NON-GAUSSIAN TAILS OF THE CURVATURE PERTURBATION IN ULTRA-SLOW-ROLL INFLATION: STOCHASTIC-δN CONSTRAINTS ON PRIMORDIAL BLACK HOLE ABUNDANCE IN THE ASTEROID-MASS WINDOW

Matheus Silva Dias

University of Sao Paulo
Sao Paulo, Federative Republic of Brazil
E-mail: matheus.sd@usp.br

Original Scientific Article

Received: 22.01.2026. Revised: 11.04.2026. Accepted: 18.05.2026.

DOI: https://doi.org/10.65932/CAR-2026-1-1 **UDC**: 524.882:524.83:524.85

Abstract: The number of primordial black holes (PBHs) formed from inflationary perturbations depends exponentially on the upper tail of the probability distribution of the curvature perturbation ζ, which makes every abundance estimate hostage to the statistical assumptions entering that tail. This article quantifies, within a single analytic framework, how the estimated PBH abundance changes when the standard perturbative Gaussian calculation is replaced by the stochastic-δN treatment of ultra-slow-roll (USR) inflation, in which the tail decays exponentially, $P(\zeta) \propto \exp(-\Lambda\zeta)$, rather than as a Gaussian. Using the logarithmic mapping between ζ and its Gaussian precursor, with benchmark decay rate $\Lambda = 3$ and collapse threshold $\zeta c = 0.65$, I derive a closed-form amplitude-remapping factor $R\_A = [\Lambda\zeta c/(1 - \exp(-\Lambda\zeta c))]^2 \approx 5.2$, which measures how much smaller the coarse-grained variance σ^2 must be for the mass fraction β to stay fixed once exponential tails are switched on. Applied across the asteroid-mass window (10^{17}–10^{23} g), the calculation shows that a Gaussian-calibrated amplitude $\sigma^2 \approx (6.3–7.9)\times10^{-3}$ overproduces PBHs by a factor of 5×10^9 to 1.2×10^{12} when the exponential tail operates at fixed amplitude, while the amplitude required for $f\_PBH = 1$ falls to $\sigma^2 \approx (1.2–1.5)\times10^{-3}$. The local sensitivity of the abundance to the amplitude, $d \ln \beta/d \ln \sigma^2 \approx 31$ at the calibration point, is nonetheless invariant under the change of statistics, so the notorious fine-tuning of PBH dark matter scenarios survives the transition intact. The asteroid-mass window itself remains a viable candidate for the totality of dark matter: its boundaries are set by evaporation and microlensing physics that do not depend on formation statistics, although the accompanying induced gravitational-wave signal weakens by roughly a factor of 27.

Keywords: *primordial black holes, ultra-slow-roll inflation, stochastic inflation, δN formalism, non-Gaussian tails, curvature perturbation, dark matter, asteroid-mass window.*

INTRODUCTION

Half a century separates the first suggestion that black holes formed in the early universe from the present debate about whether they constitute the dark matter, and the suggestion has aged remarkably well (Carr & Hawking, 1974). Most candidate masses have been excluded. What remains is a single contiguous interval, commonly called the asteroid-mass window, spanning roughly 10^{17} to 10^{23} g, in which no current observation forbids primordial black holes (PBHs) from making up all of the dark matter (Carr et al., 2021; Green & Kavanagh, 2021). Careful re-

Cosmological and Astrobiological Review; Journal for the Study of the Universe, Life and the Natural Sciences, Year 2026, Volume 4, Issue 1, Serial number 5, p. 9 – 29.

examination of older bounds widened rather than narrowed this interval: Montero-Camacho et al. (2019) placed the open range at approximately 3.5×10^{-17} to 4×10^{-12} solar masses after revisiting capture and destruction arguments, while Smyth et al. (2020) showed that finite-source effects erase optical microlensing sensitivity near the upper edge. A decade of constraint-building has thus left one door open, and it is a wide one (Carr & Kühnel, 2020).

The window sits where it does for reasons of physics, not of observational neglect. Below it, Hawking radiation betrays the population: a hole lighter than about 5×10^{14} g has evaporated entirely within the age of the universe, and holes up to a few times 10^{17} g radiate photons of order 100 keV and below, brightly enough for the Galactic and extragalactic backgrounds to notice (Carr et al., 2021; Auffinger, 2023). Above it, gravity betrays them instead, because objects heavier than roughly 10^{23} g magnify background stars detectably as they transit the line of sight (Niikura et al., 2019). In between, a PBH is a nearly perfect nonparticipant: too cold to shine, too light to lens, too rare to collide with anything on human timescales (Green & Kavanagh, 2021). Montero-Camacho et al. (2019) closed off the last proposed astrophysical vetoes in this range, from white-dwarf ignition to neutron-star capture, and the survey instruments that might have reached into it succumb to finite-source suppression (Smyth et al., 2020). Six decades of mass with nothing to say against them. That silence is what makes the statistical question of this article consequential: the only handle on the window's occupancy, for now, is the calculation that predicts it.

Producing PBHs in this window requires the primordial curvature power spectrum to grow from its measured large-scale amplitude of 2.1×10^{-9} to roughly 10^{-2} on scales that re-enter the horizon when the horizon mass is of order 10^{17}–10^{23} g, an enhancement of nearly seven orders of magnitude (Karam et al., 2023). Single-field inflation can deliver such growth through a transient ultra-slow-roll (USR) phase, in which the inflaton crosses a very flat stretch of its potential and the would-be decaying mode of the curvature perturbation instead grows (Byrnes et al., 2019). Explicit constructions with polynomial potentials achieve the required amplification while respecting large-scale observations (Ballesteros et al., 2020). The price is steep parametric sensitivity: percent-level shifts in potential parameters move the predicted abundance by orders of magnitude, a problem catalogued in detail by Cole et al. (2023).

The origin of that sensitivity is statistical, not dynamical. The fraction β of horizon patches that collapse into black holes is an integral over the extreme upper tail of the probability distribution of ζ, evaluated seven to nine standard deviations from the mean for the abundances of interest, $\beta \sim 10^{-16}$–10^{-13} (Carr et al., 2021). At such depths the integral is controlled entirely by the functional form of the tail. Nearly all constraint analyses assume that form to be Gaussian, sometimes corrected by the nonlinear relation between ζ and the density contrast (Young et al., 2019) or by the shape dependence of the collapse threshold (Germani & Musco, 2019; Musco, 2019). Yet a Gaussian tail at nine sigma is an extrapolation, not a measurement. Perturbative non-Gaussianity parameters bend this tail appreciably even at modest coupling, as shown for quadratic local models (Atal & Germani, 2019; Ferrante et al., 2023) and in peak theory (Kitajima et al., 2021). Stochastic inflation replaces the extrapolation with a calculation. In this framework, long-wavelength field modes evolve as a random walk sourced by short-wavelength quantum noise, an idea whose equilibrium limit was solved analytically long ago (Starobinsky & Yokoyama, 1994) and whose extension beyond slow roll is now under quantitative control (Pattison et al., 2019; Cruces & Germani, 2022; Cruces, 2022). Combined with the δN correspondence between expansion histories and curvature perturbations, the framework yields the full one-point distribution of ζ as a first-passage-time problem (Pattison et al., 2021; Tada & Vennin, 2022). The generic outcome is striking: the tail of $P(\zeta)$ decays exponentially, $\exp(-\Lambda\zeta)$, no matter how Gaussian the core looks (Ezquiaga et al., 2020). Numerical simulations of USR inflation confirmed the exponential tail and traced its consequences for PBH counts (Figueroa et al., 2021, 2022), and the same structure emerges from an exact logarithmic duality of the long-wavelength equations (Pi & Sasaki, 2023). The tail is not a small correction to a Gaussian. It is a different function.

Doubts about the perturbative pipeline have arrived from a second direction at the same time. Kristiano and Yokoyama (2024a, 2024b) argued that one-loop corrections from small-scale USR modes could contaminate observable scales strongly enough to constrain PBH-producing models, a claim disputed on the grounds that superhorizon perturbations are protected (Inomata, 2024), supported path-integral proofs of loop smallness (Kawaguchi et al., 2024), and separate loop calculations finding no obstruction (Firouzjahi & Riotto, 2024). I take no side in that exchange here. Its relevance for present purposes is diagnostic: both the loop debate and the tail problem signal that perturbation theory is being asked to operate at the edge of its domain precisely where PBH predictions are made. This article addresses a deliberately narrow question sitting at that edge: by how much, quantitatively, does the estimated PBH abundance in the asteroid-mass window change when the perturbative Gaussian computation is replaced by the stochastic-δN computation with its exponential tail, and does the window survive as a candidate for the totality of dark matter once the replacement is made? Three hypotheses organize the analysis. The first (H1) states that at fixed power-spectrum amplitude the change in the predicted mass fraction β exceeds eight orders of magnitude everywhere in the window. The second (H2) states that the correction to the amplitude required for a fixed abundance is nonetheless modest, below one order of magnitude, because the abundance responds exponentially to amplitude in both statistical frameworks. The third (H3) states that the viability of the asteroid-mass window is unaffected by the choice of statistics, since the observational bounds that define its edges act on the present-day PBH population rather than on the formation calculation.

