

First observational tests of eternal inflation: Analysis methods and WMAP 7-year resultsStephen M. Feeney,^{1,*} Matthew C. Johnson,^{2,3,†} Daniel J. Mortlock,^{4,‡} and Hiranya V. Peiris^{1,5,§}¹*Department of Physics and Astronomy, University College London, London WC1E 6BT, United Kingdom*²*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*³*California Institute of Technology, Pasadena, California 91125, USA*⁴*Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London, SW7 2AZ, United Kingdom*⁵*Institute of Astronomy and Kavli Institute for Cosmology, University of Cambridge, Cambridge CB3 0HA, United Kingdom*

(Received 17 December 2010; published 8 August 2011)

In the picture of eternal inflation, our observable universe resides inside a single bubble nucleated from an inflating false vacuum. Many of the theories giving rise to eternal inflation predict that we have causal access to collisions with other bubble universes, providing an opportunity to confront these theories with observation. We present the results from the first observational search for the effects of bubble collisions, using cosmic microwave background data from the WMAP satellite. Our search targets a generic set of properties associated with a bubble-collision spacetime, which we describe in detail. We use a modular algorithm that is designed to avoid *a posteriori* selection effects, automatically picking out the most promising signals, performing a search for causal boundaries, and conducting a full Bayesian parameter estimation and model selection analysis. We outline each component of this algorithm, describing its response to simulated CMB skies with and without bubble collisions. Comparing the results for simulated bubble collisions to the results from an analysis of the WMAP 7-year data, we rule out bubble collisions over a range of parameter space. Our model selection results based on WMAP 7-year data do not warrant augmenting Λ CDM with bubble collisions. Data from the Planck satellite can be used to more definitively test the bubble-collision hypothesis.

DOI: [10.1103/PhysRevD.84.043507](https://doi.org/10.1103/PhysRevD.84.043507)

PACS numbers: 98.80.Es, 98.70.Vc, 98.80.Cq

I. INTRODUCTION

Observations of the cosmic microwave background (CMB) radiation have ushered in a new era of precision cosmology. Full-sky temperature maps produced by the Wilkinson Microwave Anisotropy Probe (WMAP) [1] have confirmed with high precision that the observed temperature fluctuations are consistent with a nearly Gaussian and scale invariant primordial power spectrum, as predicted by inflation. The recently launched Planck satellite [2] has a resolution 3 times better than that of WMAP, with an order of magnitude greater sensitivity, and significantly wider frequency coverage (allowing for far more robust foreground removal, and therefore reduced systematics). These high quality data sets allow for the possibility of observing deviations from the standard inflationary paradigm, some of which could have drastic consequences for our understanding of the Universe and its origins.

Perhaps the largest gap in our description of the early Universe lies in an understanding of its initial conditions. One possibility, motivated by the proliferation of vacua in compactifications of string theory (known as the string theory landscape [3]), is that our observable Universe is only a tiny piece of a vast multiverse, the majority of which is still inflating. This picture of *eternal inflation* (for a

review, see, e.g., Ref. [4]) arises when the rate at which local regions exit an inflating phase is outpaced by the accelerated expansion of the inflating background. Eternal inflation is a fairly generic consequence of any theory containing positive vacuum energy and multiple vacua, highlighting the importance of understanding how this scenario might be confronted with observational tests.

The first attempts to embed our cosmology inside an eternally inflating universe led to “open inflation” [5,6]; see Ref. [7] for a review. In this scenario, a scalar field (or set of scalar fields) has a potential with a high energy metastable minimum that drives the eternally inflating phase. Transitions out of this vacuum proceed via the Coleman–De Luccia (CDL) instanton [8,9], resulting in expanding bubbles inside which the scalar field rests on an inflationary plateau. The symmetries of the CDL instanton ensure that there is a very nearly homogeneous and isotropic open universe inside the bubble; inflation, reheating, and standard cosmological evolution follow.

In any given bubble, the future light cone of the nucleation event forms the “big bang” (where the scale factor vanishes) of an open Friedmann–Robertson–Walker (FRW) universe. The eternally inflating phase outside our bubble can therefore be thought of as a pre-big bang epoch, and one might expect inflation to erase any of the scant observational evidence of our parent vacuum. In single-bubble open inflation, various anomalies are induced in the CMB temperature power spectrum (see Ref. [7] and references therein), but unfortunately, the size of these effects

*stephen.feeney.09@ucl.ac.uk

†mjohnson@perimeterinstitute.ca

‡mortlock@ic.ac.uk

§h.peiris@ucl.ac.uk

decreases with the present energy density in curvature (related to the number of inflationary e -folds), rendering them negligible at all but the lowest multipoles where cosmic variance dominates. However, our bubble does not evolve in isolation. There are other nucleation events from the false vacuum, containing a phase that might be identical to ours, or perhaps very different. If one of these secondary nucleation events occurs close enough to our bubble wall, then a collision inevitably results. In fact, since our bubble grows to reach infinite size, there are an infinite number of collisions [10–13] (a finite subset of which are causally accessible to any one observer). This raises the possibility that if such collisions are both survivable and only small perturbations on top of standard cosmology, they might leave observable signatures of eternal inflation [14]; it is these signatures that our analysis targets.

If we are to detect such bubble collisions, their predicted signatures must be consistent with our observed cosmology, but sufficiently distinct to be differentiated from other possible signals in the CMB. In addition, the theory must predict that we expect to have causal access to bubble collisions. While these criteria are not met in every model of eternal inflation, recent work [10–26] (for a review, see Ref. [27]) has established that bubble collisions could in some theories be both expected and detectable. Bubble collisions produce a fairly characteristic set of inhomogeneities in the very early Universe, which are processed into temperature anisotropies in the CMB. From the spherical symmetry of the colliding bubbles, the collision possesses azimuthal symmetry, and by causality must be confined to a disk on the sky. The CMB temperature and its derivatives need not be constant across the causal boundary. Therefore, the signals we are searching for are localized, and because they are primordial, consist of a long-wavelength modulation of the standard inflationary density fluctuations inside the affected region [20]. The amplitude and angular scale of the signal are dependent upon the underlying model and kinematics of the collision.

These general features suggest a set of strategies for data analysis. The localization of the collision implies that wavelet analysis could be a sensitive tool for picking out both the location and angular scale of a candidate signal. The causal boundary, across which the temperature and its derivatives need not be constant, suggests the use of edge-detection algorithms similar to those used in searches for cosmic strings [28–31]. Finally, the prediction that the temperature modulation induced by the collision is a rather long wavelength yields a sufficiently generic template to perform a full Bayesian parameter estimation and model selection analysis.

In this paper, we describe a modular analysis algorithm designed to look for the signatures of eternal inflation, and apply it to the WMAP 7-year data [32]. This algorithm can easily be adapted to test any model that predicts a population of spatially localized sources in addition to the

standard fluctuations predicted by Λ CDM. A summary of our results was presented in Ref. [33]; in this paper we describe our analysis in detail. Currently available full-sky CMB data are rather limited in their sensitivity to the signatures of bubble collisions listed above; the main current limitation is the low resolution. Therefore, we apply our algorithm to current data mainly as a validation exercise; to exploit its full power would require future high resolution data, e.g., from Planck.

The individual steps of our analysis pipeline are calibrated using realistic simulations of the WMAP experiment with and without bubble collisions. The calibrated pipeline applied to data is fully automated, identifying the candidate signals and processing them without any human intervention. This removes any *a posteriori* choices from our analysis, which must be carefully avoided in any analysis of a large data set such as the WMAP 7-year data [34].

The plan of the paper is as follows. In Sec. II, we review some of the background on bubble collisions in eternal inflation and outline the predicted observable signatures. Our analysis pipeline is summarized in Sec. III. We describe some properties of the WMAP experiment and our simulations in Sec. IV and detail our analysis tools in Sec. V. Section VI summarizes the results of our analysis of the WMAP 7-year data, and we conclude in Sec. VII.

II. THE OBSERVABLE EFFECTS OF BUBBLE COLLISIONS

The simplest model of eternal inflation involves a single scalar field in four dimensions, with a double-well potential. In many models (as long as the average curvature of the potential between the minima is small compared to the Planck scale), the CDL instanton [8,9] mediates a transition from the false (higher energy) to the true (lower energy) vacuum. This tunneling event corresponds to the appearance of an expanding bubble of the true vacuum embedded in the false. As long as the probability that a bubble nucleates in each horizon volume of the false vacuum during a Hubble time is less than 1 (so that the background expansion of the false vacuum on average prevents bubbles from merging), the phase transition never completes and inflation is eternal [10] (see Ref. [35] for a modern treatment of the percolation problem in eternal inflation). The $O(4)$ invariance of the instanton guarantees that the bubble interior possesses $SO(3,1)$ symmetry, and therefore contains an infinite open FRW universe. Although homogeneity is ensured by the symmetries of the instanton, if the interior of a bubble is to resemble our own Universe, a second epoch of inflation inside the bubble is necessary to dilute the negative curvature and provide the correct spectrum of primordial density perturbations to seed structure. Models of this type are known in the literature as open inflation and have been explored in detail (see Ref. [7]).

The signatures of single-bubble open inflation include negative curvature and modifications to the power spectrum. These modifications are most important at large angular scales (see Ref. [7], and references therein), where cosmic variance is dominant and would be very difficult to detect. Since curvature alone would not be a very distinguishing prediction, we do not consider these signals further.

A less ambiguous signature of eternal inflation would be the visible remnants of collisions between bubbles. Although the bubbles formed during eternal inflation do not percolate, there are many (in fact, an infinite number of) collisions. These collisions lead to inhomogeneities in the inner-bubble cosmology, perhaps leaving observable signatures in the CMB. Assessing the observational consequences of bubble collisions in an eternally inflating universe has been an active area of research [13,14,17–26] (for a review, see Ref. [27]). These studies have established that a number of criteria necessary for the observation of bubble collisions [14] can be satisfied, at least in some models:

Compatibility: In order to satisfy this criterion, there must be a bubble we can collide with that only minimally disturbs the homogeneity of the observable portion of the surface of last scattering. Such collisions do seem to exist, as evidenced by thin-wall junction condition analysis [17,19] as well as numerical simulations [18] and a study of the inflaton field in a background thin-wall collision geometry [20].

Probability: We should expect to have a collision in our causal past. The number of collisions in our past is $\bar{N} = \lambda V_4^F$, where λ is the bubble nucleation probability per unit four-volume and V_4^F is the four-volume outside the bubble to which we have causal access. The expected number (assuming the original open FRW foliation) is formally infinite [14]; however, collisions that contribute to this divergence only produce very long-wavelength fluctuations at last scattering, and so would not be observable [18,22] (this is similar to the infrared divergence found in models of slow-roll inflation). The average number of collisions that affect the *observable* portion of the surface of last scattering is finite [22,27] and is given by

$$\bar{N} \simeq \frac{16\pi}{3} \lambda H_F^{-4} \left(\frac{H_F}{H_I}\right)^2 \Omega_k^{1/2}, \quad (1)$$

where H_F is the false vacuum Hubble constant, H_I is the inflationary Hubble constant inside our bubble, and Ω_k is the current component of energy in curvature. For the expected number of observable collisions to be 1 or larger, the separation of scales between H_F and H_I must be large enough to compensate for the low probability λ (which is exponentially suppressed because this is a tunneling process) and the observational constraint $\Omega_k \lesssim .0084$ [36]. Given a particular scalar potential underlying eternal

inflation, \bar{N} for each possible type of collision is fixed. However, in a theory with a complicated potential landscape for the scalar field(s), it makes sense to think of \bar{N} as a continuous parameter with some prior probability distribution.¹ Without detailed knowledge of the theory underlying eternal inflation and an associated measure, it is difficult to assess how likely it is to have $\bar{N} > 1$, but see Refs. [22,27] for some speculative comments. There is also an exponential pressure from the nucleation rates toward $\bar{N} \gg 1$ or $\bar{N} \ll 1$. In the following, we assume \bar{N} can be of the order of 1.

Observability: Since the effects of a collision must pass through the entire inner-bubble cosmology, they can be thought of as perturbations of the big bang in an FRW cosmology. As such, they are stretched by inflation, and we expect the strength of most signatures to scale with (some power of) Ω_k . We therefore must require that there are not too many more e -folds than required to satisfy the observational bound on curvature. Given a field theory model, the number of e -folds of inflation inside the bubble is uniquely determined by the instanton. However, if we consider a landscape of scalar potentials, then it is necessary to find a measure over the number of inflationary e -folds (or equivalently Ω_k). For some work in this direction, see, e.g., Refs. [37,38].

Much remains to be learned about the full spectrum of possible outcomes of a bubble collision and the exact details of the associated observational signatures. Nevertheless, all potentially observable bubble collisions involving two bubbles share a sufficiently general set of properties to allow for a meaningful observational search even in the absence of a detailed model. In summary, we expect all such observable bubble collisions to possess the following:

- (i) Azimuthal symmetry: A collision leaves an imprint on the CMB sky that has azimuthal symmetry. This is a consequence of the SO(2,1) symmetry of the spacetime describing the collision of two vacuum bubbles [13,14,16].
- (ii) A causal boundary: The surface of last scattering can only be affected inside the future light cone of a collision event. The intersection of our past light cone, the future light cone of a collision, and the surface of last scattering is a ring. This is the causal boundary of the collision on the CMB sky. The temperature and its derivatives need not be continuous across this boundary. Neglecting the backreaction of the collision on the geometry of the bubble interior, the distribution of ring sizes was found in Ref. [22] to be

$$\frac{d\bar{N}}{d\theta_{\text{crit}}} \sim 4\pi\lambda H_F^{-4} \left(\frac{H_F}{H_I}\right)^2 \Omega_k^{1/2} \sin(\theta_{\text{crit}}), \quad (2)$$

¹We return to this point in Sec. VC when discussing the Bayesian framework for testing bubble-collision models.

where θ_{crit} is the angular radius measured from the center of the disk to the causal boundary and the other quantities are as defined in Eq. (1).

- (iii) An overall modulation of the background fluctuations: We assume that the temperature fluctuations, including the effects of the collision, at a location on the sky $\hat{\mathbf{n}}$ can be written as [20,39]

$$\frac{\delta T(\hat{\mathbf{n}})}{T_0} = (1 + f(\hat{\mathbf{n}}))(1 + \delta(\hat{\mathbf{n}})) - 1, \quad (3)$$

where $f(\hat{\mathbf{n}})$ is the modulation induced by the collision and $\delta(\hat{\mathbf{n}})$ are the temperature fluctuations induced by modes set down during inflation. This is motivated by the observation that the main effect of the bubble collision is to slightly advance or retard the inflaton inside our bubble. The modulation is multiplicative under the assumption that the normal inflationary density fluctuations simply “paint” the perturbed surface of last scattering and have identical statistical properties in the regions both affected and unaffected by the collision.

- (iv) Long-wavelength modulation: A collision is a pre-inflationary relic. The effects of a collision inside the causal boundary are stretched by inflation, and so we can expect that the relevant fluctuations are large scale. As we describe below, this implies that the temperature modulation due to a collision centered on the north pole ($\theta = 0$) has the form

$$f(\hat{\mathbf{n}}) = (c_0 + c_1 \cos\theta + \mathcal{O}(\cos^2\theta))\Theta(\theta_{\text{crit}} - \theta), \quad (4)$$

where the c_i are constants related to the properties of the collision, θ is the angle measured from the center of the affected disk, and $\Theta(\theta_{\text{crit}} - \theta)$ is a step function at the causal boundary θ_{crit} . Truncating the sum at $\mathcal{O}(\cos\theta)$, the constants c_0 and c_1 can be expressed in terms of a central amplitude z_0 and edge discontinuity z_{crit} :

$$c_0 = \frac{z_{\text{crit}} - z_0 \cos\theta_{\text{crit}}}{1 - \cos\theta_{\text{crit}}}, \quad c_1 = \frac{z_0 - z_{\text{crit}}}{1 - \cos\theta_{\text{crit}}}, \quad (5)$$

as shown in Fig. 1. Allowing the collision to be centered on an arbitrary location $\{\theta_0, \phi_0\}$ on the celestial sphere, the induced temperature modulation can be expressed as a function of five parameters: $\{z_0, z_{\text{crit}}, \theta_{\text{crit}}, \theta_0, \phi_0\}$. A modulation of this form was first derived in Ref. [20], where it was obtained from the observed modulation of a field representing the inflaton inside our bubble, numerically evolved in a background thin-wall bubble-collision geometry. These authors did not predict the existence of a temperature discontinuity z_{crit} . While further work is needed to better predict the precise form of the template, in our analysis we




FIG. 1. The radial temperature modulation Eq. (4) induced by a bubble collision centered on the north pole ($\theta = 0$).

allow bubble collisions to produce modulations with and without discontinuities.

In Fig. 2, we show the Poincaré-disk representation (see Ref. [14] for the details of this construction) of the surface of last scattering inside our parent bubble. The collision affects the shaded portion of this surface. The observed CMB is formed at the intersection of our past light cone (dashed circle) with the surface of last scattering, which in this case includes regions both affected and unaffected by the collision. From the underlying azimuthal symmetry, the




FIG. 2 (color online). A Poincaré-disk representation of the surface of last scattering inside our parent bubble, with one dimension suppressed. The future light cone of the collision at this time is denoted by the boundary of the shaded region; the latter represents the portions of the surface of last scattering that are to the future of the collision. Our past light cone at last scattering is represented by the dashed circle. From the present bounds on curvature, the size of our past light cone must be much smaller than 1 curvature radius. Zooming in on the portion of the surface of last scattering that we have causal access to (inset), the Universe is very close to flat, and the region affected by the collision has approximate planar symmetry. The region affected by the collision appears as a disk of angular radius θ_{crit} on the CMB sky.

collision appears as a disk on the observer's CMB sky. Zooming in on the neighborhood of our past light cone (inset), we can treat the Universe as being flat. In addition, because we have causal access to much less than one curvature radius at last scattering (again from the observational bound on Ω_k), the collision has an approximate planar symmetry.

