

Constraining Inflationary Particle Production with CMB Polarization

Luca H. Abu El-Haj*

Department of Physics, Columbia University, NY 10027, USA

Oliver H.E. Philcox

Leinweber Institute for Theoretical Physics at Stanford, 382 Via Pueblo, Stanford, CA 94305, USA

Kavli Institute for Particle Astrophysics and Cosmology, 382 Via Pueblo, Stanford, CA 94305, USA

Simons Society of Fellows, Simons Foundation, New York, NY 10010, USA and

Department of Physics, Columbia University, NY 10027, USA

J. Colin Hill

Department of Physics, Columbia University, New York, NY 10027, USA

Inflation generates correlated temperature (T) and polarization (E -mode) fluctuations in the cosmic microwave background (CMB), which differ only by their transfer functions. Following Philcox *et al.* [1], we investigate a scenario with a massive partner to the inflaton ($\mathcal{O}(100)$ times the inflationary Hubble scale), in which particles are produced during a narrow time period, leaving characteristic hot- or cold-spots in the CMB. Using tools developed for thermal Sunyaev-Zel'dovich cluster-finding, we search component-separated *Planck* PR4 E -mode maps for these hotspots, and compare to analogous results in T . Our analysis pipeline is validated on simulated observations and gives unbiased constraints for sufficiently large and bright hotspots. We find no strong evidence for primordial hotspots and thereby place novel bounds on the marginal and relevant coupling between the inflaton and massive scalars during inflation, probing physics at energies many orders of magnitude above any feasible terrestrial collider. Due to the large noise in *Planck* polarization, the temperature data dominate constraints for small hotspots, but for sufficiently large hotspots, our bounds improve on those of [1] by $\sim 25\%$. We also forecast the inferred bounds on inflationary physics for a search using Atacama Cosmology Telescope (ACT) data, and from an optimistic cosmic-variance-limited experiment (CV), for which E -mode data provide stronger constraints than T on nearly all scales. ACT should improve on the *Planck* constraints by $\gtrsim 25\%$, nearing the CV limit on intermediate scales. Finally, we compare the constraining power of localized searches to that of a power spectrum analysis, and demonstrate that for sufficiently few produced particles the localized search performed herein is dominant.

I. INTRODUCTION

A critical feature of inflation is that it generates fluctuations in the temperature (T) and E -mode polarization of the cosmic microwave background (CMB) in correlated ways, differing only by their respective transfer functions. Inflationary physics can thus be constrained by both observations of the CMB intensity and polarization, which crucially differ in their noise and systematic properties. In this work, we explore this in the context of a specific inflationary scenario involving an extremely massive scalar coupled to the inflaton.

Given the enormous number of microphysical inflationary scenarios, a natural way to constrain the physics of the early universe is to consider fairly generic properties of these models and to explore their phenomenological implications in the context of contemporary and future cosmological surveys. The existence of many fields during inflation, even extremely massive ones, is a common prediction of UV-complete physical models [2, 3]; the inflationary production of such particles thus provides a powerful phenomenological probe of many multi-field

models. In practice, this production can be sourced by interactions between the inflaton and the massive field, which can induce a time-dependent effective mass for the former. If this mass drops rapidly, particle production can be favorable [4–7]. Moreover, if the minimum effective mass is sufficiently large when particles are produced, the particles can measurably modify the local gravitational potential, or equivalently induce a curvature perturbation [8], which in turn evolves to produce localized hot- or cold-spots in the CMB¹. The observational consequences of the production of massive particles have been studied in CMB temperature data [1, 4, 9, 10], but have not been extensively explored in polarization.

If these hotspots exist, how can we best search for them in observational data? In the context of axion monodromy inflation, Münchmeyer and Smith [9] demonstrate that a matched-filter estimator approach for local profile-finding performs better in constraining inflationary particle production than N -point correlation function analyses in the large- N limit, so long as particle

¹Following [1], we refer to both hotspots and coldspots as “hotspots” for concision. As we will discuss in Sec. II, given the complicated profile structure for E -modes there is some ambiguity about whether particle production anisotropy is “hot” or “cold”.

* lha2126@columbia.edu

production is rare. In this work, we follow a similar path, as outlined in [1], applying a matched-filter estimator to the SEVEM and SMICA component-separated E -mode maps built from the *Planck* PR4 data [11–13]. We exploit the fact that, while it is uncommon to employ localized searches in polarization, we can form scalar E -mode maps analogous to conventional T fields, so we may take advantage of data analysis parallels with searches for galaxy clusters in CMB maps via the thermal Sunyaev-Zel’dovich (tSZ) effect (the inverse-Compton scattering of CMB photons off hot electrons in the intracluster medium [14]). In both the tSZ and the primordial hotspot production case, there is a well-defined angular profile of the signal of interest, characterized by a few free parameters. We compute these angular profiles for polarization and then construct a matched filter to search for them.

Polarization provides an attractive probe with which to constrain inflationary physics, for a wide variety of reasons. Most obviously, polarization gives a (partially) statistically independent signal to temperature, from which we can glean new information and cross-check temperature results. Strikingly, hotspot profiles for temperature and polarization are very phenomenologically distinct, due to the differences in their respective transfer functions. For this reason, coupled with the different properties of noise in T and E , we will later find that, for future experiments (and existing ground-based CMB experiments), polarization data will provide better constraints on inflationary hotspots than temperature on nearly all scales. Although this may be model-specific,² it emphasizes the power of a polarization search, and makes clear that, as we enter a generation of CMB experiments where polarization sensitivity is comparable to temperature [16–20], localized polarization searches are powerful. Even beyond future interest, we show from simulations that our *Planck* polarization search performs better on large angular scales than that of temperature and actually exceeds the constraining power of a temperature-only cosmic-variance-limited search on these scales. Additionally, foreground contamination is much less significant in polarization than in temperature, particularly on small scales (noting that tSZ clusters, which can cause false detections in temperature, are unpolarized to leading order). Polarization comes with an additional built-in cross-check: knowing that noise impacts E - and B -mode polarization similarly, and foregrounds can produce both, we could explicitly check for a strong hotspot candidate that it *did not* appear in a B -mode map. While we do not find any sufficiently strong candidates in this work to motivate such a test, this could be very powerful in the event of a future detection.

²However, we expect it is generically true for such particle production signatures, given works such as [15], which demonstrate that CMB polarization constrains cosmological parameters better than temperature in the cosmic-variance limit.

The remainder of this work is organized as follows. In Sec. II, we discuss particle production hotspots and their phenomenology. In Sec. III, we explain our pipeline for the polarization search. We then verify our pipeline on simulations in Sec. IV and apply it to *Planck* SEVEM and SMICA component-separated maps in Sec. V, comparing to the search using temperature data performed in [1]. In Sec. VI, we carefully explore the consequences of polarization versus temperature data in localized hotspot searches in general and robustly forecast the detection capabilities of both probes for current and future CMB experiments. In Sec. VII, we explore the constraining power of a power spectrum analysis, and compare it to our profile-finding method for both temperature and polarization, before concluding in Sec. VIII. Throughout, we adopt a flat Λ CDM cosmology based on [21]: $\{H_0 = 67.32 \text{ km/s/Mpc}, \omega_b = 0.022383, \omega_{\text{cdm}} = 0.12011, \tau_{\text{reio}} = 0.0543, \sum m_\nu = 0.06 \text{ eV}, A_s = 2.1006 \times 10^{-9}, n_s = 0.96605\}$. We use the $(-, +, +, +)$ metric signature convention.

II. THEORETICAL BACKGROUND

A. Particle Production

Let us consider an extremely massive field σ whose effective mass momentarily passes through a minimum at some inflaton field value φ_* , yielding the expansion:

$$M_\sigma(\varphi) = M_0 + \left. \frac{d^2 M_\sigma}{d\varphi^2} \right|_{\varphi_*} \frac{1}{2}(\varphi - \varphi_*)^2 + \mathcal{O}((\varphi - \varphi_*)^3). \quad (1)$$

The validity and interest of this approximation can equivalently be thought of as the inclusion of only relevant and marginally relevant terms in the effective mass term of the Lagrangian for the σ field, which is given by

$$\mathcal{L} \supset -\frac{1}{2}(\partial\sigma)^2 - \frac{1}{2}(M_0^2 + (g\varphi - \mu)^2)\sigma^2, \quad (2)$$

where we have defined coupling constants $g^2 \equiv M_0 \frac{d^2 M_\sigma}{d\varphi^2} \big|_{\varphi_*}$ and $\mu \equiv g\varphi_*$. Employing the slow-roll expansion, we may write $\varphi - \varphi_* \simeq \dot{\varphi}_0(t - t_*)$, with the slow-rolling velocity given by $\dot{\varphi}_0 \simeq (57H_I)^2$ from the large-scale CMB anisotropy normalization, where H_I is the inflationary Hubble scale. Under this approximation, the effective time-dependent squared mass is given by:

$$M_\sigma^2 \simeq M_0^2 + g^2 \dot{\varphi}_0^2 (t - t_*)^2. \quad (3)$$

In conformal time, defined by $\eta \equiv \int \frac{dt}{a}$ where a is the scale factor (during standard quasi-de Sitter inflation $a \sim e^{H_I t}$), we may simplify by noting that $(t - t_*) = -(1/H_I) \ln(\eta/\eta_*)$, where η_* denotes the conformal time




FIG. 1. The number of particle-production hotspots $N_{\text{HS}}(g, M_0)$, where we take characteristic values for η_* and $\Delta\eta$ equal to 100 Mpc to give a sense of the nontrivial dependence of particle production on the physical parameters in the model.

at production³:

$$M_\sigma^2 \simeq M_0^2 + \frac{g^2 \dot{\varphi}_0^2}{H_I^2} \ln^2 \left| \frac{\eta}{\eta_*} \right|. \quad (4)$$

For a simple estimate, if we take $M_0 = \mathcal{O}(100H_I)$ and $g = \mathcal{O}(10)$, consistent with the values explored in this work, it is clear that the second term dominates, and in particular as $\eta \rightarrow \eta_*$ the time-dependent term vanishes so there is a significant decrease in the effective mass, thus allowing σ particle production.