The original contribution of this article lies in a closed-form amplitude-remapping factor, R_A = $[\Lambda\zeta c/(1 - \exp(-\Lambda\zeta c))]^2$, which converts any Gaussian-calibrated variance of the curvature perturbation into its exponential-tail equivalent at fixed PBH abundance, evaluated here for the first time consistently across the entire asteroid-mass window together with the accompanying fixed-amplitude abundance error. A secondary and, to my knowledge, previously unreported observation follows from the same algebra: the local logarithmic sensitivity of β to the variance is exactly invariant under the change from Gaussian to exponential-tail statistics at matched abundance, so quantum diffusion does not relieve the fine-tuning of PBH dark matter even as it reshapes every number in the calibration. The argument proceeds as follows. The next section reviews the literature on which the analysis stands and specifies the methodology, including all assumptions and their known weaknesses. A results section then reports the numerical findings without interpretation. Three analytical sections follow: the first examines why exponential tails are a structural property of quantum diffusion rather than a model artifact, the second interprets the recalibration numbers and their consequences for model building and induced gravitational waves, and the third confronts the asteroid-mass window with the revised statistics. A conclusion returns verdicts on the three hypotheses, states limitations, and identifies what observations could discriminate between the two statistical frameworks.

LITERATURE REVIEW AND METHODOLOGY

Literature Review

The cartography of PBH constraints has been redrawn several times since 2019, and the asteroid-mass window owes its present shape to that redrawing. The encyclopedic review of Carr et al. (2021) assembles evaporation, lensing, dynamical, and accretion bounds and locates the fully open interval between about 10^{17} and 10^{23} g, in agreement with the independent synthesis of Green and Kavanagh (2021). At the light edge, Hawking evaporation dominates: the Galactic 511 keV positron line excludes PBHs lighter than about 10^{17} g from being all of the dark matter (Laha, 2019), a bound sharpened by later analyses of Galactic and extragalactic emission (Korwar & Profumo, 2023) and refined again with improved propagation modelling of MeV photons and electron-positron pairs (De la Torre Luque et al., 2024). Auffinger (2023) reviews the entire evaporation program and its instrumental future. At the heavy edge, the Subaru/HSC survey of M31 provides the strongest lensing limit, excluding f_PBH = 1 above roughly 10^{23} g (Niikura et al., 2019), though the sensitivity claimed below that mass was later shown to dissolve once finite source sizes and wave optics are treated consistently (Smyth et al., 2020). Clustering does not rescue the excluded regions: for the Poissonian initial conditions relevant here, small-scale clustering leaves microlensing constraints essentially unchanged (Gorton & Green, 2022). Adjacent to the window, the OGLE high-cadence survey of the Magellanic Clouds now bounds planetary-mass PBHs at the percent level (Mróz et al., 2024).

A separate strand of work asks how the window could ever be closed, or the hypothesis confirmed. Proposals include direct detection of Hawking radiation from a single nearby PBH with next-generation MeV telescopes (Coogan et al., 2021), statistical detection of the Galactic evaporation glow, which such instruments could push to 10^{18} g (Ray et al., 2021), continuous gravitational waves from light PBH binaries (Miller et al., 2022), and, more speculatively, crater statistics on airless Solar system bodies (Santarelli et al., 2025). On the production side, any PBH population formed from enhanced curvature perturbations is accompanied by scalar-induced gravitational waves whose spectrum tracks the square of the power-spectrum amplitude (Domènech, 2021). The LISA Cosmology Working Group treats this counterpart as a primary science target for the window (Bagui et al., 2025), and dedicated studies show that the non-Gaussianity accompanying USR-type tails modifies both the PBH count and the induced-wave signal (Inui et al., 2025). The importance of non-Gaussianity for the induced-wave interpretation of pulsar-timing data has been demonstrated separately (Franciolini et al., 2023).

Model building for the window has consolidated around transient USR episodes in single-field inflation. Byrnes et al. (2019) established that the power spectrum cannot grow faster than k^4 toward small scales, which constrains how abruptly the enhancement can be engineered; Karam et al. (2023) dissected the anatomy of USR models and the mapping between potential features and spectral peaks; Ballesteros et al. (2020) provided concrete polynomial realizations compatible with CMB observations. Two systematic pathologies emerged from this program. The first is fine-tuning: Cole et al. (2023) quantified how percent-level parameter variations translate into order-of-magnitude abundance shifts. The second is statistical fragility, the subject of this article. Germani and Musco (2019) showed early on that the abundance depends on the full shape of the peaked spectrum rather than its amplitude alone, an observation that already gestured toward the inadequacy of one-number characterizations.

What counts as a collapsing perturbation has itself been renegotiated. Musco (2019) demonstrated that the critical threshold depends on the radial profile of the perturbation, with δc ranging from about 0.41 to 0.66 across realistic shapes, and Musco et al. (2021) compressed that dependence into an analytic prescription accurate to a few percent. Escrivà et al. (2020) found a shape-independent formulation in terms of the averaged compaction function. Young et al. (2019) established that the nonlinear relation between curvature and density perturbations must be retained when converting statistics of ζ into collapse statistics, which by itself alters β at fixed spectrum. More recent work extends the criterion to type II fluctuations with non-Gaussian statistics

(Shimada et al., 2025) and, notably for this article, computes the compaction function directly from stochastic USR fluctuations (Raatikainen et al., 2024). A broader synthesis of formation physics is given by Escrivà et al. (2024). These refinements matter quantitatively, but they act multiplicatively on thresholds while the tail question acts on the distribution being thresholded; the two can, to first approximation, be studied separately, which is the approach I adopt.

Perturbative treatments of non-Gaussianity anticipated much of the qualitative story. Quadratic local corrections skew the tail enough to shift required amplitudes appreciably (Atal & Germani, 2019), and the machinery for propagating local non-Gaussianity of arbitrary order into PBH abundances now exists in closed form (Ferrante et al., 2023). Peak-theory implementations reach compatible conclusions (Kitajima et al., 2021). The limitation shared by these approaches is structural: a truncated expansion in ζ_G reproduces polynomial distortions of a Gaussian but cannot generate a genuinely exponential tail, which is a non-perturbative object (Gow et al., 2023). Sudden-transition models illustrate the difference concretely, with the tail switching between Gaussian-like and exponential regimes depending on a single parameter of the transition (Cai et al., 2022).

The stochastic program supplies the non-perturbative statistics. Its foundations date to the equilibrium de Sitter solution of Starobinsky and Yokoyama (1994), but the modern phase began when the formalism was extended beyond slow roll (Pattison et al., 2019) and its validity conditions were established order by order in slow-roll parameters (Cruces & Germani, 2022; Cruces, 2022). Within the stochastic-δN framework, the curvature perturbation on a given scale is the fluctuation in the number of e-folds needed to exit a region of the potential, a first-passage-time quantity whose full distribution is calculable (Pattison et al., 2021; Tada & Vennin, 2022; Ando & Vennin, 2021). Ezquiaga et al. (2020) proved that this distribution generically possesses an exponential tail whose decay rate is the lowest eigenvalue of the associated boundary problem. Figueroa et al. (2021) simulated stochastic USR inflation and found the exponential tail numerically, with PBH abundances differing from Gaussian estimates by several orders of magnitude; the follow-up quantified the effect across model parameters (Figueroa et al., 2022). Independent numerical machinery has matured in parallel, including importance-sampled simulations that resolve the deep tail (Jackson et al., 2022), noise modelling beyond slow roll (Jackson et al., 2025), and simulations constrained by frozen noise (Tomberg, 2023). Rigopoulos and Wilkins (2021) showed that diffusion-dominated regimes overproduce PBHs so efficiently that viable models must remain drift-dominated, while effective-field-theory extensions confirm exponential tails at next-to-next-to-next-to-leading order (Choudhury et al., 2025). The tail structure also propagates into spatial statistics, where quantum diffusion imprints a universal clustering signature on PBHs (Animali & Vennin, 2024), and into spectator-field scenarios that generate heavy tails without any large power spectrum (Chen et al., 2025). Exponential tails may even have left marks at cluster scales, where objects like El Gordo are otherwise hard to accommodate (Ezquiaga et al., 2023). On the analytic side, Pi and Sasaki (2023) derived a logarithmic duality showing that the mapping $\zeta = -(1/3)\ln(1 - 3\zeta_G)$ is exact for USR in the drift limit, which supplies the benchmark distribution used throughout this article.

Reading these literatures side by side reveals the gap this article fills. The stochastic papers demonstrate that exponential tails change PBH abundances dramatically, but each does so inside a specific simulated model, at a specific scale, with conclusions expressed in model-internal parameters (Figueroa et al., 2022; Jackson et al., 2025). The constraint papers map the asteroid-mass window with ever finer instruments but calibrate the required inflationary input with Gaussian statistics almost without exception (Carr et al., 2021; Green & Kavanagh, 2021). Nobody, as far as I can determine, has written down the window-wide conversion between the two calibrations in closed form, nor asked whether the fine-tuning diagnosed under Gaussian statistics (Cole et al., 2023) is relieved, worsened, or untouched by the exponential tail. Those are the questions answered below.

