

Harvesting primordial black holes from stochastic trees with FOREST

Chiara Animali,^{a,b} Pierre Auclair,^a Baptiste Blachier,^{a,b} Vincent Vennin^b

^aCosmology, Universe and Relativity at Louvain (CURL), Institute of Mathematics and Physics, University of Louvain, 2 Chemin du Cyclotron, 1348 Louvain-la-Neuve, Belgium

^bLaboratoire de Physique de l'École Normale Supérieure, ENS, CNRS, Université PSL, Sorbonne Université, Université Paris Cité, F-75005 Paris, France

E-mail: chiara.animali@uclouvain.be, pierre.auclair@uclouvain.be,
baptiste.blachier@uclouvain.be, vincent.vennin@ens.fr

Abstract. We introduce a novel framework to implement stochastic inflation on stochastic trees, modelling the inflationary expansion as a branching process. Combined with the δN formalism, this allows us to generate real-space maps of the curvature perturbation that fully capture quantum diffusion and its non-perturbative backreaction during inflation. Unlike lattice methods, trees do not proceed on a fixed background since new spacetime units emerge dynamically as trees unfold, naturally incorporating metric fluctuations. The recursive structure of stochastic trees also offers remarkable numerical efficiency, and we develop the FORtran Recursive Exploration of Stochastic Trees (**FOREST**) tool and demonstrate its performance. We show how primordial black holes blossom at unbalanced nodes of the trees, and how their mass distribution can be obtained while automatically accounting for the “cloud-in-cloud” effect. In the “quantum-well” toy model, we find broad mass distributions, with mild power laws terminated by exponential tails. We finally compare our results with existing approximations in the literature and discuss several prospects.

Contents

1	Introduction	1
2	Stochastic trees for inflation	4
2.1	The tree structure of inflating spacetimes	4
2.2	Curvature perturbation on the final hypersurface	6
2.3	Mapping the curvature perturbation in comoving space	8
3	Harvesting primordial black holes	9
3.1	Coarse-shelled curvature perturbation	10
3.2	Unbalance index	12
3.3	Mass distribution	13
3.4	Cloud-in-cloud	13
4	Application: quantum-well model	13
4.1	Probability distributions over the trees	14
4.2	Probability distributions over the leaves	15
4.3	Distribution of primordial black holes	16
5	Discussion	17
5.1	Volume weighting	17
5.2	Large-volume approximation	18
5.3	Discretization artefacts	21
6	Conclusion	25
A	Scaling tests	28
B	Convergence tests	29
C	Counting function	31

1 Introduction

Observations of the universe at large cosmological distances reveal that density fluctuations have a small amplitude at large scales, and that they feature Gaussian, phase-coherent and scale-invariant statistics [1–3]. This is consistent with scenarios where cosmological perturbations arise from vacuum quantum fluctuations, stretched to astrophysical distances and amplified by gravitational instability [4–9] during an early era of accelerated expansion called inflation [10–15].

However, there have been recent hints for the existence of large (though rare) density fluctuations at small scales. These include the existence of very massive galaxy clusters [16], the presence of galaxies and quasars at extremely high redshifts [17–20], and the observation in JWST [21–28] of galaxies at very high redshift with large star-formation rate, advanced stellar maturity, or even harbouring black holes at a stage where they are not yet expected to do so if primordial perturbations are Gaussian and quasi scale invariant across all scales.

Large primordial density fluctuations may also lead to the formation of primordial black holes (PBHs), which could account for a fraction or all the dark matter, act as seeds for early structure formation, explain the existence of supermassive black holes in galactic nuclei, and constitute low-spin progenitors for the black-hole mergers detected by gravitational-wave experiments [29–31]. Finally, large scalar perturbations can induce the production of gravitational waves through non-linear effects [32, 33], which could explain the recent pulsar-timing array detections [34, 35].

If inflation is driven by a scalar field (the inflaton) rolling down a smooth potential function, in the attractor regime known as “slow roll” and if its dynamics is dominated by classical drift, then it is expected to give rise to small, quasi-Gaussian curvature perturbations. However, features in the inflationary potential may make the inflaton deviate from the slow-roll attractor towards the end of inflation, or subject it to significant quantum diffusion, leading to enhanced density fluctuations at small scales. In these cases, since large fluctuations are produced, standard cosmological perturbation theory often breaks down in the super-Hubble regime and non-perturbative techniques are instead required. One such approach is the separate-universe picture [36–43], in which the causal structure of inflating spacetimes is used to describe it as an ensemble of independent Hubble-sized patches, locally homogeneous and isotropic, hence evolving according to local background equations of motion but with different initial conditions. In this framework, the curvature perturbation ζ is identified with the local amount of expansion

$$\zeta(t, \vec{x}) = N(t, \vec{x}) - \bar{N}(t) \equiv \delta N, \quad (1.1)$$

where $N(t, \vec{x}) = \ln[a(t, \vec{x})]$ is the number of e -folds, $a(t, \vec{x})$ being the local value of the scale factor, and $\bar{N}(t)$ is the unperturbed number of e -folds. This is the δN formalism [6, 37, 41, 44–46]. In practice, the local amount of expansion, and its statistical properties, can be computed in the stochastic-inflation formalism [47], where quantum fluctuations at small scales source the local background evolution as they cross out the Hubble radius during inflation, in the effective form of a random noise. This makes the dynamics of background fields stochastic, and gives rise to the stochastic- δN formalism [48–51], where the curvature perturbation can be inferred from first-passage-time analysis.

This framework has been used to study the statistics of large fluctuations, which have been shown to feature exponential tails [51–66]. These are substantially heavier than Gaussian tails, which leads to more abundant large-curvature regions than what cosmological perturbation theory would predict. However, determining the real-space profile of the curvature perturbation within the stochastic formalism remains challenging. This difficulty arises because, when simulating realizations of the Langevin equation that governs the stochastic dynamics of the background fields, only a single worldline is tracked until inflation ceases. Repeating this procedure generates an ensemble of such worldlines, but it provides no information about the spatial arrangement of the Hubble patches that terminate them on the end-of-inflation hypersurface.

In principle, in the separate-universe picture, the distance between two final Hubble patches is correlated with the time at which the worldlines they emerge from became independent. Indeed, owing to the accelerated expansion, there is a time prior to which two such worldlines are distant by less than the Hubble radius, hence they belong to the same Hubble patch and follow the same Langevin realization. When their distance grows larger than the Hubble radius they become independent, hence they follow separate realizations of the Langevin equation. This is why patches that are closer one to the other on the

end-of-inflation hypersurface tend to be more strongly correlated, since their worldlines became independent later, and they had less time to evolve away from their common past. This can be used to compute correlation functions of the curvature perturbation, relying only on simulations of single realizations of the Langevin equation, first-passage-time analysis and conditional probabilities [67–69]. Such methods have the advantage to provide analytical or semi-analytical results in some models, but they are so far limited to a subset of cosmological observables. They also rely on approximations regarding the link between physical distances at the end of inflation and the field-space configuration when those scales emerged from a Hubble patch during inflation. Indeed, the physical distance between the endpoints of two worldlines on the end-of-inflation hypersurface depends not only on the two trajectories, but also on the integrated expansion of all the patches in between them. In classical setups, there is a one-to-one correspondence between this physical distance and the field-space configuration within the last common past patch of such worldlines, which does not hold in stochastic systems. Different approximation schemes have been proposed, such as the “backward-distribution approximation” in Refs. [67, 68] and the “large-volume approximation” in Ref. [69], but the validity of these approximations remains to be further tested, and methods to go beyond are still missing.

In this article, we address these issues by proposing to implement the stochastic- δN program on stochastic trees. The idea is that, when a Hubble patch lying along one Langevin realization grows into two Hubble patches, it gives rise to two independent Langevin realizations. This branching process is the elementary structure of a stochastic tree [70, 71]. We show that the tree coordinate of a patch (that is, the direction its worldline followed at each branching process) can be mapped onto physical coordinates on the final hypersurface. In application of the δN formalism, the curvature perturbation on the leaves (namely, on the patches lying on the end-of-inflation hypersurface), is directly related to the time it took for them to reach the end of inflation. Similarly, the time associated with branching nodes yields the curvature perturbation coarse-grained over its set of descendant leaves. This makes the calculation of coarse-grained quantities straightforward, while automatically accounting for cloud-in-cloud effects in the context of structure (including PBH) formation.

Stochastic trees are commonly used in a number of contexts such as genetics and evolutionary biology where phylogenetic trees are often employed, in epidemiology where they are used to model the spread of infectious diseases, in computer science and machine learning where decision trees and random forests are used for classification and regression tasks, in finance and economics to model the pricing of options and other derivatives, in environmental science where trees are employed to simulate e.g. the growth of forests and the spread of wildfires, in ecology to model the population dynamics of species within an ecosystem, in network theory, etc. [72–76]. This implies that a variety of mathematical and numerical tools have been developed to study and simulate stochastic trees. By framing stochastic inflation as a tree problem, we can leverage these existing tools.

Note that other numerical approaches can be followed to simulate non-perturbative dynamics during inflation, such as numerical relativistic methods [77–82] (mostly employed to investigate the onset of inflation from non-homogeneous initial conditions) or lattice codes [83–87].

These approaches are numerically expensive, and do not include the effect of quantum diffusion, which is nonetheless essential in predicting the statistics of large fluctuations during inflation. Recently however, lattice codes have been extended to include stochastic noise

and perform the stochastic- δN program. In particular, the code **STOLAS** [88] delivers three-dimensional maps of the curvature perturbation, where the effect of quantum diffusion and of non-linear evolution is included at large scales. It thus represents a significant advance.

In lattice codes, simulations proceed on a fixed grid, i.e. on flat hypersurfaces. Metric perturbations can therefore not be included, and the stochastic lattice simulations only keep track of matter-field fluctuations. The way the curvature perturbation is extracted in Ref. [88] is by evolving the grid until a time (well before the end of inflation) where the grid spacing exceeds the Hubble radius. At that point, all nodes in the grid have become independent. One can then solve individual stochastic realizations starting from each node, and record the number of e -folds until inflation ends. Although this method is well adapted to some models, it relies on the assumption that inflation stops nowhere on the grid until nodes are evolved independently. This either requires curvature fluctuations to be small, or to terminate the grid sufficiently before the end of inflation, hence to remove a sufficiently large range of small scales. We will see that stochastic trees allow one to overcome these limitations: they are numerically much less expensive, given that the nodes never interact, and since they proceed on a non-rigid grid structure, the local (and coarse-grained) curvature perturbation can be straightforwardly computed.

This article is organized as follows. In section 2, we explain how stochastic trees can be used to describe the dynamics of inflating spacetimes, how they can be numerically implemented, and how the curvature perturbation can be extracted from them. In section 3, we propose a criterion for PBH formation inspired by the well-known compaction function [89–91] and expressed in terms of the quantities accessible during the production of a stochastic tree. We test for PBH collapse at every node in the stochastic tree, automatically accounting for cloud-in-cloud. As an illustration, we apply this formalism in section 4 to reconstruct the mass fraction and mass distribution of PBHs in the quantum-well model of inflation, using populations of 10^{10} to 10^{11} trees. In section 5, we consider how our results compare with previous analytical approximations and discuss the impact of discretization artefacts on our results. We summarize our results and mention a few prospects in section 6, and we end the paper with a few appendices to which technical details are deferred.

2 Stochastic trees for inflation

Although the stochastic-inflation formalism can be applied to multiple-field setups [92], and to models featuring deviations from the slow-roll attractor [93–97], for explicitness we will consider single-field slow-roll models, where the inflaton field ϕ coarse grained at the scale $R_\sigma = (\sigma H)^{-1}$ follows the Langevin and Friedmann equations

$$\frac{d\phi}{dN} = -\frac{V'(\phi)}{3H^2} + \frac{H}{2\pi}\xi(N) \quad \text{and} \quad H^2 = \frac{V(\phi)}{3M_{\text{Pl}}^2}. \quad (2.1)$$

In this expression, where the number of e -folds N is used as a time variable [98], $V(\phi)$ is the potential energy stored in the inflaton, and $\xi(N)$ is a normalized white Gaussian noise, such that $\langle \xi(N) \rangle = 0$ and $\langle \xi(N)\xi(N') \rangle = \delta(N - N')$. Along the slow-roll attractor the effective phase-space is one-dimensional [99], but the extension of stochastic trees to higher-dimensional systems is straightforward.

2.1 The tree structure of inflating spacetimes

In eq. (2.1), ϕ represents the inflaton averaged over a physical region of size $R_\sigma = (\sigma H)^{-1}$ where $\sigma \ll 1$ is a fixed parameter that sets the ratio between the coarse-graining radius and




Figure 1. Elementary vertex of a stochastic tree, where one parent Hubble patch i gives rise to two independent Hubble patches ℓ and m after expanding for $\Delta N = \ln(2)/3$.




Figure 2. Example of a stochastic tree, made of multiple elementary vertices of the type shown in fig. 1. The leaves are Hubble patches where inflation ends, and the patch labels correspond to their topological coordinates in the tree (see main text).

the Hubble radius. We call such a region a Hubble patch. As the expansion proceeds, the physical volume of such a patch grows, and after a time

$$\Delta N = \ln(2)/3 \tag{2.2}$$

it doubles in volume. The situation is depicted in fig. 1 where a parent patch i gives rise to two children patches ℓ and m . In accelerating backgrounds, these two patches have no future causal contact, hence they evolve independently. This gives rise to the “separate-universe” picture mentioned in section 1. This also implies that the field values in the patches ℓ and m , denoted ϕ_ℓ and ϕ_m respectively, have to be evolved from the one in patch i , denoted ϕ_i , using two distinct realizations of the Langevin equation (2.1). In practice, the Langevin equation is thus solved twice over a duration ΔN , from the same initial condition ϕ_i , in order to get ϕ_ℓ and ϕ_m .

When this procedure is iterated, the children patches give rise to two grand-children patches each, so on and so forth, leading to a binary tree structure. An example of such a tree is displayed in fig. 2. It is made of “nodes” and “branches” that implement the causal structure of inflating spacetimes.

If inflation never stops, the tree grows indefinitely. In practice however, inflation has to stop and the nodes where this occurs correspond to the “leaves” of the tree. In the tree sketched in fig. 2, the leaves are the nodes with labels 4, 20, 21, 11, 6, 14 and 15. In single-field slow-roll models, inflation most often stops by slow-roll violation, i.e. when ϕ reaches a certain field value ϕ_{end} where the potential function becomes too steep for inflation to proceed. This is why, when evolving ϕ from a parent node i to one of its child nodes ℓ , if ϕ crosses ϕ_{end} then the process is terminated and a leaf L_ℓ is created. In units of the Hubble volume $V_\sigma = 4\pi R_\sigma^3/3$ at the end of inflation, the physical volume of that leaf is given by (see footnote 7)

$$V(L_\ell) = e^{3(\mathcal{N}_{i \rightarrow \ell} - \Delta N)} \in [1/2, 1[, \quad (2.3)$$

where $\mathcal{N}_{i \rightarrow \ell} < \Delta N$ denotes the number of e -folds realized along the final branch connecting the patches i and ℓ .