The particle production associated with this Lagrangian can be computed using standard Bogoliubov techniques. A pedagogical review of the calculation can be found in [4]. The approximate number of hotspots produced in a shell of thickness $\Delta\eta$ around the surface of last scattering is given by [4, 10]

$$N_{\text{HS}} \simeq 4 \times 10^8 \times g^{\frac{3}{2}} \left(\frac{\Delta\eta}{100 \text{ Mpc}} \right) \left(\frac{100 \text{ Mpc}}{\eta_*} \right)^3 \times e^{-\pi(M_0^2 - 2H_I^2)/(g|\dot{\varphi}_0|)}. \quad (5)$$

The number of hotspots is clearly dependent on the mass M_0 of the particle. For $M_0 \gg H_I$, one can assume that the produced particles become non-relativistic very quickly, and as such induce localized features on the CMB. In Figure 1, we plot N_{HS} as a function of g and M_0/H_I , for representative values $\eta_* = 100 \text{ Mpc} = \Delta\eta$. As we discuss in Sec. VII, for a large number of hotspots

³While formally this is the relevant definition, more intuitively η_* defines a characteristic size for the hotspot — see Figure 2.

(shown in the yellow region of parameter space in Figure 1, for example) a power spectrum analysis becomes more optimal than a localized search. Our profile-finding search is optimal in the upper left corner of parameter space shown in Figure 1.

B. CMB E -mode Particle-Production Hotspots

The hotspots induced by massive particle production can be thought of intuitively as follows. The production of a very massive particle induces a local gravitational potential, which, in turn, causes a non-zero vacuum expectation value for the curvature perturbation, which propagates to the observable CMB anisotropies. Equivalently, one can consider the Lagrangian in Equation (2) as containing an interaction term between φ and σ . As such, σ exerts a force on the inflaton field, slowing it down and locally causing inflation to end later. This leaves overdensities, which, after CMB decoupling, induce hot or cold spots depending on the CMB transfer function. These hotspot profiles for temperature have been computed and explored extensively in [1, 4, 10]; here we present the analogous computation for E -mode polarization.

Adopting a similar approach to previous works, one can analytically compute the position-space profiles for these hotspots. The curvature perturbation associated with a massive particle is computed in [4] and reproduced for convenience in Appendix A; this yields

$$\langle \zeta_{\text{HS}} \rangle = e^{-i\mathbf{k} \cdot \mathbf{x}_{\text{HS}}} \frac{gH_I^2}{\dot{\varphi}_0} \frac{\text{Si}(k\eta_*) - \sin k\eta_*}{k^3}. \quad (6)$$

Here, $\text{Si}(x) \equiv \int_0^x \frac{\sin t}{t} dt$ and we also define $f(x) \equiv \frac{gH_I^2}{\dot{\varphi}_0} (\text{Si}(x) - \sin x)$. To transform from curvature to CMB observables, we must take into account the E -mode transfer function $\mathcal{T}_\ell^E(k)$, which is dominated by Doppler contributions. The E -mode anisotropy induced by a curvature perturbation from a Fourier mode of wavevector $\mathbf{k}$ is [22, 23]

$$\delta E(\mathbf{k}, \hat{\mathbf{n}}) = \sum_{\ell=2}^{\infty} i^\ell \mathcal{P}_\ell(\hat{\mathbf{n}} \cdot \mathbf{k}) \mathcal{T}_\ell^E(k) \langle \zeta_{\text{HS}}(k) \rangle, \quad (7)$$

where $\mathcal{P}_\ell(x)$ are the Legendre polynomials. Summing over Fourier modes (following [10]) leaves us with

$$\delta E(\hat{\mathbf{n}}, \hat{\mathbf{n}}_{\text{HS}}, \eta_*, \eta_{\text{HS}}) = \frac{1}{2\pi^2} \sum_{\ell=2}^{\infty} (2\ell+1) \mathcal{P}_\ell(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}_{\text{HS}}) \times \int \frac{dk}{k} f(k\eta_*) j_\ell(k\chi_{\text{HS}}) \mathcal{T}_\ell^E(k), \quad (8)$$

where $j_\ell(x)$ are the spherical Bessel functions. Here we have approximated the position of the hotspot as $\mathbf{x}_{\text{HS}} - \mathbf{x}_0 \simeq (\chi_{\text{HS}}, \hat{\mathbf{n}}_{\text{HS}})$, where $\chi_{\text{HS}} = \eta_0 - \eta_{\text{HS}}$ and η_{HS} and η_0 are the sizes of the comoving horizon when the hotspot is produced and today, respectively. We assume

that η_{HS} is close to the comoving horizon size at last scattering η_{rec} , such that the massive particles leave detectable impacts on the observed CMB. We compute the transfer function using CAMB⁴ and cut off the sum above $\ell_{\text{max}} = 3500$, which is sufficient to extract all information in contemporary and near-future CMB experiments. Note critically the linearity of the profile in g , which is the basis of our analysis pipeline.

We provide exemplar E -mode hotspot profiles in Figure 2. The E -mode profiles have a number of salient features. Beyond just a central amplitude, the profiles have a characteristic trough around 36 arcmin (0.01 radians). This trough is η_{HS} -dependent; for these calculations, we fix $\eta_{\text{HS}} = \eta_{\text{rec}}$, though other values will be discussed below. There is also a characteristic, though significantly less dramatic, peak at a small angular distance past the trough. Notice also that these are phenomenologically quite different from the temperature hotspot profiles (also shown in Figure 2), where the characteristic feature is the central temperature and the profile cuts off around the characteristic scale $\theta_* = \sqrt{4\pi\eta_*/\eta_0}$.

For a better sense of the phenomenology, we also present in Figure 3 the central amplitude of the E -mode hotspot profile, which is computed by taking $\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}_{\text{HS}} \rightarrow 1$, which means that $\mathcal{P}_\ell \rightarrow 1$. Thus,

$$\delta E_{\text{cent}} = \frac{1}{2\pi^2} \sum_{\ell=2}^{\infty} (2\ell+1) \int \frac{dk}{k} f(k\eta_*) j_\ell(k\chi_{\text{HS}}) \mathcal{T}_\ell^E(k). \quad (9)$$

We note two phenomenological features in Figure 3. First, we find a strongly peaked behavior for the polarization central amplitude close to where the visibility function peaks, *i.e.*, $\eta_* \approx 150$ Mpc. By contrast, the central amplitude in T is fairly constant for $\eta_* \gtrsim 150$ Mpc but strongly decreasing for smaller η_* .

Also, we find the temperature anisotropy to be characteristically $\sim 10\times$ larger than the polarization, similar to that of the ratio between polarization and temperature from a standard single-field inflationary scenario. Considering the variation with respect to η_{HS} , one should also observe that the central amplitude is sharply peaked around $\eta_{\text{HS}} = \eta_{\text{rec}}$. This arises from the E -mode polarization transfer function, which is sharply peaked around recombination.

III. ANALYSIS METHODOLOGY

To search for inflationary hotspots in E -mode data, we use a methodology similar to [1]. By conservation of momentum in a homogeneous inflationary background, our particle production mechanism should produce pairs of particles separated by a distance $\lesssim \eta_*$ [4]. An optimal search would thus involve a pairwise profile, which

is computationally intensive since (a) the pairwise profile breaks isotropy, and (b) we do not know the separation precisely, given its heavy dependence on the microphysics of the problem. We simplify this by performing a search for single hotspots, which is shown to be near-optimal by [1, 10].

To constrain the coupling parameter g , we use a matched-filter analysis, which is implemented following the methodology of tSZ cluster searches [24]. We make minor modifications to the SZIFI code described in [1, 25, 26], enhancing the primordial non-Gaussianity functionality added in [1], allowing computation of a massive particle hotspot template in both T and E . Additionally, we add functionality for the E -mode *Planck* SEVEM and SMICA component-separated maps, using the *Planck* SMICA polarization beam.

Our analysis is performed on SEVEM [11] and SMICA [12, 13] *Planck* PR4 component-separated E -mode polarization maps produced using a spin-two harmonic transform from Q/U to real space E/B -modes in HEALPY [27] with $N_{\text{side}} = 2048$. This transformation allows for a scalar treatment of real space E -mode maps. We then apply the *Planck* component-separated common mask to the E -mode real-space maps, which removes the Galactic plane, and we inpaint a polarized point source mask using a diffusive algorithm [28]. Once the maps have been produced and masked, we split the full sky into 768 tiles, each $14.8^\circ \times 14.8^\circ$. Each tile contains 1024^2 pixels following the procedure of [25, 26]. We then use a matched filter estimator for the coupling constant:


$$\hat{g}(\boldsymbol{\theta}) = \sigma_g^2 \int \frac{d^2\mathbf{l}}{(2\pi)^2} \frac{t^*(\mathbf{l}, \mathbf{0})d(\mathbf{l})}{C^{EE}(\mathbf{l})}, \quad (10)$$

$$\sigma_g = \left(\int \frac{d^2\mathbf{l}}{(2\pi)^2} \frac{|t(\mathbf{l}, \mathbf{0})|^2}{C^{EE}(\mathbf{l})} \right)^{-\frac{1}{2}}. \quad (11)$$


Here, $t(\mathbf{l}, \boldsymbol{\theta})$ is the Fourier-space hotspot profile centered at $\boldsymbol{\theta}$, and $d(\boldsymbol{\theta})$ is the masked and inpainted component-separated E -mode map. By isotropy, we can fix the hotspot template at the origin. The power spectrum $C^{EE}(\mathbf{l})$, which includes both signal and noise, is directly computed from the data on each tile via $C^{EE}(\mathbf{l}) = \langle d(\mathbf{l})d^*(\mathbf{l}) \rangle$. The SNR of a hotspot candidate is defined by $\text{SNR} \equiv \hat{g}/\sigma_g$. Although we include transfer functions up to $\ell_{\text{max}} = 3500$ in the computation of our profiles, when computing $\hat{g}$ and σ_g we only integrate in the range $\ell = [30, 3000]$.