The collision introduces preinflationary inhomogeneities into our bubble. The exact nature of these inhomogeneities depends on the specific model underlying the formation of our bubble and the subsequent epoch of slow-roll inflation, as well as the specifics of the collision. In dramatic cases, the collision ends slow-roll inflation everywhere within its future light cone [18], induces the transition to another vacuum state [23,40,41], or produces a postcollision domain wall that eats into our bubble interior [17,19]. These scenarios are obviously in conflict with observation, and we do not consider them further. In mild cases, which will be our focus in the remainder of this paper, collisions satisfy the ‘‘compatibility’’ criterion defined above: The observable portion of the surface of last scattering is only minimally disturbed by the collision. Thin-wall analysis [17] and numerical simulations [18,20] indicate that it is indeed possible to find situations where the effects of a collision are compatible with our observed cosmology.

The disturbance caused by a collision is a preinflationary relic and thus is stretched by the period of inflation inside the bubble. From the current bound on curvature [36], we can infer that our past light cone encompasses less than one horizon volume at the onset of inflation. This implies that the initial disturbances caused by a collision, which is smeared out on the scale of the inflationary horizon after a few e -folds of inflation, has a wavelength today that is larger than the current horizon size. Together with the planar symmetry of the collision at last scattering (by convention along the y - z plane), this implies that we can Taylor expand the Newtonian potential (see Ref. [26] for a translation between the Newtonian potential and the originally postulated temperature modulation presented in Ref. [20]) about the causal boundary of the collision at $x = x_{\text{crit}}$ as

$$\begin{aligned} \Phi_{\text{coll}} = & \Phi(a)(\bar{c}_0 + \bar{c}_1(x - x_{\text{crit}}) \\ & + \mathcal{O}((x - x_{\text{crit}})^2))\Theta(x - x_{\text{crit}}), \end{aligned} \quad (6)$$

where $\Phi(a)$ encodes the evolution of the potential with scale factor a and the \bar{c}_i are model-dependent constants.²

²We are modeling the collision as a collection of supermodes truncated at the causal boundary, and our treatment is therefore very similar to the so-called ‘‘tilted universe’’ scenario [42,43]. The important distinction in the case of bubble collisions is that the perturbation vanishes at the causal boundary x_{crit} . Because the collision entered our past light cone only relatively recently, we are still comoving with respect to the undisturbed FRW foliation, and the cancellation of the dipolar temperature modulation seen in Refs. [42–44] does not occur.

There are contributions to the observed temperature modulation from the Sachs-Wolfe effect, the integrated Sachs-Wolfe effect, and a Doppler effect (coming from the induced bulk peculiar velocity \mathbf{v} of the fluid in the region affected by the collision):

$$\frac{\delta T}{T} \simeq \frac{\Phi_{\text{coll}}(a_{\text{ls}})}{3} + 2 \int_{a_{\text{ls}}}^1 da \frac{d\Phi_{\text{coll}}}{da} + (\mathbf{v} \cdot \hat{\mathbf{n}} + \mathcal{O}(v^2)), \quad (7)$$

where a_{ls} is the scale factor at last scattering, $a = 1$ today, and

$$\mathbf{v} \propto \nabla \Phi_{\text{coll}} + a \frac{d}{da} \nabla \Phi_{\text{coll}}. \quad (8)$$

To leading order in \mathbf{v} , the temperature induced by the collision is linear in Φ_{coll} and its derivatives. Therefore, since $x = x_{\text{ls}} \cos\theta$ (where x_{ls} is the comoving distance out to which we can see on the surface of last scattering), the temperature fluctuations induced by a collision are generally of the form Eq. (4). Further, even if the Newtonian potential is continuous across $x = x_{\text{crit}}$, the resulting temperature fluctuations need not be continuous across the causal boundary at θ_{crit} . This discontinuity arises from the integrated Sachs-Wolfe and Doppler contributions to the observed temperature fluctuation. Effects that we have neglected, including the finite thickness of the surface of last scattering and uncertainties about how the perturbations caused by a bubble collision propagate through our bubble interior, are encapsulated by the higher order terms in Eq. (4). These effects could smear out the causal boundary enclosing the collision on subdegree scales. These corrections could be incorporated into our analysis as theoretical understanding improves.

Given a specific model for the scalar fields making up the bubbles and driving eternal inflation, the kinematics of a particular collision, and our position inside our bubble, it is in principle possible to determine the free parameters in Eq. (4). Treating the colliding bubbles in the thin-wall approximation, some measure of the strength of a collision can be specified in terms of the vacuum energies in the bubbles, wall tensions, and kinematics as in Ref. [17]. The kinematics will induce a probability distribution for the free parameters in Eq. (4). However, an accurate treatment requires a calculation of the backreaction of the collision on the behavior of the inflaton inside our bubble. Preliminary work in this direction has been done [18,20,23,40], providing a handful of examples. However, a systematic investigation has not yet been performed. This is distinct from the case where an ensemble of field theory models is considered, representing the string theory landscape. In this case, the fundamental parameters governing the structure of the colliding bubbles (wall tensions and vacuum energies) and the properties of the inner-bubble cosmology (including the number of inflationary e -folds, etc.) are drawn from some probability distribution. This




FIG. 3 (color online). On the left, we show a bubble-collision template with $\{z_0 = 5.0 \times 10^{-5}, z_{\text{crit}} = -5.0 \times 10^{-5}, \theta_{\text{crit}} = 10.0^\circ, \theta_0 = 57.7^\circ, \phi_0 = 99.2^\circ\}$. On the right we add simulated background fluctuations, smoothing, and instrumental noise.

again will induce a probability distribution for the free parameters in Eq. (4), whose nature is presently poorly understood.

What would a bubble collision embedded in a CMB temperature map look like? In Fig. 3 we show a large-amplitude collision with and without background CMB fluctuations. In the following sections, we apply the various stages of our analysis pipeline to this example to illustrate the algorithm. We make extensive use of such simulations in calibrating our analysis pipeline, and the details of their construction are presented in Sec. IV. Although there could conceivably be many overlapping collisions, the predicted observational signatures of this scenario have yet to be explored, and we focus on simulations of distinct individual bubble collisions. Again, as theoretical understanding improves, our analysis could be extended to include the possibility of overlapping collisions.

What would the detection, or absence, of a bubble collision tell us about the underlying theory of eternal inflation? To examine what the answer to this question might be, let us make some further assumptions about the temperature modulations caused by a bubble collision. First, assume that the potential induced by the collision [Eq. (6)] is composed mostly of a single long-wavelength mode of physical wave number k . Second, assume that the Sachs-Wolfe effect is the dominant contribution to the observed temperature modulation. Under these assumptions, the amplitude of an observed temperature modulation is

$$z_0 \simeq \frac{2}{3} \frac{k}{H_0} \Phi(a_{\text{ls}}) (1 - \cos\theta_{\text{crit}}), \quad (9)$$

where a_{ls} is the scale factor at last scattering. If the initial wavelength of the disturbance was of the order of 1 inflationary Hubble length $k \sim H_I$ (since any fine structure in the collision would be smeared within the first few e -folds of inflation), then $\Phi(a_{\text{ls}}) = \Phi(a = 0)$, and the physical size of such a mode at last scattering is given by

$$k \simeq \Omega_k^{1/2} H_0. \quad (10)$$

In this case, we have

$$z_0 \simeq \Phi(0) \Omega_k^{1/2} (1 - \cos\theta_{\text{crit}}). \quad (11)$$

If a bubble collision is detected, and a similar set of assumptions is proven correct in a specific model, the measured values of z_0 and θ_{crit} allow one to infer the value of Ω_k .³ In the absence of a detected collision, Eq. (11) can be turned into a bound on a combination of $\Omega_k^{1/2}$ and $\Phi(0)$:

$$\Omega_k^{1/2} \Phi(0) < [z_0 / (1 - \cos\theta_{\text{crit}})]_{\text{observational upper bound}}. \quad (12)$$

This analysis can be recognized as an example of the Grishchuk-Zel'dovich effect in an open universe [42,45].

Determining the detailed properties of the theory underlying eternal inflation through the observation of bubble collisions is likely to be a messy business. However, any model will predict an expectation value for the number of observable bubble collisions \bar{N} , making this a very useful phenomenological parameter. Any constraints on \bar{N} from data will also yield interesting information about our parent vacuum through Eq. (1). The most naive application of such a constraint, where we have evidence that $\bar{N} > 1$ or $\bar{N} < 1$, would yield the inequalities

$$\begin{aligned} \lambda H_F^{-2} &< \frac{3H_I^2}{16\pi\Omega_k^{1/2}} \quad (\text{no detected collisions}), \\ \lambda H_F^{-2} &> \frac{3H_I^2}{16\pi\Omega_k^{1/2}} \quad (\text{collision detection}). \end{aligned} \quad (13)$$

These bounds would be most useful if we detect Ω_k and/or B -mode polarization (the amplitude of which can be

³The values of Ω_k that one might be able to infer are conceivably below both the observational bound $\Omega_k \leq 0.0084$ and the theoretical observational bound $\Omega_k \leq 10^{-5}$. For example, assuming $z_0 \sim 10^{-5}$ and $\Phi(0) \sim 1$ (since the collision involves a relatively large release of energy), if a collision were observed at large angular scale (where $\cos\theta_{\text{crit}} \sim 0$), we can infer that $\Omega_k \sim 10^{-10}$. This implies that a collision is in principle observable even when a curvature is not. We thank Lam Hui for elucidating this point.

related to H_I) in future data. In the most optimistic scenario,⁴ if primordial B -mode polarization is detected by the Planck satellite, we can infer that $H_I \sim 10^{11} - 10^{13}$ GeV. Further, if curvature is detected at the level $\Omega_k \sim 10^{-3}$, then in Eq. (13), λH_F^{-2} would be bounded from above or below by $\sim 10^{26}$ GeV². The condition for eternal inflation is $\lambda H_F^{-4} < 1$. Any application of Eq. (13) must be consistent with this inequality. For example, assuming a Planckian false vacuum energy ($H_F \sim 10^{19}$ GeV), the nucleation probability λH_F^{-4} could be bounded from above or below by $\sim 10^{-4}$, remaining consistent with the condition for eternal inflation.

III. SUMMARY OF THE ANALYSIS PIPELINE

Before providing a detailed description of our analysis pipeline, we motivate and summarize its various components. Eternal inflation can arise from a wide range of inflationary potentials, each producing a different expected number of detectable collisions on the CMB sky, \bar{N}_s . We will therefore use \bar{N}_s as a continuous parameter that characterizes particular models of eternal inflation. The standard cosmological model is given by the special case in which $\bar{N}_s = 0$. Our primary goal is to determine, given the WMAP 7-year data, what constraints can be placed on \bar{N}_s and whether models predicting $\bar{N}_s > 0$ should be preferred over models predicting $\bar{N}_s = 0$.

The optimal approach to achieving this goal would be to construct the full posterior for \bar{N}_s from Bayes's theorem given full-resolution CMB data on the whole sky. Unfortunately, this would require inverting the full-sky full-resolution CMB covariance matrix as well as integrating the bubble-collision likelihood over a many-dimensional parameter space. These tasks are computationally intractable. However, taking advantage of the fact that bubble collisions produce discrete localized effects on the CMB sky, it is possible to approximate the full-sky Bayesian analysis by a patchwise analysis if the most promising candidate signatures can be identified in advance. The implementation of such an approximation scheme requires two assumptions. First, we assume that the likelihood of models predicting $\bar{N}_s > 0$ is peaked in the regions of the sky containing the candidate collisions, and that the integral over the likelihood can therefore be estimated by concentrating on these regions, which make the largest contribution to the full integral. Second, we assume that these regions are separated widely enough to be un-

correlated with each other, so much smaller local covariance matrices can be used. These assumptions allow the results of a small number of localized (and therefore computationally feasible) Bayesian model selection tests to be combined into estimates of the required full-sky statistics. Put simply, our algorithm implements a conservative approximation to the required numerical integral. A complete treatment of the full-sky analysis and the assumptions on which it is formed can be found in the Appendix. In addition, once a set of candidates have been identified, it is possible to apply further tests of the data in parallel.

The full-sky approximation necessitates the development of an algorithm that identifies the most promising regions of the CMB sky and then processes them individually. Upon segmenting the full data set, it is important to avoid biasing oneself with *a posteriori* selection effects [34], and it is therefore critical to minimize human intervention in choosing what portions of the sky to analyze. Thus our analysis pipeline is fully automated, tested, and calibrated on realistic simulations of the data and frozen before being applied to the real data. The final pipeline contains no algorithmic choices tunable via human intervention. As discussed in the Appendix, missing a bubble-collision candidate which makes a significant contribution to the full integral leads to a *conservative* bias toward models predicting $\bar{N}_s = 0$. This alleviates the worry that selection effects might lead to a spurious detection.

Our analysis pipeline consists of a candidate identification step, followed by two parallel verification procedures:

- (i) Blob detection: To begin, we attempt to locate the most promising candidate signals using wavelets. Wavelet analysis is a compromise between working purely in position or harmonic space, and therefore yields information both about the location and angular scale of particular features in the temperature map. Specifically, we employ standard [46–50] and Mexican [51] spherical needlets, two classes of wavelets defined on the sphere. The statistics of the needlet representation of a purely Gaussian CMB temperature map (expected in the absence of a bubble collision at large scales where WMAP is cosmic variance dominated), combined with simulations of a bubble-free masked CMB sky, can be used to quantify the significance of various features. A set of significance thresholds are then defined to ensure a manageable number of “false detections” in the end-to-end simulation of the WMAP experiment (see Sec. IV A). Regions of the sky passing these thresholds are sewn into “blobs,” whose size and location is determined by the needlet responses and passed onto the next stages of the pipeline.
- (ii) Edge detection: Once a set of candidate signatures is found, we look for circular edges across which the temperature is discontinuous. As discussed above, such causal edges are expected to be a generic

⁴In models where a collision is expected to be in our past, there might be a good reason to expect a correlation between observed B modes and the observation of a bubble collision [18]. This is because large-field models of inflation, which generically predict a larger value for the tensor to scalar ratio, are much more robust in the presence of a bubble collision. In models of small-field inflation, a bubble collision can end inflation everywhere in its future light cone, implying that collisions in such models are not compatible with our observed cosmology.

signature of bubble collisions. We use the Canny algorithm [52], adding an adaptation of the circular Hough transform (CHT) [53], to focus our search on circular edges. The algorithm consists of identifying the most likely center for a noisy circular edge. The significance of this response is calibrated from a detailed analysis of bubble-collision simulations including cosmic variance, spatially varying WMAP instrumental noise, and smoothing due to the instrumental beam. We verify that this step produces no false detections in the WMAP end-to-end simulation.

- (iii) Bayesian parameter estimation and model selection: The regions highlighted by the blob-detection step can be used to construct an approximation to the full-sky posterior probability distribution for \bar{N}_s using the methods outlined in the Appendix. We first perform a pixel-based evaluation of the likelihood and Bayesian evidence in each blob for bubble-collision templates of the form given in Eq. (4), sampling the parameter space using the nested sampler MULTINEST [54]. The likelihood analysis includes cosmic variance, spatially varying WMAP instrumental noise, and the smoothing due to the instrumental beam. Combining the evidences from each blob we obtain the posterior probability distribution for \bar{N}_s , which is used to derive constraints on \bar{N}_s and perform model selection to determine if a theory with $\bar{N}_s \neq 0$ is preferred over a theory with no predicted collisions. The significance of a detection is again calibrated from an analysis of simulated collisions and an end-to-end collision-free simulation of the experiment.

The most important output of our pipeline is the approximation to the full-sky posterior probability distribution for \bar{N}_s . This allows us to derive marginalized constraints on \bar{N}_s and perform model selection between theories with $\bar{N}_s = 0$ and $\bar{N}_s \neq 0$. In addition, for each blob identified by the first set of the pipeline, we obtain a set of marginalized posterior constraints on the model parameters $\{z_0, z_{\text{crit}}, \theta_{\text{crit}}, \theta_0, \phi_0\}$, a maximum needlet significance, CHT score, and a local Bayesian evidence ratio with respect to the no-bubble-collision model.

IV. SIMULATIONS

Our analysis pipeline is general, but each step must be calibrated using simulations of the particular data set under consideration, in this case the WMAP 7-year data release [32]. WMAP has measured the intensity and polarization of the microwave sky in five frequency bands. The resolution of the instrument in each band is limited by the detectors' beams and is highest at 0.22° in the 94 GHz W band. We perform our analysis on the foreground-subtracted W-band WMAP temperature map, as this combines the highest resolution full-sky data currently available with the least

foreground contamination. To minimize the effects of the residual foregrounds we cut the sky with the conservative KQ75 mask, leaving 70.6% of the sky unmasked.

We carry out extensive simulations to quantify the thresholds at which areas of the sky are passed from one step to another. To find the best approximation to the full-sky Bayesian analysis, we process as much of the sky as is computationally feasible.