Stochastic trees of that type can be efficiently generated numerically using recursive methods, since each vertex gives rise to two subtrees that can be processed independently. Once the routine implementing the elementary vertex depicted in fig. 1 is programmed, it shall call itself from each child node, until inflation ends at the tip of every branch and the entire tree is scanned. We have developed **FOREST** (FOrtan Recursive Exploration of Stochastic Trees), a parallel Modern Fortran code using MPI+OpenMP that proceeds along these lines, and which is used to produce all numerical results presented in this paper. **FOREST** takes as an input an inflationary potential function $V(\phi)$, its derivative $V'(\phi)$, and boundary conditions for the field configuration. The Langevin equation (2.1) is solved using the Euler–Maruyama method [100] with a varying step δN . Close to the end-of-inflation absorbing boundary, we enforce $\delta N < 3[2\pi(\phi - \phi_{\text{end}})]^2/V(\phi)/5$ to limit the probability of barrier crossing to 5σ and avoid double crossings that spoil estimations of the first-passage time (see section B for more details). The value of the Hubble function is updated at each step δN using Friedmann’s equations. We have successfully tested **FOREST** with populations of more than 10^{12} stochastic trees, shared across 2×256 cores using MPI+OpenMP. We report excellent scaling as a function of the number of cores as well as the number of trees (see section A). **FOREST** is publicly accessible under GPLv3 license [101].

For practical purposes, each node of a given tree can be uniquely labelled using its “trajectory” in the binary tree. Starting from the root node 1, one may label the descendants of node n as $2n$ for the leftward node and $2n + 1$ for the rightward node. As a consequence, the position of a node n in the tree can be directly read from its representation in binary. Consider for instance the patch labelled 11 in fig. 2, which reads as 1011_2 in binary. Starting from the parent patch labelled 1, the path goes left (0), then right (1), then right (1) again. This path may be represented as $(1, 0, 1, 1)$, which is indeed the representation of 11 in binary. In this way, a tree is entirely characterized by a set of node labels, and by the volume (2.3) associated to each leaf.

2.2 Curvature perturbation on the final hypersurface

Following the δN formalism, the curvature perturbation corresponds to the local fluctuation in the amount of expansion, see section 1. In the tree structure introduced above, it can be computed as follows.

For each node i in the tree, let V_i be the physical volume emerging from that node, and let W_i be the number of e -folds realized from that node, volume-averaged over all its child

leaves. In other words, if $\mathcal{L}_i$ denotes the set of leaves descending from node i , then one has

$$V_i = \sum_{j \in \mathcal{L}_i} V(L_j) \quad \text{and} \quad W_i = \frac{1}{V_i} \sum_{j \in \mathcal{L}_i} V(L_j) \mathcal{N}_{i \rightarrow j}. \quad (2.4)$$

Here, $V(L_j)$ is the volume of each leaf, computed according to eq. (2.3), and $\mathcal{N}_{i \rightarrow j}$ is the number of e -folds realized on the path that connects the nodes i and j .

Let us consider a late-time observer performing measurements in a bounded region of the final hypersurface. That region contains a subset of leaves of the whole inflationary tree, and the lowest common ancestor of those leaves defines the primeval patch of the observer. The leaves emerging from the primeval patch correspond to the “observable universe” of that observer, while the other leaves are inaccessible to them. For explicitness, let us consider the case where the primeval patch is the root of the tree and is given the label 1. When measuring the curvature perturbation, the observer detects deviations of $\mathcal{N}$ from the mean value computed within its observable universe. In each leaf L_j they thus measure

$$\zeta_{V_j}(\vec{x}_j) = \mathcal{N}_{1 \rightarrow j} - W_1. \quad (2.5)$$

Here, $\vec{x}_j$ labels the position of the leaf L_j , and ζ_{V_j} corresponds to the curvature perturbation coarse-grained over a volume $V_j = V(L_j)$.

The curvature perturbation can also be coarse-grained over volumes larger than individual leaves. Consider for instance the set of leaves $\mathcal{L}_i$ emerging from a node i . The volume-averaged value of ζ_{V_j} across that set corresponds to the curvature perturbation coarse-grained over the volume V_i . In other words,

$$\begin{aligned} \zeta_i \equiv \zeta_{V_i}(\vec{x}_i) &= \frac{1}{V_i} \sum_{j \in \mathcal{L}_i} V_j (\mathcal{N}_{1 \rightarrow j} - W_1) = \frac{1}{V_i} \sum_{j \in \mathcal{L}_i} V_j (\mathcal{N}_{1 \rightarrow i} + \mathcal{N}_{i \rightarrow j} - W_1) \\ &= \mathcal{N}_{1 \rightarrow i} + W_i - W_1, \end{aligned} \quad (2.6)$$

where $\vec{x}_i$ labels the position of the region emerging from node i .

In practice, when exploring the tree, the value of V_i and W_i is computed on the leaves and propagated upwards as follows. On the leaves, $W_j = 0$, while $V_j = V(L_j)$ is obtained as explained above. Then, when a node i gives rise to two child nodes ℓ and m , one has

$$V_i = V_\ell + V_m \quad \text{and} \quad W_i = \frac{1}{V_i} [V_\ell (\mathcal{N}_{i \rightarrow \ell} + W_\ell) + V_m (\mathcal{N}_{i \rightarrow m} + W_m)]. \quad (2.7)$$

In this way, using eq. (2.6), the curvature perturbation coarse-grained over the regions emerging from each node in the tree can be computed.¹

Let us note that, since this propagation proceeds upwards (i.e. from the leaves to the primeval patch), the value of W_1 is computed at the very end only, while it is needed to evaluate the curvature perturbation in eq. (2.6). This implies that the tree has to be explored twice: first to compute W_i and V_i , and second to compute ζ_i . In practice, the recursive embedding of the tree, as detailed in section 2.1, does not rely on storing the whole tree once its exploration is complete. This is one of the reasons why the stochastic-tree procedure is numerically efficient, as it does not require a large amount of dynamic memory. The tree is not created and then explored; instead, the two processes are simultaneous and intertwined.

¹Combining eqs. (2.6) and (2.7), one obtains $\zeta_i = (V_\ell \zeta_\ell + V_m \zeta_m)/V_i$, which is of course expected.




Figure 3. Construction of maps in comoving coordinates by recursively dividing growing Hubble patches.

In practice, to enable us to explore the same tree realization multiple times, we store the seed from which the Pseudo-Random Number Generator (PRNG) is initiated at the onset of the tree. As a consequence, each tree realization is determined uniquely by its seed and can be reconstructed at will. For our PRNG, we use the `xoshiro256**` pseudorandom number generator included in GFortran. This generator is thread-safe when using OpenMP and has a period of $2^{256} - 1 \sim 10^{77}$.

2.3 Mapping the curvature perturbation in comoving space

So far, we have explained how to compute the curvature perturbation at each leaf (i.e. within each Hubble patch belonging to the final hypersurface), as well as its value coarse-grained across the descendant leaves of a given node. To derive actual maps of the curvature perturbation, physical coordinates need to be assigned to the leaves, such that their relative position in real space is prescribed.

In a given tree, there is no unambiguous way to assign physical coordinates to the leaves, or physical distances between them. Consider for instance the one-dimensional tree depicted in fig. 2, and swap the nodes 10 and 11. Leaves 11 and 6 are not neighbour leaves anymore, and they end up being distant by three (instead of one) Hubble diameters across the end-of-inflation hypersurface. However, by swapping two sibling nodes one does not modify the tree (one is still dealing with the same stochastic realization of the tree process), one only changes its graphical representation. The problem is that, at each branching node, i.e. at each volume doubling, two independent branches emerge but the geometry of the regions they describe has been left ambiguous so far.

This ambiguity needs to be lifted somewhat arbitrarily. In practice, we proceed according to the prescription sketched in fig. 3. We start from the primeval patch, i.e. the root of the tree, pictured as a cube, and we let it evolve until its volume doubles after $\Delta N = \ln(2)/3$ (or until it reaches the end of inflation). If the patch is still inflating, we split it along the x -axis, i.e. we produce two Hubble patches separated by an $y-z$ plane. We carry this procedure recursively, and permute at each level of the tree the direction of the splitting as (xyz) . In this way, comoving coordinates are assigned to each leaf, corresponding to their position within the cube.

This prescription breaks the permutation symmetry between sibling nodes, which is ultimately related to homogeneity and isotropy of Friedmann-Lemaître-Robertson-Walker

spacetimes. This symmetry is broken for each individual tree but at large distances it re-emerges in stochastic averages over tree realizations. This is shown explicitly in section 5.3 below, where we find that standard results from quantum field theory (QFT) in curved spacetimes for a light test field can be recovered with stochastic trees, and where we further discuss the impact and mitigation of discretization artefacts.

Note that other spatial discretization schemes may be adopted. For instance, if the orientation of the splitting plane between two sibling Hubble patches is chosen randomly at each step, translational invariance is recovered even at small scales when ensemble average is performed over the trees. However, even in that case it remains true that known QFT results are recovered only at large distances. In fact, in Ref. [71] it is found that the convergence towards QFT results is slower than with the naïve prescription of fig. 3, and that the size of the inflating region at a given time is also not well accounted for (contrary to the prescription of fig. 3, which performs very well). Given that the stochastic- δN formalism provides approximations for the large scales anyway, we will thus stick with the prescription of fig. 3.

In fig. 4, we show the two-dimensional comoving map associated with the tree of fig. 5, where we restrict to two dimensions for the sake of clarity. In practice, we carry even splittings across the x -axis – left child is on the left and right child is on the right – and the odd splittings across the y -axis – left child is at the top and right child is at the bottom.

Let us stress that the above procedure allows one to construct *comoving* maps, which do not directly display the physical distance between leaves. However, the value of the scale factor within each leaf j , as measured with respect to the primeval patch, is nothing but $a(\vec{x}_j) = e^{\mathcal{N}_{1 \rightarrow j}}$, where $\mathcal{N}_{1 \rightarrow j}$ is precisely the quantity being displayed in fig. 4. As a consequence the metric $d\ell^2 = a^2(\vec{x})d\vec{x}^2$ on the end-of-inflation hypersurface is known, and physical distances across that hypersurface can be computed by solving for geodesics and integrating $d\ell$ along them. We do not need maps with physical coordinates in the present article and thus leave their analysis for future work.

3 Harvesting primordial black holes

The formation of primordial black holes usually takes place in regions of large curvature. However, using the value of the curvature perturbation as a criterion for PBH formation is not always well justified, in particular when quantum diffusion plays an important role. The reason is that, as explained in section 2.2, ζ is defined with respect to a given observer, and relative to a “background” that encloses their whole observable universe. However, gravitational collapse into a black hole is a local process, and should not depend on averages performed over regions much larger than the Hubble radius when it takes place. In other words, ζ_i in eq. (2.6) is not a fully local quantity since it is affected by the volume produced in regions of the tree that are very distant from node i , via their contributions to W_1 . Using ζ_i to decide whether the region emerging from node i collapses into a black hole thus seems problematic.

For this reason, other quantities are often considered for PBH formation criteria, such as the comoving density contrast [102, 103], or its non-perturbative generalization, the compaction function [89, 91, 104]. The reason why the comoving density contrast is better suited than the curvature perturbation is that, through Poisson equation, it is proportional to the Laplacian of the curvature perturbation, hence in Fourier space it features an additional k^2 factor. This suppresses large-scale contributions in the coarse-grained density contrast, which is thus a more “local” tracer than the curvature perturbation.




Figure 4. Example of a two-dimensional comoving map with five PBHs in the flat-well toy model discussed in section 4, with $\mu = 0.8$ and starting from $\phi = \Delta\phi$. The colour scale indicates the elapsed number of e -folds of the different patches on the end-of-inflation hypersurface. This is the same realization as in fig. 5. On the right panel, we mark the five PBHs by superimposing black patches, according to the PBH formation criterion presented in section 3.2. We single out this realization for illustrative and pedagogical purposes, it is not particularly representative of the tree population, in which PBH formation remains a rare event.

3.1 Coarse-shelled curvature perturbation

In the stochastic trees, we implement an idea similar to the “coarse-shelled” curvature perturbation proposed in Ref. [68]. Consider a given node i , giving birth to two nodes ℓ and m (this is the situation sketched in fig. 1). Our goal is to use the node i as a local background for the node ℓ , and to evaluate the curvature perturbation in node ℓ relative to its local environment i ,

$$\zeta_{\ell i} = \zeta_{\ell} - \zeta_i. \quad (3.1)$$

This “coarse-shelled” curvature perturbation can be related to the Laplacian of the curvature perturbation, or more generally to the compaction function, as follows. The region emerging from node i is composed of the regions stemming from its two child nodes ℓ and m , but the geometry of these subregions is not specified.² Since the Laplacian is a spherically symmetric operator, let us picture the regions emerging from nodes i and ℓ as being two concentric spheres. In this case, $\zeta_{\ell i}$ is the curvature perturbation coarse grained over the shell located between the spheres of radius R_{ℓ} and R_i , where $V_i = 4\pi R_i^3/3$ and likewise for V_{ℓ} , which is why it is called the “coarse-shelled” curvature perturbation. For the noise appearing in the Langevin equation (2.1) to be white, i.e. uncorrelated over time, coarse graining has to be performed by means of a sharp window function in Fourier space, $\zeta_R(\vec{k}) = \theta(a/R - k)\zeta(\vec{k})$,

²When constructing the maps of section 2.3, a choice was made that consists in splitting volumes along orthogonal planes, the orientation of which alternates as the time elapsed from the primeval patch increases, such that statistical isotropy is recovered at large scales, see section 5.3. It can thus be used to compute quantities coarse-grained over large volumes. The coarse-shelled curvature perturbation being a local quantity, one should rather employ a scheme that is isotropic already at the level of an individual leaf and its parent. Note that if the “concentric-sphere” scheme were embedded globally it would lead to some Hubble patches being pictured as thin layers of large radius, which would be incompatible with the geometry of inflating space-times. Therefore, different schemes should be employed for local and global quantities, such that the space-time symmetries are correctly implemented at both levels.

where $\theta(x)$ is the Heaviside function. In this case,

$$\zeta_{\ell i}(\vec{k}) = \mathcal{W}_{R_\ell, R_i}(k) \zeta(\vec{k}) \quad \text{where} \quad \mathcal{W}_{R_\ell, R_i}(k) = \theta(k - a/R_i) \theta(a/R_\ell - k). \quad (3.2)$$

The window function for the coarse-shelled curvature perturbation, $\mathcal{W}_{R_\ell, R_i}(k)$, removes all scales larger than R_i , hence it ensures that $\zeta_{\ell i}$ is indeed a local quantity.

