit mass for baryonic matter is $\sigma/m \sim 10^{24} \text{ cm}^2/\text{g}$ for Thomson scattering at typical temperatures—many orders of magnitude larger than the SIDM cross sections considered in the literature ($\sigma/m \sim 0.1\text{--}10 \text{ cm}^2/\text{g}$). However, the effective self-interaction cross section for the sequestered baryons as perceived by the visible sector depends on how the internal dynamics of the sequestered region project onto the gravitational potential in our region. If the sequestered baryons thermalise internally, but interact with us only gravitationally, the effective “self-interaction” is mediated by the throat geometry and may be very different from the microscopic baryon-baryon cross section. A detailed analysis of this “effective self-interaction” is an important open problem. We note that current SIDM constraints from cluster observations and halo shapes require $\sigma/m \lesssim 1 \text{ cm}^2/\text{g}$ on cluster scales, and allow $\sigma/m \sim 1\text{--}10 \text{ cm}^2/\text{g}$ on dwarf galaxy scales [52]. If the effective cross section of sequestered baryons (as projected through throats) is velocity-dependent—larger at low velocities (dwarf galaxies) and smaller at high velocities (clusters)—this could naturally reproduce the velocity-dependent SIDM phenomenology that has been proposed to resolve small-scale CDM tensions.

E. Halo formation and galaxy-halo correlation

A challenge for any dark matter proposal is explaining the observed tight correlation between luminous galaxies and their dark matter halos. In Λ CDM, CDM halos form first from gravitational collapse seeded by primordial density fluctuations, and visible baryons subsequently fall into these potential wells, cool radiatively, and form galaxies. The correlation is a consequence of the shared gravitational dynamics: galaxies form where the dark matter density is highest.

In the geometric-sequestration scenario, the same causal chain operates. The sequestered baryons contribute to the gravitational potential through the monopole channel, producing potential wells that are indistinguishable from CDM halos at scales $r \gg r_0$ (Sec. II). Visible baryons fall into these wells, cool, and form galaxies in the standard way. The galaxy-halo correlation is therefore inherited directly from the gravitational dynamics, just as in Λ CDM—the only difference is the physical origin of the potential well (sequestered baryonic mass behind throats, rather than non-baryonic CDM particles).

What must be explained is why the throat spatial distribution tracks the large-scale density field. If the throat network formed from primordial density fluctuations—as would be the case for any formation mech-

anism tied to a cosmological phase transition or to the primordial curvature perturbation spectrum—then overdense regions naturally acquire more throats and deeper potential wells, while underdense regions acquire fewer. The galaxy-halo correlation then follows from the same shared-primordial-origin argument as in Λ CDM: galaxies and “dark matter halos” both trace the same primordial density peaks, because the throats that constitute the halos were seeded by those peaks. A quantitative test of this argument requires a specific formation model (the subject of WP5 in the research programme of Sec. VIII) that predicts the spatial correlations, mass function, and radial profile of the throat population and allows comparison with observed halo profiles, galaxy correlation functions, and abundance-matching constraints.

F. Dissipative dynamics within sequestered regions

An important consequence of the sequestered matter being baryonic is that it can undergo dissipative processes within its own region: radiative cooling, fragmentation, star formation, and supernova feedback. If the sequestered baryons cool efficiently, they could form disk-like structures within their sequestered volume, analogous to dissipative dark matter models [53]. The gravitational potential of such a disk, transmitted through the monopole channel, would be distinguishable from a spherical halo—it would produce oblate (rather than prolate or spherical) gravitational potential contours in our region, potentially affecting halo shapes inferred from gravitational lensing and satellite dynamics.

Current constraints on dissipative dark matter require that no more than $\sim 5\%$ of the dark matter can reside in a thin disk component [53]. In the geometric sequestration scenario, whether the sequestered baryons form disks depends on their cooling rate, which is determined by their internal temperature, density, and composition. If the sequestered regions are cold ($T_{\text{seq}} \lesssim 100 \text{ K}$, as required by the Lyman- α constraint (41)), the cooling rate is strongly suppressed. For a hydrogen gas at the mean cosmological dark-matter density

$$n = \Omega_c \rho_{\text{crit}} / m_p \approx 1.4 \times 10^{-6} \text{ cm}^{-3}$$

at $z = 0$, the dominant cooling mechanism at $T \sim 100 \text{ K}$ is molecular hydrogen (H_2) line emission. We assume the sequestered sector is chemically unprocessed—it has no stellar nucleosynthesis history, no metals, and no dust grains, so the cooling channels available to enriched gas (metal-line cooling, dust thermal emission, molecular cooling beyond H_2) are absent. Under this assumption, H_2 is the only relevant cooling channel below 10^4 K , and atomic hydrogen line cooling becomes available only for $T \gtrsim 10^4 \text{ K}$ [54]. The H_2 cooling function for a molecular fraction $f_{\text{H}_2} \sim 10^{-3}$ gives a volumetric cooling rate [55]:

$$\Lambda \sim f_{\text{H}_2} n^2 \Lambda_{\text{H}_2}(T) \sim 10^{-47} \text{ erg cm}^{-3} \text{ s}^{-1},$$

and a cooling time

$$t_{\text{cool}} = \frac{3nk_B T}{\Lambda} \sim 9 \times 10^{10} \text{ Gyr} \gg t_{\text{Hubble}}. \quad (42)$$

At the cosmic mean density, the cooling time exceeds the Hubble time by ten orders of magnitude. This means sequestered baryons at mean density cannot cool, fragment, or form disk-like structures, eliminating the “dark disk” concern [53].

Cooling becomes efficient only in overdense regions ($\delta \gtrsim 10^3$, where $t_{\text{cool}} < t_{\text{Hubble}}$), i.e., within collapsed halos. If sequestered baryons participate in gravitational collapse (which they do, since they couple gravitationally), the resulting dense cores could in principle cool to lower temperatures, fragment, and form compact objects within the sequestered volume. The collapse physics is potentially richer than the atomic-hydrogen threshold treatment used here—three-body H_2 formation, HD cooling, fine-structure transitions in light elements produced during BBN, and gravitational instability in marginally cooling gas can each open additional pathways. A quantitative treatment of which sequestered halos cool and what fraction of their mass collapses to compact objects depends on the halo mass function, the internal density profile, and the detailed cooling chemistry of a metal-free sector, none of which are modelled here. We expect, on the basis of analogous calculations for primordial-gas collapse in the visible sector, that the cooled fraction is small but nonzero. Compact objects formed by sequestered cooling would project onto our sector as localised mass concentrations at the corresponding throat mouths, and would be subject to the microlensing constraints of Sec. VII E. A complete quantitative treatment of dissipative dynamics within the sequestered sector is deferred to future work (Sec. VIII).

VII. Observational signatures and falsifiable predictions

A. Electromagnetic opacity as a function of frequency

The EM suppression is frequency-dependent: $T_{\text{EM}}(\omega)$ decreases as a power law below the barrier-top frequency $\omega_{\text{max}} = \sqrt{2}/r_0$. High-frequency photons ($\omega \gg \omega_{\text{max}}$) propagate freely through the throat—only $\omega < \omega_{\text{max}}$ is suppressed. Complete electromagnetic invisibility across the observed spectrum (from radio to gamma-ray) requires $r_0 \lesssim \sqrt{2} \lambda_{\text{min}}/(2\pi)$, where λ_{min} is the shortest wavelength of interest (the condition $\omega < \omega_{\text{max}}$ in SI units becomes $2\pi c/\lambda < \sqrt{2} c/r_0$). The required throat radius is $r_0 \lesssim 2 \times 10^{-13}$ m for gamma rays ($\lambda \sim 1$ pm), $\lesssim 10^{-7}$ m for optical ($\lambda \sim 500$ nm), and $\lesssim 0.2$ m for radio ($\lambda \sim 1$ m).

If r_0 is not sufficiently small, there would be a frequency-dependent leakage—high-frequency photons

would be partially transmitted while low-frequency photons are blocked. This would produce a distinctive spectral signature—suppression increasing toward lower frequencies—that is in principle detectable and would distinguish this scenario from particulate dark matter.

B. Gravitational wave suppression

Two distinct physical questions arise for gravitational waves in this scenario, and we treat them separately. First, can GWs propagate through a throat from the sequestered side to the visible side? This is the cross-throat transmission question studied in [1], which establishes that the $\ell = 2$ tensor modes face the same centrifugal barrier as the $\ell \geq 1$ EM modes, and are therefore strongly suppressed at all astrophysically relevant frequencies. Second, are visible-sector GWs from astrophysical sources attenuated by scattering off intervening throats as they propagate through a throat-populated universe? This is a different process—Rayleigh scattering off compact obstacles much smaller than the wavelength—and has a different answer, derived below.

Gravitational waves are $\ell \geq 2$ tensor perturbations and see a barrier similar to the EM one, with a curvature correction that lowers the barrier height [1]. For $\ell = 2$ on the EB throat, the barrier-top frequency is $f_{\text{max}}^{(\text{GW})} = \sqrt{3} c/(2\pi r_0)$. Modes with frequency below f_{max} are in the sub-barrier regime and are strongly suppressed, whereas modes above f_{max} propagate freely.

For $r_0 \sim 10^{-13}$ m (the gamma-ray constraint), $f_{\text{max}}^{(\text{GW})} \sim 8 \times 10^{20}$ Hz. All astrophysical GW frequencies—including the band probed by the ground-based LIGO [56], Virgo [57], and KAGRA [58] interferometers (~ 10 – 10^3 Hz) and the band targeted by the planned space-based Laser Interferometer Space Antenna (LISA) [59] ($\sim 10^{-4}$ – 1 Hz)—are far below this barrier top and are therefore strongly suppressed, just like EM radiation.

For submicron throats, the sequestered matter is invisible in all propagating channels (EM and GW) across the entire astrophysically observed spectrum. Only the static gravitational monopole ($\ell = 0$) transmits. Gravitational wave signals from mergers, inspirals, or supernovae occurring in the sequestered sector would not be detectable by any current or planned GW observatory. We show below (Sec. VII B 2) that this conclusion extends to fermionic fields, including the $j = \frac{1}{2}$ s-wave channel.

1. Quantitative GW transmission at astrophysical frequencies

For each detector band and throat radius in the viable window, the relevant dimensionless frequency is $\omega r_0/c = 2\pi f r_0/c$. In the deep sub-barrier regime ($\omega r_0/c \ll 1$), the GW $\ell = 2$ transmission follows the

asymptotic scaling $T_{\text{GW}} \sim K(\omega r_0/c)^{2\ell+1} = K(\omega r_0/c)^5$, set by Bessel-function matching for the centrifugal $1/\sigma^2$ tail of the EB barrier (Appendix B). The empirical exponent $\nu \approx 6$ fitted in the WKB-dominated range $\omega r_0/c \in [0.003, 0.1]$ [1] is steeper than the asymptotic value; the strict $\nu = 5$ scaling controls the deep sub-barrier regime relevant to astrophysical frequencies, where $\omega r_0/c \lesssim 10^{-12}$. We adopt $\nu = 5$ as the conservative (less-suppressed) estimate, with the prefactor $K \sim \mathcal{O}(1)$ bounded above by the throat-profile constant from the parametric family in [1].

TABLE IV. Estimated GW $\ell = 2$ transmission across astrophysical frequency bands and throat radii in the viable window, using the asymptotic scaling $T_{\text{GW}} \sim K(\omega r_0/c)^5$ with $K \sim \mathcal{O}(1)$. The viable window for r_0 spans roughly seven orders of magnitude (Sec. III): 10^{-13} m (gamma-ray opacity edge) to 10^{-7} m (lower-count microlensing edge).

Detector	f (Hz)	$r_0 = 10^{-13}$ m	$r_0 = 10^{-9}$ m	$r_0 = 10^{-7}$ m
PTA	10^{-9}	10^{-148}	10^{-128}	10^{-118}
LISA	10^{-2}	10^{-113}	10^{-93}	10^{-83}
LIGO/Virgo	10^3	10^{-88}	10^{-68}	10^{-58}

The transmission is far below any conceivable observational sensitivity across the entire (f, r_0) plane spanned by current and planned GW observatories and the throat radii admitted by the lensing and exotic-energy bounds. The least-suppressed entry, $T \sim 10^{-58}$ at LIGO/Virgo frequencies for $r_0 = 10^{-7}$ m, is some forty orders of magnitude below any plausible detector sensitivity. Sequestered-sector merger or inspiral signals are therefore undetectable in any propagating GW channel; the only gravitational signature available is the static monopole field of the enclosed total mass, captured by lensing and dynamical observables.

2. Fermion transmission: the $j = \frac{1}{2}$ s-wave

Standard Model neutrinos are spin- $\frac{1}{2}$ fermions. The Dirac equation on a spherically symmetric background decomposes into partial waves with total angular momentum $j = \frac{1}{2}, \frac{3}{2}, \dots$. The $j \geq \frac{3}{2}$ channels encounter centrifugal barriers analogous to the EM case and are suppressed. The critical question is whether the $j = \frac{1}{2}$ channel—the lowest fermionic partial wave, which contains the $\ell = 0$ (s-wave) orbital angular momentum component—is also suppressed, or whether it provides an unsuppressed energy-exchange route across the throat.

