# Primordial black hole dark matter: A quantitative parameter sensitivity comparison across formation mechanisms and particle candidates

Stefano Profumo*

*Department of Physics and Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, California 95064, USA*

(Received 13 June 2026; accepted 6 August 2026; published 10 September 2026)

Primordial black holes (PBHs) in the asteroid-mass window ($10^{17}$–$10^{22}$ g) can account for all of the dark matter without violating any observational constraint, yet are routinely dismissed as fine-tuned. I put that dismissal to the test by applying three complementary sensitivity measures uniformly across a broad landscape: three noninflationary PBH production mechanisms, six classes of inflationary PBH models, and seven particle dark matter benchmarks, all evaluated against the same observable target. Three distinct naturalness universality classes emerge, determined entirely by the analytic structure of the abundance map rather than by the nature of the dark matter candidate. Biased-domain-wall PBHs, in their least model-dependent (free-$V_b$) form, have the same low sensitivity, $\Delta = 4.5$, as off-resonance weakly interacting massive particles and freeze-in particles ($\Delta = 2$), a sensitivity that, because it is constant over the entire parameter space of the construction, also coincides trivially with its own Wilson-normalized average within that parameter space (Section Definition and conventions), an equivalence that concerns only the space over which $\Delta$ is computed and is not a naturalness statement about the construction as a whole; a further reduction to $\Delta = 2$ is possible only under the additional, independently motivated but not required, assumption that the domain-wall bias is generated by Planck-suppressed operators; early matter-domination PBHs occupy an intermediate tier alongside coannihilating weakly interacting massive particles (WIMPs), unified by a structural identity in which the sensitivity measure equals the logarithm of the ratio of the formation scale to the matter–radiation equality scale; first-order phase transition PBHs, once the more accurate super-exponential collapse probability is used in place of the single-exponential approximation, instead belong to the same highly sensitive tier as resonant WIMP annihilation and single-field inflationary collapse, for a structurally distinct reason; single-field ultraslow-roll inflationary collapse is severely tuned for a distinct reason: a double exponential in which the power spectrum amplitude is itself exponentially sensitive to the inflaton potential coefficients, on top of the exponential collapse sensitivity of the abundance map. My main conclusion is that the claim that PBH dark matter is generically fine-tuned conflates the worst case with a landscape spanning every naturalness tier. The Barbieri-Giudice sensitivity computed here and the Wilson-normalized measure of Iovino and Riotto answer distinct and mutually consistent questions about the same construction, a distinction I clarify within the two-layer decomposition.

DOI: 10.1103/nk1q-5k51

## I. INTRODUCTION

The identity of dark matter remains one of the central open questions of physics. Among the candidates that have attracted sustained interest are weakly interacting massive particles (WIMPs) [1,2], axions [3–5], feebly interacting ("frozen-in") particles [6,7], and asymmetric dark matter [8,9]. More recently, primordial black holes (PBHs) have reemerged as a viable dark matter candidate [10,11], particularly in the asteroid-mass window ($10^{17} \lesssim M_{\rm PBH}/{\rm g} \lesssim 10^{22}$) where Hawking evaporation is negligible, microlensing and disruption constraints are weak, and the PBH fraction $f_{\rm PBH} \equiv \Omega_{\rm PBH}/\Omega_{\rm DM}$ can equal unity [12–14].

A common objection to PBH dark matter is that it requires fine-tuning: in inflationary PBH production the probability of collapse is exponentially sensitive to the primordial curvature power spectrum, which must be arranged to produce the right number density in the right mass window. This argument is frequently made qualitatively, but rarely subjected to the same quantitative scrutiny applied to WIMP dark matter, where a similar exponential sensitivity is implicitly encoded, for instance, in the fabled

*Contact author: profumo@ucsc.edu

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

thermal freeze-out mechanism. The WIMP miracle, the coincidence that a particle with electroweak-scale mass and coupling naturally freezes out with $\Omega h^2 \approx 0.12$, is itself a statement about naturalness, with a well-defined quantitative expression that deserves comparison.

The fine-tuning of dark matter relic abundances has attracted quantitative attention most prominently in the supersymmetric WIMP context. Grothaus, Lindner, and Takanishi [15] applied the Barbieri-Giudice (BG) measure [16], which quantifies how sensitively a physical observable responds to fractional variations of the underlying input parameters, $\Delta \equiv \max_i |\partial \ln \Omega / \partial \ln x_i|$, to neutralino dark matter in the minimal supersymmetric Standard Model (MSSM), showing that direct-detection exclusions preferentially eliminate untuned parameter regions. The most comprehensive MSSM treatment was given by Cabrera *et al.* [17], who studied all major neutralino annihilation mechanisms under a unified BG framework; their finding that funnel and coannihilation scenarios carry the highest combined tuning is fully consistent with the WIMP benchmarks of the present paper. King and Roberts [18] coined the term "supernatural dark matter" for MSSM regions where the relic abundance is achieved with no fine-tuning at all.

For PBH dark matter, the fine-tuning discussion has been framed almost entirely around the inflationary-collapse mechanism. Several studies have quantified the sensitivity of the PBH abundance to the inflaton potential shape in the ultraslow-roll regime [19,20]. Kalaja *et al.* [21] showed that the curvature power-spectrum variance $\sigma$ must lie within a $\sim$7% fractional window to produce all of the dark matter as PBHs under Gaussian statistics, a result broadly consistent with the quantitative fine-tuning estimates of Refs. [19,20]. Iovino and Riotto [22], by contrast, apply Wilson's naturalness criterion to single-field inflationary PBH models and conclude that they are not technically unnatural when the measure is formulated symmetrically around the target abundance. This is not, properly understood, in tension with the sensitivity estimates above: a bare logarithmic sensitivity and a Wilson-normalized naturalness verdict answer distinct questions (Sec. II A), and the present paper reports the former, organized into the two-layer decomposition of Sec. III D.

Three important gaps remain in the existing literature that this paper fills. First, no prior work has compared PBH and particle dark matter fine-tuning under a common measure applied to a common observable target ($\Omega_{\rm DM} h^2 = 0.120$); the two communities have developed parallel but disconnected naturalness vocabularies. Second, the noninflationary PBH formation mechanisms, including biased domain walls, first-order phase transitions, and early matter domination, have never been analyzed from a quantitative fine-tuning perspective. Third, no systematic comparison of fine-tuning across inflationary PBH model classes, from curvaton and spectator-field models through hybrid inflation to single-field ultraslow-roll, has been presented within a unified, internally consistent sensitivity measure framework. The present paper addresses all three gaps.

In what follows I apply the BG measure to the parameter spaces of three noninflationary PBH formation mechanisms, six classes of inflationary PBH model, and seven particle dark matter benchmarks, using an identical protocol for each: $\Delta = 1$ means the abundance scales linearly with the parameter (perfectly natural), $\Delta \gg 1$ means it is exponentially sensitive (fine-tuned).

The principal results of this analysis are

(1) The fine-tuning of any dark matter construction is set by the analytic structure of its abundance map, *not* by whether the dark matter is a particle or a gravitational relic. Three universality classes emerge, and PBH constructions appear in all three.

(2) Biased-domain-wall PBHs exhibit the same low Barbieri-Giudice sensitivity, $\Delta = 2$–$4.5$, as off-resonance WIMPs, freeze-in particles, and post-inflationary axions, and a substantially lower sensitivity than the coannihilating or resonant-funnel WIMPs that survive current direct-detection limits, a statement about sensitivity within the $(\eta, V_b)$ parameter space (Sec. VI B), not a naturalness verdict on the mechanism as a whole. Comparing the most favorable noninflationary parametrization considered here (gravity-fixed domain walls, $\Delta = 2$) with the most pessimistic single-field inflationary estimate (Sec. III D, $\Delta \gtrsim 10^9$–$10^{10}$), the PBH paradigm alone spans more than seven orders of magnitude in $\Delta$; this is a comparison between the extremes of distinct formation mechanisms, not a uniform, like-for-like scan within a single model class.

(3) Constructions with a single exponential factor in their abundance universally satisfy $\Delta \approx 14$–$50$ on the observed dark matter contour, independently of the underlying microscopic physics; the precise value is set by the logarithm of the ratio of the formation temperature to the matter–radiation equality temperature. Coannihilating WIMPs and early matter-domination PBHs occupy this universality class robustly. First-order phase transition PBHs also fall here within a standard single-exponential approximation for the collapse probability, but a more accurate treatment raises the fine-tuning estimate significantly, and tracing the abundance back to the underlying scalar potential reveals a further double-exponential structure with $\Delta \sim 10^3$–$10^4$. Curvaton inflationary PBH models occupy this intermediate tier only for low reheating temperatures; for standard high-reheating scenarios they are elevated to the most-tuned class, with the reheating temperature setting the boundary between the two.

The off-resonance WIMP "miracle," by contrast, falls in the natural class: when the freeze-out exponential is absorbed into the measured annihilation cross section, the abundance map reduces to a pure power law with $\Delta = 2$.

(4) Resonance- or cancellation-dependent constructions (e.g., Higgs-funnel WIMP) and single-field inflationary collapse models are both highly tuned, but for distinct reasons: the former through a narrow annihilation pole, the latter through a double exponential in which the power spectrum amplitude is itself exponentially sensitive to inflaton potential coefficients, on top of an additional exponential sensitivity in the condition for collapse. Multifield and spectator-field inflationary models avoid the second exponential and fall into the intermediate universality class.

Throughout, the sensitivity measure $\Delta$ quantifies local parameter sensitivity in the sense of Barbieri and Giudice [16]. Where this sensitivity is not constant over the physically accessible parameter space (class II and III, Sec. V B), a Wilson-normalized naturalness verdict would additionally require specifying a reference ensemble, which I do not attempt here.

This paper is organized as follows. In Sec. II I define the sensitivity measures and establish the protocol applied uniformly to all scenarios. In Sec. III I derive the BG measure for three noninflationary PBH formation mechanisms (Secs. III A–III C), presenting analytic results and parameter-space heatmaps, and then map the framework onto six classes of inflationary PBH model (Sec. III D), including a systematic comparison of their two-layer tuning structure. In Sec. IV I apply the same analysis to seven particle dark matter benchmarks. In Sec. V I compare the full set of results, establish the three universality classes, and prove the structural identity $\Delta \approx \ln(\text{scale ratio})$ that unifies the class II scenarios; $\Delta_{\rm BG}$ is compared against two related quantities, the Strumia-Rattazzi measure $\Delta_{\rm SR}$ and the island half-width $\epsilon$, which track $\Delta_{\rm BG}$ closely for the benchmarks considered here rather than provide an independent robustness test. Section VI places the results in broader context, discusses limitations, and identifies directions for follow-up. Conclusions are in Sec. VII.

## II. THE BARBIERI-GIUDICE MEASURE APPLIED TO RELIC ABUNDANCES

### A. Definition and conventions

Fine-tuning in the context of new physics models was first quantified systematically by Barbieri and Giudice [16], who sought to measure how sensitively the electroweak scale depends on parameters in the supersymmetric spectrum. Given a physical observable $\mathcal{O}$ and a set of independent input parameters $\{x_i\}$, the Barbieri-Giudice (BG) measure is defined as

$$\Delta \equiv \max_i \Delta_i, \qquad \Delta_i \equiv \left| \frac{\partial \ln \mathcal{O}}{\partial \ln x_i} \right|. \tag{1}$$

The interpretation is direct: $\Delta_i = n$ means that a fractional change $\epsilon$ in the parameter $x_i$ produces a fractional change $n\epsilon$ in the observable $\mathcal{O}$. A construction is considered to be "natural" if $\Delta \sim \mathcal{O}(1)$, "moderately tuned" if $\Delta \sim \mathcal{O}(10)$–$\mathcal{O}(100)$, and "highly tuned" if $\Delta \gg 100$; the requirement $\Delta < 10$ is commonly adopted as a naturalness criterion in the SUSY literature [16,23], though no sharp threshold carries fundamental significance.

I stress that $\Delta$ is, strictly speaking, a measure of local logarithmic sensitivity: a large value does not automatically imply fine-tuning. A meaningful naturalness verdict requires comparing $\Delta_i$ to a reference sensitivity $\bar{\Delta}$, as in Wilson's criterion in the normalized form $\gamma \equiv \Delta_i/\bar{\Delta}$ advocated by Anderson and Castano [24] and adopted by Iovino and Riotto [22]. Specifying $\bar{\Delta}$ requires a choice of ensemble over which the average is taken; this choice is a convention rather than a derived quantity, in the same sense that the flat prior on the axion misalignment angle discussed in Sec. V E is a convention. For the constructions in this paper for which $\Delta$ is constant over the entire physically accessible parameter space (class I, Sec. V B), this convention dependence is immaterial: $\Delta$ trivially equals its own average for any choice of ensemble, so $\gamma = 1$ and the sensitivity and Wilson-normalized statements coincide without further input. For class II and class III constructions I report $\Delta$ as a sensitivity and do not attach an unqualified naturalness verdict to it.

In this paper the observable is the dark matter relic abundance, $\mathcal{O} \equiv \Omega h^2$, evaluated at the benchmark values of the model parameters and compared against the Planck 2018 measurement $\Omega_{\rm DM}h^2 = 0.1200 \pm 0.0012$ [25]. The sensitivity $\Delta_i$ measures how precisely one must specify $x_i$ in order for the predicted abundance to land within, say, a factor of two of $\Omega_{\rm DM}h^2$: a value $\Delta_i = 35$ means that a $\sim$2% variation in $x_i$ takes the predicted abundance outside this window.[1]

The BG measure is a *local* quantity: it characterizes the sensitivity of $\Omega_{\rm DM}h^2$ to small variations around a chosen benchmark point in parameter space. It should be distinguished from related but inequivalent notions of fine-tuning, including the prior probability of landing on the $\Omega = \Omega_{\rm DM}$ surface from a random draw of parameters, and the fractional volume of the natural region relative to the prior volume. I return to the distinction between these

[1] A fractional shift $\epsilon$ in $x_i$ produces a fractional shift $\Delta_i \cdot \epsilon$ in $\Omega h^2$; requiring the abundance to remain within a factor of two of the target imposes $\Delta_i \cdot \epsilon \lesssim \ln 2 \approx 0.69$, hence $\epsilon \lesssim 0.69/\Delta_i$. For $\Delta_i = 35$ this gives $\epsilon \lesssim 2\%$; more generally, the fractional precision requirement is $\epsilon \approx 69\%/\Delta_i$.

concepts, and to the specific case of misalignment-axion dark matter where it is particularly consequential, in Sec. V E.

The BG measure is symmetric in the sense that $\Delta_i$ takes the same value whether one uses $\mathcal{O} = \Omega h^2$ or $\mathcal{O} = 1/(\Omega h^2)$ as the observable, since $|\partial \ln \mathcal{O}/\partial \ln x_i|$ is unchanged by this inversion. For PBH scenarios, replacing $\Omega h^2$ with $f_{\rm PBH}$ also leaves $\Delta$ unchanged, since $f_{\rm PBH} = \Omega_{\rm PBH} h^2/\Omega_{\rm DM} h^2$ differs only by the constant $\Omega_{\rm DM} h^2$. It is *not* symmetric under reparametrizations of the parameter space: $\Delta$ evaluated with respect to parameters $\{x_i\}$ and $\{y_j = f_j(x)\}$ will generally differ if $f_j$ is nonlinear. The physical content of $\Delta$ therefore depends on the identification of fundamental parameters, which I address in Sec. II B.

For reference, a pure power-law abundance map of the form

$$\Omega h^2 \propto \prod_i x_i^{\alpha_i} \tag{2}$$

yields $\Delta_i = |\alpha_i|$ identically everywhere in parameter space. For an abundance set by a single exponential,

$$\Omega h^2 \propto e^{-Ax} \cdot x^{\beta}, \tag{3}$$

the dominant contribution on the $\Omega h^2 = \Omega_{\rm DM} h^2$ surface is

$$\Delta_x \approx Ax = \ln\left[\frac{\Omega_{\rm natural}}{\Omega_{\rm DM} h^2}\right], \tag{4}$$

where $\Omega_{\rm natural}$ is the abundance that would obtain if the exponential were replaced by unity. The full derivative is $\Delta_x = Ax + \beta$, so the approximation $\Delta_x \approx Ax$ overestimates the true value by $-\beta/(Ax)$; for the early matter-domination benchmark where $\beta = 5$ and $Ax \approx 9$–$19$, the correction is 20%–35% and is the reason the benchmark $\Delta_\sigma = 14.4$ (which includes the prefactor term $+5$) is larger than the pure-exponential approximation $Ax \approx 9$–$10$ at $\sigma = 0.055$. This result, which is exact when $\beta = 0$ and approximate otherwise, is central to understanding why the three "single-exponential" dark matter constructions discussed in Sec. V yield a universal value $\Delta \approx \ln(T_{\rm form}/T_{\rm eq}) \approx 14$–$50$. Constructions with a Breit–Wigner resonance yield $\Delta \sim M^2/(\Gamma M) \gg \ln(\text{scale ratio})$, a qualitatively distinct and much more severe tuning.

Two complementary measures are computed alongside BG throughout this paper. The Strumia-Rattazzi (SR) measure [26] replaces the maximum with a quadrature sum over all parameter directions:

$$\Delta_{\rm SR} \equiv \sqrt{\sum_i \Delta_i^2} = |\nabla_{\ln x} \ln \Omega h^2|. \tag{5}$$

This equals the magnitude of the log-gradient vector and is sensitive to fine-tuning distributed across multiple parameters simultaneously. For single-parameter scenarios, obviously, $\Delta_{\rm SR} = \Delta_{\rm BG}$; for scenarios where all parameters are equally sensitive (e.g. ADM with $\Delta_M = \Delta_R = 1$), $\Delta_{\rm SR} = \sqrt{n}\Delta_{\rm BG}$ where $n$ is the number of parameters.

The island half-width $\epsilon$ in the most sensitive parameter direction provides the most operationally transparent measure of tuning. Rather than asking how steeply the abundance responds to a parameter shift (BG) or summing sensitivities in quadrature (SR), $\epsilon$ asks directly: what is the fractional range of the most sensitive parameter that is compatible with the observed dark matter density? Formally,

$$\epsilon \equiv \left.\frac{\Delta x_{\rm sens}}{x_{\rm sens}}\right|_{\rm island} \approx \frac{\ln 2}{\Delta_{\rm BG}} \approx \frac{0.693}{\Delta_{\rm BG}}, \tag{6}$$

where $\Delta x_{\rm sens}$ is the full fractional width of the natural island in the direction of the most sensitive parameter $x_{\rm sens}$ (the parameter with the largest $|\Delta_i|$), and the approximation holds for a factor-of-2 abundance band. The equivalence $\epsilon \approx \ln 2/\Delta_{\rm BG}$ follows directly from the definition of $\Delta_{\rm BG}$: a fractional shift $\delta x/x$ changes $\ln \Omega$ by $\Delta_{\rm BG} \cdot \delta x/x$, so the island boundary is reached when $\Delta_{\rm BG} \cdot \epsilon \approx \ln 2$ (a factor of 2 in $\Omega$). The result is a single percentage that answers the most experimentally concrete version of the naturalness question. To illustrate with two benchmarks defined below: benchmark B3 (coannihilation, Sec. IV A 3, $\Delta = 48$) requires $M_{\rm NLSP}$ to be specified to within $\epsilon = 1.4\%$ to remain in the natural island; benchmark B5 (asymmetric dark matter, Sec. IV C, $\Delta = 1$) tolerates a factor of two in either direction ($\epsilon = 75\%$).

The three measures, BG, SR, and $\epsilon$, are not independent; they are related by construction and probe the same underlying sensitivity from different angles. BG isolates the worst-case direction; SR weights all directions equally; $\epsilon$ converts the worst-case sensitivity into an experimentally interpretable precision requirement. In single-parameter scenarios all three agree exactly. In multiparameter scenarios BG and SR diverge by at most $\sqrt{n}$ (where $n$ is the number of equally sensitive parameters), and $\epsilon$ always tracks BG through Eq. (6). The three measures have been applied and compared in the context of supersymmetric dark matter and electroweak naturalness by several authors. Anderson and Castano [24] introduced a normalized variant of BG that accounts for the expected sensitivity of the prior; Strumia [26] proposed the quadrature measure adopted here as $\Delta_{\rm SR}$; Cabrera *et al.* [27] applied both BG and Bayesian posterior-volume measures to MSSM dark matter and found them to agree on tier assignments while disagreeing on absolute values by factors of a few.

A separate line of work has addressed fine-tuning measures specifically in the context of inflationary PBH production. Iovino and Riotto [22] identified a pathology of the BG measure when applied to the inflaton potential coefficients $\{c_k\}$: as the target PBH abundance is reduced

below the observed dark matter density (i.e. as $f_{\rm PBH} \to 0$), the BG measure grows without bound even though a model predicting fewer PBHs is physically less constrained, not more. This occurs because BG measures the local log-slope $|\partial \ln \Omega / \partial \ln c_k|$ without reference to whether $\Omega$ is above or below the target; as the exponential suppression is tuned further, the log-slope at any fixed point in parameter space continues to grow. Iovino and Riotto proposed Wilson's naturalness criterion as a remedy: a parameter is natural if its value is stable under small variations of the UV completion, independent of the abundance it produces. Their measure assigns a value close to unity to essentially all inflationary PBH models, in sharp contrast to the BG values of $10^2$–$10^8$ discussed in Sec. III D.

Evaluating all three measures (BG, SR, $\epsilon$) *on* the $f_{\rm PBH} = 1$ contour, rather than at an arbitrary point in parameter space, removes this specific divergence: the bare BG derivative, evaluated off the target contour, grows without bound as $f_{\rm PBH} \to 0$ even for parameter variations that do not affect the physical viability of the model. The relevant question is not "how sensitive is the model at some generic parameter point?" but "how precisely must each parameter be specified to maintain the observed dark matter density?"; and this question is only meaningful on the $f_{\rm PBH} = 1$ contour. Evaluated there, the BG measure is finite and well defined for all scenarios in the comparison, including inflationary collapse [Eq. (44)], and the Iovino-Riotto divergence does not arise. This is not the same as supplying the reference sensitivity $\bar{\Delta}$ that a Wilson-normalized measure requires, and I do not claim that it is. The Wilson question, "does a technically natural UV completion exist for the required parameter values?," and the sensitivity question asked here, "how precisely must a given parameter be specified to remain on the observed contour?," are complementary and, properly stated, not in tension; Sec. V E returns to this distinction in the context of the prior-probability vs local-sensitivity debate. To the best of my knowledge, this is the first study to apply all three measures simultaneously to a cross-paradigm comparison that includes PBH formation mechanisms alongside particle dark matter candidates.

### B. Choice of fundamental parameters

Because the BG measure is not reparametrization-invariant, the choice of what to call a "fundamental" parameter must be specified and physically motivated. I adopt the following convention throughout this paper: the fundamental parameters are the independently specified inputs to the UV Lagrangian or to the cosmological initial conditions of the model under consideration, evaluated at the relevant renormalization scale. For a particle physics model these are the pole masses and Lagrangian coupling constants; in supersymmetric models they are the soft supersymmetry-breaking masses at the mediation scale; for cosmological constructions they are the scales and couplings of the relevant phase transition or symmetry-breaking potential, together with any field-space initial conditions.

This convention has concrete implications for several of the scenarios analyzed below.

*Coannihilation.* In the MSSM [28], supersymmetry is broken in a hidden sector and the breaking is transmitted to the visible sector through a mediation mechanism (gravity, gauge, or anomaly mediation) at a high scale $M_{\rm med}$. The low-energy spectrum is then determined by a set of soft-breaking mass terms in the Lagrangian; crucially, each soft mass is an independent parameter at $M_{\rm med}$, and there is no symmetry or dynamics that correlates them in the absence of additional model structure. For the coannihilation scenario, the two relevant soft masses are $M_1$ (the bino mass parameter, which sets $M_\chi \equiv |M_1|$ at low energies) and the soft mass $m_{\tilde{f}}$ of the lightest sfermion (stop, stau, or sneutrino), which sets $M_{\rm NLSP}$ (I take the bino as the lightest neutralino and the lightest sfermion as the next-to-lightest supersymmetric particle (NLSP); the analysis is qualitatively identical for any spectrum where the two lightest states are nearly degenerate). Coannihilation is relevant when $M_{\rm NLSP}$ and $M_\chi$ are nearly degenerate: the NLSP is present in the thermal bath at freeze-out and its annihilation cross section $\langle \sigma v \rangle_{\rm NLSP}$ contributes to the effective rate that sets the relic density. The contribution is Boltzmann-suppressed by a factor $e^{-x_F \delta}$ [Eq. (52)], where $x_F \equiv m_\chi / T_F \approx 25$ is the ratio of the dark matter mass to the freeze-out temperature $T_F$ [29,30] and $\delta \equiv (M_{\rm NLSP} - M_\chi)/M_\chi$ is the fractional mass splitting. Since $T_F \simeq m_\chi / x_F$ is set by the freeze-out condition $\langle \sigma v \rangle n_\chi = H$, and $x_F \approx$ 20–30 for weak-scale dark matter regardless of the exact cross section [29], the exponential suppression can be substantial even for splitting $\delta \sim 0.05$–$0.10$.

The fractional mass splitting $\delta$ is a derived quantity: it is the difference of two Lagrangian parameters, not itself a fundamental input. Taking BG derivatives with respect to $\delta$ at fixed $M_\chi$ gives

$$\Delta_\delta = x_F \delta + \mathcal{O}(\delta^2) \approx 1 - 2 \quad (x_F \delta \ll 1), \tag{7}$$

which substantially underestimates the tuning because the derivative $\partial \ln \Omega / \partial \ln \delta$ vanishes as $\delta \to 0$ by construction: the coordinate $\delta$ crosses zero at the threshold where the two masses are degenerate, and multiplicative log-derivatives are pathologically small near any additive zero. Taking derivatives with respect to the independent Lagrangian parameters $(M_\chi, M_{\rm NLSP})$ at fixed mediation-scale boundary conditions instead yields

$$\Delta_{M_{\rm NLSP}} = x_F \frac{M_{\rm NLSP}}{M_\chi} \approx 26\text{–}50, \tag{8}$$

which correctly captures the sensitivity of the Boltzmann-suppressed coannihilation rate to the NLSP pole mass: a

2% shift in $M_{\rm NLSP}$ at fixed $M_\chi$ changes $\delta$ by 2%, which changes $e^{-x_F\delta}$ by a factor $e^{\pm 0.5} \approx 1.6$, shifting $\Omega h^2$ by ∼50% and taking the model outside the natural island. The MSSM coannihilation benchmark B3 uses the parametrization $(M_\chi, M_{\rm NLSP})$ and yields $\Delta \approx 48$ as a consequence [Eq. (8); see also Sec. IV A 3]. This distinction between derived and fundamental parameters has been noted in the electroweak fine-tuning literature [31] but is rarely applied consistently to dark matter relic calculations.

*Axion misalignment.* The initial misalignment angle $\theta_i \in [-\pi, \pi]$ is a cosmological initial condition, not a Lagrangian parameter, and is independently specified from $f_a$, the axion decay constant. It is therefore treated as a separate fundamental input. The important distinction between the BG measure evaluated with respect to $\theta_i$ (which gives $\Delta \approx 2$ for $\theta_i \sim 1$) and the prior probability of $\theta_i$ being small (which gives a prior tuning $\sim 1/\theta_i$ for a flat distribution) is discussed in Sec. V E.

*PBH formation rates.* For the PBH mechanisms analyzed in Sec. III, the fundamental parameters are the scales and couplings appearing in the potential or power spectrum of the underlying model, e.g. the symmetry-breaking scale $\eta$ and bias energy density $V_b$ for domain walls, the modulus mass $m_\phi$ and curvature perturbation amplitude $\sigma$ for early matter domination, and the transition temperature $T_*$, nucleation rate parameter $\beta/H$, and latent-heat parameter $\alpha$ for the FOPT. All three sets of parameters are taken as independent inputs, and BG derivatives are evaluated holding all other parameters fixed.