To determine the response of our pipeline to bubble collisions over the range of possible parameters, we generate simulations containing a variety of bubble collisions plus CMB, realistic noise and Gaussian beam smoothing. However, we also wish to ensure that we have a method to guard against systematic effects (e.g., foreground residuals and any map-making artifacts that may be present) that we do not have capability to simulate. These effects might lead to false detections in the ‘‘blob-detection’’ stage, or critically, the edge-detection and Bayesian analysis stages. It is impossible to claim a detection without first ensuring that there are no such false detections due to systematics.

A. WMAP end-to-end simulation

A realistic simulation of a WMAP-quality data set that does not contain a bubble collision is an important tool for calibrating and quantifying the expected false detection rate of our analysis pipeline when applied to data which may include systematics (such as foreground residuals) that are not captured in our simulations or likelihood function. For this purpose we use a complete end-to-end simulation of the WMAP experiment provided by the WMAP science team.⁵ The temperature maps in this simulation are produced from a simulated time-ordered data stream, which is processed using the same algorithm as the actual data. The data for each frequency band are obtained separately from simulated sources including diffuse Galactic foregrounds, CMB fluctuations, realistic noise, smearing from finite integration time, finite beam size, and other instrumental effects. In our analysis, we utilize the foreground-reduced W-band simulation.

B. Simulated bubble collisions

The temperature fluctuations observed in the CMB, including the effects of a bubble collision (originally found in Refs. [20,39]), can be written as

$$\begin{aligned} \delta T(\hat{\mathbf{n}}) = & [T'_0(1 + f(\hat{\mathbf{n}}))(1 + \delta(\hat{\mathbf{n}})) - T_0]_{\text{smoothed}} \\ & + \delta T_{\text{noise}}(\hat{\mathbf{n}}), \end{aligned} \quad (14)$$

where $\hat{\mathbf{n}} = \{\theta, \phi\}$ is the position on the sky,⁶ T_0 is the average temperature of the map including the modulation,

⁵http://lambda.gsfc.nasa.gov/product/map/dr4/sim_maps_info.cfm.

⁶These angular positions can be expressed in terms of Galactic coordinates through longitude $l = \phi$ and latitude $b = 90^\circ - \theta$.

T'_0 is the average temperature without the modulation, δT_{noise} is the contribution from instrumental noise, and $f(\hat{\mathbf{n}})$ and $\delta(\hat{\mathbf{n}})$ are defined as in Eq. (3). The quantities in the brackets are smoothed with a Gaussian beam of 0.22° (approximating the beam size of the WMAP experiment in the W band). We use the WMAP best-fit 7-year power spectrum [55] in the multipole range $2 \leq \ell \leq 1024$ to generate fluctuation maps $\delta T_{\text{syn}}(\hat{\mathbf{n}}) = T'_0 \delta(\hat{\mathbf{n}})$ at the full WMAP resolution of $N_{\text{side}} = 512$ (with 3 145 728 pixels). The noise term δT_{noise} is generated from WMAP 7-year noise variances at the same resolution. Since the templates we consider add a relatively small temperature excess/deficit in one location, the features do not cause the power spectrum to deviate from that measured by WMAP [20]. Additionally, we can replace $T_0 \approx T'_0$, which gives

$$\delta T(\hat{\mathbf{n}}) = [(1 + f(\hat{\mathbf{n}}))(T_0 + \delta T_{\text{syn}}(\hat{\mathbf{n}})) - T_0]_{\text{smoothed}} + \delta T_{\text{noise}}(\hat{\mathbf{n}}). \quad (15)$$

We consider collisions with $\theta_{\text{crit}} = 5^\circ, 10^\circ, 25^\circ$ and choose centers in a high-noise region ($\theta_0 = 56.6^\circ, \phi_0 = 193.0^\circ$) and a low-noise region ($\theta_0 = 57.7^\circ, \phi_0 = 99.2^\circ$) of the sky that remain significantly outside of the main body of the WMAP KQ75 7-year mask. The regions of the sky affected by 5° and 10° collisions are overplotted in Fig. 4 on a masked map of the instrumental noise variance. For each θ_{crit} and location, we generate 35 simulated collisions with parameter values logarithmically spaced in the ranges $10^{-6} \leq z_{\text{crit}} \leq 10^{-4}$ and $-10^{-4} \leq z_0 \leq -10^{-6}$. The response of our pipeline depends only on the absolute value of z_0 and z_{crit} , so the choice of sign for z_0 and z_{crit} is arbitrary. We repeat this for three realizations of the background CMB fluctuations, yielding a total of 210 simulated sky maps for each of the three collision sizes.




FIG. 4 (color online). The two locations chosen for our simulated bubble collisions are overplotted on the WMAP 7-year noise variances with the KQ75 7-year mask applied. The regions encompassed by the 5° and 10° simulated collisions are shaded black and gray, respectively. Bubbles are centered in an unmasked high-noise $\{\theta_0 = 56.6^\circ, \phi_0 = 193.0^\circ\}$ (left) and low-noise $\{\theta_0 = 57.7^\circ, \phi_0 = 99.2^\circ\}$ region (right).

V. ANALYSIS TOOLS

We now describe in detail the analysis tools which make up our pipeline and how they are calibrated with simulations before being applied to the data. Readers wanting to skip these details may want to study Figs. 11 and 17 and turn to Sec. VD for a summary of the outputs of the pipeline at each stage and the conditions under which a detection can be claimed.

A. Needlets

Wavelet analysis is a powerful tool for identifying features localized on the sky, the type of signal expected for bubble collisions. There exist families of wavelets that are defined on the sphere known as standard [46–50] and Mexican [51] spherical needlets. These functions form what is known as a “tight frame,” allowing for a well-defined forward and reverse needlet transform. As in other forms of wavelet analysis, decomposing the temperature map into a sum over such functions yields information both about the location and angular scale of specific features. For a purely Gaussian temperature field, the statistical properties of the needlet transform can be straightforwardly related to the power spectrum, allowing a quantification of the significance of a possible detection. In addition, the spatial localization properties of the standard and Mexican needlets make it possible to avoid many of the problems associated with working on a cut sky. In this section we outline the properties of needlet transforms and analyze their utility in searching for bubble collisions.

1. Definition of the spherical needlet transform

The needlet transform is defined as

$$T(\hat{\mathbf{n}}) = \sum_{j,k} \beta_{jk} \psi_{jk}(\hat{\mathbf{n}}), \quad (16)$$

where $\hat{\mathbf{n}}$ denotes a direction $\{\theta, \phi\}$ on the sky, β_{jk} are constant needlet coefficients, and $\psi_{jk}(\hat{\mathbf{n}})$ are the needlet functions. The members of this family of functions are labeled by the index k of the pixel at which they are centered, and their “frequency” j , which is related to the spatial extent of the needlet profile in real space. The sum in the needlet transform is over all pixels k , and all frequencies $j = 0, 1, 2, \dots, \infty$. For fixed j , there is one needlet coefficient β_{jk} for each pixel k , allowing us to represent the needlet coefficients at fixed j as a map on the pixelated sky. The needlet functions are defined in terms of spherical harmonics $Y_{\ell m}(\hat{\mathbf{n}})$ as

$$\psi_{jk}(\hat{\mathbf{n}}) = \sqrt{\lambda_{jk}} \sum_{\ell} b(\ell, B, j) \sum_{m=-\ell}^{\ell} Y_{\ell m}^*(\hat{\mathbf{n}}) Y_{\ell m}(\hat{\mathbf{n}}_k). \quad (17)$$

Here, λ_{jk} are the cubature weights, which are related to the area of each pixel. In the equal-area HEALPix pixelization [56] we employ, all cubature weights are equal to

$\lambda_{jk} = 4\pi/N_{\text{pix}}$, where N_{pix} is the number of pixels, and we absorb this constant into the needlet coefficients β_{jk} . The function $b(\ell, B, j)$ acts as a filter in harmonic space, where B is a constant bandwidth parameter. It is chosen such that the family of functions $\psi_{jk}(\hat{\mathbf{n}})$ forms a tight frame (see, e.g., Ref. [46]), which guarantees the existence of an inverse needlet transform given by

$$\beta_{jk} = \int T(\hat{\mathbf{n}}) \psi_{jk}(\hat{\mathbf{n}}) d\Omega. \quad (18)$$

There are a number of possible choices for the function $b(\ell, B, j)$, which distinguish the standard and Mexican needlets. A description of the explicit form of the function $b(\ell, B, j)$ can be found for standard needlets in Ref. [46] and Mexican needlets in Ref. [51]. We plot b as a function of the multipole moments ℓ in Fig. 5. For standard needlets, b only has support for values of ℓ between $B^{j-1} < \ell < B^{j+1}$. The bandwidth parameter B controls the width of each window function in harmonic space. Mexican needlets have support over all ℓ at each frequency j , and again have a bandwidth parameter B which controls the localization properties of the functions in harmonic space.

In Fig. 6 we plot the wavelet functions in pixel space. As is to be expected, increasing the width of the function $b(\ell, B, j)$ in harmonic space corresponds to improved localization in pixel space. In the limit of large j , there is an extremely small overlap of the needlet functions at nearby

pixels. The compact support of $b(\ell, B, j)$ in harmonic space for standard needlets leads to slightly poorer localization in pixel space than is enjoyed by the Mexican needlets.

If we decompose the temperature map into spherical harmonics:

$$T(\hat{\mathbf{n}}) = \sum_{\ell, m} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \quad (19)$$

then Eq. (18), together with the inverse transform

$$a_{\ell m} = \int T(\hat{\mathbf{n}}) Y_{\ell m}^* d\Omega \quad (20)$$

leads to

$$\beta_{jk} = \sqrt{\lambda_{jk}} \sum_{\ell} b(\ell, B, j) \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}_k), \quad (21)$$

where $a_{\ell m}$ are the spherical harmonic coefficients. This formula allows us to transform directly from the $a_{\ell m}$'s to the spherical needlet coefficients β_{jk} . In our analysis, the needlet transforms are accelerated by generating the $a_{\ell m}$'s at full WMAP resolution ($N_{\text{side}} = 512$) but limiting the reconstruction multipoles to $2 \leq \ell \leq 256$, and needlet positions k to the pixels at $N_{\text{side}} = 128$. This retains the resolution required to reconstruct features from half-sky to half-degree scales, encompassing the range of all detectable collisions.




FIG. 5. The filter function $b(\ell, B, j)$ for standard (left panel) and Mexican (right panel) needlets with $B = 2.5$ for $j = 0, 1, 2, 3$ (solid, dashed, dot-dashed, dotted curves).




FIG. 6 (color online). Standard needlets in pixel space. On the left, we show standard needlets ψ_{jk} with $B = 2.5$ for $j = 1, 2, 3$ (dot-dashed, dashed, solid curves) at fixed k as a function of the polar angle θ . On the right, we show standard needlets ψ_{jk} for fixed $j = 3$ at three pixels k (dashed, solid, dot-dashed curves) as a function of the polar angle θ (note: since we are projecting onto θ , the needlets appear asymmetric).

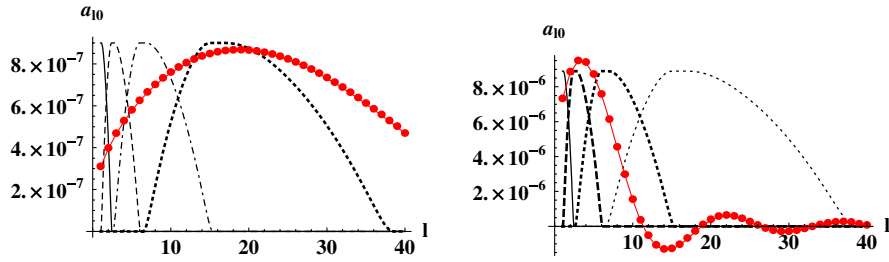


FIG. 7 (color online). The spherical harmonic transform (connected dots) of a $\theta_{\text{crit}} = 5^\circ$ (left panel) and $\theta_{\text{crit}} = 25^\circ$ (right panel) bare collision template centered on the north pole on top of the filter function $b(\ell, B, j)$ for standard needlets with $B = 2.5$ for $j = 0, 1, 2, 3$ (solid, dashed, dot-dashed, dotted curves). The overlap of the filter function $b(\ell, B, j)$ with the spherical harmonic transform of the bubble template [see Eq. (21)] determines for which value of j the needlet transform yields the largest signal. In these examples, the 5° collision has the largest needlet response at $j = 3$ and the 25° collision at $j = 2$. The needlet response as a function of angle for the 25° collision is plotted in Fig. 8.

2. Needlet response to the bubble-collision templates

We now quantify the sensitivity of the needlet transform to the presence of a collision. In the absence of the background Gaussian fluctuations, we perform the needlet transform on a set of bare collision templates. As an illustration, in Fig. 7 we plot the spherical harmonic coefficients as a function of ℓ (all coefficients for $m \neq 0$ vanish by symmetry if we center the template on the north pole) for 25° and 5° collision templates, overlaid on top of the rescaled filter function $b(\ell, B, j)$ for spherical needlets. The spherical harmonic coefficients for the collision templates peak at a value of ℓ related to the angular scale of the causal boundary. Therefore, the needlet coefficients are largest at a frequency j that is directly related to the angular scale of the collision. This can be seen in Fig. 7, where the 5° collision has a maximum response at $j = 3$ and the 25° collision has a maximum response at $j = 2$.

In Fig. 8, we plot the needlet coefficients for the 25° template at a variety of polar angles, for $0 \leq j \leq 3$. The needlet coefficients are largest in the center of the region affected by the collision (here, the template is centered on $\theta = 0$), and at a frequency j correlated with the angular scale of the collision ($j = 2$). As expected, the needlet response is sensitive to both the location and angular scale of the collision.




FIG. 8 (color online). $B = 2.5$ for standard needlets and a 25° collision. $j = 0$ (circles), $j = 1$ (squares), $j = 2$ (diamonds), and $j = 3$ (triangles).

By studying the needlet response to a variety of bare collision templates given by Eq. (4), we can find optimal values of the bandwidth parameter B for each needlet type. Larger-bandwidth needlets produce stronger signals, but also respond to a greater range of bubble sizes. We have found that the values $B = 1.8$ and $B = 2.5$ for the standard and $B = 1.4$ for the Mexican needlets are a good compromise between signal strength and angular localization of response. In our analysis, we use this suite of three needlet transforms, ensuring that we are sensitive to temperature modulations with a variety of profiles and allowing us to cross-check any candidate signals.

As an important step in our analysis pipeline, we build lookup tables containing the possible range of bubble-collision scales, $\theta_{\text{crit,min}} \leq \theta_{\text{crit}} \leq \theta_{\text{crit,max}}$, to which each needlet type and frequency is sensitive. We first generate a set of 100 templates at each integer θ_{crit} between 1° and 89° by randomly sampling z_{crit} between $-5 \times 10^{-5} \leq z_{\text{crit}} \leq 5 \times 10^{-5}$ with $z_0 = 5 \times 10^{-5}$ for $z_{\text{crit}} > 0$, and $z_0 = z_{\text{crit}} + 5 \times 10^{-5}$ for $z_{\text{crit}} < 0$. This creates a set of templates with uniform total amplitude [[i.e., constant $f(\theta, \phi) - f(\theta_{\text{crit}}, \phi)$] but a variety of profiles at each angular scale. Next, we calculate the needlet coefficients for each of the three members of our needlet suite and record the frequency generating the maximum central needlet response for each template. The range in θ_{crit} recorded at each frequency is used to generate the lookup tables in Table I.

3. Needlet coefficients on a cut sky

The CMB is completely dominated by foreground emission in the region of the Galactic plane and is also affected by bright point sources. These issues are typically handled by applying a mask which covers the Galactic plane and known point sources. The needlet transform can be applied directly on the masked temperature maps, and because the needlet functions are localized in pixel space, needlet coefficients far from the mask for sufficiently high-frequency j are not significantly affected.

TABLE I. Angular scale lookup tables for standard needlets with $B = 2.5$ (top), standard needlets with $B = 1.8$ (middle), and Mexican needlets with $B = 1.4$ (bottom). For a needlet frequency j , the needlet transform is sensitive to bubble collisions on scales $\theta_{\text{crit,min}} \leq \theta_{\text{crit}} \leq \theta_{\text{crit,max}}$. No results are shown for the standard needlets with $B = 1.8$, $j = 0$ as they have no support over the range of angular scales considered.

j	$\theta_{\text{crit,min}}$	$\theta_{\text{crit,max}}$
0	60°	90°
1	33°	71°
2	12°	32°
3	5°	14°
4	2°	5°
5	1°	2°

j	$\theta_{\text{crit,min}}$	$\theta_{\text{crit,max}}$
1	56°	90°
2	28°	64°
3	17°	38°
4	10°	21°
5	6°	12°
6	3°	7°
7	2°	4°
8	1°	2°

j	$\theta_{\text{crit,min}}$	$\theta_{\text{crit,max}}$
0	86°	90°
1	78°	90°
2	55°	90°
3	36°	71°
4	27°	48°
5	19°	37°
6	14°	27°
7	10°	20°
8	8°	16°
9	5°	12°
10	4°	8°
11	3°	6°
12	2°	4°
13	1°	2°

These high-frequency needlets are mainly composed of high- ℓ spherical harmonics, and so cut-sky $a_{\ell m}$'s can safely be used to calculate the needlet coefficients through Eq. (21). Unfortunately, the low-frequency needlets are quite sensitive to the presence of the mask. To partially mitigate this sensitivity, we calculate the optimal unbiased maximum-likelihood estimators of the $a_{\ell m}$'s [57] at low ℓ .