We show here that the $j = \frac{1}{2}$ mode is barrier-suppressed on the EB throat, with power-law transmission qualitatively similar to (but quantitatively weaker than) the EM $\ell = 1$ mode. The result follows from the structure of the radial Dirac equation on the ultrastatic EB background.

On the metric $ds^2 = -dt^2 + d\sigma^2 + a(\sigma)^2 d\Omega^2$ with $a(\sigma) = \sqrt{\sigma^2 + r_0^2}$, the Dirac equation for a fermion of mass m separates into radial and angular parts. For a mode with energy ω and angular quantum number $\kappa = \pm(j + \frac{1}{2})$, the two radial spinor components F and G satisfy the coupled first-order system

$$F' + \frac{\kappa}{a} F = (\omega + m) G, \quad (43)$$

$$G' - \frac{\kappa}{a} G = -(\omega - m) F. \quad (44)$$

For $j = \frac{1}{2}$, the two parity channels are $\kappa = -1$ ($\ell = 0$, the s-wave) and $\kappa = +1$ ($\ell = 1$). The s-wave channel $\kappa = -1$ is the one that might, a priori, evade the centrifugal barrier because its orbital angular momentum is zero.

Eliminating G from the coupled system (43)–(44) yields a second-order Schrödinger-form equation for F :

$$F'' + [\omega^2 - m^2 - V_F(\sigma)] F = 0, \quad (45)$$

with effective potential

$$V_F(\sigma) = \frac{\kappa^2}{a^2} + \frac{\kappa a'}{a^2}. \quad (46)$$

Eliminating F instead gives a Schrödinger-form equation for G with the partner potential

$$V_G(\sigma) = \frac{\kappa^2}{a^2} - \frac{\kappa a'}{a^2}. \quad (47)$$

These are supersymmetric partner potentials related by the superpotential $W(\sigma) = \kappa/a(\sigma)$:

$$V_F = W^2 + W', \quad V_G = W^2 - W'. \quad (48)$$

A standard result of supersymmetric quantum mechanics is that partner potentials have identical transmission spectra, $T_F(\omega) = T_G(\omega)$, for all ω [60]. The physical power transmission coefficient of the coupled Dirac system can therefore be computed from either equation.

For the s-wave channel $\kappa = -1$ on the EB throat ($a' = \sigma/a$), the two potentials are

$$V_F|_{\kappa=-1} = \frac{1 - \sigma/a}{a^2}, \quad V_G|_{\kappa=-1} = \frac{1 + \sigma/a}{a^2}. \quad (49)$$

At the throat ($\sigma = 0$), both equal $1/r_0^2$. Their asymptotic behaviours differ: $V_F \sim r_0^2/(2\sigma^4)$ (fast decay), while $V_G \sim 2/\sigma^2$ (the same $1/\sigma^2$ centrifugal tail that governs the EM $\ell = 1$ mode). The SUSY equivalence $T_F = T_G$ ensures that the physical transmission is controlled by the $1/\sigma^2$ tail regardless of which component equation is used for the computation—the s-wave fermion sees an effective centrifugal barrier with the same long-range structure as the EM dipole.

We compute the transmission coefficient by Numerov integration of Eq. (45) (equivalently G) for massless fermions ($m = 0$), using the same numerical method as the EM and GW computations in [1]. The integration

domain is $\sigma \in [-L_{\max}, +L_{\max}]$ with $L_{\max} = 50 r_0 - 400 r_0$ (scaled with ω to ensure convergence), and the transmission coefficient is extracted from the asymptotic plane-wave decomposition at $\sigma = \pm L_{\max}$. We verify that $T_F = T_G$ to better than 10^{-4} at every tabulated frequency, confirming the SUSY partner-potential identity numerically.

The results are shown in Table V. In the sub-barrier regime ($\omega r_0 \lesssim 0.5$), the transmission follows a power law $T \sim K (\omega r_0)^\nu$ with best-fit exponent $\nu \approx 4.2$ and prefactor $K \approx 0.4$. For comparison, the EM $\ell = 1$ mode on the same background has $\nu \approx 6.0$ [1]. The s-wave fermion is therefore less strongly suppressed than the EM dipole—by approximately three orders of magnitude at $\omega r_0 = 0.1$ —but is still strongly suppressed: $T \approx 3 \times 10^{-5}$ at $\omega r_0 = 0.1$, falling to $\sim 10^{-6}$ at $\omega r_0 = 0.05$.

TABLE V. Power transmission coefficient for the Dirac $j = \frac{1}{2}$ s-wave ($\kappa = -1$, massless) on the EB throat, compared with the EM $\ell = 1$ mode. Numerov integration, $r_0 = 1$.

ωr_0	$T_{\text{Dirac}}^{(j=1/2)}$	$T_{\text{EM}}^{(\ell=1)}$
1.5	9.7×10^{-1}	7.9×10^{-1}
1.0	5.9×10^{-1}	8.8×10^{-2}
0.5	2.8×10^{-2}	5.4×10^{-4}
0.3	2.7×10^{-3}	1.9×10^{-5}
0.1	2.6×10^{-5}	2.1×10^{-8}
0.05	1.6×10^{-6}	3.2×10^{-10}

The barrier height for the s-wave fermion is $V_F(0) = V_G(0) = 1/r_0^2$, which is half the EM $\ell = 1$ barrier height $V_{\text{EM}}^{(\ell=1)}(0) = 2/r_0^2$. This accounts for the weaker suppression: the barrier-top frequency is $\omega_{\max} = 1/r_0$ (fermion) versus $\sqrt{2}/r_0$ (EM $\ell = 1$), so the fermion enters the sub-barrier regime at a lower frequency and accumulates less WKB tunnelling action at any given $\omega < \omega_{\max}$. The difference in the power-law exponent ($\nu \approx 4.2$ versus 6.0) reflects both the lower barrier and the different shape of the partner potential V_G relative to the EM potential $V_{\text{EM}} = 2/a^2$ —although both have $1/\sigma^2$ tails, the coefficients differ (2 for V_G versus 2 for V_{EM} at leading order), and the sub-leading structure near the throat core contributes differently to the WKB action.

The physical consequence is that the $j = \frac{1}{2}$ s-wave is the least suppressed channel for any Standard Model field on the EB throat. It sets the floor for the barrier suppression. All other channels—EM ($\ell \geq 1$), GW ($\ell \geq 2$), massive scalars ($\ell \geq 0$ with additional mass barrier), and higher fermionic partial waves ($j \geq \frac{3}{2}$)—are more strongly suppressed. For the washout estimate of Sec. IV D, the s-wave fermion transmission should therefore replace the EM $\ell = 1$ value as the controlling channel: the universal washout parameter (25) gives $\omega r_0/c \approx 3 \times 10^{-10}$ for $r_0(t_0) = 10^{-13}$ m, at which the

s-wave transmission factor is $T \sim (3 \times 10^{-10})^{4.2} \sim 10^{-40}$. For submicron throats ($r_0 \lesssim 10^{-7}$ m), $\omega r_0/c \sim 3 \times 10^{-4}$ and $T \sim 10^{-15}$. The washout suppression is weakened relative to the EM-based estimate but remains far below any level at which re-equilibration could operate.

We note that for massive fermions ($m > 0$), the effective potential (46) acquires an additional mass term m^2 that raises the barrier and strengthens the suppression. For Standard Model neutrinos ($m_\nu \lesssim 0.1$ eV), the mass contribution is negligible compared to the geometric barrier $1/r_0^2$ for all submicron throat radii, so the massless result above applies directly. For heavier fermions (electrons, quarks), the mass term further suppresses transmission, making the massless s-wave neutrino the most conservative (least-suppressed) case.

Even in the hypothetical limit where the $j = \frac{1}{2}$ mode is transmitted freely ($T = 1$), the cosmological impact would be bounded by the N_{eff} constraint (Sec. IV)—any transmitted neutrino flux from the sequestered sector contributes at most $\Delta N_{\text{eff}} \sim (T_{\text{seq}}/T_{\text{vis}})^4 \times 3.046$, which is already constrained to be < 0.34 . The actual transmission factor ($\sim 10^{-15}$ or below for submicron throats) renders this channel negligible by a further fifteen or more orders of magnitude.

This strengthens the proposal—sequestered baryons are completely dark to all propagating radiation, coupling to our sector exclusively through the static gravitational potential. The only possible exception arises for very large throats ($r_0 \gtrsim 10$ km, giving $f_{\max}^{(\text{GW})} \lesssim 10^3$ Hz), where some GW frequencies could exceed the barrier top and propagate through. However, such large throats would also transmit optical and radio EM radiation, making the sequestered matter visible.

3. Propagation of visible-sector GWs through a throat-populated universe

A separate question is whether visible-sector GWs— for instance, those from binary mergers observed by LIGO-Virgo-KAGRA, or the stochastic background from cosmological sources—are attenuated as they propagate through a universe populated by submicron throats along their path. The answer is that they are not, to any observationally relevant precision. A wave of wavelength λ is effectively insensitive to scatterers of size $r_0 \ll \lambda$ —the differential scattering cross-section of a compact obstacle in this regime is Rayleigh-suppressed by a factor $(r_0/\lambda)^4$ relative to the geometric cross-section [61]. For LIGO-band GWs ($\lambda \sim 3000$ km) and submicron throats ($r_0 \sim 10^{-7}$ m), $(r_0/\lambda)^4 \sim 10^{-54}$, giving a per-throat scattering cross-section of order $r_0^2 (r_0/\lambda)^4 \sim 10^{-68}$ m². The cosmological-average throat number density is

$$n_{\text{throat}} \sim \rho_{\text{DM}}/M_{\text{throat}} \sim 2 \times 10^{-47} \text{ m}^{-3}$$

for the favoured asteroid-mass regime of Sec. VII E ($M_{\text{throat}} \lesssim 5 \times 10^{-11} M_\odot$), giving an inverse mean free

path of order $3 \times 10^{-115} \text{ m}^{-1}$ and a mean free path of order $3 \times 10^{114} \text{ m}$, more than 88 orders of magnitude larger than the Hubble radius. The corresponding optical depth to scattering across a Hubble volume is $\tau \sim 4 \times 10^{-89}$, far below any conceivable detection threshold. A throat-populated universe is optically thin to the entire astrophysical and cosmological GW spectrum — stochastic GW background measurements by LIGO, LISA, and pulsar timing arrays place no useful constraint on the scenario and are unaffected by the presence of submicron throats along propagation paths.

4. *Non-appearance of sequestered compact-object mergers*

The scenario also predicts no GW detection from mergers occurring within the sequestered sector, but this prediction has two independent structural origins and is not by itself a distinguishing signature of the proposal. First, as established above, GWs emitted in the sequestered sector cannot propagate through submicron throats to reach a visible-sector detector — the axial $\ell = 2$ mode faces the same centrifugal barrier as EM modes, and the sub-barrier suppression is essentially complete at LIGO and LISA frequencies. Second, and independently, the bulk of the sequestered matter is expected to be too cold to form compact objects. At the cosmic mean density, the cooling time exceeds the Hubble time by ten orders of magnitude (Sec. VIF), and the sequestered gas remains diffuse and pressure-supported. Compact-object formation is not excluded in overdense regions where cooling becomes efficient ($\delta \gtrsim 10^3$), but the fraction of sequestered mass affected is expected to be small (Sec. VIF). With sequestered-sector temperatures $T_{\text{seq}} \lesssim 100 \text{ K}$ (Sec. VI), the sequestered gas is well below the hydrogen cooling threshold ($\sim 10^4 \text{ K}$ at which atomic hydrogen radiative transitions become efficient [54]), and the dissipative cooling required to form stars, neutron stars, or black holes through gravitational collapse is severely suppressed. The sequestered matter is expected to remain diffuse and pressure-supported rather than forming stellar or compact remnants (Sec. VIF). Both effects — barrier suppression of GW transmission and absence of compact-object formation — predict a null result in the “GW from the dark sector” channel, and both are structural consequences of the scenario rather than phenomenological claims about small event rates.

However, this null prediction is not a distinguishing signature against standard cold dark matter candidates. Particle CDM (WIMPs, axions, sterile neutrinos) likewise predicts no compact-object mergers and no dark-sector GW background, because particle dark matter does not form compact objects either. The GW channel therefore cannot be used to discriminate between the sequestered-baryon scenario and particle CDM — both predict the same null result, albeit for very different structural reasons. The scenario’s observationally distinguishing signatures lie instead in the structural gravita-

tional observables discussed in the remainder of this section — the Jeans-scale cutoff in small-scale clustering, the absence of dark-sector substructure, the smooth halo profiles, and the polynomial lensing correction at the throat scale.

C. **The Bullet Cluster**

The Bullet Cluster (1E 0657-56) [6] exhibits a spatial offset of order $\sim 200 \text{ kpc}$ between the gravitational lensing signal (tracing total mass) and the X-ray emission (tracing hot gas), establishing that the dominant component of the cluster mass is collisionless on Mpc scales. In our scenario, the sequestered baryons satisfy this constraint structurally rather than through a small interaction cross-section. Two effects are responsible.

First, sequestered baryons in different clusters cannot interact electromagnetically with one another: the constraint-wave asymmetry blocks all $\ell \geq 1$ propagation through the throat, so the sequestered components of two colliding clusters cannot exchange momentum via radiation pressure or Coulomb scattering. The only sector-to-sector coupling is the static gravitational monopole, which carries no momentum flux.