### C. Numerical protocol

All BG derivatives in this paper are evaluated numerically via central finite differences in log-space,

$$\Delta_i \approx \frac{|\ln\Omega(x_i e^h) - \ln\Omega(x_i e^{-h})|}{2h}, \quad h = 10^{-4}. \tag{9}$$

This step size is small enough to suppress second-order corrections to better than 1 part in $10^6$ for all smooth abundance functions considered, yet large enough to avoid floating-point cancellation errors. Analytic closed-form expressions for $\Delta_i$ are derived independently in Appendix A for all cases where the abundance map is sufficiently simple; numerical and analytic results agree to better than 0.1% in every case.

Parameter-space heatmaps are computed on grids of $400 \times 400$ points, logarithmically spaced over the ranges specified in the captions of Figs. 1–5. The color scale is identical across all panels: $\log_{10}\Delta \in [0, 3.5]$ on the `viridis_r` color map, saturating to dark purple for $\Delta > 3000$. This shared scale is here the principal visual device for comparing fine-tuning across production paradigms.

The natural island in each panel is defined as the region of parameter space satisfying

$$0.06 \leq \Omega h^2 \leq 0.24 \tag{10}$$

(i.e. within a factor of two of $\Omega_{\rm DM} h^2 = 0.120$) together with all constraints listed in Table I. The factor-of-two band is chosen to include representative theoretical uncertainties in the relic-abundance calculation (QCD uncertainties in $g_*$, perturbative corrections to annihilation cross sections, etc.) without enlarging the island to the point of obscuring the tuning structure. Results are qualitatively insensitive to widening this band to a factor of 3 or narrowing it to 30%. For PBH mechanisms, $\Omega h^2$ is replaced by $f_{\rm PBH}$ (both measure the fraction of dark matter), and the gray dashed contours in the PBH heatmaps show the wider ±1 dex sidebands $f_{\rm PBH} \in [0.1, 10]$ as visual guides; all BG derivatives and island half-widths are computed using the factor-of-two band of Eq. (10) applied to $f_{\rm PBH}$ (i.e., $f_{\rm PBH} \in [0.06/0.12, 0.24/0.12] \equiv [0.5, 2]$), consistent with the particle DM protocol.

The direct-detection constraint for WIMPs is implemented via the spin-independent cross section per nucleon,

$$\sigma_{\rm SI}(M_S, \lambda) = \frac{\lambda^2 f_N^2 m_N^2 \mu_{SN}^2}{\pi M_S^2 M_h^4}, \tag{11}$$

where $f_N = 0.30$ is the nucleon Higgs form factor [44], $m_N = 0.939$ GeV, and $\mu_{SN} = m_N M_S/(m_N + M_S)$ is the singlet-nucleon reduced mass. The LZ 2024 sensitivity curve is approximated by a piecewise power law calibrated to the published limit [40]; this prescription overestimates exclusion at $M_S \lesssim 10$ GeV and underestimates it above ∼10 TeV, but neither regime contains any of the benchmarks analyzed in this paper. The Lyman-$\alpha$ forest constraint on the freeze-in DM mass ($M_{\rm DM} \gtrsim 3$ keV for production from a 100-GeV mediator [41,42]) is implemented as a hard boundary in the FIMP parameter plane.

## III. PRIMORDIAL BLACK HOLE DARK MATTER

The formation of PBHs in the asteroid-mass window has been studied in a wide variety of cosmological scenarios. Three noninflationary mechanisms have received particular attention for their ability to produce $f_{\rm PBH} \approx 1$ without additional cosmological ingredients: a network of biased domain walls [45–47], an early matter-dominated era driven by a decaying modulus field [48–51], and a strongly first-order phase transition in a hidden sector [52–54]. The idea that biased domain walls can collapse to form a PBH dark matter population was developed in Ref. [45] in the context of QCD axion domain wall networks, building on the long-known instability of biased networks [55]; it was subsequently revisited and extended by several authors, including the network-collapse treatments of Refs. [46,47]. The most extensively discussed mechanism in the literature is single-field inflation with an enhanced small-scale power spectrum, generated by an

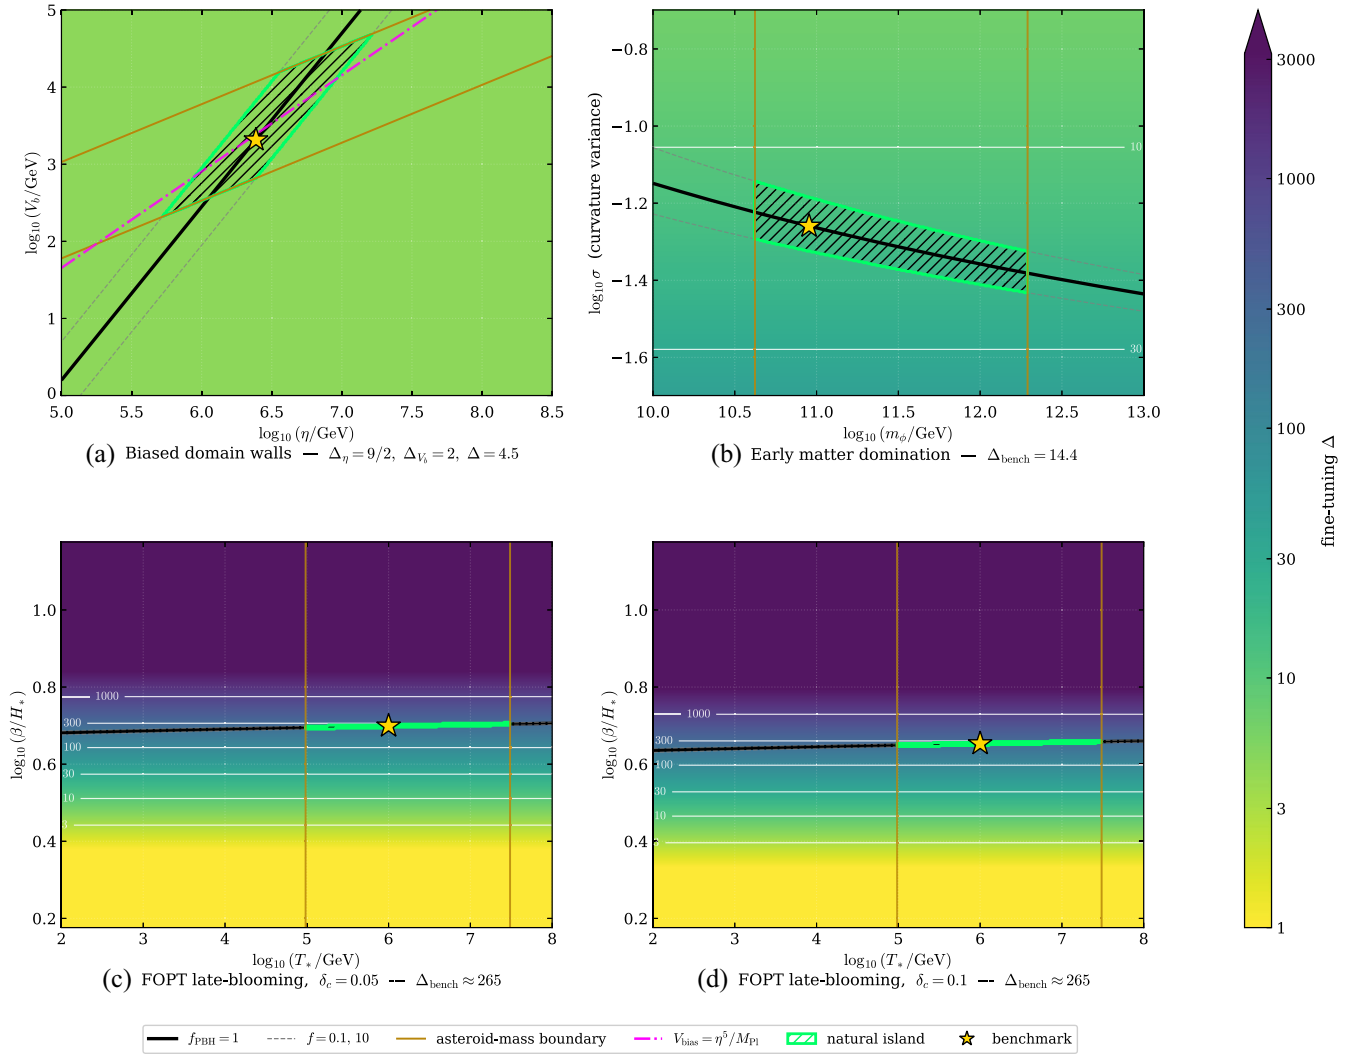



FIG. 1. Fine-tuning of the three PBH dark-matter formation mechanisms analyzed in this paper. Background color: $\log_{10}\Delta$ on a shared scale $\Delta \in [1, 3000]$ (pale yellow = natural, dark purple = highly tuned; see color bar at right). Thick black contour: $f_{\rm PBH} = 1$. Gray dashed contours: $f_{\rm PBH} = 0.1$ and 10 ($\pm 1$ dex sidebands). Gold/orange vertical lines: asteroid-mass boundaries $M_{\rm PBH} = 10^{17}$ g and $10^{22}$ g. Green hatching: natural island, defined as $0.1 \le f_{\rm PBH} \le 10$ within the asteroid-mass window and satisfying all observational constraints. Gold star: benchmark point. (a) Biased domain walls in the $(\eta, V_b)$ plane [$\log_{10}$ axes; $\eta$ = symmetry-breaking scale, $V_b$ = bias energy scale; Eq. (17)]. The magenta dash-dotted line is the gravity-induced-bias locus $V_b = (\eta^5/M_{\rm Pl})^{1/4}$. The panel is uniformly pale yellow, reflecting $\Delta = 9/2$ everywhere [Eq. (21)]. Benchmark: $\eta = 2.4 \times 10^6$ GeV, $V_b = 2$ TeV. (b) Early matter domination in the $(m_\phi, \sigma)$ plane [$m_\phi$ = modulus mass, $\sigma$ = curvature perturbation variance; Eqs. (23)–(25)]. The color gradient from yellow-green to teal reflects the rising $\Delta_\sigma = 5 + 0.196/\sigma^{4/3}$ [Eq. (27)] as $\sigma$ decreases. The asteroid-mass boundaries correspond to $m_\phi \in [4 \times 10^{10}, 2 \times 10^{12}]$ GeV [Eq. (24)]. Benchmark: $m_\phi = 9 \times 10^{10}$ GeV, $\sigma = 0.055$, giving $M_{\rm PBH} \approx 10^{21}$ g; $\Delta = 14.4$. (c),(d) FOPT late-blooming in the $(T_*, \beta/H_*)$ plane, now computed with the Gouttenoire-Vitagliano (GV) superexponential collapse probability Eq. (70), for collapse threshold $\delta_c = 0.05$ [panel (c)] and $\delta_c = 0.10$ [panel (d)]. The GV formula contains no $\alpha$ dependence; the two panels illustrate instead the sensitivity to $\delta_c$. Background color: $\log_{10}\Delta$ on the shared viridis_r scale. The color field transitions sharply near $\beta/H_* \approx 5$ [panel (c)] and $\approx 4.5$ [panel (d)]: yellow-green below ($f_{\rm PBH} > 1$, overproduction), deep teal-to-purple at and above the $f_{\rm PBH} = 1$ contour ($\Delta \approx 265$–$300$, class III). The $f_{\rm PBH} = 1$ contour (thick black) is nearly flat, running at $\beta/H_* \approx 4.9$–$5.1$ [panel (c)] and $\approx 4.4$–$4.6$ (panel (d) across the full asteroid-mass temperature range. The natural island (green hatching, $0.1 \le f_{\rm PBH} \le 10$) is an extremely narrow strip of width $\sim 0.1$ in $\beta/H_*$ (corresponding to $\epsilon \approx 0.26\%$, class III), visible only as a thin band; its narrowness is itself a visual demonstration of the severe fine-tuning. The previous analytic structural identity $\Delta_{\beta/H} = \ln(T_*/0.84 T_{\rm eq}) \approx 35$ [Eq. (37)], derived from the single-exponential approximation Eq. (30), does not apply here; see Sec. VI B for the full comparison. Benchmark: $T_* = 10^6$ GeV, $\beta/H_* = 5.0$ ($\delta_c = 0.05$) and $\beta/H_* = 4.5$ ($\delta_c = 0.10$); both at $\Delta \approx 265$ [Eq. (72)].

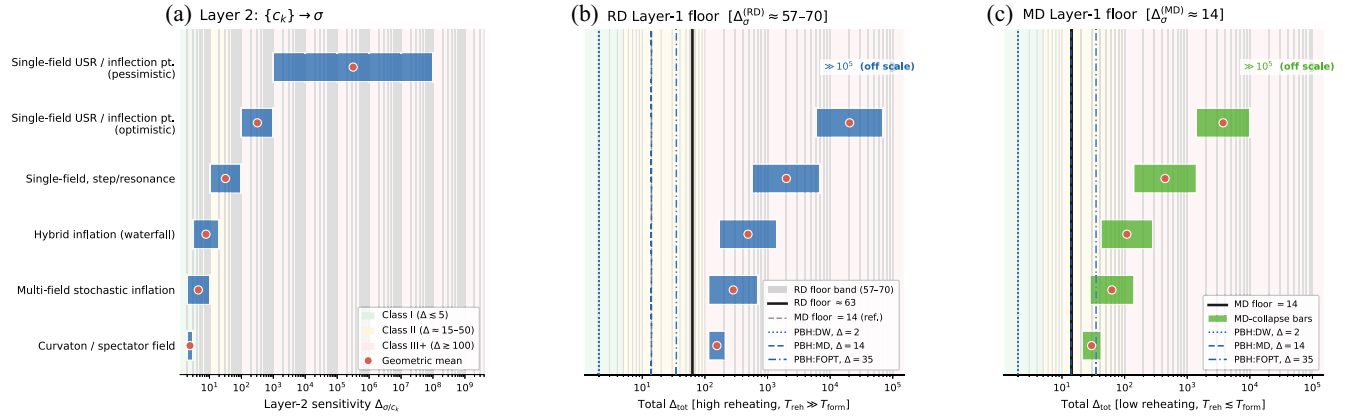



FIG. 2. Comparison of layer 2 sensitivity [$\Delta_{\sigma/c_k}$, panel (a)] and total fine-tuning [$\Delta_{tot} = \Delta_\sigma \times \Delta_{\sigma/c_k}$, panels (b) and (c)] across six inflationary PBH model classes, ordered from most natural (bottom) to most tuned (top). Bars show the estimated range within each class; the dot marks the geometric mean. Background shading indicates qualitative fine-tuning tier bands (class I, II, and III; defined in Sec. V B). Panel (b): total fine-tuning with the matter-dominated HYKN layer 1 floor $\Delta_\sigma^{(MD)} \approx 14$ [Eq. (38)], applicable when $T_{reh} \lesssim T_{form}$ (low reheating, inflaton-oscillation MD era at reentry). Panel (c): total fine-tuning with the radiation-dominated Press-Schechter layer 1 floor $\Delta_\sigma^{(RD)} \approx 57$–$70$ [Eq. (41)], applicable when $T_{reh} \gg T_{form}$ (standard high reheating, RD background at reentry). The $x$ axes of panels (b) and (c) are identical (capped at $10^5$) to make the factor-of-five shift in the layer 1 floor directly visible; the top row (single-field USR pessimistic, $\Delta_{tot} \gtrsim 10^9$–$10^{10}$) is off scale and indicated by an arrow. Thick vertical gray lines mark the respective layer 1 floor in each total-tuning panel; in panel (c) the gray band spans the 57–70 range. Blue vertical lines show the three noninflationary PBH benchmarks (dotted: DW gravity-fixed, $\Delta = 2$; dashed: early MD, $\Delta = 14$; dash-dotted: FOPT, $\Delta = 35$, single-exponential approximation; see Sec. VI B for the more accurate estimate $\Delta \approx 265$). All ranges are estimates based on analytic scaling arguments and the results of Refs. [19,32–39]; precise values are model-dependent.

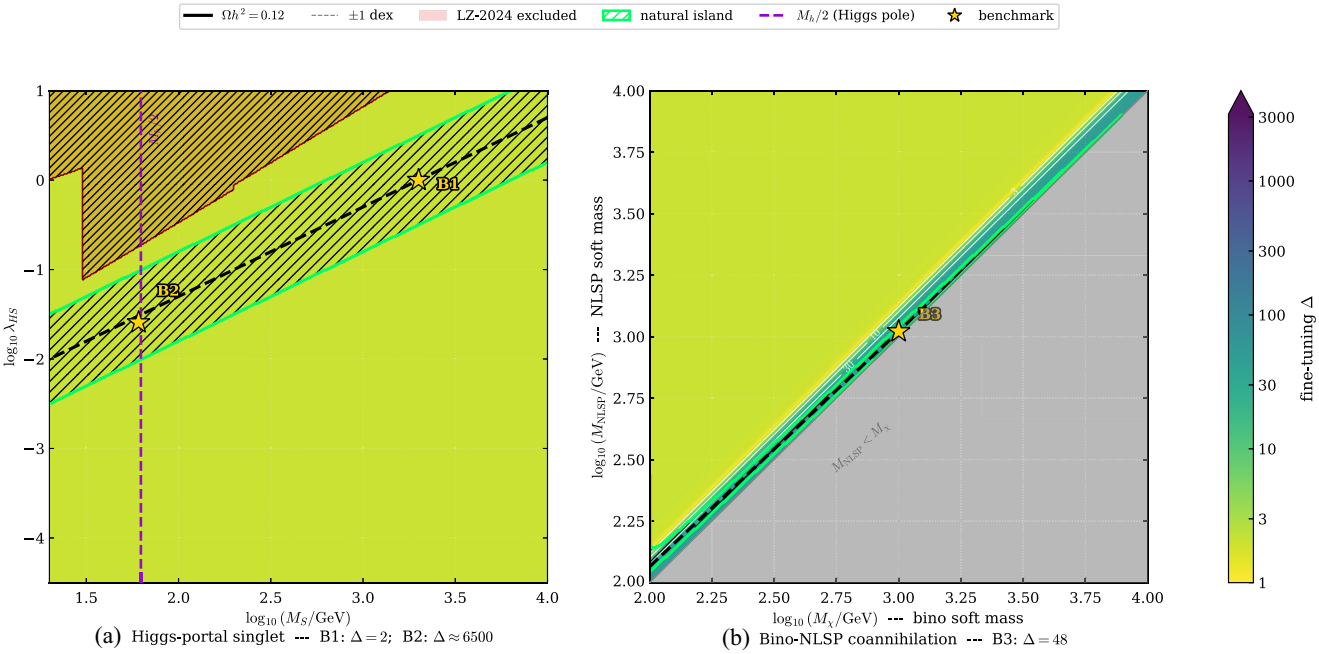
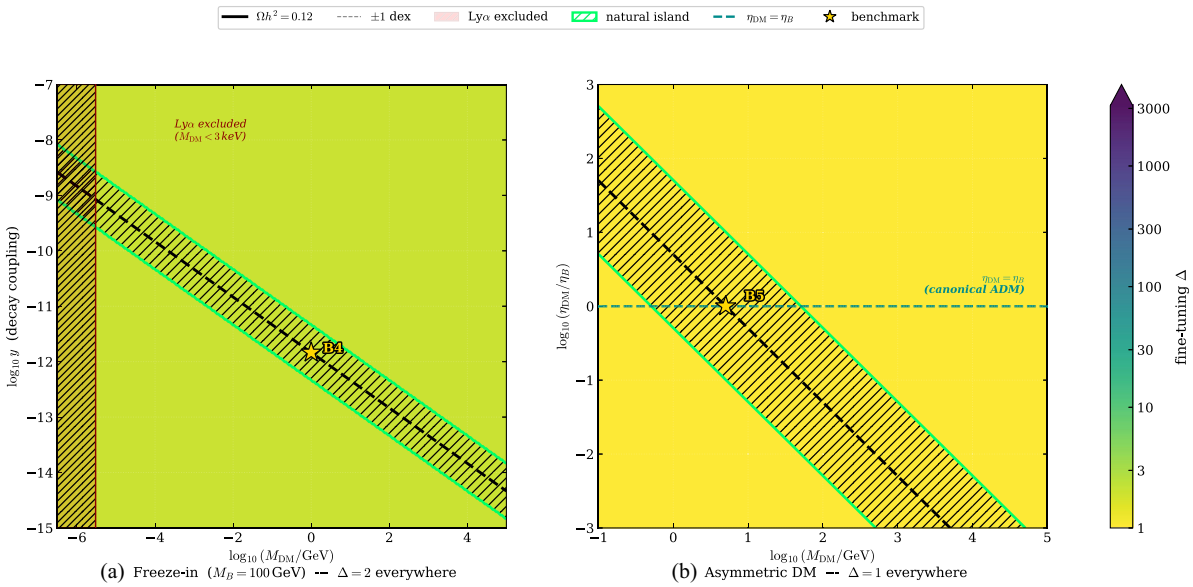
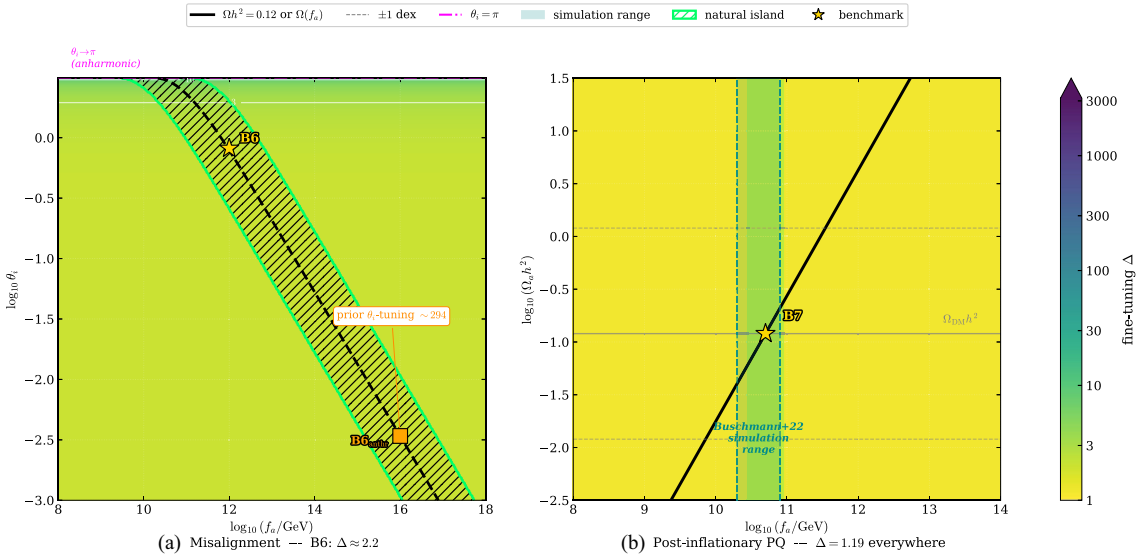



FIG. 3. Fine-tuning of thermal-relic WIMP dark matter. Color scale, contours, and natural island (green hatching) as in Fig. 1, with $\Omega h^2 = 0.12$ replacing $f_{PBH} = 1$ as the abundance target. The red-shaded region is excluded at 90% CL by LZ 2024 spin-independent direct detection [40]. (a) Real singlet Higgs-portal WIMP in the $(M_S, \lambda_{HS})$ plane [Eq. (46)]. The cyan dotted vertical line marks $M_S = M_h/2$ (Higgs pole). Benchmark B1 (off-resonance, gold star): $M_S = 2$ TeV, $\lambda_{HS} = 1$; $\Delta = 2.0$. Benchmark B2 (Higgs funnel, gold star): $M_S = 60.7$ GeV, $\lambda_{HS} = 0.026$; $\Delta \approx 6500$ (color saturated to dark purple). The contrast between B1 (uniformly yellow stripe) and B2 (deep purple near the resonance) illustrates the class I vs class III distinction. (b) MSSM bino-NLSP coannihilation in the fundamental soft-mass plane $(M_\chi, M_{NLSP})$, using Eqs. (52)–(54). The gray-shaded triangle is the unphysical region $M_{NLSP} < M_\chi$. The natural island (narrow teal strip) hugs the diagonal at constant $\Delta \approx 48 = x_F M_{NLSP}/M_\chi$. Benchmark B3: $M_\chi = 1$ TeV, $M_{NLSP} \approx 1.0515$ TeV; $\Delta = 48$.




(a) Freeze-in ($M_B = 100$ GeV) -- $\Delta = 2$ everywhere (b) Asymmetric DM -- $\Delta = 1$ everywhere

FIG. 4. Fine-tuning of nonthermal dark matter production. Color scale and conventions as in Fig. 1. Both panels are uniformly pale yellow, reflecting power-law abundance maps with $\Delta \leq 2$ everywhere. (a) Freeze-in via heavy-mediator decay $B \to \mathrm{DM} + X$ with fixed mediator mass $M_B = 100$ GeV, in the $(M_{\mathrm{DM}}, y)$ plane [Eqs. (55) and (56)]. The red-shaded region at $M_{\mathrm{DM}} \lesssim 3$ keV is excluded by Lyman-$\alpha$ forest observations [41,42]. The natural island is the diagonal stripe $y \propto M_{\mathrm{DM}}^{-1/2}$ (from $\Omega \propto M_{\mathrm{DM}} y^2$). $\Delta = 2$ throughout the entire accessible parameter space [Eq. (57)]. Benchmark B4: $M_{\mathrm{DM}} = 1$ GeV, $y = 1.46 \times 10^{-12}$; $\Delta = 2$. (b) Asymmetric dark matter in the $(M_{\mathrm{DM}}, \eta_{\mathrm{DM}}/\eta_B)$ plane [Eq. (58)]. The cyan dashed horizontal line marks the "canonical ADM" locus $\eta_{\mathrm{DM}} = \eta_B$. The natural island is the hyperbola $M_{\mathrm{DM}}(\eta_{\mathrm{DM}}/\eta_B) = 4.7$ GeV. $\Delta = 1$ throughout [Eq. (60)], the minimum possible BG measure for a nontrivial abundance map. Benchmark B5: $M_{\mathrm{DM}} = 5$ GeV, $\eta_{\mathrm{DM}}/\eta_B = 1$; $\Delta = 1$.




(a) Misalignment -- B6: $\Delta \approx 2.2$ (b) Post-inflationary PQ -- $\Delta = 1.19$ everywhere

FIG. 5. Fine-tuning of QCD axion dark matter for two cosmological histories of PQ symmetry breaking. Color scale and conventions as in Fig. 1. (a) Misalignment mechanism (preinflationary PQ breaking) in the $(f_a, \theta_i)$ plane [Eq. (61)]. The magenta dash-dotted horizontal line at $\theta_i = \pi$ marks the onset of the anharmonic divergence $F(\theta_i) \to \infty$. The color gradient near this line reflects the diverging $\Delta_{\theta_i}$. The natural island (green hatching) runs diagonally from small $f_a$, $\theta_i \sim 1$ to large $f_a$, small $\theta_i$. Gold star B6: $f_a = 10^{12}$ GeV, $\theta_i \approx 0.787$; $\Delta = 2.2$. Orange square $\mathrm{B6_{anthr}}$: $f_a = 10^{16}$ GeV, $\theta_i \approx 3 \times 10^{-3}$; local $\Delta = 2.0$, but prior tuning $\approx 300$ (see Sec. V E). (b) Postinflationary PQ (string-network scenario), showing $\Omega_a h^2$ vs $f_a$ (black curve) against the uniform $\Delta = 1.19$ background [Eq. (64)]. The entire panel is uniformly pale yellow: $\Delta_{f_a} = 1.19$ everywhere, the lowest fine-tuning of any construction in the comparison other than pure asymmetric dark matter. The cyan shaded band marks the preferred $f_a$ range $(2-8) \times 10^{10}$ GeV from the large-scale string-network simulations of Ref. [43]. The green shaded band marks the natural island. Gray horizontal lines show $\Omega_{\mathrm{DM}} h^2 = 0.12$ (solid) and $\pm 1$ dex (dashed). Gold star B7: $f_a = 5 \times 10^{10}$ GeV; $\Delta = 1.19$.