Research Methodology

The methodological design is analytic rather than simulational, by deliberate choice. I compare two probability distributions for the coarse-grained curvature perturbation ζ on the scale entering the horizon at each mass in the window. The Gaussian branch takes ζ normally distributed with variance σ^2. The stochastic branch takes the logarithmic mapping $\zeta = -(1/3)\ln(1 - 3\zeta_G)$, where ζ_G is Gaussian with the same variance parameter; this is the exact drift-limit solution for USR found through the duality of Pi and Sasaki (2023), it reproduces the exponential tail $P(\zeta) \propto \exp(-3\zeta)$ that stochastic-δN analyses derive from first-passage theory (Ezquiaga et al., 2020; Pattison et al., 2021), and it matches the tail behaviour observed in full numerical simulations (Figueroa et al., 2021). The benchmark decay rate is therefore $\Lambda = 3$, with the general-Λ dependence retained in all closed-form expressions so that readers can substitute the value their preferred model produces. Model-specific simulations report Λ between roughly 0.5 and 5 depending on the potential (Figueroa et al., 2022); the benchmark sits comfortably inside that range.

Collapse is decided by threshold statistics applied directly to ζ, with $\zeta c = 0.65$, a value consistent with the threshold literature once the shape dependence quantified by Musco (2019) and the analytic prescription of Musco et al. (2021) are averaged over the narrow spectral peaks typical of USR models. The mass fraction at formation is $\beta = \int P(\zeta)d\zeta$ over $\zeta > \zeta c$. For the Gaussian branch this is an error-function integral; for the logarithmic branch it maps to an error-function integral between the transformed threshold $\zeta_{G,c} = (1 - \exp(-3\zeta c))/3 \approx 0.2859$ and the wall at 1/3, which gives every result below in closed form. My first implementation coupled the exponential tail to a compaction-function criterion instead, following the logic of Young et al. (2019) and Raatikainen et al. (2024); I abandoned it after finding that the additional threshold ambiguity it introduces, of order tens of percent in σ^2, is larger than several of the effects I set out to isolate, and a transparent criterion applied identically to both branches serves the comparative purpose better. The absolute abundances carry that crudeness; the ratios between branches, which are the article's actual results, largely do not.

Conversion from β to the present-day dark-matter fraction f_{PBH} uses the standard radiation-era relation in the convention of Carr et al. (2021), with collapse efficiency $\gamma = 0.2$ and 106.75 relativistic degrees of freedom at formation, under which $f_{PBH} = 1$ corresponds to $\beta \approx 8\times10^{-16}$ at a horizon mass of 5×10^{18} g and scales as the square root of the mass across the window. Monochromatic mass functions are assumed at each evaluation mass, in line with how the constraint literature reports bounds (Green & Kavanagh, 2021); extended mass functions would smear but not erase the comparisons made here (Gow et al., 2021). All integrals and root-findings were evaluated numerically in Python with SciPy to relative tolerances of 10^{-13}, and every number quoted in the results section was cross-checked against the closed-form expressions to machine precision. The evaluation grid covers seven masses spaced decade by decade from 10^{17} to 10^{23} g, the sensitivity analysis perturbs σ^2 by $\pm1\%$, $\pm5\%$, and $\pm10\%$, and the threshold-dependence analysis varies ζc from 0.50 to 1.00. Three limitations are flagged now and taken up again in the conclusion: the ζ-threshold criterion is cruder than compaction-based criteria, the drift-limit mapping omits the diffusion-dominated regime where even exponential tails fail (Rigopoulos & Wilkins, 2021), and the analysis holds Λ fixed across the window although realistic potentials would vary it with scale.

Verification proceeded on three tracks. Analytically, every numerical output was rederived from the closed-form expressions and required to agree to machine precision, which caught two transcription errors during development. Internally, the constancy of the recalibration ratio across seven masses and eight abundance targets served as a consistency test of the implementation, since the algebra predicts exact constancy; any drift would have flagged a coding fault. Externally, the outputs were checked for compatibility with published anchors: the Gaussian-calibrated variances land in the amplitude range that Gaussian analyses of the window standardly report (Gow et al., 2021; Karam et al., 2023), the direction and rough size of the tail correction agree with the simulation literature (Figueroa et al., 2022; Gow et al., 2023), and the target mass fractions $\beta \sim 10^{-16}$–10^{-13} match the conversion tables of the constraint reviews (Carr et al., 2021; Green & Kavanagh, 2021). Scales were mapped to masses through the standard horizon-crossing relation, under which the window corresponds to comoving wavenumbers of order 10^{12}–10^{14} Mpc^{-1}, far below any scale probed by the CMB or large-scale structure (Gow et al., 2021). No result depends on cosmological parameters beyond the radiation-era conversion already described, and the full computation runs in seconds, which is the point of closed forms: anyone can repeat it, vary any assumption, and see the consequences without simulation infrastructure.

RESEARCH RESULTS

The computation generated three blocks of findings, one per hypothesis. Throughout, the Gaussian branch and the exponential-tail branch are evaluated with identical threshold $\zeta_c = 0.65$, identical conversion conventions (Carr et al., 2021), and identical numerical tolerances, so that every difference reported below is attributable to the change in tail statistics alone. The mapped Gaussian-variable threshold in the exponential branch is $\zeta_{G,c} = (1 - \exp(-1.95))/3 = 0.2859$, against the mapping wall at 1/3.

The first block addresses H1, the fixed-amplitude comparison. For each evaluation mass, the variance σ^2_G required for $f_{PBH} = 1$ under Gaussian statistics was determined first; the exponential-tail mass fraction was then evaluated at that same variance. At 10^{20} g, the Gaussian calibration gives $\sigma^2_G = 6.977\times10^{-3}$ for the target $\beta = 3.58\times10^{-15}$, while the exponential-tail branch at the identical variance yields $\beta = 2.77\times10^{-4}$. The ratio is 7.7×10^{10}. Expressed as a dark-matter fraction, the same amplitude that produces exactly the observed dark-matter density under Gaussian statistics implies $f_{PBH} \approx 7.7\times10^{10}$ under exponential-tail statistics, an overproduction of nearly eleven orders of magnitude. The error factor is largest at the light edge of the window, 1.24×10^{12} at 10^{17} g, and smallest at the heavy edge, 4.8×10^{9} at 10^{23} g, declining monotonically because the target β itself rises with mass. Table 1 reports the full grid. Every value exceeds the eight-orders-of-magnitude criterion of H1 by at least one and a half additional orders.

Two supplementary checks belong to this block. The threshold dependence of the fixed-amplitude comparison was mapped at 10^{20} g: lowering ζ_c to 0.50 reduces the overproduction factor to 7.8×10^{9}, raising it to 0.80 increases the factor to 2.8×10^{11}, so the qualitative statement of H1 is stable across the entire threshold interval that the collapse literature considers defensible (Musco, 2019; Musco et al., 2021). The geometric origin of the factors is also worth recording as a bare number. At the Gaussian calibration point for 10^{20} g, the threshold $\zeta_c = 0.65$ lies 7.78 standard deviations from the mean; the mapped threshold $\zeta_{G,c} = 0.2859$ lies 3.42 standard deviations out. The overproduction factor of Table 1 is, to within the wall correction, the ratio of Gaussian tail probabilities at those two depths. The wall term itself, the subtracted error function at 1/3 in the exponential branch, contributes at the level of 1.6×10^{-5} relative to the leading term at the recalibrated variance and is numerically irrelevant everywhere in the window, although it is retained in all computations for completeness.

M [g]	β required for f_PBH = 1	σ²_G (Gaussian calibration)	β with exponential tail at σ²_G	Overproduction factor
10^{17}	1.13×10^{-16}	6.27×10^{-3}	1.40×10^{-4}	1.24×10^{12}
10^{18}	3.58×10^{-16}	6.49×10^{-3}	1.76×10^{-4}	4.92×10^{11}
10^{19}	1.13×10^{-15}	6.73×10^{-3}	2.21×10^{-4}	1.95×10^{11}
10^{20}	3.58×10^{-15}	6.98×10^{-3}	2.77×10^{-4}	7.74×10^{10}
10^{21}	1.13×10^{-14}	7.25×10^{-3}	3.47×10^{-4}	3.07×10^{10}
10^{22}	3.58×10^{-14}	7.54×10^{-3}	4.35×10^{-4}	1.22×10^{10}
10^{23}	1.13×10^{-13}	7.86×10^{-3}	5.45×10^{-4}	4.81×10^{9}

Table 1. Fixed-amplitude comparison across the asteroid-mass window: Gaussian-calibrated variance and the resulting exponential-tail overproduction. *Note.* Source: author's computation; conversion convention and target β follow Carr et al. (2021) with $\gamma = 0.2$ and $g^* = 106.75$.

The second block addresses H2, the recalibration. Solving the inverse problem, the variance required for f_PBH = 1 under exponential-tail statistics runs from $\sigma^2\_NG = 1.21\times10^{-3}$ at 10^{17} g to 1.52×10^{-3} at 10^{23} g, against the Gaussian values of 6.27×10^{-3} to 7.86×10^{-3} reported above. The ratio of the two calibrations is 5.17 at every mass in the grid and at every abundance target tested between $\beta = 10^{-20}$ and $\beta = 10^{-12}$, to the third decimal place. This constancy is not numerical coincidence but algebra: because both branches reduce to complementary error functions of their respective thresholds divided by σ, and because the subtracted wall term contributes negligibly at these depths, the required variances stand in the fixed ratio $R\_A = (\zeta c/\zeta\_G,c)^2$, which evaluates to $[\Lambda\zeta c/(1 - \exp(-\Lambda\zeta c))]^2$ in general and to $(0.65/0.2859)^2 = 5.17$ at the benchmark. Table 2 records how R_A moves with the two inputs it depends on, and an algebraic simplification appears immediately: R_A is a function of the single product Λζc, so a soft tail with a high threshold and a hard tail with a low threshold recalibrate identically. Across the plausible threshold range 0.50–1.00 identified in the collapse literature (Musco, 2019; Musco et al., 2021), R_A at the benchmark $\Lambda = 3$ spans 3.7 to 10.0, staying below one order of magnitude as H2 requires. Only the most extreme corner examined, $\Lambda = 5$ combined with $\zeta c = 1.00$, pushes R_A to 25.3, or 1.4 orders of magnitude.