We generate 100 position-space templates, using 10 logarithmically spaced values of $\eta_* \in [10, 1000]$ Mpc and similarly 10 linearly spaced values of $\chi_{\text{HS}} \in [\chi_{\text{rec}} - \eta_*, \chi_{\text{rec}} + \eta_*]$. We enforce a causality condition on the comoving distance by demanding $0 \leq \chi_{\text{HS}} \leq \eta_0$, where $\chi(\eta) = \eta_0 - \eta$. Our lower bound is justified by noting that the *Planck* beam makes hotspots smaller than $\eta_* \approx 10$ Mpc undetectable. Additionally, we choose an upper bound of 1 Gpc based on the angular size of our tiles. Searching for larger hotspots would require going beyond the flat-space Fourier transform approximation.


⁴<https://camb.info>



Variation with respect to η_* , with $\eta_{\text{HS}} = \eta_{\text{rec}}$.



Variation with respect to η_{HS} , with $\eta_* = 100$ Mpc.



Comparison of T and E profiles.

FIG. 2. Temperature and polarization hotspot templates for $g = 1$, demonstrating the variability of the polarization hotspot profiles with respect to both η_* (encoding the hotspot formation time) and η_{HS} (parametrizing the hotspot distance). **(Top Left)**: We fix $\eta_{\text{HS}} = \eta_{\text{rec}}$ and show profiles for $\eta_* = 30, 100, 300$ Mpc in blue, orange, and green, respectively. **(Top Right)**: We fix $\eta_* = 100$ Mpc and show the variability of the profile with respect to η_{HS} . Each hotspot has a significant “cold ring” at $\theta \simeq 0.5\text{--}0.6^\circ$ from the center. **(Bottom)**: We compare the profiles for temperature (dashed) and polarization (solid). We multiply the polarization profiles by a factor of 10 for visual clarity. Notice that the temperature amplitude is characteristically around an order of magnitude larger than that in polarization.

We compute position-space profiles using CAMB [29] transfer functions, across a logarithmically-spaced grid with $k \in [10^{-6}, 1] h^{-1}\text{Mpc}$, using a Simpson integrator. We compute each profile on a 1024^2 pixel grid, using a flat-space distance $\theta = \cos^{-1}(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}_{\text{HS}})$. We then convolve each profile with the *Planck* SMICA beam. We compute each profile to a maximum $\theta_{\text{max}} = \max[0.1 \text{ rad}, \sqrt{4\pi}(\eta_*/\eta_0)]$. This avoids the edges of tiles, and also generally removes regions where the profile is less than 1% of its peak [1]. While T and E have phenomenologically different profiles, the angular scale of the hotspots should be comparable (see Figure 2).

For the E -mode search, we restrict our analysis to

component-separated data, rather than considering an analysis on the individual frequency maps using a multi-frequency matched filter (MMF). One reason for this is that to leading order there is no tSZ contamination in polarization. Also, in general, there is significantly less small-scale foreground contamination in polarization than in temperature. Note that our use of the common mask excludes point sources in polarization with $\text{SNR} \geq 5$.

We apply a more conservative Galactic mask than our original *Planck* common polarization mask to the output candidates. We adopt a mask that is frequently used for tSZ cluster searches and removes the brightest $\simeq 40\%$ of




FIG. 3. Central amplitude of the hotspots. **(Left)**: We show a comparison between the central amplitudes of hotspot profiles in temperature (blue) and in E -mode polarization (orange), where we have fixed $\eta_{\text{HS}} = \eta_{\text{rec}}$. Negative values are plotted with dashed curves. **(Right)**: We show the central amplitude for polarization computed for various η_{HS} values, where we have chosen 100 Mpc as a characteristic scale for the variation of η_{HS} away from η_{rec} , while preserving causality. Note that the central amplitude of the anisotropy is strongest for hotspots located close to the surface of last scattering ($\eta_{\text{HS}} = \eta_{\text{rec}}$).

the sky, as detailed in [1, 24]. This procedure leaves us with an output catalog of hotspot candidates, which we then further process by applying the *Planck* HFI temperature point source mask [11], which excludes candidates within 10 arcmin of 5σ point sources in temperature. We do this to avoid potential contamination associated with the slight polarization of a point source that is bright in temperature. In practice, as seen below in Sec. V, we find no additional exclusions from this procedure.

IV. SIMULATION RESULTS AND PIPELINE VALIDATION

Before jumping into the analysis on real *Planck* maps, we verify our methodology on a single NPIPE simulation representing the full component-separated *Planck* PR4 dataset. Unlike [1], we only consider simulations for our analysis using injected maps of single hotspots, *i.e.*, we do not inject pairwise hotspot signals. This is justified given that [1] found consistent recovery results between single and pairwise searches.

We follow a similar simulation procedure to [1], where we generate 300 E -mode hotspots with $\eta_{\text{HS}} = \eta_{\text{rec}}$ and consider ten logarithmically spaced values of $\eta_* \in [10, 1000]$ Mpc, convolving all profiles with the *Planck* SMICA beam. We also demand all the injected hotspots to be more than 3° apart, so as to avoid interference effects between two nearby hotspots. We then add this hotspot injection map to simulated NPIPE SEVEM E -mode polarization maps for $g \in [10, 20, 30, 40, 50]$, and run our analysis pipeline. This requires around 20 CPU-hours for each value of g .

In Figure 4, we show the primary results of our simulation analysis. For sufficiently high g and η_* , the com-

pleteness is very high, *i.e.*, we detect the hotspots we inject. Note that because the mask leaves $\simeq 60\%$ of the sky unmasked, we cannot obtain perfect completeness relative to the unmasked catalog.

In general, we find that the efficacy of our polarization search is consistent with that of the previous temperature search [1]. We find that for $g > 20, \eta_* > 25$ Mpc, the completeness ratio is $62.9 \pm 2.5\%$, while [1] found a completeness of $64 \pm 7\%$. As for the temperature search, the completeness is poor for $g = 10$ and does not exceed $\sim 50\%$, even for large η_* , implying that these weak couplings are hard to practically detect. As η_* decreases (and thus the hotspots get smaller), consistent with [1], the completeness reduces significantly. We also present our pipeline's recovery of g and η_* , finding that η_* is well recovered for $\eta_* \gtrsim 50$ Mpc and g is successfully recovered for $\eta_* \gtrsim 100$ Mpc. This is consistent with the results of [1].

Our simulation test results can be simply understood: for large g the hotspots are brighter and easier to detect, while for large η_* , they are bigger and thus have no trouble being resolved by the *Planck* beam. Our conclusion of this simulation test is that our *Planck* pipeline can robustly detect E -mode hotspots with $g \geq 20$ and $\eta_* \geq 70$ Mpc. We also note that our E -mode search presents unbiased estimates on g, η_* for $g \geq 20$ and $\eta_* \geq 70$ Mpc.

Notably, the g values discussed above would naively break the perturbativity bounds of the calculation done in [4], which should require $g \ll \sqrt{4\pi}$. Generically, one would expect large radiative corrections when this bound is violated. In situations with large amounts of symmetry (e.g., supersymmetry [4]) one can attain effective values of relatively high g . Such scenarios can also arise in extra-dimensional models, scenarios with tachyonic Higgs




FIG. 4. Summary of the analysis procedure detailed in Sec. III applied to the simulations detailed in Sec. IV. **(Left)**: estimated values of η_* from the injected hotspots. For sufficiently high g and η_* , the recovered values reproduce the injected values more accurately among hotspots than found in the temperature search of [1]. **(Center)**: fraction of injected sources recovered at $\text{SNR} \geq 5$ as a function of the coupling parameter g and hotspot size η_* . Note that due to masking (primarily of the Galaxy), our result asymptotes to $\sim 60\%$; relative to the masked input catalog, the completeness is effectively 100% for sufficiently large values of g and η_* . Larger hotspots are easier to detect: we are limited on small scales by the *Planck* beam and on large scales by the size of the tiles used in our search and by the Galactic mask. **(Right)**: recovered values of g for the injected sources. Once again, for $g \geq 20$, we recover g accurately for sufficiently large hotspot sizes.

production, or quantum diffusion effects [30–32].

V. APPLICATION TO *PLANCK*

Next, we apply our pipeline to *Planck* component-separated *E*-mode polarization maps using the common and point-source inpainting masks described in Sec. III. Our search uses 100 position-space templates with varying η_* , η_{HS} . For robustness, we apply the procedure to both SEVEM and SMICA component-separated maps [11–13]. Generating a candidate hotspot catalog using our matched-filter method takes ~ 300 CPU-hours.

To avoid double-counting, we merge any two hotspots candidate within η_* of one another. While this may seem dangerous, we adopt this approach for two reasons. First, this follows the convention used in tSZ cluster searches [25, 26]. Second, bright candidates near the edges of our flat-space tiling are very likely to be detected in multiple tiles, and would otherwise be double-counted. Additionally, because of the four free parameters of our template profile (η_{HS} , η_* , and the angular coordinates of the hotspot center), when one finds a true source in a simulation, many candidates are detected locally in a small region around the injected hotspot. As such, our method avoids false over-detections. We further emphasize that this is an important reason why visual inspection plays a role in our final analysis. Because physically the hotspots are produced in pairs, upon the detection of a real hotspot, we should find a similar signal nearby (within distance η_*), which would likely have been merged in our analysis, but would remain visible to visual inspection.