A correspondence between $\zeta_{\ell i}$ and the density contrast or the compaction function can be established by matching their effective window functions. At linear order in perturbation theory, the comoving density contrast is related to the curvature perturbation as

$$\delta = \frac{2(1+w)}{5+3w} \frac{1}{a^2 H^2} \Delta \zeta, \quad (3.3)$$

where $w = p/\rho$ denotes the equation-of-state parameter of the background fluid. If it is coarse grained with a Gaussian window function [105, 106], $\delta_R(\vec{k}) = e^{-(kR/a)^2/2} \delta(\vec{k})$, its relation to the curvature perturbation is given by

$$H^2 R^2 \frac{5+3w}{2(1+w)} \delta_R(\vec{k}) = \mathcal{W}_R(k) \zeta(\vec{k}), \quad \text{where} \quad \mathcal{W}_R(k) = \frac{k^2 R^2}{a^2} \exp\left[-\frac{1}{2} \left(\frac{kR}{a}\right)^2\right]. \quad (3.4)$$

In practice, PBH formation criteria need to be evaluated at the time when the relevant scale R re-enters the Hubble radius, thus we assume $HR = 1$ in the above.

The compaction function $\mathcal{C}(r)$ is defined as the difference between the Misner–Sharp mass contained in the sphere of comoving radius r , and the background mass within the same areal radius. Around spherically-symmetric peaks, it is related to the curvature perturbation via

$$\mathcal{C}(r) = \frac{3(1+w)}{5+3w} \left\{ 1 - [1 + r\zeta'(r)]^2 \right\}. \quad (3.5)$$

If r_m denotes the value of r where $\mathcal{C}$ is maximum, PBH formation thresholds have been derived on $\mathcal{C}(r_m)$, where the mass of the resultant black hole is related to the one contained within r_m . When the Misner–Sharp mass is defined with a Gaussian smoothing function [107], in Ref. [68] it is shown that

$$R\zeta'(R) \simeq -\frac{1}{3} \mathcal{W}_R(k) \zeta(\vec{k}), \quad (3.6)$$

which features the same window function as the comoving density contrast. A correspondence between the coarse-shelled curvature perturbation, the density contrast and the compaction function can thus be established provided the replacement

$$\mathcal{W}_{R_\ell, R_i}(k) \sim \alpha \mathcal{W}_{\beta R_i}(k) \quad (3.7)$$

can be made, where α and β should be set such that both hands of eq. (3.7) are as similar as possible. In practice, we require that when employed with a logarithmic measure both window functions peak at the same scale and that they share the same “volume” $\int \mathcal{W}(k) d \ln k$, which leads to $\alpha = \ln(R_i/R_\ell)$ and $\beta = 1$. As a consequence, a PBH threshold δ_c on the comoving density contrast can be translated into a threshold $\zeta_{\ell i, c}$ on the coarse-shelled curvature perturbation according to

$$\zeta_{\ell i, c} = \frac{5+3w}{2(1+w)} \ln\left(\frac{R_i}{R_\ell}\right) \delta_c. \quad (3.8)$$

Likewise, a PBH threshold $\mathcal{C}_c$ on the compaction function can be translated into

$$\zeta_{\ell i, c} = 3 \ln\left(\frac{R_i}{R_\ell}\right) \left[1 - \sqrt{1 - \left(\frac{5+3w}{3+3w}\right) \mathcal{C}_c} \right]. \quad (3.9)$$




Figure 5. Example of a tree realization with five PBHs in the flat-well toy model, with $\mu = 0.8$ and starting from $\phi = \Delta\phi$. This is the same realization as fig. 4. The black patches are those for which the coarse-shelled curvature perturbation, used as a proxy for the compaction function, exceeds the PBH formation threshold.

3.2 Unbalance index

In practice, we use eq. (3.9) to evaluate $\zeta_{\ell i, c}$ at each branching. The nodes ℓ such that $\zeta_{\ell i} > \zeta_{\ell i, c}$ are assumed to collapse into PBHs. Expressed in terms of the compaction function, we take the value $\mathcal{C}_c = 0.5$ as a threshold for PBH formation, assuming that PBHs form in a radiation era with $w = 1/3$.³ This leads to $\zeta_{\ell i, c} = \ln(V_i/V_\ell)/2$. In fig. 5, we display a tree realization generated in the flat-well toy model (see section 4), where the nodes collapsing to PBHs are coloured in black.

One can see that the nodes giving rise to PBHs are those for which the two emerging branches generate substantially different volumes. This can be understood as follows. By inserting eq. (2.6) into eq. (3.1), one finds

$$\frac{\zeta_{\ell i}}{\zeta_{\ell i, c}} = \frac{2}{\ln(V_i/V_\ell)} \frac{V_m}{V_i} (W_\ell - W_m), \quad (3.10)$$

where we have used that $\mathcal{N}_{1 \rightarrow \ell} = \mathcal{N}_{1 \rightarrow i} + \mathcal{N}_{i \rightarrow \ell}$ together with eq. (2.7). We have also assumed that inflation does not end in the nodes ℓ and m , hence $\mathcal{N}_{i \rightarrow \ell} = \mathcal{N}_{i \rightarrow m} = \Delta N$. If the tree

³Numerical simulations [91, 103, 108, 109] show that the critical threshold depends sensitively on the details of the density profile around the peak and may vary between 2/5 and 2/3. Given the approximation made when matching the window functions in section 3.1, we adopt a fixed value for $\mathcal{C}_c$, on which we have checked that our results do not depend much.

is perfectly balanced between the nodes ℓ and m , then $W_\ell = W_m$ and $\zeta_{\ell i} = \zeta_{mi} = 0$. In the opposite limit where the subtree emerging from node ℓ is much larger than the one emerging from node m , $V_\ell \gg V_m$, one can expand $\ln(V_i/V_\ell) \simeq V_m/V_i$, which leads to

$$\frac{\zeta_{\ell i}}{\zeta_{\ell i, c}} \simeq 2(W_\ell - W_m). \quad (3.11)$$

If $V_\ell \gg V_m$, then in most cases $W_\ell \gg W_m$ and the condition $\zeta_{\ell i} > \zeta_{\ell i, c}$ is met.

The coarse-shelled curvature perturbation of eq. (3.10) may thus be seen as an “unbalance index”: PBHs form at nodes where the tree exhibits a high level of unbalancing. This allows one to interpret fig. 5 readily, since the black nodes are indeed the most asymmetric ones.

3.3 Mass distribution

Once a node giving rise to a PBH has been identified, its mass can be computed. As a first approximation, it is given by the mass $M_H(R_i)$ comprised in a Hubble volume at the time when the scale R_i re-enters the Hubble radius. Assuming that reheating proceeds instantaneously at the end of inflation, one has

$$M_H(R_i) \simeq M_{\text{Pl}}^2 R_i^2 H_{\text{end}}. \quad (3.12)$$

A more precise estimate can be obtained accounting for critical collapse [110], but the associated scaling law is valid only when $\mathcal{C} - \mathcal{C}_c \lesssim \mathcal{O}(10^{-2})$ [107]. Since most super-critical nodes are found to exceed the threshold by more than 10^{-2} , critical scaling is not more accurate than the simple estimate (3.12), to which we thus stick.

3.4 Cloud-in-cloud

As can be seen in the example displayed in fig. 5, when a subtree deviates from the average trajectory and produces a large volume, it induces a substantial unbalancing not only of the node it starts from, but also of some of its parent nodes. The reason is that all parent nodes give rise to one branch that holds this subtree and one branch that does not, hence they are likely unbalanced by its large volume. As a consequence, it is common that PBHs arise at multiple nodes along a path segment, hence that they form in a nested way.

Nested PBH formation is analogous to the well-known cloud-in-cloud problem [111], in which a hierarchy of overdense regions leads to PBH formation inside larger PBHs. Fortunately, the recursive structure of stochastic trees is perfectly suited to treat this issue. Indeed, in **FOREST**, PBH formation is evaluated dynamically at each node while the tree is created. Information such as V_i, W_i , the number and masses of PBHs is always propagated upwards and saved in the stack as we progress in the tree structure. As a consequence, if the algorithm finds a new PBH at a node, then all its inner PBHs are discarded and do not propagate further. “Cloud-in-cloud” is therefore automatically accounted for in this approach.

4 Application: quantum-well model

As a direct application of the framework we presented and to showcase the power of **FOREST**, we revisit the common toy model of quantum-well inflation [51, 52, 67–69]. In this framework, the inflaton field leaves its classical trajectory, dominated by the potential-induced drift, when

it enters a flat region of the potential $V(\phi) = 24\pi^2 M_{\text{Pl}}^2 v_0$, where the dynamics is driven by the stochastic noise.

We assume that the flat well ends at ϕ_{end} ⁴, which thus features an absorbing boundary for our Langevin trajectories. The quantum well also exhibits a reflective boundary at $\phi_{\text{end}} + \Delta\phi$, with $\Delta\phi$ being the width of the well, separating the region of quantum diffusion from the region of classical evolution prior to it.

It is convenient to parametrize the model in terms of the following dimensionless quantities

$$\mu^2 \equiv \frac{\Delta\phi^2}{v_0 M_{\text{Pl}}^2} \quad \text{and} \quad x \equiv \frac{\phi - \phi_{\text{end}}}{\Delta\phi} \in [0, 1]. \quad (4.1)$$

Qualitatively, μ is an indicator of the time scales associated to the flat well. For instance, the mean number of e -folds across the well is given by [51] $\langle \mathcal{N} \rangle = \mu^2/2$. One can also show that the mean volume emerging from a patch with field value x , in Hubble units, is given by [69]

$$\langle V \rangle = \langle e^{3\mathcal{N}} \rangle = \frac{\cos[\sqrt{3}\mu(1-x)]}{\cos(\sqrt{3}\mu)}, \quad (4.2)$$

which is well-defined only up to a critical value $\mu < \mu_c \equiv \pi/(2\sqrt{3})$.⁵ For values of $\mu \geq \mu_c$, the tail of the first-passage-time distribution decays more slowly than $e^{3\mathcal{N}}$, and the mean volume is divergent, which may lead to an “eternal inflation” problem (see discussion in section 5.1). For this reason, we only work with values of μ below the critical value.

4.1 Probability distributions over the trees


Starting from a patch with field configuration $x_* \in [0, 1]$, we let it grow for $\Delta N = \ln(2)/3$ then split it recursively. This produces a tree configuration of which we compute the volume V and the volume-averaged number of e -folds W . By simulating a population of such trees, we are able to sample the distributions of V and W over the ensemble of trees. We show these distributions for the flat-well model and for several values of μ in fig. 6.

First, we observe that the tree populations have a modest mean volume $\langle V \rangle = \mathcal{O}(10)$, as well as a modest mean value for the volume-averaged expansion $\langle W \rangle = \mathcal{O}(1)$. Both distributions, however, exhibit exponential tails, which become progressively heavier as μ increases. This is to be expected, given that quantum-diffusion effects are larger for increasing values of μ . In particular, the volume distribution shows the behaviour $P(V|\phi_*) \propto e^{z_* V} V^{-3/2}$, where $z_* < 0$ is the rightmost pole of the generating function for V [112]. Above its mean value, the volume distribution follows a power-law behaviour, which is then overtaken by an exponential decay at the far tail.

Increasing μ shifts the onset of this exponential phase to larger volumes, resulting in heavier tails. For $\mu \rightarrow \mu_c$, $z_* \rightarrow 0$, which implies $P(V) \propto V^{-3/2}$: this marks the onset of the eternal-inflation regime, which is represented by the black dashed line in the left panel of fig. 6. This is in agreement with previous studies in which a similar behaviour for the volume distribution at the end of inflation has been found [112, 113]. Also note that the slight kinks when V and W are of order one are discretization artefacts, see section 5.3.1, which do not affect the large-scale (i.e. large V or large W) statistics.

⁴The field value ϕ_{end} marks with the end of inflation in our setting. If the exit from the well does not coincide with the end of inflation, but instead a classical phase follows, the number of inflationary e -folds would be shifted by a constant value. However, this shift does not affect the amplitude of the curvature perturbations, although it does shift the scales and the masses of collapsed objects, by a constant multiplicative factor.

⁵The first equality $\langle V \rangle = \langle e^{3\mathcal{N}} \rangle$ is not trivial but can be proven rigorously using branching processes.



(a) Probability distribution of the final volume. (b) Probability distribution of the volume-averaged expansion.

Figure 6. Probability distribution of the final volume (a) and of the volume-averaged number of e -folds (b) over the tree population in the flat well for $\mu = 0.6, 0.7, 0.8, 0.85, 0.89$ and $x_* = 1$. These were obtained with, respectively, $10^{11}, 10^{11}, 5 \times 10^{10}, 2 \times 10^{10}, 10^{10}$ trees.

Finally, let us stress that exponential tails are one of the key features of the stochastic- δN formalism [51, 52, 60, 63, 69, 114]. This is of particular interest in the context of PBH formation since the relative abundance of anomalously large subtrees is what determines the fraction of the end-of-inflation hypersurface that eventually collapses into PBHs.

4.2 Probability distributions over the leaves

In addition to global properties of the trees, one may be interested in properties of their leaves, i.e. compute statistics over the set of leaves of the entire tree population. For instance, starting from the root patch 1 having field configuration x_* , one can register for each leaf j its number of e -folds from the root, weighted by its final volume. This gives rise to the volume-weighted first-passage-time distribution through the end-of-inflation hypersurface, $P_{\text{FPT}, x_*}^V(\mathcal{N})$, which is computed with **FOREST** and shown in fig. 7 for several values of μ and with $x_* = 1$.

We find that $P_{\text{FPT}, x_*}^V(\mathcal{N})$ displays a transient power-law behaviour followed by a heavy exponential tail, showing that a relatively abundant number of branches expands for a very long time compared to the average. Note that PBH formation is generically dominated by these outliers in the tail of the distribution.

In principle, from the first-passage-time distribution one can also derive the distribution of the curvature perturbation coarse-grained at the scale R_σ at the end of inflation, ζ_{R_σ} . Indeed, the two quantities differ only by the subtraction of the volume-averaged expansion W_1 in eq. (2.5). However, as mentioned above, W_1 is to be averaged over the region of the final hypersurface a given observer has access to. That region is different for every tree, hence different values of W_1 need to be subtracted from leaves belonging to different trees, and the distribution for ζ_{R_σ} is not simply a shifted version of the distribution for $\mathcal{N}$. Since this issue does not arise for the coarse-shelled curvature perturbation, which is relevant for PBH formation that we next study, we limit ourselves to presenting the distribution of $\mathcal{N}$.




Figure 7. Volume-weighted probability distribution of the first-passage time $\mathcal{N}_{x_* \rightarrow x_{\text{end}}}$ through the end-of-inflation hypersurface, in the flat well for $\mu = 0.6, 0.7, 0.8, 0.85, 0.89$ and $x_* = 1$. These were obtained with, respectively, $10^{11}, 10^{11}, 5 \times 10^{10}, 2 \times 10^{10}, 10^{10}$ trees. Full coloured lines stand for the results of FOREST, whereas black dashed lines represent the analytical formula (5.1).