ζc	R_A at Λ = 1	R_A at Λ = 3 (benchmark)	R_A at Λ = 5
0.50	1.61	3.73	7.42
0.60	1.77	4.65	9.97
0.65	1.85	5.17	11.43
0.70	1.93	5.73	13.02
0.80	2.11	6.97	16.60
1.00	2.50	9.97	25.34

Table 2. The amplitude-remapping factor $R\_A = [\Lambda\zeta c/(1 - \exp(-\Lambda\zeta c))]^2$ as a function of collapse threshold and tail decay rate. *Note.* Source: author's computation from the closed-form expression; threshold range follows Musco (2019), decay-rate range follows Figueroa et al. (2022). R_A depends on Λ and ζc only through their product.

The third block addresses H3 and the sensitivity structure. Three findings belong here. First, the recalibrated amplitudes $\sigma^2\_NG \approx (1.2\text{–}1.5)\times10^{-3}$ remain within the range that explicit USR constructions generate without qualitative redesign, corresponding to peak power-spectrum amplitudes of order 10^{-3}–10^{-2} (Ballesteros et al., 2020; Karam et al., 2023). Second, the observational bounds that delimit the window are functions of the present-day PBH mass and abundance only:

the evaporation limits at the light edge (Laha, 2019; Korwar & Profumo, 2023; De la Torre Luque et al., 2024) and the microlensing limits at the heavy edge (Niikura et al., 2019; Smyth et al., 2020) contain no reference to formation statistics, so the window's boundaries are numerically identical in both branches. Third, and unexpectedly, the local sensitivity of the abundance to the variance is invariant: at the f_PBH = 1 calibration point for 10^{20} g, d ln β/d ln σ^2 = 30.8 in the Gaussian branch and 30.8 in the exponential-tail branch. A 1% upward perturbation of σ^2 multiplies β by 1.36 in both branches; a 10% perturbation multiplies it by 16.4 in both. The invariance follows from the same algebra as the constancy of R_A: at matched abundance the error-function arguments are equal, and the local logarithmic derivative depends only on that argument. One further number completes the block: since the scalar-induced gravitational-wave background scales as the square of the power-spectrum amplitude (Domènech, 2021), the recalibration reduces the expected induced-wave signal from an f_PBH = 1 population by $R_A^2 \approx 26.7$.

THE EXPONENTIAL TAIL AS A STRUCTURAL FEATURE OF QUANTUM DIFFUSION

The findings in Table 1 would matter little if exponential tails were a curiosity of one simulated potential. The case for taking them as generic rests on three independent lines of argument, and it is worth walking through each, because the overproduction factors above inherit their credibility from this foundation. The first line is spectral. When the curvature perturbation is computed as a first-passage time of a diffusing field, its distribution admits an expansion over the eigenfunctions of the associated Fokker-Planck boundary problem, and the large-ζ behaviour is dominated by the slowest-decaying mode, giving $P(\zeta) \rightarrow A \exp(-\Lambda\zeta)$ with Λ set by the lowest eigenvalue (Ezquiaga et al., 2020). Nothing in that derivation invokes a particular potential; flatness of the USR stretch only makes the eigenvalue small and the tail correspondingly heavy. Characteristic-function methods reach the same conclusion by locating the poles of the generating function, with the nearest pole to the origin fixing the decay rate (Pattison et al., 2021). The structure is as universal as the spectral decomposition itself, which is to say, very.

It helps to be precise about what quantity carries the tail, because the δN construction is where the stochastic and observational languages meet. In the separate-universe picture, each super-horizon patch evolves as an independent homogeneous cosmology, and the curvature perturbation on a comoving scale equals the fluctuation in the number of e-folds that patch accumulates before a fixed final hypersurface, δN = ζ (Tada & Vennin, 2022). Under slow roll the e-fold count is a smooth function of the field value and inherits near-Gaussian statistics from the field fluctuations; the stochastic formalism keeps this bookkeeping but promotes the field evolution to a Langevin process, so the e-fold count becomes a first-passage time whose distribution must be solved for, not assumed (Pattison et al., 2021; Cruces, 2022). The load-bearing point is that the noise amplitude entering that process is itself calculable from the linear mode functions, and when this is done self-consistently for USR the resulting power spectrum agrees with the standard perturbative result at leading order (Ando & Vennin, 2021; Jackson et al., 2025). The two frameworks therefore disagree about almost nothing measurable at the two-point level. They part company exactly where the observable of interest lives, in the far tail, which is why the disagreement stayed hidden for as long as PBH cosmology borrowed its statistics from CMB practice, where the tail never mattered (Ezquiaga et al., 2020; Figueroa et al., 2021).

The second line is the exact mapping. Pi and Sasaki (2023) showed that the long-wavelength USR dynamics admits a logarithmic duality under which ζ is an explicit nonlinear function of a Gaussian variable, $\zeta = -(1/3)\ln(1 - 3\zeta_G)$. Two features of this mapping do the interpretive work. Near the origin it is linear, so the core of the distribution looks Gaussian and low-order statistics such as the power spectrum barely register the difference (Ando & Vennin, 2021). Far from the origin it diverges logarithmically as ζ_G approaches the wall at 1/3, which stretches modest Gaussian excursions into arbitrarily large curvature perturbations and manufactures the exponential tail. The transformation of the collapse threshold is then the entire story of Table 1 in one line: a 0.65 threshold in ζ corresponds to only 0.2859 in the Gaussian precursor, and moving a threshold inward from 7.8σ to 3.4σ at fixed variance buys ten to twelve orders of magnitude of probability. Seen through the duality, the enormous overproduction factor is unmysterious. It is the price of having previously measured distance to threshold in the wrong variable.

The third line is numerical, and arguably the most persuasive because it dispenses with the approximations of the other two. Lattice-free stochastic simulations of USR inflation, run to the depths relevant for PBH counts, display the exponential tail directly in the histogram of e-fold fluctuations (Figueroa et al., 2021), and the follow-up study mapped how the decay rate responds to the duration and exit of the USR phase (Figueroa et al., 2022). Importance-sampling techniques, which steer simulations into the rare-event region instead of waiting for it, reproduce the same tails with far smaller variance (Jackson et al., 2022), and their extension to realistic noise beyond slow roll confirms the picture when the noise amplitude is computed self-consistently rather than assumed (Jackson et al., 2025). Simulations with frozen noise histories agree (Tomberg, 2023). Where the semiclassical expansion around the stochastic mean is valid at all, it is drift domination that keeps it so; in diffusion-dominated regimes the abundance runs away entirely, which is itself a stochastic prediction with no perturbative counterpart (Rigopoulos & Wilkins, 2021). The formal footing of the framework, including which slow-roll orders the stochastic equations capture, has been secured separately (Cruces & Germani, 2022; Cruces, 2022).

Against this triple foundation, the perturbative treatment of non-Gaussianity looks like what it is: a polynomial approximation to an exponential. A truncated local expansion, however carefully resummed, distorts the Gaussian tail by powers of ζ_G and cannot reproduce $\exp(-\Lambda\zeta)$ at any finite order, a point made quantitatively by comparing quadratic-expansion abundances against the full non-perturbative distribution, where the quadratic template misestimates the required amplitude even when tuned to match the same skewness (Gow et al., 2023). All-order machinery for local non-Gaussianity narrows the gap when the full nonlinear mapping happens to be known (Ferrante et al., 2023), and peak-theory studies with heavy tails point the same way (Kitajima et al., 2021), but these are workarounds that concede the underlying issue. The formation probability is a tail functional, and tail functionals do not commute with truncation (Biagetti et al., 2021). Sudden-transition models sharpen the moral: a single parameter controlling the post-USR relaxation switches the tail between Gaussian-like and exponential families, so even the qualitative form of the statistics is a dynamical output, not an assumption one is free to make (Cai et al., 2022).

How far the exponential family extends beyond the single-field benchmark is an active question, and the answer so far seems to be: further than expected. Effective-field-theory formulations of stochastic single-field inflation, pushed to third subleading order, retain the exponential tail while shifting its parameters (Choudhury et al., 2025). Spectator scenarios generate heavy tails in the compaction function from an axion-like field that never dominates the power spectrum, decoupling the tail question from the spectral-amplitude question altogether (Chen et al., 2025). At the opposite end of the mass scale, quantum-diffusion tails have been invoked to accommodate massive high-redshift clusters such as El Gordo within standard cosmology (Ezquiaga et al., 2023). And the stochastic program has begun to deliver the quantity that formation criteria actually want: the full compaction function of a stochastic USR realization, computed field-theoretically rather than reconstructed from ζ alone (Raatikainen et al., 2024), with the corresponding threshold analysis for the exotic type II configurations that heavy tails make more common

(Shimada et al., 2025). None of this settles the value of Λ for a given potential; simulations put it anywhere between roughly 0.5 and 5 depending on the exit dynamics (Figueroa et al., 2022). What it does settle, on the available evidence, is that the exponential family is the correct family, and that is all the closed-form results of this article require.