Second, sequestered baryons in different clusters are also gravitationally collisionless on the relevant scales. The sequestered population associated with each cluster sits behind that cluster’s own throat network; sector-to-sector gravitational coupling is mediated by the $\mathcal{O}(r_0/d)$ monopole transmission of Sec. II, which for submicron throats and any astrophysical separation d is fractionally below 10^{-19} . The two sequestered populations therefore pass through each other with no measurable interaction during the $\sim 200 \text{ Myr}$ crossing time of the collision — exactly the behaviour expected of CDM, and consistent with the observed lensing-vs-gas offset.

The argument requires that each cluster’s sequestered baryons be gravitationally bound to the throats associated with that cluster, rather than being distributed across a globally shared network. This is the natural expectation for any throat-formation mechanism that imprints throats from primordial density peaks (Sec. VIE), but a quantitative analysis of the throat topology of merging clusters — and the corresponding lensing-vs-gas offset prediction — would require the formation model of WP5 (Sec. VIII) and is left to future work. Within the conditional framework of this paper, the Bullet Cluster offset is reproduced by the same structural collisionlessness that yields the CDM phenomenology elsewhere.

D. **21-cm cosmology**

The 21-cm hyperfine transition of neutral hydrogen provides a probe of the dark ages ($z \sim 30\text{--}200$) and the epoch of reionisation ($z \sim 6\text{--}20$) [62]. In the geometric-sequestration scenario, the sequestered baryons include

neutral hydrogen that emits 21-cm radiation—but this radiation cannot cross the throat. The 21-cm signal observed by experiments (e.g. HERA [63], SKA [64], and EDGES [65]) would therefore reflect only the visible hydrogen, with the gravitational effects of the sequestered baryons appearing as enhanced structure in the 21-cm fluctuation power spectrum.

If the sequestered baryons have a different thermal history from the visible sector (as required by the N_{eff} constraint, Sec. IV), their 21-cm emission would be at a different frequency (redshifted differently by the local expansion rate within the sequestered region) and would not contribute to the observed 21-cm signal. This is consistent with current upper limits on extra-galactic 21-cm backgrounds.

E. Gravitational lensing and microlensing constraints

Each throat mouth acts as a point-mass gravitational lens. We derive this from the weak-field deflection integral. The gravitational potential of total mass M behind the throat is

$$\Phi(\sigma) = -GM/a(\sigma) = -GM/\sqrt{\sigma^2 + r_0^2},$$

which reduces to $-GM/r$ at $r \gg r_0$, but is regularised at the throat. For a null geodesic at impact parameter b in the weak-field regime, the deflection angle from the standard integral $\alpha = (2/c^2) \int \nabla_{\perp} \Phi d\ell$ (doubled for the spatial metric contribution) evaluates to

$$\alpha = \frac{4GMb}{c^2(b^2 + r_0^2)}. \quad (50)$$

The ratio to the point-mass result $\alpha_{\text{PM}} = 4GM/(bc^2)$ is

$$\frac{\alpha}{\alpha_{\text{PM}}} = \frac{1}{1 + (r_0/b)^2}, \quad (51)$$

which approaches unity for $b \gg r_0$. For submicron r_0 and astrophysical impact parameters ($b \sim \text{AU}$ for microlensing), $r_0/b \sim 10^{-24}$ and the correction is negligible at $\sim 10^{-48}$. Eq. (50) also shows that the deflection is bounded at the throat ($b \rightarrow r_0$), avoiding the point-mass divergence, but this regime is irrelevant for astrophysical lensing.

1. Robustness beyond the Ellis-Bronnikov model

Although Eq. (50) was derived on the specific EB profile $a(\sigma) = \sqrt{\sigma^2 + r_0^2}$, its leading behaviour in the astrophysically relevant regime $b \gg r_0$ is universal across throat geometries. This follows from matched asymptotics combined with the multipole structure established in Sec. II. At distances $b \gg r_0$, the gravitational field

on our side of the throat is dominated by its $\ell = 0$ component, because the source-side multipoles with $\ell \geq 1$ are suppressed by the scaling (6), $(r_0/b)^{2\ell+1}$, relative to the monopole. The monopole itself is determined by Gauss's law, and on any throat background with $e^{2\alpha} > 0$ at the throat the conserved flux gives $\Phi(r) \rightarrow -GM/r$ at $r \gg r_0$, with corrections of order r_0/r [1]. Substituting this far-field form into the standard weak-field deflection integral then reproduces the point-mass result $\alpha_{\text{PM}} = 4GM/(bc^2)$ with fractional corrections of order $(r_0/b)^2$, regardless of the specific profile $a(\sigma)$. The $1/(1 + (r_0/b)^2)$ resummation visible in Eq. (50) is an EB-specific feature of the all-orders matching but is irrelevant at astrophysical impact parameters where r_0/b is many orders of magnitude below any detectable threshold. The matched-asymptotics argument also makes explicit that the lensing formula is insensitive to the detailed structure of the sequestered source—the sequestered sector could contain disks, compact objects, dense substructure, or any other matter configuration, and only its total enclosed mass contributes to the lensing signal at astrophysical impact parameters.

2. Internal structure of sequestered sources

The same multipole suppression implies a qualitative prediction about the shape of dark-matter lensing and dynamical signals. Within a single sequestered region, baryonic matter can organise itself dissipatively into structures with nontrivial multipole moments—for example, rotating disks, globular-cluster-like condensations, or compact remnants—but the higher multipoles of these structures are filtered by Eq. (6) before they reach our sector. A sequestered rotating disk of quadrupole moment Q projects onto our side a fractional amplitude $(r_0/d)^5$ relative to its monopole, which for $r_0 \lesssim 10^{-6}$ m and any astrophysical source distance $d \gtrsim 1$ AU is below 10^{-130} . The sequestered sector therefore appears to us, via the gravitational channel alone, as a collection of pointlike monopoles located at the centres of the throat mouths, even if the sequestered matter has elaborate internal structure. This is a feature of the scenario: it is the reason smooth Navarro-Frenk-White (NFW)-like halo profiles and scale-free lensing predictions survive the sequestered-baryon hypothesis even if sequestered galaxies are internally clumpy. It also suggests a distinguishing observational test—sequestered matter cannot produce the tidal or higher-multipole lensing signatures that internally structured luminous matter would, so any robust detection of galaxy-scale anisotropic lensing shear tracing clumpy dark substructure would rule out geometric sequestration as the dominant dark matter channel.

The mass per throat mouth is $M_{\text{throat}} = M_{\text{halo}}/N_{\text{throat}}$, where N_{throat} is the number of throats per halo. Microlensing surveys (EROS [66], MACHO, OGLE [67], Subaru/HSC [68]) and quasar microlensing/millilensing

studies constrain the fraction of dark matter in compact objects as a function of mass. The composite constraints [69] leave only two allowed windows for $f_{\text{DM}} = 1$: the asteroid-mass window at $M \sim 5 \times 10^{-18}$ to $5 \times 10^{-11} M_{\odot}$, and a high-mass window at $M \sim 10^{14}$ to $5 \times 10^{18} M_{\odot}$ that corresponds to less than one object per MW-scale halo and is therefore not physically realised for halo-bound throats (Fig. 7).

The scenario requires M_{throat} to lie within one of the allowed windows of Fig. 7. The high-mass window ($M \sim 10^{14}$ to $5 \times 10^{18} M_{\odot}$) corresponds to fewer than one object per MW-scale halo and is not physically realised for halo-bound throats. The asteroid-mass window ($M \sim 5 \times 10^{-18}$ to $5 \times 10^{-11} M_{\odot}$) spans seven orders of magnitude in throat mass and corresponds to throat numbers

$$2 \times 10^{22} \lesssim N_{\text{throat}} \lesssim 2 \times 10^{29} \quad (\text{per MW-like halo}). \quad (52)$$

The viable regime is therefore a “many light throats” scenario with per-throat mass below the Subaru/HSC microlensing limit and above the Hawking-evaporation edge. Scenarios with smaller N_{throat} — including the “few massive throats” regime ($N \sim 10^6$, $M_{\text{throat}} \sim 10^6 M_{\odot}$) sometimes considered for strongly-lensed dark-matter populations — are excluded by background-effects constraints (CMB μ -distortion, second-order gravitational waves, and BBN) by more than ten orders of magnitude in f_{DM} , and are not viable within the geometric dark-matter scenario.

Combining the microlensing lower bound (52) with the exotic-energy budget (10) of Sec. III places a non-trivial joint constraint on the throat radius:

$$r_0 \lesssim 10^{-7} \text{ m} \quad \text{for } N_{\text{throat}} = 2 \times 10^{22}, \quad (53)$$

tightening for larger N : $r_0 \lesssim 2 \times 10^{-10}$ m at $N_{\text{throat}} = 10^{25}$ and $r_0 \lesssim 2 \times 10^{-13}$ m at $N_{\text{throat}} = 10^{28}$. The microlensing constraint therefore selects throat radii spanning roughly seven orders of magnitude, from $r_0 \sim 10^{-7}$ m (submicron) at $N_{\text{throat}} = 2 \times 10^{22}$ down to $r_0 \sim 10^{-14}$ m (sub-picometre) at $N_{\text{throat}} \sim 2 \times 10^{29}$, still many orders of magnitude above the Planck length and mapping onto distinct formation scales from near the submicron range down to well above the electroweak epoch. This gives a concrete, falsifiable prediction: any future tightening of the Subaru/HSC microlensing bound, the femtolensing limit, or the evaporation edge that closes the asteroid-mass window would rule out the geometric dark-matter hypothesis entirely.

The inter-throat spacing across the halo in this regime is $d \sim (N/V_{\text{halo}})^{-1/3} \sim 10^{-5} - 10^{-2}$ pc, small compared to galactic scale-heights, so the sequestered mass distribution is effectively smooth on all scales probed by current lensing and dynamical observations.

3. Event rates and timescales: quantitative comparison with surveys

The Carr et al. [69] composite plot summarises the bounds derived by individual surveys; here we recast the analysis in terms of the directly observed quantities (event rate and event timescale) for the throat scenario, and confront the result with the published rate limits from each survey separately. The standard microlensing optical depth toward a source at distance D_s is

$$\tau = \frac{4\pi G}{c^2} \int_0^{D_s} \rho(D_l) \frac{D_l(D_s - D_l)}{D_s} dD_l, \quad (54)$$

which depends only on the dark-matter mass density ρ along the line of sight and is independent of the per-throat mass M_{throat} — a generic feature of point-mass lensing that holds equally well for our throat-mouth lenses, by the matched-asymptotics argument of Sec. VII E. For the canonical Milky-Way halo and the Large Magellanic Cloud line of sight ($D_s \approx 50$ kpc), the standard result is $\tau_{\text{LMC}} \approx 5 \times 10^{-7}$ [70]; HSC observations toward M31 yield a similar value. The mean event rate per source is then

$$\Gamma = \frac{2\tau}{\pi \langle t_E \rangle}, \quad t_E = \frac{R_E}{\langle v_T \rangle}, \quad R_E = \sqrt{\frac{4GM_{\text{throat}}}{c^2} \frac{D_l(D_s - D_l)}{D_s}}, \quad (55)$$

with the Einstein radius R_E and characteristic transverse velocity $\langle v_T \rangle \sim 200$ km/s evaluated at the canonical lens distance $D_l = D_s/2$. The event rate per source therefore scales as $\Gamma \propto M_{\text{throat}}^{-1/2}$, and the timescale as $t_E \propto M_{\text{throat}}^{1/2}$.

Table VI evaluates these quantities at three mass points spanning the asteroid-mass window from (52), assuming $f_{\text{DM}} = 1$ and the LMC line of sight.

TABLE VI. Predicted microlensing event timescales and rates for sequestered-baryon throats across the asteroid-mass window ($f_{\text{DM}} = 1$, LMC line of sight, $D_l = D_s/2 = 25$ kpc, $\langle v_T \rangle = 200$ km/s). The rate per 10^7 stars per year is for direct comparison with EROS-2 ($\sim 7 \times 10^6$ stars) and OGLE ($\sim 10^7$ stars). All three mass points produce predicted event timescales below the cadence of any current ground-based survey (1 day) and below the HSC fast-cadence limit of 90 s for $M_{\text{throat}} < 1.4 \times 10^{-10} M_{\odot}$, so although the predicted rates are large, the events are not detectable.