In Figure 5, we show the sky distribution of hotspot candidates with $\text{SNR} \geq 5$. Notably, there is no obvious clustering of the spots in position. Also, due to the four free parameters, the SNR is actually not a true detection

significance (*i.e.*, $\text{SNR} = 5$ does not correspond to a 5σ detection). As such, for a confident detection a more conservative threshold would be $\text{SNR} \geq 6$, which was used in [1].

In Figure 6 we show the distribution of SNRs for hotspot candidates. Applying the procedure to the SEVEM (SMICA) maps we find 20 (23) candidates after the application of our more conservative Galactic mask, with a maximum SNR of 5.4 (5.5). The distributions of SNR values are consistent between the two maps.

We find fewer candidates than found in the temperature analysis of [1]: 20 (23) *E*-mode versus 48 (35) *T* candidates for SEVEM (SMICA). Notably, there are no candidates above an SNR of 6. We expect that many of the $\text{SNR} \geq 5$ candidates are random fluctuations (we present a visual inspection below in Appendix B, which appears consistent with noise). We also present in Table I the recovered parameters for each of the SEVEM *E*-mode candidates; as shown in Appendix C, we find similar inferred parameters when analyzing the SMICA maps. Most of the hotspot candidates appear near the edge of masked regions (see Appendix B), so we unsurprisingly find similar candidate locations in the SEVEM and SMICA maps, as well as similar recovered candidates. Twelve of these hotspots appear in regions of parameter space excluded by the temperature search, and as such we would expect them to have been seen in [1]. As such, they are unlikely to be physical, though they are still visually inspected in Appendix B as an additional cross-check. The remaining eight would likely not have been detected in the temperature search due to large scatter at the relevant η_* seen in Sec. IV. In summary, we find no hotspots with $\text{SNR} > 6$, and none with $\text{SNR} > 5$ with a distinctive profile that is obviously consistent with a primordial hotspot.

The above non-detections can be recast as bounds on inflationary physics, as shown in the exclusion plot of Figure 7. We attain the bounds by computing σ_g from Equa-




FIG. 5. Visualization of the hotspot detection candidates from the *Planck* SEVEM *E*-mode polarization maps. The dark and light shaded areas show the regions removed by the analysis mask and the post-processing mask, respectively. Each of the 20 circles represents a single detection candidate with $\text{SNR} \geq 5$, whose size is proportional to the SNR and whose color indicates the inferred value of η_* . A histogram of the detection significances is shown in Figure 6, with the parameters of all 20 of the $\text{SNR} \geq 5$ detections given in Table I. We find no strong candidates.




FIG. 6. PDF of the SNR values of our *E*-mode hotspot candidates. The SNR is not a true significance due to the four free parameters in the template (η_* , η_{HS} , and the hotspot coordinate center). We also emphasize the consistency between the SEVEM (orange) and SMICA (blue) results, which show similar hotspot numbers and SNR distributions.

tion (11) using our simulations. We present 5σ bounds based on the scatter of recovered g values for hotspots at a given η_* .

Our constraints in the (g, η_*) plane using polarization are strong for large η_* , though due to the size of the *Planck* polarization beam, they become significantly weaker for small η_* . We also compare the bounds from

our polarization search with those attained using temperature in [1]. We find the *E*-mode search improves on the bounds from the *T* search for $\eta_* \geq 300$ Mpc by $\sim 10 - 25\%$, but for $\eta_* \leq 50$ Mpc *T* dominates the constraints. Our bounds from the *E*-mode search are comparable with those from *T* on intermediate η_* scales ($50 \text{ Mpc} \leq \eta_* \leq 300 \text{ Mpc}$). We can then relate these bounds in the (g, η_*) plane to bounds on M_0/H_I via Equation (6). These bounds are consistent with those from [1] in that they only strongly cover the region where $M_0/H_I \gtrsim 500$, *i.e.*, we are only sensitive to *very* massive particles.

VI. FORECASTS FOR LOCALIZED HOTSPOT SEARCHES IN TEMPERATURE AND POLARIZATION

Figure 7 makes clear that our polarization hotspot search provides independent information from that obtained using temperature data in [1]. As noted above, we see improvement in the parameter constraints, which is most pronounced for large η_* . In particular, we find stronger bounds on particle production from *Planck* polarization data for $\eta_* \gtrsim 100$ Mpc.

It is interesting to explore the constraining power of *E*-mode polarization in these localized searches for existing and future CMB experiments. Here, we detail our procedure for analytically computing a forecast exclusion curve for both *T* and *E*-mode hotspot profiles. Our approach follows from Equation (11), though we work




FIG. 7. **(Left)**: Exclusion plot inferred from the analysis of the *Planck* SEVEM component-separated *E*-mode polarization data. The data points show 5σ error bars attained from the simulation in Sec. IV. The shaded region above the data points is the parameter space excluded by our analysis. The colored regions indicate reasonable regions of parameter space corresponding to different values of the particle mass at particle production M_0 . The solid line lower bound is set by requiring that at least one hotspot be produced (see Equation (6)), and the upper limit is set by backreaction constraints from [4]. We also plot on the top axis the characteristic scale ℓ_* , estimated by $\ell_* \simeq \eta_0/\eta_*$. Note that this is an approximate scale for the profile in harmonic space, *not* the scale at which the signal-to-noise is maximized. **(Right)**: We show forecast results for both ACT and an optimal cosmic-variance-limited experiment, where *E* (*T*)-modes are represented by solid (dashed) lines. These forecasts are computed analytically using Equation (11). Notice that for *T*, ACT comes very close to the CV limit for $\eta_* \geq 30$ Mpc, and for *E* ACT improves on the *Planck* *E*-mode constraints by between a factor of 1.2 – 2.75.

in the full-sky limit for computational efficiency (with $t(1) \rightarrow t_{\ell m}$). This can be done due to the isotropy of our profile, which allows us to reduce our problem to calculating $t_{\ell 0} \equiv t_\ell$. From our definition of the position-space profiles in Equation (8), with $X \in \{T, E\}$, we can derive the harmonic space templates using the $m = 0$ spherical harmonics $Y_{\ell 0} = \sqrt{(2\ell + 1)/4\pi} P_\ell$. Defining

$$\mathcal{I}_\ell^X \equiv \int_0^\infty \frac{dk}{k} f(k\eta_*) \mathcal{T}_\ell^X(k) j_\ell(k\chi_{\text{HS}}), \quad (12)$$

we have:

$$\begin{aligned} t_\ell^X &= \int_{S^2} d^2\hat{\mathbf{n}} Y_{\ell 0}^* \delta X(\theta, \eta_*, \eta_{\text{HS}}) \\ &= \sum_{\ell'=2}^\infty \sqrt{\frac{(2\ell+1)^3}{4\pi^3}} \mathcal{I}_{\ell'}^X \int_{-1}^1 d\cos\theta P_\ell(\cos\theta) P_{\ell'}(\cos\theta) \\ &= \sqrt{\frac{2\ell+1}{\pi^3}} \int_0^\infty \frac{dk}{k} f(k\eta_*) \mathcal{T}_\ell^X(k) j_\ell(k\chi_{\text{HS}}). \end{aligned} \quad (13)$$

In the third equality we have used the orthogonality of Legendre polynomials, $\int_{-1}^1 dx P_\ell(x) P_{\ell'}(x) = 2/(2\ell + 1) \delta_{\ell\ell'}$. We can then compute the variance of the matched-filter estimator:

$$\sigma_g^X = \left(\sum_\ell \frac{[t_\ell^X(\eta_*, \chi_{\text{HS}})]^2}{C_{\ell, \text{Exp}}^{XX}} \right)^{-\frac{1}{2}}, \quad (14)$$

where $C_{\ell, \text{Exp}}^{XX} = C_{\ell, \Lambda\text{CDM}}^{XX} + \mathcal{N}_{\ell, \text{Exp}}^{XX}$ denotes the beam-deconvolved power spectrum for a given experiment. As

in our *Planck* data analysis, we define the exclusion curve at $\text{SNR} = 5$ assuming a null hotspot signal, *i.e.*, we consider the bound $g \geq 5\sigma_g$.

It is computationally straightforward to extend this approach to a joint analysis of *T* and *E*, by analogy to the multi-frequency matched-filter formalism [25, 26, 33, 34]. In the joint analysis, we must account for the non-zero covariance of the *T* and *E* fields arising from C_ℓ^{TE} . The variance of the resulting joint estimator is given by

$$(\sigma_g^{T \times E})^{-2} = \sum_{X, X' \in \{T, E\}} \sum_\ell t_\ell^X C_{\ell, \text{Exp}}^{-1, XX'} t_\ell^{X'}. \quad (15)$$

Here, $C_{\ell, \text{Exp}}^{-1, XX'}$ is the matrix inverse of the full (T, E) covariance matrix, *i.e.*, the matrix containing diagonal blocks of C_ℓ^{TT} and C_ℓ^{EE} and an off-diagonal block of C_ℓ^{TE} . In practice, while this joint $T \times E$ analysis would put minor improvements on the *Planck* bounds (10 – 20%), they are not nearly as striking as the improvement associated with using higher-resolution CMB data (for which one could well do a joint analysis).