TABLE I. Observational constraints applied in this paper and the scenarios to which they are relevant. The stochastic gravitational-wave background from the FOPT itself is not listed here because it constrains the GW signal rather than $f_{\rm PBH}$ directly; this constraint is discussed qualitatively in Secs. III C and VI A. All PBH abundance constraints are already encoded in the $f_{\rm PBH} = 1$ requirement and the asteroid-mass window.

| Constraint | Value | Applies to |
|---|---|---|
| LZ-2024 SI cross section | see [40] | WIMPs |
| Planck 2018 $\Omega_{\rm DM}h^2$ | $0.1200 \pm 0.0012$ | All |
| PBH asteroid window | $10^{17}$–$10^{22}$ g | PBHs |
| Ly$\alpha$ forest | $M_{\rm DM} \geq 3$ keV | FIMPs |

inflection point or ultraslow-roll phase in the inflaton potential [19,56,57], following the foundational treatments of García-Bellido and Ruiz Morales [58] and Germani and Prokopec [59]; multifield and spectator-field variants offer partial relief from the severe inflationary fine-tuning [20]. Concretely, Refs. [60,61] demonstrated that the USR potential shape arises naturally in multifield models with nonminimal couplings to gravity, and performed the first full MCMC analysis of CMB constraints and PBH formation in this setting, identifying a degeneracy direction along which parameter shifts of ∼10% preserve the fit to Planck data. More recently, Refs. [62,63] showed that adding a noninteracting spectator field to a single-field USR potential changes the formation mechanism entirely: for example, the system no longer enters USR, and the resulting model predictions become exponentially more robust to small parameter changes, substantially reducing the layer 2 fine-tuning. Within the framework of this paper, these constructions populate the lower end of the multifield stochastic inflation range in Fig. 2, and move toward the boundary between class II and class III for low-reheating scenarios. A distinct inflationary mechanism is the collapse of Higgs field fluctuations near the electroweak vacuum instability scale [64].

For the three noninflationary mechanisms, I derive the abundance map $f_{\rm PBH}(\{x_i\})$ from the underlying model parameters, compute $\Delta$ analytically, and display the result as a parameter-space heatmap (Secs. III A–III C). Inflationary mechanisms are treated separately in Sec. III D: rather than a single heatmap, I map the BG framework onto six representative inflationary model classes and compare their two-layer tuning structure systematically (Fig. 2). This distinction in treatment reflects a genuine physical difference. For the noninflationary mechanisms, the fundamental Lagrangian parameters map onto $f_{\rm PBH}$ through a single, well-calibrated chain with no hidden exponentials; the BG measure can therefore be computed precisely and displayed as a smooth landscape. For inflationary models, the mapping from inflaton potential coefficients $\{c_k\}$ to $\sigma$ to $f_{\rm PBH}$ involves two successive exponentials whose ratio depends sensitively on the potential shape, leading to estimates that span several orders of magnitude depending on the model class. Furthermore, as noted by Iovino and Riotto [22], applying BG to the potential coefficients off the $f_{\rm PBH} = 1$ contour raises a subtlety: $\Delta$ grows without bound as $f_{\rm PBH} \to 0$ even for variations that are physically harmless. The approach adopted throughout this paper, i.e. evaluating all measures *on* the $f_{\rm PBH} = 1$ contour, avoids this pathology by construction (see Sec. V D).

All four mechanisms produce PBHs whose mass is set by the Hubble horizon at formation,

$$M_{\rm PBH} \simeq \frac{3.8 M_{\rm Pl}^3}{T_{\rm form}^2}, \tag{12}$$

where $T_{\rm form}$ is the temperature at which PBH formation occurs (in natural units $c = k_B = 1$). The asteroid-mass window $10^{17}$–$10^{22}$ g requires formation temperatures $10^4 \lesssim T_{\rm form}/{\rm GeV} \lesssim 10^{6.5}$, shown as gold vertical lines in Fig. 1.

### A. Biased domain walls

#### 1. Formation mechanism and abundance map

Consider a scalar field $\phi$ with a discrete $\mathbb{Z}_2$ symmetry $\phi \to -\phi$ that is spontaneously broken when $\phi$ acquires a vacuum expectation value $\pm\eta$ at a phase transition in the early Universe. The resulting potential has two degenerate minima separated by a potential barrier, and the field takes different values in causally disconnected regions; the boundaries between these regions form a network of two-dimensional topological defects called domain walls [55,65]. Each wall carries a surface energy density (tension) $\sigma_w \sim \eta^3$, where $\eta$ is the symmetry-breaking scale; the exact coefficient is model-dependent and of order unity, and is absorbed into $\eta$ throughout.

In the exact $\mathbb{Z}_2$ limit the two vacua are energetically degenerate and the wall network is cosmologically stable [65], which would lead to domain-wall domination and conflict with cosmology [55]. A small explicit breaking of the $\mathbb{Z}_2$ symmetry lifts this degeneracy by an energy density $V_{\rm bias}$ between the two vacua. Such a bias arises naturally from Planck-suppressed operators in the scalar potential [66,67], or can be introduced by hand; in either case it can be parametrized as $V_{\rm bias} = V_b^4$, where $V_b$ is an effective bias energy scale. The pressure difference $V_{\rm bias}$ across each wall provides a volume force that drives the lower-energy vacuum to expand, causing the network to collapse and annihilate at a temperature $T_{\rm ann}$ set by the competition between this volume pressure and the Hubble friction that resists wall motion [45,55]:

$$T_{\rm ann} \simeq \left(\frac{90}{\pi^2 g_*}\right)^{1/4} \sqrt{\frac{c_t M_{\rm Pl} V_b^4}{\eta^3}}, \tag{13}$$

where $g_* \approx 100$ counts relativistic degrees of freedom and $c_t \approx 0.30$ is a numerical coefficient characterizing the wall-network annihilation dynamics (see Ref. [45], which estimated the resulting PBH abundance using the earlier field-theoretic wall-network simulations of Ref. [68]; a dedicated set of network-collapse simulations was performed subsequently in Ref. [47]). At $T_{\rm ann}$, a fraction $p \approx 10^{-3}$ of horizon-sized wall segments collapse gravitationally into PBHs [45,46]. This collapse probability reflects the small fraction of Hubble volumes in which the converging wall geometry generates a local overdensity sufficient to exceed the gravitational collapse threshold; as discussed in Ref. [45] and quantified by the simulations of Ref. [47], the PBH mass is robust but the abundance itself carries large uncertainties (see Sec. VI B). The resulting PBH mass and formation fraction are

$$M_{\rm PBH} \simeq \frac{3.8 M_{\rm Pl}^3}{T_{\rm ann}^2} \propto \frac{\eta^3}{V_b^4}, \tag{14}$$

$$\beta_{\rm PBH} \equiv \left.\frac{\rho_{\rm PBH}}{\rho_{\rm tot}}\right|_{T_{\rm ann}} \simeq \frac{p V_b^4}{\frac{\pi^2}{30} g_* T_{\rm ann}^4}. \tag{15}$$

The present-day PBH dark matter fraction is obtained by redshifting $\beta_{\rm PBH}$ from $T_{\rm ann}$ to matter-radiation equality at $T_{\rm eq} = 0.75$ eV:

$$f_{\rm PBH} \simeq \beta_{\rm PBH} \frac{T_{\rm ann}}{0.84 T_{\rm eq}}, \tag{16}$$

where the factor 0.84 accounts for the change in $g_*$ between $T_{\rm ann}$ and $T_{\rm eq}$ [12]. Substituting Eqs. (13)–(15) into Eq. (16) and collecting powers, one finds

$$f_{\rm PBH} \propto \frac{\eta^{9/2}}{V_b^2}, \tag{17}$$

a pure power-law in the two fundamental parameters $(\eta, V_b)$ with no exponential factor. This is the central structural result for this construction: the abundance map is a monomial, and the BG fine-tuning measure is therefore constant across the entire parameter space.

The asteroid-mass window $10^{17}$–$10^{22}$ g requires, via Eq. (14),

$$T_{\rm ann} \in [10^5, 3 \times 10^7]\ {\rm GeV} \approx [10^{4.9}, 10^{7.5}]\ {\rm GeV}. \tag{18}$$

These boundaries follow directly from inverting $M_{\rm PBH} = 3.8 M_{\rm Pl}^3 / T_{\rm ann}^2$: at $T_{\rm ann} = 10^5$ GeV one obtains $M_{\rm PBH} \approx 10^{22}$ g (upper edge), and at $T_{\rm ann} = 3 \times 10^7$ GeV one obtains $M_{\rm PBH} \approx 10^{17}$ g (lower edge).

The benchmark $\eta = 2.4 \times 10^6$ GeV, $V_b = 2$ TeV is chosen as the simultaneous intersection of three conditions. First, $f_{\rm PBH} \approx 1$, so PBHs account for all of the dark matter. Second, $T_{\rm ann} \approx 5 \times 10^5$ GeV, placing $M_{\rm PBH} \approx 4 \times 10^{20}$ g well within the asteroid window. Third, $V_b = 2$ TeV is within 20% of the gravity-induced bias value $V_b^{\rm grav} = (\eta^5 / M_{\rm Pl})^{1/4} \approx 2.4$ TeV at this $\eta$, so the bias is naturally generated by Planck-suppressed operators of the form $\phi^5 / M_{\rm Pl}$ evaluated at $\phi = \eta$, i.e. $V_{\rm bias} \sim \eta^5 / M_{\rm Pl}$ [66,67], without introducing a new free scale.

$$V_b^{\rm grav}(\eta) \equiv (\eta^5 / M_{\rm Pl})^{1/4}. \tag{19}$$

These three conditions together identify $\eta \approx 2.4 \times 10^6$ GeV as a theoretically motivated "sweet spot": a single symmetry-breaking scale that simultaneously accounts for all of the dark matter, targets the asteroid-mass window, and requires no bias scale beyond what quantum gravity generates automatically. A particularly clean subclass is the "gravity-induced bias" scenario in which one sets $V_b = V_b^{\rm grav}$ exactly, reducing the construction to a single-parameter family parametrized by $\eta$ alone.

### 2. Analytic sensitivity structure

From Eq. (17) the BG sensitivities follow immediately by log-differentiation:

$$\Delta_\eta = \left|\frac{\partial \ln f_{\rm PBH}}{\partial \ln \eta}\right| = \frac{9}{2}, \qquad \Delta_{V_b} = \left|\frac{\partial \ln f_{\rm PBH}}{\partial \ln V_b}\right| = 2. \tag{20}$$

Both are exact constants, independent of the values of $\eta$ and $V_b$ and of the location in parameter space. The overall measure is

$$\Delta_{\rm DW} = \max(\Delta_\eta, \Delta_{V_b}) = \frac{9}{2} \quad \forall\ (\eta, V_b). \tag{21}$$

For the gravity-induced-bias subcase, substituting $V_b \propto \eta^{5/4} M_{\rm Pl}^{-1/4}$ into Eq. (17) gives $f_{\rm PBH} \propto M_{\rm Pl}^{1/2} \eta^2$, so the single-parameter measure reduces to $\Delta_\eta^{(\rm grav)} = 2$, the minimum value compatible with a nontrivial dependence on any dimensionful scale.

As shown below, this result places biased-domain-wall PBHs in the same naturalness tier as off-resonance heavy WIMPs and freeze-in particles (all with $\Delta \approx 2$), and substantially below the coannihilation and FOPT constructions discussed below. The quantitative answer to the title question is clear for this mechanism: biased-domain-wall PBHs are *not* fine-tuned by any reasonable standard.

I have checked explicitly that this result is stable under a change of fundamental variables, in the spirit of the check performed for the coannihilation benchmark in Sec. II B. Unlike the derived mass splitting $\delta = (M_{\rm NLSP} - M_\chi)/M_\chi$, which vanishes additively at the degeneracy threshold and produces a spuriously suppressed multiplicative log-derivative near that point, the domain-wall abundance map [Eq. (17)] is an exact monomial in $(\eta, V_b)$, and the BG sensitivity of a monomial is invariant under any

multiplicative (power-law) reparametrization of its arguments; no analogous additively vanishing derived coordinate appears in this construction. I regard this as a structural distinction between the two cases, not merely a definitional one. This invariance concerns only the parameters $(\eta, V_b)$ over which $\Delta$ is computed; it says nothing about the physical origin of $V_b$ itself, nor about the calibration coefficients $p \approx 10^{-3}$ and $c_t \approx 0.30$, which I return to in Sec. VI B.

#### 3. Parameter-space heatmap

Figure 1(a) shows $\log_{10}\Delta$ in the $(\log_{10}\eta, \log_{10} V_b)$ plane. The entire panel is uniformly pale yellow, the lowest color on the shared scale, reflecting the analytic result $\Delta = 9/2$ everywhere, a direct consequence of the pure power-law structure $f_{\rm PBH} \propto \eta^{9/2}/V_b^2$ [Eq. (17)]. The thick black contour marks $f_{\rm PBH} = 1$; gray dashed contours at $f_{\rm PBH} = 0.1$ and 10 define the $\pm 1$ dex sidebands that bound the natural island (green hatching). The two gold-orange lines show the asteroid-mass boundaries $M_{\rm PBH} = 10^{17}$ g and $10^{22}$ g; these are nearly parallel to the $f_{\rm PBH} = 1$ contour because both $f_{\rm PBH}$ and $M_{\rm PBH}$ are controlled by the same parameter combination $\eta^3/V_b^4$ [Eqs. (14) and (17)]. The magenta dash-dotted line shows the gravity-induced-bias locus $V_b = (\eta^5/M_{\rm Pl})^{1/4}$ [Eq. (19)], along which the bias is generated by Planck-suppressed operators rather than freely tuned. This locus passes through the natural island and through the benchmark point ($\eta = 2.4 \times 10^6$ GeV, $V_b = 2$ TeV; gold star), confirming that the Planck-suppressed scenario naturally populates the asteroid-mass window without any additional parameter adjustment.

### B. Early matter domination

#### 1. Formation mechanism and abundance map

A natural setting for an early matter-dominated era is provided by string-theoretic moduli: scalar fields $\phi$ with only gravitational-strength couplings to ordinary matter that are ubiquitous in compactifications of extra dimensions [69,70]. More broadly, any scalar field that begins to oscillate coherently in the early Universe when the Hubble rate drops below its mass, $H \sim m_\phi$, behaves as a pressureless fluid and temporarily dominates the energy density before decaying [70]. Such a field decays with a rate set by dimensional analysis for a gravitationally coupled particle, $\Gamma_\phi \sim c_\phi m_\phi^3/M_{\rm Pl}^2$, where $c_\phi \sim 10^{-3}$ encodes loop and phase-space suppression factors [69]. When $H$ drops to $\Gamma_\phi$, the field decays and reheats the Universe to a temperature [48,49,70]

$$T_{\rm rh} \simeq 0.55 \left(\frac{c_\phi m_\phi^3}{M_{\rm Pl}}\right)^{1/2}. \tag{22}$$

The period of matter domination preceding reheating alters cosmological perturbation growth in a way that is crucial for PBH production. In a radiation-dominated background, sub-Hubble density perturbations oscillate and decay; in a matter-dominated background, the same perturbations grow as $\delta\rho/\rho \propto a$ (where $a$ is the scale factor), and sufficiently large fluctuations can overcome pressure support and collapse gravitationally before reheating occurs [48,71]. The relevant perturbations are curvature perturbations inherited from inflation, characterized by their root-mean-square amplitude $\sigma$ on the scale that reenters the Hubble horizon during the matter-dominated era. The fraction of the energy density that collapses into PBHs depends sensitively on how many Hubble patches contain a perturbation large enough to collapse; this collapse criterion in a pressureless medium was derived analytically and verified numerically by Harada, Yoo, Kohri, and Nakama (HYKN) [72], who found

$$\beta_{\rm AM}(\sigma) \simeq 4.6 \times 10^{-6} \sigma^5 \exp\left(-\frac{0.147}{\sigma^{4/3}}\right). \tag{23}$$

Here $\beta_{\rm AM}$ is the mass fraction of the Universe that collapses into PBHs at the time of formation. The $\sigma^5$ prefactor reflects the five-dimensional phase-space volume for a fluctuation to exceed the collapse threshold in a pressureless fluid (compared to the exponential-only Press–Schechter form in radiation domination [73]), while the exponential encodes the probability tail of a Gaussian distribution above the threshold density contrast $\delta_c \simeq 0.05$ appropriate to matter domination [72]. The two contributions are equal at $\sigma \approx 0.050$, close to the natural-island benchmark, so neither can be neglected.

The PBH mass equals the Hubble-horizon mass at the time of collapse, which for this mechanism occurs just before reheating:

$$M_{\rm PBH} \simeq \frac{3.8 M_{\rm Pl}^3}{T_{\rm rh}^2} \propto m_\phi^{-3}, \tag{24}$$

where the proportionality uses Eq. (22). The asteroid-mass window $M_{\rm PBH} \in [10^{17}, 10^{22}]$ g therefore maps to a modulus-mass range $m_\phi \in [4 \times 10^{10}, 2 \times 10^{12}]$ GeV, as shown by the gold vertical lines in Fig. 1(b). The present-day dark matter fraction is obtained by redshifting $\beta_{\rm AM}$ from the formation epoch to matter-radiation equality at $T_{\rm eq} = 0.75$ eV:

$$f_{\rm PBH}(m_\phi, \sigma) \simeq \beta_{\rm AM}(\sigma) \frac{T_{\rm rh}(m_\phi)}{0.84 T_{\rm eq}}. \tag{25}$$

The factor 0.84 accounts for the change in relativistic degrees of freedom $g_*$ between $T_{\rm rh}$ and $T_{\rm eq}$ [12]. The dependence on $m_\phi$ enters solely through $T_{\rm rh}$ and is a pure power law, $f_{\rm PBH} \propto m_\phi^{3/2} \beta_{\rm AM}(\sigma)$, while the $\sigma$ dependence through Eq. (23) is nontrivial, containing both the polynomial prefactor and the exponential that will give rise to

the characteristic class II fine-tuning structure analyzed in Sec. III B 2.

### 2. Analytic sensitivity structure

Differentiating Eq. (25) with respect to $\ln m_\phi$ and $\ln\sigma$:

$$\Delta_{m_\phi} = \left|\frac{\partial \ln f_{\rm PBH}}{\partial \ln m_\phi}\right| = \frac{3}{2}, \tag{26}$$

$$\Delta_\sigma = \left|\frac{\partial \ln f_{\rm PBH}}{\partial \ln \sigma}\right| = 5 + \frac{0.196}{\sigma^{4/3}}. \tag{27}$$

The modulus-mass sensitivity $\Delta_{m_\phi} = 3/2$ is a constant, reflecting the simple power-law connection $T_{\rm rh} \propto m_\phi^{3/2}$. The curvature-variance sensitivity $\Delta_\sigma$ is dominated by the exponential factor in $\beta_{\rm AM}$ at small $\sigma$ and by the prefactor at large $\sigma$; the two contributions are equal at $\sigma^{4/3} = 0.196/5 \approx 0.039$, i.e. $\sigma \approx 0.050$. Since the natural island (defined by $f_{\rm PBH} \approx 1$ in the asteroid window) requires $\sigma \in [0.037, 0.071]$ at the benchmark $m_\phi = 9 \times 10^{10}$ GeV, the relevant range is

$$\Delta_\sigma^{(\rm island)} \in [12, 21], \tag{28}$$

with a benchmark value $\Delta_\sigma = 14.4$ at $\sigma = 0.055$.

The benchmark value $\sigma = 0.055$ is chosen because it sits near $\sigma \approx 0.050$ where the exponential and polynomial terms in Eq. (23) contribute equally ($0.147/\sigma^{4/3} \approx 5$), making it a physically representative point at which neither term dominates and the value $\Delta_\sigma = 14.4$ is typical of the natural island. It corresponds to a curvature power spectrum amplitude $\mathcal{P}_\mathcal{R}(k_{\rm form}) \approx \sigma^2 \approx 3 \times 10^{-3}$ on the asteroid-mass scale, a factor of $\sim 10^6$ above the CMB amplitude. Values $\sigma \lesssim 0.04$ exponentially suppress $\beta_{\rm AM}$ and require an unrealistically high $m_\phi$ to compensate; values $\sigma \gtrsim 0.10$ overproduce PBHs and push $m_\phi$ below the asteroid-mass lower boundary. The natural island in $\sigma$ is therefore bounded from below by the exponential tail of the HYKN formula and from above by the asteroid-mass constraint, precisely the two-sided confinement that makes $\Delta_\sigma \approx 14$ a robust result within this window.

The overall measure $\Delta_{\rm MD} = \max(3/2, \Delta_\sigma) = \Delta_\sigma$ across the entire natural island, placing this construction firmly in the intermediate "single-exponential" class despite the $\sigma^5$ prefactor moderating the growth rate.

The value $\sigma \approx 0.055$ required for $f_{\rm PBH} = 1$ in the asteroid-mass window corresponds to a primordial curvature power spectrum amplitude $\mathcal{P}_\mathcal{R} \approx \sigma^2 \approx 3 \times 10^{-3}$ at the formation scale $k_{\rm form}$, some seven orders of magnitude above the CMB amplitude $\mathcal{P}_\mathcal{R}^{\rm CMB} \approx 2 \times 10^{-9}$. Generating such an enhancement without disturbing CMB scales requires additional model ingredients—a spectator field, a feature in the inflationary potential, or resonant particle production [56,57,74]—whose model-building cost I fold into the tuning through the single effective parameter $\sigma$, rather than through the inflaton potential shape. This treatment is appropriate for the early matter-domination mechanism, where $\sigma$ is a free parameter of the cosmological initial conditions; for inflationary models, the additional cost of generating the required $\sigma$ from the inflaton potential is treated separately in Sec. III D as a second layer of tuning.

### 3. Parameter-space heatmap

Figure 1(b) shows $\log_{10}\Delta$ in the $(m_\phi, \sigma)$ plane, with the x-axis spanning the modulus masses relevant for asteroid-mass PBH production, $m_\phi \in [10^{10}, 10^{13}]$ GeV. In contrast to the domain-wall panel, a clear color gradient is visible: the panel darkens from yellow-green at large $\sigma$ (low tuning, $\Delta \approx 7$) toward teal at small $\sigma$ (higher tuning, $\Delta \approx 20$), reflecting the rising exponential sensitivity $\Delta_\sigma = 5 + 0.196/\sigma^{4/3}$ [Eq. (27)] as the Gaussian tail that controls collapse is pushed further into the rare-event regime. The $\Delta = 10$ and $\Delta = 30$ white contours bracket the natural island, confirming the analytic range $\Delta_\sigma \in [12, 21]$ of Eq. (28). The vertical gold lines mark the asteroid-mass boundaries $M_{\rm PBH} = 10^{17}$ g and $10^{22}$ g, which correspond via Eq. (24) to $m_\phi \in [4 \times 10^{10}, 2 \times 10^{12}]$ GeV; the constraint is purely horizontal because $M_{\rm PBH}$ depends on $m_\phi$ alone and not on $\sigma$. The natural island (green hatching) is a diagonal strip running along the $f_{\rm PBH} = 1$ contour. Its tilt reflects the interplay between the two parameters: increasing $\sigma$ raises $\beta_{\rm AM}$ through the HYKN formula, which would overproduce PBHs unless compensated by a larger $m_\phi$ (hence higher $T_{\rm rh}$ and stronger dilution from formation to equality). The benchmark $m_\phi = 9 \times 10^{10}$ GeV, $\sigma = 0.055$ (gold star) lies within the asteroid-mass natural island at $M_{\rm PBH} \approx 10^{21}$ g and $\Delta = 14.4$.

A guaranteed observational consequence of the enhanced curvature power spectrum $\mathcal{P}_\mathcal{R}(k_{\rm form}) \approx \sigma^2 \approx 3 \times 10^{-3}$ required for asteroid-mass PBH production is a stochastic gravitational-wave background (SGWB) sourced at second order in perturbation theory by the same scalar perturbations [75,76]. The peak frequency of this scalar-induced GW (SIGW) background is set by the formation temperature, $f_{\rm peak} \approx 3 \times 10^{-3}$ Hz$(T_{\rm form}/10^6$ GeV$)$, placing it in the mHz–Hz band for the asteroid-mass window, precisely the target range of LISA and DECIGO [77]. The SIGW amplitude scales as $\Omega_{\rm GW}h^2 \sim \mathcal{P}_\mathcal{R}^2 \sim 10^{-6} - 10^{-4}$ at the benchmark $\sigma = 0.055$, which is within the projected LISA sensitivity for transition temperatures $T_{\rm form} \gtrsim 10^5$ GeV [77]. Current bounds from Planck on the effective number of relativistic species, $\Delta N_{\rm eff} \lesssim 0.3$, translate into an integrated upper bound $\int \Omega_{\rm GW} d\ln f \lesssim 5 \times 10^{-6}$ that does not constrain the benchmark point but begins to disfavor $\sigma \gtrsim 0.15$ across the full asteroid-mass range [77]. The SIGW signal is therefore a robust and near-future testable

prediction of the early matter-domination PBH mechanism, and constitutes an independent observational channel complementary to microlensing searches.

### C. First-order phase transition

#### 1. Formation mechanism and abundance map

A first-order phase transition (FOPT) occurs when a field $\Phi$ in a hidden sector tunnels from a metastable false vacuum to the true vacuum of its potential, rather than rolling continuously as in a second-order transition. The transition proceeds by the stochastic nucleation of bubbles of the true vacuum within the false-vacuum background; these bubbles expand at close to the speed of light, collide, and eventually convert the entire Universe to the new phase [78,79]. I consider a FOPT occurring at a temperature $T_*$ in a hidden sector that is thermally coupled to the Standard Model plasma, so that the transition temperature is a physical parameter that sets the PBH mass through the Hubble horizon at that time.

Two dimensionless parameters characterize the transition completely for the purposes of PBH production. The transition strength $\alpha$ is the ratio of the latent heat released during nucleation to the ambient radiation energy density,

$$\alpha \equiv \frac{\Delta V - T_* \partial \Delta V/\partial T_*}{\rho_{\rm rad}(T_*)}, \tag{29}$$

where $\Delta V$ is the difference in free-energy density between the two phases [80]. A small $\alpha$ corresponds to a weak transition in which latent heat is a minor perturbation to the radiation bath; a large $\alpha$ corresponds to a strong, potentially runaway transition that can dramatically reheat the Universe. The transition rate parameter $\beta/H_*$ measures how rapidly the nucleation rate grows relative to the Hubble rate at $T_*$; equivalently, $H_*/\beta$ is the duration of the transition in Hubble units, with small $\beta/H_*$ corresponding to a slow, prolonged transition and large $\beta/H_*$ to an almost instantaneous one.