$M_{\text{throat}} (M_{\odot})$	R_E (m)	t_E (s)	Γ ($10^{-7} \text{ src}^{-1} \text{ yr}^{-1}$)
5×10^{-18}	3.4×10^3	1.7×10^{-2}	6.0×10^9
1×10^{-14}	1.5×10^5	7.5×10^{-1}	1.3×10^8
5×10^{-11}	1.1×10^7	5.3×10^1	1.9×10^6

The numbers show that the asteroid-mass window is unconstrained by the kinematic reason often emphasised in the PBH literature: the predicted event rates per

source are large for $f_{\text{DM}} = 1$, but the predicted timescales are below the temporal sensitivity of all current surveys. EROS-2 [66], MACHO [71], OGLE-IV [67], and Subaru/HSC [68] all become insensitive to events with t_E shorter than their respective sampling cadences—typically one day for ground-based surveys (corresponding to $M_{\text{throat}} \gtrsim 1.3 \times 10^{-4} M_{\odot}$) and 90 s for HSC fast-cadence M31 monitoring ($M_{\text{throat}} \gtrsim 1.4 \times 10^{-10} M_{\odot}$). The throat scenario predicts no detectable events at any of the three tabulated mass points, in agreement with the non-detection by every survey at $M \lesssim 5 \times 10^{-11} M_{\odot}$ —which is precisely the upper edge of the asteroid window. Future cadence improvements (Roman Space Telescope at ~ 15 -min cadence, or proposed millisecond-cadence GRB-monitoring instruments) would directly probe the upper end of the asteroid window from below.

a. Wave-optics regime A sharper observation is that the entire asteroid-mass window lies in the wave-optics regime at optical wavelengths. The dimensionless wave-optics parameter is

$$w \equiv \frac{8\pi G M_{\text{throat}} \nu_{\text{phot}}}{c^3},$$

with $\nu_{\text{phot}} \sim 6 \times 10^{14}$ Hz at visible wavelengths. Geometric optics applies for $w \gg 1$; wave optics dominates for $w \ll 1$. At our three mass points, $w \approx 4 \times 10^{-7}$, 7×10^{-4} , and 4 respectively. Only the upper edge of the asteroid window approaches the geometric-optics limit; the bulk of the window is in the deep wave-optics regime, where the on-axis magnification is bounded by [72]

$$\mu_{\text{max}} - 1 \approx \frac{(\pi w)^2}{3} \quad \text{for } w \ll 1,$$

giving $\mu_{\text{max}} - 1 \approx 5 \times 10^{-13}$ at $M = 5 \times 10^{-18} M_{\odot}$ and $\approx 2 \times 10^{-6}$ at $M = 10^{-14} M_{\odot}$. These deviations from unit magnification are far below the photometric sensitivity of any current or planned imaging survey, so even if the cadence were short enough to resolve the predicted timescale, the events at the lower and middle parts of the asteroid window would be photometrically undetectable. Only at the upper edge of the window ($M \sim 5 \times 10^{-11} M_{\odot}$, $w \sim 4$) does the magnification reach $\mathcal{O}(1)$ —the regime where, by construction, the HSC cadence becomes the binding constraint.

b. Throat-specific effects Two features distinguish throat lenses from idealised PBH lenses, neither of which strengthens the constraint:

- The deflection-angle cap of Eq. (50) matters only when the impact parameter b is comparable to r_0 . For all of $R_E \gtrsim 10^3$ m (Table VI) and viable throat radii $r_0 \leq 10^{-7}$ m, the ratio r_0/R_E is at most 10^{-10} , and the correction to the deflection at $b \sim R_E$ is below 10^{-20} . The cap is observationally irrelevant for microlensing.
- Finite-throat-radius corrections to the wave-optics magnification become relevant when r_0 is compa-

rable to the optical wavelength. For the largest viable throats ($r_0 \sim 10^{-7}$ m), r_0 approaches $\lambda_{\text{phot}} \sim 5 \times 10^{-7}$ m, and an additional finite-extent diffraction correction of order $(r_0/\lambda)^2 \sim 0.04$ enters the magnification calculation. For smaller throats this correction vanishes. Crucially, this is a small modulation of an already-undetectable signal (the wave-optics magnification at this mass scale is already $\sim 1 + w \sim 5$, marginal to detect even with infinite cadence): the throat-specific correction does not produce a new observational signature within current survey sensitivities.

The microlensing analysis therefore confirms the asteroid-mass window of Eq. (52) as the viable parameter space, with the conclusion robust against both the kinematic argument (timescale below cadence) and the wave-optics argument (magnification below detection threshold). Quantitative confrontation with future short-cadence surveys remains an important falsifiability test of the scenario; we identify this as the single most direct observational route to ruling out (or detecting) the geometric dark-matter hypothesis (Sec. VIII, WP5).

c. Halo granularity constraints A distributed population of $N_{\text{throat}} \gtrsim 2 \times 10^{22}$ point masses per halo could in principle produce observable stochastic signatures—Poisson fluctuations in the local dark matter density, dynamical heating of stellar populations, and timing noise in pulsar timing arrays—that are absent in a genuinely smooth CDM distribution. We estimate each of these in turn and find that all are comfortably below current sensitivities for the regime (52).

d. Poisson density fluctuations The fractional shot-noise variance in the sequestered mass density, averaged over a coarse-graining volume V_{smooth} , is

$$\sigma_{\rho}/\rho = 1/\sqrt{n_{\text{throat}} V_{\text{smooth}}},$$

where $n_{\text{throat}} = N_{\text{throat}}/V_{\text{halo}}$ is the halo-averaged throat number density. For the lower-count bound $N_{\text{throat}} = 2 \times 10^{22}$ in a MW-like halo ($V_{\text{halo}} \approx 10^{66} \text{ m}^3$), the mean throat density is $n_{\text{throat}} \approx 2 \times 10^{-44} \text{ m}^{-3}$, and

$$\begin{aligned} \sigma_{\rho}/\rho &\approx 1 \times 10^{-3} & (V_{\text{smooth}} = 1 \text{ pc}^3), \\ \sigma_{\rho}/\rho &\approx 4 \times 10^{-4} & (V_{\text{smooth}} = 10 \text{ pc}^3), \\ \sigma_{\rho}/\rho &\approx 1 \times 10^{-4} & (V_{\text{smooth}} = 100 \text{ pc}^3). \end{aligned} \quad (56)$$

At stellar-stream coarse-graining scales (~ 10 – 100 pc, the widths of observed cold streams such as GD-1), density fluctuations are at the 10^{-4} level, comfortably below the few-percent sensitivity of current stream-perturbation analyses [73, 74]. For the upper-count bound $N_{\text{throat}} = 2 \times 10^{29}$, the fluctuations are smaller still by a factor of $\sqrt{10^7} \approx 3000$.

e. Dynamical heating The energy injection rate into a stellar population from encounters with dark perturbers of mass M_{pert} and number density n_{pert} scales as $dE/dt \propto G^2 M_{\text{pert}}^2 n_{\text{pert}} \ln \Lambda/v_{\text{rel}}$ [75]. Comparing to

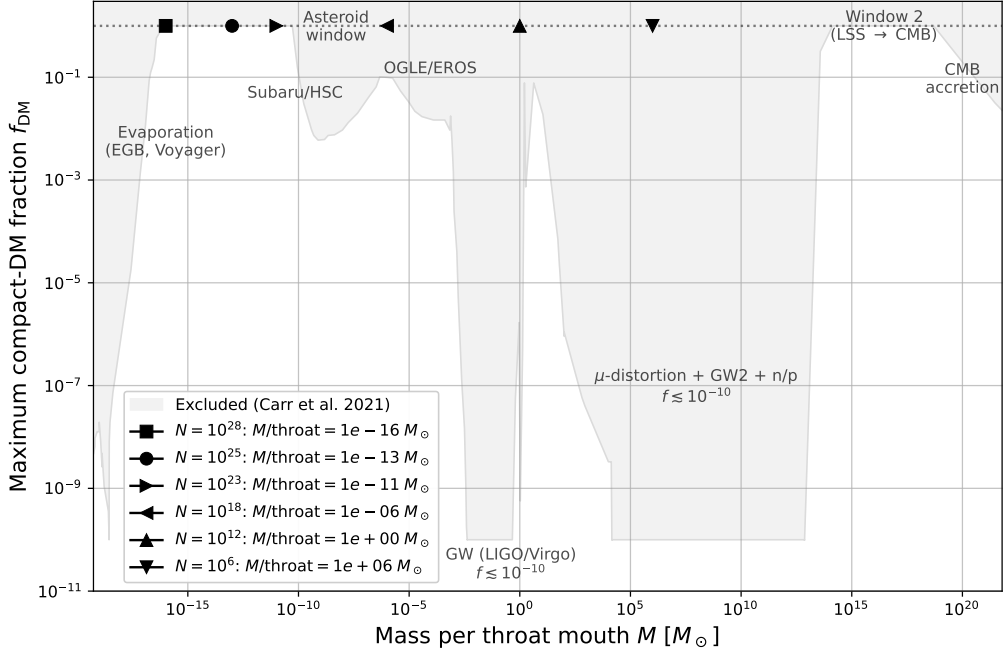


FIG. 7. Observational constraints on the compact dark-matter fraction f_{DM} as a function of individual object mass M , based on Figure 10 of Carr, Kohri, Sendouda, and Yokoyama (2021) [69]. The shaded region is excluded by the combined bounds from evaporation (EGB, Voyager), microlensing (Subaru/HSC, OGLE, EROS, MACHO, Kepler), gravitational waves (LIGO-Virgo-KAGRA), CMB μ -distortions from FIRAS, second-order gravitational waves, the BBN neutron-to-proton ratio, CMB accretion anisotropies, dynamical constraints on stars and clusters, and large-scale structure. Two allowed windows are visible: the asteroid-mass window ($M \sim 5 \times 10^{-18}$ to $5 \times 10^{-11} M_{\odot}$) and a high-mass window between the background-effects and CMB-accretion bounds ($M \sim 10^{14}$ to $5 \times 10^{18} M_{\odot}$). Markers show the mass per throat mouth $M/\text{throat} = M_{\text{halo}}/N$ for a Milky-Way-like halo ($M_{\text{halo}} = 10^{12} M_{\odot}$) at several values of the throat number N . Scenarios with $N \lesssim 10^{18}$ lie in the excluded region; scenarios with $N \gtrsim 2 \times 10^{22}$ fall within the asteroid-mass window and are observationally allowed. The incredulity limit of Carr et al., which corresponds to fewer than one object per Hubble volume, is not shown: it is a geometric cosmological limit rather than an observational constraint, and does not apply to geometric dark-matter throats, which are bound to individual halos with integer multiplicity N .

the heating rate expected from the $\sim 10^5$ subhalos of mass $\sim 10^8 M_{\odot}$ predicted by Λ CDM simulations in a MW-like halo, we find the ratio

$$\begin{aligned} \frac{(dE/dt)_{\text{throats}}}{(dE/dt)_{10^8 M_{\odot} \text{ subhalos}}} &\sim \frac{N_{\text{throat}} M_{\text{throat}}^2}{N_{\text{subhalo}} M_{\text{subhalo}}^2} \\ &= \frac{M_{\text{halo}} M_{\text{throat}}}{N_{\text{subhalo}} M_{\text{subhalo}}^2} \\ &\lesssim 5 \times 10^{-20} \end{aligned} \quad (57)$$

for the microlensing upper bound $M_{\text{throat}} \sim 5 \times 10^{-11} M_{\odot}$ from (52), with the ratio decreasing monotonically for smaller per-throat masses (larger N_{throat}). The scaling $M_{\text{pert}}^2 n_{\text{pert}} \propto M_{\text{throat}} M_{\text{halo}}$ favours few massive perturbers over many light ones, so the many-light-throats regime is strongly preferred by dynamical-heating arguments independently of the lensing bound. The dynamical heating contribution from the throat population is therefore negligible compared to the heating from the already-present CDM substructure population—and in

any case well below the observational limits on disk thickening and stream stirring, which already accommodate the CDM substructure contribution.

f. Pulsar timing arrays A mass $M_{\text{throat}} \lesssim 5 \times 10^{-11} M_{\odot}$ produces a per-encounter Shapiro-delay timing residual of order $\delta t \sim GM/c^3 \lesssim 5 \times 10^{-16}$ s, roughly seven orders of magnitude below the per-pulse timing sensitivity of current pulsar timing array experiments [76]. Even accounting for the high encounter rate from the dense throat population, the cumulative stochastic timing noise is far below detection thresholds. Gravitational wave backgrounds from the throat population (if any such mechanism operated) would face the Rayleigh cross-section suppression quantified in Sec. VII and are similarly unobservable.

g. Cosmological shot noise in the matter power spectrum A population of discrete point-mass perturbers contributes a scale-independent shot-noise term $P_{\text{shot}} = 1/\bar{n}_{\text{throat}}$ to the total matter power spectrum,

where $\bar{n}_{\text{throat}} = \rho_{\text{DM}}/M_{\text{throat}}$ is the cosmic-mean throat number density. For the microlensing upper bound $M_{\text{throat}} = 5 \times 10^{-11} M_{\odot}$,

$$\begin{aligned}\bar{n}_{\text{throat}} &\approx 1.4 \times 10^{21} (h/\text{Mpc})^3, \quad \text{giving} \\ P_{\text{shot}} &\approx 7 \times 10^{-22} (\text{Mpc}/h)^3.\end{aligned}$$

This is approximately twenty-one orders of magnitude below the linear Λ CDM matter power at Lyman- α scales ($P_{\text{lin}}(k \sim 10 h/\text{Mpc}) \sim 1 (\text{Mpc}/h)^3$ at $z = 0$), so the Poisson contribution to the total power is negligible and cannot be constrained by any current or foreseeable large-scale-structure survey. For smaller per-throat masses—i.e., larger N_{throat} within the range (52)—the shot noise is smaller still.