Such maximum-likelihood estimators are computationally very expensive, and we must balance accuracy against limited computational resources. Another minor complication arises from the smoothing that is necessary to band limit the data when performing the maximum-likelihood reconstruction algorithm of Ref. [57]: information leaks

from inside the mask. Comparing the reconstruction on masked and unmasked simulated temperature maps using 10° -FWHM Gaussian smoothing, we have determined that a reasonably small bias is obtained when maximum-likelihood $a_{\ell m}$'s are used for $\ell < 10$.

This set of hybrid $a_{\ell m}$'s—maximum-likelihood reconstructed $a_{\ell m}$'s for $\ell \leq 10$ and cut-sky $a_{\ell m}$'s for $\ell > 10$ —is used in Eq. (21) to calculate the needlet coefficients in the analysis that follows.

4. Statistical properties of needlet coefficients

For a Gaussian CMB without sky cuts, the statistical properties of the spherical harmonic coefficients [46] are

$$\langle a_{\ell m} \rangle = 0, \quad \langle |a_{\ell m}|^2 \rangle = C_\ell. \quad (22)$$

These are related to the statistical properties of the β_{jk} in a straightforward way by

$$\langle \beta_{jk} \rangle = 0, \quad \langle \beta_{jk}^2 \rangle = \sum_{\ell} b(\ell, B, j) \frac{2\ell + 1}{4\pi} C_\ell, \quad (23)$$

which are identical at each pixel k . Thus, in a full-sky analysis, comparison with the Gaussian variance yields a measure of how likely it is to find a particular needlet coefficient in a purely Gaussian realization of the CMB sky.

In the presence of foregrounds, however, it is necessary to work on a cut sky, which introduces a j - and a k -dependent bias. Following Ref. [48], we determine the significance of a needlet coefficient on the cut sky by⁷

$$S_{jk} = \frac{|\beta_{jk} - \langle \beta_{jk} \rangle_{\text{Gauss,cut}}|}{\sqrt{\langle \beta_{jk}^2 \rangle_{\text{Gauss,cut}}}}, \quad (25)$$

where the average, $\langle \beta_{jk} \rangle_{\text{Gauss,cut}}$, and variance, $\langle \beta_{jk}^2 \rangle_{\text{Gauss,cut}}$, are calculated at each pixel from the needlet coefficients of 3000 collision-free Gaussian CMB realizations with the WMAP 7-year KQ75 sky cut applied. Simulating only cosmic variance is sufficient here because the measurements made by WMAP are cosmic variance limited on the scales of interest, $\theta_{\text{crit}} \gtrsim 5^\circ$.

Maps of the needlet variances obtained from simulations are shown in Fig. 9 for examples with low (top row) and high needlet (bottom row) frequencies. On the left are the

⁷One can also define composite significances, involving needlet coefficients at multiple frequencies. An example is

$$S_{jj'} = \frac{|\beta_{jk}\beta_{j'k} - \langle \beta_{jk}\beta_{j'k} \rangle|}{\sqrt{\langle (\beta_{jk}\beta_{j'k})^2 \rangle}}. \quad (24)$$

We have evaluated this statistic on a variety of collision templates modulating Gaussian realizations of the background CMB fluctuations and found that it returns about half the significance given by Eq. (25).

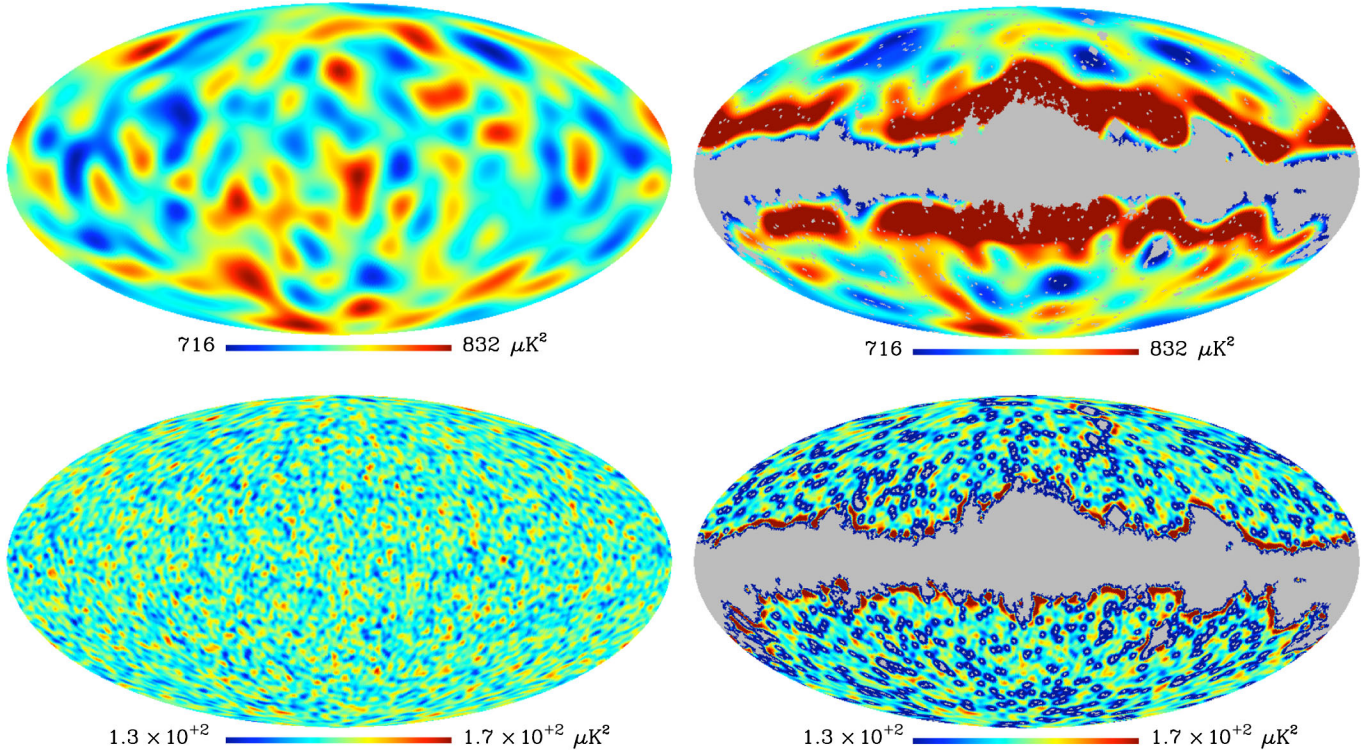


FIG. 9 (color online). Needlelet-coefficient variance maps for standard needlelets with $B = 2.5$, $j = 2$ (top row) and Mexican needlelets with $B = 1.4$, $j = 11$ (bottom row), generated using an ensemble of 3000 Gaussian CMB realizations. The map-averaged variances of the full-sky maps (left column) agree well with expectations from theory [Eq. (23) predicts $774 \mu\text{K}^2$ and $150 \mu\text{K}^2$ respectively], as do regions of sky sufficiently distant from the 7-year KQ75 mask (right column). The low-frequency needlelets are affected predominantly by the Galactic cut, whereas the high-frequency needlelets are affected by point-source masking.

needlelet variances calculated without a mask, which agree at the 5% level with the expected variances from Eq. (23). On the right are the masked needlelet variances, which are clearly biased within a certain distance from the mask in both cases. Variances in the low-frequency example are affected predominantly by the Galactic cut. Variances in the high-frequency example are affected in a much smaller region of the Galactic cut (reflecting the increased spatial localization of needlelets at high frequencies), but are much more significantly affected by the point-source masks.

To illustrate the expected response to a bubble collision, in Fig. 10 we show the temperature map of our illustrative example of a simulated bubble collision on the cut sky, and the significances of its needlelet coefficients calculated from Eq. (25) at $j = 3$ using standard needlelets with $B = 2.5$. The location of the collision is clearly highlighted in the map of needlelet coefficients. The significance of the needlelet coefficients in pixels in the center of the collision forms a global maximum on the entire map.

In order to identify a set of needlelet coefficients with a particular feature, we sew regions with 5 or more pixels whose needlelet coefficients exceed a frequency-dependent threshold into blobs (we discuss in more detail below how these thresholds are set). The core of each blob contains all adjacent pixels that pass the significance threshold. This

core is then extended by first finding the average position $\hat{\mathbf{n}}_0$ of the pixels in a blob and modeling it as a disk of radius $\Delta\theta/2$ (where $\Delta\theta$ is the maximum separation between any two pixels in the blob) centered on $\hat{\mathbf{n}}_0$. The blob is then extended to a radius of $1.1 \times (\theta_{\text{crit,max}} + \Delta\theta/2)$, where $\theta_{\text{crit,max}}$ is found from Table I (which is dependent upon the needlelet type and frequency at which a feature is found) to ensure we include all related pixels. All pixels not contained in a blob are masked, and this new temperature map is passed to subsequent steps in the analysis pipeline. Eliminating irrelevant pixels allows us to drastically reduce the computational effort needed in the subsequent analysis.

5. Analysis of the WMAP end-to-end simulation

As there are many independent needlelet coefficients over the sky, it is inevitable that highly significant features will be detected in even a purely Gaussian CMB temperature map. In addition, residual foregrounds and instrumental artifacts may give rise to features which are misclassified as candidate bubble collisions by the pipeline. To understand how these effects might contribute to the needlelet significances, we ran the suite of needlelet transforms on the end-to-end simulation of the WMAP experiment

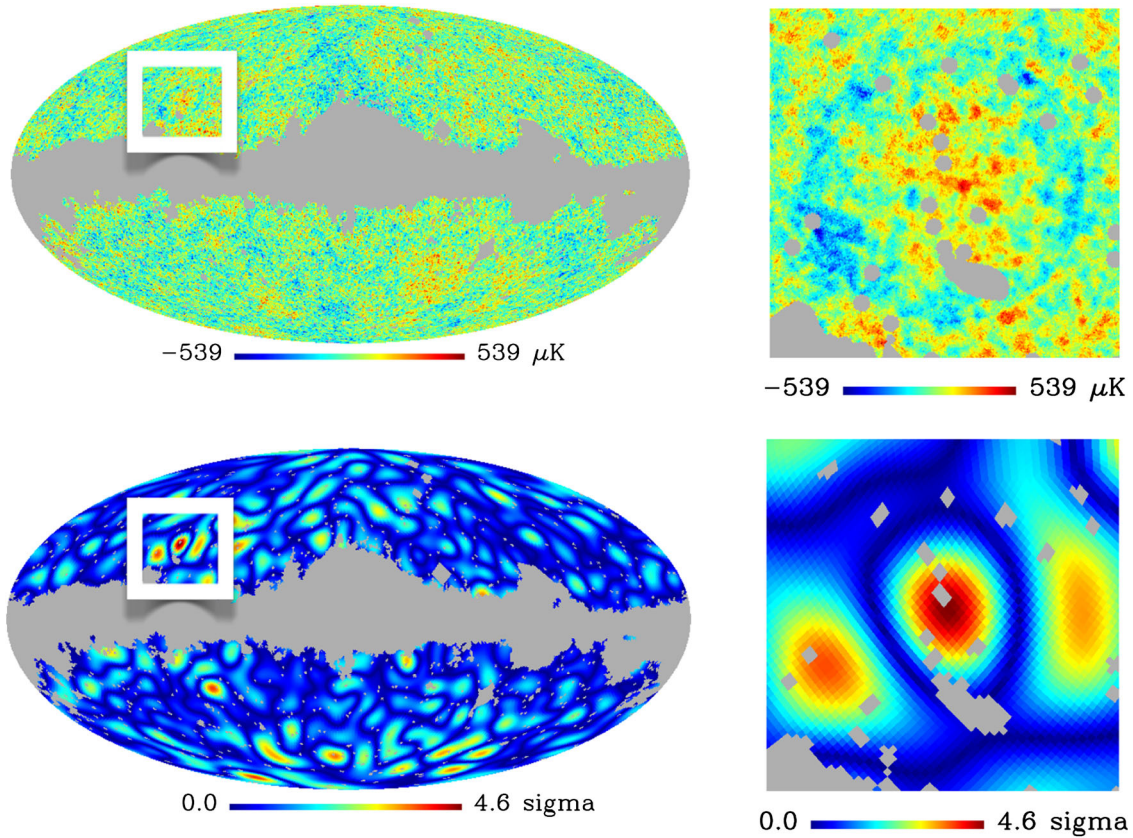


FIG. 10 (color online). The temperature (top panels) and needlet-coefficient significance (bottom panels) maps for a simulated bubble collision with $z_0 = 5 \times 10^{-5}$, $z_{\text{crit}} = -5 \times 10^{-5}$, $\theta_{\text{crit}} = 10^\circ$ on the CMB sky with the WMAP 7-year KQ75 mask applied. We show the map of needlet coefficients which gives the largest needlet response, in this case $j = 3$ for standard needlets with $B = 2.5$. The right-hand panels provide close-ups of the collision region.

(described in Sec. IV) with the 7-year KQ75 mask applied. As an illustration of our results, in Table II we give the number of blobs of varying significance found in the masked end-to-end simulation using standard needlets with $B = 2.5$. At increasingly high frequency, for which there are more independent needlet coefficients, more and more blobs are found that pass relatively large significance thresholds.

We use the results of Table II (and similar tables for other members of the needlet suite) to define a set of needlet frequency-dependent significance thresholds that allow a manageable number of false positives, while

TABLE II. The number of blobs found in the masked WMAP end-to-end simulation above significances ranging from $3 \leq S \leq 4$ for standard needlets with $B = 2.5$.

j	$S = 3$	3.25	3.5	3.75	4
0	0	0	0	0	0
1	0	0	0	0	0
2	0	0	0	0	0
3	10	4	2	1	0
4	23	10	4	0	0

retaining sensitivity to a fairly large range of collision template parameters. The significance thresholds we use in our analysis are shown in Table III. There are a total of 17 blobs in the masked end-to-end simulation that pass these thresholds. Comparing their locations on the sky, we can associate these blobs with 13 features (if a feature is picked up by multiple needlet types or frequencies, it can have multiple blobs associated with it). For three of these features, the set of pixels that pass the needlet threshold intersect the Galactic cut. We associate these with a response to the mask and do not consider them further. For the other ten features, the needlet type and frequency which yielded the maximum significance are recorded in Table IV.

6. Analysis of bubble-collision simulations

To assess how robustly the needlet transform can pick out a collision in the temperature map, we have performed an analysis of the simulated collisions described in Sec. IV. If the later steps of the analysis are to have a chance at detecting a simulated collision, it must be contained within the set of pixels defined by the blob. To determine if the needlet analysis has detected a bubble collision, we

TABLE III. Sensitivity thresholds S_{\min} and the number of significant blobs detected in the end-to-end simulations N_{blobs} for standard needlets with $B = 2.5$ (top), standard needlets with $B = 1.8$ (middle), and Mexican needlets with $B = 1.4$ (bottom). Only blobs that do not intersect the Galactic cut are reported. No results are shown for the standard needlets with $B = 1.8$, $j = 0$ as they have no support over the range of angular scales considered.

j	S_{\min}	N_{blobs}
0	3.0	0
1	3.0	0
2	3.5	0
3	3.5	1
4	3.75	0

j	S_{\min}	N_{blobs}
1	3.0	0
2	3.0	0
3	3.25	0
4	3.25	1
5	3.25	3
6	3.5	1
7	3.5	3

j	S_{\min}	N_{blobs}
0	3.0	0
1	3.0	0
2	3.0	0
3	3.0	0
4	3.0	0
5	3.0	0
6	3.5	0
7	3.5	0
8	3.5	0
9	3.5	0
10	3.5	1
11	3.5	1
12	3.75	1

therefore require that a blob exists which fully contains the region affected by the collision, and that the true center of the collision lies within the set of pixels passing the significance threshold.

The results of this analysis for the 5° and 10° collisions are shown in Fig. 11. We define the “exclusion region” of these plots as the part of parameter space for which all six realizations/locations yield a detection. If there were a bubble collision in the WMAP 7-year data with these parameters, it would be detected with high significance regardless of its location on the sky (as long as the collision was not significantly masked). The “sensitivity region” is defined as the part of parameter space for which *any* of the six realizations/locations yields a detection. A bubble

TABLE IV. Blobs found by the needlet transform in the WMAP end-to-end simulations with the 7-year KQ75 mask.

Feature	Blob	B	j	S
1	1	2.5	3	3.83
1	2	1.8	5	3.55
2	1	1.8	4	3.99
3	1	1.8	5	3.28
4	1	1.8	5	3.33
4	2	1.4	10	3.77
5	1	1.8	6	3.96
6	1	1.8	7	4.13
7	1	1.8	7	3.97
8	1	1.8	7	4.34
9	1	1.4	11	3.71
10	1	1.4	12	4.14

collision in this range of parameter space would be detected only for a favorable location or realization of the background fluctuations. The exclusion and sensitivity regions for the 25° collisions are identical to those for the 10° collisions.

Looking at the simulations in detail, there are a few general trends. First, for the needlets to pick out a collision, it is sufficient to have either a relatively large central amplitude z_0 or a relatively large temperature discontinuity z_{crit} at the causal edge. This is clear from the shape of the exclusion region in Fig. 11. From the size of the sensitivity region in this plot, one can also see that the amount of instrumental noise and particular realization of the background fluctuations can greatly affect the significance of the needlet coefficients in the vicinity of a collision. Many more collisions were detected in the low-noise region than the high-noise region of the sky. For collisions in the exclusion region, there is a significant needlet response for all three needlet types over a range of frequencies, with an average significance exceeding $S > 5$. Collisions in the sensitivity region are typically detected by only one needlet type and frequency, with an average significance near $S \approx 4$. As our ability to detect 5° , 10° , and 25° collisions is nearly the same, we conclude that these results are fairly representative of our detection limits over all angular scales $\geq 5^\circ$.