FROM PERTURBATIVE CALIBRATION TO STOCHASTIC RECALIBRATION

What should a reader do with an overproduction factor of 7.7×10^{10}? The temptation is to conclude that published Gaussian estimates of PBH abundance are wrong by that factor, and the temptation should be resisted. In practice, almost no one predicts an abundance from a fixed, independently measured amplitude; the logic runs backwards. Model builders adjust potential parameters until the Gaussian pipeline returns $f_{PBH} \approx 1$, then report the model as viable (Ballesteros et al., 2020; Karam et al., 2023). The fixed-amplitude error of Table 1 therefore does not describe a mistake anyone has published. It describes the size of the correction hiding inside the calibration step, and the operationally meaningful number is the recalibration factor $R_A \approx 5.2$: every Gaussian-calibrated model in the asteroid window is, if the stochastic-δN statistics are right, overshooting the required variance by a factor of five and thereby overproducing dark matter by ten orders of magnitude. Five in amplitude, ten orders in outcome. That asymmetry between cause and effect is the entire pathology of tail-dominated observables, compressed into one pair of numbers.

The magnitude of R_A fits sensibly among published model-specific results, which is reassuring given how differently it was obtained. Figueroa et al. (2022) report that Gaussian estimates misjudge the abundance in their simulated models by several orders of magnitude, corresponding to order-unity-to-factor-of-a-few shifts in required amplitude, with the exact value depending on the potential's exit phase. The quadratic-template comparison of Gow et al. (2023) finds required-amplitude shifts of comparable size, and peak-theory analyses with exponential-type tails report abundance changes bracketing the fixed-amplitude factors found here (Kitajima et al., 2021; Ferrante et al., 2023). The closed-form factor should be read as the transparent limit of this literature: what one obtains when the tail is exactly exponential with $\Lambda = 3$, the threshold acts on ζ directly, and everything else is stripped away. Its virtue is portability. A reader who prefers $\zeta_c = 0.55$, or whose model simulation returns $\Lambda = 2$, reads their own R_A off Table 2 in a few seconds, because the factor depends on those two choices only through the product $\Lambda\zeta_c$. I would caution against using it for precision work near the diffusion-dominated regime, where no fixed-Λ description survives (Rigopoulos & Wilkins, 2021; Pattison et al., 2021); its home territory is the drift-dominated bulk where most viable models live.

For a model builder, the results reduce to a short procedure, and spelling it out seems more useful than leaving it implicit. Given a candidate potential tuned under Gaussian statistics to yield $f_{PBH} = 1$ at some mass in the window, one first extracts the tail exponent, either by running an inexpensive stochastic simulation of the USR stretch (Tomberg, 2023; Jackson et al., 2022) or, in the drift limit with a rapid exit, by reading Λ off the exit dynamics analytically (Pi & Sasaki, 2023; Cai et al., 2022). One then forms the product $\Lambda\zeta_c$ with whatever threshold prescription the analysis uses (Musco et al., 2021), takes R_A from the closed form, and divides the model's peak variance by it. The potential parameters are finally retuned to hit the reduced target, a retuning whose difficulty is unchanged, per the invariance result, from the tuning already performed. The whole correction costs an afternoon, not a simulation campaign. What it does not deliver is precision beyond the ζ-threshold approximation: for publication-grade abundances one still hands the recalibrated model to the full stochastic machinery with compaction-based criteria (Raatikainen et al., 2024; Shimada et al., 2025). The factor's role is triage, separating models that survive the statistics correction from those that were artifacts of the Gaussian tail, and at that job a closed form beats a cluster allocation.

The invariance of the fine-tuning is, I think, the finding most likely to be misread, so it deserves unpacking. One might have hoped that a fatter tail, by making large fluctuations cheaper, would soften the exquisite parameter sensitivity that Cole et al. (2023) documented for Gaussian pipelines, where moving f_PBH across its entire observationally interesting range requires amplitude control at the percent level. The computation says otherwise: at matched abundance, d ln β/d ln σ^2 = 30.8 in both branches, and a 10% amplitude error costs a factor of 16 in abundance either way. The reason is that matching the abundance forces the two branches to place their effective thresholds at the same number of standard deviations, and the local steepness of an error-function tail depends on that number alone. Quantum diffusion relocates the calibration point; it does not flatten the cliff around it. Fine-tuning of PBH dark matter is thus not an artifact of Gaussian statistics but a structural feature of demanding $\beta \sim 10^{-15}$ from any distribution with a steep tail, and audits of that tuning performed under Gaussian assumptions (Cole et al., 2023) remain, somewhat fortuitously, quantitatively valid after the statistics are corrected. What does change is the sensitivity to the tail parameter itself: since β responds to Λ through the exponent, an uncertainty of ± 0.5 in Λ at $\zeta c = 0.65$ moves the required variance by tens of percent (Table 2) and the fixed-variance abundance by more than an order of magnitude. The fine-tuning problem has not shrunk. It has acquired a second axis.

The one-loop controversy intersects these results at the level of diagnosis rather than arithmetic. Kristiano and Yokoyama (2024a) argued that USR amplification feeds back on long-wavelength modes at one loop strongly enough to endanger the consistency of PBH-producing models, with the companion calculation extending the claim through the bispectrum (Kristiano & Yokoyama, 2024b); the counter-literature holds that superhorizon conservation protects observable scales (Inomata, 2024), that path-integral treatments show the loops cancelling (Kawaguchi et al., 2024), and that explicit loop computations find no fatal correction (Firouzjahi & Riotto, 2024). Whichever side prevails, the episode and the tail problem are two symptoms of one condition: the observables of PBH cosmology are hostage to regimes where the leading-order, free-field description of ζ is least trustworthy. The stochastic-δN framework is not a spectator in that dispute; by resumming the leading infrared behaviour non-perturbatively, it offers statistics that remain finite and calculable precisely where the loop expansion becomes contentious (Tada & Vennin, 2022; Choudhury et al., 2025). It seems plausible, though it remains unproven, that a fully consistent treatment will vindicate the stochastic tail while defusing the loop divergences; the two programs are converging on the same missing physics from opposite directions.

The recalibration also reaches into gravitational-wave phenomenology, and here it cuts against observability. The scalar-induced background that accompanies any curvature-peak scenario scales with the square of the spectral amplitude (Domènech, 2021), so reducing the required variance by R_A = 5.17 reduces the expected background from an f_PBH = 1 population by a factor of about 27. For the asteroid-mass window, whose induced signal falls near the millihertz band, this pushes the guaranteed counterpart accordingly deeper relative to forecasts built on Gaussian calibration, a shift directly relevant to the detection strategies assembled by the LISA Cosmology Working Group (Bagui et al., 2025). The non-Gaussianity of the tail feeds the induced spectrum a second way, through the connected contributions to the source, and dedicated computations for logarithmic non-Gaussianity of exactly the form used here find modifications to both the amplitude and the shape of the signal (Inui et al., 2025). The lesson from pulsar-timing analyses, where non-Gaussianity decides whether a PBH interpretation of the common-spectrum process survives at all (Franciolini et al., 2023), transfers to the window intact: statistics and phenomenology can no longer be calibrated in separate rooms.

THE ASTEROID-MASS WINDOW UNDER STOCHASTIC STATISTICS

The question posed in the title of this article, whether the asteroid-mass window survives the replacement of perturbative by stochastic statistics, can now be answered by inspecting which quantities in the constraint chain the replacement touches. The window's edges are built from present-day astrophysics. The light edge stands on Hawking evaporation: the 511 keV line budget of the Galactic bulge (Laha, 2019), the broader gamma-ray and cosmic-ray accounting that has progressively hardened it (Korwar & Profumo, 2023; Auffinger, 2023), and the propagation-corrected Galactic emission analysis that currently sets the strongest light-edge bounds (De la Torre Luque et al., 2024). The heavy edge stands on microlensing: the Subaru/HSC campaign toward M31 (Niikura et al., 2019), corrected for finite-source and wave-optics effects (Smyth et al., 2020), with the adjacent planetary-mass decades policed by OGLE (Mróz et al., 2024). Every one of these limits is a functional of the PBH mass function and abundance as they exist today, thirteen billion years after formation. None contains a single quantity that depends on whether the progenitor fluctuations were Gaussian. The window survives. Its boundaries do not move by a single gram.

What the stochastic statistics do move is the inflationary side of the ledger, and there the movement is, on balance, favourable to the scenario. The recalibrated variance σ^2_NG ≈ 1.3×10^{-3} corresponds to a peak enhancement over the CMB amplitude of about 6×10^5, compared with 3×10^6 under Gaussian calibration, which relaxes, modestly, the demands placed on the steepness of the spectral rise that single-field dynamics caps at k^4 (Byrnes et al., 2019). Models that struggled to reach the Gaussian target while respecting large-scale constraints gain headroom (Ballesteros et al., 2020; Karam et al., 2023). There is a second, subtler relaxation: since the exponential tail does the probabilistic heavy lifting, the spectral peak no longer needs to be tuned as close to the collapse cliff, although the invariance result of the previous section warns that the tuning, once the abundance target is fixed, is exactly as sharp as before. Tentatively, one might say that stochastic statistics make the window easier to populate but no easier to populate precisely.