4.3 Distribution of primordial black holes

As described in sections 3.1 and 3.2, FOREST identifies and records all the PBHs that appear in our populations of stochastic trees using an “unbalance index”, i.e. a proxy for the compaction function. As per the name suggests, the unbalance index picks up nodes from which two sibling branches experience very different expansions, which indicates strong gradients in the curvature perturbation. When we explore our stochastic trees, we only keep track of the utmost/largest PBHs in the tree (see section 3.4) and we assume that the downstream branch completely collapses, thus setting the mass of the resulting PBH. This is illustrated in figs. 4 and 5, which display a tree with five distinct PBHs of relatively large masses. When averaged over a large ensemble of trees, we report our findings for the flat well and $\mu = 0.6, 0.7, 0.8, 0.85, 0.89$ in fig. 8.

First, in fig. 8a we show $f_{\text{PBH, end}}$, the fraction of the universe at the end of inflation that will eventually collapse into PBHs, as a function of μ . We recover the expected result that PBH production is more abundant when quantum diffusion is more prominent, that is when μ increases. As μ gets closer to the critical value μ_c , nearly all the universe is contained in PBHs.

At a first glance, one may be surprised to have $f_{\text{PBH, end}} \rightarrow 1$ close to the critical point, when our criterion for PBH requires having contiguous regions of short and long expansion. This can be understood by the fact that the overall volume of the universe becomes dominated by a limited number of very large branches leading to PBHs, separated by numerous regions of short expansion whose volume is exponentially suppressed close to the critical value.

Second, we show the mass distribution of these PBHs in fig. 8b. We find that the distributions peak at a value of the mass that increases with μ , that they are then followed by a power-law range whose width also increases with μ , before being interrupted by an




Figure 8. Statistics of PBH production in the flat quantum well. Panel (a): fraction of the universe comprised in (regions that will eventually collapse into) PBHs at the end of inflation, as a function of μ , compared to the simple estimate $P(\zeta_{R_\sigma} > \zeta_c)$ derived in the large-volume approximation in section 5.2. Panel (b): distribution of masses for the produced PBHs for $\mu = 0.6, 0.7, 0.8, 0.85, 0.89$ and $x_* = 1$. These were obtained with, respectively, $10^{11}, 10^{11}, 5 \times 10^{10}, 2 \times 10^{10}, 10^{10}$ trees.

exponential tail. The existence of these wide and mild power laws is a direct consequence of the “cloud-in-cloud” mechanism described in section 3.4, by which smaller PBHs are concealed within larger ones.

When μ approaches μ_c , the power-law behaviour extends to larger and larger masses, with $df/d \ln M \propto M^{-\alpha}$, with $\alpha \approx 2/3$. This can be seen clearly in the case $\mu = 0.89$, which is indeed close to criticality. PBH formation close to μ_c thus displays some properties one expects from critical phenomena: PBH form at all scales and cover nearly all the space in the universe, only separated by infinitely small regions of void.

5 Discussion

5.1 Volume weighting

Previous attempts to analytically describe the curvature perturbation and PBH formation in stochastic inflation are based on first-passage-time analysis. In this section we review why first-passage-time distributions need to be volume weighted, and we check that they may account for some (but not all) of the results presented above.

In the stochastic-inflation picture, distinct regions of the universe inflate by different amounts, hence they give rise to different numbers of Hubble patches at the end of inflation and contribute differently to ensemble averages over these patches. These ensemble averages correspond to observable quantities: if a hypothetical observer located on the end-of-inflation hypersurface is tasked with computing the average expansion, they will measure the expansion in each patch contained within their observable region and then perform an average over the ensemble of such measurements. Averaging over the set of final Hubble patches, i.e. over the set of leaves, is nothing but volume averaging over the set of trees or subtrees, i.e. weigh each subtree by the volume it generates [64, 69]. In other words, spatial averaging in physical coordinates is equivalent to volume-weighted spatial averaging in comoving coordinates.

For the first-passage-time distribution associated to the Langevin problem (2.1), the volume-weighted version can be defined as [69]

$$P_{\text{FPT},\phi}^{\text{V}}(\mathcal{N}) = \frac{P_{\text{FPT},\phi}(\mathcal{N})e^{3\mathcal{N}}}{\int_0^\infty P_{\text{FPT},\phi}(\mathcal{N})e^{3\mathcal{N}}d\mathcal{N}}, \quad (5.1)$$

which is well-defined only when the integral in the denominator converges. If the first-passage-time distribution features an exponential tail $P_{\text{FPT},\phi}(\mathcal{N}) \propto e^{-\Lambda_0\mathcal{N}}$ at large $\mathcal{N}$, then it does not converge when $\Lambda_0 < 3$, giving rise to eternal inflation. In that case the mean tree volume is always infinite, regardless of the initial condition.

For the quantum-well model considered in section 4, the first-passage-time problem has an exact solution in terms of the derivative of the second elliptic theta function [51], $P_{\text{FPT},x_*}(\mathcal{N}) = -\pi/(2\mu^2)\vartheta'_2(x_*\pi/2, e^{-\pi^2/\mu^2\mathcal{N}})$. Its tail is of the exponential type, $P_{\text{FPT},x_*} \propto e^{-\pi^2/(4\mu^2)\mathcal{N}}$, which gives rise to the critical value $\mu < \mu_c \equiv \pi/(2\sqrt{3})$ mentioned below eq. (4.2). The volume-weighted version (5.1) of this first-passage-time distribution is shown in fig. 7, where it is compared with the distribution reconstructed from stochastic-tree simulations. We observe an excellent agreement between the two, for all values of μ scanned and for any value of $\mathcal{N}$.

The reason why this agreement is not so obvious is that, in fig. 7, the leaves over which the first-passage-time distribution is computed are not the end-point of independent stochastic realizations of the Langevin equation. Indeed, two leaves belonging to the same tree share the same path until their splitting node and thus give rise to values of $\mathcal{N}$ that are correlated. As a consequence, the distribution displayed in fig. 7 is not obtained by sampling independent realizations of the Langevin process: realizations are gathered in trees within which they may be highly correlated. However, when the number of trees is large, these correlations become sparse since the fraction of the leaves to which a given leaf is correlated becomes arbitrarily small (it cannot exceed the ratio between the volume of the tree it lies in and the total volume of all trees). In this limit, leaves are mostly independent and the mere volume weighting (5.1) coincides with the ensemble average over leaves.

This serves both as a robustness test for the numerical code and, more importantly, as a consistency check for the volume-weighting procedure (5.1), at least as far as the duration of inflation is concerned. This is important since such procedures are ubiquitous when designing analytical approximations in terms of first-passage-time distributions, such as the large-volume approximation that we now discuss.

5.2 Large-volume approximation

As mentioned in section 1, the stochastic evolution blurs the classical relationship between a comoving scale and the value of the inflaton field when that scale crosses out the Hubble radius. This led to the development of different approximation schemes.

The framework that most closely aligns with the idea of stochastic trees is the one presented in Ref. [69], which relies on the so-called “large-volume approximation” that can be summarized as follows. Let us consider a region $\mathcal{B}$ of radius R within the end-of-inflation hypersurface, whose parent patch during inflation is denoted as $\mathcal{P}_*$. Field values in the latter are denoted Φ_* , where the bold notation indicates that several fields may be at play. The physical volume of that region is given by

$$V = \frac{4}{3}\pi R^3 = \int_{\mathcal{P}_*} e^{3\mathcal{N}_{\mathcal{P}_*}(\vec{x})} d^3\vec{x}, \quad (5.2)$$

where $\vec{x}$ denotes the comoving coordinate of each spatial point inside the patch $\mathcal{P}_*$, and $\mathcal{N}_{\mathcal{P}_*}(\vec{x})$ is the first-passage time of the worldline attached to $\vec{x}$ through the end-of-inflation hypersurface.

If $R \gg (\sigma H_{\text{end}})^{-1}$, the volume of $\mathcal{B}$ is much larger than the volume of the parent patch $V_* = 4/3\pi[\sigma H(\Phi_*)]^{-3}$, which, in the language of stochastic trees, corresponds to considering configurations in which a certain parent node generates a large number of leaves. In this case, according to the central limit theorem, one may replace $e^{3\mathcal{N}(\vec{x})}$ by its mean value

$$e^{3\mathcal{N}(\vec{x})} \rightarrow \langle e^{3\mathcal{N}_{\Phi_*}} \rangle = \int P_{\text{FPT},\Phi_*}(\mathcal{N}) e^{3\mathcal{N}} d\mathcal{N}, \quad (5.3)$$

where P_{FPT,Φ_*} denotes the first-passage time distribution with initial condition Φ_* .

As a result, the volume distribution takes the form

$$P(V|\Phi_*) \simeq \delta_D(V - V_* \langle e^{3\mathcal{N}_{\Phi_*}} \rangle), \quad (5.4)$$

where δ_D is the Dirac distribution, and the same holds for the volume-averaged expansion

$$W = \frac{V_*}{V} \int_{\mathcal{P}_*} e^{3\mathcal{N}_{\mathcal{P}_*}(\vec{x})} \mathcal{N}_{\mathcal{P}_*}(\vec{x}) d^3x \rightarrow \frac{V_*}{V} \langle \mathcal{N}_{\Phi_*} e^{3\mathcal{N}_{\Phi_*}} \rangle. \quad (5.5)$$

This allows one to establish a correspondence between field values during inflation and physical scales at the end-of-inflation hypersurface, since $V_* \langle e^{3\mathcal{N}_{\Phi_*}} \rangle = V \propto R^3$. Further considering single-clock models of inflation, hypersurfaces of constant forward volume reduce to single points in field phase space, and therefore backward-field values become deterministic quantities.

The large-volume approximation thus differs from the backward approximation [67, 68], where a single representative trajectory ending on the region $\mathcal{B}$ is considered to approximate the spatial average of the amount of expansion within that region, and which gives rise to a probability distribution for backward fields. They may be seen as approximations in opposite regimes.

Within the large-volume approximation framework, and in single-clock models of inflation like the one considered in this work, eq. (2.6) can be approximated as

$$\zeta_R \simeq \mathcal{N}_{\phi_0 \rightarrow \phi_*} + \langle \mathcal{N}_{\phi_*} \rangle_V - \langle \mathcal{N}_{\phi_0} \rangle_V, \quad (5.6)$$

where ϕ_0, ϕ_* are the inflaton values at nodes 1 and i respectively and R is the physical radius of the region emerging from node i . In the above expression, the first term denotes a first-passage time, while the latter two terms are volume-weighted stochastic averages, $\langle \mathcal{N}_{\phi} \rangle_V = \int_0^\infty d\mathcal{N} \mathcal{N} P_{\text{FPT},\phi}^V(\mathcal{N})$. This ultimately leads to expressing the one-point distribution of the curvature perturbation, coarse-grained over a region of size R on the end-of-inflation hypersurface, as

$$P(\zeta_R) = P_{\text{FPT},\phi_0 \rightarrow \phi_*}^V(\zeta_R - \langle \mathcal{N}_{\phi_*} \rangle_V + \langle \mathcal{N}_{\phi_0} \rangle_V). \quad (5.7)$$

Let us now see how predictions of the large-volume approximation compare with the results from stochastic trees, focusing on quantities relevant for the study of primordial black holes.

Although, as explained in section 3, the curvature perturbation is not the optimal cosmological field when it comes to discussing PBHs, let us consider it for the sake of illustration, and let us assume that black holes form when it exceeds a certain critical value, if coarse

grained over a region of a certain size R . In the Press-Schechter approximation [115], the abundance of PBHs is determined by the integral

$$P[\zeta_R > \zeta_c] = \int_{\zeta_c}^{\infty} P(\zeta_R) d\zeta_R, \quad (5.8)$$

where ζ_c is the PBH formation threshold, here assumed $\zeta_c = 1$. The above is the probability of a certain region of size R to lie above the threshold, or equivalently, the probability to collapse into a structure of size at least R . The large-volume approximation provides access to the probability (5.8), from which the mass fraction and mass distribution of PBHs can thus be derived.

However, this approximation is valid only when certain conditions are met. The main requirement is that large volumes are generated, in order for the central-limit replacement (5.3) to be performed. In other words, the mean volume emerging from a given backward field value ϕ_* must be much larger than one (in Hubble units). Yet, as shown in the left panel of fig. 6, see also eq. (4.2), for values of $\mu \leq 0.89$ considered with **FOREST** the mean volume is only of order a few. This implies that the large-volume approximation may not provide a reliable description of the model considered here.

In fact, even in models yielding large volumes, doubts may be cast regarding the ability of the large-volume approximation to describe extreme events such as PBH formation. The reason is that, as mentioned above, PBHs form when a subtree grows anomalously large. This can be readily seen in fig. 8b, where most PBH masses exceed by far the maximal mean volume (respectively given by 2.0, 2.9, 5.4, 10.2 and 34.2 for the values of μ displayed in fig. 8b, in increasing order, see eq. (4.2)). In this case, approximating volumes by their mean value is incorrect, and the replacement (5.3) is not justified. PBHs remain outlier events living in the tail of the distribution, and cannot be described directly with a central limit theorem.

As a consequence, here we do not compare directly the mass distributions, but instead we focus on the abundance of PBHs at the end of inflation. In practice, this corresponds to the limiting case where $\phi_* \rightarrow \phi_{\text{end}}$, such that eq. (5.6) becomes $\zeta_{R_{\text{end}}} \simeq \mathcal{N}_{\phi_0 \rightarrow \phi_{\text{end}}} - \langle \mathcal{N}_{\phi_0} \rangle_V$, where $R_{\text{end}} = (\sigma H_{\text{end}})^{-1}$, and we are coarse-graining in a single leaf. In the flat-well model, this leads to the one-point distribution [60]

$$P(\zeta_{R_{\text{end}}}) = -\frac{\pi}{2\mu^2} \cos(\sqrt{3}\mu) \vartheta'_2 \left\{ \frac{\pi}{2}, e^{-\frac{\pi^2}{\mu^2} \left[\zeta_{R_{\text{end}}} + \frac{1}{2\sqrt{3}} \mu \tan(\sqrt{3}\mu) \right]} \right\} e^{3\zeta_{R_{\text{end}}} + \frac{\sqrt{3}}{2} \mu \tan(\sqrt{3}\mu)}, \quad (5.9)$$

and the corresponding mass fraction, according to eq. (5.8), is shown as a function of μ as the orange line in fig. 8a.

This estimate exhibits the same behaviour as the full result obtained with **FOREST**, and is even in quantitative agreement to within about one order of magnitude. This is surprising, given the significant differences between the two approaches: **FOREST** addresses the cloud-in-cloud problem, which is not considered in the large-volume approximation. Moreover, the latter employs the curvature perturbation, rather than the coarse-shelled curvature perturbation, as the criterion for PBH formation. Finally, through eq. (5.6), the large-volume approximation also provides a prescription for the averaged expansion to subtract, namely $W_1 \simeq \langle \mathcal{N}_{\phi_0} \rangle_V$, though, as discussed above, in principle this is a non-local, tree-dependent quantity. Interestingly, this implies that qualitative estimates of the mass fraction generally made using the proxy (5.8) in the literature prove to be fairly reasonable estimates. They would fail to predict the details of the mass distributions (or even the typical masses at which




Figure 9. Different prescriptions for the branching time.