Regions of the false vacuum that are not reached by any bubble wall before the transition ends are compressed by the surrounding expanding walls. If the resulting overdensity exceeds the gravitational collapse threshold, these regions form PBHs [52–54,81–83]. The probability that a given Hubble-volume region remains in the false vacuum until the end of the transition is determined by the nucleation statistics; for an exponentially growing nucleation rate $\Gamma \propto e^{\beta t}$ (the standard approximation [77]), the fraction of the Universe that collapses into PBHs is

$$\beta_{\rm PBH} = \exp\left[-S(\alpha)\frac{\beta}{H_*}\right], \tag{30}$$

where $S(\alpha)$ encodes how efficiently the transition converts false-vacuum regions into PBHs as a function of its strength.[2] $S(\alpha)$ is model-dependent, but can be computed analytically for a given scalar potential; for the hidden-sector benchmark model of Ref. [54] the result is well approximated by the interpolating formula

$$S(\alpha) \simeq 5.5 - \frac{2\alpha}{\alpha + 0.3}, \tag{31}$$

which decreases monotonically from $S(0) = 5.5$ in the weak-transition limit (where the radiation bath is barely disturbed and the false-vacuum survival probability is relatively insensitive to $\alpha$) to $S(\infty) = 3.5$ in the ultrastrong limit (where the latent heat dominates and bubble walls run away [77,80]). The coefficients in Eq. (31) are specific to that model; different bubble-wall velocities, plasma couplings, or reheating dynamics would yield different values. As I show below in Eq. (37), however, the key fine-tuning identity $S(\alpha)\beta/H_*|_{f_{\rm PBH}=1} = \ln[T_*/(0.84T_{\rm eq})]$ follows from setting $f_{\rm PBH} = 1$ in Eq. (32) and holds for *any* monotone $S(\alpha)$; the fine-tuning result is therefore model-independent even though $S(\alpha)$ is not. The PBH mass is set by the Hubble horizon at $T_*$, $M_{\rm PBH} \simeq 3.8 M_{\rm Pl}^3/T_*^2$, and the present-day dark matter fraction is

$$f_{\rm PBH}(T_*, \beta/H_*, \alpha) = \exp\left[-S(\alpha)\frac{\beta}{H_*}\right] \cdot \frac{T_*}{0.84 T_{\rm eq}}. \tag{32}$$

The parameter space of interest for this mechanism is defined entirely by the requirement $f_{\rm PBH} \approx 1$ within the asteroid-mass window, which fixes $\beta/H_* \sim 5$–$15$ at the benchmark temperatures $T_* \sim 10^5$–$10^7$ GeV. It is a nontrivial consequence that this same range of $\beta/H_*$ lies within the sensitivity reach of LISA and next-generation ground-based gravitational-wave detectors for the stochastic background produced by the transition itself [77,85]: FOPT-PBH dark matter is therefore simultaneously a target of gravitational-wave observatories and microlensing surveys, not because the two observations are by construction

[2]Equation (30) is the standard phenomenological approximation in which the false-vacuum survival fraction is written as a single exponential in the nucleation-rate parameter $\beta/H_*$, treating $\beta/H_*$ as a fundamental input. A more complete analysis of the collapse probability in the supercooled limit, tracking the full past-light-cone percolation integral, shows that $\mathcal{P}_{\rm coll}$ is super-exponential in $\beta/H$ [46]: the BG derivative evaluated from that formula at the benchmark $\beta/H_* \approx 8$ gives $\Delta_{\beta/H} \approx 265$ rather than $\approx$35, placing FOPT-PBH in class III even at the phenomenological level [Sec. VI B, Eq. (72)]. At a deeper level, tracing $f_{\rm PBH}$ back to the microphysical scalar potential through the Euclidean action $S_3(T_p)/T_p$ yields a double-exponential ("superexponential") structure, $f_{\rm PBH} \simeq M \exp(-Q \exp(-S_3(T_p)/T_p))$, with $Q \sim 10^{77}$ [84], giving $\Delta \sim 10^3$–$10^4$. The fine-tuning analysis in Eqs. (33)–(37) uses the single-exponential approximation and therefore gives a lower bound on the true BG measure; see Sec. VI B for the full assessment.

linked, but because both probe the same underlying parameter $\beta/H_*$.

#### 2. Analytic sensitivity structure

Differentiating Eq. (32):

$$\Delta_{T_*} = \left| \frac{\partial \ln f_{\rm PBH}}{\partial \ln T_*} \right| = 1, \tag{33}$$

$$\Delta_{\beta/H} = \left| \frac{\partial \ln f_{\rm PBH}}{\partial \ln(\beta/H_*)} \right| = S(\alpha)\frac{\beta}{H_*}, \tag{34}$$

$$\Delta_{\alpha} = \left| \frac{\partial \ln f_{\rm PBH}}{\partial \ln \alpha} \right| = \frac{\beta}{H_*}\cdot\frac{0.6\alpha}{(\alpha+0.3)^2}. \tag{35}$$

The temperature sensitivity $\Delta_{T_*} = 1$ is a constant, since $f_{\rm PBH} \propto T_*$ through the dilution factor alone. The dominant contribution is $\Delta_{\beta/H}$, which grows linearly with $\beta/H_*$.

The key result emerges upon imposing the condition $f_{\rm PBH} = 1$, that is, solving Eq. (32) for $\beta/H_*$ at fixed $T_*$ and $\alpha$:

$$S(\alpha)\frac{\beta}{H_*}\bigg|_{f_{\rm PBH}=1} = \ln\left[\frac{T_*}{0.84T_{\rm eq}}\right]. \tag{36}$$

Substituting Eq. (36) into Eq. (34),

$$\Delta_{\beta/H}|_{f_{\rm PBH}=1} = \ln\left[\frac{T_*}{0.84T_{\rm eq}}\right] \approx 30-37 \tag{37}$$

for $T_* \in [10^4, 10^{6.5}]$ GeV (the asteroid-mass formation window). This result is exact for any $\alpha$ and any $T_*$ in the relevant range, and is independent of the specific value of $S(\alpha)$: the fine-tuning of FOPT-PBH dark matter is determined entirely by the logarithm of the dilution ratio from formation to equality, and is insensitive to the strength or dynamics of the phase transition. The nucleation-rate parameter $\beta/H_*$ must be tuned to ∼3% precision to land on the $f_{\rm PBH} = 1$ contour.

Equation (37) is the gravitational-sector analog of the thermal-freeze-out tuning $\Delta \approx x_F \simeq 25$ for standard WIMPs (see Sec. IV A 3); both express the cost of exponentially suppressing a natural formation rate to match the observed dark matter density. This structural identity is the central observation of Sec. V.

#### 3. Parameter-space heatmap

Figures 1(c) and 1(d) show $\log_{10}\Delta$ in the $(T_*, \beta/H_*)$ plane computed from the GV superexponential collapse probability Eq. (70), for collapse threshold $\delta_c = 0.05$ [panel (c)] and $\delta_c = 0.10$ [panel (d)]. The GV formula contains no $\alpha$ dependence; the two panels illustrate instead the sensitivity to $\delta_c$.

Both panels display a sharply stratified color structure. Below the $f_{\rm PBH} = 1$ contour (small $\beta/H_*$), the color is yellow-green, corresponding to $\Delta \sim 1$–10 (class I); $f_{\rm PBH}$ is large here (PBH overproduction) and fine-tuning in $\beta/H_*$ is mild. At and above the contour, the color transitions steeply to teal and then dark purple: $\Delta$ jumps from ∼100 to ∼265 within a window $\Delta(\beta/H_*) \approx 0.1$ around the contour. The steep transition reflects the superexponential sensitivity of $P_{\rm coll}$ to $\beta/H_*$.

The $f_{\rm PBH} = 1$ contour (thick black) runs nearly horizontally, at $\log_{10}(\beta/H_*) \approx 0.70$ [panel (c), $\delta_c = 0.05$] and ≈0.65 [panel (d), $\delta_c = 0.10$], spanning the full asteroid-mass temperature range $\log_{10}(T_*/\text{GeV}) \in [4.7, 7.6]$. The contour's horizontal position shifts by $\Delta\log_{10}(\beta/H_*) \approx 0.05$ between panels as $\delta_c$ changes from 0.05 to 0.10; a higher collapse threshold requires fewer surviving false-vacuum patches, achieved at slightly lower $\beta/H_*$.

The natural island (green hatching, $0.1 \le f_{\rm PBH} \le 10$ within the asteroid-mass window) is an extremely narrow, nearly horizontal strip of total width ≈0.09 in $\beta/H_*$, corresponding to an island half-width $\epsilon \approx 0.26\%$ [Eq. (6)], consistent with $\Delta \approx 265$ in both panels. The narrowness of the hatched strip, compared with the broader islands visible in Figs. 1(a) and 1(b), is itself a visual demonstration of the class III fine-tuning severity.

The benchmark $T_* = 10^6$ GeV is marked by gold stars in both panels; the required nucleation rate is $\beta/H_* = 5.0$ [panel (c)] and 4.5 [panel (d)], and both stars lie precisely on the $f_{\rm PBH} = 1$ contour at $\Delta \approx 265$ [Eq. (72)]. The earlier structural identity $\Delta_{\beta/H}|_{f_{\rm PBH}=1} = \ln[T_*/(0.84T_{\rm eq})] \approx 35$ [Eq. (37)] was an exact result of the single-exponential formula and has no direct analog in the GV formula; no white $\Delta = 30$ contour coincides with the $f_{\rm PBH} = 1$ contour in these panels.

### D. Inflationary collapse: A two-layer structure

The three mechanisms analyzed above (biased domain walls, early matter domination, and first-order phase transitions) do not exhaust the list of PBH formation channels. The most extensively studied mechanism in the literature is the collapse of large-amplitude curvature perturbations generated during single-field inflation. In the standard picture, inflation produces a nearly scale-invariant power spectrum of curvature perturbations $\mathcal{P}_\mathcal{R}(k)$ whose amplitude on CMB scales is $\mathcal{P}_\mathcal{R} \approx 2\times10^{-9}$. PBH formation requires $\mathcal{P}_\mathcal{R} \sim 10^{-2} - 10^{-3}$ on sub-CMB scales corresponding to the asteroid-mass Hubble horizon, an enhancement of $\mathcal{O}(10^6$–$10^7)$ relative to the CMB amplitude [56,57]. Such an enhancement is typically achieved through a feature in the single-field inflaton potential, most commonly an inflection point or a phase of ultraslow-roll (USR), during which the inflaton velocity is briefly reduced to near zero, allowing quantum fluctuations to grow to large amplitude on the scales that cross the Hubble horizon at that moment [19,57]. The curvature perturbation variance $\sigma$

that enters the collapse formula [Eq. (23)] is related to the power spectrum by $\sigma^2 \approx \int_{k_{\rm form}} \mathcal{P}_{\mathcal{R}}(k) W^2(k) {\rm d}\ln k$, where $W(k)$ is a window function on the formation scale $k_{\rm form}$; for a narrow peak at $k_{\rm form}$, $\sigma \approx \mathcal{P}_{\mathcal{R}}^{1/2}(k_{\rm form})$, though more generally $\sigma$ depends on the full spectral shape through the window-function integral. Since the present analysis treats $\sigma$ as a free effective parameter, the precise form of this relation does not affect the fine-tuning conclusions; the model-building cost of generating the required large $\sigma$ from the inflaton potential is instead captured by the layer 2 analysis below.

I do not analyze inflationary PBH production in the same detail as the previous three mechanisms, partly because its fine-tuning has already received dedicated attention in the literature [19–22], and partly for the methodological reason explained in the opening of this section. Instead, building on the fine-tuning audit of Cole, Gow, Byrnes, and Patil [32], who established for single-field models the chain-rule relation between the sensitivity of $f_{\rm PBH}$ and that of the curvature power-spectrum amplitude $\sigma$ (discussed further below), I organize the inflationary fine-tuning into the same two layers: a first layer capturing the sensitivity of the PBH abundance to the curvature power spectrum amplitude $\sigma$, whose value depends critically on the reheating history, and a second layer capturing the additional sensitivity of $\sigma$ to the inflaton potential coefficients. I extend this decomposition systematically across the six inflationary model classes considered here and two reheating scenarios, clarifying why their fine-tuning spans more than seven orders of magnitude.

The inflationary mechanism involves two successive parameter-to-observable maps, each contributing independently to the total fine-tuning.

*Layer 1:* $\sigma \to f_{\rm PBH}$. The form of the collapse fraction $\beta(\sigma)$ depends on the equation of state at the epoch when the enhanced perturbations reenter the Hubble horizon, and this distinction is physically important for the fine-tuning estimate.

For the early matter-domination mechanism of Sec. III B, collapse occurs during a pressureless era and is described by the HYKN formula (23) with threshold $\delta_c \approx 0.05$, giving $\Delta_\sigma^{(\rm MD)} \approx 14$ [Eq. (27)]. I refer to this as the layer 1 floor for the MD case: it sets a minimum fine-tuning cost independent of whatever mechanism generates the required $\sigma$.

For inflationary PBH models the background at reentry depends critically on the reheating temperature $T_{\rm reh}$. Two limiting cases arise.

*Low reheating* ($T_{\rm reh} \lesssim T_{\rm form}$). If the inflaton decay rate is small, the inflaton oscillates coherently after inflation ends and the universe is matter-dominated at the epoch when the formation-scale perturbations reenter the horizon [49,70]. In this case the collapse occurs in an MD background and the HYKN formula with $\delta_c \approx 0.05$ applies, giving a layer 1 floor identical to the noninflationary early MD mechanism:

$$\Delta_\sigma^{(\rm MD)} \approx 14, \tag{38}$$

identical to Eq. (27). For asteroid-mass PBHs with $T_{\rm form} \sim 10^4$–$10^7$ GeV, this requires $T_{\rm reh} \lesssim 10^7$ GeV, achievable in models with a weakly coupled inflaton sector. Note that while the layer 1 floor is the same as in the noninflationary early MD case, the two scenarios are physically distinct: here $\sigma$ is not a free initial condition but is generated by the inflaton potential, and the additional cost of achieving the required $\sigma$ is captured by layer 2.

*High reheating* ($T_{\rm reh} \gg T_{\rm form}$). In the standard scenario where reheating completes well before the formation scale reenters the horizon, the background is radiation-dominated at reentry [11]. The collapse fraction follows the Press-Schechter form with $\delta_c \approx 0.45$–$0.55$ [21]:

$$\beta_{\rm RD}(\sigma) = \frac{1}{2}\,{\rm erfc}\left(\frac{\delta_c}{\sqrt{2}\sigma}\right) \approx \frac{\sigma}{\delta_c\sqrt{2\pi}} \exp\left(-\frac{\delta_c^2}{2\sigma^2}\right), \tag{39}$$

where the approximation holds for $\delta_c/(\sqrt{2}\sigma) \gg 1$. Log-differentiating Eq. (39):

$$\Delta_\sigma^{(\rm RD)} = \left|\frac{\partial \ln \beta_{\rm RD}}{\partial \ln \sigma}\right| = 1 + \frac{\delta_c^2}{\sigma^2}. \tag{40}$$

For asteroid-mass PBH production one requires $\beta_{\rm RD} \sim 10^{-9} - 10^{-15}$, giving $\delta_c^2/(2\sigma^2) \approx 20$–$35$, so $\sigma \approx 0.05$–$0.07$ for $\delta_c = 0.45$–$0.55$. Evaluating Eq. (40):

$$\Delta_\sigma^{(\rm RD)} \approx 57 - 70 \qquad (\sigma \approx 0.06, \delta_c \approx 0.45 - 0.55), \tag{41}$$

approximately five times larger than the matter-dominated value of Eq. (38), reflecting the much higher collapse threshold in radiation domination.

The two layer 1 floors bracket the layer 1 contribution to the fine-tuning of inflationary models:

$$\Delta_\sigma^{(\rm inf)} \in [\Delta_\sigma^{(\rm MD)}, \Delta_\sigma^{(\rm RD)}] \approx [14, 57 - 70], \tag{42}$$

depending entirely on the reheating history. Figure 2 displays $\Delta_{\rm tot}$ for both cases side by side, with the same $x$-axis scale to make the factor-of-five shift in the layer 1 floor directly visible; note that for model classes with large layer 2 sensitivity the total tuning is dominated by layer 2 and the factor-of-five shift between panels becomes a small fractional effect. The enhanced power spectrum required for inflationary PBH production carries the same unavoidable scalar-induced gravitational-wave (SIGW) background prediction discussed in Sec. III B, with the additional feature that the peak frequency and amplitude directly encode the inflaton potential shape [56,75,76], making it a particularly powerful discriminant between inflationary PBH models.

*Layer 2:* $\{c_k\} \to \sigma$. The additional tuning specific to inflationary models enters through the mapping from the inflaton potential coefficients $\{c_k\}$—e.g., the height of the inflection point, its slope, and the field value at which USR begins—to the power spectrum amplitude $\sigma$ on the formation scale. Near an inflection point, the inflaton velocity $\dot{\phi}$ passes through a local minimum whose depth is controlled by the local slope $c_{\rm slope}$ of the potential. The power spectrum amplitude depends on $\dot{\phi}^{-2}$ at horizon crossing, so it is exponentially sensitive to $c_{\rm slope}$:

$$\sigma \propto \exp\left(-\frac{\text{const}}{c_{\rm slope}}\right). \tag{43}$$

A fractional change $\delta c_{\rm slope}/c_{\rm slope}$ therefore produces a fractional change $\delta\sigma/\sigma \propto \delta c_{\rm slope}/c_{\rm slope}^2$, an exponentially amplified response. Hertzberg and Yamada [19] quantified this explicitly for polynomial inflation potentials: a change in the potential coefficients by one part in $10^2$–$10^8$ (depending on the specific model) shifts the power spectrum peak amplitude by an order-one factor, implying $\Delta_{\sigma/c_k} \sim 10^2$–$10^8$. A systematic "fine-tuning audit" of single-field PBH models by Cole, Gow, Byrnes, and Patil [32] reached the same conclusion across three representative classes of single-field potential (superposed-feature, polynomial, and nonpolynomial), finding that the potential coefficients must be tuned to between one part in $\sim 10^2$ and one part in $\sim 10^8$ to hold the power-spectrum peak amplitude fixed, in direct agreement with the layer 2 range adopted here. Notably, Ref. [32] also separates this from the additional sensitivity of the PBH abundance to the power-spectrum amplitude, reporting that the latter contributes a further factor of $\rho \approx 22$–$38$ [for radiation-dominated collapse, with a value $\rho \approx \delta_c^2/2\sigma^2$ that maps onto the layer 1 floor $\Delta_\sigma^{(\rm RD)} \approx 57$–$70$ of Eq. (41), reduced to $\rho \approx 5/4$ in an early matter-dominated era]. Their two-stage decomposition of the tuning is precisely the layer 2 × layer 1 structure formalized in this paper, and provides an independent cross-check of both factors. The total BG measure for the potential coefficients is therefore the product of the two layers,

$$\Delta_{c_k} \sim \Delta_\sigma \times \Delta_{\sigma/c_k} \sim (14-70)\times(10^2-10^8), \tag{44}$$

placing single-field inflationary collapse in class III or beyond (as defined in Sec. V B) for essentially all realistic potentials [20,21], with the layer 1 factor ranging from 14 (low-reheating MD collapse) to 57–70 (high-reheating RD collapse) as discussed above. This double-exponential structure has no analog in the three mechanisms analyzed above, where the fundamental parameters ($m_\phi$, $V_b$, $T_*$, $\beta/H_*$) map onto $f_{\rm PBH}$ through at most one exponential, and not through any hidden exponential in the parameter-to-$\sigma$ map.

Note that the magnitude of layer 2 varies enormously across inflationary model classes, as summarized in Fig. 2 and reviewed in detail in the next section. For single-field USR and inflection-point potentials $\Delta_{\sigma/c_k} \sim 10^2$–$10^8$, placing the total fine-tuning firmly in class III or beyond (Sec. V B) regardless of the reheating history. For multi-field and spectator-field models $\Delta_{\sigma/c_k} \sim 2$–$10$; the total tuning then depends on whether collapse occurs in RD ($\Delta_{\rm tot} \approx 115$–$700$, class III for standard high reheating) or MD ($\Delta_{\rm tot} \approx 21$–$140$, class II to lower class III for low reheating with $T_{\rm reh} \lesssim T_{\rm form}$). The severe fine-tuning of single-field USR inflationary PBH production is a specific pathology of models that rely on an exponentially sensitive slow-roll feature to amplify the power spectrum; the reheating history then sets the overall tuning scale for all other inflationary model classes.

The recent analysis of Iovino and Riotto [22] reaches a more optimistic conclusion by applying Wilson's naturalness criterion to the potential coefficients rather than to $\sigma$. Their measure assigns a value close to unity even for fine-tuned inflaton potentials because it is constructed to be insensitive to the overall normalization of the abundance: a potential that predicts $f_{\rm PBH} = 0.001$ and one that predicts $f_{\rm PBH} = 1$ are treated as equally natural. The BG measure is more conservative in this respect: it registers the full sensitivity of $f_{\rm PBH}$ to each parameter along the observable contour $f_{\rm PBH} = 1$, and therefore captures the double-exponential cost of maintaining that condition via an inflection-point potential. Both approaches are internally consistent, and answer complementary questions. The BG question is: "how precisely must the potential coefficient $c_k$ be specified to keep $f_{\rm PBH} \approx 1$?" The Wilson question is: "does a technically natural UV completion exist in which $c_k$ takes its current value?" For the cross-paradigm comparison of this paper, the BG question is the relevant one, since it is the same question being asked of all other dark matter mechanisms in the comparison.

### 1. Fine-tuning across inflationary model classes

The analysis above shows that the double-exponential structure of layer 2 is a specific property of single-field inflection-point or USR models, not a universal feature of inflationary PBH production. Figure 2 quantifies layer 2 [$\Delta_{\sigma/c_k}$, panel (a)] and the total fine-tuning $\Delta_{\rm tot}$ for two choices of layer 1 floor [panels (b) and (c)]: the matter-dominated HYKN floor $\Delta_\sigma^{(\rm MD)} \approx 14$ that applies when $T_{\rm reh} \lesssim T_{\rm form}$ (low reheating, inflaton-oscillation era at reentry), and the radiation-dominated Press–Schechter floor $\Delta_\sigma^{(\rm RD)} \approx 57$–$70$ that applies when $T_{\rm reh} \gg T_{\rm form}$ (standard high reheating). The $x$-axes of panels (b) and (c) are identical to facilitate direct comparison; the pessimistic USR row ($\Delta_{\rm tot} \gtrsim 10^9$–$10^{10}$) is off the shared scale and marked by an arrow. For reference, the three noninflationary PBH benchmarks are shown as vertical blue lines in both total-tuning panels at $\Delta = 2$, 14, and 35.

I now discuss each model class in turn.

(a) *Curvaton and spectator-field models* ($\Delta_{\sigma/c_k} \approx 2$–$3$; $\Delta_{\rm tot} \approx 115$–$210$ *[RD] or* $21$–$42$ *[MD]*). In the curvaton mechanism, a light spectator field $\chi$ with mass $m_\chi \ll H_{\rm inf}$ acquires perturbations $\delta\chi \simeq H_{\rm inf}/(2\pi)$ during inflation and later converts them into curvature perturbations when it decays [86,87]. PBH production requires an enhanced curvature perturbation on sub-CMB scales, which can be arranged if $\chi$ rolls along a self-interacting potential that generates a non-Gaussian or scale-dependent signal on those scales [33,34]. The key structural property is that the curvature power spectrum scales as $\mathcal{P}_\mathcal{R} \propto (H_{\rm inf}/m_\chi)^{2n}$ for some integer $n \geq 1$ depending on the interaction structure. The BG sensitivity of $\sigma$ to the curvaton mass is therefore $\Delta_{\sigma/m_\chi} = n \approx 1$–$2$, and to the self-coupling $\Delta_{\sigma/\lambda} \approx 1$. The combined layer 2 is $\Delta_{\sigma/c_k} \approx 2$–$3$. Multiplied by the RD layer 1 floor $\Delta_\sigma^{(\rm RD)} \approx 57$–$70$ [Eq. (41)], the total fine-tuning is $\Delta_{\rm tot} \approx 115$–$210$ (class III, Sec. V B) for standard high-reheating scenarios. For low-reheating scenarios where collapse occurs during the inflaton-oscillation matter-dominated era, the MD layer 1 floor applies and the total reduces to $\Delta_{\rm tot} \approx 21$–$42$ (class II, Sec. V B), comparable to the early matter-domination mechanism of Sec. III B. The reheating temperature is therefore a critical parameter that determines whether curvaton PBH models belong to class II or class III.

(b) *Multifield stochastic inflation* ($\Delta_{\sigma/c_k} \approx 2$–$10$; $\Delta_{\rm tot} \approx 115$–$700$ *[RD] or* $28$–$140$ *[MD]*). When the inflaton traverses an approximately flat direction, quantum diffusion can compete with or dominate the classical drift, and the statistics of perturbations must be described by a Fokker-Planck equation rather than the standard linear perturbation theory [88]. In multifield models where a secondary field undergoes stochastic diffusion, the power spectrum amplitude on a given scale is controlled by the ratio $H_{\rm inf}/M_{\rm flat}$, where $M_{\rm flat}$ is the effective mass of the flat direction [35,89]. Since $\mathcal{P}_\mathcal{R} \propto (H_{\rm inf}/M_{\rm flat})^4$, the layer 2 sensitivity is $\Delta_{\sigma/M_{\rm flat}} \approx 2$ per parameter, giving a combined $\Delta_{\sigma/c_k} \approx 2$–$10$ when multiple parameters of the flat-direction potential contribute. Multiplied by the RD layer 1 floor $\Delta_\sigma^{(\rm RD)} \approx 57$–$70$ [Eq. (41)], the total fine-tuning range is $\Delta_{\rm tot} \approx 115$–$700$ (class III, Sec. V B) for high-reheating scenarios; with the MD floor the range reduces to $\Delta_{\rm tot} \approx 28$–$140$, straddling the class II/III boundary. The upper end of this range arises when the stochastic enhancement requires the flat direction to remain nearly degenerate with the inflationary trajectory for a precise number of e-folds [35], introducing additional sensitivity that stretches the layer 2 estimate toward its upper bound.

(c) *Hybrid inflation with waterfall transition* ($\Delta_{\sigma/c_k} \approx 3$–$20$; $\Delta_{\rm tot} \approx 170$–$1400$ *[RD] or* $42$–$280$ *[MD]*). In hybrid inflation, the inflaton $\phi$ rolls toward a critical value $\phi_c$ at which a "waterfall" field $\psi$ becomes tachyonic and triggers the end of inflation [90]. Near $\phi_c$, the perturbations of $\psi$ are dramatically amplified because its effective mass $m_\psi^2 \approx g^2(\phi - \phi_c)M^{-2} - \mu^2$ approaches zero, and regions where the transition is delayed can collapse into PBHs [36]. The curvature power spectrum amplitude scales as a power law in the coupling $g$ and the mass parameter $\mu$: $\mathcal{P}_\mathcal{R} \propto g^{-2n}\mu^{2m}$ for integers $n$, $m$ depending on the potential shape, giving $\Delta_{\sigma/g} \approx n$ and $\Delta_{\sigma/\mu} \approx m$. For a mild waterfall ($|\phi - \phi_c|$ not exponentially small), $n + m \approx 3$–$5$ and the layer 2 tuning is modest. However, as $\phi \to \phi_c$ more closely, the enhancement duration becomes exponentially sensitive to the proximity to the critical point, stretching $\Delta_{\sigma/c_k}$ toward 10–20 at the upper end. Multiplied by the RD layer 1 floor, the resulting total tuning is $\Delta_{\rm tot} \approx 170$–$1400$ (class III, Sec. V B) for high-reheating scenarios; with the MD floor the range is $\Delta_{\rm tot} \approx 42$–$280$, in the lower part of class III.

(d) *Single-field with step or resonance feature* ($\Delta_{\sigma/c_k} \approx 10$–$100$; $\Delta_{\rm tot} \approx 570$–$7000$ *[RD] or* $140$–$1400$ *[MD]*). A sharp feature such as a step, a bump, or a periodic modulation in the single-field inflaton potential can resonantly amplify perturbations as modes cross the feature [37,91]. For a step of height $h$ at field value $\phi_s$, the peak power spectrum scales roughly as $\mathcal{P}_\mathcal{R}^{\rm peak} \propto h^2$, a polynomial relationship giving $\Delta_{\sigma/h} \approx 1$. The dominant source of layer 2 tuning is therefore not the feature amplitude but its position: the resonance condition requires that precisely the modes corresponding to the asteroid-mass scale cross the horizon when the inflaton is at the step location. A small change $\delta\phi_s$ in the step position shifts the resonant scale by $\delta k/k \sim \delta\phi_s/(\phi_s - \phi_{\rm CMB})$, changing which physical scales are enhanced. In terms of the inflaton potential coefficients, $\Delta_{\sigma/\phi_s} \sim 10$–$50$ because $\phi_s$ must be tuned at the 2%–10% level to target the asteroid-mass window specifically. The combined layer 2 is $\Delta_{\sigma/c_k} \approx 10$–$100$; multiplied by the RD layer 1 floor this gives $\Delta_{\rm tot} \approx 570$–$7000$ (class III, Sec. V B) for high-reheating scenarios, or $\Delta_{\rm tot} \approx 140$–$1400$ with the MD floor. A key advantage of this class of models is that the prediction is falsifiable: the power spectrum feature also produces a distinctive oscillatory signature in the SIGW background [91] that is detectable by LISA for formation temperatures in the asteroid-mass range.