We now apply this procedure to an optimal cosmic-variance-limited experiment (denoted CV, with $\mathcal{N}_\ell = 0$), with results shown in Figure 7. We find dramatic differences between the constraints from temperature and polarization, where for $\eta_* \leq 150$ Mpc or $\eta_* \geq 500$ Mpc the constraints from *E* are $\sim 1.6\times$ better than those from *T*, and on intermediate scales the improvement factor varies between 0.8 (in the small region of parameter space where *T* is superior) to 1.4. Strikingly, polarization

SNR	Longitude [°]	Latitude [°]	$\hat{g}$	η_* [Mpc]	η_{HS} [Mpc]
5.0	150.7	77.5	20*	77.4	217.1
5.0	72.5	72.0	17	77.4	245.2
5.1	123.5	72.5	16	77.4	259.3
5.3	214.6	64.6	29*	46.4	259.3
5.1	265.4	71.2	12	129.2	292.2
5.4	78.5	62.4	16	215.4	378.4
5.1	224.6	43.1	83*	27.8	293.1
5.3	338.9	48.5	1556*	10.0	272.2
5.3	299.3	38.8	14	215.4	339.2
5.1	33.7	40.5	1246*	10.0	283.2
5.0	82.8	30.5	25	46.4	242.4
5.0	210.2	26.6	31*	46.4	259.3
5.1	37.4	-24.4	77*	27.8	272.8
5.2	275.1	-24.4	801*	10.0	279.5
5.1	242.5	-29.3	32*	46.4	318.4
5.0	116.2	-33.7	1160*	10.0	272.2
5.4	279.2	-38.9	187	16.7	266.8
5.1	46.9	-39.8	21	77.4	329.7
5.2	292.7	-57.1	885*	10.0	272.2
5.1	181.8	-67.5	93*	27.8	293.1

TABLE I. Inferred parameters for hotspot candidates with $\text{SNR} \geq 5$ from the *Planck* SEVEM component-separated E -mode maps. Any of the candidates marked with * are within the parameter space covered by the temperature search [1], but were not detected in the former work. For the remaining eight candidates, all of them are either at low enough η_* or g that our matched-filter pipeline cannot reliably recover hotspots because the *Planck* beam is too large, as shown in Figure 4. By contrast, in the SEVEM temperature component-separated maps, [1] found 48 candidates with $\text{SNR} \geq 5$, and SNR values up to 8. This may be due to stronger foregrounds in temperature than in polarization, in particular given that the high-SNR candidates were matched to point sources, and masking effects.

provides almost uniformly stronger bounds than temperature across the full parameter space. This is emphasized by Figure 7, where we see that for $\eta_* \geq 300$ Mpc, *Planck* E -mode constraints are stronger than those inferred from a CV-limited temperature experiment.

We should also note that the CV limit is very sensitive to the value of ℓ_{max} used in the sum in Equations (14) and (15). In Figure 8 (black curves), we present results for $\ell_{\text{max}} = 3500$, which provides a reasonable comparison with current and proposed future experiments [20, 35–38]. We also demonstrate in Figure 8 that in the optimal case a joint $T \times E$ search only provides marginal improvement over a search using polarization alone ($\sim 5 - 25\%$, where the maximum is reached in the small regime where temperature constraints overtake those from polarization). This occurs since the temperature and polarization are only weakly correlated, and the constraints due to polarization are almost uniformly better than those from temperature. The fact that polarization dominates the constraints is quite interesting and is a striking consequence of the differing noise (including noise from the standard-model CMB fluctuations) and transfer functions in T and E .




FIG. 8. Exclusion forecasts on g (at 5σ) computed with Fisher matrix techniques for a CV-limited experiment. The blue and purple lines show forecasts for a power spectrum analysis, assuming $N_{\text{HS}} = 1$ and 100, respectively. The black curve shows the exclusion region for our matched-filter analysis. Solid lines represent constraints from an E -mode analysis, dots represent T , and dashes represent a joint $T \times E$ analysis. For the power spectra, dot-dashes represent constraints from the TE power spectrum on its own. Note that the power-spectrum-based bounds only become comparable to those from the matched filter for $N_{\text{HS}} \gtrsim 100$. Here we are being as optimistic as possible for the power spectrum and have scaled by $f_{\text{sky}} = 0.6$, which is the same fraction of the sky considered in our hotspot search.

Next, we directly forecast the constraints on inflationary hotspots that could be derived using existing data from the Atacama Cosmology Telescope (ACT) [16, 17, 35, 36, 39]. For this forecast, we use the beam-deconvolved noise power spectra for both temperature and E -mode polarization from ACT Data Release 6 [36]. Our forecast suggests that the bounds of this work can be significantly enhanced with the ACT data on small scales (see Figure 7). For temperature, ACT should nearly reach the CV limit (at $\eta_* \gtrsim 30$ Mpc). We also find that polarization provides significantly stronger bounds on intermediate scales ($25 \text{ Mpc} \lesssim \eta_* \lesssim 300 \text{ Mpc}$) than temperature, by a factor of $\sim 1.25 - 1.5$. For large η_* , we find that *Planck* constraints are comparable to those forecast for ACT; this is not surprising, since $\eta_* = 1000$ Mpc corresponds to a characteristic $\ell_* = \eta_0/\eta_* \simeq 14$, a scale sufficiently large that atmospheric noise renders the ACT data much less powerful than *Planck* [16, 36]. In general, our forecasts (Figure 7) indicate that ACT will provide improved bounds on extremely massive particle production during inflation. It should also be noted that very small hotspots are generically difficult to constrain because of finite resolution effects from the beam.

We should also emphasize that, for this style of matched-filter search, we do not need to correct for ACT's smaller sky coverage than *Planck* (roughly 35% versus 60%, respectively, after masking), because the detection significance is independent of the number of hotspots. To reiterate, our matched-filter pipeline should

detect any individual hotspot for appropriate values of g, η_* , independent of the sky coverage. This differs from a power-spectrum-based approach, as discussed below.

VII. POWER SPECTRUM SEARCHES

Inflationary massive particles also induce changes to the CMB two-point function, which provide an alternative observational probe. Here we provide a derivation⁵ of the induced correction to the power spectrum from massive particle production (see also Appendix C of Ref. [32]). We then compare the bounds to those inferred from a power spectrum-based search to those from a matched-filter search. Throughout, we consider an optimistic power spectrum search, considering a CV-limited experiment without experimental noise up to $\ell_{\max} = 3500$, without marginalizing over Λ CDM parameters. Even in such an optimistic case, we will find that the bounds from a profile-finding search are dominant until large N_{HS} values are reached.

Using the harmonic coefficients from Equation (13), the contribution to the EE autospectrum takes a simple form:

$$C_{\ell, \text{HS}}^{EE} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |t_{\ell m}^E|^2 = \frac{|\mathcal{I}_{\ell}^E|^2}{\pi^3}. \quad (16)$$

Equivalently, we can take $\mathcal{D}_{\ell, \text{HS}}^{EE} \equiv \ell(\ell + 1)/(2\pi) C_{\ell, \text{HS}}^{EE} = \ell(\ell + 1)/(2\pi^4) |\mathcal{I}_{\ell}^E|^2$. An example of the hotspot contribution to the TT, EE , and TE power spectra is shown in Figure 9, setting $\eta_* = 150$ Mpc, $\eta_{\text{HS}} = \eta_{\text{rec}}$, $g = 10$, and $N_{\text{HS}} = \{1, 10, 100\}$. Note that the hotspot contribution to the power spectrum is quadratic in $f(k\eta_*)$, so it scales as $g^2 N_{\text{HS}}$, where the N_{HS} factor enters because each hotspot is assumed to be uncorrelated, so we may add their contributions to the power spectrum linearly. By contrast, the hotspot profiles used in the previous sections scale with g , without dependence on N_{HS} . This presents an interesting difference in the formalism. In the profile-finding case, we look for individual bright spots and expect to find each one. By contrast, for a power spectrum search we are looking for population-averaged effects, and as such it is advantageous to have many hotspots.

We find that the fractional contribution of the hotspots to the polarization power spectrum is comparable to that found for temperature. Additionally, as noted above, the contribution to the power spectrum goes as $g^2 N_{\text{HS}}(g, M_0, \eta_*)$. For example, to get to 1% of the fiducial power spectrum, one needs $g^2 N_{\text{HS}} \simeq 10^4$.

⁵We make the simplifying assumption that single hotspots are produced in an uncorrelated manner. Of course, from momentum conservation, hotspots are produced pairwise, and as such are correlated on distance scales $\lesssim \eta_*$. This effect is however subleading and could be explored in further work.

To proceed more rigorously, we can use Fisher forecast techniques to compute the 1σ sensitivity to $g^2 N_{\text{HS}}$ from the power spectrum. We can then find where the constraints become competitive with our localized search techniques. For our simple estimate, given that we have not specified M_0 , we will treat N_{HS} as an independent variable, noting that the signal goes linearly with $g^2 N_{\text{HS}}$. We consider the power spectrum of a CV-limited experiment (for which there is no noise contribution to the power spectrum), and treat g^2 as the sole parameter in our forecast, simply scaling the overall signal by N_{HS} . The Fisher matrix is thus one-dimensional:

$$F_{g^2} = \sum_{\ell} (2\ell + 1) \frac{f_{\text{sky}} N_{\text{HS}}^2 \left(\frac{dC_{\ell, \text{HS}}}{dg^2} \right)^2}{(C_{\ell}^{\Lambda\text{CDM}})^2} \quad (17)$$

$$= \sum_{\ell} (2\ell + 1) \frac{f_{\text{sky}} N_{\text{HS}}^2 (C_{\ell, \text{HS}}(g = 1))^2}{(C_{\ell}^{\Lambda\text{CDM}})^2} \quad (18)$$

The 1σ error on g^2 is then given by $\sigma_{g^2} = (F_{g^2})^{-1/2}$. As such, we can then transform in the standard way to the error on g and compute the 5σ exclusion curve on g via $g \geq (50/F_{g^2})^{1/4}$.