PBHs form), and presumably other quantities such as the typical distances over which PBHs are correlated, but for their overall abundance they seem reliable.

5.3 Discretization artefacts

In the framework introduced above, space is described as a tessellation of Hubble patches, and branches split in half at fixed time intervals ΔN . Stochastic trees are thus objects that are discrete both over space and time. This implies that spacetime properties below and around the Hubble scale are not properly captured by stochastic trees, which can only provide insight into the structure of inflating spacetimes over long times and large distances. In this section, we discuss artefacts coming from the discretization procedure, check that they become indeed suppressed at large scales and show that we can reproduce known results from quantum field theory in de Sitter.

For the sake of simplicity, we focus on a light test scalar field in a fixed binary tree.⁶ The field being test, inflation ends at the same time in all the branches, i.e. the tree is perfectly balanced. This makes it easier to relate topological distances in the tree and physical distances on the final hypersurface. Although it may be seen as restrictive, it will allow us to show explicitly how discretization artefacts disappear at large scales, and to argue why this is in fact a generic result.

5.3.1 Branching times

In the elementary vertex depicted in fig. 1, as soon as a patch grows larger than twice the Hubble volume, it is divided into two independent patches. In practice however, the two children patches do not become independent instantaneously, since the decay of gradient interactions at super-Hubble scales is a gradual process.

In stochastic trees, branching is an instantaneous event, but the precise time at which it takes place is not unambiguous. For instance, instead of splitting a node i and letting the branches leading to the nodes ℓ and m grow independently, one could let the node i grow to twice its volume and then split it. More generally, one could consider a hybrid process where the node i grows up to a time $\alpha\Delta N$, then splits, and the two child branches are evolved independently for $(1 - \alpha)\Delta N$, where $0 \leq \alpha \leq 1$. These three possibilities are depicted in fig. 9. The “split-then-grow” procedure corresponds to $\alpha = 0$, while the “grow-then-split” procedure corresponds to $\alpha = 1$.⁷

⁶Instead of a binary tree, one may consider trees with more branches attached to each node (for instance, trees where Hubble patches split into 8 patches after $\Delta N = \ln(2)$ e -folds in 3 dimensions). This would not change the nature of the arguments developed here.

⁷Note that eq. (2.3) was written with the “split-then-grow” convention for explicitness.




Figure 10. Example tree considered in section 5.3.

In order to show explicitly that large-scale properties do not depend on the detailed choice of the α parameter, let us consider the tree depicted in fig. 10. Note that a subscript T signals quantities associated with the whole tree. After a time $N_T = q_T \Delta N$ has elapsed from the parent node 1, consider two nodes i and j separated by a tree distance $d = |j - i|$. By tree distance, we mean a one-dimensional distance on the final hypersurface in the representation of the binary tree in fig. 10, where two sibling nodes are separated by $d = 1$. The maximal such distance one can consider is $D = 2^{q_T} - 1$.

The paths leading to i and j coincide until node n_* , and then become independent. In the example depicted in fig. 10, for $i = 20$ and $j = 23$, one has $q_T = 4$, $d = 3$ and the last common ancestor has label $n_* = 5$. First, consider the case where, from n_* , the path leading to i always proceeds leftwards, while the path leading to j always proceeds rightwards (we will show how to go beyond this assumption in section 5.3.2). In this case, the last common ancestor has label $n_* = i/(d + 1)$.⁸ The duration of the common path between nodes i and j is thus

$$N_* = N_T - \Delta N_* \quad \text{where} \quad \Delta N_* = \left[\frac{\ln(d + 1)}{\ln(2)} - \alpha \right] \Delta N. \quad (5.10)$$

Let us now denote the solution of the Fokker-Planck equation associated to eq. (2.1) by $P(\phi|\phi_{\text{in}}, N)$. This is the probability that, starting from ϕ_{in} , the test field takes value ϕ at time N . The joint probability for the field values in the nodes i and j is given by

$$P(\phi_i, \phi_j) = \int d\phi_* P(\phi_*|\phi_1, N_*) P(\phi_i|\phi_*, \Delta N_*) P(\phi_j|\phi_*, \Delta N_*) \quad (5.11)$$

where ϕ_* denotes the field value at the branching point. When $1 \ll d \ll D$, changing α amounts to changing N_* and ΔN_* by a small amount and this has little effect on the correlations between ϕ_i and ϕ_j . This can be more easily seen by rewriting the above result in terms of the three-dimensional physical distance d_P between the nodes i and j . The set of nodes $i, i + 1, \dots, j - 1, j$, represents a volume $(d + 1)V_\sigma = 4\pi(d_P/2)^3/3$, which leads to the following relationship between d and d_P

$$d + 1 = \left(\frac{d_P \sigma H}{2} \right)^3. \quad (5.12)$$

⁸More generally, it is convenient to use the binary representation of the nodes to identify the last common ancestor. For example, $i = 20 = 10100_2$ and $j = 23 = 10111_2$, their last common ancestor is the longest sequence of identical bits in the two numbers, so $101_2 = 5$.

Hence eq. (5.10) can be rewritten as

$$\Delta N_* = \ln(Hd_P) + \ln\left(2^{-1-\frac{\alpha}{3}}\sigma\right). \quad (5.13)$$

As a consequence, since $P(\phi_i, \phi_j)$ depends on α only through ΔN_* , α can be entirely reabsorbed in the definition of σ , hence the choice of the branching time is degenerate with the choice of the coarse-graining scale.

For explicitness, let us consider a light test field with $V(\phi) = m^2\phi^2/2$ in a de-Sitter universe with Hubble parameter H . In that case the stochastic problem has Gaussian solution [116]

$$P(\phi|\phi_{\text{in}}, N) = \frac{e^{-\frac{[\phi - \bar{\phi}(N, \phi_{\text{in}})]^2}{2s^2(N)}}}{\sqrt{2\pi s^2(N)}} \quad (5.14)$$

where

$$\bar{\phi}(N, \phi_{\text{in}}) = \phi_{\text{in}} e^{-\frac{m^2}{3H^2}N} \quad \text{and} \quad s^2(N) = \frac{3H^4}{8\pi^2 m^2} \left(1 - e^{-\frac{2m^2}{3H^2}N}\right). \quad (5.15)$$

From this expression, the two-point distribution function can be obtained with eq. (5.11) and is found to be also Gaussian,

$$P(\phi_i, \phi_j) = \frac{1}{\sqrt{(2\pi)^2 \det \Sigma}} e^{-\frac{1}{2}(\Delta\phi_i, \Delta\phi_j) \cdot \Sigma^{-1} \cdot \begin{pmatrix} \Delta\phi_i \\ \Delta\phi_j \end{pmatrix}} \quad (5.16)$$

where we have introduced $\Delta\phi_i = \phi_i - \bar{\phi}(N_T, \phi_1)$ and the covariance matrix Σ reads

$$\Sigma_{ii} = \langle \Delta\phi_i^2 \rangle = \Sigma_{jj} = \langle \phi_j^2 \rangle = s^2(N_T), \quad (5.17)$$

$$\Sigma_{ij} = \langle \Delta\phi_i \Delta\phi_j \rangle = s^2(N_*) e^{-\frac{2m^2}{3H^2}\Delta N_*}. \quad (5.18)$$

As mentioned above, changing α yields corrections that are suppressed by $\alpha/\ln(d)$ and which can thus be neglected at large distances (when $N_* \gg H^2/m^2$, these corrections are further suppressed by m^2/H^2).

In passing, it is worth noting that the two-point correlation function of a free test field in a de-Sitter background can be computed in full quantum-field theory [117–119]. Once proper renormalization is performed, at late time one finds $\Sigma_{ii} = 3H^4/(8\pi^2 m^2)$ in the coincident limit and $\Sigma_{ij} \propto (Hd_P)^{-\frac{2m^2}{3H^2}}$ at large distance, which matches the above results.

5.3.2 Branching surfaces

When space is described as a tessellation of Hubble patches, one has to introduce hypersurfaces that separate the patches, but the position and the geometry of those hypersurfaces is somewhat arbitrary. For instance, two nearby points located on each side of a branching hypersurface may be less distant than the Hubble radius, so treating them as disconnected is not strictly valid. More fundamentally, branching hypersurfaces single out specific regions in physical space, which breaks homogeneity of spacetime. This creates artefacts in our discretization scheme.

This can be readily seen by considering neighbour leaves in the tree represented in fig. 10, and their degree of parenthood $\Delta N_* = \Delta q_* \Delta N$. The leaves 16 and 17 are such that

$\Delta q_* = 1$, i.e. they are sibling leaves. However, leaves 17 and 18 are such that $\Delta q_* = 2$, i.e. they are cousin leaves, and leaves 19 and 20 have $\Delta q_* = 3$ while leaves 23 and 24 have $\Delta q_* = 4$. In other words, $\Delta N_*(i, j)$ is not just a function of $|i - j|$, hence $P(\phi_i, \phi_j)$ is not a function of $|i - j|$ either, which manifestly breaks space-translation invariance. The problem is that the topological distance Δq_* in the tree is not directly mapped to the geometrical distance d on the final hypersurface.

As we shall now argue, one way to solve this issue is to perform ensemble averages over the leaves (in addition to performing stochastic averages over the tree realizations). For instance, the two-point correlation function of the field fluctuation at physical distance d_P shall be defined as the ensemble average over all pairs of two leaves distant by d on the final hypersurface (there are $2^{q_T} - d$ such pairs)

$$\Sigma(d_P) = \frac{1}{2^{q_T} - d} \sum_{i=2^{q_T}}^{2^{q_T+1}-d-1} \Sigma_{i,i+d} \quad (5.19)$$

where d and d_P are related through eq. (5.12). In the example discussed in section 5.3.1, this can be computed as follows. Amongst the $2^{q_T} - d$ pairs of leaves separated by d at time $N_T = q_T \Delta N$, let $\beta(d, q)$ denote the number of pairs whose Δq_* is such that $\Delta q_* = q_T - q$. For instance, when $d = 1$, there is only one pair of leaves such that $\Delta q_* = q_T$, namely the central pair of leaves. Likewise, there are two pairs of leaves such that $\Delta q_* = q_T - 1$, so on and so forth, hence $\beta(1, q) = 2^q$. One can show that in general, the number of pairs reads

$$\beta(d, q) = \begin{cases} 2^q d & \text{if } d \leq 2^{q_T - q - 1} \\ 2^{q_T} - 2^q d & \text{if } 2^{q_T - q - 1} \leq d \leq 2^{q_T - q} \\ 0 & \text{if } d \geq 2^{q_T - q} \end{cases} \quad (5.20)$$

A demonstration involving the binary representation of the nodes is given in section C.

This counting function is maximal when $d = 2^{q_T - q - 1} = 2^{\Delta N_* / \Delta N - 1}$, which is close to the configuration (5.10) to which the analysis of section 5.3.1 was restricted. However, all values of Δq_* are encountered in a given tree, and they contribute to the spatially averaged correlation functions.

Since $\Sigma_{i,i+d}$ is only a function of ΔN_* given by eq. (5.18), one can organize eq. (5.19) according to

$$\Sigma(d_P) = \frac{1}{2^{q_T} - d} \sum_{q=0}^{q_T-1} \beta(d, q) \Sigma[(q_T - q) \Delta N] \quad (5.21)$$

Inserting eqs. (5.18) and (5.20) into eq. (5.21) leads to

$$\Sigma(d_P) = \frac{3H^4}{8\pi^2 m^2} \frac{e^{-aq_T}}{2^{q_T} - d} \left[2^{q_T} (e^{aq_*} - 1) - \frac{2(e^a - 1)d(2^{q_*} e^{aq_*} - 1)}{2e^a - 1} \right] \quad (5.22)$$

where we have introduced $a = 2m^2 \Delta N / (3H^2)$ and $q_* = \lfloor q_T - \ln(d) / \ln(2) \rfloor$. At large distances, $1 \ll d \ll D$, the above can be expanded in the limit $q_* \gg 1$, and it reduces to

$$\Sigma(d_P) \simeq \frac{3H^4}{8\pi^2 m^2} \left[\frac{e^{3a}}{2e^a - 1} - e^{3a - aq_*} \right] (\sigma H d_P)^{-\frac{2m^2}{3H^2}}, \quad (5.23)$$

where we have further used eq. (5.12) to relate d and d_P . In the limit $a \ll 1$, which is required for the slow-roll treatment of the test field to be valid, the term in the square brackets can be

written as $1 - e^{-\frac{2m^2}{3H^2}N_*}$, where we have used eq. (5.13) to write $q_* = N_*/\Delta N$.⁹ This coincides with the result obtained in section 5.3.1, which is also consistent with quantum-field-theory calculations.

This shows that, even though $\Sigma(d_P)$ receives contribution from pairs of leaves with $\ln(\sigma H d_P) \leq \Delta N_*(i, j) \leq N_T$, in practice those with $\Delta N_*(i, j) \simeq \ln(\sigma H d_P)$ dominate, and the correct result is recovered at large distance. The discretization scheme could be improved, for instance by randomly drawing the position of the branching hypersurfaces within the patches. This was proposed in Ref. [71] and makes the trees statistically homogeneous indeed. However, it never produces more “correctly correlated” pairs than the standard tree, hence it does not improve the convergence towards the correct result at large distance.

Note that in this work, we have not computed the two-point correlation functions discussed above, but rather consider one-point distributions of the coarse-grained or coarse-shelled curvature perturbation. Since the method outlined in section 3.1 allows one to perform coarse-shelling between two nodes that are direct siblings only, only such pairs of nodes are accounted for in our stochastic averages. Given that these pairs are “correctly correlated” in the language developed above, we expect our results for the statistics of PBHs to be even more robust to discretisation artefacts than the two-point correlation function. In any case, in the simple example discussed here (a test free field), Gaussian statistics allows one to express the one-point distribution entirely in terms of the two-point correlation functions, hence the convergence of the latter under proper averaging ensures the convergence of the former.

6 Conclusion

In this work, we have implemented the stochastic- δN formalism on stochastic trees, conceptualizing the inflationary expansion as a branching process.

In essence, inflationary spacetime can be imagined as a branching tree, starting with a single Hubble patch that after a certain amount of stochastic evolution, described by Langevin equations, doubles in volume, giving rise to two new nodes, each representing a descendant Hubble patch then evolving independently. This process repeats recursively, with each vertex splitting into further branches until inflation ends. The final nodes, where inflation terminates, form the leaves of the tree, where the curvature perturbation and other cosmological fields are measured. The statistical properties of these fields are embedded in the tree structure: the distance between two leaves corresponds to the recursive level below which their paths separated, i.e. to the depth of their latest common ancestor.