The mass distribution within the window is a different matter, and here the stochastic treatment is not a bookkeeping correction but a qualitative change. Because the exponential tail populates the region far above threshold much more generously than a Gaussian does, the distribution of overdensity magnitudes among collapsing patches shifts upward, which feeds through the formation process into the mass function; the first field-theoretic computations of stochastic compaction functions find exactly such reshaping, along with a larger population of high-amplitude configurations (Raatikainen et al., 2024). Those high-amplitude configurations include type II fluctuations, whose collapse behaviour and thresholds differ from the standard case and whose statistics under non-Gaussianity have only recently been simulated (Shimada et al., 2025; Escrivà et al., 2024). On the clustering side, the news cuts both ways: quantum diffusion imprints a universal correlation structure on PBH positions that Gaussian-calibrated forecasts miss entirely (Animali & Vennin, 2024), yet for the microlensing and evaporation limits that define the window, realistic small-scale clustering has been shown to leave the bounds effectively unchanged (Gorton & Green, 2022). The constraints are safe; the predictions inside them are not.

Coming to this problem from early-universe theory rather than from survey astronomy, I read the constraint literature with particular attention to which bounds could one day discriminate between the two statistical frameworks rather than merely bracket the abundance, and the list is encouragingly long. A next-generation MeV telescope pointed at a single nearby PBH would measure a mass and a local density, tying down one point of the mass function that the two calibrations shape differently (Coogan et al., 2021), and the same instruments observing the Galactic evaporation glow would reach f_PBH sensitivities of 10^{-3} at 10^{17} g, deep enough to test tail-shifted mass functions (Ray et al., 2021). Continuous gravitational-wave searches already bound planetary and asteroid-mass binaries and will improve with detector sensitivity (Miller et al., 2022). Solar-system cratering statistics, if the systematics can be tamed, would probe the window at masses where nothing else does (Santarelli et al., 2025). The sharpest discriminant, though, is likely the induced gravitational-wave background, whose amplitude the recalibration suppresses by a factor of 27 and whose shape the logarithmic non-Gaussianity modifies in a calculable way (Inui et al., 2025; Domènech, 2021). A LISA detection at the Gaussian-forecast level would, within this framework, argue against the exponential tail; a detection 1.4 orders of magnitude below it would argue for it (Bagui et al., 2025). Absence of a detection, unhelpfully, would argue for nothing. Statistics of rare objects offer a further, cheaper test: exponential tails raise the abundance of extreme collapsed structures at all scales, which is the logic behind reading massive early clusters as circumstantial evidence for quantum diffusion (Ezquiaga et al., 2023).

It is worth being explicit about what could still go wrong with the survival verdict, because the edges of the window are themselves under active renegotiation for reasons unrelated to formation statistics. The light edge depends on modelling positron propagation and the bulge emission budget, where systematic shifts of a factor of a few in either direction remain possible (De la Torre Luque et al., 2024; Auffinger, 2023); the heavy edge depends on source-size distributions in M31 and on lensing efficiency calculations that have already been revised once (Smyth et al., 2020; Montero-Camacho et al., 2019). If future revisions narrow the window from both sides, the interaction with formation statistics becomes relevant in a way this article has not needed to consider: a narrowed window imposes structure on the mass function, and the mass function is precisely what the exponential tail reshapes (Raatikainen et al., 2024; Gow et al., 2021). My own expectation, for what it is worth, is that the central decades of the window will remain open for the foreseeable future, since the proposed probes capable of closing it (Coogan et al., 2021; Ray et al., 2021; Santarelli et al., 2025) are the same ones capable of discovering its occupant, and either outcome would count as success.

CONCLUSION

The first hypothesis is confirmed, and with room to spare. At fixed power-spectrum amplitude, replacing the Gaussian tail with the stochastic-δN exponential tail changes the predicted mass fraction by a factor of 4.8×10^{9} at the heavy edge of the asteroid-mass window and 1.24×10^{12} at the light edge, exceeding the hypothesized eight orders of magnitude everywhere on the evaluation grid. The mechanism is fully exposed by the logarithmic duality: the collapse threshold sits 7.8 standard deviations out in the Gaussian branch but only 3.4 standard deviations out in the mapped variable, and nearly eleven decades of probability live between those two depths. The second hypothesis is confirmed within the benchmark statistics and qualified outside them. The amplitude recalibration factor is R_A = 5.17 at the benchmark, ranges over 3.7 to 10.0 across the defensible threshold interval, and stays below one order of magnitude except in the single most extreme parameter corner examined, where it reaches 25.3; a hypothesis stated for the benchmark family survives, but readers working with very hard tails and high thresholds should carry the larger figure. The third hypothesis is confirmed without qualification. The observational boundaries of the window are functionals of the present-day population only, the recalibrated variance remains within reach of explicit single-field constructions, and the scenario in which PBHs constitute all of the dark matter is, if anything, marginally easier to realize under the corrected statistics.

The principal original contribution of this article is the closed-form amplitude-remapping factor $R\_A = [\Lambda\zeta c/(1 - \exp(-\Lambda\zeta c))]^2$, together with its window-wide evaluation and the demonstration that it depends on the tail decay rate and the collapse threshold only through their product. The factor converts any Gaussian-calibrated curvature variance into its exponential-tail equivalent at fixed abundance, in one line, for any model whose tail parameter is known. A second contribution, unplanned at the outset, is the invariance result: the local logarithmic sensitivity of the abundance to the variance is identical in the two statistical frameworks at matched abundance, so the fine-tuning of PBH dark matter, quantified under Gaussian assumptions by earlier audits, is neither relieved nor worsened by quantum diffusion but is joined by a new sensitivity to the tail exponent itself. Neither result, so far as I have been able to establish, appears in the prior literature, which has treated the recalibration model by model and the fine-tuning question only on the Gaussian side.

Beyond its immediate numbers, the exercise carries a message about how the dark-matter question in the asteroid window should be argued. The viability of PBH dark matter is often debated as though it were a single claim, when it is in fact a chain: a potential produces a spectrum, a spectrum produces statistics, statistics produce an abundance, and an abundance confronts constraints. This article has shown that the third link can silently absorb a factor of 10^{10} without any change to the first, second, or fourth, and that the correction, once made explicit, is small on the input side and enormous on the output side. Any argument for or against the scenario that does not state which statistics it assumes is therefore incomplete in a quantifiable way (Figueroa et al., 2022; Gow et al., 2023). The same discipline seems advisable for the constraint literature in reverse: bounds on the power-spectrum amplitude inferred from the absence of PBHs inherit the full factor R_A when translated into statements about inflationary potentials (Gow et al., 2021; Cole et al., 2023).

The limitations of this analysis are ones I chose, and they should be weighed accordingly. I applied threshold statistics directly to ζ rather than to the compaction function, accepting a cruder collapse criterion in exchange for exact closed forms; the ratios between branches are more trustworthy than any absolute abundance quoted here, and readers wanting absolute numbers should look to the simulation literature. I held the tail exponent fixed at $\Lambda = 3$ across six decades of mass, although realistic potentials would vary it with scale, and I worked in the drift-dominated limit throughout, so nothing here applies to models that lean on diffusion domination. The mass functions are monochromatic at each evaluation mass. I also relied on published simulations for the plausible range of Λ rather than running my own, a dependence I would prefer to remove in follow-up work. And I suspect, though I cannot yet show, that coupling the exponential tail to a full compaction-function criterion would lower R_A somewhat, since part of the tail's excess probability lands in configurations that collapse to type II structures rather than adding to the standard count.

Three directions follow naturally. The first is computational: window-wide stochastic simulations with compaction-based criteria, extending recent single-scale results to the full mass grid, would replace the benchmark Λ with a measured profile $\Lambda(k)$ and turn the closed-form factor into a calibrated one. The second is statistical: joint likelihoods over the PBH abundance and the induced gravitational-wave background, with the tail parameter shared between them, would let millihertz observations constrain formation statistics directly rather than through forecasts. The third is theoretical: a unified treatment of the loop corrections and the non-perturbative tail, which the current literature handles in separate formalisms, would settle whether the two problems are, as conjectured above, one problem. For observational programs, the practical implications are immediate: forecasts of the induced-wave counterpart of an f_PBH = 1 asteroid-window population calibrated on Gaussian statistics are optimistic by roughly a factor of 27 in energy density, while evaporation and lensing search strategies need no revision at all. The window is open, the tail is heavy, and the difference between those two statements is now a number.