(e) *Single-field USR / inflection point, optimistic* ($\Delta_{\sigma/c_k} \approx 10^2$–$10^3$; $\Delta_{\rm tot} \approx 6000$–$7 \times 10^4$ *[RD] or* $1400 - 10^4$ *[MD]*). For a single inflaton field with an inflection point, the slow-roll velocity passes through a local minimum controlled by the local slope $c_{\rm slope}$ of the potential. In the USR regime, quantum

perturbations grow as $\delta\phi \propto e^{3Ht}$ rather than remaining constant, amplifying the power spectrum by a factor $e^{6\Delta N_{\rm USR}}$ over $\Delta N_{\rm USR}$ e-folds of USR evolution [38]. The number of USR e-folds required to achieve $\mathcal{P}_\mathcal{R} \sim 10^{-2}$ from a CMB-scale amplitude $\mathcal{P}_\mathcal{R}^{\rm CMB} \sim 2 \times 10^{-9}$ is $\Delta N_{\rm USR} \approx 2.7$, which in turn requires the potential slope to satisfy $c_{\rm slope} \approx {\rm const}/\Delta N_{\rm USR}$. In the optimistic case, corresponding to potentials where the inflection point is only mildly fine-tuned and the coefficients are of order unity, $\Delta_{\sigma/c_{\rm slope}} \approx {\rm const}/c_{\rm slope} \sim 10^2$–$10^3$ (the Hertzberg-Yamada lower-end result [19]). The systematic polynomial-potential analysis of Ballesteros *et al.* [92], performed independently of and consistent with Hertzberg and Yamada [19], confirms that the required tuning of the potential coefficients depends sensitively on the polynomial order and the number of terms retained. Multiplied by the RD layer 1 floor, the total tuning is $\Delta_{\rm tot} \approx 6000$–$7 \times 10^4$ (high-reheating) or $\Delta_{\rm tot} \approx 1400$–$10^4$ (low-reheating), in either case deep in class III (Sec. V B); a change of order 0.1% in the inflaton potential slope is sufficient to miss the target abundance by a factor of order unity.

(f) *Single-field USR / inflection point, pessimistic ($\Delta_{\sigma/c_k} \approx 10^3$–$10^8$; $\Delta_{\rm tot} \gtrsim 10^{10}$ [RD] or $\gtrsim 10^9$ [MD]).* For a generic polynomial inflaton potential, the relation between the potential coefficients at the inflection point and the resulting power spectrum amplitude is far more sensitive than the optimistic case suggests. A detailed analysis of polynomial potentials of the form $V(\phi) = V_0[1 + a_2(\phi/M)^2 + a_3(\phi/M)^3 + \ldots]$ by Hertzberg and Yamada [19] and by Kannike *et al.* [39] shows that the required amplitude $\mathcal{P}_\mathcal{R} \sim 10^{-2}$ at $k_{\rm form}$ is consistent with the CMB constraint $\mathcal{P}_\mathcal{R} \sim 2 \times 10^{-9}$ at $k_{\rm CMB}$ only if the potential coefficients satisfy a set of near-cancellation conditions. In the generic case, achieving both constraints simultaneously requires tuning the coefficients to within one part in $10^3$–$10^8$, giving $\Delta_{\sigma/c_k} \sim 10^3$–$10^8$. Multiplied by the RD layer 1 floor, the resulting total fine-tuning reaches $\Delta_{\rm tot} \gtrsim 10^{10}$ (high-reheating) or $\Delta_{\rm tot} \gtrsim 10^9$ (low-reheating), far exceeding any other scenario considered in this paper in either case. This pessimistic range represents the majority of the parameter space for generic polynomial inflaton potentials; the optimistic range is accessible only in models with additional symmetry or structure that correlates the potential coefficients in a way that reduces the cancellation requirement.

The overall message of Fig. 2 is stark: the fine-tuning of inflationary PBH production spans more than seven orders of magnitude depending on the model class, in either reheating scenario. For low reheating ($T_{\rm reh} \lesssim T_{\rm form}$, MD layer 1 floor), curvaton and spectator-field models ($\Delta_{\rm tot} \approx 21$–$42$) fall in class II, comparable to the noninflationary early MD mechanism and to coannihilating-WIMP dark matter. For high reheating ($T_{\rm reh} \gg T_{\rm form}$, RD layer 1 floor), every inflationary class falls in class III or beyond; the most natural class (curvaton/spectator, $\Delta_{\rm tot} \approx 115$–$210$) is roughly five times more tuned than the corresponding low-reheating estimate, reflecting the factor-of-five ratio of the two layer 1 floors—a shift that is clearly visible for model classes with small layer 2 sensitivity, but becomes a minor correction when layer 2 dominates. In both cases the layer 2 sensitivity $\Delta_{\sigma/c_k}$ spans five orders of magnitude from curvaton models to single-field USR potentials, so the mechanism for enhancing the power spectrum remains the most physically informative quantity regardless of the reheating history.

## IV. PARTICLE DARK MATTER BENCHMARKS

I now apply the same BG protocol to seven particle dark matter benchmarks spanning four production paradigms: thermal-relic WIMPs (B1–B3), freeze-in (B4), asymmetric dark matter (B5), and axions (B6, B7). The purpose is not to provide a comprehensive survey of the particle dark matter landscape, but to place specific, representative benchmarks into the same quantitative framework used for the PBH constructions in Sec. III. The benchmarks are chosen to span the full range of BG fine-tuning observed in the particle literature and to correspond to physically well-motivated, phenomenologically viable scenarios.

### A. Thermal-relic WIMPs

Thermal relics freeze out when their annihilation rate falls below the Hubble rate at temperature $T_{\rm fo} = M_\chi/x_F$ with $x_F \approx 25$ [1,93]. The relic abundance is then determined by the Lee-Weinberg relation

$$\Omega_\chi h^2 \simeq \frac{1.07 \times 10^9\ {\rm GeV}^{-1} x_F}{g_*^{1/2} M_{\rm Pl} \langle\sigma v\rangle_{\rm fo}}, \tag{45}$$

where $g_* \approx 86$ counts the effective relativistic degrees of freedom at freeze-out and $\langle\sigma v\rangle_{\rm fo}$ is the thermally averaged annihilation cross section evaluated at $T_{\rm fo}$. The "WIMP miracle" is the observation that for $M_\chi \sim 100$ GeV and gauge-coupling-strength $\langle\sigma v\rangle \sim \alpha_W^2/M_W^2 \sim 6.7 \times 10^{-9}$ GeV$^{-2}$, Eq. (45) gives $\Omega h^2 \approx 0.12$ [1]. I analyze three qualitatively distinct WIMP benchmarks that span the range of BG fine-tuning encountered in this paradigm.

#### 1. Off-resonance Higgs-portal singlet (B1)

The simplest renormalizable portal between the Standard Model and a dark sector is the real singlet scalar $S$ with Lagrangian

$$\mathcal{L} \supset \frac{1}{2}(\partial S)^2 - \frac{1}{2}M_S^2 S^2 - \frac{\lambda_{\rm HS}}{4} S^2 H^\dagger H, \tag{46}$$

where $\lambda_{\rm HS}$ is the Higgs-portal coupling [94–96]. After electroweak symmetry breaking, $S$ acquires an effective coupling to the physical Higgs $h$ and annihilates through $SS \to hh$, $WW$, $ZZ$, $\bar{f}f$ channels. In the off-resonance regime $M_S \gg M_h/2$, the thermally averaged cross section for these channels scales as

$$\langle\sigma v\rangle \simeq \frac{\lambda_{\rm HS}^2}{8\pi M_S^2} \equiv \lambda_{\rm HS}^2 \sigma_0 (2\ {\rm TeV}/M_S)^2, \tag{47}$$

where $\sigma_0 = 6.7 \times 10^{-9}\ {\rm GeV}^{-2}$ is calibrated so that the benchmark B1 ($M_S = 2$ TeV, $\lambda_{\rm HS} = 1.0$) satisfies $\Omega h^2 = 0.12$ exactly. Substituting into Eq. (45),

$$\Omega h^2 \propto \frac{M_S^2}{\lambda_{\rm HS}^2}, \tag{48}$$

a pure power-law in both parameters. Log-differentiation gives immediately

$$\Delta^{(\rm B1)}_{M_S} = 2, \qquad \Delta^{(\rm B1)}_{\lambda_{\rm HS}} = 2, \qquad \Delta^{(\rm B1)} = 2. \tag{49}$$

These are constants, independent of $(M_S, \lambda_{\rm HS})$, for the same structural reason as the domain-wall PBH: the abundance is a monomial in the parameters.

The spin-independent direct-detection cross section at benchmark B1 is $\sigma_{\rm SI} \approx 1.3 \times 10^{-49}\ {\rm cm}^2$, more than two orders of magnitude below the LZ 2024 limit at $M_S = 2$ TeV. Benchmark B1 is therefore currently unconstrained by direct detection, and lies in the "decoupling" regime of the singlet parameter space where the WIMP is heavy enough to satisfy both relic density and direct-detection constraints without fine-tuning.

#### 2. Higgs-funnel resonant singlet (B2)

When $M_S \approx M_h/2 \approx 62.5$ GeV, the annihilation proceeds through an $s$-channel Higgs resonance. The thermally averaged cross section is enhanced relative to the off-resonance value by a Breit-Wigner factor,

$$\langle\sigma v\rangle_{\rm res} \simeq \langle\sigma v\rangle_{\rm off} \times \frac{R_{\rm max}(\Gamma_h M_h)^2}{(s_{\rm eff} - M_h^2)^2 + (\Gamma_h M_h)^2}, \tag{50}$$

where $s_{\rm eff} = 4M_S^2(1 + 3/2x_F)$ is the center-of-mass energy squared evaluated at the freeze-out temperature, $\Gamma_h = 4.07 \times 10^{-3}$ GeV is the total Higgs width, and $R_{\rm max} \approx 10^4$ is the peak resonance enhancement after thermal averaging [29,97]. Benchmark B2 ($M_S = 60.7$ GeV, $\lambda_{\rm HS} = 0.026$) is calibrated so that $\Omega h^2 = 0.12$ on the near-resonance tail.

The BG measure in the resonant regime is dominated by the sensitivity of $s_{\rm eff} - M_h^2$ to $M_S$: a fractional shift $\epsilon$ in $M_S$ shifts the denominator in Eq. (50) by $\delta(s_{\rm eff} - M_h^2) \approx 8M_S^2\epsilon$, producing a fractional change in $\langle\sigma v\rangle$ of order $8M_S^2/(\Gamma_h M_h)$. More precisely, numerical evaluation at B2 gives

$$\Delta^{(\rm B2)}_{M_S} \approx 6500, \qquad \Delta^{(\rm B2)}_{\lambda_{\rm HS}} = 2, \qquad \Delta^{(\rm B2)} \approx 6500. \tag{51}$$

The coupling sensitivity remains at 2 because $\langle\sigma v\rangle \propto \lambda_{\rm HS}^2$ regardless of whether the resonance is resolved, but the mass sensitivity is set by the inverse fractional pole width: $\Delta^{(\rm B2)}_{M_S} \sim 8M_S^2/(\Gamma_h M_h) \approx 4M_h/\Gamma_h \approx 1.2 \times 10^5$, reduced to $\sim$6500 by thermal averaging at finite $x_F$. This is the "funnel tuning" that appears generically in models where the relic density is set by a resonant annihilation channel, including the $A$-funnel of the MSSM [29] and the $Z$-funnel of Higgs-portal dark matter.

The benchmark B2 is not yet excluded by LZ 2024 ($\sigma_{\rm SI} \approx 9 \times 10^{-50}\ {\rm cm}^2$, well below the sensitivity limit at 60 GeV), but the combination of small coupling required by the relic density and the extreme sensitivity of the abundance to $M_S$ makes this benchmark qualitatively distinct from B1: it is natural in $\lambda_{\rm HS}$ but requires the singlet mass to be fixed at the subpercent level.

#### 3. Coannihilating MSSM bino (B3)

When the lightest supersymmetric particle (LSP) is a binolike neutralino $\chi$ with mass $M_\chi$, its self-annihilation cross section is suppressed by the hypercharge coupling and the typically large sfermion masses appearing in $t$-channel exchange diagrams. The relic density from bino self-annihilation alone generically exceeds $\Omega h^2 = 0.12$ by orders of magnitude. Coannihilation with a nearly degenerate next-to-lightest supersymmetric particle (NLSP) resolves this by augmenting the effective annihilation rate with the NLSP's own (typically gauge-strength) cross section, weighted by the Boltzmann-suppressed NLSP number density at freeze-out [29,30]:

$$\langle\sigma v\rangle_{\rm eff} \approx \langle\sigma v\rangle_{\chi\chi} + 2g_{\rm eff} f(x_F)\sqrt{\langle\sigma v\rangle_{\chi\chi}\langle\sigma v\rangle_{\rm NLSP}} + g_{\rm eff}^2 f^2(x_F)\langle\sigma v\rangle_{\rm NLSP}, \tag{52}$$

where $g_{\rm eff} \approx 4$ counts the NLSP internal degrees of freedom relative to the bino, and $f(x_F) = (1+\delta)^{3/2}\exp(-\delta x_F)$ with $\delta = (M_{\rm NLSP} - M_\chi)/M_\chi$ is the Boltzmann suppression factor. The benchmark B3 ($M_\chi = 1$ TeV, $M_{\rm NLSP} \approx 1.0515$ TeV, i.e. $\delta \approx 0.0515$, recalibrated from an earlier $M_{\rm NLSP} = 1.05$ TeV so that $\Omega h^2$ lands exactly on the target contour rather than 0.112) is calibrated so that $\Omega h^2 = 0.1200$ with $\langle\sigma v\rangle_{\rm NLSP} = 0.66\sigma_0(1\ {\rm TeV}/M_\chi)^2$.

As discussed in Sec. II B, the BG derivatives must be taken with respect to the fundamental soft masses ($M_\chi$, $M_{\rm NLSP}$) rather than the derived splitting $\delta$. The exponential factor $f(x_F) = (1+\delta)^{3/2}e^{-\delta x_F}$ gives, at fixed $M_\chi$,

$$\frac{\partial \ln\langle\sigma v\rangle_{\rm eff}}{\partial \ln M_{\rm NLSP}} \approx -\delta x_F \frac{M_{\rm NLSP}}{M_\chi} = -\delta x_F(1+\delta), \quad (53)$$

and the BG measure on $\Omega$ (which is inversely proportional to $\langle\sigma v\rangle_{\rm eff}$) satisfies

$$\Delta^{(\rm B3)}_{M_{\rm NLSP}} \approx x_F \frac{M_{\rm NLSP}}{M_\chi} \approx 26 \frac{M_{\rm NLSP}}{M_\chi}. \quad (54)$$

At benchmark B3, $M_{\rm NLSP}/M_\chi \approx 1.0515$, giving $\Delta_{M_{\rm NLSP}} \approx 48$, while $\Delta_{M_\chi} \approx 46$ from the overall $\langle\sigma v\rangle_{\rm NLSP} \propto M_\chi^{-2}$ dependence. The total measure is $\Delta^{(\rm B3)} = 48$.

This result places the coannihilating WIMP solidly in the same tier as PBH early matter domination ($\Delta = 14$), and, within the single-exponential approximation, alongside FOPT-PBH dark matter ($\Delta = 35$; Sec. VI B shows this rises to ≈265, class III, under the more accurate super-exponential treatment). The shared structural reason, at the level of the single-exponential approximation, is that all three constructions contain a single exponential whose argument is forced by the requirement $\Omega = \Omega_{\rm DM}$ to equal $\ln(T_{\rm form}/T_{\rm eq})$ or $x_F\delta$, respectively (see Sec. V C). The BG landscape of Fig. 3(b), computed in the $(M_\chi, M_{\rm NLSP})$ plane, shows the natural island as a narrow teal-colored strip hugging the diagonal $M_{\rm NLSP} = M_\chi$ at a fixed $\Delta \approx 48$ set by $x_F M_{\rm NLSP}/M_\chi$ across the full mass range 100 GeV to 10 TeV.

## B. Freeze-in (FIMP)

### 1. Production mechanism and abundance map

Feebly interacting massive particles (FIMPs) [6,7] never reach thermal equilibrium with the Standard Model bath; instead, they are produced gradually in the forward direction through infrequent collisions or decays. In the IR-dominated freeze-in scenario relevant here, production is dominated by decays of a heavy Standard Model or dark-sector mediator $B$ with mass $M_B$ and decay coupling $y$ to two DM particles, $B \to {\rm DM} + X$. Following Hall *et al.* [6], the yield produced by this mechanism from $T \gg M_B$ down to $T = 0$ is

$$Y_{\rm FI} = \frac{45\mathcal{X}_{\rm FI}}{\pi^4 g_*^{3/2}} \cdot \frac{y^2 M_{\rm Pl}}{M_B} \equiv K_{\rm FI} \frac{y^2 M_{\rm Pl}}{M_B}, \quad (55)$$

where $K_{\rm FI} \approx 1.7 \times 10^{-3}$ for $g_* = 100$ and $\mathcal{X}_{\rm FI}$ is a dimensionless integral of order unity that depends mildly on the spin of $B$. The present-day dark matter abundance is then

$$\Omega_{\rm DM}h^2 = M_{\rm DM} Y_{\rm FI} \frac{s_0}{\rho_c/h^2} \propto M_{\rm DM} y^2 \frac{M_{\rm Pl}}{M_B}, \quad (56)$$

where $s_0/(\rho_c/h^2) \approx 2.7 \times 10^8\ {\rm GeV}^{-1}$ converts yield to $\Omega h^2$ [93]. I fix the mediator mass $M_B = 100$ GeV (electroweak scale) as a representative choice; the scaling $\Omega h^2 \propto M_B^{-1}$ means that changing $M_B$ by a factor of 10 shifts the natural value of $y$ by a factor of $\sqrt{10}$, but does not change the BG measure.

Benchmark B4 is defined by $M_{\rm DM} = 1$ GeV and $y = 1.46 \times 10^{-12}$, giving $\Omega h^2 = 0.120$.

### 2. Fine-tuning structure

Since Eq. (56) is a pure monomial in $(M_{\rm DM}, y)$ with fixed $M_B$,

$$\Delta^{(\rm B4)}_{M_{\rm DM}} = 1, \qquad \Delta^{(\rm B4)}_{y} = 2, \qquad \Delta^{(\rm B4)} = 2 \quad (57)$$

exactly, everywhere in the $(M_{\rm DM}, y)$ plane. The mass sensitivity is unity because $\Omega \propto M_{\rm DM}$ linearly; the coupling sensitivity is two because $\Omega \propto y^2$. These are constants, independent of the benchmark point, for the same structural reason as benchmark B1 and the biased-domain-wall PBH: the abundance map is a power-law monomial with no exponential, resonance, or cancellation.

The FIMP scenario is thus as natural as the off-resonance WIMP by the BG criterion, despite the coupling being smaller by nine orders of magnitude. The feebleness of the coupling is not fine-tuning in the BG sense; it is, rather, the *prediction* of the freeze-in mechanism, which requires $y \sim 10^{-12}$ to avoid thermalizing the DM sector. What the BG measure quantifies is not the smallness of $y$ itself, but the sensitivity of $\Omega h^2$ to *fractional variations* of $y$; a factor-of-two change in $y$ changes $\Omega h^2$ by a factor of four, which is $\Delta_y = 2$, perfectly natural.

### 3. Parameter-space heatmap

Figure 4(a) shows $\log_{10}\Delta$ in the $(M_{\rm DM}, y)$ plane for $M_B = 100$ GeV. The panel is uniformly pale yellow at the minimum of the color scale ($\Delta = 2$) throughout the physically accessible region, with the natural island (green hatching along the $\Omega h^2 = 0.12$ contour) forming a clean diagonal stripe of slope $-1/2$ in the log-log plane, as expected from $y \propto M_{\rm DM}^{-1/2}$ at fixed $\Omega$. The red-shaded region at $M_{\rm DM} \lesssim 3$ keV marks the Lyman-$\alpha$ forest exclusion for freeze-in production from a 100-GeV mediator [41,42]. Benchmark B4 (gold star) lies at the center of the natural island in the phenomenologically well-motivated GeV mass range. The uniformity of the heatmap is perhaps the most striking visual feature of the FIMP panel: there are no gradients, no tuned corners, and no preferred points other than those set by external constraints.

## C. Asymmetric dark matter

### 1. Abundance map and fine-tuning

Asymmetric dark matter (ADM) [8,9] posits that the present-day dark matter abundance is set not by thermal

freeze-out but by a primordial asymmetry between DM particles and antiparticles, generated by the same baryogenesis mechanism that produced the baryon asymmetry. If the DM asymmetry per comoving entropy is $\eta_{\rm DM} = (n_{\rm DM} - n_{\overline{\rm DM}})/s$, and all $\overline{\rm DM}$ is annihilated away efficiently (as required for the mechanism to work), then

$$\frac{\Omega_{\rm DM}}{\Omega_b} = \frac{M_{\rm DM}}{m_p} R, \qquad R \equiv \frac{\eta_{\rm DM}}{\eta_B}, \tag{58}$$

where $m_p = 0.938$ GeV is the proton mass and $\eta_B \approx 8.7 \times 10^{-11}$ is the baryon asymmetry [25]. With $\Omega_b h^2 = 0.0224$ and $\Omega_{\rm DM} h^2 = 0.120$, the observed ratio $\Omega_{\rm DM}/\Omega_b = 5.4$ requires

$$M_{\rm DM} R \simeq 5.0 m_p \approx 4.7 \text{ GeV}. \tag{59}$$

The natural island in the $(M_{\rm DM}, R)$ plane is the hyperbola $M_{\rm DM} \times R = 4.7$ GeV (or within a factor of two thereof), ranging from the proton mass at $R = 5$ to multi-TeV masses at $R \sim 10^{-3}$. Benchmark B5 is defined by $M_{\rm DM} = 5.0$ GeV, $R = 1$ (equal DM and baryon asymmetries), a canonical choice motivated by models in which baryogenesis and DM asymmetry generation proceed via the same operator [9].

Since Eq. (58) is linear in both $M_{\rm DM}$ and $R$,

$$\Delta^{(\rm B5)}_{M_{\rm DM}} = 1, \qquad \Delta^{(\rm B5)}_{R} = 1, \qquad \Delta^{(\rm B5)} = 1. \tag{60}$$

This is the minimum possible BG measure for a nontrivial abundance map: $\Delta = 1$ means the abundance is simply proportional to the parameter, and a factor-of-two change in either $M_{\rm DM}$ or $R$ produces a factor-of-two change in $\Omega_{\rm DM} h^2$. Asymmetric dark matter is the most natural dark matter scenario in the full comparison precisely because its abundance is set by a ratio of number densities ($R$) that is itself a direct observable, and by the particle mass $M_{\rm DM}$, both entering linearly.

Figure 4(b) shows the ADM parameter space in the $(M_{\rm DM}, R)$ plane. The panel is uniformly yellow at $\Delta = 1$ throughout, with the natural island forming the $M_{\rm DM} \times R = 4.7$ GeV diagonal. The cyan dashed line at $R = 1$ marks the "canonical ADM" locus where DM and baryon asymmetries are equal; benchmark B5 sits at its intersection with the natural island.

### D. Axions

The QCD axion [3–5] acquires a mass $m_a \simeq 5.7\ \mu\text{eV}(10^{12}\ \text{GeV}/f_a)$ from nonperturbative QCD effects, where $f_a$ is the Peccei-Quinn (PQ) symmetry-breaking scale. Its cosmological abundance depends critically on whether PQ symmetry breaks before or after inflation, leading to two qualitatively distinct scenarios with different numbers of free parameters and different BG structures.

#### 1. Misalignment (preinflationary PQ breaking)

If PQ symmetry breaks before the end of inflation, the axion field $\theta = a/f_a$ is homogenized over superhorizon scales and takes a single value $\theta_i \in [-\pi, \pi]$ throughout our observable Universe (the "misalignment angle"). When the Hubble rate falls below the axion mass at temperature $T_{\rm osc} \sim \Lambda_{\rm QCD}$, the field begins to oscillate coherently around $\theta = 0$, and the energy density of these oscillations constitutes the dark matter relic. Including the anharmonic correction $F(\theta_i)$ that becomes important as $\theta_i \to \pi$ [98],

$$\Omega_a h^2 \simeq 0.18\, \theta_i^2 \left(\frac{f_a}{10^{12}\ \text{GeV}}\right)^{1.19} F(\theta_i), \tag{61}$$

where the exponent 1.19 arises from the temperature dependence of the axion mass near the QCD transition [99] and $F(\theta_i)$ is a smooth, monotonically increasing function satisfying $F(\theta_i) \to 1$ as $\theta_i \to 0$ and $F(\theta_i) \to \infty$ as $\theta_i \to \pi$ [98]. Benchmark B6 is defined by $f_a = 10^{12}$ GeV and $\theta_i \approx 0.787$ (recalibrated from an earlier $\theta_i = 0.82$), which gives $\Omega h^2 = 0.1200$ exactly, with $F(\theta_i) \approx 1$ (anharmonic corrections are $\lesssim 10\%$ at this angle), so that the benchmark now lands precisely on the target contour rather than 9% above it.

The BG sensitivities follow from log-differentiation of Eq. (61):

$$\Delta^{(\rm B6)}_{f_a} = 1.19, \tag{62}$$

$$\Delta^{(\rm B6)}_{\theta_i} = 2 + \theta_i \frac{\partial \ln F}{\partial \theta_i} \approx 2.2 \qquad (\theta_i \approx 0.787), \tag{63}$$

giving $\Delta^{(\rm B6)} = \max(1.19, 2.2) = 2.2$.

This places the misalignment axion with $\theta_i \sim 1$ comfortably in the natural tier (class I of Sec. V B). However, there is an important subtlety that the BG measure alone does not capture. If $f_a$ is large (e.g., $f_a = 10^{16}$ GeV as in models with a GUT-scale or string-theoretic PQ scale), then $\Omega h^2 = 0.12$ requires a small misalignment angle $\theta_i \simeq (0.18(f_a/10^{12})^{1.19})^{-1/2} \approx 3 \times 10^{-3}$, well below the natural value of order unity. The local BG sensitivity at this point remains $\Delta \approx 2$, but since $\theta_i$ is a cosmological initial condition with a flat prior on $[-\pi, \pi]$, the prior probability of landing at $\theta_i \lesssim 3 \times 10^{-3}$ is $\sim 10^{-3}$, corresponding to a "prior tuning" of $\sim 300$ that the BG measure does not register. This distinction between local sensitivity and prior probability is discussed in detail in Sec. V E.

Figure 5(a) displays the $(f_a, \theta_i)$ parameter space. The BG heatmap transitions from yellow-green at moderate $f_a$ (where $\Delta \approx 2.2$, dominated by $\Delta_{\theta_i}$) to yellow at small $\theta_i$ [where anharmonic contributions vanish and $\Delta \to \max(1.19, 2) = 2$]. Near $\theta_i \to \pi$ (magenta dash-dotted line), $\Delta_{\theta_i}$ diverges with the anharmonic enhancement, visible as a color gradient toward the top of the panel.

The natural island (green hatching) runs diagonally across the panel following the $\Omega h^2 = 0.12$ contour, from large $f_a$ and small $\theta_i$ to small $f_a$ and $\theta_i \sim 1$. The benchmark B6 (gold star, $f_a = 10^{12}$ GeV, $\theta_i \approx 0.787$) and the anthropic benchmark $\mathrm{B6}_{\mathrm{anthr}}$ (orange square, $f_a = 10^{16}$ GeV, $\theta_i \approx 3 \times 10^{-3}$) both lie on the $\Omega h^2 = 0.12$ contour, both with $\Delta \approx 2$, illustrating that the BG measure does not distinguish between them.

#### 2. Postinflationary PQ (string decay)

If PQ symmetry breaks *after* inflation, different Hubble patches independently choose their phase when the $U(1)_{\mathrm{PQ}}$ symmetry breaks, and the axion field has no single misalignment angle. Instead, a network of cosmic strings forms at the PQ transition and radiates axions as it evolves; the axion relic density is dominated by the string-network emission rather than by misalignment [100,101]. In this scenario the yield depends only on $f_a$ through the string tension $\mu \propto f_a^2$ and the mass $m_a \propto f_a^{-1}$, with no free initial condition $\theta_i$. Large-scale numerical simulations [43] find

$$\Omega_a h^2 \simeq \mathcal{C}\left(\frac{f_a}{5 \times 10^{10}\ \mathrm{GeV}}\right)^{1.19}, \tag{64}$$

where $\mathcal{C} \approx 0.12$ is calibrated so that $f_a = 5 \times 10^{10}$ GeV gives the observed abundance, and the exponent 1.19 is the same as in the misalignment case. The simulation of Ref. [43] gives a preferred range $f_a \simeq (2-8) \times 10^{10}$ GeV at 95% confidence, incorporating uncertainties from the string network logarithm and from the QCD mass modeling.