This intricate structure can be numerically realized with **FOREST** [101], a cutting-edge, parallel code that generates and scans vast populations of stochastic trees to deliver relevant statistics. We showed how **FOREST** returns the first-passage time, in terms of e -folds, through the end-of-inflation hypersurface, averaged over the volume emerging from the initial node, and how the coarse-grained curvature perturbation at the Hubble scale at the end of inflation, or at any arbitrary scale, can be reconstructed.

Stochastic trees tessellate spacetime into Hubble patches that evolve and emerge from branches, making them discrete objects in both time and space. We presented how the information contained in stochastic trees can be translated into real-space maps of the end-of-inflation hypersurface, unfolding a comoving volume by progressively dividing it along

⁹One may also absorb the a -dependence in the redefinition of σ .

orthogonal planes with alternating orientations at each volume-doubling time. Despite arbitrariness in the convention employed to discretize space and in the choice of a branching time, translation invariance and consistency with quantum field theory in curved spacetime are recovered at large distances when large populations of trees are considered.

Stochastic trees are also ideal tools for harvesting primordial black holes. In addition to the curvature perturbation, we have shown how they allow for the reconstruction of the coarse-shelled curvature perturbation, namely, the curvature perturbation averaged between two concentric spheres. This serves as a proxy for the compaction function, a cosmological field that is better suited for studying the formation of primordial black holes due to its local nature. We have shown how it can be viewed as an “unbalance index”: primordial black holes tend to form at the most asymmetric nodes of the tree. Moreover, the cloud-in-cloud problem, in which a hierarchy of overdense regions could lead to PBH formation inside larger PBHs, is directly implemented within **FOREST**, which first scans the smallest scales to check for PBH formation, propagating from the leaves to the root across the tree and discarding black holes when a larger one is found in the same path.

We applied **FOREST** to a simple toy model, namely the “quantum well”, which had previously been studied using other methods. We reconstructed the probability distribution for the final volume and for the volume-averaged expansion. Both exhibit a transient power-law behaviour turning to exponential in the far tail. Tails become heavier as μ – the order parameter measuring the amount of quantum diffusion – increases, up to a critical value which sets the onset of eternal inflation. Close to the critical point, the volume distribution features a power-law tail and the mean volume diverges.

When scanning the set of leaves, we also observed that exponential tails appear in their first-passage-time distribution. Remarkably, this distribution is perfectly matched by the volume-weighted first-passage-time distribution drawn from independent Langevin processes, which can be computed analytically and which is central to various formulations and approximations of stochastic inflation in terms of first-passage-time analysis.

We have reconstructed the mass distribution of PBHs in the quantum well. We found that the typical mass of PBHs increases with μ and that the distribution itself exhibits mild power laws terminated by exponential tails and that become broader as μ increases. Near the critical point, nearly all the universe is contained into PBHs.

Finally, we discussed how these results compare with those obtained using the previously developed large-volume approximation, which replaces the final volume with its stochastic average. We found that, for the total PBH mass fraction, the simple Press-Schechter estimate based on the curvature perturbation compares well with the full numerical results (in spite of neglecting cloud in cloud, using the large-volume approximation and relying on the curvature perturbation instead of the density contrast). This shows that the usual estimates of the PBH abundance performed in the stochastic-inflation literature provide reasonable approximations.

However, we also uncovered that the details of the mass distribution cannot be correctly captured by these approximations. The main reason is that the large-volume approximation connects physical scales at the end of inflation to field values during inflation via the mean volume. Yet, PBHs form at the root of subtrees that generate an anomalously large volume, very far from the mean volume indeed. As a consequence, they form at scales and masses that are much larger than what the large-volume approximation predicts. Since PBHs are extreme phenomena living in the “tail” of distribution functions, it should probably not come as a surprise that they cannot be described by a central limit theorem. Note that this does

not preclude the large-volume approximation to perform well for bulk statistics such as the two- or three-point correlation functions, and we defer its comparison with stochastic trees in models producing large volumes to future investigations.

Before concluding, let us highlight several points that undoubtedly deserve attention and offer potential for future development.

By simulating stochastic trees, **FOREST** reconstructs statistics through direct sampling, enabling, for instance, the investigation of the first-passage-time distribution down to values of probability as small as 10^{-10} , improving on the numerical efficiency of direct Langevin simulations and lattice codes. To investigate the tails of the distributions even more efficiently, methods based on importance sampling can be employed [97, 120], and we leave this possibility for future versions of **FOREST**.

In this work, we focused on the simplest quantum-well toy model, which, despite its unrealistic simplicity, served as an ideal playground for testing and demonstrating the power of **FOREST**. We also confined our study to slow-roll inflation. However, achieving enhanced perturbations at small scales that could lead to primordial black hole formation typically requires departing from the slow-roll regime and incorporating non-attractor phases, such as ultra-slow roll. We thus plan to extend **FOREST** to higher-dimensional phase spaces that include the field’s momentum, and to investigate more realistic models in the future.

In this study, we employed **FOREST** to reconstruct one-point statistics and make predictions for the abundance and mass distribution of PBHs. However, its potential extends beyond this, being capable to tackle multiple-point statistics. The information embedded in the topological structure of stochastic trees can map real-space correlations of cosmological fields at the end of inflation. This would allow for a non-perturbative reconstruction of the power spectrum and higher-order correlations, and offer a way to explore the spatial distribution of PBHs at formation. In this way we could check whether PBHs created from the collapse of large non-Gaussian fluctuations are naturally born clustered, as suggested by recent analytical investigations [69], and if they display an exclusion effect on small-scales [121]. This is the subject of a forthcoming publication.

Finally, N -body simulations heavily rely on initial conditions drawn from inflationary models. In this context, stochastic trees emerge as an ideal tool to provide non-perturbative initial conditions in the form of ready-to-use maps. This will help us understanding how PBHs [122], and more generally heavy-tailed statistics [123], affect structure formation. This offers a fascinating perspective for bridging the physics of the early and late universe through comprehensive, full-scale simulations.

Acknowledgments

The authors thank Marie-Paule Faber, Christophe Ringeval and Juan García-Bellido for insightful discussions. Pierre Auclair is a Postdoctoral Researcher of the Fonds de la Recherche Scientifique – FNRS. Baptiste Blachier is a Research Fellow of the Fonds de la Recherche Scientifique – FNRS. Chiara Animalì is supported by the ESA Belgian Federal PRODEX Grant N°4000143201. Computational resources have been provided by the CURL’s development cluster, the supercomputing facilities of the Université catholique de Louvain (CISM/UCL) and the Consortium des Équipements de Calcul Intensif en Fédération Wallonie Bruxelles (CÉCI) funded by the Fond de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under convention 2.5020.11 and by the Walloon Region.

One node			Two nodes (MPI)		
# of threads	Wall-time	Speed-up	# of threads	Wall-time	Speed-up
1	02:23:30	—	—	—	—
2	01:29:46	1.6	2×1	01:13:56	1.9
4	00:44:26	3.2	2×2	00:44:30	3.2
8	00:22:15	6.5	2×4	00:22:12	6.5
16	00:11:14	13	2×8	00:11:10	13
32	00:05:45	25	2×16	00:05:38	25
64	00:03:00	48	2×32	00:02:52	50
128	00:01:48	80	2×64	00:01:30	95
256	00:00:56	153	2×128	00:00:54	159
—	—	—	2×256	00:00:28	308

Table 1. Strong-scaling test of **FOREST** in the flat-well model with $\mu = 0.8$ and 10^8 trees. We test the code on two EPYC CPUs communicating through MPI. On each computing node, we use OpenMP to distribute the work on the 128 cores / 256 threads. Times are reported in the format HH:MM:SS, which is hours, minutes and seconds.

# of threads	# of trees	Wall-time	Efficiency
1	10^6	01:27	—
2	2×10^6	01:47	81%
4	4×10^6	01:47	81%
8	8×10^6	01:47	81%
16	16×10^6	01:49	79%
32	32×10^6	01:51	79%
64	64×10^6	01:55	75%
128	128×10^6	02:17	63%
256	256×10^6	02:25	60%

Table 2. Weak-scaling test of **FOREST** in the flat-well model with $\mu = 0.8$ on a single EPYC CPU containing 128 cores and 256 threads. Times are reported in the format MM:SS, which is minutes and seconds.

A Scaling tests

The main task of **FOREST** is to simulate a large population of independent stochastic trees to sample the statistics of the end-of-inflation hypersurface. Our strategy to parallelize the code is straightforward: we first distribute our samples evenly across the different nodes on the cluster using MPI. On each node, we use OpenMP to distribute the workload on the different threads. This hybrid approach using MPI+OpenMP is becoming a standard, and combines the advantages of MPI – communication between an array of distinct computers – and of OpenMP – better optimization on machines with shared memory and advanced load balancing. Each thread contains its own copy of the variables and objects in the code to avoid a memory access bottleneck. Threads run independently without any communication and results are combined at the end of the simulation using a sum reduction.

We report *strong* scaling results in table 1, in which we increase the number of threads for a given task. Strong scaling is an indicator of whether the wall-time of a given simulation can be reduced by increasing the number of parallel processes. In the ideal limit, the speed-up




Figure 11. Convergence test for different values of a fixed step δN_{fix} in the flat quantum well. Black dashed line corresponds to the analytical formula of eq. (5.1).

factor should equal the number of processes. As a benchmark, we use the flat-well potential with $\mu = 0.8$ and 10^8 samples. We find that, apart from a small overhead due to the usage of OpenMP, our code scales consistently from 2 to 64 threads per node. In contrast, combining different nodes with MPI does not produce any overhead. The reason for such a difference in behaviour is that on a single node, OpenMP has to distribute the shared memory to the different threads. Depending on the node topology and on the memory allocation strategy, the various processes on the same node can see their memory access times increase significantly. At 128 threads per node and above, we still gain a substantial speed-up but at a slower rate. This behaviour is expected as we saturate the core count of our nodes and may encounter hardware limitations such as power and temperature.

Then, we report *weak* scaling results in table 2, in which we increase both the number of threads and the size of the population at the same rate. Weak scaling is an indicator of whether larger populations can be simulated efficiently by increasing the number of parallel processes. In the ideal limit, all the tests should have the same wall-time, and we define the efficiency as the ratio between the ideal and the real wall-time. As a benchmark, we again use the flat-well potential with $\mu = 0.8$ and increasing the number of samples from 10^6 . The conclusions are similar to our strong-scaling tests. We experience a small drop in efficiency due to OpenMP and a stable efficiency from 2 to 64 cores. Above this threshold, we also note a drop in efficiency.

B Convergence tests

At the core of **FOREST**, we solve the Stochastic Differential Equation (SDE) (2.1) using the Euler-Maruyama method. As briefly mentioned in section 2.1, we decided to use a varying step δN such that there is at most a $\kappa\sigma$ probability to cross the end-of-inflation hypersurface,

$$\delta N = \min \left\{ \delta N_{\text{fix}}, \frac{3 [2\pi M_{\text{Pl}}(\phi - \phi_{\text{end}})]^2}{\kappa V(\phi)} \right\}, \quad (\text{B.1})$$

in which δN_{fix} is a fixed value. The purpose of this appendix is to show, with two quantitative examples, the limitations of using a fixed time step $\delta N = \delta N_{\text{fix}}$ and why taking $\kappa = 5$ is




Figure 12. Convergence test for different values of a varying step δN following eq. (B.1) in the flat quantum well. For illustrative purposes, we add the constraint of eq. (B.1) to the worst corresponding scenario of fig. 11. That is $\delta N_{\text{fix}} = \Delta N/10$ for the left panel and $\delta N_{\text{fix}} = \Delta N/100$ on the right panel. Black dashed line corresponds to the analytical formula of eq. (5.1).

optimal. As benchmarks, we use the flat-well model, with $\mu = 0.6, 0.8$. Despite its simplicity, the flat-well model can be considered as the limiting scenario in which quantum diffusion dominates completely the dynamics of the system, hence a perfect candidate to study the convergence of our SDE solver. We focus on the distribution of the volume-averaged elapsed number of e -folds, $P_{\text{FPT}, \Phi_0}^V(\mathcal{N})$, for two reasons. First, because a value of δN that is too large has a very visible impact on this distribution. Second, because we have an exact analytical formula (5.1) at our disposal to compare our simulations with.

In fig. 11, we show the effect of a fixed step δN_{fix} given as a fraction of the splitting time $\Delta N = \ln(2)/3$. If the value of δN is too large, then the solver is blind to trajectories that may have exited the quantum well and re-entered during δN . Therefore, this tends to over-estimate the first-passage time, as can be seen in fig. 11.

In fig. 12, we show how our prescription of eq. (B.1) converges when added to the worst-case scenario of fig. 11. In both our benchmarks, we find that the prescription (B.1) substantially improves the convergence of the distribution even for small values of κ . This convergence is robust and less sensitive to the specific value of μ one considers.

Notably, adding our condition (B.1) to an existing fixed time step does not increase the computation time by more than a factor of 2 or 3. This should be compared with the cost of dividing the size of the time step by several orders of magnitude, and hence multiplying the wall-time by the same amount. Qualitatively for $\mu = 0.8$, it takes about the same time to obtain the best curve in fig. 11b with 10^5 trees and the best curve in fig. 12b with 10^7 trees, a performance boost of a factor 100. The reason for such a performance gain is that reducing the time step increases the resolution of the SDE solver across the whole quantum well. To the contrary, eq. (B.1) increases the resolution close to the end-of-inflation hypersurface only. Eventually, what we are interested in is the duration of inflation in the different patches and not the full history of the Langevin trajectories.

In production, we impose $\delta N_{\text{fix}} = 10^{-3}$, a value more stringent than the two benchmarks presented in this appendix. The remarkable agreement between our large numerical simulations and the corresponding analytical formula can be seen in fig. 7.

$$\begin{array}{rcl}
i & & \boxed{1} \boxed{101} \boxed{001} \\
i + d & & \boxed{1} \boxed{101} \boxed{111}
\end{array}$$

Figure 13. Demonstration for eq. (5.20) using of the binary representation of nodes i and $j = i + d$. In this example, $i = 1101001_2 = 105$ and $j = 1101111_2 = 111$ are distant by $d = 6$. One has $q = 3$ (three green bits) and $q_T - q = 3$ (three red bits).

C Counting function

We consider the balanced binary trees studied in section 5.3. The goal of this appendix is to compute the counting function $\beta(d, q)$, which corresponds to the number of pairs of leaves i and j such that $j = i + d$ and such that the last common ancestor to the two leaves lies at $N = q\Delta N$ from the root.

Following the example displayed in fig. 13, the binary representations of i and j are given by strings of bits that contain three segments. The first segment, displayed in blue in fig. 13, is a single bit, 1, which represents the root node. It is always present and equal to one.