BIBLIOGRAPHY

Ando, K., & Vennin, V. (2021). Power spectrum in stochastic inflation. *Journal of Cosmology and Astroparticle Physics*, *2021*(04), 057. https://doi.org/10.1088/1475-7516/2021/04/057

Animali, C., & Vennin, V. (2024). Clustering of primordial black holes from quantum diffusion during inflation. *Journal of Cosmology and Astroparticle Physics*, *2024*(08), 026. https://doi.org/10.1088/1475-7516/2024/08/026

Atal, V., & Germani, C. (2019). The role of non-gaussianities in primordial black hole formation. *Physics of the Dark Universe*, *24*, 100275. https://doi.org/10.1016/j.dark.2019.100275

Auffinger, J. (2023). Primordial black hole constraints with Hawking radiation—A review. *Progress in Particle and Nuclear Physics*, *131*, 104040. https://doi.org/10.1016/j.ppnp.2023.104040

Bagui, E., Clesse, S., De Luca, V., et al. (2025). Primordial black holes and their gravitational-wave signatures. *Living Reviews in Relativity*, *28*(1), 1. https://doi.org/10.1007/s41114-024-00053-w

Ballesteros, G., Rey, J., Taoso, M., & Urbano, A. (2020). Primordial black holes as dark matter and gravitational waves from single-field polynomial inflation. *Journal of Cosmology and Astroparticle Physics*, *2020*(07), 025. https://doi.org/10.1088/1475-7516/2020/07/025

Biagetti, M., De Luca, V., Franciolini, G., Kehagias, A., & Riotto, A. (2021). The formation probability of primordial black holes. *Physics Letters B*, *820*, 136602. https://doi.org/10.1016/j.physletb.2021.136602

Byrnes, C. T., Cole, P. S., & Patil, S. P. (2019). Steepest growth of the power spectrum and primordial black holes. *Journal of Cosmology and Astroparticle Physics*, *2019*(06), 028. https://doi.org/10.1088/1475-7516/2019/06/028

Cai, Y.-F., Ma, X.-H., Sasaki, M., Wang, D.-G., & Zhou, Z. (2022). One small step for an inflaton, one giant leap for inflation: A novel non-Gaussian tail and primordial black holes. *Physics Letters B*, *834*, 137461. https://doi.org/10.1016/j.physletb.2022.137461

Carr, B. J., & Hawking, S. W. (1974). Black holes in the early Universe. *Monthly Notices of the Royal Astronomical Society*, *168*(2), 399–415. https://doi.org/10.1093/mnras/168.2.399

Carr, B., Kohri, K., Sendouda, Y., & Yokoyama, J. (2021). Constraints on primordial black holes. *Reports on Progress in Physics*, *84*(11), 116902. https://doi.org/10.1088/1361-6633/ac1e31

Carr, B., & Kühnel, F. (2020). Primordial black holes as dark matter: Recent developments. *Annual Review of Nuclear and Particle Science*, *70*, 355–394. https://doi.org/10.1146/annurev-nucl-050520-125911

Chen, C., Ghoshal, A., Tasinato, G., & Tomberg, E. (2025). Stochastic axionlike curvaton: Non-Gaussianity and primordial black holes without a large power spectrum. *Physical Review D*, *111*(6), 063539. https://doi.org/10.1103/PhysRevD.111.063539

Choudhury, S., Karde, A., Padiyar, P., & Sami, M. (2025). Primordial black holes from effective field theory of stochastic single field inflation at NNNLO. *European Physical Journal C*, *85*(1), 21. https://doi.org/10.1140/epjc/s10052-024-13644-2

Cole, P. S., Gow, A. D., Byrnes, C. T., & Patil, S. P. (2023). Primordial black holes from single-field inflation: A fine-tuning audit. *Journal of Cosmology and Astroparticle Physics*, *2023*(08), 031. https://doi.org/10.1088/1475-7516/2023/08/031

Coogan, A., Morrison, L., & Profumo, S. (2021). Direct detection of Hawking radiation from asteroid-mass primordial black holes. *Physical Review Letters*, *126*(17), 171101. https://doi.org/10.1103/PhysRevLett.126.171101
Cruces, D. (2022). Review on stochastic approach to inflation. *Universe*, *8*(6), 334. https://doi.org/10.3390/universe8060334
Cruces, D., & Germani, C. (2022). Stochastic inflation at all order in slow-roll parameters: Foundations. *Physical Review D*, *105*(2), 023533. https://doi.org/10.1103/PhysRevD.105.023533
De la Torre Luque, P., Koechler, J., & Balaji, S. (2024). Refining Galactic primordial black hole evaporation constraints. *Physical Review D*, *110*(12), 123022. https://doi.org/10.1103/PhysRevD.110.123022
Domènech, G. (2021). Scalar induced gravitational waves review. *Universe*, *7*(11), 398. https://doi.org/10.3390/universe7110398
Escrivà, A., Germani, C., & Sheth, R. K. (2020). Universal threshold for primordial black hole formation. *Physical Review D*, *101*(4), 044022. https://doi.org/10.1103/PhysRevD.101.044022
Escrivà, A., Kühnel, F., & Tada, Y. (2024). Primordial black holes. In *Black holes in the era of gravitational-wave astronomy* (pp. 261–377). Elsevier. https://doi.org/10.1016/B978-0-32-395636-9.00012-8
Ezquiaga, J. M., García-Bellido, J., & Vennin, V. (2020). The exponential tail of inflationary fluctuations: Consequences for primordial black holes. *Journal of Cosmology and Astroparticle Physics*, *2020*(03), 029. https://doi.org/10.1088/1475-7516/2020/03/029
Ezquiaga, J. M., García-Bellido, J., & Vennin, V. (2023). Massive galaxy clusters like El Gordo hint at primordial quantum diffusion. *Physical Review Letters*, *130*(12), 121003. https://doi.org/10.1103/PhysRevLett.130.121003
Ferrante, G., Franciolini, G., Iovino, A. J., & Urbano, A. (2023). Primordial non-Gaussianity up to all orders: Theoretical aspects and implications for primordial black hole models. *Physical Review D*, *107*(4), 043520. https://doi.org/10.1103/PhysRevD.107.043520
Figueroa, D. G., Raatikainen, S., Räsänen, S., & Tomberg, E. (2021). Non-Gaussian tail of the curvature perturbation in stochastic ultraslow-roll inflation: Implications for primordial black hole production. *Physical Review Letters*, *127*(10), 101302. https://doi.org/10.1103/PhysRevLett.127.101302
Figueroa, D. G., Raatikainen, S., Räsänen, S., & Tomberg, E. (2022). Implications of stochastic effects for primordial black hole production in ultra-slow-roll inflation. *Journal of Cosmology and Astroparticle Physics*, *2022*(05), 027. https://doi.org/10.1088/1475-7516/2022/05/027
Firouzjahi, H., & Riotto, A. (2024). Primordial black holes and loops in single-field inflation. *Journal of Cosmology and Astroparticle Physics*, *2024*(02), 021. https://doi.org/10.1088/1475-7516/2024/02/021
Franciolini, G., Iovino, A. J., Vaskonen, V., & Veermäe, H. (2023). Recent gravitational wave observation by pulsar timing arrays and primordial black holes: The importance of non-Gaussianities. *Physical Review Letters*, *131*(20), 201401. https://doi.org/10.1103/PhysRevLett.131.201401
Germani, C., & Musco, I. (2019). Abundance of primordial black holes depends on the shape of the inflationary power spectrum. *Physical Review Letters*, *122*(14), 141302. https://doi.org/10.1103/PhysRevLett.122.141302
Gorton, M., & Green, A. M. (2022). Effect of clustering on primordial black hole microlensing constraints. *Journal of Cosmology and Astroparticle Physics*, *2022*(08), 035. https://doi.org/10.1088/1475-7516/2022/08/035

Gow, A. D., Assadullahi, H., Jackson, J. H. P., Koyama, K., Vennin, V., & Wands, D. (2023). Non-perturbative non-Gaussianity and primordial black holes. *Europhysics Letters*, *142*(4), 49001. https://doi.org/10.1209/0295-5075/acd417

Gow, A. D., Byrnes, C. T., Cole, P. S., & Young, S. (2021). The power spectrum on small scales: Robust constraints and comparing PBH methodologies. *Journal of Cosmology and Astroparticle Physics*, *2021*(02), 002. https://doi.org/10.1088/1475-7516/2021/02/002

Green, A. M., & Kavanagh, B. J. (2021). Primordial black holes as a dark matter candidate. *Journal of Physics G: Nuclear and Particle Physics*, *48*(4), 043001. https://doi.org/10.1088/1361-6471/abc534

Inomata, K. (2024). Superhorizon curvature perturbations are protected against one-loop corrections. *Physical Review Letters*, *133*(14), 141001. https://doi.org/10.1103/PhysRevLett.133.141001

Inui, R., Joana, C., Motohashi, H., et al. (2025). Primordial black holes and induced gravitational waves from logarithmic non-Gaussianity. *Journal of Cosmology and Astroparticle Physics*, *2025*(03), 021. https://doi.org/10.1088/1475-7516/2025/03/021

Jackson, J. H. P., Assadullahi, H., Gow, A. D., et al. (2025). Stochastic inflation beyond slow roll: Noise modelling and importance sampling. *Journal of Cosmology and Astroparticle Physics*, *2025*(04), 073. https://doi.org/10.1088/1475-7516/2025/04/073

Jackson, J. H. P., Assadullahi, H., Koyama, K., Vennin, V., & Wands, D. (2022). Numerical simulations of stochastic inflation using importance sampling. *Journal of Cosmology and Astroparticle Physics*, *2022*(10), 067. https://doi.org/10.1088/1475-7516/2022/10/067

Karam, A., Koivunen, N., Tomberg, E., Vaskonen, V., & Veermäe, H. (2023). Anatomy of single-field inflationary models for primordial black holes. *Journal of Cosmology and Astroparticle Physics*, *2023*(03), 013. https://doi.org/10.1088/1475-7516/2023/03/013

Kawaguchi, R., Tsujikawa, S., & Yamada, Y. (2024). Proving the absence of large one-loop corrections to the power spectrum of curvature perturbations in transient ultra-slow-roll inflation within the path-integral approach. *Journal of High Energy Physics*, *2024*(12), 95. https://doi.org/10.1007/JHEP12(2024)095

Kitajima, N., Tada, Y., Yokoyama, S., & Yoo, C.-M. (2021). Primordial black holes in peak theory with a non-Gaussian tail. *Journal of Cosmology and Astroparticle Physics*, *2021*(10), 053. https://doi.org/10.1088/1475-7516/2021/10/053