The many-light-throats regime is therefore consistent not only with the microlensing constraints that define the range (52), but also with all current observational probes of halo granularity—density fluctuations, dynamical heating, and timing noise—precisely because the per-throat masses are small enough that individual encounters are undetectable while the collective Poisson noise is suppressed by the enormous throat number. This “many light throats” scenario is self-consistently safe against all currently accessible halo-substructure observables.

Within the halo, throats are presumably concentrated in higher-density regions, tracing the sequestered-matter distribution rather than filling the halo volume uniformly. We do not attempt to model this spatial distribution in this paper—the combination of Eq. (52) with the NFW-like sequestered-matter profile inferred from rotation curves and weak lensing constitutes an effective description sufficient for the leading-order predictions discussed in this section.

h. Formation mechanism The physical process that might produce throat numbers in the range (52) remains an open question, and we do not resolve it here—a quantitative formation model is deferred to future work (Sec. VIII). We sketch two broad classes of candidate mechanisms to indicate the kind of work required:

(i) *Primordial density fluctuations*

If throats form from the collapse of large-amplitude primordial density fluctuations at a specific cosmological epoch—analogueous to the primordial-black-hole formation mechanism but producing geometric bottlenecks rather than horizons—the number density is set by the amplitude and spectral index of the fluctuations on the relevant scale. A quantitative model specifying the initial conditions, the dynamical collapse mechanism, and the resulting throat mass function is a central open problem of the scenario (Sec. VIII).

(ii) *Phase-transition production*

Cosmological phase transitions at the QCD scale ($\xi \sim 1 \text{ fm}$) or higher produce topological defects at a density set by the correlation length, which after

cosmological dilution would give throat counts potentially far exceeding (52). The exotic-energy budget of Eq. (10) rules out the naïve phase-transition densities on cosmological grounds: a QCD-scale formation scenario producing $n_0 \sim 4 \times 10^9 \text{ m}^{-3}$ after dilution would give $N \sim 10^{75}$ throats per halo and an exotic-energy budget $\sim 5 \times 10^{52}$ times the halo rest mass, a catastrophic violation of the backreaction constraint. Phase-transition scenarios can contribute viable throat populations only if the formation efficiency is suppressed by many orders of magnitude relative to the naïve estimate $n_{\text{throat}} \sim \xi^{-3}$ —for example, by an efficiency factor $\epsilon \lesssim 2 \times 10^{-53}$ selecting only rare fluctuations above a high threshold, or by a formation process that is active only during a narrow cosmological window.

Neither mechanism has been worked out quantitatively, and a complete throat formation model—including the mass function and the spatial distribution within dark-matter haloes—remains an important open problem (Sec. VIII).

Within the asteroid-mass window (52), femtolensing of gamma-ray bursts [77] provides the most direct constraint in the range 10^{-17} – $10^{-13} M_{\odot}$, but the bounds are model-dependent and do not close the window at $f_{\text{DM}} = 1$ [69]. The inclusion of the asteroid-mass window in Carr et al. [69] already accounts for these femtolensing limits. Dynamical heating constraints from stellar streams and disk thickening [78] apply primarily at $M \gtrsim 10^6 M_{\odot}$ and do not affect the “many light throats” regime. The convergence κ and shear γ produced by sequestered baryons would be identical to those produced by CDM of the same mass distribution. Galaxy-galaxy lensing, cluster lensing, and cosmic shear measurements should therefore be consistent with Λ CDM predictions, provided the spatial distribution of sequestered matter matches the CDM profile.

A distinguishing prediction arises if the throat-connected regions have a characteristic spatial scale. In that case, the lensing signal of “dark” matter would show a preferred scale in the convergence power spectrum, distinguishing it from the scale-free CDM prediction. Detecting such a feature would require high-precision weak lensing surveys at the level of the Euclid space mission [79] or the ground-based Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) [80, 81].

F. Structural impossibility of direct detection

The sequestered-baryon scenario makes a strong claim about direct dark matter searches: their null results are not a phenomenological consequence of small interaction cross sections but a structural consequence of the geometry. Particle dark matter searches—XENON [82], LUX-ZEPLIN [83], PandaX [84], axion haloscopes, and their descendants—look for interactions between detector nuclei (or photons, in the axion case) and dark matter par-

ticles residing in the same spatial region as the detector. In the sequestered-baryon scenario, there are no such particles—the dark matter is ordinary baryonic matter located in spatial regions that our detectors cannot access, separated from our region by geometric throats through which no matter field with angular momentum $\ell \geq 1$ can propagate.

The argument extends the constraint-wave asymmetry of Sec. II from massless fields (EM, GW) to massive matter fields. A massive Klein-Gordon or Dirac field ϕ on a throat background satisfies a wave equation of the same structural form as the EM and GW sectors—the angular part of the Laplacian produces a centrifugal term $\ell(\ell + 1)/a^2$ that dominates near the throat for every $\ell \geq 1$ multipole of the field. The effective potential for a mode of mass m and multipole ℓ takes the form

$$V_\ell^{(m)}(\sigma) = e^{2\alpha} \left[\frac{\ell(\ell + 1)}{a^2} + m^2 \right] + \text{curvature corrections}, \quad (58)$$

which for $\ell \geq 1$ is peaked at the throat and produces sub-barrier suppression identical in structure to the massless case. A sequestered dark matter “particle” of any mass m , modelled as a wave packet constructed from modes of this field, faces the same exponential suppression as an EM photon when attempting to traverse the throat—the transmission coefficient $T \sim \exp(-2S)$, where S is the WKB tunnelling action, is far below any conceivable detection threshold for submicron r_0 and any physically plausible particle mass. The $\ell = 0$ mode of a massive field is additionally suppressed by the mass term itself, which prevents the conservation-law argument that works for the static gravitational monopole from applying to matter fields: $(a^2\phi')' = m^2 a^2\phi$ is not a conservation law when $m \neq 0$.

The scenario therefore predicts that all direct dark matter detection experiments, independent of improvements in sensitivity, will return null results. This is a sharper prediction than the WIMP scenario’s “interactions exist but are very small,” and it is the feature that most distinguishes the geometric-sequestration proposal from particle dark matter candidates. The only channels through which the scenario can be confirmed or refuted are gravitational ones: small-scale structure and the Jeans-scale cutoff in the matter power spectrum (Sec. VI), anomalies in gravitational lensing at the throat scale (Sec. VII E), frequency-dependent electromagnetic leakage for throats near the gamma-ray opacity threshold (Sec. VII), and—in principle—any precision test of Newtonian gravity that could detect the $\mathcal{O}(r_0/d)$ polynomial correction to the static monopole potential near an isolated throat. A direct detection of dark matter particles in any conventional experiment would falsify our scenario.

G. Summary of predictions

The geometric-sequestration scenario makes the following qualitative predictions:

- (i) no direct detection of dark matter particles (because there are none);
- (ii) frequency-dependent electromagnetic leakage if r_0 is not too small (a spectral signature distinct from CDM);
- (iii) a small-scale cutoff in the matter power spectrum at the Jeans scale of the sequestered baryons (potentially addressing CDM small-scale tensions);
- (iv) standard gravitational lensing signals on large scales;
- (v) gravitational waves from the sequestered sector are suppressed for submicron r_0 , along with all EM radiation; only the static gravitational monopole couples to our sector.

Falsification is possible if (a) dark matter particles are directly detected, (b) the matter power spectrum shows no small-scale cutoff and CDM-like clustering extends to arbitrarily small scales, or (c) precision CMB/LSS data require features that are incompatible with a sequestered-baryon component.

VIII. A research programme

The preceding sections have laid out a conditional proposal: if submicron throat geometries exist in the spatial fabric of the universe, and if the sequestered sector can be arranged to be sufficiently cold at recombination, then ordinary baryonic matter behind such throats reproduces the observational phenomenology of cold dark matter across BBN, the CMB, large-scale structure, and lensing. Neither of these conditional premises has been established. What we have demonstrated is that neither is excluded by current observations, and that the scenario as a whole constitutes a well-defined research programme whose viability can be tested by addressing a modest number of concrete physical questions. In this section we lay out that programme as a structured agenda, grouped into five work packages corresponding to the five main open questions raised by the preceding analysis. We order the packages by logical dependency rather than by priority: WP1 and WP4 address the two premises (existence and thermal history), while WP2, WP3, and WP5 develop the phenomenological consequences that would become rigorously testable once those premises are addressed.

Specifically, the cosmological analysis of Sections IV–VII rests on the following assumptions, each of which is addressed by one of the work packages described below:

1. *Throat existence*—that the spatial geometry of our universe contains a network of submicron throat geometries. This is the central unproven premise. Static throats require NEC-violating sources; whether such sources exist is addressed in WP1.
2. *Throat stability*—that throats persist over cosmological timescales. The illustrative LTB construction of Sec. III E suggests this is plausible for dust dynamics, but a rigorous treatment with realistic matter sources is part of the numerical-relativity strand of WP1.
3. *Constraint-wave asymmetry on realistic backgrounds*—that the kinematic asymmetry derived in [1] on static, spherically symmetric throats with a phantom-scalar source continues to hold on whatever modified backgrounds the actual support mechanism produces, and on time-dependent FLRW-embedded geometries. The robustness question is addressed in WP1 (background dependence) and WP2 (time-dependent extension).
4. *Cold sequestered sector*—that the sequestered baryons satisfy $T_{\text{seq}} \lesssim 100$ K at late times, which requires asymmetric reheating. Whether the asymmetry can be both established and maintained against inflaton-mediated washout is addressed in WP4.
5. *Sufficient sequestered baryonic mass*—that the total mass of sequestered baryons equals the observed dark matter mass ($\Omega_{\text{seq}} \approx 0.25$). The mechanism that places this mass behind throats is part of the formation question addressed in WP5.

A. Work packages

1. WP1: Throat existence and stability

The decisive question is whether submicron throat geometries can be sustained by physics within or near the Standard Model. We describe three distinct routes for investigation:

- *Semiclassical constructions*

Computing the renormalised stress-energy tensor for quantum fields on throat backgrounds and determining whether Casimir-type negative energy densities can self-consistently sustain the geometry at submicron scales, along the lines of quantum energy inequality analyses [18, 21] but applied to closed-throat rather than open-throat configurations. The decisive technical step is a quantitative application of the smeared and double-smeared null energy conditions (SNEC/DSNEC), as developed in the wormhole context by Kontou [21] for the

Maldacena-Milekhin-Popov geometry, to our scenario’s (N_{throat}, r_0) viable window: such an analysis would convert the heuristic per-throat exotic-energy estimate of Eq. (9) and the cumulative bound of Eq. (10) into a sharp constraint on the field content and UV cutoff required to sustain the throat population, and would either select a region of parameter space compatible with semiclassical gravity or identify the geometric-sequestration scenario as one that requires physics beyond the semiclassical regime.

- *Modified gravity*

Exact traversable wormhole solutions in $f(R)$, scalar-tensor, and Gauss-Bonnet theories that avoid classical NEC violation, with the transmission analysis of Ref. [1] repeated on each such background to verify that the constraint-wave asymmetry persists.

- *Numerical relativity*

Numerical 3+1 evolution of throat initial data in a full general-relativistic code, with realistic matter sources, to determine persistence timescales and identify instability channels.

A separate, logically prior, activity is to clarify the status of topological censorship theorems [28, 29] for sub-topological bottlenecks—simply connected geometries with a minimal two-sphere—which do not obviously fall under the assumptions of the standard theorems. The deliverable for this package is at least one explicit, controlled throat construction with quantified stability, or a rigorous no-go result delimiting the parameter space. The latter would itself be a significant contribution, since it would rule out an entire class of geometric dark matter scenarios.

2. WP2: Higher multipoles, tidal fields, and lensing

The gravitational analysis developed in Sec. II and Sec. VII E establishes the monopole transmission law and the point-mass lensing approximation from matched asymptotics, but several refinements are needed to connect these results rigorously to structure-formation phenomenology. The connection coefficients C_ℓ reported in [1] for the Ellis-Bronnikov background should be computed for the broader parametric and non-ultrastatic throat families, to quantify how the $(r_0/d)^{2\ell+1}$ scaling depends on the throat profile. The extension of the monopole conservation law to fully time-dependent, FLRW-embedded throat networks, beyond the quasistatic regime identified in Sec. II, would close the last gap between the static derivation and the cosmological application. On the observational side, matched-asymptotics derivations of shear, convergence, and flexion predictions for generic throat geometries (extending Eq. (50) beyond EB) would enable direct comparison with weak lensing surveys. Finally, the predicted

absence of sequestered-source higher multipoles gives a distinguishing signature: sequestered matter cannot produce the tidal-field and flexion signals associated with internally clumpy luminous sources, which opens a route to falsification via high-resolution shear and flexion measurements from Euclid, LSST, and future surveys.