The second segment, displayed in green, represents the common part of the paths leading to i and j , prior to their last common ancestor. Its contains q bits that are identical for i and j , there are thus 2^q possible green segments. In the tree representation, these correspond to the paths leading to the 2^q nodes located q levels below the root node.

The third segment, displayed in red, represents the part of the two paths that lies below the last common ancestor. It contains $q_T - q$ bits in the representation of i , which fully determine the bits in j since $j = i + d$. Therefore, there are at most $2^{q_T - q}$ possibilities for the red segment. However, two conditions need to be fulfilled by the red segment.

First, one must ensure that no green bit is being flipped when summing i and d , since by definition the green segments have to remain identical for i and j . This limits the last $q_T - q$ bits of i to be strictly smaller than $2^{q_T - q} - d$, since larger values would induce a bit flip. Therefore, there are $2^q \times (2^{q_T - q} - d) = 2^{q_T} - d2^q$ pairs sharing *at least* the same first $1 + q$ bits, that is having a common ancestor at $N = q\Delta N$. In terms of the counting function,

$$\sum_{q' \geq q} \beta(d, q') = 2^{q_T} - d2^q. \quad (\text{C.1})$$

Second, we have imposed that i and j share the same path *at least* until a node at level q , but this does not guarantee this node is the *last* common ancestor. One must also ensure that i and j follow different branches below this node. In terms of the binary decomposition of i and j , we need to impose that the first bit of the red segment differs between i and j . Three cases need to be distinguished:

- If $d \geq 2^{q_T - q}$, then the binary representation of d contains more bits than the red segment, hence adding d to i necessarily flips bits in the green segment. Since this is impossible, $\beta(d, q) = 0$ in that case. Consistently, eq. (C.1) does not give a positive number of pairs in this case.
- If $d < 2^{q_T - q - 1}$, then we obtain the number of pairs having their last common ancestor at level q using eq. (C.1),

$$\beta(d, q) = \sum_{q' \geq q} \beta(d, q') - \sum_{q' \geq q+1} \beta(d, q') = 2^q d. \quad (\text{C.2})$$

- Finally, if $2^{q_T-q-1} \leq d < 2^{q_T-q}$, then it is not possible to have $q+1$ identical bits in the green segment, hence $\beta(d, q+1) = 0$ and

$$\beta(d, q) = \sum_{q' \geq q} \beta(d, q') = 2^{q_T} - d2^q. \quad (\text{C.3})$$

To summarize, we have found that

$$\beta(d, q) = \begin{cases} 2^q d & \text{if } d < 2^{q_T-q-1} \\ 2^{q_T} - d2^q & \text{if } 2^{q_T-q-1} \leq d < 2^{q_T-q} \\ 0 & \text{if } d \geq 2^{q_T-q} \end{cases}. \quad (\text{C.4})$$

This is identical to eq. (5.20), where strict inequalities on d have been changed to weak inequalities after checking that the counting function is continuous at the two pivotal points $d = 2^{q_T-q-1}$ and $d = 2^{q_T-q}$.

As a consistency check, one can compute the total number of pairs i and $j = i + d$, regardless of their topological distance q . This number is given by

$$\sum_{q=0}^{q_T-1} \beta(d, q) = \sum_{q=0}^{q_T-2-\lfloor \frac{\ln(d)}{\ln(2)} \rfloor} 2^q d + \sum_{q=q_T-1-\lfloor \frac{\ln(d)}{\ln(2)} \rfloor}^{q_T-1-\lfloor \frac{\ln(d)}{\ln(2)} \rfloor} (2^{q_T} - d2^q) + \sum_{q=q_T-\lfloor \frac{\ln(d)}{\ln(2)} \rfloor}^{q_T-1} 0 \quad (\text{C.5})$$

$$= 2^{q_T} - d, \quad (\text{C.6})$$

where the first sum is over terms of a geometric series and can thus be readily evaluated. One recovers the total $2^{q_T} - d$ pairs of leaves separated by d at time N_T .

References

- [1] SDSS collaboration, D. J. Eisenstein et al., *Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies*, *Astrophys. J.* **633** (2005) 560–574, [[astro-ph/0501171](#)].
- [2] PLANCK collaboration, N. Aghanim et al., *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, *Astron. Astrophys.* **641** (2020) A1, [[1807.06205](#)].
- [3] M. M. Ivanov, M. Simonović and M. Zaldarriaga, *Cosmological Parameters from the BOSS Galaxy Power Spectrum*, *JCAP* **05** (2020) 042, [[1909.05277](#)].
- [4] V. F. Mukhanov and G. V. Chibisov, *Quantum Fluctuations and a Nonsingular Universe*, *JETP Lett.* **33** (1981) 532–535.
- [5] V. F. Mukhanov and G. Chibisov, *The Vacuum energy and large scale structure of the universe*, *Sov. Phys. JETP* **56** (1982) 258–265.
- [6] A. A. Starobinsky, *Dynamics of Phase Transition in the New Inflationary Universe Scenario and Generation of Perturbations*, *Phys. Lett. B* **117** (1982) 175–178.
- [7] A. H. Guth and S. Pi, *Fluctuations in the New Inflationary Universe*, *Phys. Rev. Lett.* **49** (1982) 1110–1113.
- [8] S. Hawking, *The Development of Irregularities in a Single Bubble Inflationary Universe*, *Phys. Lett. B* **115** (1982) 295.
- [9] J. M. Bardeen, P. J. Steinhardt and M. S. Turner, *Spontaneous Creation of Almost Scale - Free Density Perturbations in an Inflationary Universe*, *Phys. Rev. D* **28** (1983) 679.

- [10] A. A. Starobinsky, *A New Type of Isotropic Cosmological Models Without Singularity*, *Adv. Ser. Astrophys. Cosmol.* **3** (1987) 130–133.
- [11] K. Sato, *First-order phase transition of a vacuum and the expansion of the Universe*, *Mon. Not. Roy. Astron. Soc.* **195** (1981) 467–479.
- [12] A. H. Guth, *The Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems*, *Adv. Ser. Astrophys. Cosmol.* **3** (1987) 139–148.
- [13] A. D. Linde, *A New Inflationary Universe Scenario: A Possible Solution of the Horizon, Flatness, Homogeneity, Isotropy and Primordial Monopole Problems*, *Adv. Ser. Astrophys. Cosmol.* **3** (1987) 149–153.
- [14] A. Albrecht and P. J. Steinhardt, *Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking*, *Adv. Ser. Astrophys. Cosmol.* **3** (1987) 158–161.
- [15] A. D. Linde, *Chaotic Inflation*, *Phys. Lett. B* **129** (1983) 177–181.
- [16] E. Asencio, I. Banik and P. Kroupa, *A massive blow for Λ CDM – the high redshift, mass, and collision velocity of the interacting galaxy cluster El Gordo contradicts concordance cosmology*, *Mon. Not. Roy. Astron. Soc.* **500** (2020) 5249–5267, [2012.03950].
- [17] A. H. Gonzalez, S. A. Stanford, M. Brodwin, C. Fedeli, A. Dey, P. R. M. Eisenhardt et al., *IDCS J1426.5+3508: Cosmological Implications of a Massive, Strong Lensing Cluster at $z = 1.75$* , *ApJ* **753** (July, 2012) 163, [1205.3788].
- [18] S. L. Finkelstein et al., *A Rapidly Star-forming Galaxy 700 Million Years After the Big Bang at $z=7.51$* , *Nature* **502** (2013) 524, [1310.6031].
- [19] A.-C. Eilers, R. A. Simcoe, M. Yue, R. Mackenzie, J. Matthee, D. Durovckova et al., *EIGER. III. JWST/NIRCam Observations of the Ultraluminous High-redshift Quasar J0100+2802*, *Astrophys. J.* **950** (2023) 68, [2211.16261].
- [20] C. T. Donnan, D. J. McLeod, J. S. Dunlop, R. J. McLure, A. C. Carnall, R. Begley et al., *The evolution of the galaxy UV luminosity function at redshifts $z \approx 8 - 15$ from deep JWST and ground-based near-infrared imaging*, *MNRAS* **518** (Feb., 2023) 6011–6040, [2207.12356].
- [21] Y. Harikane, M. Ouchi, M. Oguri, Y. Ono, K. Nakajima, Y. Isobe et al., *A Comprehensive Study of Galaxies at $z \sim 9-16$ Found in the Early JWST Data: Ultraviolet Luminosity Functions and Cosmic Star Formation History at the Pre-reionization Epoch*, *Astrophys. J. Suppl.* **265** (2023) , [2208.01612].
- [22] R. P. Naidu, P. A. Oesch, P. van Dokkum, E. J. Nelson, K. A. Suess, G. Brammer et al., *Two Remarkably Luminous Galaxy Candidates at $z \approx 10-12$ Revealed by JWST*, *ApJ* **940** (Nov., 2022) L14, [2207.09434].
- [23] S. L. Finkelstein, M. B. Bagley, P. Arrabal Haro, M. Dickinson, H. C. Ferguson, J. S. Kartaltepe et al., *A Long Time Ago in a Galaxy Far, Far Away: A Candidate $z \sim 12$ Galaxy in Early JWST CEERS Imaging*, *ApJ* **940** (Dec., 2022) L55, [2207.12474].
- [24] H. Atek, M. Shuntov, L. J. Furtak, J. Richard, J.-P. Kneib, G. Mahler et al., *Revealing galaxy candidates out to $z 16$ with JWST observations of the lensing cluster SMACS0723*, *MNRAS* **519** (Feb., 2023) 1201–1220, [2207.12338].
- [25] G. Roberts-Borsani, T. Morishita, T. Treu, G. Brammer, V. Strait, X. Wang et al., *Early Results from GLASS-JWST. I: Confirmation of Lensed $z \geq 7$ Lyman-break Galaxies behind the Abell 2744 Cluster with NIRISS*, *ApJ* **938** (Oct., 2022) L13, [2207.11387].
- [26] M. Castellano, A. Fontana, T. Treu, P. Santini, E. Merlin, N. Leethochawalit et al., *Early Results from GLASS-JWST. III. Galaxy Candidates at $z 9-15$* , *ApJ* **938** (Oct., 2022) L15, [2207.09436].

- [27] P. Santini, A. Fontana, M. Castellano, N. Leethochawalit, M. Trenti, T. Treu et al., *Early Results from GLASS-JWST. XI. Stellar Masses and Mass-to-light Ratio of $z \gtrsim 7$ Galaxies*, *ApJ* **942** (Jan., 2023) L27, [[2207.11379](#)].
- [28] N. J. Adams, C. J. Conselice, L. Ferreira, D. Austin, J. A. A. Trussler, I. Juodžbalis et al., *Discovery and properties of ultra-high redshift galaxies ($9 \lesssim z \lesssim 12$) in the JWST ERO SMACS 0723 Field*, *MNRAS* **518** (Jan., 2023) 4755–4766, [[2207.11217](#)].
- [29] A. Escrivà, F. Kuhnel and Y. Tada, *Primordial Black Holes*, [2211.05767](#).
- [30] B. Carr, S. Clesse, J. Garcia-Bellido, M. Hawkins and F. Kuhnel, *Observational evidence for primordial black holes: A positivist perspective*, *Phys. Rept.* **1054** (2024) 1–68, [[2306.03903](#)].
- [31] LISA COSMOLOGY WORKING GROUP collaboration, E. Bagui et al., *Primordial black holes and their gravitational-wave signatures*, [2310.19857](#).
- [32] K. N. Ananda, C. Clarkson and D. Wands, *The Cosmological gravitational wave background from primordial density perturbations*, *Phys. Rev. D* **75** (2007) 123518, [[gr-qc/0612013](#)].
- [33] D. Baumann, P. J. Steinhardt, K. Takahashi and K. Ichiki, *Gravitational Wave Spectrum Induced by Primordial Scalar Perturbations*, *Phys. Rev. D* **76** (2007) 084019, [[hep-th/0703290](#)].
- [34] EPTA, INPTA collaboration, J. Antoniadis et al., *The second data release from the European Pulsar Timing Array - IV. Implications for massive black holes, dark matter, and the early Universe*, *Astron. Astrophys.* **685** (2024) A94, [[2306.16227](#)].
- [35] NANOGrav COLLABORATION collaboration, A. Afzal et al., *The NANOGrav 15 yr Data Set: Search for Signals from New Physics*, *The Astrophysical Journal Letters* **951** (2023) L11, [[2306.16219](#)].
- [36] D. S. Salopek and J. R. Bond, *Nonlinear evolution of long wavelength metric fluctuations in inflationary models*, *Phys. Rev. D* **42** (1990) 3936–3962.
- [37] M. Sasaki and E. D. Stewart, *A General analytic formula for the spectral index of the density perturbations produced during inflation*, *Prog. Theor. Phys.* **95** (1996) 71–78, [[astro-ph/9507001](#)].
- [38] D. Wands, K. A. Malik, D. H. Lyth and A. R. Liddle, *A New approach to the evolution of cosmological perturbations on large scales*, *Phys. Rev. D* **62** (2000) 043527, [[astro-ph/0003278](#)].
- [39] D. H. Lyth and D. Wands, *Conserved cosmological perturbations*, *Phys. Rev. D* **68** (2003) 103515, [[astro-ph/0306498](#)].
- [40] G. I. Rigopoulos and E. P. S. Shellard, *The separate universe approach and the evolution of nonlinear superhorizon cosmological perturbations*, *Phys. Rev. D* **68** (2003) 123518, [[astro-ph/0306620](#)].
- [41] D. H. Lyth and Y. Rodriguez, *The Inflationary prediction for primordial non-Gaussianity*, *Phys. Rev. Lett.* **95** (2005) 121302, [[astro-ph/0504045](#)].
- [42] D. Artigas, J. Grain and V. Vennin, *Hamiltonian formalism for cosmological perturbations: the separate-universe approach*, *JCAP* **02** (2022) 001, [[2110.11720](#)].
- [43] J. H. P. Jackson, H. Assadullahi, A. D. Gow, K. Koyama, V. Vennin and D. Wands, *The separate-universe approach and sudden transitions during inflation*, *JCAP* **05** (2024) 053, [[2311.03281](#)].
- [44] A. A. Starobinsky, *Multicomponent de Sitter (Inflationary) Stages and the Generation of Perturbations*, *JETP Lett.* **42** (1985) 152–155.
- [45] M. Sasaki and T. Tanaka, *Superhorizon scale dynamics of multiscalar inflation*, *Prog. Theor. Phys.* **99** (1998) 763–782, [[gr-qc/9801017](#)].