These general trends can be contrasted with the response to features found in the end-to-end simulations. Here, blobs are typically detected with a single needlet type and are near the significance threshold (not surprisingly, as the threshold was chosen to have this property). A feature detected in the data by multiple needlet types and/or frequencies at a significance of $S \geq 4$ would be a good bubble-collision candidate. However, we stress that many different underlying features could give rise to such a signal. The following steps in the analysis pipeline, which we now describe, are designed to verify if these candidates are in fact bubble collisions.




FIG. 11. Exclusion (black) and sensitivity (gray) regions for the needlet step of the analysis pipeline applied to a set of 5° (left panel) and 10° (right panel) simulated bubble collisions. If all realizations at the high- and low-noise locations yield a detection, we include the collision in the exclusion region; such collisions would certainly be detected as long as they were not significantly masked. If a detection is made in any realization/location, we include the collision in the sensitivity region; such collisions could be found if they were in a favorable location of the sky (i.e., low noise, or a region with a specific realization of background CMB fluctuations).

B. Edge detection

The first of the two parallel verification steps of the pipeline tests whether features highlighted by the blob-detection stage have circular temperature discontinuities. The unambiguous detection of a circular temperature discontinuity would strongly suggest that a particular feature highlighted by the needlets is in fact a bubble collision. We employ the Canny edge-detection algorithm [52], whereby the gradient of an image is generated and thinned into single-pixel proto-edges, the best of which are stitched together into “true” edges. We also use an adaptation of the CHT algorithm [53], which assigns a “score” dependent on how many edge pixels lie on circles of varying centers and radii. In this section, we describe our edge-detection algorithm, and study its performance on an end-to-end simulation of the experiment, as well as on simulated bubble collisions.

1. The Canny edge-detection algorithm

The Canny edge detector is *the* standard edge-detection algorithm in image-processing software and has recently been used to search for cosmic strings [30,31]. Designed as the optimal algorithm for localized, duplicate-free detection of edges within a noisy image, it uses three steps—smoothed gradient generation, nonmaximal suppression, and hysteresis thresholding—to extract contiguous edge sections. In Fig. 12, we depict each of these three steps as applied to a temperature map containing a circular discontinuity; each of these steps are in turn described as follows:

- (1) Smoothed gradient generation: The gradient of a Cartesian image is traditionally generated by con-
- (2) Nonmaximal suppression: The second step of the Canny algorithm reduces the smoothed gradients

volving the image with two small symmetric filters, each determining one orthogonal component of the gradient. A number of simple filters—typically 3×3 pixels—perform the job adequately, but the optimally adaptable filters are the first partial derivatives of the two-dimensional Gaussian [52]. Using these filters is equivalent to first convolving the image with a small Gaussian filter (and thus smoothing out the effects of small-scale noise on the gradient calculation, an important step given the small number of pixels involved in the calculation) and then finding its gradient components.

Unfortunately it is impossible to construct symmetric pixel-based gradient filters that cover the whole sphere. We therefore carry out both the Gaussian smoothing and gradient generation steps in harmonic space, making use of HEALPix’s in-built `alm2map_qr` subroutine to calculate the magnitude and direction of the gradient at each pixel. We smooth with a Gaussian filter of FWHM 0.22° —approximately two pixels’ width at the resolution of our input maps—to minimize features due to pixel noise while retaining longer edges.

The gradient maps are generated before masking to reduce “ringing” from the sharp sky cut back into the map. The smoothing step ensures that any leakage from masked features is restricted to areas a few pixels within the sky cut and affects only areas a few pixels outside of the cut. Nevertheless, any features found close to the mask should be carefully examined to check if they are generated from within the mask.




FIG. 12 (color online). An illustration of the Canny algorithm for edge detection. The input temperature map contains a circular discontinuity, which can be seen in a map of the gradient magnitude as a local maximum. Nonmaximal suppression selects for local maxima in the gradient map. The hysteresis thresholding step finds stitched edges by comparing the local direction of the gradient at adjacent pixels.

produced by the first step into local maxima. At this stage, all pixels are assumed to belong to a local edge, whose direction is defined to be perpendicular to the local gradient direction. Taking each pixel at a time, the two direct neighbors which lie closest to the local edge are found. The gradient magnitudes of the three pixels are compared, and the central pixel's gradient magnitude is set to zero unless it is the largest of the three. Processing each pixel in turn reduces the gradient magnitude map to only the local maxima (see the central panels of Fig. 12).

As an example, consider the simplest case of crossing a sharp discontinuity along a perpendicular path. The gradient direction is constant at each pixel, whereas the smoothed gradient magnitudes increase until the edge is crossed, when they start to decrease. A nonmaximal suppression algorithm tracks along the path setting all of the gradient magnitudes to zero apart from the pixel on (or closest to) the edge.

- (3) Hysteresis thresholding: At this stage of the Canny algorithm, we have gradient magnitudes and directions for a set of local maxima of varying amplitude, some corresponding to true edges (which may have been disrupted by noise) and others to runs of noisy pixels or to more slowly varying boundaries of CMB patches. To filter out true edges from the noise, and stitch together any edges that have been split, the final step of the algorithm takes advantage of the fact that, unlike randomly oriented noise, true edges conserve their gradient magnitude and direction (to an extent affected by the shape of the edge, the pixelization scheme, and the noise level) over their path.

Hysteresis thresholding involves first setting an upper threshold for the gradient magnitude: Any pixels surpassing this threshold are considered to be part of true edges, and a new "true edge" map is created with these pixels' positions marked. A

second, lower threshold is then set. Hysteresis thresholding then proceeds as follows:

- (a) The map is scanned until a true edge pixel is found.
- (b) The next potential edge pixel is defined to be the direct neighbor closest to 90° clockwise from the local gradient direction. The gradient direction of this pixel is compared to the current pixel's. To do so, the local phase angles to the current pixel's nearest neighbors are found and used to define the neighbor closest to the current gradient direction. The neighbors adjacent to this pixel are then determined. The gradient direction at the next potential edge pixel is required to lie between the phase angles of these neighbors. This rather loose requirement allows the algorithm to step along pixelated curved edges.
 - (i) If the two pixels' gradient directions match within the tolerance:
 - (A) If the neighbor's gradient magnitude passes the low threshold but not the high, it is considered to be part of a potential true edge. Its position is marked in a history array; then the algorithm "steps onto" this pixel and the process is repeated from step (b) until one of the other conditions is met.
 - (B) If the neighbor's gradient magnitude passes the high threshold, all pixels found on the way from the source pixel are confirmed as a true edge. Their positions are marked on the true edge map, and the algorithm returns to step (a).
 - (C) If the neighbor's gradient magnitude fails both thresholds, the edge is considered to be false: The history of potential edge pixels found on the way from the source pixel is erased, and the algorithm returns to step (a).
 - (ii) If the two pixels' gradient directions do not match within the required tolerance:
 - (A) If the neighbor's gradient magnitude passes the high threshold or the neighbor is already marked in the history array, all pixels found on the way from

the source pixel are confirmed as a true edge. Their positions are marked on the true edge map, and the algorithm returns to step (a). This ensures simple branched and looped edges can be reconstructed.

- (B) If the neighbor’s magnitude does not pass the high threshold and the pixel is not already marked in the history array, the edge is considered to be false, the history of potential edge pixels found on the way from the source pixel is erased, and the algorithm returns to step (a).

The entire process is then repeated, choosing the neighbor closest to 90° counterclockwise to the local gradient direction in step (b).

The end product of the Canny algorithm is a Boolean map of stitched true edges (see the right-hand panel of Fig. 12). To reduce computation time, the pipeline from hysteresis thresholding onwards is restricted to the blobs produced by extending the regions passing the needlet significance test to ensure any discontinuity is fully contained.

Care must be taken when setting the thresholds used in the hysteresis thresholding step. If either threshold is set too high, very few edges are confirmed. If the low threshold is set too low, a huge amount of potential edges are considered, and the algorithm runs extremely slowly. As the edges we could potentially find must be comparable in amplitude to the CMB signal and detector noise (as they have not yet been discovered by eye) and have been smoothed by the WMAP beam, we set low thresholds to err on the side of caution. Low and high thresholds of 30% and 40% of the maximum gradient magnitude found in each search region are found empirically to confirm edges generated in simulations in feasible computation time scales. This means that the gradients associated with the strongest CMB features—typically $\sim 1^\circ$ in scale—are classified as edges, as shown in Fig. 12.

2. Circular Hough transform

The maps of stitched candidate edges found using the Canny algorithm are processed using the circular Hough transform to search for the presence of circular edges. The basic idea, as shown in Fig. 13, is to count the number of intersections between circular arcs of varying radii centered on each of the candidate edge pixels and oriented along the local gradient direction. If there is a circular edge in the map, the number of intersections will be maximized at the center of the circular edge when the radius of the circular arc matches that of the edge.

Assuming an edge pixel forms part of a circular edge of angular radius θ_{crit} , one can define a prescription for the set of pixels that are the potential centers of the edge. The two most likely candidates in this set are the pixel θ_{crit} away in the direction of the local gradient, and the pixel the same distance in the opposite direction; the edge could be a step up or a step down. Building in flexibility to cope with pixelation and noise effects, this set is expanded to two annular arcs of radius θ_{crit} , oriented in the direction of the local gradient and centered on the edge pixel.

The CHT works by assuming that *all* edge pixels are part of circular edges. The most likely circle center at a given radius is defined to be the pixel that is included in the greatest number of these arcs, counted using an accumulator array. If the search radius matches the radius of a circular edge within the map, we expect all of the circular edge pixels’ arcs to include the true center, and the CHT accumulator will show a single clear peak. If the search and true radii are discrepant, fewer of the circular edge pixels’ arcs will intersect, and this peak will appear as a ring with decreased amplitude (see Fig. 13). When the search and true radii are very discrepant, any rings will disappear beneath the background due to randomly oriented noise and other noncircular edges. Note that “noncircular edges” will include the $\sim 1^\circ$ CMB fluctuations that are qualified as edges by the hysteresis thresholding step.




FIG. 13 (color online). A depiction of the circular Hough transform (CHT). On the left is a Boolean map of edge pixels, as output by the Canny algorithm. Centering a pair of arcs oriented in the direction of the local gradient about each edge pixel, the CHT counts the number of times each pixel is intersected. The presence of a circular edge is indicated by a maximum in the CHT score—the hit count divided by the arc radius—as the arc radii are varied. On the left, we show a set of arcs, centered on four pixels on the circular edge we wish to detect; there will be no clear peak in the CHT score for this radius. Increasing the arc radius to match that of the circular edge (center), there will be a large number of hits at the true center of the edge. On the right, we show the actual map of the CHT score at each pixel for this example. As this has been scanned at the correct angular scale, there is a large peak at the center of the circular edge.

To compare the CHT results at different search radii, one must divide out the approximately linear growth with angular radius of the number of pixels in each annular arc. We call this normalized quantity the ‘‘CHT score.’’ The most likely center and radius of a circular edge within a map can therefore be found by scanning the map with the CHT at a range of radii and determining the maximum CHT score.

The blob-detection step provides the range of scales $\theta_{\text{crit,min}} \leq \theta_{\text{crit},i} \leq \theta_{\text{crit,max}}$ of potential circular edges present in each blob. To determine whether a blob contains a circular edge, we compare the CHT scores obtained by scanning at every 0.25° increment within this range, using annular arcs that are 0.25° thick and which cover 45° of phase about each edge pixel. The annular arcs are therefore approximately two pixels thick and are fairly wide to account for the effects of pixelation on the gradient direction. The thickness of the CHT annular arcs leads to an uncertainty in the CHT radius of 0.25° and position of 0.50° . If a circular edge is detected, we expect a clear peak in the CHT results for a particular blob.

In Fig. 14 we show the output of the edge-detection algorithm on our illustrative example bubble-collision simulation (see Fig. 10). On the left is the portion of the temperature map containing the collision. On the right we plot the CHT score in the pixels that passed the needlet significance threshold for $\theta_{\text{crit},i} = 10^\circ$ (the input value). There is a clear peak at the location of the true center of the simulated bubble collision, which is ~ 3 times the average response at other pixels. Since this feature was flagged in the blob-detection step for standard needlets with $B = 2.5$ at $j = 3$, the range of radii scanned during the CHT step is determined from Table I to be $5^\circ \leq \theta_{\text{crit}} \leq 14^\circ$. This range contains the true radius $\theta_{\text{crit}} = 10^\circ$. In Fig. 15 we plot the maximum CHT score found in the map for each circular radius, which contains a clear peak at the true radius of the causal boundary. This signal is a clear and unambiguous




FIG. 14 (color online). The temperature map (left panel) and CHT score (right panel) for our illustrative example of a 10° bubble-collision simulation. The CHT score is recorded at each pixel passing the needlet significance threshold. For a search radius of 10° there is a clear peak in the CHT score at the center of the simulated collision.




FIG. 15. The global maximum of the CHT score at each circle radius for the collision simulation shown in Fig. 14. The collision has a maximum response for standard needlets with $B = 2.5$ at $j = 3$, from which Table I sets the search range to be $5^\circ \leq \theta_{\text{crit}} \leq 14^\circ$. The peak of the CHT score at 10° clearly identifies the correct angular scale of the simulated collision.

signature of a bubble collision. From visual inspection of the temperature map, it can be seen that we are able to clearly detect the edge even though the background fluctuations, noise, and beam drastically reduce the sharpness of the observed temperature discontinuity.

3. Analysis of the WMAP end-to-end simulation

We expect strong circular edges to be extremely rare in a purely Gaussian CMB temperature map. However, it is possible that foregrounds, instrumental noise, the mask, and other experimental artifacts could lead to a spurious detection of a circular edge. To evaluate this, we have performed the edge-detection step of our analysis pipeline on the features that passed the significance threshold in the WMAP end-to-end simulation (see Table IV) with the KQ75 mask applied.

Comparing each feature in the end-to-end simulation with the bubble-collision example studied above, the peak structure of the CHT score as a function of angle and morphology in pixel space are both drastically different. Examining the maximum CHT score as a function of circular radius, although there are several peaks, the clearest of which is shown in Fig. 16, their amplitude relative to the average score is nowhere near that of the collision example shown in Fig. 15. In addition, from the plots of the CHT score at each pixel, there are typically a number of fairly broad local maxima at different locations with approximately the same score. This is in contrast to the collision example of Fig. 14, which yields a highly peaked score around a small number of pixels.

4. Analysis of bubble-collision simulations

To better understand the response of our edge-detection algorithm to the signal from a bubble collision in WMAP-quality data, we have analyzed the simulations described in Sec. IV. We use as inputs the blobs found using the first step

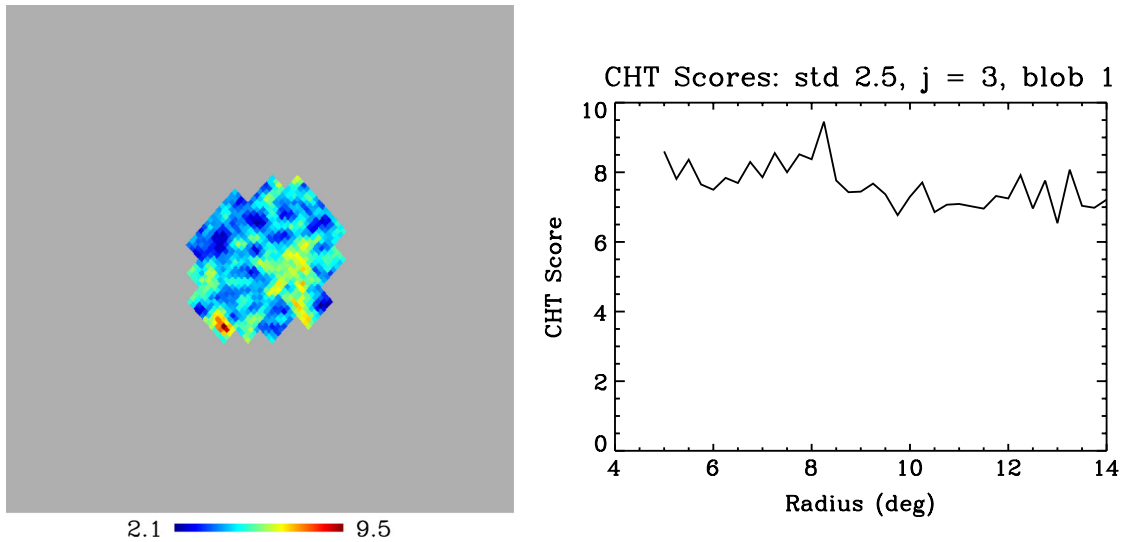


FIG. 16 (color online). The most edgelike feature in the WMAP end-to-end simulation. The contrast in scores as a function of position (left panel) and radius (right panel) is greatly reduced compared to the collision example (Figs. 14 and 15).

of the pipeline and search for circular edges over the range of angular scales appropriate to the needlet type and frequency for each blob (see Table I). We conclude that a true causal edge has been detected if there is a global maximum for the CHT score at the radius of the true edge and the pixel with the highest score is within a typical CHT error (0.5°) of the actual center. We again present our results in the form of a contour plot denoting exclusion and sensitivity regions in the parameter space of z_0 and z_{crit} . This is shown in Fig. 17 for simulated bubbles with $\theta_{\text{crit}} = 5^\circ$ and 10° . The plot for the 25° collisions is identical to the plot for 10° collisions.