In practice, one would likely do a joint power spectrum analysis considering the joint $TT/EE/TE$ signal data vector, $\mathbf{d}(\ell) = (C_{\ell, \text{HS}}^{TT}, C_{\ell, \text{HS}}^{EE}, C_{\ell, \text{HS}}^{TE})$. The Gaussian covariance between these spectra is (with $X, Y \in [T, E]$)

$$\text{Cov}[C_{\ell}^{XY}, C_{\ell}^{X'Y'}] = \frac{(C_{\ell}^{XX'} C_{\ell}^{YY'} + C_{\ell}^{XY'} C_{\ell}^{YX'})}{(2\ell + 1) f_{\text{sky}}} \quad (19)$$

The Fisher forecast for the joint T and E analysis is then

$$\begin{aligned} F_{g^2}^{T \times E} &= N_{\text{HS}}^2 \sum_{\ell, X, X'} \mathbf{d}(\ell) \text{Cov}^{-1} \mathbf{d}(\ell) \\ &= N_{\text{HS}}^2 \sum_{\ell, X, X'} C_{\ell, \text{HS}}^X \text{Cov}_{XX'}^{-1} C_{\ell, \text{HS}}^{X'}, \end{aligned} \quad (20)$$

where we have taken $X = \{TT, EE, TE\}$ and we have noted that the derivative of the signal with respect to g^2 is simply given by the fiducial power spectra (evaluated at $g = 1$).

In Figure 8, we show the exclusion curves computed for $N_{\text{HS}} = \{1, 100\}$ compared to the results from our matched filter analysis. Note that our forecast is for a CV-limited experiment comparable to *Planck* in sky coverage, but with zero noise, to present the best that a power spectrum search can do. From the figure, it is clear that until large N_{HS} values are reached (with $N_{\text{HS}} \geq 100$ corresponding roughly to $M_0 \leq 450 H_I$, for $g \simeq 10$ — see Figure 1), the localized matched-filter search dwarfs the bounds on g from a power spectrum search. We have also fixed a specific fiducial Λ CDM cosmology here and ignored the variations in the power spectrum associated with shifts in the cosmological parameters. In a true power spectrum search these would be partially degenerate with an inflationary hotspot signal, and a full parameter analysis would be required, which would lead




FIG. 9. Power spectra for TT (top left), EE (top right), and TE (bottom center), including contributions from particle-production hotspots, with the absolute value of the fractional hotspot contributions shown in the bottom panels. The shaded region in each bottom panel represents the uncertainty from cosmic variance. In each top panel we show the power spectra for $N_{\text{HS}} = 0, 1, 10, 100$ hotspots, as labeled (with $g = 10$ and $\eta_* = 150$ Mpc). Note that we have fixed fiducial values for our ΛCDM parameters and do not show variations in the power spectra arising from uncertainties on the parameters themselves.

to further degradation of the constraints, particularly in cases where the hotspot contributions are nearly in phase with the ΛCDM power spectrum (as in the example in Figure 9).⁶ This illustrates another advantage of the direct hotspot search, *i.e.*, we are not limited by parameter degeneracies.

VIII. SUMMARY AND DISCUSSION

Non-adiabatic production of massive particles is a feature appearing in many multi-field inflationary scenarios, and is known to leave observational signatures in the CMB.

⁶Note that the hotspot contributions are nearly in phase with the ΛCDM power spectrum for $\eta_* \approx \eta_{\text{rec}}$, but we find that this behavior does not persist for much smaller or larger η_* values.

Ref. [1] found that such signatures can be effectively constrained by applying a localized search to *Planck* temperature data, but could easily be missed by standard correlation function analyses. In this work we expand the horizons of localized CMB searches from temperature to polarization.

The use of polarization data has several advantages. Beyond providing an independent probe of inflationary physics, polarization comes with different noise properties, and because the Sunyaev-Zel'dovich effect is unpolarized to leading order there are many fewer possible false detections. Additionally, unlike the foregrounds, the hotspots should not have counterparts in B -mode maps, which gives us an additional method to verify any candidates. We also perform forecasting for current and future ground-based CMB experiments, which suggests that while the *Planck* T and E -mode constraints are comparable, in future experiments polarization will become a superior channel in which to look for such hotspots. Al-

though this may be a model-dependent phenomenon, we expect that this is a relatively generic property of profile-finding searches for inflationary signatures given that it has long been known that polarization can constrain parameters better in the CV limit (e.g., [15]).

We validate our E -mode search pipeline by injecting hotspots into simulated component-separated data, and successfully recover the input parameters for sufficiently high values of η_* and g . From the *Planck* polarization data, we find no evidence of new physics, with no evidence for hotspots with $\text{SNR} \geq 6$ (we impose a high SNR threshold due to the four free parameters in the template used in the search). This yields a relatively strong bound on couplings between the inflaton field, φ , and a much heavier massive field with $M/H_I \gtrsim 100$: $g \gtrsim 30$ for $\eta_* \gtrsim 30$ Mpc.

Our analysis of *Planck* E -mode polarization data is verified on both SEVEM and SMICA maps, which produce consistent results, with neither finding strong hotspot candidates. Like the temperature analysis of [1], our polarization analysis puts bounds on massive scalars with mass $M_0 \gtrsim 500H_I$, where H_I is the inflationary Hubble scale. The inflationary Hubble scale is not known, but the current upper limit from primordial B -mode searches is $H_I < 4.8 \times 10^{13}$ GeV [40]. This implies that our analysis may directly probe physics near the Grand Unified Theory scale, if H_I is near the current upper limit.

We also present a direct comparison of our profile-finding analysis method versus a power spectrum analysis. Noting that the corrections to the power spectrum from inflationary hotspots scale as N_{HS} , whereas our profile-finding search is independent of N_{HS} , we find that for sufficiently few produced particles ($N_{\text{HS}} \lesssim 100$) a profile-finding approach dominates over correlation function searches (which would be dominant for lighter, more numerous produced particles, e.g., [41, 42]). We emphasize a generic point, which is that our analysis and forecasts point to the fact that for rare dramatic events local searches are preferable, where by contrast, for common events with smaller individual impacts, the best way forward is the measurement of low-order correlators.

This work can be extended in several important ways. Minor improvements ($\sim 10 - 20\%$) to the *Planck* constraints could be obtained by jointly analysing the *Planck* PR4 temperature and E -mode polarization data. However, a much more exciting way forward is to apply our methods to higher-resolution CMB data. Figure 7 shows that ACT is close to CV-limited for $\eta_* \gtrsim 30$ Mpc, and would yield a factor of improvement between $1.2 - 2.75$ for E and ~ 1.2 for T , providing a clear motivation for a search using these data. Similarly, one could explore analyses using other high-resolution experiments, such as the South Pole Telescope [18, 37] or the Simons Observatory [19, 20]. Additionally, our pipeline can be simply extended to search for a wide variety of primordial features [6, 9], and even for localized signatures induced by novel later-time physics such as patchy screening by axions [43, 44] or dark photons [45, 46].

ACKNOWLEDGMENTS

We thank Soubhik Kumar for insightful discussions and comments on a draft of this work. LHA acknowledges support from the Columbia College Summer Research Fellowship, Columbia University. OHEP was a Junior Fellow of the Simons Society of Fellows. JCH acknowledges support from DOE grant DE-SC0011941, NASA grants 80NSSC22K0721 [ATP] and 80NSSC23K0463 [ADAP], the Sloan Foundation, and the Simons Foundation. This work utilized `numpy` [47], `matplotlib` [48], `healpy` [49], and `HEALPix` [50]. We acknowledge computing resources from Columbia University's Shared Research Computing Facility project, which is supported by NIH Research Facility Improvement Grant 1G20RR030893-01, and associated funds from the New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract C090171, both awarded April 15, 2010.

Appendix A: In-in computation of the curvature perturbation due to a massive particle

Here we derive the curvature perturbation due to a massive particle. Our treatment is very similar to that of [4]; here our primary purpose is to be pedagogical. We work in units with $c = 1$, and the derivative with respect to cosmic time is represented by a dot.

Given the natural assumption that particle production happens when the heavy field's time-dependent mass reaches its minimum, the mass should rapidly increase after production, and we thus assume that the particle quickly becomes non-relativistic. This leads to the action:

$$S_{\text{particle}} = - \int dt \sqrt{-g_{00}} M(t), \quad (\text{A1})$$

where $g_{\mu\nu}$ is the metric. We parametrize the metric fluctuations in the standard ADM variables [51], writing

$$ds^2 = -N^2 dt^2 + h_{ij}(dx^i + N^i dt)(dx^j + N^j dt) \quad (\text{A2})$$

$$h_{ij} = e^{2Ht}[(1 - 2\psi)\delta_{ij} + \gamma_{ij}],$$

where we choose γ_{ij} to be transverse and traceless. In this gauge there are no inflaton fluctuations, $\delta\varphi = 0$. It is proven in [52] that N and N^i are time-independent at leading order. Furthermore, N obeys the algebraic identity

$$N = 1 - \frac{\dot{\psi}}{H_I} = 1 + \frac{\dot{\zeta}}{H_I}, \quad (\text{A3})$$

where ζ is the comoving curvature perturbation:

$$\zeta = -\psi - H_I \frac{\delta\varphi}{\dot{\varphi}} = -\psi \quad (\text{A4})$$

in our gauge. We can then write our particle action more explicitly,

$$S_{\text{particle}} = - \int dt N M(t) = - \int dt \left(M(t) + \dot{\zeta} \frac{M(t)}{H_I} \right). \quad (\text{A5})$$

The second term induces the contributing curvature perturbation in our model. It will be more convenient to do our calculation in conformal time η ($d\eta \equiv dt/a$), noting that we may write the relevant term in the action as

$$\begin{aligned} S_{\text{particle}} \supset & - \int d\eta \partial_\eta \zeta \frac{M(\eta)}{H_I} \\ = & - \int_{\eta_*}^0 \frac{M(\eta)}{H_I} \int \frac{d^3 \mathbf{k}}{(2\pi)^3} \partial_\eta \zeta_{\mathbf{k}} e^{i\mathbf{k} \cdot \mathbf{x}_{\text{HS}}} \end{aligned} \quad (\text{A6})$$