- [46] D. H. Lyth, K. A. Malik and M. Sasaki, *A General proof of the conservation of the curvature perturbation*, *JCAP* **05** (2005) 004, [[astro-ph/0411220](#)].
- [47] A. A. Starobinsky, *Stochastic De Sitter (inflationary) stage in the early universe*, *Lect. Notes Phys.* **246** (1986) 107–126.
- [48] K. Enqvist, S. Nurmi, D. Podolsky and G. Rigopoulos, *On the divergences of inflationary superhorizon perturbations*, *JCAP* **04** (2008) 025, [[0802.0395](#)].
- [49] T. Fujita, M. Kawasaki, Y. Tada and T. Takesako, *A new algorithm for calculating the curvature perturbations in stochastic inflation*, *JCAP* **12** (2013) 036, [[1308.4754](#)].
- [50] V. Vennin and A. A. Starobinsky, *Correlation Functions in Stochastic Inflation*, *Eur. Phys. J. C* **75** (2015) 413, [[1506.04732](#)].
- [51] C. Pattison, V. Vennin, H. Assadullahi and D. Wands, *Quantum diffusion during inflation and primordial black holes*, *JCAP* **10** (2017) 046, [[1707.00537](#)].
- [52] J. M. Ezquiaga, J. García-Bellido and V. Vennin, *The exponential tail of inflationary fluctuations: consequences for primordial black holes*, *JCAP* **03** (2020) 029, [[1912.05399](#)].
- [53] G. Panagopoulos and E. Silverstein, *Primordial Black Holes from non-Gaussian tails*, [1906.02827](#).
- [54] D. G. Figueroa, S. Raatikainen, S. Rasanen and E. Tomberg, *Non-Gaussian Tail of the Curvature Perturbation in Stochastic Ultraslow-Roll Inflation: Implications for Primordial Black Hole Production*, *Phys. Rev. Lett.* **127** (2021) 101302, [[2012.06551](#)].
- [55] C. Pattison, V. Vennin, D. Wands and H. Assadullahi, *Ultra-slow-roll inflation with quantum diffusion*, *JCAP* **04** (2021) 080, [[2101.05741](#)].
- [56] E. Tomberg, *A numerical approach to stochastic inflation and primordial black holes*, *J. Phys. Conf. Ser.* **2156** (2021) 012010, [[2110.10684](#)].
- [57] G. Rigopoulos and A. Wilkins, *Inflation is always semi-classical: diffusion domination overproduces Primordial Black Holes*, *JCAP* **12** (2021) 027, [[2107.05317](#)].
- [58] A. Achúcarro, S. Cespedes, A.-C. Davis and G. A. Palma, *The hand-made tail: non-perturbative tails from multifield inflation*, *JHEP* **05** (2022) 052, [[2112.14712](#)].
- [59] J. M. Ezquiaga, J. García-Bellido and V. Vennin, *Massive Galaxy Clusters Like El Gordo Hint at Primordial Quantum Diffusion*, *Phys. Rev. Lett.* **130** (2023) 121003, [[2207.06317](#)].
- [60] C. Animali and V. Vennin, *Primordial black holes from stochastic tunnelling*, *JCAP* **02** (2023) 043, [[2210.03812](#)].
- [61] Y.-F. Cai, X.-H. Ma, M. Sasaki, D.-G. Wang and Z. Zhou, *Highly non-Gaussian tails and primordial black holes from single-field inflation*, *JCAP* **12** (2022) 034, [[2207.11910](#)].
- [62] A. D. Gow, H. Assadullahi, J. H. P. Jackson, K. Koyama, V. Vennin and D. Wands, *Non-perturbative non-Gaussianity and primordial black holes*, *EPL* **142** (2023) 49001, [[2211.08348](#)].
- [63] V. Briaud and V. Vennin, *Uphill inflation*, *JCAP* **06** (2023) 029, [[2301.09336](#)].
- [64] V. Vennin and D. Wands, *Quantum diffusion and large primordial perturbations from inflation*, [2402.12672](#).
- [65] R. Inui, H. Motohashi, S. Pi, Y. Tada and S. Yokoyama, *Constant roll and non-Gaussian tail in light of logarithmic duality*, [2409.13500](#).
- [66] D. Sharma, *Stochastic inflation and non-perturbative power spectrum beyond slow roll*, [2411.08854](#).
- [67] K. Ando and V. Vennin, *Power spectrum in stochastic inflation*, *JCAP* **04** (2021) 057, [[2012.02031](#)].

- [68] Y. Tada and V. Vennin, *Statistics of coarse-grained cosmological fields in stochastic inflation*, *JCAP* **02** (2022) 021, [[2111.15280](#)].
- [69] C. Animali and V. Vennin, *Clustering of primordial black holes from quantum diffusion during inflation*, *JCAP* **08** (2024) 026, [[2402.08642](#)].
- [70] A. D. Linde, D. A. Linde and A. Mezhlumian, *From the Big Bang theory to the theory of a stationary universe*, *Phys. Rev. D* **49** (1994) 1783–1826, [[gr-qc/9306035](#)].
- [71] M. Jain and M. P. Hertzberg, *Statistics of Inflating Regions in Eternal Inflation*, *Phys. Rev. D* **100** (2019) 023513, [[1904.04262](#)].
- [72] J. Felsenstein, *Inferring Phylogenies*. Sinauer Associates, Sunderland, Massachusetts, 2004.
- [73] R. M. Anderson and R. M. May, *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, 05, 1991, [10.1093/oso/9780198545996.001.0001](#).
- [74] L. Breiman, *Random Forests*, *Machine Learning* **45** (2001) 5–32.
- [75] P. Glasserman, *Monte Carlo Methods in Financial Engineering*. Springer, 2003, [10.1007/978-0-387-21617-1](#).
- [76] I. Hanski, *Metapopulation Ecology*. Oxford University Press, 1999.
- [77] W. E. East, M. Kleban, A. Linde and L. Senatore, *Beginning inflation in an inhomogeneous universe*, *JCAP* **09** (2016) 010, [[1511.05143](#)].
- [78] K. Clough, E. A. Lim, B. S. DiNunno, W. Fischler, R. Flauger and S. Paban, *Robustness of Inflation to Inhomogeneous Initial Conditions*, *JCAP* **09** (2017) 025, [[1608.04408](#)].
- [79] J. K. Bloomfield, P. Fitzpatrick, K. Hilbert and D. I. Kaiser, *Onset of inflation amid backreaction from inhomogeneities*, *Phys. Rev. D* **100** (2019) 063512, [[1906.08651](#)].
- [80] C. Joana and S. Clesse, *Inhomogeneous preinflation across Hubble scales in full general relativity*, *Phys. Rev. D* **103** (2021) 083501, [[2011.12190](#)].
- [81] M. Corman and W. E. East, *Starting inflation from inhomogeneous initial conditions with momentum*, *JCAP* **10** (2023) 046, [[2212.04479](#)].
- [82] Y. L. Launay, G. I. Rigopoulos and E. P. S. Shellard, *Stochastic inflation in general relativity*, *Phys. Rev. D* **109** (2024) 123523, [[2401.08530](#)].
- [83] G. N. Felder and I. Tkachev, *LATTICEEASY: A Program for lattice simulations of scalar fields in an expanding universe*, *Comput. Phys. Commun.* **178** (2008) 929–932, [[hep-ph/0011159](#)].
- [84] A. V. Frolov, *DEFROST: A New Code for Simulating Preheating after Inflation*, *JCAP* **11** (2008) 009, [[0809.4904](#)].
- [85] R. Easther, H. Finkel and N. Roth, *PSpectRe: A Pseudo-Spectral Code for (P)reheating*, *JCAP* **10** (2010) 025, [[1005.1921](#)].
- [86] D. G. Figueroa, A. Florio, F. Torrenti and W. Valkenburg, *CosmoLattice: A modern code for lattice simulations of scalar and gauge field dynamics in an expanding universe*, *Comput. Phys. Commun.* **283** (2023) 108586, [[2102.01031](#)].
- [87] A. Caravano, E. Komatsu, K. D. Lozanov and J. Weller, *Lattice simulations of inflation*, *JCAP* **12** (2021) 010, [[2102.06378](#)].
- [88] Y. Mizuguchi, T. Murata and Y. Tada, *STOLAS: STOchastic LAttice Simulation of cosmic inflation*, [2405.10692](#).
- [89] M. Shibata and M. Sasaki, *Black hole formation in the Friedmann universe: Formulation and computation in numerical relativity*, *Phys. Rev. D* **60** (1999) 084002, [[gr-qc/9905064](#)].

- [90] T. Harada, C.-M. Yoo, T. Nakama and Y. Koga, *Cosmological long-wavelength solutions and primordial black hole formation*, *Phys. Rev. D* **91** (2015) 084057, [[1503.03934](#)].
- [91] I. Musco, *Threshold for primordial black holes: Dependence on the shape of the cosmological perturbations*, *Phys. Rev. D* **100** (2019) 123524, [[1809.02127](#)].
- [92] L. Pinol, S. Renaux-Petel and Y. Tada, *A manifestly covariant theory of multifield stochastic inflation in phase space: solving the discretisation ambiguity in stochastic inflation*, *JCAP* **04** (2021) 048, [[2008.07497](#)].
- [93] H. Firouzjahi, A. Nassiri-Rad and M. Noorbala, *Stochastic Ultra Slow Roll Inflation*, *JCAP* **01** (2019) 040, [[1811.02175](#)].
- [94] C. Pattison, V. Vennin, H. Assadullahi and D. Wands, *Stochastic inflation beyond slow roll*, *JCAP* **07** (2019) 031, [[1905.06300](#)].
- [95] G. Ballesteros, J. Rey, M. Taoso and A. Urbano, *Stochastic inflationary dynamics beyond slow-roll and consequences for primordial black hole formation*, *JCAP* **08** (2020) 043, [[2006.14597](#)].
- [96] S. Raatikainen, S. Räsänen and E. Tomberg, *Primordial Black Hole Compaction Function from Stochastic Fluctuations in Ultraslow-Roll Inflation*, *Phys. Rev. Lett.* **133** (2024) 121403, [[2312.12911](#)].
- [97] J. H. P. Jackson, H. Assadullahi, A. D. Gow, K. Koyama, V. Vennin and D. Wands, *Stochastic inflation beyond slow roll: noise modelling and importance sampling*, [2410.13683](#).
- [98] F. Finelli, G. Marozzi, A. A. Starobinsky, G. P. Vacca and G. Venturi, *Stochastic growth of quantum fluctuations during slow-roll inflation*, *Phys. Rev. D* **82** (2010) 064020, [[1003.1327](#)].
- [99] J. Grain and V. Vennin, *Stochastic inflation in phase space: Is slow roll a stochastic attractor?*, *JCAP* **05** (2017) 045, [[1703.00447](#)].
- [100] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*. Springer Berlino, Heidelberg, 1992, [10.1007/978-3-662-12616-5](#).
- [101] P. Auclair, *Forest: Fortran recursive exploration of stochastic trees*, Apr., 2025. [10.5281/zenodo.15235932](#).
- [102] S. Young, C. T. Byrnes and M. Sasaki, *Calculating the mass fraction of primordial black holes*, *JCAP* **07** (2014) 045, [[1405.7023](#)].
- [103] I. Musco, V. De Luca, G. Franciolini and A. Riotto, *Threshold for primordial black holes. II. A simple analytic prescription*, *Phys. Rev. D* **103** (2021) 063538, [[2011.03014](#)].
- [104] A. Kehagias, D. Perrone and A. Riotto, *Why the Universal Threshold for Primordial Black Hole Formation is Universal*, [2405.05208](#).
- [105] S. Young, *The primordial black hole formation criterion re-examined: Parametrisation, timing and the choice of window function*, *Int. J. Mod. Phys. D* **29** (2019) 2030002, [[1905.01230](#)].
- [106] K. Tokeshi, K. Inomata and J. Yokoyama, *Window function dependence of the novel mass function of primordial black holes*, *JCAP* **12** (2020) 038, [[2005.07153](#)].
- [107] A. Kalaja, N. Bellomo, N. Bartolo, D. Bertacca, S. Matarrese, I. Musco et al., *From Primordial Black Holes Abundance to Primordial Curvature Power Spectrum (and back)*, *JCAP* **10** (2019) 031, [[1908.03596](#)].
- [108] C. Germani and I. Musco, *Abundance of Primordial Black Holes Depends on the Shape of the Inflationary Power Spectrum*, *Phys. Rev. Lett.* **122** (2019) 141302, [[1805.04087](#)].
- [109] A. Escrivà, C. Germani and R. K. Sheth, *Universal threshold for primordial black hole formation*, *Phys. Rev. D* **101** (2020) 044022, [[1907.13311](#)].

- [110] J. C. Niemeyer and K. Jedamzik, *Near-critical gravitational collapse and the initial mass function of primordial black holes*, *Phys. Rev. Lett.* **80** (1998) 5481–5484, [[astro-ph/9709072](#)].
- [111] K. Jedamzik, *The Cloud in cloud problem in the Press-Schechter formalism of hierarchical structure formation*, *Astrophys. J.* **448** (1995) 1–17, [[astro-ph/9408080](#)].
- [112] S. Winitzki, *Reheating-volume measure for random-walk inflation*, *Phys. Rev. D* **78** (2008) 063517, [[0805.3940](#)].
- [113] S. Dubovsky, L. Senatore and G. Villadoro, *The Volume of the Universe after Inflation and de Sitter Entropy*, *JHEP* **04** (2009) 118, [[0812.2246](#)].
- [114] V. Vennin, *Stochastic inflation and primordial black holes*, *Habilitation thesis* (2020) , [[2009.08715](#)].
- [115] W. H. Press and P. Schechter, *Formation of galaxies and clusters of galaxies by selfsimilar gravitational condensation*, *Astrophys. J.* **187** (1974) 425–438.
- [116] A. A. Starobinsky and J. Yokoyama, *Equilibrium state of a selfinteracting scalar field in the de Sitter background*, *Phys. Rev. D* **50** (1994) 6357–6368, [[astro-ph/9407016](#)].
- [117] N. A. Chernikov and E. A. Tagirov, *Quantum theory of scalar fields in de Sitter space-time*, *Ann. Inst. H. Poincaré A Phys. Theor.* **9** (1968) 109.
- [118] E. A. Tagirov, *Consequences of field quantization in de Sitter type cosmological models*, *Annals Phys.* **76** (1973) 561–579.
- [119] T. S. Bunch and P. C. W. Davies, *Quantum Field Theory in de Sitter Space: Renormalization by Point Splitting*, *Proc. Roy. Soc. Lond. A* **360** (1978) 117–134.
- [120] J. H. P. Jackson, H. Assadullahi, K. Koyama, V. Vennin and D. Wands, *Numerical simulations of stochastic inflation using importance sampling*, *JCAP* **10** (2022) 067, [[2206.11234](#)].
- [121] P. Auclair and B. Blachier, *Small-scale clustering of primordial black holes: Cloud-in-cloud and exclusion effects*, *Phys. Rev. D* **109** (2024) 123538, [[2402.00600](#)].
- [122] J. F. N. n. Siles and J. García-Bellido Capdevila, *Primordial Black Hole clusters, phenomenology & implications*, [2405.06391](#).
- [123] W. R. Coulton, O. H. E. Philcox and F. Villaescusa-Navarro, *The Impact of Non-Gaussian Primordial Tails on Cosmological Observables*, [2406.15546](#).