Again, from the size of the sensitivity region, we conclude that our ability to make a detection is dependent on the location of the collision and the particular realization of the background CMB fluctuations. The exclusion region

for the 10° (and 25°) collisions is far larger than for the 5° collisions. We attribute this to the proliferation of $\sim 1^\circ$ -sized features in the background fluctuations, which can both disrupt a significant fraction of the edge pixels in a small collision and swamp the collision signal with their own strong gradients. We therefore expect our sensitivity to edges at small angular scales $\theta_{\text{crit}} \lesssim 5^\circ$ to be quite poor at WMAP resolution. As the performance of the edge-detection algorithm for 10° and 25° collisions is identical, we conclude that the 10° results are fairly representative of our sensitivities over a wide range of angular scales $\theta_{\text{crit}} \gtrsim 10^\circ$. Most of the collisions we mark as a detection have a clear peak in the CHT score of the type seen in Fig. 15. If a collision has parameters in the exclusion region, it would be reliably detected. Based on these results, the first two steps of our pipeline can detect bubble

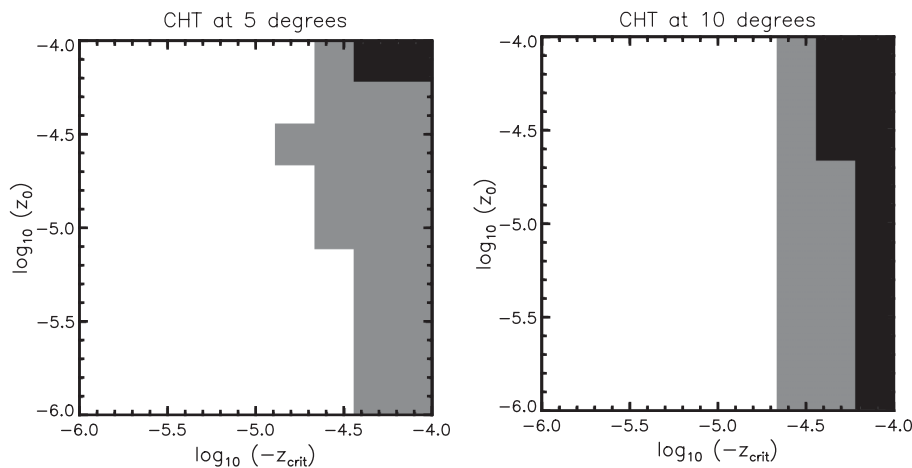


FIG. 17. Exclusion (black) and sensitivity (gray) regions (see Fig. 11) for the edge-detection step of the analysis pipeline applied to a set of 5° (left panel) and 10° (right panel) simulated bubble collisions.

collisions with central modulations $|z_0| \gtrsim 3 \times 10^{-5}$ and causal edges $|z_{\text{crit}}| \gtrsim 3 \times 10^{-5}$ at $\theta_{\text{crit}} \gtrsim 5^\circ$ in WMAP-quality data.

C. Parameter estimation and model selection

In many CMB anomaly analyses (but not all—see, e.g., Refs. [58–60]), the significance of a signal is quantified by calculating the frequentist P value of some relevant statistic. This typically involves doing a large number of Monte Carlo realizations of the standard cosmological model (i.e., the “null hypothesis”), calculating the above statistic for each, and finding the fraction for which the statistic has a “more extreme” value than was actually observed. There are several problems with this approach. First, the calculated P value is only a measure of how (un)likely the measured data were given the null hypothesis of the standard model; no comparison is made to an alternative model (which is what we are primarily interested in here). Second, the notion of more extreme is fundamentally ambiguous—both “more discrepant” (i.e., values of the statistic which are farther from some fiducial expected value than that which was measured) or “less likely” are common choices.⁸ The heart of the problem is that all such P values are integrals over the likelihood, whereas it is only the likelihood of the actual data that is relevant. The fact that the likelihood and its integral generally have a similar qualitative dependence in the tail(s) of the distribution (i.e., both tend to zero for extreme values of the statistic) can mask this problem. In particular, if the tails of the likelihood are Gaussian then the integral that gives the P value falls off more rapidly than the likelihood itself, and so the resultant P values are unreasonably harsh on the null hypothesis. A related problem is that many attempts to identify CMB anomalies using frequentist P values are overly sensitive to *a posteriori* selection effects (see, e.g., Refs. [34,61] for a discussion of this effect). Here the issue that the statistics being applied to the data are often chosen on the basis of interesting features initially identified in the same data is not intrinsic to frequentist methods (which, correctly, do not permit any data to be used more than once); but the need to invent a statistic from which to calculate a P value can make it hard to avoid this trap. For these reasons we do not use P values in our analysis.

Instead, we adopt a Bayesian approach. Bayes’s theorem provides a prescription for parameter estimation. In addition, given that we have two well-defined hypotheses, we can utilize Bayesian model comparison to make probabilistic statements about the degree to which the available data (and theoretical prior information) imply that bubble collisions have been observed. As shown by Ref. [62],

⁸For simple, single-peaked, likelihoods these two definitions are at least equivalent, but in some cases (e.g., a likelihood that is constant over some finite range) neither definition is satisfactory.

Bayesian methods are the only self-consistent framework for such calculations. The optimal Bayesian calculation would be to evaluate the likelihood of the entire WMAP data set under the two models; however, this is not computationally feasible at present. In the Appendix, we outline a set of simplifications that allow us to approximate the optimal Bayesian result. As outlined in Sec. III, we utilize the information on the location and scales of the most probable bubble-collision sites obtained in the blob-detection step of the pipeline to implement this procedure. Even this reduced problem is computationally demanding: Analysis of the blobs detected in the WMAP 7-year data during the first steps of the pipeline requires three days’ processing on 28 cores. Working at full resolution is necessary to ensure that any possible circular temperature discontinuities are examined.

These computational limitations also mean we are only able to process a limited number of simulated temperature maps with and without bubble collisions. The WMAP end-to-end simulation provides a great asset at this stage, giving the best possible measure of what false detections are to be expected from experimental effects and any systematic errors that are not included in our likelihood. We also analyze a small number of representative bubble-collision simulations to obtain an estimate of the strength of signal we are looking for.

We now describe our methods and the results from simulations in greater detail.

1. Bayesian formalism

A model of eternal inflation predicts the average number of collisions \bar{N}_s that are, in principle, detectable by our pipeline on the full sky.⁹ The ultimate goal of our Bayesian analysis is to evaluate the full posterior probability distribution for \bar{N}_s , given a CMB data set \mathbf{d} covering a sky fraction f_{sky} . Using Bayes’s theorem, this can be written as

$$\Pr(\bar{N}_s | \mathbf{d}, f_{\text{sky}}) = \frac{\Pr(\bar{N}_s) \Pr(\mathbf{d} | \bar{N}_s, f_{\text{sky}})}{\Pr(\mathbf{d} | f_{\text{sky}})}. \quad (26)$$

The form of the posterior depends on the model prior $\Pr(\bar{N}_s)$ and the *evidence* (also known as the model likelihood) $\Pr(\mathbf{d} | \bar{N}_s, f_{\text{sky}})$. The evidence is defined by marginalizing the likelihood, $\Pr(\mathbf{d} | \mathbf{m}, \bar{N}_s, f_{\text{sky}})$, over the n parameters describing a collision, as specified by the model \mathbf{m} . Once the shape of the posterior has been determined, it is normalized using $\Pr(\mathbf{d} | f_{\text{sky}})$. The posterior leads directly to constraints on the values of \bar{N}_s consistent with a CMB data set.

In a landscape scenario, \bar{N}_s can be considered as a continuous parameter and the prior $\Pr(\bar{N}_s)$ would be determined from a measure over the possible values of \bar{N}_s .

⁹The number of detectable sources \bar{N}_s is a subset of the total number of sources on the sky \bar{N} [Eq. (1)].

We can also view \bar{N}_s as a proxy for different models of eternal inflation (i.e., selecting a single value of \bar{N}_s), as described further in Sec. II. The standard cosmological model without bubble collisions is specified by the case $\bar{N}_s = 0$. Using Eq. (26), the probability of a model which predicts \bar{N}_s collisions (on average) relative to that of the standard cosmological model is

$$\frac{\Pr(\bar{N}_s|\mathbf{d}, f_{\text{sky}})}{\Pr(0|\mathbf{d}, f_{\text{sky}})} = \frac{\Pr(\bar{N}_s) \Pr(\mathbf{d}|\bar{N}_s, f_{\text{sky}})}{\Pr(0) \Pr(\mathbf{d}|0, f_{\text{sky}})}. \quad (27)$$

The model priors and the evidence values play an equal role in this relationship, but in the absence of a detailed understanding of the former, it is often useful to proceed under the assumption that the two models are equally probable *a priori*. A theory predicting an expected \bar{N}_s collisions is favored over the standard model when the relative probability on the left-hand side of Eq. (27) is greater than unity.

It is also useful to provide heuristic conversions between the Bayesian evidence ratio and other commonly used model comparison quantities. The number of “sigma” of an anomaly statistic, $N\sigma$, is often used to characterize the deviation from a null model, but it is unambiguously defined only in the case in which the null distribution of the chosen statistic is Gaussian with zero mean. In such a case the probability of measuring an $N\sigma$ deviation is $P(N) \propto \exp(-N^2/2)$, which can be identified approximately with the inverse of the ratio in Eq. (27), so that, e.g., a 3σ detection is comparable to a ratio of approximately 100. However we emphasize that both the number of sigma and related statistics such as $\Delta\chi^2$ are of limited utility in the context of all but the most trivial model comparison problems.

Computing $\Pr(\mathbf{d}|\bar{N}_s, f_{\text{sky}})$ by marginalizing over the likelihood for the full prior volume is an immense computational task, requiring the inversion of the full-sky WMAP covariance matrix at full resolution and marginalizing over all possible numbers, locations, and sizes of collisions. However, taking advantage of the fact that bubble collisions produce discrete localized effects on the CMB sky, it is possible to approximate the full-sky Bayesian analysis by a patchwise analysis if the most promising candidate signatures can be identified in advance. We describe in detail in the Appendix an algorithm to perform such a patchwise approximation to this full multidimensional integral.

The key ingredient is determining the regions of parameter space where the likelihood is significantly peaked, and hence gives the most significant contributions to the evidence. If these regions can be identified, the integral need only be performed over the restricted ranges to obtain an estimate of the evidence at greatly reduced computational cost. We use the results of the blob-detection step of the analysis pipeline to identify the patches which likely make the most significant contributions to the integral. Assuming

that the bubble-collision model likelihood is peaked in the N_b detected blobs, we show in the Appendix that the unnormalized posterior can be approximated as

$$\Pr(\bar{N}_s|\mathbf{d}, f_{\text{sky}}) \propto \Pr(\bar{N}_s) e^{-f_{\text{sky}}\bar{N}_s} \sum_{N_s=0}^{N_b} \frac{(f_{\text{sky}}\bar{N}_s)^{N_s}}{N_s!} \times \sum_{b_1, b_2, \dots, b_{N_s}=1}^{N_b} \left[\prod_{s=1}^{N_s} \rho_{b_s} \prod_{i,j=1}^{N_s} (1 - \delta_{s_i, s_j}) \right], \quad (28)$$

where the prefactors reflect the fact that the number of collisions present on the observable sky, N_s , is the realization of a Poisson-like process (of mean $f_{\text{sky}}\bar{N}_s$), and ρ_b is the evidence ratio evaluated within a candidate collision region (with data subset \mathbf{d}_b) using a single bubble-collision template

$$\rho_b = \frac{\Pr(\mathbf{d}_b|1)}{\Pr(\mathbf{d}_b|0)}. \quad (29)$$

The posterior can therefore be built from local measures of how well the data are described by the standard model with and without a collision template. Once Eq. (28) is obtained in a particular case, it can be normalized, although this is not strictly necessary to perform the parameter estimation and model selection analyses.

To illustrate some possibilities, in Fig. 18 we plot the normalized posterior assuming $f_{\text{sky}} = 0.7$ (from the KQ75 mask) and a uniform prior on \bar{N}_s , for the case where there is a single detected blob (left panel), and four detected blobs (right panel). A theory predicting a particular value of \bar{N}_s will be preferred to a theory without bubble collisions if the ratio in Eq. (27) is larger than 1. This amounts to comparing the posterior at some value of \bar{N}_s to the posterior at $\bar{N}_s = 0$ (dashed line). To prefer *any* theory with bubble collisions, in the one-blob case it is necessary for the blob to yield a local evidence ratio larger than 1 (here, we plot the posterior assuming $\rho_b = 4$). This is not true when there are multiple blobs, as can be seen in the right panel of Fig. 18, where we plot the posterior assuming each blob has a local evidence ratio $\rho_b = 0.5$. The bubble-collision hypothesis (for some values of \bar{N}_s) is preferred even when the local evidence ratios are less than 1: A number of marginal detections can be significant when considered together. We can also obtain any desired confidence intervals on \bar{N}_s by examining the shape of the posterior (although it is always the whole distribution that is the full answer to any parameter estimation problem).

When the local evidence ratios are large, the posterior can be approximated by Eq. (A16), appropriately normalized. In Fig. 19, we plot the posterior in the limit of large evidence ratios (again assuming $f_{\text{sky}} = 0.7$) for no blobs, two blobs, and four blobs. Even in the presence of large local evidence ratios, it can be seen that the posterior has a




FIG. 18. The normalized posterior $\Pr(\bar{N}_s | N_b, f_{\text{sky}})$ [see Eq. (A16)] assuming $f_{\text{sky}} = 0.7$. In the left panel, we show the posterior obtained for one blob $N_b = 1$ for a local evidence ratio $\rho_b = 4$. Comparing with the posterior at $\bar{N}_s = 0$ (dashed line), we see that any theory predicting $\bar{N}_s \leq 4$ will be preferred over the theory without bubble collisions. In the right panel, we show the posterior obtained for four blobs with identical local evidence ratios $\rho_b = 1/2$. Again, comparing with the posterior at $\bar{N}_s = 0$, any theory with $\bar{N}_s \leq 7$ will be preferred over the theory without bubble collisions. When there are multiple blobs, the bubble-collision hypothesis can be supported even when the evidence ratio for each blob is less than 1.

significant spread due to cosmic variance: We have access to only one realization of bubble collisions on the CMB sky. Note that this is true even when there are no detected blobs. When there are multiple decisively detected blobs, the posterior correctly assigns a very small probability to $\bar{N}_s = 0$.

Our analysis also provides constraints on the parameter values of each candidate collision. The constraints on the n template parameters \mathbf{m} are encoded in their joint posterior distribution

$$\Pr(\mathbf{m} | \mathbf{d}_b, 1) = \frac{\Pr(\mathbf{m}) \Pr(\mathbf{d}_b | \mathbf{m}, 1)}{\Pr(\mathbf{d}_b, 1)}. \quad (30)$$

The marginal distribution of any subset of the parameters is given by integrating $\Pr(\mathbf{m} | \mathbf{d}_b, 1)$ over the remaining parameters which are not of interest. For the bubble-collision model the parameters should include both those describing the collision and the global cosmological parameters; marginalizing over the latter would give constraints on the properties of a (putative) detected collision. We now discuss the analysis of the likelihood and evidence ratios for a patch in greater detail.




FIG. 19. The full posterior $\Pr(\bar{N}_s | N_b, f_{\text{sky}})$ [Eq. (A16)] that would be obtained from a conclusive detection (i.e., $\rho_b \gg 1$) of $N_b = 0, 2, 4$ (solid, dashed, and dot-dashed curves) blobs containing bubble collisions assuming $f_{\text{sky}} = 0.7$. The presence of a sky cut skews the distribution toward $\bar{N}_s > N_b$. Note that even when features are conclusively detected, there is an intrinsic uncertainty in \bar{N}_s ; this is a form of cosmic variance.

2. Analysis of candidate bubble-collision patches

At the heart of the above formalism for assessing the full posterior for \bar{N}_s is the evaluation of the patch likelihood for a single collision, $\Pr(\mathbf{d}_b | \mathbf{m}, 1)$. Here the data, \mathbf{d}_b , are the measured temperature values of the pixels in the vicinity of the detected blob that are not in the sky cut. The bubble-collision model parameters, \mathbf{m} , should include both those that describe the collision, $\{z_0, z_{\text{crit}}, \theta_{\text{crit}}, \theta_0, \phi_0\}$, as well as the cosmological parameters which determine the CMB power spectrum. However any plausible bubble collision would be sufficiently localized so that the cosmological parameters are essentially uncorrelated with them; moreover they are sufficiently tightly constrained from CMB measurements so that their uncertainties are minimal in the context of a template-matching problem such as this. Hence we fix the cosmological parameters to their best-fit WMAP values [55] and only the bubble-collision parameters are varied. Hence $\mathbf{m} = \{z_0, z_{\text{crit}}, \theta_{\text{crit}}, \theta_0, \phi_0\}$ for the bubble-collision model, and there are no free parameters in the null model. Indeed, the no-collision model can be treated as a special case of the collision model in which the collision has zero amplitude.