3. WP3: Precision cosmological constraints

The CLASS Boltzmann computation presented in Sec. VF establishes consistency of the scenario with Planck at the 5×10^{-5} numerical precision floor of our modified code, and places pivot-scale Lyman- α bounds of $T_{\text{seq}} \lesssim 10^3$ K (direct criterion) to $\lesssim 10^2$ K (WDM-mapping criterion) in the constant- c_s^2 approximation. The adiabatic CLASS runs of Sec. VF exclude the warm-today/hot-in-the-past branch of the adiabatic family, and the worked example of Sec. VG 2 shows that the most natural Scenario-B candidate produces a Scenario-A-like outcome. Whether the diffuse phase is described by Scenario A (passive, mK-scale today) or Scenario B (active, requiring an unspecified low-redshift heating mechanism) determines the precise interpretation of the constant- c_s^2 bounds. Converting these scoping estimates into rigorous constraints requires three further steps:

1. A full Planck likelihood analysis using the modified CLASS module to report $\Delta\chi^2$ and marginalised posteriors for T_{seq} alongside the standard Λ CDM parameters, using Planck TT/TE/EE + lensing + low- ℓ likelihoods and complementary BOSS/eBOSS data.
2. A dedicated Lyman- α flux-power analysis using modern emulators (LaCE [47], Lyman-ABGD [85], ForestFlow, or equivalent), implemented for the constant- c_s^2 thermal history (the only viable limit, Sec. VF). The Jeans-pressure cutoff shape differs characteristically from the free-streaming WDM shape — the Jeans model preserves more power at intermediate k and cuts off more steeply — so a shape-aware emulator analysis is needed to determine the precise bound on T_{seq} . Our pivot-scale estimate almost certainly overstates the constraint, and a shape-aware analysis is expected to relax the bound.
3. Cross-correlations with ISW, CMB lensing, and 21-cm measurements to identify any distinctive signatures of the sequestered-baryon component relative to Λ CDM.

The deliverable is a reproducible set of constraints, with public code and posterior chains, that definitively locates the scenario within the observationally allowed parameter space.

4. WP4: Early-universe thermal history

The washout estimate of Sec. IVD establishes that the centrifugal barrier suppresses cross-sector energy transfer by at least $\sim 10^{-18}$ for submicron coexpanding throats, independent of the reheating temperature (Eq. (25)). Combined with inflaton decoherence for early-forming throats with $r_0 < \lambda_\phi$, this makes re-equilibration of the sectors strongly suppressed; a full nonlinear treatment is required to establish whether equilibration is avoided once the asymmetry $\xi \ll 1$ has been established. The central question for this work package is therefore not whether ξ survives, but what value of ξ each establishment mechanism produces. The three mechanisms identified in Sec. IVD — inflaton decoherence at narrow throats, post-reheating throat formation, and geometric suppression of preheating resonances — each require a concrete embedding in an inflaton model specifying masses, couplings, and branching ratios, to determine whether $\xi \lesssim 0.39$ is a generic outcome or a tuned one.

A second, logically independent, question is the late-time thermal evolution of the sequestered gas. The adiabatic CLASS runs of Sec. VF exclude the warm-today/hot-in-the-past branch of the adiabatic family, and the worked example of Sec. VG 2 shows that the most natural Scenario-B candidate — dark-photon Compton coupling to a primordial sequestered radiation bath — delivers a Scenario-A-like mK-scale diffuse phase, not the warm $10\text{--}10^3$ K phase required to relieve Λ CDM small-scale tensions. The substantive open question is therefore whether any low-redshift energy-injection mechanism — beyond a primordial thermal reservoir — can sustain the diffuse phase at galactic-Jeans-relevant temperatures. Three classes of candidate mechanism merit quantitative investigation: impulsive gravitational scattering off visible-sector substructure, with explicit treatment of the bulk-velocity-to-thermal-velocity conversion; tidal energy transfer through the throat network from visible-sector structure formation; and slowly-decaying sequestered relic species. Each mechanism would feed back into the modified CLASS implementation through a non-trivial $c_s(z)$ history, with detectable signatures in the matter power spectrum at Lyman- α scales. Idealised SPH or N-body simulations of structure formation in a self-gravitating, metal-free gas with H_2 cooling would be needed to map the in-halo phase, in either scenario.

A further subtlety concerns baryogenesis — whether the matter-antimatter asymmetry is necessarily equal on both sides of a throat, and what observational consequences follow from any mismatch. These questions are logically independent of the throat existence question of WP1, and either package could be pursued in parallel.

5. WP5: Formation scenarios and observational predictions

The microlensing analysis of Sec. VII E favours the “many light throats” regime with $N_{\text{throat}} \gtrsim 2 \times 10^{22}$ per

MW-like halo, with $r_0 \lesssim 10^{-7}$ m required by the exotic-energy budget. A quantitative formation model requires specifying the initial conditions (phase transition parameters or primordial fluctuation spectrum), the dynamical mechanism, and the resulting mass function and spatial correlations of the throat population. This model would then feed directly into halo-profile reconstruction (matching sequestered monopoles to observed rotation curves and NFW profiles), extension of the microlensing and millilensing rate predictions of Sec. VII E 3 to a model-derived mass function, plus millilensing rate predictions for higher-mass throats outside the asteroid-mass window, GW opacity predictions as a function of throat size and frequency for LISA and Einstein Telescope, and frequency-dependent EM leakage predictions for throats near the gamma-ray opacity threshold ($r_0 \sim 10^{-13}$ m) comparable to Fermi Large Area Telescope (Fermi-LAT) and Cherenkov Telescope Array (CTA) sensitivity. The NEC-violating energy budget summed over the full throat population is a further output of such a formation model and constitutes a sharp internal consistency check: a formation scenario that requires integrated NEC violation parametrically larger than what WP1 can supply is ruled out on thermodynamic grounds alone.

B. Programme status

The scenario explored in this paper is speculative. What we have shown is that a sharply defined geometric mechanism produces the right observational phenomenology at leading order — what we have not shown, and do not claim to show, is that the spatial geometry of our universe actually realises this mechanism. The above five work packages define the concrete physical calculations and constructions whose outcomes would determine whether geometric sequestration is a viable alternative to particle dark matter or a mathematically instructive dead end. In the first case, the scenario becomes a quantitatively testable alternative to the particle dark matter programme, with distinctive observational signatures in lensing, structure formation, and — for throats near the gamma-ray opacity threshold — electromagnetic leakage. In the second case, the constraint-wave asymmetry remains a mathematical result about field propagation on throat backgrounds, the cosmological analysis remains a quantitative exploration of a particular hypothesis space, and the no-go result itself places bounds on what space-time geometry can and cannot do in the dark matter problem.

IX. Conclusion

We have proposed that dark matter phenomenology could arise from a constraint-wave transmission asymmetry in spacetimes containing narrow geometric

throats. The asymmetry — strong electromagnetic suppression versus smooth gravitational monopole transmission — is derived from the four-dimensional field equations [1], and is a universal feature of static, spherically symmetric throat geometries with $e^{2\alpha} > 0$ at the throat.

The cosmological viability of the proposal requires:

- (i) throat formation before or during reheating, so that the sequestered regions have radiation temperatures below $\sim 39\%$ of the visible sector’s (the N_{eff} constraint);
- (ii) the existence of an NEC-violating source — semiclassical, modified-gravitational, or quantum-gravitational — capable of sustaining the throat geometry, which we do not establish in this paper but for which several candidate mechanisms are identified in Sec. III B and analysed as WP1 in Sec. VIII; and
- (iii) a quantitative demonstration that the scenario reproduces the CMB power spectrum and matter power spectrum.

We have provided a first-principles calculation addressing (iii): the numerical solution of the modified Boltzmann equations shows that the matter transfer function is indistinguishable from CDM for $T_{\text{seq}} \leq 10$ K and marginally consistent with Lyman- α bounds for $T_{\text{seq}} \sim 100\text{--}10^3$ K (Fig. 6), while the CMB power spectrum is identical to Λ CDM at the 5×10^{-5} fractional precision of our modified CLASS code at all measured multipoles (Fig. 3). The binding constraint is the Lyman- α forest, for which a half-mode-to-WDM mapping gives $T_{\text{seq}} \lesssim 100$ K (conservative), while a direct $P(k)$ criterion at the Lyman- α pivot is satisfied up to $T_{\text{seq}} \sim 10^3$ K. These bounds apply in the constant- c_s limit; the adiabatic CLASS runs of Sec. V F exclude the warm-today/hot-in-the-past branch of the adiabatic family. Whether the actual diffuse-phase thermal history is closer to passive Scenario A (mK-scale today, indistinguishable from CDM at every observable scale) or active Scenario B (requiring a low-redshift heating mechanism beyond a primordial radiation bath) is a substantive open question (Sec. V G 2); the H_2 cooling time exceeds the Hubble time by ten orders of magnitude at the cosmic mean density (Sec. V I F), so cooling does not deplete any heating that does operate.

The sequestered baryons behave like CDM on large scales (pressureless, non-relativistic, gravitationally clustering). On small scales, the minimal scenario — in which the diffuse phase is heated only by a primordial sequestered radiation bath — predicts a matter power spectrum indistinguishable from Λ CDM at all observational scales: $P(k)/P_{\Lambda\text{CDM}} > 0.999$ for $k \leq 10 h/\text{Mpc}$, with the eventual high- k turnover from the brief Compton-coupled epoch occurring at $k \gtrsim 30 h/\text{Mpc}$, far beyond any current observation (Sec. V G 2, Fig. 4). A galactic-scale Jeans cutoff capable of relieving the cusp-core, missing-satellites, and too-big-to-fail tensions would require an

additional diffuse-phase heating mechanism beyond the minimal scenario, and is left to future work (Sec. VIII).

This paper has explored a speculative, but quantitatively defined, proposal—that dark matter could be ordinary baryonic matter geometrically sequestered behind narrow throats in the spatial geometry. Static throats require NEC-violating sources (Sec. III), which may arise from semiclassical effects, modified gravity, or quantum gravity—we do not resolve this existence question here. What we have demonstrated is that given such geometries, the cosmological consequences are quantitatively consistent with all current observational constraints, and that the same centrifugal barrier that renders the sequestered matter dark also prevents thermal re-equilibration of the sectors after reheating (Sec. IV D), with a suppression factor that is independent of the reheating temperature for coexpanding throats. The self-consistent CLASS Boltzmann computation (Sec. V F) shows CMB deviations below the 5×10^{-5} numerical precision floor, the Lyman- α comparison shows compatibility at the pivot-scale level across the full $T_{\text{seq}} \leq 10^3$ K range studied here, and the lensing analysis confirms point-mass behaviour for astrophysical impact parameters. The scenario constitutes a well-defined research programme whose viability can be established or ruled out by addressing the geometric existence question.

If viable, the geometric-sequestration mechanism would resolve the dark matter problem without invoking new particle species, would, in extensions of the minimal scenario, produce a Jeans-scale cutoff in the matter power spectrum at galactic scales, and would predict that sequestered matter couples to our sector exclusively through the static gravitational potential—invisible to all propagating fields (EM, GW, and fermionic, including the $j = \frac{1}{2}$ s-wave; Sec. VII B 2) for submicron throats. For larger throats, frequency-dependent electromagnetic leakage would provide a distinctive observational signature.

The central open question is the existence and stability of the required throat geometries. Promising directions include: numerical evolution of throat initial data in full 3+1 general relativity; exact wormhole solutions in modified gravity theories ($f(R)$, scalar-tensor, Gauss-Bonnet) that avoid classical NEC violation; and semiclassical constructions exploiting Casimir-type negative energy densities at submicron scales. If any of these programmes produces a stable, submicron throat geometry, the cosmological phenomenology presented here applies immediately, and the scenario becomes a quantitatively testable alternative to particle dark matter.

Acknowledgments

The cosmological computations in this work were performed using the Boltzmann code CLASS [37]; the author thanks Julien Lesgourgues, Thomas Tram, and Nils Schoeneberg for making it publicly available. The

modified CLASS module implementing the sequestered-baryon fluid, together with the `.ini` files and analysis script used to generate Fig. 3, is publicly available at https://github.com/jeffriley/class_sequestered_baryons, tagged as *seq-baryon-paper-v1*. The repository is an independent clone of CLASS [37], not a fork of the upstream project, and is maintained solely as a reproducibility artefact for this paper. Refer to the README file in the *sequestered_baryons* directory in the repository for a description of the modifications.

A. Derivation of the effective coupling kernel

We derive the effective coupling kernel $W(k, r_0)$ that governs how cosmological density perturbations on the sequestered side source gravitational potentials on our side through a throat.

A plane-wave density perturbation $\delta(\mathbf{x}) = \delta_0 e^{i\mathbf{k}\cdot\mathbf{x}}$ can be expanded in spherical harmonics centred on a throat at the origin using the Rayleigh formula:

$$e^{i\mathbf{k}\cdot\mathbf{x}} = \sum_{\ell=0}^{\infty} (2\ell+1) i^\ell j_\ell(kr) P_\ell(\cos\theta). \quad (\text{A1})$$

The fraction of the perturbation in multipole ℓ at the throat radius $r = r_0$ is

$$f_\ell(kr_0) = \frac{(2\ell+1) |j_\ell(kr_0)|^2}{\sum_{\ell'} (2\ell'+1) |j_{\ell'}(kr_0)|^2}. \quad (\text{A2})$$

In the regime $kr_0 \ll 1$ (perturbation wavelength much larger than throat radius), the spherical Bessel functions have the asymptotic behaviour $j_0(x) \approx 1 - x^2/6$ and $j_\ell(x) \approx x^\ell/(2\ell+1)!!$ for $\ell \geq 1$, giving

$$f_0 = 1 - \mathcal{O}((kr_0)^2), \quad f_1 = \frac{(kr_0)^2}{3} + \mathcal{O}((kr_0)^4). \quad (\text{A3})$$

The perturbation is almost entirely in the monopole.