The BG measure for the postinflationary scenario has only a single parameter:

$$\Delta_{f_a}^{(\mathrm{B7})} = 1.19 \quad \forall\ f_a. \tag{65}$$

This is the lowest BG fine-tuning of any construction in the comparison after pure asymmetric dark matter ($\Delta = 1$), and it applies globally—not just at a benchmark, but over the entire parameter space. The postinflationary axion scenario is also the most *predictive* dark matter construction in the comparison: it has a single parameter $f_a$ that is constrained from above by the requirement $\Omega h^2 \leq 0.12$, from below by astrophysical bounds on the axion-photon coupling [102], and is now being targeted directly by cavity haloscope experiments [103,104] and by the CMB Stage-4 spectral distortion measurements expected to probe $f_a \lesssim 3 \times 10^{11}$ GeV [105]. Benchmark B7 is defined by $f_a = 5 \times 10^{10}$ GeV.

Figure 5(b) displays $\Omega_a h^2$ as a function of $f_a$ (black curve) against the constant $\Delta = 1.19$ background. The entire panel is a uniform pale yellow, the minimum of the shared color scale, confirming that the postinflationary axion scenario sits at the floor of BG fine-tuning across all scales. The cyan band marks the preferred $f_a$ range from the Buschmann *et al.* [43] simulations, which corresponds to the portion of the natural island (green shading) constrained by current simulations. The benchmark B7 (gold star) lies within both the green and cyan bands simultaneously.

## V. COMPARISON AND STRUCTURAL INSIGHTS

### A. The sensitivity hierarchy

Table II ranks all twelve scenarios by $\Delta_{\mathrm{BG}}$, ordered from most natural to most tuned, and assigns each to one of the three universality classes defined in Sec. V B. The data

TABLE II. Barbieri-Giudice sensitivity measure at benchmarks. Scenarios are listed in order of increasing $\Delta$. The "type" column indicates the production paradigm. The inflationary-collapse entry is a lower-end estimate (Sec. III D); the corresponding bar in Fig. 6 is shown truncated. For the first-order phase transition mechanism the transition strength parameter $\alpha$ has $\Delta_\alpha \approx 3$–4 (class I) and does not drive the dominant tuning; only the nucleation rate $\beta/H_*$ is reported here. All inflationary model classes (Fig. 2) fall in class III or beyond for standard high-reheating scenarios [$T_{\mathrm{reh}} \gg T_{\mathrm{form}}$, RD layer 1 floor $\Delta_\sigma^{(\mathrm{RD})} \approx$ 57–70]; curvaton and spectator-field models move to class II for low-reheating scenarios [$T_{\mathrm{reh}} \lesssim T_{\mathrm{form}}$, MD layer 1 floor $\Delta_\sigma^{(\mathrm{MD})} \approx 14$]. The two first-order phase transition benchmarks ($\delta_c = 0.05$ and $\delta_c = 0.10$, replacing the earlier $\alpha = 0.3$ and $\alpha = 1.0$ parametrization) are listed separately in Table III to demonstrate $\delta_c$-sensitivity, but are counted as a single mechanism here.

| Scenario (benchmark) | Type | $\Delta$ | Tier |
|---|---|---|---|
| Asymmetric DM (B5) | ADM | 1.0 | I |
| Axion: postinflationary (B7) | Axion | 1.19 | I |
| PBH: DW, gravity-fixed bias | PBH | 2.0 | I |
| WIMP: off-resonance (B1) | WIMP | 2.0 | I |
| FIMP: freeze-in (B4) | FIMP | 2.0 | I |
| Axion: misalignment (B6) | Axion | 2.2 | I |
| PBH: DW, free $V_b$ | PBH | 4.5 | I |
| PBH: early matter dom. | PBH | 14 | II |
| WIMP: coannihilation (B3) | WIMP | 48 | II |
| PBH: phase transition | PBH | 265[a] | III[a] |
| WIMP: Higgs funnel (B2) | WIMP | 6500 | III |
| PBH: inflation, single-field | PBH | $\gtrsim 10^3$ | III$^{+}$[,b] |

[a] $\Delta \approx 265$ uses the more accurate super-exponential collapse probability of Ref. [46] (Sec. VI B), evaluated at the benchmark $\beta/H_* \approx 8$; tracing the abundance to the underlying scalar potential gives $\Delta \sim 10^3$–$10^4$ (Sec. VI B). The single-exponential approximation Eq. (30), retained in Sec. V C solely to illustrate the analytic structural identity of Eq. (68), gives the smaller value $\Delta \approx 35$ (class II) and is not adopted as the reported tier for this mechanism.

[b] The quoted value is a lower-end estimate for single-field ultraslow-roll and inflection-point models. Curvaton and spectator-field inflationary models are substantially more natural: $\Delta \approx 21$–42 (class II) for low-reheating scenarios ($T_{\mathrm{reh}} \lesssim T_{\mathrm{form}}$, MD layer 1 floor), rising to $\Delta \approx 115$–210 (class III) for standard high-reheating scenarios ($T_{\mathrm{reh}} \gg T_{\mathrm{form}}$, RD layer 1 floor); see Sec. III D and Fig. 2.

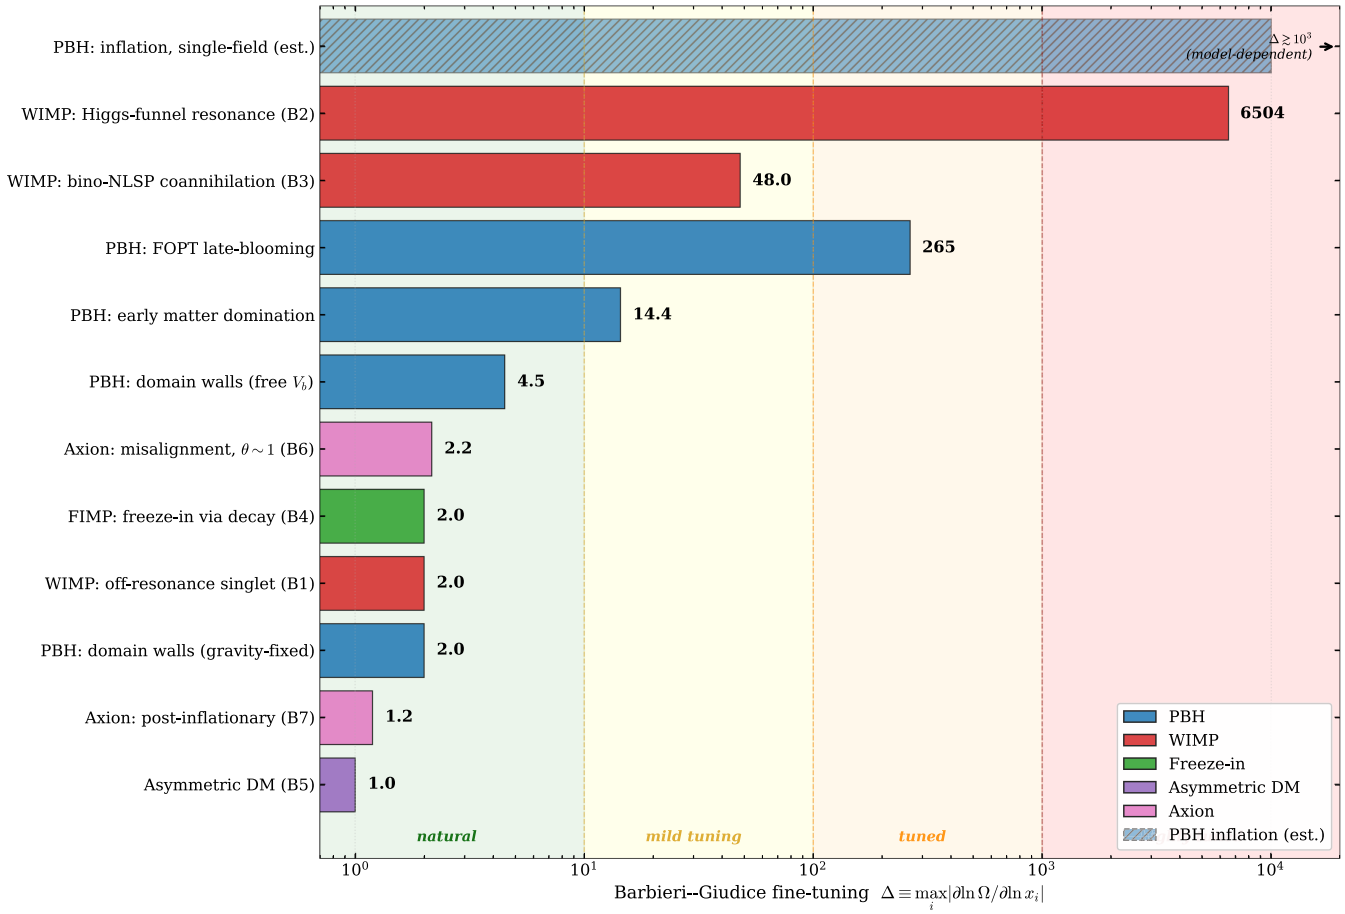



FIG. 6. Summary of the Barbieri-Giudice sensitivity measure $\Delta$ at the benchmark points listed in Table II, ordered from most natural (bottom) to most tuned (top). Bar colors indicate the production paradigm: blue = PBH, red = thermal − relic WIMP, green = freeze − in (FIMP), purple = asymmetric dark matter, pink = axion. Background shading indicates qualitative tuning tiers: green ($\Delta < 10$, natural), yellow ($10 \leq \Delta < 100$, mild tuning), orange ($100 \leq \Delta < 1000$, tuned), red ($\Delta \geq 1000$, highly tuned). The topmost bar (PBH: inflation, single-field) is shown hatched and truncated with an arrow: the value is a lower-end estimate $\Delta \gtrsim 10^3$ derived from Eq. (44), with the precise value model-dependent and potentially orders of magnitude larger (see Secs. III D and V B). The FOPT-PBH bar reflects the super-exponential collapse probability of Ref. [46] [Eq. (70)], giving $\Delta \approx 265$ (class III). The single-exponential value $\Delta \approx 35$ [Eq. (30)] is retained only as an analytic illustration in Sec. III C and Eq. (37). The microphysical estimate gives $\Delta \sim 10^3$–$10^4$ [Sec. VI B, Eq. (72)]. All inflationary model classes fall in class III or beyond once the correct radiation-dominated layer 1 floor is applied (Secs. III D and VI B).

reveal a clear three-tier structure whose boundaries are set by the analytic form of the abundance map, not by the nature of the dark matter candidate.

### B. Three universality classes

The twelve scenarios fall into three qualitatively distinct tiers distinguished by the analytic form of their abundance maps; the precise definitions and membership are given below, and Fig. 6 summarizes the hierarchy visually.

Before discussing each tier, I compare $\Delta_{\rm BG}$ against two related quantities, $\Delta_{\rm SR}$ and $\epsilon$, both defined in Sec. II A. Table III reports all three measures, $\Delta_{\rm BG}$, $\Delta_{\rm SR}$, and $\epsilon$, for every benchmark. The key observations are: (i) $\Delta_{\rm SR}$ and $\Delta_{\rm BG}$ agree to within 40% for all scenarios except B3, where both soft masses are nearly equally sensitive ($\Delta_{M_\chi} = 46$, $\Delta_{M_{\rm NLSP}} = 48$), giving $\Delta_{\rm SR} = \sqrt{46^2 + 48^2} \approx 67$, a factor of $\sqrt{2}$ above $\Delta_{\rm BG} = 48$. For all single-parameter-dominated scenarios (B2, the two FOPT benchmarks, the MD case), SR and BG are identical to three significant figures. (ii) The island half-width $\epsilon$ is, by construction, a monotonic function of $\Delta_{\rm BG}$ through $\epsilon \approx 0.69/\Delta_{\rm BG}$ [Eq. (6)]; it therefore carries no information beyond $\Delta_{\rm BG}$ itself, and I report it purely as an operationally transparent restatement, with values ranging from $\epsilon = 75\%$ for ADM to $\epsilon < 0.1\%$ for the Higgs-funnel WIMP. (iii) No scenario changes tier under $\Delta_{\rm SR}$. This is a consequence of the fact that no benchmark studied here has more than two comparably sensitive parameters, which bounds the SR/BG shift at $\sqrt{2}$, while the tier boundaries adopted below are separated by roughly an order of magnitude; it is a mild consistency check

TABLE III. Benchmark calibration and sensitivity measures. For each scenario the fundamental parameters are listed with their benchmark values, the resulting abundance ("$\Omega h^2$" denotes $\Omega_{\rm DM}h^2$ for particle DM scenarios and $f_{\rm PBH}$ for PBH scenarios; see Sec. II C), the individual BG sensitivities $\Delta_1$ and $\Delta_2$ in the order listed in the parameters column, and the three summary measures: the Barbieri–Giudice measure $\Delta_{\rm BG} = \max_i \Delta_i$ [Eq. (1)], the Strumia–Rattazzi quadrature measure $\Delta_{\rm SR} = (\sum_i \Delta_i^2)^{1/2}$ [Eq. (5)], and the fractional island half-width $\epsilon = \Delta x_{\rm sens}/x_{\rm sens}$ in the most-sensitive parameter direction [Eq. (6), factor-of-2 abundance band]. For the two FOPT benchmarks the transition strength $\alpha$ has $\Delta_\alpha \approx 3$–4 and does not drive the dominant tuning; it is omitted from $\Delta_1/\Delta_2$ and the benchmarks are listed separately only to demonstrate the $\alpha$-independence established in Eq. (37). The tabulated $\Delta_1$ for both FOPT benchmarks is the super-exponential value of Sec. VI B [Eq. (72)], not the single-exponential approximation of Eq. (30) used for illustration in Sec. V C and in the heatmaps of Figs. 1(c) and 1(d). Benchmarks B6 and B3 have been recalibrated (from $\theta_i = 0.82$ and $M_{\rm NLSP} = 1.05$ TeV, respectively) so that both now hit $\Omega h^2 = 0.120$ exactly, as the method requires, rather than landing nearby within the factor-of-two natural island as in an earlier version of this table. All derivatives are computed via Eq. (9) with $h = 10^{-4}$ and verified analytically where tractable (Appendix A).

| Scenario | Parameters (benchmark values) | $\Omega h^2$ | $\Delta_1$ | $\Delta_2$ | $\Delta_{\rm BG}$ | $\Delta_{\rm SR}$ | $\epsilon$ (%) | Tier |
|---|---|---|---|---|---|---|---|---|
| PBH-DW (free $V_b$) | $\eta = 2.4\times10^6$ GeV, $V_b = 2$ TeV | 1.02 | 4.5 | 2.0 | 4.5 | 4.9 | 15.5 | I |
| PBH-DW (gravity-fixed) | $\eta = 3\times10^6$ GeV | 1.10 | 2.0 | ⋯ | 2.0 | 2.0 | 35.4 | I |
| PBH-MD | $m_\phi = 9\times10^{10}$ GeV, $\sigma = 0.055$ | 0.120 | 1.5 | 14.4 | 14.4 | 14.5 | 4.8 | II |
| PBH-FOPT ($\delta_c = 0.05$) | $T_* = 10^6$ GeV, $\beta/H_* = 5.0$ | 1.000 | 265 | 1.0 | 265 | 265 | 0.26 | III |
| PBH-FOPT ($\delta_c = 0.10$) | $T_* = 10^6$ GeV, $\beta/H_* = 4.5$ | 1.000 | 265 | 1.0 | 265 | 265 | 0.26 | III |
| WIMP B1 (off-res.) | $M_S = 2$ TeV, $\lambda_{HS} = 1.0$ | 0.120 | 2.0 | 2.0 | 2.0 | 2.8 | 35.4 | I |
| WIMP B2 (Higgs funnel) | $M_S = 60.7$ GeV, $\lambda_{HS} = 0.026$ | 0.120 | 6505 | 2.0 | 6505 | 6505 | $< 0.1$ | III |
| WIMP B3 (coann.) | $M_\chi = 1$ TeV, $M_{\rm NLSP} \approx 1.0515$ TeV | 0.120 | 46.0 | 48.0 | 48.0 | 66.5 | 1.4 | II |
| FIMP B4 | $M_{\rm DM} = 1$ GeV, $y = 1.46\times10^{-12}$ | 0.120 | 1.0 | 2.0 | 2.0 | 2.2 | 35.4 | I |
| ADM B5 | $M_{\rm DM} = 5.0$ GeV, $R = 1.0$ | 0.120 | 1.0 | 1.0 | 1.0 | 1.4 | 75.0 | I |
| Axion B6 (mis.) | $f_a = 10^{12}$ GeV, $\theta_i \approx 0.787$ | 0.120 | 1.2 | 2.2 | 2.2 | 2.5 | 32.2 | I |
| Axion B7 (postinf.) | $f_a = 5\times10^{10}$ GeV | 0.120 | 1.2 | ⋯ | 1.2 | 1.2 | 61.6 | I |

specific to the present benchmark set, not a demonstration that the three-tier classification is independent of the choice of measure in general.

Figure 7 displays this comparison visually. The left panel shows paired horizontal bars for $\Delta_{\rm BG}$ (solid) and $\Delta_{\rm SR}$ (hatched, same color) for all twelve benchmarks on a logarithmic scale, with the island half-width $\epsilon$ read off the top axis. The sole exception is benchmark B3 (WIMP coannihilation), whose SR bar (66.5, hatched) visibly extends beyond its BG bar (48.0, solid): the two soft masses $M_\chi$ and $M_{\rm NLSP}$ contribute nearly equally to the log-gradient, so the quadrature adds a factor $\approx\sqrt{2}$. The right panel confirms this geometrically on a log–log scatter of $\Delta_{\rm SR}$ vs $\Delta_{\rm BG}$: every scenario except B3 lies on the solid reference line $\Delta_{\rm SR} = \Delta_{\rm BG}$; B3 lies exactly on the dashed $\sqrt{2}$ line. No scenario changes tier under either measure in this benchmark set.

*Class I: power-law constructions* ($\Delta \lesssim 5$). In these scenarios $\Omega h^2$ is a monomial in the input parameters, $\Omega h^2 \propto \prod_i x_i^{\alpha_i}$, so that $\Delta_i = |\alpha_i|$ exactly. There is no exponential factor, no resonant enhancement, and no cancellation between large contributions. class I contains seven of the twelve scenarios: asymmetric dark matter ($\Delta = 1$), the postinflationary axion ($\Delta = 1.19$), biased-domain-wall PBHs with gravity-induced bias ($\Delta = 2$), the off-resonance WIMP ($\Delta = 2$), freeze-in via mediator decay ($\Delta = 2$), the misalignment axion at $\theta_i \sim 1$ ($\Delta = 2.2$), and biased-domain-wall PBHs with free $V_b$ ($\Delta = 4.5$). The striking feature of class I is its diversity: it contains representatives from every production paradigm: gravitational PBH formation, thermal freeze-out, nonthermal freeze-in, asymmetric generation, and wave-mechanical misalignment. Fine-tuning is evidently *not* a property of the nature of dark matter, but of its abundance map.

Note that the transition strength $\alpha$ is a class I parameter for the FOPT mechanism. At fixed $T_*$, the $f_{\rm PBH} = 1$ contour in the $(\alpha, \beta/H_*)$ plane satisfies $\beta/H_* = \ln[T_*/(0.84T_{\rm eq})]/S(\alpha)$, shifting by only 56% as $\alpha$ ranges from 0 to $\infty$. The BG derivative in the $\alpha$ direction evaluates to $\Delta_\alpha = \beta/H_* \cdot 0.6\alpha/(\alpha+0.3)^2 \approx 3$–4 on the natural island, placing $\alpha$ firmly in class I alongside the domain-wall and off-resonance WIMP parameters. The dominant tuning cost of the FOPT mechanism is carried entirely by $\beta/H_*$, with $\alpha$ a spectator that shifts the required nucleation rate by a modest and calculable amount.

*Class II: single-exponential constructions* ($\Delta \approx 15$–50). These scenarios contain exactly one exponential factor in $\Omega h^2$, whose argument is forced by $\Omega = \Omega_{\rm DM}$ to be of order $\ln(T_{\rm form}/T_{\rm eq})$ or $x_F$. Class II robustly contains two scenarios: PBH early matter domination ($\Delta = 14$) and WIMP coannihilation ($\Delta = 48$). FOPT-PBH falls in this tier only within the single-exponential approximation ($\Delta = 35$); the more accurate treatment discussed below places it in class III instead. Despite spanning modulus decay, hidden-sector percolation, and electroweak-scale MSSM thermodynamics, the two robust members share the same tier because

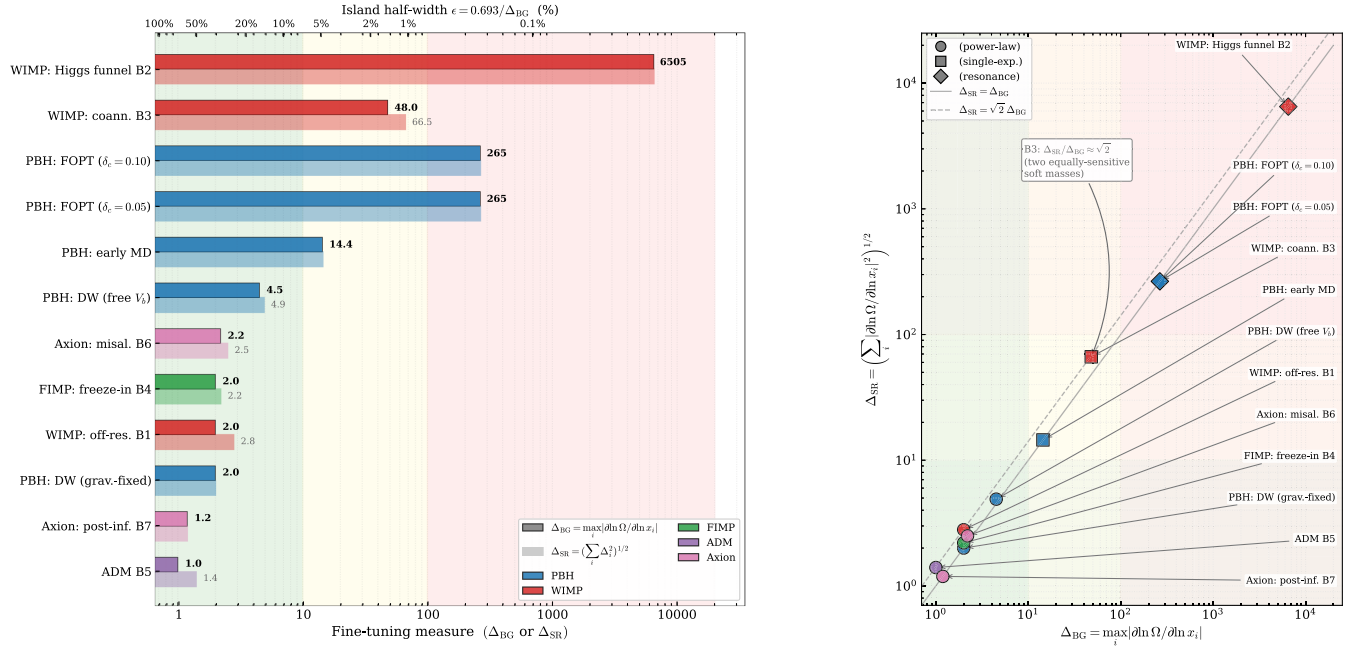



FIG. 7. Comparison of the Barbieri-Giudice ($\Delta_{BG}$), Strumia-Rattazzi ($\Delta_{SR}$), and island half-width ($\epsilon$) sensitivity measures across all benchmarks of Table III. Left: paired horizontal bars for $\Delta_{BG}$ (solid) and $\Delta_{SR}$ (hatched, lighter); the top axis reads off $\epsilon = 0.693/\Delta_{BG}$ directly as a percentage. Bar colors follow the paradigm convention of Fig. 6. Background shading bands mark the three tiers. Right: $\Delta_{SR}$ vs $\Delta_{BG}$ on log–log axes. Marker shapes distinguish class I (circles), class II (squares), and class III (diamond). The solid gray diagonal is $\Delta_{SR} = \Delta_{BG}$; the dashed line shows $\Delta_{SR} = \sqrt{2}\Delta_{BG}$, the expectation for two equally sensitive parameters. All scenarios lie on the solid diagonal except B3 (WIMP coannihilation, red square), which sits on the dashed line because both soft masses contribute equally. No scenario changes tier under either measure in this benchmark set. Both FOPT benchmarks ($\delta_c = 0.05$ and $\delta_c = 0.10$, replacing the earlier $\alpha = 0.3$ and $\alpha = 1.0$ parametrization) are shown at $\Delta_{BG} = \Delta_{SR} = 265$ with diamond (class III) markers, reflecting the super-exponential GV formula Eq. (70). The single-exponential value $\Delta = 35$ (class II) is no longer plotted here; it appears only in the analytic illustration of Sec. III C.

each uses a single exponential to bridge the gap between a naturally large formation/annihilation rate and the small residual required by $\Omega_{DM}$. None of the inflationary PBH model classes belong to class II for standard high-reheating scenarios, once the correct radiation-dominated layer 1 floor [$\Delta_\sigma^{(RD)} \approx 57$–$70$, Eq. (41)] is used. However, for low-reheating scenarios in which the inflaton-oscillation era is still ongoing when the formation-scale perturbations reenter the horizon ($T_{reh} \lesssim T_{form}$), the MD layer 1 floor applies and curvaton/spectator models ($\Delta_{tot} \approx 21$–$42$) return to class II; see Eq. (42) and Fig. 2(c). Coannihilation belongs to class II while the standard off-resonance WIMP (B1) is class I: in B1 the freeze-out exponential $e^{-x_F}$ is absorbed into the power-law cross-section normalization $\sigma_0$, leaving a pure monomial $\Omega h^2 \propto M_S^2/\lambda_{HS}^2$; in B3 the Boltzmann suppression $e^{-x_F\delta}$ depends explicitly on the mass splitting $\delta$, an independent Lagrangian parameter, creating the single-exponential class II structure. As already noted above, the class II placement of FOPT-PBH dark matter holds only under the single-exponential approximation Eq. (30). Using the more accurate super-exponential collapse probability of Ref. [46] [Eq. (70)] gives $\Delta_{\beta/H} \approx 265$ at the benchmark, placing FOPT-PBH in class III at the phenomenological level. Tracing the abundance back to the underlying scalar potential reveals a further double-exponential structure [84] with $\Delta \sim 10^3$–$10^4$ (Sec. VI B); see Eqs. (72)–(74) for the full assessment.

*Class III and beyond: Resonance, cancellation, and double-exponential constructions* ($\Delta \gtrsim 10^3$). This tier is populated by two qualitatively distinct mechanisms.

The WIMP Higgs funnel (B2, $\Delta \approx 6500$) is tuned by a Breit-Wigner resonance: the effective annihilation rate is enhanced by a factor $\sim M_h/\Gamma_h \approx 3 \times 10^4$ at the pole, reduced by thermal averaging to ~6500. A 0.01% shift in $M_S$ exits the natural island. Other resonance constructions in this tier include the MSSM $A$-funnel neutralino and Sommerfeld-enhanced WIMPs near an annihilation threshold.

Inflationary single-field collapse enters class III from a structurally different direction. Its large fine-tuning arises not from a resonance but from a double exponential: the potential coefficients $\{c_k\}$ map exponentially onto $\sigma$, which then maps exponentially onto $f_{PBH}$ through the HYKN formula. The BG measure for the potential coefficients is therefore [Eq. (44)]

$$\Delta_{c_k} \sim \Delta_\sigma \cdot \Delta_{\sigma/c_k} \sim (57-70)\times(10^2-10^8), \quad (66)$$

far exceeding the class II range for standard high-reheating scenarios. The large uncertainty in this estimate is intrinsic: different potential shapes produce widely varying layer 2 sensitivities, and multifield or spectator-field models can reduce $\Delta_{\sigma/c_k}$ enough to bring the total into class II for low-reheating scenarios ($T_{\rm reh} \lesssim T_{\rm form}$, MD layer 1 floor); for standard high-reheating all inflationary classes remain in class III or beyond (Sec. III D).

The contrast within the PBH paradigm is striking: domain-wall and gravity-fixed PBH mechanisms sit in class I alongside the most natural particle dark matter candidates, while single-field inflationary collapse sits in class III or beyond. Both the lowest and the highest fine-tuning in the entire dark matter landscape belong to PBH dark matter.