As both the CMB signal and the WMAP noise are Gaussian, the likelihood has the form

$$\begin{aligned} \Pr(\mathbf{d}_b | \mathbf{m}, 1) &\propto \exp\left(-\frac{1}{2}\chi^2\right) \\ &= \exp\left\{-\frac{1}{2}[\mathbf{d}_b - \mathbf{t}(\mathbf{m})]^T \mathbf{C}_b^{-1} [\mathbf{d}_b - \mathbf{t}(\mathbf{m})]\right\}, \end{aligned} \quad (31)$$

where $\mathbf{t}(\mathbf{m})$ is the temperature modulation caused by the collision and \mathbf{C}_b is the pixel-pixel covariance matrix. The temperature modulation of the p th pixel is given from Eq. (3) as $t_p = 1 + f(\hat{\mathbf{n}}_p)$, where $\hat{\mathbf{n}}$ is the position on the sky. The covariance matrix includes CMB cosmic variance, Gaussian smoothing approximating the WMAP W-band beam, and the pixel-dependent WMAP noise. The covariance between two pixels p and q with angular positions $\hat{\mathbf{n}}_p$ and $\hat{\mathbf{n}}_q$ is hence given by

$$C_{p,q} = N_{p,q} + \sum_{\ell} \frac{2\ell + 1}{4\pi} \bar{C}_{\ell} P_{\ell}(\hat{\mathbf{n}}_p \cdot \hat{\mathbf{n}}_q), \quad (32)$$

where \bar{C}_{ℓ} is the best-fit WMAP CMB power spectrum convolved with a Gaussian beam of FWHM 0.22° , $P_{\ell}(x)$ is the Legendre polynomial of degree ℓ , and $N_{p,q}$ is the noise covariance between pixels. This is taken to be

$$N_{p,q} = \delta_{p,q} \frac{\sigma_w^2}{N_{\text{obs},p}}, \quad (33)$$

where $\delta_{p,q}$ is the Kronecker delta function, $\sigma_w = 6.549$ mK is the rms noise of the W-band detectors, and $N_{\text{obs},p}$ is the number of times WMAP has observed the p th pixel. To preserve any edges, we must invert $\mathbf{C}b$ at full resolution. Given available computational resources, the maximum area of the sky we can study at any one time is limited to patches of radius $\sim 11^{\circ}$ surrounding the center of each detected blob.

The evaluation of the evidence integral Eq. (A12) and the full characterization of the posterior distribution of the parameters are both computationally challenging—even when restricted to small patches—as they require a large number of likelihood evaluations. In all but the simplest of cases it is fatally inefficient to evaluate the likelihood over a multidimensional grid, and so a variety of sampling algorithms have been developed in which the likelihood is evaluated only in the high posterior regions that are of most interest. For both parameter estimation and evidence calculations we use the nested sampling algorithm [63] as implemented in the publicly available MULTINEST package [54]. MULTINEST performs numerical integration in order to estimate the evidence values; the required convergence of the integration can be adjusted to balance computation speed with accuracy. At the settings we use, the evidence values returned by MULTINEST are accurate to $\sim 10\%$. We use the `getdist` routine included in COSMOMC [64] to extract parameter estimates and uncertainties.

The parameter prior $\text{Pr}(\mathbf{m})$ in Eq. (30) is derived from theory, previous experimental results, and the limitations of the data set and pipeline: It encompasses the full prior understanding of what defines a detectable collision. Because we lack a detailed theoretical prediction for the amplitude parameters in each template (as discussed in Sec. II), we assume a uniform prior on z_0 and z_{crit} over the ranges $-10^{-4} \leq z_0 \leq 10^{-4}$ and $-10^{-4} \leq z_{\text{crit}} \leq 10^{-4}$, set by the observed temperature fluctuations in the CMB. Bubbles with larger values of these parameters would have been visible to the naked eye in any existing CMB data set. Bubble collisions are expected to be distributed isotropically on the CMB sky, and so we assume uniform priors on the full ranges of $\{\cos\theta_0, \phi_0\}$ to ensure that the probability of finding a bubble per unit area is constant across the sphere. Theory predicts that bubble-collision radii should range from 0° to half-sky, but our pipeline’s sensitivity is restricted by CMB power at small

scales and computational requirements at large scales. The nonzero prior range for *detectable* bubble collisions is accordingly restricted, and we assume uniform priors on θ_{crit} values in the range $2^{\circ} \leq \theta_{\text{crit}} \leq 11^{\circ}$.¹⁰

To minimize computation time, the evidence integrals are calculated only over the parameter ranges within which the priors are nonzero and the likelihood is peaked. For each feature, the angular scale lookup tables (Table I) indicate the range of interest for θ_{crit} . Merging all of the sets of significant pixels found for each feature yields the ranges for $\{\theta_0, \phi_0\}$. As the needlets are equally sensitive to cold and hot features with varying profiles, little information about the ranges of interest for $\{z_0, z_{\text{crit}}\}$ can be extracted from the blob-detection results, so the full prior volume must be considered. The accuracy of this procedure has been verified by performing the integral on the same patch of sky using different parameter ranges for $\{\theta_{\text{crit}}, \theta_0, \phi_0\}$. As long as the likelihood peak is encompassed by the parameter ranges given to MULTINEST, the returned evidence values agree to within numerical accuracy.

For the fiducial collision example shown in Fig. 10, our analysis yields an evidence ratio of $\ln\rho_b = 119.8 \pm 0.3$: The collision model is a very good fit to the data. The full-sky posterior would favor any theory predicting bubble collisions over a large range of \bar{N}_s . The marginalized bounds on the parameters are compared to the input parameters in the first row of Table VI: The agreement is excellent. However, to make a final judgment about a detection, we must ask what types of evidence ratios we get for false detections in the WMAP end-to-end simulation.

3. Analysis of the WMAP end-to-end simulation

We have performed the full Bayesian parameter estimation and model selection analysis on the blobs found in the WMAP end-to-end simulation (see Table IV). The total processing time for the full pipeline to run on this single map is on the order of 12 hours on 28 cores. Our results for the evidence ratios and marginalized parameter constraints for $\{z_0, z_{\text{crit}}, \theta_{\text{crit}}, \theta_0, \phi_0\}$ for each feature are recorded in Table V.

The evidence ratios for the features identified in the blob-detection step of the pipeline are all significantly less than 1. We can therefore approximate the full posterior for \bar{N}_s by Eq. (A20), and rule out $\bar{N}_s \gtrsim 1.6$ at the 68% confidence level. The posterior is maximized at $\bar{N}_s = 0$,

¹⁰Equation (2) predicts that the angular scale distribution for all bubbles falling within our past light cone varies with $\sin\theta_{\text{crit}}$. However, this is derived under the assumption that collisions do not affect our bubble interior, and a more careful treatment might lead to a correlation between the values of z_0 , z_{crit} , and θ_{crit} . To retain consistency with our uniform priors on z_0 and z_{crit} , we assume a uniform prior on θ_{crit} . Regardless, both choices for the prior lead to identical conclusions for the WMAP 7-year data.

TABLE V. Results of the Bayesian parameter estimation and model selection analysis for the WMAP end-to-end simulation. The ranges of θ_{crit} are determined from the needlet response (see Table I). By computational necessity, the evidence integral is truncated at 11° . Reported error bars are at 68% C.L. Angular quantities are quoted in degrees.

Feature	θ_{crit} range	$\ln\rho_b$	z_0	z_{crit}	θ_{crit}	θ_0	ϕ_0
1	5–14	-7.9 ± 0.1	$-3.3^{+1.1}_{-1.0} \times 10^{-5}$	$0.0^{+0.4}_{-0.4} \times 10^{-5}$	$8.6^{+1.5}_{-1.5}$	$120.4^{+1.5}_{-1.3}$	$77.8^{+1.8}_{-1.4}$
2	10–21	-9.9 ± 0.1	$-2.4^{+1.3}_{-1.2} \times 10^{-5}$	$0.0^{+0.4}_{-0.4} \times 10^{-5}$	$10.6^{+0.5}_{-0.6}$	$46.9^{+3.2}_{-2.4}$	$152.0^{+3.0}_{-2.8}$
3	6–12	-11.9 ± 0.1	$-2.9^{+1.0}_{-1.0} \times 10^{-5}$	$0.0^{+0.4}_{-0.5} \times 10^{-5}$	$8.5^{+1.9}_{-1.9}$	$13.9^{+0.3}_{-0.3}$	$72.1^{+1.6}_{-1.5}$
4	4–12	-6.9 ± 0.2	$4.6^{+1.2}_{-1.2} \times 10^{-5}$	$-0.1^{+0.7}_{-0.6} \times 10^{-5}$	$5.2^{+0.7}_{-1.2}$	$50.3^{+0.6}_{-0.6}$	$221.7^{+0.8}_{-0.9}$
5	3–7	-10.7 ± 0.1	$4.1^{+1.3}_{-1.2} \times 10^{-5}$	$0.0^{+0.7}_{-0.7} \times 10^{-5}$	$4.4^{+0.5}_{-0.6}$	$80.5^{+0.3}_{-0.3}$	$218.3^{+0.3}_{-0.4}$
6	2–4	-11.3 ± 0.1	$-4.2^{+1.5}_{-1.5} \times 10^{-5}$	$0.1^{+1.0}_{-1.0} \times 10^{-5}$	$2.7^{+0.4}_{-0.5}$	$62.6^{+0.4}_{-0.3}$	$146.1^{+0.3}_{-0.4}$
7	2–4	-6.6 ± 0.1	$6.5^{+1.5}_{-1.4} \times 10^{-5}$	$-0.1^{+0.9}_{-0.8} \times 10^{-5}$	$3.0^{+0.3}_{-0.3}$	$62.6^{+0.2}_{-0.2}$	$142.0^{+0.3}_{-0.3}$
8	2–4	-8.2 ± 0.1	$6.1^{+1.7}_{-1.7} \times 10^{-5}$	$-0.2^{+1.1}_{-1.1} \times 10^{-5}$	$2.5^{+0.2}_{-0.3}$	$111.5^{+0.2}_{-0.2}$	$69.7^{+0.2}_{-0.2}$
9	3–6	-9.7 ± 0.1	$5.0^{+1.3}_{-1.3} \times 10^{-5}$	$0.0^{+0.6}_{-0.6} \times 10^{-5}$	$3.7^{+0.5}_{-0.4}$	$36.4^{+0.3}_{-0.3}$	$131.0^{+0.5}_{-0.5}$
10	2–4	-9.9 ± 0.2	$5.6^{+1.6}_{-1.5} \times 10^{-5}$	$-0.8^{+0.9}_{-0.9} \times 10^{-5}$	$2.6^{+0.2}_{-0.4}$	$160.2^{+0.2}_{-0.2}$	$235.9^{+0.6}_{-0.7}$

TABLE VI. The input and marginalized 68% C.L. parameter bounds for the representative sample of simulated 10° collisions. All simulated collisions are located at $\theta_0 = 57.7^\circ$, $\phi_0 = 99.2^\circ$. Hatted quantities are estimates, and unhatted quantities are inputs. No errors are quoted for the estimated central positions and radii for the cases where there was an extremely strong detection. This is due to the pixelization of the map: variations in collision- center coordinates or radius of much less than a pixel's width will not affect the pixelated template and hence will not affect the likelihood. Angular quantities are quoted in degrees.

Example	z_0	z_{crit}	$\ln\rho_b$	\hat{z}_0	\hat{z}_{crit}	$\hat{\theta}_{\text{crit}}$	$\hat{\theta}_0$	$\hat{\phi}_0$
1	5.0×10^{-5}	-5.0×10^{-5}	119.8 ± 0.3	$5.1^{+1.0}_{-1.0} \times 10^{-5}$	$-5.0^{+0.3}_{-0.3} \times 10^{-5}$	10.0	57.7	99.2
2	1.0×10^{-5}	-5.6×10^{-5}	136.0 ± 0.2	$2.3^{+1.1}_{-1.1} \times 10^{-5}$	$-5.4^{+0.3}_{-0.3} \times 10^{-5}$	10.0	57.7	99.2
3	1.0×10^{-4}	-1.0×10^{-6}	28.9 ± 0.3	$9.4^{+0.6}_{-0.4} \times 10^{-5}$	$-0.2^{+0.8}_{-0.8} \times 10^{-5}$	$10.0^{+0.9}_{-0.9}$	$57.7^{+0.5}_{-0.7}$	$99.8^{+0.5}_{-0.6}$
4	3.2×10^{-5}	-3.2×10^{-5}	9.0 ± 0.3	$6.4^{+1.0}_{-1.1} \times 10^{-5}$	$-2.0^{+0.3}_{-0.3} \times 10^{-5}$	$8.92^{+0.09}_{-0.02}$	$57.01^{+0.03}_{-0.03}$	$100.27^{+0.04}_{-0.04}$
5	1.0×10^{-5}	-3.2×10^{-5}	-1.0 ± 0.3	$4.2^{+0.6}_{-0.6} \times 10^{-5}$	$-1.2^{+0.3}_{-0.4} \times 10^{-5}$	$8.6^{+0.4}_{-0.9}$	$57.3^{+0.1}_{-0.3}$	$100.1^{+0.3}_{-0.1}$
6	1.0×10^{-5}	-1.8×10^{-5}	-7.2 ± 0.2	$4.8^{+1.2}_{-1.1} \times 10^{-5}$	$-0.1^{+0.7}_{-0.7} \times 10^{-5}$	$6.5^{+0.4}_{-0.7}$	$58.2^{+0.3}_{-0.3}$	$99.9^{+0.3}_{-0.3}$

and we therefore correctly conclude that the data from the end-to-end simulation does *not* warrant augmenting Λ CDM with bubble collisions.

These results from the end-to-end simulation yield quantitative information on the degree to which systematics and foregrounds could mimic the signal from a bubble collision. Reassuringly, no features yield evidence ratios greater than 1. To be distinguishable from systematics and foregrounds, we require the evidence ratios that we find for any feature to at least exceed the evidence ratios found in the end-to-end simulation at similar needlet frequencies.

4. Analysis of bubble-collision simulations

The long processing time, even for a single map, prohibits us from running the Bayesian parameter estimation and model selection analysis on the full set of bubble-collision simulations. We therefore choose a small number of representative examples from the set of simulated collisions passing the needlet significance threshold (drawn from the exclusion and sensitivity regions of Fig. 11). Six 10° collision simulations were chosen to sample distinct areas of our parameter space, specifically collisions with (1) a large central amplitude and edge, (2) a small central

amplitude but large edge, (3) a large central amplitude but small edge, (4) a medium central amplitude and medium edge, (5) a small central amplitude and medium edge, and (6) a small central amplitude and small edge. The first two collisions lie in the CHT exclusion zone, the third in the needlets exclusion zone, and the others in the sensitivity region. All collisions were placed at the low-noise location to maximize the chance of a detection.

The results of the Bayesian analysis of the collision simulations are displayed in Table VI. The first example corresponds to the collision in Fig. 10 and is clearly a highly significant detection with an evidence ratio of $\ln\rho_b \approx 120$. The second example is, again, an extremely clear detection, with $\ln\rho_b \approx 136$. While the evidence for the third example is numerically lower than for the strongly discontinuous cases, at $\ln\rho_b \approx 29$, it is again a conclusive detection. In each of these examples, the full-sky posterior assuming $N_b = 1$ [which is well approximated by Eq. (A18)], would prefer models with bubble collisions over a wide range of \bar{N}_s .

For the collisions sampled in the sensitivity region, the maximum needlet significance recorded in each case was around $S \approx 4$, which is on the upper end of the

significances found in the end-to-end simulation: Similar features in the data are passed to the Bayesian analysis section of the pipeline. The evidence ratios were found to be $\ln \rho_b \approx 9$ for the collision with a medium central amplitude and a medium edge, $\ln \rho_b \approx -1$ for the collision with the medium edge but a smaller central amplitude, and $\ln \rho_b \approx -7$ for the collision with a small edge. Since the latter two templates differ only by the value of z_{crit} , this is further proof that the presence of a detectable causal boundary increases our ability to distinguish a collision. In addition, comparing examples 3 and 4, it can be seen that changing the central amplitude by a bit more than a factor of 2 yields an evidence ratio that is orders of magnitude larger. Apparently, there are rather sharply defined limits of detection. For these marginal cases, the parameter uncertainties are significantly underestimated due to the relative strength of the CMB and noise. Only in the case of the collision with a medium amplitude and medium edge could we conclude that models with bubble collisions are preferred over those without over a modest range in \bar{N}_s .

In conclusion, for the simulated collisions in the needlet and CHT exclusion regions of parameter space, our pipeline can clearly determine that the bubble-collision hypothesis is favored for a variety of \bar{N}_s . In the other cases we have studied, where the collision lies in the needlet sensitivity region, the conclusion is less clear. The evidence ratios are higher than most of those in the end-to-end simulation, but not much greater. They are also small in magnitude and therefore do not yield full-sky posteriors that favor the bubble-collision hypothesis. Thus, while we might rule these features out as being systematics or foregrounds, better data would be needed to definitively establish the bubble-collision hypothesis. Furthermore, the bounds on parameter values in detections associated with the sensitivity regions of parameter space should be regarded as rough estimates only. Note also that since the data sets we consider for each blob are restricted to patches of the sky smaller than 11° , the gain in sensitivity that arises from the existence of a circular temperature discontinuity will not be present for modulations with $\theta_{\text{crit}} \gtrsim 11^\circ$. For large features with an edge, the evidence ratios we obtain would therefore be an *underestimate*.

D. Summary of the analysis pipeline and conditions for claiming a detection

We now summarize the analysis pipeline and the interpretation of its outputs. First, the analysis pipeline segments the sky into blobs, each of which corresponds to a region which, for some needlet type and frequency, passes our needlet significance threshold. A specific region of the temperature map can be covered by multiple blobs if there is a response for multiple needlet types/frequencies at the same location. The output of this first step in our pipeline is the location, size, and maximum significance associated with each blob. The edge-detection step of our pipeline

finds the CHT score as a function of assumed circle size and pixel. If there is a clearly peaked global maximum for the CHT score, this can be processed into the most likely circle center and angular scale. In parallel, we calculate the marginalized constraints on the parameters $\{z_0, z_{\text{crit}}, \theta_{\text{crit}}, \theta_0, \phi_0\}$ and Bayesian evidence ratio ρ_b for each feature. These evidence ratios are then used to construct the full-sky posterior $\text{Pr}(\bar{N}_s | N_b, f_{\text{sky}})$ [Eq. (28)] which is a function of \bar{N}_s .