Each multipole ℓ transmits through the throat with a factor \mathcal{T}_ℓ , where $\mathcal{T}_0 = 1$ (the monopole conservation law, Eq. (4)) and $\mathcal{T}_{\ell \geq 1} \ll 1$ (sub-barrier suppression, Eq. (3)). The effective coupling kernel is

$$W(k, r_0) = \sum_{\ell=0}^{\infty} f_\ell(kr_0) \mathcal{T}_\ell = f_0 + \sum_{\ell \geq 1} f_\ell \mathcal{T}_\ell. \quad (\text{A4})$$

Since $\mathcal{T}_\ell \leq 1$ for all ℓ , we have the rigorous bounds

$$f_0 \leq W(k, r_0) \leq 1, \quad (\text{A5})$$

and therefore

$$W(k, r_0) = 1 - \mathcal{O}((kr_0)^2). \quad (\text{A6})$$

This result is independent of the specific transmission factors $\mathcal{T}_{\ell \geq 1}$ —it depends only on the smallness of kr_0 , which ensures that the perturbation is dominated by the monopole component at the throat.

For a statistical ensemble of N throats distributed through a cosmological volume, each throat independently transmits the monopole component of the ambient perturbation. The ensemble-averaged coupling is $\langle W \rangle = W(k, r_0)$ (the same for each throat), so the result (A6) holds for the full network. This establishes that sequestered baryons couple gravitationally to our sector with the same strength as CDM, at all observable cosmological scales, for any throat radius $r_0 \lesssim 1$ m.

B. Transmission coefficients on the Ellis-Bronnikov throat

For self-containment we summarise here the quantitative transmission results derived in the companion paper [1], with cross-references to the corresponding equations and tables of that work. The cosmological analysis of the present paper depends on these results in three places: the reheating-washout estimate of Sec. IV D, the high-multipole gravitational suppression invoked in Sec. VII E, and the qualitative claim that sequestered matter is invisible to all propagating radiation (Sec. VII). The numbers below are reproduced from [1] for reader convenience; full derivations, convergence analyses, and figures are deferred to that paper.

1. Effective potentials and barrier-top frequencies

On the EB throat $ds^2 = -dt^2 + d\sigma^2 + (\sigma^2 + r_0^2)d\Omega^2$, decomposition of the four-dimensional Maxwell equations into vector spherical harmonics yields a Schrödinger-form equation for each multipole $\ell \geq 1$ with effective potential

$$V_\ell^{(\text{EM})}(\sigma) = \frac{\ell(\ell+1)}{\sigma^2 + r_0^2}, \quad (\text{B1})$$

peaked at the throat with maximum $V_\ell^{(\text{EM})}(0) = \ell(\ell+1)/r_0^2$. The corresponding barrier-top frequency is $\omega_{\text{max}}^{(\ell)} = \sqrt{\ell(\ell+1)}/r_0$. For the lowest physical EM mode $\ell = 1$, $\omega_{\text{max}}^{(\ell=1)} = \sqrt{2}/r_0 \approx 1.414/r_0$.

For axial gravitational waves, the corresponding effective potential is

$$V_\ell^{(\text{GW})}(\sigma) = \frac{\ell(\ell+1)}{\sigma^2 + r_0^2} - \frac{3r_0^2}{(\sigma^2 + r_0^2)^2}, \quad (\text{B2})$$

where the second term is a curvature correction that lowers the barrier height to $V_\ell^{(\text{GW})}(0) = [\ell(\ell+1) - 3]/r_0^2$. The lowest GW mode is $\ell = 2$, with barrier height $3/r_0^2$ and barrier-top frequency $\omega_{\text{max}}^{(\text{GW}, \ell=2)} = \sqrt{3}/r_0 \approx 1.732/r_0$ —half the EM $\ell = 2$ barrier height of $6/r_0^2$ due to the curvature correction.

The static gravitational monopole $\ell = 0$ satisfies, in the source-free region, the conservation law

$$\frac{1}{a^2} \frac{d}{d\sigma} \left(a^2 e^{2\alpha} \frac{d\Phi}{d\sigma} \right) = 0, \quad (\text{B3})$$

which admits the exact EB solution $\Phi(\sigma) = (C/r_0) \arctan(\sigma/r_0) + \Phi_0$, with the conserved flux $\mathcal{F} = a^2 e^{2\alpha} \Phi' = C$ identified with the enclosed mass via Gauss's law. There is no centrifugal barrier and no sub-barrier suppression provided $e^{2\alpha}(0) > 0$, the condition that excludes the $\lambda \rightarrow 0$ limit of the Damour-Solodukhin construction.

2. Sub-barrier transmission: numerical values

The power transmission coefficients $T(\omega)$ for EM $\ell = 1$, EM $\ell = 2$, and GW $\ell = 2$ on the EB throat are obtained by Numerov integration of the corresponding Schrödinger-form equations, with plane-wave boundary conditions at $\sigma = \pm L_{\text{max}} = \pm 30 r_0$, on a uniform grid with $N = 10^4$ points [1]. Convergence is verified by doubling N and L_{max} , and unitarity $|R|^2 + T = 1$ is satisfied to $< 10^{-4}$ at all tabulated frequencies.

Table VII reproduces the headline transmission coefficients from [1] at benchmark sub-barrier frequencies. The entries confirm that all propagating modes with $\ell \geq 1$ are strongly suppressed below their respective barrier tops.

TABLE VII. Power transmission coefficients on the Ellis-Bronnikov throat ($r_0 = 1$), reproduced from [1] (Tables I and III of that work). The Dirac $j = \frac{1}{2}$ s-wave entries are computed in Sec. VII B 2 of the present paper. Entries marked $< 10^{-15}$ are below double-precision Numerov reliability and should be interpreted as upper bounds.

ωr_0	EM $\ell = 1$	EM $\ell = 2$	GW $\ell = 2$	Dirac $j = \frac{1}{2}$
0.05	—	—	—	1.6×10^{-6}
0.1	2.2×10^{-8}	$< 10^{-15}$	$< 10^{-15}$	2.6×10^{-5}
0.3	1.8×10^{-5}	$< 10^{-12}$	—	2.7×10^{-3}
0.5	5.4×10^{-4}	4.9×10^{-9}	3.7×10^{-5}	2.8×10^{-2}
1.0	9.5×10^{-2}	1.1×10^{-4}	3.5×10^{-2}	5.9×10^{-1}
1.5	—	1.4×10^{-3}	—	9.7×10^{-1}
2.0	0.99	0.24	0.93	—

In the deep sub-barrier regime, the transmission follows a power-law scaling $T \sim K(\omega r_0)^\nu$ with $\nu \approx 6.0 \pm 0.3$ for EM $\ell = 1$ [1], agreeing with the analytic core-plus-tail WKB estimate $\nu = 4\sqrt{\ell(\ell+1)} \approx 5.66$. The strict asymptotic value $\nu = 2\ell + 1 = 3$ predicted by Bessel-function matching for the centrifugal $1/\sigma^2$ tail is recovered only in the limit $\omega r_0 \rightarrow 0$ below the numerical range [1]. The Dirac $j = \frac{1}{2}$ s-wave (Sec. VII B 2) gives a power-law fit $\nu \approx 4.2$ in the same numerical range—less strongly suppressed than EM $\ell = 1$ but qualitatively the same power-law form. All channels enter the strongly-suppressed regime $T \leq 10^{-5}$ for $\omega r_0 \lesssim 0.1$, which is the controlling input to the reheating-washout estimate of Sec. IV D.

3. Higher static multipoles

For $\ell \geq 1$ static gravitational multipoles, the source-free Poisson equation on the spatial metric $d\sigma^2 + a^2 d\Omega^2$ reads

$$\frac{1}{a^2} \frac{d}{d\sigma} \left(a^2 \frac{d\Phi_\ell}{d\sigma} \right) - \frac{\ell(\ell+1)}{a^2} \Phi_\ell = 0, \quad (\text{B4})$$

which admits a centrifugal barrier near the throat and exhibits decaying/growing solutions $\sigma^{-(\ell+1)}/\sigma^\ell$ at large $|\sigma|$. A source on one side at $\sigma = d \gg r_0$ produces a far-side amplitude

$$\Phi_\ell(\text{far})/\Phi_\ell(\text{source}) \sim C_\ell (r_0/d)^{2\ell+1}, \quad (\text{B5})$$

where the connection coefficients C_ℓ are determined by integrating Eq. (B4) outward from the throat with even initial conditions. On the EB profile the numerical values are $C_1 = \pi/2 \approx 1.5708$, $C_2 = 3.0000$, and $C_3 = 5.8905$, with fitted growth exponents matching ℓ to better than 0.01% [1]. For $\ell = 1$ (gravitational dipole) at $d = 10 r_0$, the suppression factor is $\sim 10^{-3}$; for $\ell = 2$ (quadrupole), $\sim 10^{-5}$.

This polynomial suppression is the basis of the point-mass lensing approximation used in Sec. VII E and the effective-monopole-only coupling to internal source structure used throughout. The qualitative scaling $(r_0/d)^{2\ell+1}$ is universal across throat profiles; the prefactor C_ℓ varies by at most a factor of a few across the parametric family of throat shapes and the reflected-Schwarzschild background [1].

4. Universality across throat backgrounds

The constraint-wave asymmetry derived above is not specific to the EB profile. The qualitative pattern—strong sub-barrier suppression for all $\ell \geq 1$ propagating and static modes, polynomial transmission for the $\ell = 0$ monopole—holds on any static, spherically symmetric throat geometry with a minimal-area two-sphere and $e^{2\alpha}(0) > 0$ [1]. The companion paper verifies this explicitly across a one-parameter family of throat profiles $a(\sigma) = (\sigma^{2n} + r_0^{2n})^{1/(2n)}$, $n \in [1, \infty)$, and on the reflected-Schwarzschild Damour-Solodukhin wormhole with $\lambda > 0$. The $\lambda \rightarrow 0$ limit of the latter, in which $e^{2\alpha} \rightarrow 0$ at the throat, is excluded as discussed in Sec. II—the conserved flux then gives a non-integrable monopole potential, and the limit recovers the Schwarzschild horizon [1].

The numerical prefactors K (in $T \sim K(\omega r_0)^\nu$) and C_ℓ depend on the specific profile, but the kinematic structure underlying the cosmological application of the present paper—the existence of a centrifugal barrier with barrier-top frequency $\omega_{\text{max}} \sim 1/r_0$, the resulting strong suppression at all astrophysical and cosmological frequencies for submicron r_0 , and the unsuppressed transmission of the static gravitational monopole—is universal.

-
- [1] J. Riley, Electromagnetic, gravitational wave, and static gravitational transmission through throat spacetimes: a constraint-wave asymmetry (2026), arXiv:2604.14238 [gr-qc].
- [2] T. M. Undagoitia and L. Rauch, Dark matter direct-detection experiments, *Journal of Physics G: Nuclear and Particle Physics* **43**, 013001 (2015).
- [3] G. Arcadi, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, M. Pierre, S. Profumo, and F. S. Queiroz, The waning of the WIMP? a review of models, searches, and constraints, *The European Physical Journal C* **78**, 203 (2018).
- [4] V. C. Rubin and W. K. Ford, Jr., Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions, *The Astrophysical Journal* **159**, 379 (1970).
- [5] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard, Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc)., *The Astrophysical Journal* **238**, 471 (1980).
- [6] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, A direct empirical proof of the existence of dark matter, *Astrophys. J. Lett.* **648**, L109 (2006).
- [7] Planck Collaboration, N. Aghanim, *et al.*, Planck 2018 results. VI. cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [8] V. Springel, C. S. Frenk, and S. D. M. White, The large-scale structure of the Universe, *Nature* **440**, 1137 (2006).
- [9] A. Kirillov and E. Savelova, Dark matter from a gas of wormholes, *Physics Letters B* **660**, 93 (2008).
- [10] A. A. Kirillov and E. P. Savelova, Density perturbations in a gas of wormholes, *Monthly Notices of the Royal Astronomical Society* **412**, 1710 (2011).
- [11] R. Foot, Mirror Matter-Type Dark Matter, *Int. J. Mod. Phys. D* **13**, 2161 (2004), arXiv:astro-ph/0407623 [astro-ph].
- [12] A. R. Frey, A. Mazumdar, and R. C. Myers, Stringy effects during inflation and reheating, *Phys. Rev. D* **73**, 026003 (2006).
- [13] X. Chen and S.-H. Henry Tye, Heating in brane inflation and hidden dark matter, *Journal of Cosmology and Astroparticle Physics* **2006** (06), 011.
- [14] D.-C. Dai and D. Stojkovic, Observing a wormhole, *Phys. Rev. D* **100**, 083513 (2019).
- [15] M. S. Morris and K. S. Thorne, Wormholes in spacetime and their use for interstellar travel: A tool for teaching general relativity, *American Journal of Physics* **56**, 395 (1988).
- [16] D. Hochberg and M. Visser, Geometric structure of the generic static traversable wormhole throat, *Phys. Rev. D* **56**, 4745 (1997).