### C. The single-exponential universality identity

The coincidence among class II scenarios can be stated precisely as a theorem.

*Theorem.* Consider any dark matter construction whose relic abundance takes the form

$$\Omega h^2 = \mathcal{A}(\{x_j\}_{j\neq i})e^{-Ax_i} \quad (67)$$

for some parameter $x_i$, with $\mathcal{A}$ a function of the remaining parameters that does not vanish on the $\Omega=\Omega_{\rm DM}$ surface. Then the BG measure with respect to $x_i$, evaluated on the $\Omega h^2 = \Omega_{\rm DM}h^2$ contour, satisfies

$$\Delta_{x_i}|_{\Omega=\Omega_{\rm DM}} = Ax_i = \ln\left[\frac{\mathcal{A}}{\Omega_{\rm DM}h^2}\right] \equiv \ln\left[\frac{\Omega_{\rm natural}}{\Omega_{\rm DM}h^2}\right], \quad (68)$$

where $\Omega_{\rm natural} \equiv \mathcal{A}$ is the prefactor that the exponential must suppress to achieve $\Omega_{\rm DM}$; it is not a physically realizable relic density [for FOPT-PBH, $\mathcal{A}=T_*/(0.84T_{\rm eq}) \sim 10^{14}$–$10^{20}$], but rather the "unquenched" abundance that would result if no exponential suppression were present.

*Proof.* From Eq. (67), $\partial \ln\Omega/\partial \ln x_i = -Ax_i$, so $\Delta_{x_i} = Ax_i$. On the $\Omega=\Omega_{\rm DM}$ contour, $e^{-Ax_i} = \Omega_{\rm DM}/\mathcal{A}$, hence $Ax_i = \ln(\mathcal{A}/\Omega_{\rm DM})$. ∎

Applying it to the class II scenarios within the single-exponential approximation Eq. (30):

*FOPT-PBH.* On $f_{\rm PBH}=1$: $\Delta_{\beta/H} = \ln[T_*/(0.84T_{\rm eq})] \approx$ 30–37, reproducing Eq. (37) exactly. The more accurate superexponential formula (Sec. VI B) gives $\Delta \approx 265$ at the benchmark, so the identity holds only within the approximation and the class II placement of FOPT-PBH is not robust.

*WIMP coannihilation.* On $\Omega=\Omega_{\rm DM}$: $\Delta_{M_{\rm NLSP}} \approx x_F M_{\rm NLSP}/M_\chi = 48$ at benchmark B3 [Eq. (54)].

*PBH early matter domination.* The exponential argument $0.147/\sigma^{4/3}$ evaluates to ≈9–19 across the natural island. Multiplying by 4/3 gives the contribution $0.196/\sigma^{4/3} \approx$ 12–25, to which the prefactor term +5 is added, yielding the full $\Delta_\sigma = 5 + 0.196/\sigma^{4/3} \approx$ 12–21 [Eq. (28)]. The approximate pure-exponential value $Ax \approx$ 9–19 is smaller than the full measure because the $\sigma^5$ prefactor in Eq. (23) contributes $+\beta = +5$, a 20%–35% correction at the natural-island benchmark that accounts for the difference.

Numerically,

$$\ln\left[\frac{T_{\rm form}}{0.84T_{\rm eq}}\right] \approx 2.3\log_{10}\left[\frac{T_{\rm form}}{1\ {\rm eV}}\right] \approx 30-37 \quad (69)$$

for $T_{\rm form} \in [10^4, 10^7]$ GeV, and the WIMP coannihilation result $x_F M_{\rm NLSP}/M_\chi$ falls in the same range. Both are set by the same cosmological clock: the ratio of the formation/freeze-out scale to $T_{\rm eq}$. Note that inflationary collapse, while sharing layer 1 of this identity ($\Delta_\sigma \approx 14$ from HYKN), does not satisfy Eq. (68) as a whole: its layer 2 adds an additional large exponential that takes the total well outside the class II range.

### D. Sensitivity to parameter choice

The BG measure is not invariant under reparametrizations, so the values in Table II depend on the conventions of Sec. II B. The most consequential choice is coannihilation benchmark B3.

Parametrizing by $(M_\chi, \delta)$ with $\delta = (M_{\rm NLSP}-M_\chi)/M_\chi$ gives $\Delta_\delta = x_F\delta \approx 1.3$, misleadingly low because $\delta$ vanishes at the threshold by construction. Parametrizing by the two independent soft masses $(M_\chi, M_{\rm NLSP})$ gives the correct $\Delta_{M_{\rm NLSP}} = 48$ [Eq. (54)]. The convention is unambiguous: fundamental parameters are those independently specified at the UV scale. In the MSSM, $M_\chi = |M_1|$ and $M_{\rm NLSP}$ are set by independent soft terms; $\delta$ is not a Lagrangian parameter. Treating $g_*$ as fixed introduces $\Delta_{g_*} \lesssim 0.5$, well below the dominant contributions.

### E. Prior probability vs local sensitivity

The BG measure quantifies local sensitivity: how precisely must each parameter be specified to maintain $\Omega_{\rm DM}$? This differs from the prior-volume question: what fraction of randomly chosen parameter values satisfies $\Omega h^2 \in [0.06, 0.24]$?

The misalignment axion illustrates this most clearly. Benchmark B6 ($f_a = 10^{12}$ GeV, $\theta_i \approx 0.787$) and B6$_{\rm anthr}$ ($f_a = 10^{16}$ GeV, $\theta_i \approx 3\times10^{-3}$) both have local $\Delta \approx$ 2.0–2.2, since $\Omega h^2 \propto \theta_i^2$ has the same power-law structure regardless of benchmark. However, $\theta_i = 3\times10^{-3}$ is a factor of ∼300 below the natural $\mathcal{O}(1)$ value; on a flat prior $[-\pi, \pi]$ the probability of landing there by chance is $\sim 10^{-3}$. This prior tuning of ∼300 is in addition to, and qualitatively different from, the local BG measure of 2.

The same distinction is central to interpreting inflationary PBH collapse. The BG layer 2 measure $\Delta_{c_k} \gg 10^2$ [Eq. (66)] is a local sensitivity statement: how much does $f_{\rm PBH}$ change when $c_k$ is varied fractionally? The Wilson-normalized measure of Iovino and Riotto [22], $\gamma \equiv \Delta/\bar{\Delta}$ (Sec. II A), is instead a normalized local sensitivity, not a statement about the prior probability or phase-space volume of UV completions, and should not be conflated with the prior-volume question above, which concerns the probability, under an assumed flat prior on $\theta_i$, that the observed value is as small as it is. It asks a related but distinct question: does a technically natural UV completion exist in which $c_k$ takes its current value? Both are physically meaningful. This paper focuses on the local BG question throughout, which is why the inflationary entry carries a $\gtrsim$ rather than a precise value, and why B6 and $\text{B6}_{\rm anthr}$ receive the same tier classification.

## VI. DISCUSSION

### A. Implications for the primordial black hole paradigm

The central message for the PBH community is that the question "are PBHs fine-tuned?" does not have a single answer: it crucially depends on the formation mechanism. Among the three mechanisms analyzed in detail here, biased domain walls exhibit by far the lowest sensitivity ($\Delta = 2$–$4.5$) within the $(\eta, V_b)$ parameter space, and subject to the caveats of Sec. VI B regarding the origin of $V_b$ and the calibration coefficients; early matter domination occupies an intermediate class II tier ($\Delta \approx 14$, comparable to the coannihilation WIMP level), and FOPT-PBH, evaluated with the more accurate super-exponential collapse formula [46], gives $\Delta \approx 265$ (class III); the single-exponential approximation Eq. (30) gives the smaller, illustrative value $\Delta \approx 35$ (see Sec. VI B). When the six classes of inflationary PBH model analyzed in Sec. III D 1 are included, the picture depends on the reheating history. For standard high-reheating scenarios ($T_{\rm reh} \gg T_{\rm form}$), the tuning ranges from $\Delta \approx 115$–$210$ for curvaton and spectator-field models (class III, RD layer 1 floor), through $\Delta \sim 10^2$–$10^3$ for step and resonance features, to $\Delta \approx 6 \times 10^3$–$7 \times 10^4$ (optimistic) or $\Delta \gtrsim 10^{10}$ (pessimistic) for single-field inflection-point and ultra-slow-roll models. For low-reheating scenarios ($T_{\rm reh} \lesssim T_{\rm form}$) where the inflaton oscillation era is still active at reentry, the MD layer 1 floor applies and curvaton/spectator models reduce to $\Delta \approx 21$–$42$ (class II), with all other classes shifting down by a factor of $\sim$5 in their layer 1 contribution; for model classes with large layer 2 sensitivity this shift is a minor correction to the total. Comparing the gravity-fixed domain-wall construction at $\Delta = 2$ with the pessimistic single-field USR scenario at $\Delta \gg 10^4$, the full span of the PBH sensitivity landscape covers more than seven orders of magnitude, a larger range than all particle dark matter scenarios in this paper combined; as above, this compares the extremes of distinct formation mechanisms rather than a like-for-like scan.

The domain-wall result deserves particular emphasis because it is least expected. Domain walls are sometimes described as a "cosmological disaster" [55,65] requiring special engineering to avoid; the biased-wall construction turns this around, exploiting the collapse of the wall network to produce PBHs rather than treating it as a problem to be engineered away. Yet this analysis shows that the resulting abundance map contains *no* exponential: the dependence on $\eta$ and $V_b$ is purely power-law because the Hubble timescale at annihilation (which sets $\beta_{\rm PBH}$) and the PBH mass (which sets the dilution from formation to equality) conspire to cancel all logarithmic factors. This cancellation is not an accident or a tuning: it follows algebraically from the scalings $T_{\rm ann} \propto V_b^2/\eta^{3/2}$ and $M_{\rm PBH} \propto \eta^3/V_b^4$. The gravity-induced-bias variant ($V_{\rm bias} = \eta^5/M_{\rm Pl}$) reduces the domain-wall construction to a single-parameter family with $\Delta = 2$. Among all dark matter production mechanisms, this matches the minimum achievable tuning short of the pure linear cases (ADM and postinflationary axion), and it does so with a single technically natural small parameter: the ratio $V_{\rm bias}/\eta^4 \sim (\eta/M_{\rm Pl})$, which is small because $\eta \ll M_{\rm Pl}$, not because of any special choice. I emphasize that this class I assignment rests on the power-law abundance map of Eq. (17), in which the network-collapse fraction is treated as a fixed input. The dedicated simulations of Ref. [47] find instead an exponential sensitivity of the surviving false-vacuum fraction to the network-collapse temperature, which, together with the large remaining uncertainty in the PBH abundance, is the dominant caveat to the natural classification of this mechanism (see Sec. VI B).

The inflationary model comparison in Fig. 2 clarifies a debate that has persisted in the literature. For high-reheating scenarios, the claim that inflationary PBH production requires fine-tuning is accurate across all inflationary model classes: the RD layer 1 floor $\Delta_\sigma^{(\rm RD)} \approx 57$–$70$ places even the most natural inflationary class (curvaton/spectator) in class III. For low-reheating scenarios, curvaton and spectator-field models achieve $\Delta_{\rm tot} \approx 21$–$42$ (class II), comparable to coannihilating WIMPs and the noninflationary early MD mechanism, while single-field USR models remain in class III or beyond regardless of the reheating temperature. In both cases the layer 2 sensitivity $\Delta_{\sigma/c_k}$ spanning five orders of magnitude from curvaton to single-field USR remains the physically informative quantity: the reheating history shifts the entire landscape up or down by a factor of five in the layer 1 contribution, but does not change the relative ordering of model classes.

The quantitative comparison thus shows the biased-domain-wall PBH mechanism to have a robustly low sensitivity, and clarifies the status of inflationary PBH production: the frequently stated objection that "PBH dark matter requires fine-tuning" is correct for inflationary

mechanisms across all model classes under standard high-reheating assumptions, but the degree of sensitivity spans seven orders of magnitude depending on whether the power spectrum is generated by a single-field USR phase or by a separate spectator field. More broadly, the cross-paradigm comparison shows that fine-tuning is a property of the abundance map's analytic structure, not of whether the dark matter is a particle or a gravitational relic: the PBH paradigm alone spans all three naturalness tiers identified in this paper. Distinguishing between model classes requires precisely the framework developed here: a common measure, a common observable target, and a systematic comparison across production paradigms.

A particularly notable implication of the FOPT-PBH analysis is its multimessenger character. The same nucleation-rate parameter $\beta/H_* \sim 5$–$15$ that places the PBH abundance in the asteroid-mass window also lies within the sensitivity reach of LISA and next-generation ground-based gravitational-wave detectors for the stochastic background sourced by the phase transition itself [77,85]. This is not a coincidence engineered by model construction: it follows directly from the fine-tuning identity $S(\alpha)\beta/H_* = \ln[T_*/(0.84T_{\rm eq})] \approx 30$–$37$ [with benchmark value $\approx 35$ at $T_* = 10^6$ GeV; Eq. (37)], which fixes $\beta/H_*$ to a window that is simultaneously relevant for PBH overproduction and for gravitational-wave detection. The more accurate super-exponential collapse probability of Ref. [46] shifts the natural island to the somewhat narrower range $\beta/H_* \sim 4$–$7$ (compared to the broader 5–15 from the single-exponential approximation), but preserves the overlap with the LISA sensitivity window; the multimessenger conclusion is therefore robust, even though the fine-tuning cost is larger than the single-exponential estimate suggests, as discussed in Sec. VI B. LISA therefore has the potential to *simultaneously* constrain or confirm FOPT-PBH dark matter through two independent channels: (i) microlensing and PBH mass-function measurements in the asteroid window, and (ii) the stochastic gravitational-wave background amplitude and spectral shape. This multimessenger leverage is unique among the PBH formation mechanisms analyzed here and makes the FOPT scenario particularly testable in the near future.

A complementary implication emerges on the particle side. The LZ 2024 results [40] eliminate much of the natural (class I) off-resonance WIMP parameter space below $\sim$1 TeV, leaving as viable primarily the Higgs-funnel (class III, $\Delta \approx 6500$) and coannihilation (class II, $\Delta \approx 48$) regions. As direct detection improves, the experimentally surviving WIMP parameter space migrates toward higher tuning tiers—a trend that, if it continues, will place the most viable WIMP models at naturalness levels comparable to or worse than the FOPT-PBH scenario under its more accurate treatment (Sec. VI B).

Note that several of the class I and class II scenarios make near-term testable predictions: benchmark B7 (postinflationary axion) has a model-independent prediction for $f_a$ falsifiable by next-generation cavity haloscopes [104,106,107], and the early matter-domination PBH mechanism predicts a scalar-induced gravitational-wave background detectable by LISA at $f_{\rm GW} \sim 10^{-3} - 10^{-2}$ Hz.

### B. Limitations and caveats

The analysis presented here rests on several approximations and conventions that I enumerate for completeness.

*Local vs global fine-tuning.* The BG measure captures only the local, one-dimensional sensitivity to each parameter in isolation; it does not measure the total volume of the viable parameter region in the full multidimensional space, nor does it account for correlations between parameters. Multiparameter generalizations of BG (including the "electroweak fine-tuning" measure [31] and Bayesian naturalness [108]) exist but require a specification of the prior, which I have deliberately avoided here; Bayesian refinements are left to future work.

*Formation-rate calibration uncertainties.* The domain-wall collapse fraction $p \approx 10^{-3}$ is calibrated to the wall-network annihilation dynamics of Ref. [45] with an order-of-magnitude uncertainty; varying $p$ by a factor of 10 shifts the natural island in $(\eta, V_b)$ without changing $\Delta = 4.5$. A more important caveat concerns the analytic form of the abundance map itself. The dedicated network-collapse simulations of Ref. [47] find that the false-vacuum volume fraction surviving to collapse decays superexponentially in conformal time, $F_{\rm fv} \sim \exp[-(\eta/\eta_{\rm ann})^p]$ with $p \simeq 3$, which translates into a false-vacuum collapse fraction that is exponentially sensitive to the ratio of the annihilation temperature to the collapse temperature, schematically $\beta_{\rm PBH} \sim \exp[-(wT_{\rm ann}/T_{\rm PBH})^p]$ with $p \simeq 2$–$3$ and $w = \mathcal{O}(1)$. This is structurally analogous to the FOPT and inflationary-collapse cases: the smooth power-law $f_{\rm PBH} \propto \eta^{9/2}/V_b^2$ of Eq. (17) is recovered only when the collapse fraction $p$ is treated as a fixed, temperature-independent input, whereas the more complete treatment of Ref. [47] introduces an exponential dependence on the network-collapse temperature that would raise the BG measure substantially and could move the biased-domain-wall mechanism out of class I once this dependence is folded in. Reference [47] stresses, however, that the resulting PBH abundance is presently subject to large uncertainties (dominated by the unknown departure from spherical collapse), so a definitive tier assignment for this mechanism must await a sharper determination of the collapse fraction; I therefore quote the power-law value $\Delta = 4.5$ as the abundance map of Eq. (17) implies it, while flagging that the exponential temperature sensitivity identified in Ref. [47] is the leading source of model uncertainty in this assignment. Reference [47] also notes that, in the corresponding parameter region, the asteroid-mass window in which PBHs constitute all of the dark matter is accompanied by a correlated stochastic gravitational-wave

background in reach of LVK, LISA, and the Einstein Telescope, paralleling the multimessenger character of the FOPT mechanism discussed in Sec. VI A. I emphasize that $\Delta$ quantifies sensitivity only with respect to the parameters explicitly varied, $\eta$ and $V_b$; it does not, and is not intended to, capture whatever physical input or selection is involved in fixing the calibration coefficients ($p$, $c_t$, $g_*$) at their simulation-central values, nor the microphysical origin of the bias term $V_b$ itself (Sec. III A 2). I state this distinction explicitly to avoid the impression that $\Delta = 2$–$4.5$ constitutes a complete accounting of every physical input to the domain-wall mechanism; an analogous accounting of the microphysical origin of the small parameters entering single-field inflationary models is reflected in the larger layer-2 sensitivities of Sec. III D. The HYKN formula Eq. (23) is fitted to numerical simulations for $\sigma \in [0.03, 0.2]$; extrapolation outside this range is uncertain, but the natural island lies comfortably within it. Non-Gaussian corrections to the curvature perturbation statistics can shift $\beta_{\rm AM}$ by orders of magnitude and move the natural island in $\sigma$ [34], modifying the precise value of $\Delta_\sigma$ without changing its single-exponential class membership, since the collapse probability retains an exponential tail for any distribution with a well-defined variance. The natural island in Fig. 1(b) is also conservative in that it does not include the projected LISA constraint from the associated scalar-induced gravitational-wave background [75,76]; including this would disfavor $\sigma \gtrsim 0.10$ across the asteroid-mass window. For benchmark B2, the thermally averaged Breit-Wigner peak enhancement $R_{\rm max} \approx 10^4$ is accurate to a factor of a few [29,97]; varying it by a factor of 10 changes $\Delta^{(\rm B2)}$ by $\lesssim$30%, which does not affect the class III classification.

*Layer 1 floor: Radiation-dominated vs matter-dominated collapse.* As discussed in Sec. III D, the layer 1 BG sensitivity $\Delta_\sigma$ for inflationary PBH models depends on whether collapse occurs in a radiation-dominated [$\Delta_\sigma^{(\rm RD)} \approx$ 57–70] or matter-dominated [$\Delta_\sigma^{(\rm MD)} \approx 14$] background, determined entirely by the reheating temperature relative to the formation scale. The modern numerical treatment of RD collapse [109,110] modifies the precise value of $\delta_c$ but does not change the exponential structure $\beta_{\rm RD} \propto \exp(-\delta_c^2/2\sigma^2)$, leaving the factor-of-five enhancement over the MD result intact; the RD tier assignments (class III throughout) are robust to $\mathcal{O}(20\%)$ uncertainty in $\delta_c$.

*Two-layer fine-tuning structure of FOPT-PBH dark matter.* Treating $\beta/H_*$ as a fundamental parameter gives the single-exponential result $\Delta_{\beta/H} \approx 35$ (class II, Sec. III C). However, the single-exponential form Eq. (30) is itself an approximation. A more complete treatment of the supercooled percolation integral [46], which tracks the full past-light-cone nucleation history, gives a superexponential collapsed fraction

$$\mathcal{P}_{\rm coll} \simeq \exp\left[-a\left(\frac{\beta}{H}\right)^b (1+\delta_c)^{c\beta/H}\right], \tag{70}$$

with fitted parameters $a \simeq 1.024$, $b \simeq 0.692$, $c \simeq 0.883$, valid for $\beta/H \in [3, 8]$ and $\delta_c \in [0.4, 2/3]$. The BG derivative evaluated on the $f_{\rm PBH} = 1$ contour from Eq. (70) is

$$\Delta^{(\rm GV)}_{\beta/H} = a\left(\frac{\beta}{H}\right)^b (1+\delta_c)^{c\beta/H}\left[b + c\frac{\beta}{H}\ln(1+\delta_c)\right], \tag{71}$$

giving at the benchmark $\beta/H_* \approx 8$, $\delta_c = 0.5$:

$$\Delta^{(\rm GV)}_{\beta/H}\Big|_{\beta/H=8} \approx 265, \tag{72}$$

roughly seven times larger than the single-exponential estimate, placing FOPT-PBH in class III at the phenomenological level. The discrepancy arises because the single-exponential approximation underestimates the steepness of the suppression: at $\beta/H \approx 8$ the outer exponent in Eq. (70) equals $a(\beta/H)^b(1+\delta_c)^{c\beta/H} \approx 75$, and the BG derivative must account for the $\beta/H$-dependence of this outer exponent itself.

At a deeper level, $\beta/H_*$ is itself determined by the shape of the scalar potential. A complete analysis [84] reveals a double-exponential structure,

$$f_{\rm PBH} \simeq M \exp[-Q \exp[-S_3(T_p)/T_p]], \tag{73}$$

where $S_3(T_p)/T_p \approx 170$–$176$ is the Euclidean action-to-temperature ratio at its minimum, $Q \sim 10^{77}$, and $M \sim 4 \times 10^7$. This is structurally identical to the layer 1/layer 2 decomposition of inflationary collapse (Sec. III D): layer 2 (potential coefficients $\{c_k\} \to S_3/T \to \beta/H_*$) introduces an additional exponential with BG sensitivity $\Delta_{S_3/T, c_k} \sim \mathcal{O}(10^2)$–$\mathcal{O}(10^3)$, estimated from the result of Ref. [84] that a fractional shift of $\delta\mu_3/\mu_3 \sim 10^{-3}$ in the cubic potential coefficient moves $f_{\rm PBH}$ from $\mathcal{O}(1)$ to $\mathcal{O}(10^{-100})$. The total BG measure at the level of the potential is

$$\Delta^{\rm FOPT}_{c_k} \sim \ln M \times \Delta_{S_3/T, c_k} \sim 17 \times (10^2 - 10^3) \sim 10^3 - 10^4, \tag{74}$$

where $\ln M \approx 17$ is the value of $Qe^{-S_3/T_p}$ on the $f_{\rm PBH} = 1$ contour. In summary, whether evaluated at the phenomenological level [$\Delta \approx 265$ from Eq. (70)] or at the microphysical level [$\Delta \sim 10^3$–$10^4$ from Eq. (74)], the FOPT-PBH scenario lies in class III or beyond. The natural island in the $(T_*, \beta/H_*)$ plane is defined by $f_{\rm PBH} \approx 1$ and its location is unchanged by the choice of formula for $\mathcal{P}_{\rm coll}$; only the fine-tuning cost of maintaining that condition differs, and the multimessenger implications of Sec. VI A remain valid regardless.

*Gauge dependence of the FOPT nucleation rate.* A recent gauge-invariant calculation [111] finds that PBH production from FOPTs may be substantially suppressed relative to the gauge-dependent estimates used in Sec. III C, potentially shifting or eliminating the natural island in the $(T_*, \beta/H_*)$ plane. This gauge-dependence issue has since been independently confirmed by Wang *et al.* [112] and by Ai and Xie [113]; the latter further show that the mechanism remains viable if reheating is sufficiently slow, since the post-transition Universe can enter an early matter-dominated era during which overdensities grow and collapse into PBHs, a scenario that maps onto the class II fine-tuning structure of Sec. III B rather than the single-exponential estimate of Sec. III C.

*Axion string network uncertainty.* The simulation result $f_a \approx (2\text{–}8) \times 10^{10}$ GeV has a factor-of-~4 range from string-network logarithm uncertainty and QCD modeling [43]. This shifts the absolute prediction for $f_a$ but does not affect $\Delta_{f_a} = 1.19$, which is exact and independent of the normalization coefficient $\mathcal{C}$.

*Definition of the natural island.* The factor-of-two band $\Omega h^2 \in [0.06, 0.24]$ is conventional. I have verified that widening the band to a factor of five ([0.024, 0.60]) or narrowing it to 20% ([0.096, 0.144]) does not change the tier assignments in Table II.

*Restricted benchmark set.* This paper analyzes a curated set of twelve scenarios spanning the most widely studied production mechanisms; the two FOPT benchmarks, now parametrized by $\delta_c = 0.05$ and $\delta_c = 0.10$ (replacing the earlier $\alpha$ parametrization), are counted as a single mechanism in Table II and listed separately in Table III to illustrate $\delta_c$-sensitivity and the effect of the collapse threshold on the required $\beta/H_*$. Notable omissions include: the MSSM $A$-funnel (class III, $\Delta \sim 10^3$–$10^4$, similar to B2); gravitational production of ultralight dark matter from de Sitter fluctuations (class I, $\Delta \approx 2$, since $\Omega \propto H_{\rm inf}^2/M_{\rm DM}^2$); Hawking evaporation seeding a secondary dark matter component; and clockwork axion constructions. These are straightforward to add within the framework developed here and are left for future work.

## VII. CONCLUSIONS

Primordial black holes in the asteroid-mass window ($10^{17}$–$10^{22}$ g) are a viable dark matter candidate, and for some formation mechanisms exhibit a parameter sensitivity far lower than is often assumed. I have quantified this statement by applying the Barbieri-Giudice sensitivity measure $\Delta$ uniformly to twelve dark matter scenarios spanning noninflationary and inflationary PBH formation mechanisms, thermal-relic WIMPs, freeze-in, asymmetric dark matter, and axions. The principal conclusions are as follows.

(1) *Fine-tuning is a property of the abundance map, not of the dark matter candidate.* The BG hierarchy in Table II contains representatives from every paradigm at every tier: the most natural scenario (ADM, $\Delta = 1$) and the second most natural (postinflationary axion, $\Delta = 1.19$) are particle dark matter; class I also includes both a WIMP and a gravitational relic (PBH-DW) at $\Delta = 2$; and the most tuned precisely computed scenario (Higgs funnel, $\Delta = 6500$) is a thermal relic, not a PBH. The narrative that "PBHs are tuned while WIMPs are natural" is not supported by a quantitative comparison.

(2) *Three universality classes classify the dark matter landscape.* Class I ($\Delta \lesssim 5$) contains power-law abundance maps with no exponential factor; class II ($\Delta \approx 15$–$50$) contains constructions with a single exponential; class III and beyond ($\Delta \gtrsim 10^3$) contains resonance, cancellation, and double-exponential constructions. PBH dark matter appears in all three classes, spanning more than seven orders of magnitude in sensitivity within the PBH paradigm alone when the gravity-fixed domain-wall construction and the pessimistic single-field USR estimate are compared—a larger range than all particle dark matter scenarios in this paper combined, though this compares the extremes of distinct formation mechanisms rather than a like-for-like scan.

(3) *Biased-domain-wall PBHs exhibit the lowest Barbieri-Giudice sensitivity of any construction considered here short of the pure linear cases.* With $\Delta = 4.5$ in the free-$V_b$ case and $\Delta = 2$ in the gravity-induced-bias case, domain-wall PBHs occupy class I alongside off-resonance WIMPs and FIMPs; as discussed in Sec. VI B, this statement is restricted to the $(\eta, V_b)$ parameter space and does not, by itself, address the physical origin of $V_b$ or the calibration coefficients entering the collapse fraction.