Also note that we can write the comoving curvature perturbation in terms of creation and annihilation operators, for a simple inflationary model [2]:

$$\zeta_{\mathbf{k}} = \frac{H_I^2}{\dot{\varphi}_0 \sqrt{2k^3}} [(1 - ik\eta) e^{ik\eta} a_{\mathbf{k}}^\dagger + (1 + ik\eta) e^{-ik\eta} a_{-\mathbf{k}}] \quad (\text{A7})$$

From here we have the requisite tools to compute the average curvature perturbation associated with the massive particle in the in-in formalism. Recall that in the in-in formalism we have the master formula in the commutator form in the interaction picture [53, 54] for the expectation value of some operator, $\mathcal{O}(\eta)$, given some interaction Hamiltonian, H_{int} :

$$\begin{aligned} \langle \mathcal{O}(\eta) \rangle &= \sum_{N=0}^{\infty} i^N \int_{-\infty}^{\eta} d\eta_N \dots \int d\eta_1 \\ &\times \langle 0 | [H_{\text{int}}(\eta_1), [H_{\text{int}}(\eta_2), \dots [H_{\text{int}}(\eta_N), \mathcal{O}(\eta)] \dots]] | 0 \rangle \end{aligned} \quad (\text{A8})$$

Note that to leading order this is simply

$$\langle \mathcal{O}(\eta) \rangle = i \int_{-\infty}^{\eta} \langle 0 | [H_{\text{int}}(\eta'), \mathcal{O}(\eta)] | 0 \rangle d\eta' \quad (\text{A9})$$

One can then observe that, by hermicity, that $\langle 0 | [H_{\text{int}}(\eta'), \mathcal{O}(\eta)] | 0 \rangle = 2i \text{Im}\{\langle H_{\text{int}} \mathcal{O} \rangle\}$. From here we can note that

$$\begin{aligned} \langle \zeta_{\mathbf{k}}(\eta \rightarrow 0) \rangle &= \\ - 2 \lim_{\eta \rightarrow 0} \int_{\eta_*}^0 d\eta' \frac{M(\eta')}{H_I} \int \frac{d^3 \mathbf{p}}{(2\pi)^3} e^{i\mathbf{k} \cdot \mathbf{x}_{\text{HS}}} \text{Im}\{\langle \partial_\eta \zeta_{\mathbf{p}} \zeta_{\mathbf{k}} \rangle\} \end{aligned} \quad (\text{A10})$$

Evaluating the central matrix element yields

$$\begin{aligned} \lim_{\eta \rightarrow 0} \langle \partial_\eta \zeta_{\mathbf{p}} \zeta_{\mathbf{k}} \rangle &= \frac{H_I^4}{2\dot{\varphi}_0^2} \times \frac{p^2 \eta'}{\sqrt{p^3 k^3}} e^{-ip\eta'} \langle a_{-\mathbf{p}} a_{\mathbf{k}}^\dagger \rangle \\ &= \frac{H_I^4}{2\dot{\varphi}_0^2} \times \frac{p^2 \eta'}{\sqrt{p^3 k^3}} e^{-ip\eta'} \times (2\pi)^3 \delta^{(3)}(\mathbf{p} + \mathbf{k}) \end{aligned} \quad (\text{A11})$$

Taking the imaginary part and integrating over $\mathbf{p}$ gives us

$$\langle \zeta_{\mathbf{k}} \rangle = \frac{H_I^3}{\dot{\varphi}_0^2} \int_{\eta_*}^0 d\eta' \frac{M(\eta') \eta'}{k} \sin k\eta' e^{-i\mathbf{k} \cdot \mathbf{x}_{\text{HS}}} \quad (\text{A12})$$

This integral can be done if we approximate the mass around η_* by dropping the constant M_0 term, such that $M \simeq \frac{g\dot{\varphi}_0}{H_I} \ln |\eta_*/\eta|$. This then yields the hotspot profile from Sec. II:

$$\langle \zeta_{\mathbf{k}} \rangle = \frac{gH_I^2}{\dot{\varphi}_0} e^{-i\mathbf{k} \cdot \mathbf{x}_{\text{HS}}} \frac{\text{Si}(k\eta_*) - \sin(k\eta_*)}{k^3}, \quad (\text{A13})$$

where $\text{Si}(x) = \int_0^x \frac{\sin t}{t} dt$ is the integrated sine function.

Appendix B: Visual inspection of the candidates

In Figure 10, we present a visual inspection of the results of our polarization hotspot search. We plot the SEVEM E -mode map at the locations of all of the $\text{SNR} \geq 5$ hotspot candidates. We also show SNR maps (*i.e.*, $\hat{g}/\sigma_g$) on the top row. We note that many candidates are found near the edges of the mask, and as such are likely ringing as opposed to primordial. We also emphasize that due to the four parameters in our search template, moderately high SNR values are to be expected due to chance fluctuations.

Appendix C: SMICA hotspot candidates

We present the recovered parameters from the application of our analysis to the *Planck* SMICA maps in Table II. We emphasize that beyond the fact that the SEVEM and SMICA maps produce a similar number of candidates, 20 versus 23, the candidates are also recovered with similar parameters. This demonstrates the robustness of our pipeline.

[1] O. H. E. Philcox, S. Kumar, and J. C. Hill, Searching for inflationary particle production in Planck data, *Phys. Rev. D* **111**, 103523 (2025), arXiv:2405.03738 [astro-

ph.CO].

[2] D. Baumann, Tasi lectures on inflation (2012), arXiv:0907.5424 [hep-th].



FIG. 10. SNR maps (top row) and SEVEM component-separated E -mode maps (second row) for all $\text{SNR} \geq 5$ hotspot candidates, which are circled in green. Each image is a $5^\circ \times 5^\circ$ Cartesian projection around the centers given in the title. The SNR maps have colorbar truncated to ± 5 , while the maximal absolute value for each other plot is shown in purple text. The bounding rectangles presented are merely the flat-space tiles used in our SZIFI analysis.

- [3] P. Berglund and G. Ren, Multi-field inflation from string theory (2009), arXiv:0912.1397 [hep-th].
- [4] J. H. Kim, S. Kumar, A. Martin, and Y. Tsai, Cosmological particle production and pairwise hotspots on the CMB, JHEP **11**, 158, arXiv:2107.09061 [hep-ph].
- [5] R. Flauger, M. Mirbabayi, L. Senatore, and E. Silverstein, Productive interactions: heavy particles and non-gaussianity, Journal of Cosmology and Astroparticle Physics **2017** (10), 058–058.
- [6] M. Mirbabayi, L. Senatore, E. Silverstein, and M. Zaldarriaga, Gravitational waves and the scale of inflation, Physical Review D **91**, 10.1103/physrevd.91.063518 (2015).
- [7] L. Kofman, A. Linde, X. Liu, A. Maloney, L. McAllister, and E. Silverstein, Beauty is attractive: Moduli trapping at enhanced symmetry points, Journal of High Energy Physics **2004**, 030–030 (2004).
- [8] J. Maldacena, A model with cosmological bell inequalities, Fortschritte der Physik **64**, 10–23 (2015).
- [9] M. Münchmeyer and K. M. Smith, Higher N-point function data analysis techniques for heavy particle production and WMAP results, Phys. Rev. D **100**, 123511 (2019).

SNR	Longitude [°]	Latitude [°]	$\hat{g}$	η_* [Mpc]	η_{HS} [Mpc]
5.2	144.4	81.8	29	46.4	242.4
5.0	114.8	81.5	13	215.4	221.7
5.2	214.6	64.6	28	46.4	259.3
5.0	230.9	63.6	10	599.5	240.0
5.1	265.5	71.2	12	129.2	292.2
5.2	78.4	62.4	15	215.4	378.4
5.3	339.2	59.3	10	359.4	290.8
5.1	224.6	43.1	293	16.7	294.1
5.0	106.9	36.1	46	27.8	257.7
5.0	275.4	40.4	71	27.8	257.7
5.1	74.3	33.6	902	10	272.2
5.0	82.8	30.5	24	46.4	242.4
5.0	93.1	25.7	19	77.4	343.8
5.5	324.0	21.6	106	27.8	293.1
5.1	78.5	-34.5	28	46.4	242.4
5.4	116.2	-33.7	1194	10	272.2
5.1	270.8	-45.6	186	16.7	266.8
5.3	342.3	-41.0	19	77.4	231.2
5.1	356.5	-40.7	1052	10	272.2
5.1	29.1	-46.1	1094	10	272.2
5.0	121.3	-44.3	91	27.8	303.2
5.3	292.7	-57.1	859	10	272.2
5.2	181.8	-67.5	87	27.8	288.0

TABLE II. Inferred parameters for the hotspot candidates from *Planck* SMICA *E*-mode maps. Notice that none of them are strong candidates ($\text{SNR} \geq 6$), and that almost all have direct counterparts with candidates from the SEVEM maps.