- [17] M. Visser, *Lorentzian Wormholes: From Einstein to Hawking* (AIP Press, New York, 1995).
- [18] C. J. Fewster, Lectures on quantum energy inequalities (2012), arXiv:1208.5399 [gr-qc].
- [19] L. H. Ford and T. A. Roman, Averaged energy conditions and quantum inequalities, *Phys. Rev. D* **51**, 4277 (1995).
- [20] L. H. Ford and T. A. Roman, The quantum interest conjecture, *Phys. Rev. D* **60**, 104018 (1999).
- [21] E.-A. Kontou, Wormhole restrictions from quantum energy inequalities, *Universe* **10**, 291 (2024), arXiv:2405.05963 [gr-qc].
- [22] J. Maldacena, A. Milekhin, and F. Popov, Traversable wormholes in four dimensions, *Classical and Quantum Gravity* **40**, 155016 (2023), arXiv:1807.04726 [hep-th].
- [23] P. Kanti, B. Kleihaus, and J. Kunz, Wormholes in dilatonic einstein-gauss-bonnet theory, *Phys. Rev. Lett.* **107**, 271101 (2011).
- [24] G. Antoniou, A. Bakopoulos, P. Kanti, B. Kleihaus, and J. Kunz, Novel einstein-scalar-gauss-bonnet wormholes without exotic matter, *Phys. Rev. D* **101**, 024033 (2020).
- [25] A. B. Balakin, J. P. S. Lemos, and A. E. Zayats, Nonminimal coupling for the gravitational and electromagnetic fields: Traversable electric wormholes, *Phys. Rev. D* **81**, 084015 (2010).
- [26] F. S. N. Lobo and M. A. Oliveira, Wormhole geometries in $f(R)$ modified theories of gravity, *Phys. Rev. D* **80**, 104012 (2009).
- [27] A. G. Agnese and M. La Camera, Wormholes in the brans-dicke theory of gravitation, *Phys. Rev. D* **51**, 2011 (1995).
- [28] J. L. Friedman, K. Schleich, and D. M. Witt, Topological censorship, *Phys. Rev. Lett.* **71**, 1486 (1993).
- [29] G. J. Galloway, K. Schleich, D. M. Witt, and E. Woolgar, Topological censorship and higher genus black holes, *Phys. Rev. D* **60**, 104039 (1999).
- [30] R. H. Cyburt, B. D. Fields, K. A. Olive, and T.-H. Yeh, Big bang nucleosynthesis: Present status, *Rev. Mod. Phys.* **88**, 015004 (2016).
- [31] R. J. Cooke, M. Pettini, and C. C. Steidel, One percent determination of the primordial deuterium abundance*, *The Astrophysical Journal* **855**, 102 (2018).
- [32] T. Buchert, M. Carfora, G. F. R. Ellis, E. W. Kolb, M. A. H. MacCallum, J. J. Ostrowski, S. Räsänen, B. F. Roukema, L. Andersson, A. A. Coley, and D. L. Wiltshire, Is there proof that backreaction of inhomogeneities is irrelevant in cosmology?, *Classical and Quantum Gravity* **32**, 215021 (2015).
- [33] S. Ganguly, T. H. Jung, and S. Yun, Consistent N_{eff} fitting in big bang nucleosynthesis analysis (2025), arXiv:2507.23354 [hep-ph].
- [34] L. Kofman, A. Linde, and A. A. Starobinsky, Towards the theory of reheating after inflation, *Phys. Rev. D* **56**, 3258 (1997).
- [35] P. Adshead, Y. Cui, and J. Shelton, Chilly dark sectors and asymmetric reheating, *Journal of High Energy Physics* **2016**, 016 (2016), arXiv:1604.02458 [hep-ph].
- [36] S. Dodelson and F. Schmidt, *Modern Cosmology*, 2nd ed. (Academic Press, 2020).
- [37] D. Blas, J. Lesgourgues, and T. Tram, The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes, *Journal of Cosmology and Astroparticle Physics* **2011** (07), 034.
- [38] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, S. Cristiani, G. D. Becker, V. D'Odorico, G. Cupani, T.-S. Kim, T. A. M. Berg, S. López, S. Ellison, L. Christensen, K. D. Denney, and G. Worseck, New constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- α forest data, *Phys. Rev. D* **96**, 023522 (2017).
- [39] D. Wands, K. A. Malik, D. H. Lyth, and A. R. Liddle, New approach to the evolution of cosmological perturbations on large scales, *Phys. Rev. D* **62**, 043527 (2000).
- [40] S. Weinberg, *Cosmology* (Oxford University Press, 2008) see §5.4 for the conservation of \mathcal{R} on super-horizon scales.
- [41] J. M. Bardeen, J. R. Bond, N. Kaiser, and A. S. Szalay, The Statistics of Peaks of Gaussian Random Fields, *The Astrophysical Journal* **304**, 15 (1986).
- [42] K. Abazajian *et al.*, Cmb-s4 science case, reference design, and project plan (2019), arXiv:1907.04473 [astro-ph.IM].
- [43] A. Suzuki *et al.*, The LiteBIRD Satellite Mission: Sub-Kelvin Instrument, *Journal of Low Temperature Physics* **193**, 1048 (2018), arXiv:1801.06987 [astro-ph.IM].
- [44] P. Meszaros, The behaviour of point masses in an expanding cosmological substratum, *Astronomy and Astrophysics* **37**, 225 (1974).
- [45] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data, *Phys. Rev. D* **88**, 043502 (2013).
- [46] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto, Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman- α forest, *Phys. Rev. D* **71**, 063534 (2005).
- [47] C. Pedersen, A. Font-Ribera, and N. Y. Gnedin, Compressing the Cosmological Information in One-dimensional Correlations of the Lyman- α Forest, *The Astrophysical Journal* **944**, 223 (2023), arXiv:2209.09895 [astro-ph.CO].
- [48] N. Palanque-Delabrouille, C. Yèche, N. Schöneberg, J. Lesgourgues, M. Walther, S. Chabanier, and E. Armengaud, Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman- α and Planck full-survey data, *Journal of Cosmology and Astroparticle Physics* **2020** (04), 038.
- [49] J. S. Bullock and M. Boylan-Kolchin, Small-scale challenges to the Λ CDM paradigm, *Annual Review Astronomy and Astrophysics* **55**, 343 (2017).
- [50] P. Bode, J. P. Ostriker, and N. Turok, Halo formation in warm dark matter models, *The Astrophysical Journal* **556**, 93 (2001).
- [51] D. N. Spergel and P. J. Steinhardt, Observational Evidence for Self-Interacting Cold Dark Matter, *Phys. Rev. Lett.* **84**, 3760 (2000).
- [52] S. Tulin and H.-B. Yu, Dark matter self-interactions and small scale structure, *Physics Reports* **730**, 1 (2018).
- [53] J. Fan, A. Katz, L. Randall, and M. Reece, Double-disk dark matter, *Physics of the Dark Universe* **2**, 139 (2013).
- [54] R. Barkana and A. Loeb, In the beginning: the first sources of light and the reionization of the universe, *Physics Reports* **349**, 125 (2001), arXiv:astro-ph/0010468 [astro-ph].
- [55] D. Galli and F. Palla, The chemistry of the early Universe, *Astronomy and Astrophysics* **335**, 403 (1998), arXiv:astro-ph/9803315 [astro-ph].
- [56] LIGO Scientific Collaboration, Advanced LIGO, *Classical and Quantum Gravity* **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].
- [57] F. Acernese *et al.*, Advanced Virgo: a second-generation interferometric gravitational wave detector, *Classical and Quantum Gravity* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- [58] T. Akutsu *et al.*, Overview of kagra: Detector design and

- construction history, *Progress of Theoretical and Experimental Physics* **2021**, 05A101 (2021).
- [59] P. Amaro-Seoane *et al.*, Laser Interferometer Space Antenna (2017), submitted to ESA in response to the call for missions for the L3 slot in the Cosmic Vision Programme, arXiv:1702.00786 [astro-ph.IM].
- [60] F. Cooper, A. Khare, and U. Sukhatme, Supersymmetry and quantum mechanics, *Physics Reports* **251**, 267 (1995).
- [61] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. (Wiley, New York, 1998).
- [62] S. R. Furlanetto, S. Peng Oh, and F. H. Briggs, Cosmology at low frequencies: The 21cm transition and the high-redshift universe, *Physics Reports* **433**, 181 (2006).
- [63] D. R. DeBoer *et al.*, Hydrogen epoch of reionization array (hera), *Publications of the Astronomical Society of the Pacific* **129**, 045001 (2017).
- [64] T. L. Bourke, R. Braun, R. Fender, F. Govoni, J. Green, E. Keane, and J. Wagg, Advancing Astrophysics with the Square Kilometre Array, *PoS AASKA14*, 174 (2015), arXiv:1506.04473 [astro-ph.IM].
- [65] J. D. Bowman, A. E. E. Rogers, R. A. Monsalve, T. J. Mozdzen, and N. Mahesh, An absorption profile centred at 78 megahertz in the sky-averaged spectrum, *Nature* **555**, 67 (2018).
- [66] P. Tisserand *et al.* (EROS-2 Collaboration), Limits on the Macho content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds, *Astronomy and Astrophysics* **469**, 387 (2007), arXiv:astro-ph/0607207 [astro-ph].
- [67] A. Udalski, M. K. Szymański, and G. Szymański, OGLE-IV: Fourth Phase of the Optical Gravitational Lensing Experiment, *Acta Astronomica* **65**, 1 (2015), arXiv:1504.05966 [astro-ph.SR].
- [68] H. Niikura, M. Takada, N. Yasuda, *et al.*, Microlensing constraints on primordial black holes with Subaru/HSC Andromeda observations, *Nat. Astron.* **3**, 524 (2019).
- [69] B. Carr, K. Kohri, Y. Sendouda, and J. Yokoyama, Constraints on primordial black holes, *Reports on Progress in Physics* **84**, 116902 (2021).
- [70] K. Griest, Galactic Microlensing as a Method of Detecting Massive Compact Halo Objects, *The Astrophysical Journal* **366**, 412 (1991).
- [71] C. Alcock *et al.*, The MACHO Project: Microlensing Results from 5.7 Years of Large Magellanic Cloud Observations, *The Astrophysical Journal* **542**, 281 (2000).
- [72] T. T. Nakamura and S. Deguchi, Wave optics in gravitational lensing, *Progress of Theoretical Physics Supplement* **133**, 137 (1999), <https://academic.oup.com/ptps/article-pdf/doi/10.1143/PTPS.133.137/5283012/133-137.pdf>.
- [73] N. Banik, J. Bovy, G. Bertone, D. Erkal, and T. J. L. de Boer, Evidence of a population of dark subhaloes from Gaia and Pan-STARRS observations of the GD-1 stream, *MNRAS* **502**, 2364 (2021), arXiv:1911.02662 [astro-ph.GA].
- [74] A. Bonaca, D. W. Hogg, A. M. Price-Whelan, and C. Conroy, The spur and the gap in gd-1: Dynamical evidence for a dark substructure in the milky way halo, *ApJ* **880**, 38 (2019).
- [75] C. G. Lacey and J. P. Ostriker, Massive black holes in galactic halos ?, *ApJ* **299**, 633 (1985).
- [76] G. Agazie *et al.*, The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background, *ApJL* **951**, L8 (2023), arXiv:2306.16213 [astro-ph.HE].
- [77] A. Barnacka, J.-F. Glicenstein, and R. Moderski, New constraints on primordial black holes abundance from femtolensing of gamma-ray bursts, *Phys. Rev. D* **86**, 043001 (2012).
- [78] T. D. Brandt, Constraints on MACHO dark matter from compact stellar systems in ultra-faint dwarf galaxies, *The Astrophysical Journal Letters* **824**, L31 (2016).
- [79] Mellier, Y. *et al.* (Euclid Collaboration), Euclid - I. Overview of the Euclid mission, *A&A* **697**, A1 (2025).
- [80] Paul A. Abell *et al.* (LSST Science Collaborations), Lsst science book, version 2.0, arXiv preprint arXiv:0912.0201 (2009), arXiv:0912.0201 [astro-ph.IM].
- [81] Paul A. Abell *et al.* (The LSST Dark Energy Science Collaboration), The lsst dark energy science collaboration (desc) science requirements document, arXiv preprint arXiv:1809.01669v2 (2021), arXiv:1809.01669v2 [astro-ph.CO].
- [82] E. Aprile *et al.* (XENON Collaboration), First dark matter search with nuclear recoils from the xenonn experiment, *Phys. Rev. Lett.* **131**, 041003 (2023).
- [83] J. Aalbers *et al.* (LUX-ZEPLIN Collaboration), First dark matter search results from the lux-zeplin (lz) experiment, *Phys. Rev. Lett.* **131**, 041002 (2023).
- [84] Z. Bo *et al.* (PandaX Collaboration), Dark Matter Search Results from 1.54 Tonne · Year Exposure of PandaX-4T, *Phys. Rev. Lett.* **134**, 011805 (2025).
- [85] D. C. Hooper, N. Schöneberg, R. Murgia, M. Archidiacono, J. Lesgourgues, and M. Viel, One likelihood to bind them all: Lyman- α constraints on non-standard dark matter, *JCAP* **2022** (10), 032, arXiv:2206.08188.