(4) *Single-exponential constructions satisfy a universal identity, regardless of the underlying physics.* Any construction whose abundance is set by a single exponential satisfies, on the $\Omega = \Omega_{\rm DM}$ contour, $\Delta = \ln(\Omega_{\rm natural}/\Omega_{\rm DM}) \approx \ln(T_{\rm form}/T_{\rm eq}) \approx 14$–$50$ [Eq. (68)]. Within the single-exponential approximation, this unifies FOPT-PBH, early matter-domination PBH, and coannihilating WIMP dark matter into class II despite spanning modulus decay, hidden-sector percolation, and electroweak-scale MSSM thermodynamics. The class II assignment for FOPT-PBH is however conditional on the single-exponential approximation: the more accurate super-exponential collapse probability of Ref. [46] gives $\Delta_{\beta/H} \approx 265$ at the benchmark (class III), and tracing the abundance to the underlying scalar potential gives $\Delta \sim 10^3$–$10^4$ (Sec. VI B). Curvaton inflationary PBH models join class II only for low-reheating scenarios ($T_{\rm reh} \lesssim T_{\rm form}$, $\Delta_{\rm tot} \approx 21$–$42$); for standard high-reheating they shift to class III.

(5) *The layer 1 floor for inflationary PBH production is set by the reheating temperature.* For standard high-reheating ($T_{\rm reh} \gg T_{\rm form}$), all inflationary model classes fall in class III or beyond, with the Press-Schechter floor $\Delta_\sigma^{(\rm RD)} \approx 57$–$70$ applying universally. For low-reheating ($T_{\rm reh} \lesssim T_{\rm form}$), the HYKN floor $\Delta_\sigma^{(\rm MD)} \approx 14$ reduces the total tuning by a factor of $\sim$5, bringing curvaton and spectator-field models into class II. Single-field USR models remain in class III or beyond regardless of the reheating temperature. Joint constraints on $T_{\rm reh}$ and the inflationary power spectrum are therefore necessary to assign a definitive naturalness tier to any inflationary PBH model.

The results above connect two largely separate prior literatures. On the particle side, the MSSM naturalness studies of Grothaus *et al.* [15] and Cabrera *et al.* [17] found that Higgs-funnel and coannihilation scenarios carry the highest relic-density fine-tuning, while bino annihilation via $t$-channel exchange can be "supernatural" [18]; this is fully consistent with benchmarks B1, B2, B3 falling in classes I, III, and II respectively. On the PBH side, Hertzberg and Yamada [19], Cole *et al.* [32], and Stamou [20] established that single-field inflection-point models require tuning to one part in $10^2$–$10^8$; this paper identifies the origin as a layer 2 double-exponential sensitivity $\Delta_{\sigma/c_k} \sim 10^2$–$10^8$, distinct from the layer 1 cost whose value depends on the reheating temperature. The argument of Iovino and Riotto [22] that the bare BG measure is not, by itself, a naturalness statement for inflationary PBH models is not in tension with the results presented here: evaluating $\Delta$ on the $f_{\rm PBH} = 1$ contour removes a related but distinct pathology, the divergence of $\Delta$ as $f_{\rm PBH} \to 0$ off that contour, and the two approaches, bare sensitivity and Wilson-normalized naturalness, answer complementary questions about the same two-layer decomposition of Sec. III D. What neither prior literature contains is a cross-paradigm comparison of the kind presented here, nor any quantitative fine-tuning analysis of the noninflationary PBH mechanisms; the finding that biased-domain-wall PBHs achieve $\Delta = 2$–$4.5$, competitive with the most natural particle dark matter candidates ever identified, is among the central new results of this work.

The present analysis is deliberately local in the sense of the BG measure: it quantifies how sensitively the observed abundance responds to fractional variations of each input parameter, without committing to a prior over the parameter space. A natural and well-defined extension would replace the BG measure with a Bayesian posterior-volume measure [108], allowing direct comparison with the CMSSM and NUHM analyses in the SUSY literature and yielding probabilistic statements about the likelihood of being in the natural island for each construction. The two-layer decomposition adopted here, building on Ref. [32], provides a natural framework for such an analysis: layer 1 and layer 2 priors can be specified independently, reflecting the distinct physical origins of the two sources of tuning.

## ACKNOWLEDGMENTS

I am grateful to Gabriele Franciolini for helpful and deep suggestions, feedback, and discussions. I also thank Miguel Escudero Abenza, David Kaiser, Wenyuan Ai, Christian Byrnes, Ricardo Ferreira, Alessio Notari, Fabrizio Rompineve, and Oriol Pujolàs for useful comments and remarks. This work was supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Award No. DE-SC0010107.

## DATA AVAILABILITY

Python code available upon request.

## APPENDIX A: ANALYTIC DERIVATION OF BG MEASURES

### 1. Biased domain walls

Starting from Eqs. (13)–(16), the dark matter fraction is

$$f_{\rm PBH} = \frac{pV_b^4}{\frac{\pi^2}{30} g_* T_{\rm ann}^4} \cdot \frac{T_{\rm ann}}{0.84 T_{\rm eq}} = \frac{30p}{\pi^2 g_* \cdot 0.84 T_{\rm eq}} \cdot T_{\rm ann}^{-3} V_b^4. \tag{A1}$$

Substituting $T_{\rm ann} = (c_t M_{\rm Pl} V_b^4/\eta^3)^{1/2}$,

$$T_{\rm ann}^{-3} = c_t^{-3/2} M_{\rm Pl}^{-3/2} \eta^{9/2} V_b^{-6}, \tag{A2}$$

so that

$$f_{\rm PBH} \propto \eta^{9/2} V_b^{-6} \cdot V_b^4 = \eta^{9/2} V_b^{-2}. \tag{A3}$$

This is Eq. (17). Log-differentiation gives $\partial \ln f_{\rm PBH}/\partial \ln \eta = 9/2$ and $\partial \ln f_{\rm PBH}/\partial \ln V_b = -2$, hence $\Delta_\eta = 9/2$ and $\Delta_{V_b} = 2$ identically. ∎

### 2. Early matter domination

From $f_{\rm PBH} = \beta_{\rm AM}(\sigma) T_{\rm reh}/(0.84 T_{\rm eq})$ with $T_{\rm reh} \propto m_\phi^{3/2}$,

$$\frac{\partial \ln f_{\rm PBH}}{\partial \ln m_\phi} = \frac{\partial \ln T_{\rm reh}}{\partial \ln m_\phi} = \frac{3}{2}. \tag{A4}$$

For the $\sigma$-dependence, $\ln \beta_{\rm AM} = {\rm const} + 5 \ln \sigma - 0.147 \sigma^{-4/3}$, so

$$\frac{\partial \ln f_{\rm PBH}}{\partial \ln \sigma} = 5 + \frac{\partial}{\partial \ln \sigma}(-0.147\sigma^{-4/3}) = 5 + \frac{4}{3}\cdot 0.147\sigma^{-4/3} = 5 + \frac{0.196}{\sigma^{4/3}}, \tag{A5}$$

which is Eq. (27). ∎

### 3. FOPT late-blooming

From Eq. (32),

$$\frac{\partial \ln f_{\rm PBH}}{\partial \ln(\beta/H_*)} = -S(\alpha)\frac{\beta}{H_*}, \tag{A6}$$

$$\frac{\partial \ln f_{\rm PBH}}{\partial \ln T_*} = 1, \tag{A7}$$

$$\frac{\partial \ln f_{\rm PBH}}{\partial \ln \alpha} = -\frac{\partial S}{\partial \alpha}\cdot\alpha\cdot\frac{\beta}{H_*} = \frac{0.6\alpha}{(\alpha+0.3)^2}\cdot\frac{\beta}{H_*}, \tag{A8}$$

using $\partial S/\partial\alpha = -0.6/(\alpha+0.3)^2$. On the $f_{\rm PBH}=1$ contour, $e^{-S(\alpha)\beta/H_*} = 0.84T_{\rm eq}/T_*$, hence $S(\alpha)\beta/H_* = \ln(T_*/0.84T_{\rm eq})$, giving $\Delta_{\beta/H} = |-S(\alpha)\beta/H_*| = \ln(T_*/0.84T_{\rm eq})$ as in Eq. (37). This derivation assumes the single-exponential Eq. (30); the more accurate super-exponential formula of Ref. [46] gives $\Delta \approx 265$ at the benchmark (Sec. VI B, Eq. (72)). ∎

## APPENDIX B: BENCHMARK CALIBRATION

Table III summarizes the calibration of each benchmark: the parameter values, the resulting $\Omega h^2$, the individual BG sensitivities $\Delta_i$ for each parameter, and the total $\Delta = \max_i \Delta_i$. All quantities are computed numerically using the central-difference formula Eq. (9) with step size $h = 10^{-4}$ and independently verified against the analytic expressions derived in Appendix A and the body of the paper.

## APPENDIX C: FIGURE NOTES

All heatmap figures (Figs. 1–5 and 7) use a shared color map (`viridis_r`) with $\log_{10}\Delta \in [0.0, 3.5]$, saturating to dark purple for $\Delta > 3000$. Parameter-space grids are $400 \times 400$ points, logarithmically spaced. The relic-abundance contours are computed from the same analytic expressions used for the $\Delta$ maps; no separate numerical Boltzmann-code integration is performed. Observational constraints (LZ 2024 direct detection, Lyman-$\alpha$ forest) are implemented as described in Sec. II C and Table I. The Python source codes to generate all figures is available upon request.


[1] G. Jungman, M. Kamionkowski, and K. Griest, Supersymmetric dark matter, Phys. Rep. **267**, 195 (1996).
[2] G. Bertone, D. Hooper, and J. Silk, Particle dark matter: Evidence, candidates and constraints, Phys. Rep. **405**, 279 (2005).
[3] R. D. Peccei and H. R. Quinn, *CP* conservation in the presence of pseudoparticles, Phys. Rev. Lett. **38**, 1440 (1977).
[4] S. Weinberg, A new light boson?, Phys. Rev. Lett. **40**, 223 (1978).
[5] F. Wilczek, Problem of strong P and T invariance in the presence of instantons, Phys. Rev. Lett. **40**, 279 (1978).
[6] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, Freeze-in production of FIMP dark matter, J. High Energy Phys. 03 (2010) 080.
[7] J. McDonald, Thermally generated gauge singlet scalars as self-interacting dark matter, Phys. Rev. Lett. **88**, 091304 (2002).
[8] S. Nussinov, Technocosmology: Could a technibaryon excess provide a "natural" missing mass candidate?, Phys. Lett. B **165**, 55 (1985).
[9] D. E. Kaplan, M. A. Luty, and K. M. Zurek, Asymmetric dark matter, Phys. Rev. D **79**, 115016 (2009).
[10] B. J. Carr and S. W. Hawking, Black holes in the early universe, Mon. Not. R. Astron. Soc. **168**, 399 (1974).
[11] A. M. Green and A. R. Liddle, Constraints on the density perturbation spectrum from primordial black holes, Phys. Rev. D **56**, 6166 (1997).
[12] B. Carr and F. Kühnel, Primordial black holes as dark matter: Recent developments, Annu. Rev. Nucl. Part. Sci. **70**, 355 (2020).
[13] A. M. Green and B. J. Kavanagh, Primordial black holes as a dark matter candidate, J. Phys. G **48**, 043001 (2021).
[14] A. Escrivà, F. Kuhnel, and Y. Tada, Primordial black holes, *Black Holes in the Era of Gravitational-Wave Astronomy* (Elsevier, New York, 2024).
[15] P. Grothaus, M. Lindner, and Y. Takanishi, Naturalness of neutralino dark matter, J. High Energy Phys. 07 (2012) 094.
[16] R. Barbieri and G. F. Giudice, Upper bounds on supersymmetric particle masses, Nucl. Phys. **B306**, 63 (1988).
[17] M. E. Cabrera, J. A. Casas, A. Delgado, S. Robles, and R. Ruiz de Austri, Naturalness of MSSM dark matter, J. High Energy Phys. 08 (2016) 058.
[18] S. F. King and J. P. Roberts, Natural implementation of neutralino dark matter, J. High Energy Phys. 09 (2006) 036.
[19] M. P. Hertzberg and M. Yamada, Primordial black holes from polynomial potentials in single field inflation, Phys. Rev. D **97**, 083509 (2018).

[20] I. Stamou, Mechanisms for producing primordial black holes from inflationary models beyond fine-tuning, Universe **10**, 241 (2024).
[21] A. Kalaja, N. Bellomo, N. Bartolo, D. Bertacca, S. Matarrese, I. Musco, A. Raccanelli, and L. Verde, From primordial black holes abundance to primordial curvature power spectrum (and back), J. Cosmol. Astropart. Phys. 10 (2019) 031.
[22] A. J. Iovino and A. Riotto, Are primordial black holes truly fine-tuned?, J. Cosmol. Astropart. Phys. 07 (2026) 047.
[23] J. R. Ellis, K. Enqvist, D. V. Nanopoulos, and F. Zwirner, Observables in low-energy superstring models, Mod. Phys. Lett. A **01**, 57 (1986).
[24] G. W. Anderson and D. J. Castano, Measures of fine tuning, Phys. Lett. B **347**, 300 (1995).
[25] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, Astron. Astrophys. **641**, A6 (2020); **652**, C4(E) (2021).
[26] A. Strumia, Naturalness of supersymmetric models, in 34th Rencontres de Moriond: Electroweak Interactions and Unified Theories (1999), pp. 441–446, arXiv:hep-ph/9904247.
[27] M. E. Cabrera, J. A. Casas, and R. Ruiz de Austri, Bayesian approach and naturalness in MSSM analyses for the LHC, J. High Energy Phys. 03 (2009) 075.
[28] S. P. Martin, A supersymmetry primer, Adv. Ser. Dir. High Energy Phys. **18**, 1 (1998).
[29] K. Griest and D. Seckel, Three exceptions in the calculation of relic abundances, Phys. Rev. D **43**, 3191 (1991).
[30] J. Edsjo and P. Gondolo, Neutralino relic density including coannihilations, Phys. Rev. D **56**, 1879 (1997).
[31] H. Baer, V. Barger, and A. Mustafayev, Implications of a 125 GeV Higgs scalar for LHC SUSY and neutralino dark matter searches, Phys. Rev. D **85**, 075010 (2012).
[32] P. S. Cole, A. D. Gow, C. T. Byrnes, and S. P. Patil, Primordial black holes from single-field inflation: A fine-tuning audit, J. Cosmol. Astropart. Phys. 08 (2023) 031.
[33] K. Kohri, D. H. Lyth, and A. Melchiorri, Black hole formation and slow-roll inflation, J. Cosmol. Astropart. Phys. 04 (2007) 038.
[34] E. V. Bugaev and P. A. Klimai, Primordial black hole constraints for curvaton models with predicted large non-Gaussianity, Int. J. Mod. Phys. D **22**, 1350034 (2013).
[35] V. Vennin, Stochastic inflation and primordial black holes, arXiv:2009.08715.
[36] S. Clesse and J. García-Bellido, Massive primordial black holes from hybrid inflation as dark matter and the seeds of galaxies, Phys. Rev. D **92**, 023524 (2015).
[37] G. Ballesteros and M. Taoso, Primordial black hole dark matter from single field inflation, Phys. Rev. D **97**, 023501 (2018).
[38] H. Motohashi and W. Hu, Primordial black holes and slow-roll violation, Phys. Rev. D **96**, 063503 (2017).
[39] K. Kannike, L. Marzola, M. Raidal, and H. Veermäe, Single field double inflation and primordial black holes, J. Cosmol. Astropart. Phys. 09 (2017) 020.
[40] J. Aalbers *et al.* (LZ Collaboration), Dark matter search results from 4.2 tonne – years of exposure of the LUX-ZEPLIN (LZ) experiment, Phys. Rev. Lett. **135**, 011802 (2025).
[41] G. Ballesteros, M. A. G. Garcia, and M. Pierre, How warm are non-thermal relics? Lyman-$\alpha$ bounds on out-of-equilibrium dark matter, J. Cosmol. Astropart. Phys. 03 (2021) 101.
[42] Q. Decant, J. Heisig, D. C. Hooper, and L. Lopez-Honorez, Lyman-$\alpha$ constraints on freeze-in and superWIMPs, J. Cosmol. Astropart. Phys. 03 (2022) 041.
[43] M. Buschmann, J. W. Foster, A. Hook, A. Peterson, D. E. Willcox, W. Zhang, and B. R. Safdi, Dark matter from axion strings with adaptive mesh refinement, Nat. Commun. **13**, 1049 (2022).
[44] J. M. Alarcon, J. Martin Camalich, and J. A. Oller, The chiral representation of the $\pi N$ scattering amplitude and the pion-nucleon sigma term, Phys. Rev. D **85**, 051503 (2012).
[45] F. Ferrer, E. Massò, G. Panico, O. Pujolas, and F. Rompineve, Primordial black holes from the QCD axion, Phys. Rev. Lett. **122**, 101301 (2019).
[46] Y. Gouttenoire and E. Vitagliano, Primordial black holes and wormholes from domain wall networks, Phys. Rev. D **109**, 123507 (2024).
[47] R. Z. Ferreira, A. Notari, O. Pujolàs, and F. Rompineve, Collapsing domain wall networks: Impact on pulsar timing arrays and primordial black holes, J. Cosmol. Astropart. Phys. 06 (2024) 020.
[48] J. Georg and S. Watson, A preferred mass range for primordial black hole formation and black holes as dark matter revisited, J. High Energy Phys. 09 (2016) 138.
[49] J. Georg, B. Melcher, and S. Watson, Primordial black holes and co-decaying dark matter, J. Cosmol. Astropart. Phys. 11 (2019) 014.
[50] I. Dalianis and C. Kouvaris, Gravitational waves from density perturbations in an early matter domination era, J. Cosmol. Astropart. Phys. 07 (2021) 046.
[51] B. Carr, K. Kohri, Y. Sendouda, and J. Yokoyama, Constraints on primordial black holes, Rep. Prog. Phys. **84**, 116902 (2021).
[52] J. Liu, L. Bian, R.-G. Cai, Z.-K. Guo, and S.-J. Wang, Primordial black hole production during first-order phase transitions, Phys. Rev. D **105**, L021303 (2022).
[53] M. Lewicki, P. Toczek, and V. Vaskonen, Primordial black holes from strong first-order phase transitions, J. High Energy Phys. 09 (2023) 092.
[54] T. C. Gehrman, B. Shams Es Haghi, K. Sinha, and T. Xu, Baryogenesis, primordial black holes and MHz–GHz gravitational waves, J. Cosmol. Astropart. Phys. 02 (2023) 062.
[55] A. Vilenkin and E. P. S. Shellard, *Cosmic Strings and Other Topological Defects* (Cambridge University Press, Cambridge, England, 2000).
[56] G. Franciolini and A. Urbano, Primordial black hole dark matter from inflation: The reverse engineering approach, Phys. Rev. D **106**, 123519 (2022).
[57] O. Özsoy and G. Tasinato, Inflation and primordial black holes, Universe **9**, 203 (2023).
[58] J. García-Bellido and E. Ruiz Morales, Primordial black holes from single field models of inflation, Phys. Dark Universe **18**, 47 (2017).

[59] C. Germani and T. Prokopec, On primordial black holes from an inflection point, Phys. Dark Universe **18**, 6 (2017).
[60] S. R. Geller, W. Qin, E. McDonough, and D. I. Kaiser, Primordial black holes from multifield inflation with nonminimal couplings, Phys. Rev. D **106**, 063535 (2022).
[61] W. Qin, S. R. Geller, S. Balaji, E. McDonough, and D. I. Kaiser, Planck constraints and gravitational wave forecasts for primordial black hole dark matter seeded by multifield inflation, Phys. Rev. D **108**, 043508 (2023).
[62] D. L. Lorenzoni, S. R. Geller, Z. Ireland, D. I. Kaiser, E. McDonough, and K. A. Wittmeier, Light scalar fields Foster production of primordial black holes, arXiv:2504.13251.
[63] D. L. Lorenzoni, S. R. Geller, D. I. Kaiser, and E. McDonough, Primordial black holes from inflation with a spectator field, arXiv:2512.04199.
[64] J. R. Espinosa, D. Racco, and A. Riotto, Cosmological signature of the standard model Higgs vacuum instability: Primordial black holes as dark matter, Phys. Rev. Lett. **120**, 121301 (2018).
[65] T. W. B. Kibble, Topology of cosmic domains and strings, J. Phys. A **9**, 1387 (1976).
[66] S. M. Barr and D. Seckel, Planck-scale corrections to axion models, Phys. Rev. D **46**, 539 (1992).
[67] R. Holman, S. D. H. Hsu, T. W. Kephart, E. W. Kolb, R. Watkins, and L. M. Widrow, Solutions to the strong *CP* problem in a world with gravity, Phys. Lett. B **282**, 132 (1992).
[68] M. Kawasaki, K. Saikawa, and T. Sekiguchi, Axion dark matter from topological defects, Phys. Rev. D **91**, 065014 (2015).
[69] K. J. Bae, H. Baer, V. Barger, and R. W. Deal, The cosmological moduli problem and naturalness, J. High Energy Phys. 02 (2008) 138.
[70] R. Allahverdi *et al.*, The first three seconds: A review of possible expansion histories of the early universe, Open J. Astrophys. **4**, 1 (2021).
[71] M. Y. Khlopov and A. G. Polnarev, Primordial black holes as a cosmological test of grand unification, Phys. Lett. **97**, 383 (1980).
[72] T. Harada, C.-M. Yoo, K. Kohri, and T. Nakama, Primordial black hole formation in the matter-dominated phase of the universe, Phys. Rev. D **96**, 083517 (2017).
[73] W. H. Press and P. Schechter, Formation of galaxies and clusters of galaxies by self-similar gravitational condensation, Astrophys. J. **187**, 425 (1974).
[74] Z. Zhou, J. Jiang, Y.-F. Cai, M. Sasaki, and S. Pi, Primordial black holes and gravitational waves from resonant amplification during inflation, Phys. Rev. D **102**, 103527 (2020).
[75] K. N. Ananda, C. Clarkson, and D. Wands, The cosmological gravitational wave background from primordial density perturbations, Phys. Rev. D **75**, 123518 (2007).
[76] D. Baumann, P. J. Steinhardt, K. Takahashi, and K. Ichiki, Gravitational wave spectrum induced by primordial scalar perturbations, Phys. Rev. D **76**, 084019 (2007).
[77] C. Caprini *et al.*, Detecting gravitational waves from cosmological phase transitions with LISA: An update, J. Cosmol. Astropart. Phys. 03 (2020) 024.
[78] A. D. Linde, Fate of the false vacuum at finite temperature: Theory and applications, Phys. Lett. **100B**, 37 (1981).
[79] E. Witten, Cosmic separation of phases, Phys. Rev. D **30**, 272 (1984).
[80] J. R. Espinosa, T. Konstandin, J. M. No, and G. Servant, Energy budget of cosmological first-order phase transitions, J. Cosmol. Astropart. Phys. 06 (2010) 028.
[81] S. W. Hawking, I. G. Moss, and J. M. Stewart, Bubble collisions in the very early universe, Phys. Rev. D **26**, 2681 (1982).
[82] I. G. Moss, Singularity formation from colliding bubbles, Phys. Rev. D **50**, 676 (1994).
[83] H. Kodama, M. Sasaki, and K. Sato, Abundance of primordial holes produced by cosmological first-order phase transition, Prog. Theor. Phys. **68**, 1979 (1982).
[84] Y. Wu and S. Profumo, Superexponential primordial black hole production via delayed vacuum decay, Phys. Rev. D **111**, 103524 (2025).
[85] M. Hindmarsh, S. J. Huber, K. Rummukainen, and D. J. Weir, Shape of the acoustic gravitational wave power spectrum from a first order phase transition, Phys. Rev. D **96**, 103520 (2017).
[86] D. H. Lyth and D. Wands, Generating the curvature perturbation without an inflaton, Phys. Lett. B **524**, 5 (2002).
[87] K. Enqvist and M. S. Sloth, Adiabatic CMB perturbations in pre-big bang string cosmology, Nucl. Phys. **B626**, 395 (2002).
[88] A. A. Starobinsky and J. Yokoyama, Equilibrium state and density fluctuations in de Sitter space with and without $\lambda\phi^4$ self-interaction, Phys. Rev. D **50**, 6357 (1994).
[89] C. Pattison, V. Vennin, H. Assadullahi, and D. Wands, Quantum diffusion during inflation and primordial black holes, J. Cosmol. Astropart. Phys. 10 (2021) 046.
[90] J. García-Bellido, A. D. Linde, and D. Wands, Density perturbations and black hole formation in hybrid inflation, Phys. Rev. D **54**, 6040 (1996).
[91] Y.-F. Cai, X. Chen, M. H. Namjoo, M. Sasaki, D.-G. Wang, and Z. Wang, Revisiting non-Gaussianity from non-attractor inflation models, J. Cosmol. Astropart. Phys. 05 (2018) 012.
[92] G. Ballesteros, J. Rey, M. Taoso, and A. Urbano, Primordial black holes as dark matter and gravitational waves from single-field polynomial inflation, J. Cosmol. Astropart. Phys. 07 (2020) 025.
[93] E. W. Kolb and M. S. Turner, *The Early Universe* (Addison-Wesley, Redwood City, CA, 1990).
[94] V. Silveira and A. Zee, Scalar phantoms, Phys. Lett. **161B**, 136 (1985).
[95] J. McDonald, Gauge singlet scalars as cold dark matter, Phys. Rev. D **50**, 3637 (1994).
[96] C. P. Burgess, M. Pospelov, and T. ter Veldhuis, The minimal model of nonbaryonic dark matter: A singlet scalar, Nucl. Phys. **B619**, 709 (2001).
[97] M. Ibe, H. Murayama, and T. T. Yanagida, Breit-Wigner enhancement of dark matter annihilation, Phys. Rev. D **79**, 095009 (2009).
[98] L. Visinelli and P. Gondolo, Dark matter axions revisited, Phys. Rev. D **80**, 035024 (2009).

[99] S. Borsanyi *et al.*, Calculation of the axion mass based on high-temperature lattice quantum chromodynamics, Nature (London) **539**, 69 (2016).
[100] R. L. Davis, Goldstone bosons in string models of galaxy formation, Phys. Rev. D **35**, 3705 (1987).
[101] R. A. Battye and E. P. S. Shellard, Axion string constraints, Phys. Rev. Lett. **73**, 2954 (1994).
[102] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi, and O. Straniero, Revisiting the bound on axion-photon coupling from globular clusters, Phys. Rev. Lett. **113**, 191302 (2014).
[103] N. Du *et al.* (ADMX Collaboration), A search for invisible axion dark matter with the axion dark matter experiment, Phys. Rev. Lett. **120**, 151301 (2018).
[104] K. M. Backes *et al.* (HAYSTAC Collaboration), A quantum-enhanced search for dark matter axions, Nature (London) **590**, 238 (2021).
[105] K. Abazajian *et al.* (CMB-S4 Collaboration), CMB-S4 science book, first edition, arXiv:1610.02743.
[106] Y. Kahn, B. R. Safdi, and J. Thaler, Broadband and resonant approaches to axion dark matter detection, Phys. Rev. Lett. **117**, 141801 (2016).
[107] L. Brouwer *et al.* (DMRadio Collaboration), Proposal for a definitive search for GUT-scale QCD axions, Phys. Rev. D **106**, 112003 (2022).
[108] D. M. Ghilencea and G. G. Ross, The fine-tuning cost of the likelihood in SUSY models, Nucl. Phys. **B868**, 65 (2013).
[109] I. Musco, Threshold for primordial black holes: Dependence on the shape of the cosmological perturbations, Phys. Rev. D **100**, 123524 (2019).
[110] A. Escrivà, C. Germani, and R. K. Sheth, Universal threshold for primordial black hole formation, Phys. Rev. D **101**, 044022 (2020).
[111] G. Franciolini, Y. Gouttenoire, and R. Jinno, Curvature perturbations from first-order phase transitions: Implications to black holes and gravitational waves, Phys. Rev. Lett. **136**, 171404 (2026).
[112] X. Wang, C. Balázs, R. Ding, and C. Tian, How large are curvature perturbations from slow first-order phase transitions? A gauge-invariant analysis, Phys. Rev. D **113**, 123026 (2026).
[113] W.-Y. Ai and K.-P. Xie, Reviving primordial black hole formation in slow first-order phase transitions, arXiv:2605.11332.