- (2019), arXiv:1910.00596 [astro-ph.CO].
- [10] T. Kim, J. H. Kim, S. Kumar, A. Martin, M. Münchmeyer, and Y. Tsai, Probing cosmological particle production and pairwise hotspots with deep neural networks, *Phys. Rev. D* **108**, 043525 (2023), arXiv:2303.08869 [hep-ph].
- [11] Y. Akrami *et al.* (Planck), *Planck* intermediate results. LVII. Joint Planck LFI and HFI data processing, *Astron. Astrophys.* **643**, A42 (2020), arXiv:2007.04997 [astro-ph.CO].
- [12] J. Carron, M. Mirmelstein, and A. Lewis, Cmb lensing from planck pr4 maps, *Journal of Cosmology and Astroparticle Physics* **2022** (09), 039.
- [13] O. H. E. Philcox, Searching for inflationary physics with the CMB trispectrum. III. Constraints from Planck, *Phys. Rev. D* **111**, 123534 (2025), arXiv:2502.06931 [astro-ph.CO].
- [14] Y. B. Zeldovich and R. A. Sunyaev, The Interaction of Matter and Radiation in a Hot-Model Universe, *Astrophysics & Space Science* **4**, 301 (1969).
- [15] S. Galli, K. Benabed, F. Bouchet, J.-F. Cardoso, F. Elsner, E. Hivon, A. Mangilli, S. Prunet, and B. Wandelt, Cmb polarization can constrain cosmology better than cmb temperature, *Physical Review D* **90**, 10.1103/physrevd.90.063504 (2014).
- [16] S. Naess *et al.* (ACT), The Atacama Cosmology Telescope: DR6 Maps, arXiv preprint arXiv:2503.14451 (2025), arXiv:2503.14451 [astro-ph.CO].
- [17] E. Calabrese *et al.* (ACT), The Atacama Cosmology Telescope: DR6 Constraints on Extended Cosmological Models, arXiv preprint arXiv:2503.14454 (2025), arXiv:2503.14454 [astro-ph.CO].
- [18] E. Camphuis *et al.* (SPT-3G), SPT-3G D1: CMB temperature and polarization power spectra and cosmology from 2019 and 2020 observations of the SPT-3G Main field, arXiv preprint arXiv:2506.20707 (2025), arXiv:2506.20707 [astro-ph.CO].
- [19] P. Ade *et al.* (Simons Observatory), The Simons Observatory: Science goals and forecasts, *JCAP* **02**, 056, arXiv:1808.07445 [astro-ph.CO].
- [20] M. Abitbol *et al.* (Simons Observatory), The Simons Observatory: science goals and forecasts for the enhanced Large Aperture Telescope, *JCAP* **08**, 034, arXiv:2503.00636 [astro-ph.IM].
- [21] M. Tristram *et al.*, Cosmological parameters derived from the final Planck data release (PR4), *Astron. Astrophys.* **682**, A37 (2024), arXiv:2309.10034 [astro-ph.CO].
- [22] S. Dodelson and F. Schmidt, *Modern Cosmology* (2020).
- [23] S. Weinberg, *Cosmology* (2008).
- [24] N. Aghanim *et al.* (Planck), Planck 2015 results. XXII. A map of the thermal Sunyaev-Zeldovich effect, *Astron. Astrophys.* **594**, A22 (2016), arXiv:1502.01596 [astro-ph.CO].
- [25] Í. Zubeldia, J. Chluba, and R. Battye, Mitigating the impact of the CIB on galaxy cluster SZ detection with spectrally constrained matched filters, *Mon. Not. Roy. Astron. Soc.* **522**, 5123 (2023), arXiv:2212.07410 [astro-ph.CO].
- [26] Í. Zubeldia, A. Rotti, J. Chluba, and R. Battye, Galaxy cluster sz detection with unbiased noise estimation: an iterative approach, *Monthly Notices of the Royal Astronomical Society* **522**, 4766–4780 (2023).
- [27] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere, *The Astrophysical Journal* **622**, 759–771 (2005).
- [28] F. McCarthy and J. C. Hill, Component-separated, CIB-cleaned thermal Sunyaev-Zel'dovich maps from Planck PR4 data with a flexible public needlet ILC pipeline, *Phys. Rev. D* **109**, 023528 (2024), arXiv:2307.01043 [astro-ph.CO].
- [29] A. Lewis and A. Challinor, CAMB: Code for Anisotropies in the Microwave Background, *Astrophysics Source Code Library*, record ascl:1102.026 (2011).
- [30] J. M. Ezquiaga, J. García-Bellido, and V. Vennin, Massive Galaxy Clusters Like El Gordo Hint at Primordial Quantum Diffusion, *Phys. Rev. Lett.* **130**, 121003 (2023), arXiv:2207.06317 [astro-ph.CO].
- [31] B. Shakya, The tachyonic higgs and the inflationary universe (2023), arXiv:2301.08754 [hep-ph].
- [32] S. Kumar and N. Weiner, Early galaxies from rare inflationary processes and jwst observations (2025), arXiv:2502.08701 [astro-ph.CO].
- [33] M. G. Haehnelt and M. Tegmark, Using the kinematic Sunyaev-Zeldovich effect to determine the peculiar velocities of clusters of galaxies, *Mon. Not. Roy. Astron. Soc.* **279**, 545 (1996), arXiv:astro-ph/9507077.
- [34] J.-B. Melin, J. G. Bartlett, and J. Delabrouille, Catalog extraction in sz cluster surveys: a matched filter approach, *Astron. Astrophys.* **459**, 341 (2006), arXiv:astro-ph/0602424.
- [35] Z. Atkins *et al.*, The Atacama Cosmology Telescope: semi-analytic covariance matrices for the DR6 CMB power spectra, *JCAP* **05**, 015, arXiv:2412.07068 [astro-ph.CO].

- ph.CO].
- [36] T. Louis *et al.* (ACT), The Atacama Cosmology Telescope: DR6 Power Spectra, Likelihoods and Λ CDM Parameters, arXiv preprint arXiv:2503.14452 (2025), arXiv:2503.14452 [astro-ph.CO].
 - [37] L. E. Bleem, T. M. Crawford, B. Ansarinejad, B. A. Benson, S. Bocquet, J. E. Carlstrom, C. L. Chang, R. Chown, A. T. Crites, T. d. Haan, M. A. Dobbs, W. B. Everett, E. M. George, R. Gualtieri, N. W. Halverson, G. P. Holder, W. L. Holzapfel, J. D. Hrubes, L. Knox, A. T. Lee, D. Luong-Van, D. P. Marrone, J. J. McMahon, S. S. Meyer, M. Millea, L. M. Mocanu, J. J. Mohr, T. Natoli, Y. Omori, S. Padin, C. Pryke, S. Raghunathan, C. L. Reichardt, J. E. Ruhl, K. K. Schaffer, E. Shirokoff, Z. Staniszewski, A. A. Stark, J. D. Vieira, and R. Williamson, Cmb/ks+compton-y maps from 2500 deg² of spt-sz and planck survey data, *The Astrophysical Journal Supplement Series* **258**, 36 (2022).
 - [38] K. Abazajian *et al.*, CMB-S4 Science Case, Reference Design, and Project Plan, arXiv preprint arXiv:1907.04473 (2019), arXiv:1907.04473 [astro-ph.IM].
 - [39] B. Beringue *et al.*, The Atacama Cosmology Telescope: DR6 Power Spectrum Foreground Model and Validation, arXiv preprint arXiv:2506.06274 (2025), arXiv:2506.06274 [astro-ph.CO].
 - [40] P. A. R. Ade *et al.* (BICEP, Keck), Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season, *Phys. Rev. Lett.* **127**, 151301 (2021), arXiv:2110.00483 [astro-ph.CO].
 - [41] Planck Collaboration and Y. e. Akrami, Planck 2018 results. IX. Constraints on primordial non-Gaussianity, *Astron. Astrophys.* **641**, A9 (2020), arXiv:1905.05697 [astro-ph.CO].
 - [42] D. J. H. Chung, E. W. Kolb, A. Riotto, and I. I. Tkachev, Probing planckian physics: Resonant production of particles during inflation and features in the primordial power spectrum, *Physical Review D* **62**, 10.1103/physrevd.62.043508 (2000).
 - [43] C. Mondino, D. Pirvu, J. Huang, and M. C. Johnson, Axion-induced patchy screening of the Cosmic Microwave Background, *JCAP* **10**, 107, arXiv:2405.08059 [hep-ph].
 - [44] S. Goldstein, F. McCarthy, C. Mondino, J. C. Hill, J. Huang, and M. C. Johnson, Constraints on Axions from Patchy Screening of the Cosmic Microwave Background, *Phys. Rev. Lett.* **134**, 081001 (2025), arXiv:2409.10514 [astro-ph.CO].
 - [45] D. Pirvu, J. Huang, and M. C. Johnson, Patchy screening of the CMB from dark photons, *JCAP* **01**, 019, arXiv:2307.15124 [hep-ph].
 - [46] F. McCarthy, D. Pirvu, J. C. Hill, J. Huang, M. C. Johnson, and K. K. Rogers, Dark Photon Limits from Patchy Dark Screening of the Cosmic Microwave Background, *Phys. Rev. Lett.* **133**, 141003 (2024), arXiv:2406.02546 [hep-ph].
 - [47] C. R. Harris *et al.*, Array programming with NumPy, *Nature* **585**, 357 (2020), arXiv:2006.10256 [cs.MS].
 - [48] J. D. Hunter, Matplotlib: A 2d graphics environment, *Computing in Science & Engineering* **9**, 90 (2007).
 - [49] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski, healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python, *Journal of Open Source Software* **4**, 1298 (2019).
 - [50] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere, *Astrophys. J.* **622**, 759 (2005), arXiv:astro-ph/0409513.
 - [51] R. Arnowitt, S. Deser, and C. W. Misner, Republication of: The dynamics of general relativity, *General Relativity and Gravitation* **40**, 1997–2027 (2008).
 - [52] J. Maldacena, Non-gaussian features of primordial fluctuations in single field inflationary models, *Journal of High Energy Physics* **2003**, 013–013 (2003).
 - [53] S. Weinberg, Quantum contributions to cosmological correlations, *Physical Review D* **72**, 10.1103/physrevd.72.043514 (2005).
 - [54] M. H. G. Lee, E. Pajer, M. Giroux, H. S. Hannesdotir, S. Mizera, and C. Pasiernik, Records from the S-Matrix Marathon: A Timeless History of Time (2024) arXiv:2410.00227 [hep-th].