Korwar, M., & Profumo, S. (2023). Updated constraints on primordial black hole evaporation. *Journal of Cosmology and Astroparticle Physics*, *2023*(05), 054. https://doi.org/10.1088/1475-7516/2023/05/054

Kristiano, J., & Yokoyama, J. (2024a). Constraining primordial black hole formation from single-field inflation. *Physical Review Letters*, *132*(22), 221003. https://doi.org/10.1103/PhysRevLett.132.221003

Kristiano, J., & Yokoyama, J. (2024b). Note on the bispectrum and one-loop corrections in single-field inflation with primordial black hole formation. *Physical Review D*, *109*(10), 103541. https://doi.org/10.1103/PhysRevD.109.103541

Laha, R. (2019). Primordial black holes as a dark matter candidate are severely constrained by the Galactic Center 511 keV γ-ray line. *Physical Review Letters*, *123*(25), 251101. https://doi.org/10.1103/PhysRevLett.123.251101

Miller, A. L., Aggarwal, N., Clesse, S., & De Lillo, F. (2022). Constraints on planetary and asteroid-mass primordial black holes from continuous gravitational-wave searches. *Physical Review D*, *105*(6), 062008. https://doi.org/10.1103/PhysRevD.105.062008

Montero-Camacho, P., Fang, X., Vasquez, G., Silva, M., & Hirata, C. M. (2019). Revisiting constraints on asteroid-mass primordial black holes as dark matter candidates. *Journal of Cosmology and Astroparticle Physics*, *2019*(08), 031. https://doi.org/10.1088/1475-7516/2019/08/031

Mróz, P., Udalski, A., Szymański, M. K., et al. (2024). Limits on planetary-mass primordial black holes from the OGLE high-cadence survey of the Magellanic Clouds. *Astrophysical Journal Letters*, *976*(1), L19. https://doi.org/10.3847/2041-8213/ad8e68

Musco, I. (2019). Threshold for primordial black holes: Dependence on the shape of the cosmological perturbations. *Physical Review D*, *100*(12), 123524. https://doi.org/10.1103/PhysRevD.100.123524

Musco, I., De Luca, V., Franciolini, G., & Riotto, A. (2021). Threshold for primordial black holes. II. A simple analytic prescription. *Physical Review D*, *103*(6), 063538. https://doi.org/10.1103/PhysRevD.103.063538

Niikura, H., Takada, M., Yasuda, N., et al. (2019). Microlensing constraints on primordial black holes with Subaru/HSC Andromeda observations. *Nature Astronomy*, *3*(6), 524–534. https://doi.org/10.1038/s41550-019-0723-1

Pattison, C., Vennin, V., Assadullahi, H., & Wands, D. (2019). Stochastic inflation beyond slow roll. *Journal of Cosmology and Astroparticle Physics*, *2019*(07), 031. https://doi.org/10.1088/1475-7516/2019/07/031

Pattison, C., Vennin, V., Wands, D., & Assadullahi, H. (2021). Ultra-slow-roll inflation with quantum diffusion. *Journal of Cosmology and Astroparticle Physics*, *2021*(04), 080. https://doi.org/10.1088/1475-7516/2021/04/080

Pi, S., & Sasaki, M. (2023). Logarithmic duality of the curvature perturbation. *Physical Review Letters*, *131*(1), 011002. https://doi.org/10.1103/PhysRevLett.131.011002

Raatikainen, S., Räsänen, S., & Tomberg, E. (2024). Primordial black hole compaction function from stochastic fluctuations in ultraslow-roll inflation. *Physical Review Letters*, *133*(12), 121403. https://doi.org/10.1103/PhysRevLett.133.121403

Ray, A., Laha, R., Muñoz, J. B., & Caputo, R. (2021). Near future MeV telescopes can discover asteroid-mass primordial black hole dark matter. *Physical Review D*, *104*(2), 023516. https://doi.org/10.1103/PhysRevD.104.023516

Rigopoulos, G., & Wilkins, A. (2021). Inflation is always semi-classical: Diffusion domination overproduces primordial black holes. *Journal of Cosmology and Astroparticle Physics*, *2021*(12), 027. https://doi.org/10.1088/1475-7516/2021/12/027

Santarelli, A. D., Caplan, M. E., & Smith, N. T. (2025). Possible constraints on asteroid-mass primordial black holes from Solar system cratering. *Monthly Notices of the Royal Astronomical Society*, *538*(1), 108–113. https://doi.org/10.1093/mnras/staf300

Shimada, M., Escrivà, A., Saito, D., et al. (2025). Primordial black hole formation from type II fluctuations with primordial non-Gaussianity. *Journal of Cosmology and Astroparticle Physics*, *2025*(02), 018. https://doi.org/10.1088/1475-7516/2025/02/018

Smyth, N., Profumo, S., English, S., Jeltema, T., McKinnon, K., & Guhathakurta, P. (2020). Updated constraints on asteroid-mass primordial black holes as dark matter. *Physical Review D*, *101*(6), 063005. https://doi.org/10.1103/PhysRevD.101.063005

Starobinsky, A. A., & Yokoyama, J. (1994). Equilibrium state of a self-interacting scalar field in the de Sitter background. *Physical Review D*, *50*(10), 6357–6368. https://doi.org/10.1103/PhysRevD.50.6357

Tada, Y., & Vennin, V. (2022). Statistics of coarse-grained cosmological fields in stochastic inflation. *Journal of Cosmology and Astroparticle Physics*, *2022*(02), 021. https://doi.org/10.1088/1475-7516/2022/02/021

Tomberg, E. (2023). Numerical stochastic inflation constrained by frozen noise. *Journal of Cosmology and Astroparticle Physics*, *2023*(04), 042. https://doi.org/10.1088/1475-7516/2023/04/042

Young, S., Musco, I., & Byrnes, C. T. (2019). Primordial black hole formation and abundance: Contribution from the non-linear relation between the density and curvature perturbation. *Journal of Cosmology and Astroparticle Physics*, *2019*(11), 012. https://doi.org/10.1088/1475-7516/2019/11/012

NEGAUSOVSKI REPOVI PERTURBACIJE ZAKRIVLJENOSTI U ULTRASPOROJ INFLACIJI: STOHASTIČKA δN OGRANIČENJA NA OBILJE PRIMORDIJALNIH CRNIH RUPA U ASTEROIDNOM MASENOM PROZORU

Matheus Silva Dias

Univerzitet u Sao Paulu
Sao Paulo, Savezna Republika Brazil
E-mail: matheus.sd@usp.br

Originalni naučni članak

Primljen: 22.01.2026. Revidiran: 11.04.2026. Prihvaćen: 18.05.2026.

DOI: https://doi.org/10.65932/CAR-2026-1-1 **UDK**: 524.882:524.83:524.85

Sažetak: Broj primordijalnih crnih rupa (PBH) nastalih iz inflatornih perturbacija zavisi eksponencijalno od gornjeg repa raspodjele vjerovatnoće perturbacije zakrivljenosti ζ, zbog čega je svaka procjena obilja talac statističkih pretpostavki koje ulaze u taj rep. Ovaj članak kvantifikuje, unutar jedinstvenog analitičkog okvira, koliko se procijenjeno obilje PBH mijenja kada se standardni perturbativni Gausov račun zamijeni stohastičkim δN tretmanom ultraspore (USR) inflacije, u kojem rep opada eksponencijalno, $P(\zeta) \propto \exp(-\Lambda\zeta)$, umjesto gausovski. Koristeći logaritamsko preslikavanje između ζ i njenog Gausovog prethodnika, s referentnom stopom opadanja $\Lambda = 3$ i pragom kolapsa $\zeta c = 0.65$, izvodim faktor preslikavanja amplitude u zatvorenoj formi, $R\_A = [\Lambda\zeta c/(1 - \exp(-\Lambda\zeta c))]^2 \approx 5.2$, koji mjeri koliko manja mora biti varijansa σ^2 da bi masena frakcija β ostala fiksna kada se uključe eksponencijalni repovi. Primijenjen preko čitavog asteroidnog masenog prozora (10^{17}–10^{23} g), račun pokazuje da gausovski kalibrisana amplituda $\sigma^2 \approx (6.3–7.9)\times10^{-3}$ preproizvodi PBH za faktor od 5×10^{9} do 1.2×10^{12} kada eksponencijalni rep djeluje pri fiksnoj amplitudi, dok amplituda potrebna za $f\_PBH = 1$ pada na $\sigma^2 \approx (1.2–1.5)\times10^{-3}$. Lokalna osjetljivost obilja na amplitudu, $d \ln \beta/d \ln \sigma^2 \approx 31$ u tački kalibracije, ipak je invarijantna na promjenu statistike, pa poznato fino podešavanje scenarija PBH tamne materije preživljava tranziciju netaknuto. Sam asteroidni maseni prozor ostaje održiv kandidat za ukupnost tamne materije: njegove granice određuju fizika isparavanja i mikrosočiva, koje ne zavise od statistike formiranja, mada prateći inducirani gravitacionotalasni signal slabi za približno faktor 27.

Ključne riječi: *primordijalne crne rupe, ultraspora inflacija, stohastička inflacija, δN formalizam, negausovski repovi, perturbacija zakrivljenosti, tamna materija, asteroidni maseni prozor.*