The posterior allows us to put constraints on the possible values of \bar{N}_s that are consistent with the data. Comparing the value of the posterior at $\bar{N}_s = 0$ and some particular value of \bar{N}_s specifies whether or not the Λ CDM model should be superseded by a model that also predicts on average \bar{N}_s bubble collisions. If a large ratio of the posteriors is obtained, a conclusive detection of the bubble-collision hypothesis can be claimed (provided a model that predicts an appropriate value of \bar{N}_s exists). A clear peak in the CHT score would indicate the presence of a circular temperature discontinuity in the CMB. This is a clear signature of bubble collisions and would be nearly conclusive evidence for the eternal inflation scenario. We have found using simulations that a clear edge also yields large evidence ratios, indicating that these two tests are complementary. However, an edge is not necessary to verify the bubble-collision hypothesis. There is a clear expectation obtained from the end-to-end simulation for the contribution from false detections due to systematics and foregrounds: the absence of a clear peak in the CHT score, and evidence ratios for each blob not exceeding $\ln \rho_b \sim -6.6$ at detectable scales.

VI. ANALYSIS OF THE WMAP 7-YEAR DATA

Our analysis of the W-band WMAP 7-year foreground-reduced temperature map with the KQ75 mask produces a total of 38 blobs passing our needlet sensitivity thresholds. These blobs can be grouped into 15 distinct features, four of which either intersect or are within a few pixels of the main Galactic cut; these features are assumed to be responses to the mask, and we do not consider them further. The properties of the blobs belonging to the 11 remaining features are given in Table VII.

A number of these features have been noted previously. Feature 2 is at the same position as the famous Cold Spot [65]. In addition, features 1 and 3 are coincident with the most significant hot spots identified in the needlet analysis of Ref. [48]. The number of features we have found is consistent with the results from the WMAP end-to-end simulation, although the simulation does not contain as many high-significance features at low j . In addition, the most significant features in the WMAP 7-year data generate responses from multiple needlet types at multiple frequencies (e.g., the Cold Spot is picked out by seven needlet frequencies), whereas features in the end-to-end simulation tend to be highlighted only by a single needlet.

TABLE VII. Features found by the needlet transform in the WMAP 7-year data. Features 1 and 3 correspond to the hot spots found in Ref. [48]; feature 2 is the Cold Spot [65]. Angular quantities are reported in degrees.

Feature	Blob	B	j	θ_0	ϕ_0	Blob radius	S
1	1	2.5	2	140.1	173.7	4.5	3.76
1	2	1.8	3	140.9	174.4	5.3	3.48
2	1	2.5	3	147.8	209.5	4.1	4.49
2	2	1.8	4	148.2	207.7	5.5	4.55
2	3	1.8	5	148.5	210.2	1.4	3.37
2	4	1.4	7	147.8	209.5	2.5	3.81
2	5	1.4	8	147.8	209.5	4.1	4.58
2	6	1.4	9	147.4	208.1	2.8	4.30
2	7	1.4	10	146.6	207.5	1.3	3.74
3	1	2.5	3	123.2	321.3	2.5	4.09
3	2	1.8	4	122.8	322.4	4.9	3.82
3	3	1.4	7	122.8	321.0	1.9	3.59
3	4	1.4	8	122.8	321.0	3.2	4.01
3	5	1.4	9	122.8	321.0	2.7	4.30
3	6	1.4	10	122.4	320.6	1.5	3.78
4	1	2.5	4	145.1	33.0	0.9	4.20
4	2	1.8	6	145.5	32.4	0.7	3.72
4	3	1.4	11	145.1	33.0	0.9	3.95
5	1	1.8	5	32.2	74.0	1.2	3.41
6	1	1.8	5	128.7	91.1	1.2	3.37
7	1	1.8	5	169.8	181.6	2.3	3.82
7	2	1.8	6	169.0	187.5	0.8	3.76
7	3	1.4	10	169.4	184.7	1.5	4.12
7	4	1.4	11	168.7	187.3	1.1	4.07
8	1	1.8	6	57.9	115.7	0.7	3.78
9	1	1.8	7	152.3	241.8	0.6	4.12
10	1	1.4	10	167.2	268.7	1.0	3.99
10	2	1.4	11	166.8	271.3	1.0	4.09
11	1	1.4	11	115.0	22.5	0.5	3.80
11	2	1.4	12	114.6	22.1	0.5	4.32

Interestingly, 9 of the 11 features identified as significant are in the Southern Galactic hemisphere.

The CHT scores do not have a clear peak at any angular scale or location for any of the detected features. Indeed, the detailed outputs for the data are completely consistent with those obtained for the end-to-end simulation. The largest CHT peak found in the data is shown in Fig. 20 (which should be compared to the most peaklike feature found in the end-to-end simulation, shown in Fig. 16). We therefore find no evidence for circular temperature discontinuities in the WMAP 7-year data and can rule out bubble collisions in the CHT exclusion region defined by simulated collisions shown in Fig. 17.

The marginalized parameter constraints and local evidence ratio for each of the features is recorded in Table VIII.¹¹ Features 2, 3, 7, and 10 have evidence ratios

¹¹Since we are limited to patches of the sky 11° in radius, the evidence ratios we have obtained for features whose θ_{crit} priors extend beyond $\sim 11^\circ$ will be underestimated if a weak edge exists outside the patch of sky considered.

significantly larger than those found in the collision-free end-to-end simulation ($\ln\rho_b \sim -6.6$), specifically -4.6 , -4.1 , -5.4 , and -3.8 , respectively. Assuming $N_b = 4$, and using these values for the local evidence ratios in Eq. (28), we find that the posterior is maximized at $\tilde{N}_s = 0$, and we can constrain $\tilde{N}_s < 1.6$ at 68% C.L. One would need roughly $\ln\rho_b \sim -1$ for each of the four features to prefer the bubble-collision hypothesis for any value of \tilde{N}_s . Therefore, the WMAP 7-year data do not warrant adding bubble collisions to Λ CDM.

Although the local evidence ratios found for the WMAP 7-year data were not large enough to yield support for the bubble-collision hypothesis, they are about an order of magnitude larger than what was expected from systematics based on the end-to-end simulation. The analysis of future data sets may increase the significance of these blobs if they are indications of bubble collisions, or else they will decrease in significance if they are not; in any case they are the most significant features on our sky and thus take priority in being further investigated with better data. Thus, we now examine these four most significant features




FIG. 20 (color online). The clearest peak found during the edge-detection analysis of the WMAP data. The contrasts in scores as a function of position (left panel) and radius (right panel) are comparable to those obtained in the analysis of the end-to-end simulation (Fig. 16), and greatly reduced compared to the collision example (Figs. 14 and 15).

in more detail. The location of each of the four features on the sky is shown in Fig. 21. A closer view of each feature is shown in Fig. 22, along with plots of the needlet significances S triggering the Bayesian analysis step, collision templates for the marginalized parameter constraints found in each case, and the CMB sky as it would appear with these template contributions removed.

To confirm that these features are not due to residual foregrounds, we have also applied our suite of needlet transforms to the WMAP 7-year Q (41 GHz) and V (61 GHz) band foreground-reduced maps. Taking all of the needlets which generate a significant response for the four most significant features, we calculate the average

of the needlet coefficients within the regions described by the estimated bubble templates. The results are plotted in Fig. 23. We show, for each blob forming part of a feature, the W-band-normalized needlet-coefficient averages given by

$$\Delta\beta_{jk,Q/V} = \frac{\bar{\beta}_{jk,Q/V} - \bar{\beta}_{jk,W}}{\bar{\beta}_{jk,W}}, \quad (34)$$

where $\bar{\beta}_{jk,Q/V/W}$ is the pixel-averaged needlet-coefficient value in a given WMAP frequency band. The plots are consistent with no change in the strength of the signal with frequency, suggesting that the features are not due to foreground contamination.

TABLE VIII. Results of the Bayesian parameter estimation and model selection analysis for the WMAP 7-year data. Reported error bars are at 68% C.L. Angular quantities are reported in degrees. The ranges of θ_{crit} are determined from the needlet response (see Table I). By computational necessity, the evidence integral is truncated at 11° . Hence, an evidence ratio for feature 1 could not be calculated as its θ_{crit} range lies entirely beyond this upper bound. The angular positions $\{\theta_0, \phi_0\}$ can be related to Galactic coordinates through longitude $l_0 = \phi_0$ and latitude $b_0 = 90^\circ - \theta_0$.

Feature	θ_{crit} range	$\ln\rho_b$	z_0	z_{crit}	θ_{crit}	θ_0	ϕ_0
1	12–38						
2	4–21	-4.6 ± 0.2	$-4.8^{+1.4}_{-1.4} \times 10^{-5}$	$-0.1^{+0.7}_{-0.7} \times 10^{-5}$	$6.4^{+1.7}_{-1.1}$	$147.3^{+0.7}_{-0.6}$	$208.0^{+1.5}_{-1.4}$
3	4–21	-4.1 ± 0.2	$5.2^{+1.2}_{-1.2} \times 10^{-5}$	$0.1^{+0.5}_{-0.5} \times 10^{-5}$	$6.5^{+0.7}_{-0.7}$	$123.0^{+0.7}_{-0.7}$	$320.8^{+1.0}_{-1.0}$
4	2–7	-7.3 ± 0.1	$-5.2^{+1.3}_{-1.3} \times 10^{-5}$	$0.0^{+0.9}_{-0.9} \times 10^{-5}$	$3.2^{+0.8}_{-0.8}$	$145.3^{+0.5}_{-0.4}$	$32.8^{+0.9}_{-1.1}$
5	6–12	-9.2 ± 0.2	$3.7^{+1.2}_{-1.1} \times 10^{-5}$	$-0.1^{+0.4}_{-0.4} \times 10^{-5}$	$6.9^{+0.5}_{-0.9}$	$32.6^{+0.6}_{-0.6}$	$74.3^{+1.3}_{-1.2}$
6	6–12	-9.7 ± 0.1	$3.0^{+1.1}_{-1.2} \times 10^{-5}$	$-0.1^{+0.4}_{-0.4} \times 10^{-5}$	$7.9^{+2.2}_{-1.9}$	$128.6^{+0.7}_{-0.7}$	$91.8^{+1.2}_{-1.3}$
7	3–12	-5.4 ± 0.2	$-5.0^{+1.3}_{-1.4} \times 10^{-5}$	$0.0^{+0.6}_{-0.6} \times 10^{-5}$	$4.6^{+0.9}_{-1.1}$	$169.0^{+0.6}_{-0.6}$	$185.7^{+3.9}_{-3.9}$
8	3–7	-8.6 ± 0.1	$4.4^{+1.3}_{-1.2} \times 10^{-5}$	$0.0^{+0.7}_{-0.7} \times 10^{-5}$	$4.4^{+0.6}_{-0.6}$	$58.0^{+0.4}_{-0.4}$	$115.7^{+0.5}_{-0.6}$
9	2–4	-9.0 ± 0.2	$-6.0^{+1.6}_{-1.6} \times 10^{-5}$	$0.0^{+1.0}_{-1.0} \times 10^{-5}$	$2.3^{+0.2}_{-0.3}$	$152.1^{+0.2}_{-0.2}$	$241.9^{+0.4}_{-0.4}$
10	3–8	-3.8 ± 0.2	$-6.1^{+1.3}_{-1.3} \times 10^{-5}$	$0.1^{+0.9}_{-0.8} \times 10^{-5}$	$4.2^{+0.4}_{-0.4}$	$167.2^{+0.3}_{-0.3}$	$269.1^{+1.6}_{-1.4}$
11	2–6	-8.1 ± 0.1	$-5.9^{+1.5}_{-1.6} \times 10^{-5}$	$0.2^{+0.8}_{-0.8} \times 10^{-5}$	$2.5^{+0.4}_{-0.5}$	$114.9^{+0.4}_{-0.4}$	$22.4^{+0.4}_{-0.4}$

VII. CONCLUSIONS AND OUTLOOK

An exciting opportunity to confront the eternal inflation scenario with experiment lies in the observation of collisions between other bubble universes and our own. In this paper, we have described an algorithm to search for the imprint of bubble collisions on the cosmic microwave background and applied it to the WMAP 7-year data. Our search algorithm targets the generic signatures expected from bubble collisions: azimuthal symmetry, long-wavelength modulation of the temperature confined to disks on the sky, and circular temperature discontinuities. For this reason, we expect our analysis to be fairly robust under changing assumptions about the underlying theory, which is presently rather poorly understood.




FIG. 21 (color online). Full-sky map showing the positions and sizes of the four features with largest evidence ratios, along side the 7-year KQ75 sky cut. The features are plotted clockwise from the blob closest to the Galactic mask: feature 3 (red), feature 2 (orange), feature 7 (light blue), and feature 10 (light green).

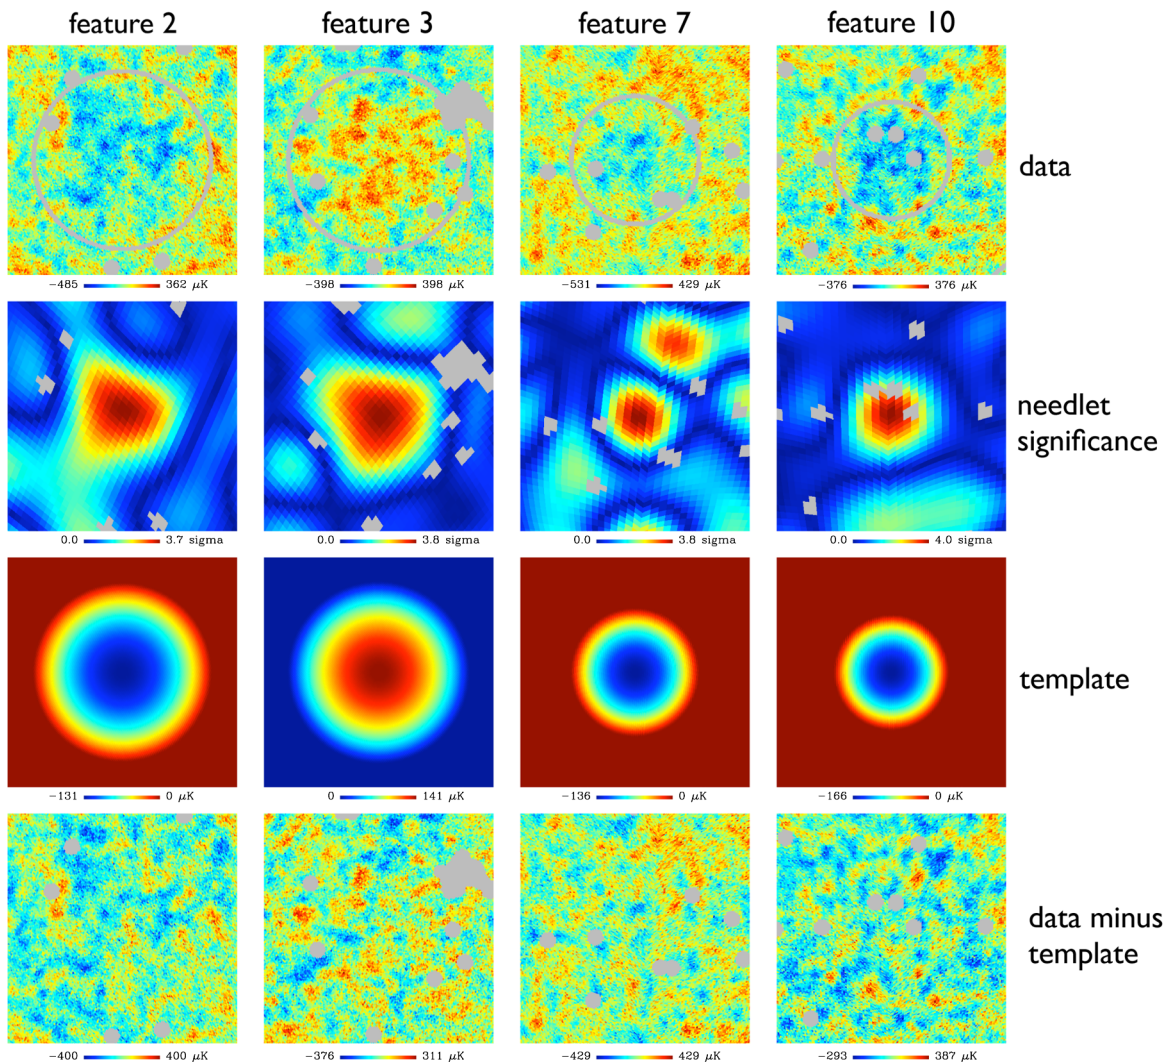


FIG. 22 (color online). Maps of the four features with largest evidence ratios. The top row shows the W-band temperature map in the locality of the four features, masked with the KQ75 mask. Overlaid are circles indicating the estimated position and angular scale found in each case. The second row contains plots of the masked needlet significances for the needlets whose θ_{crit} priors produced the largest evidence ratios. These plots appear pixelated as the blob-detection step is carried out at reduced resolution. The third row shows the bubble-collision templates corresponding to the estimated model parameters; these templates are subtracted from the W-band data in the fourth row. The width of each plot is $\sim 16.7^\circ$.

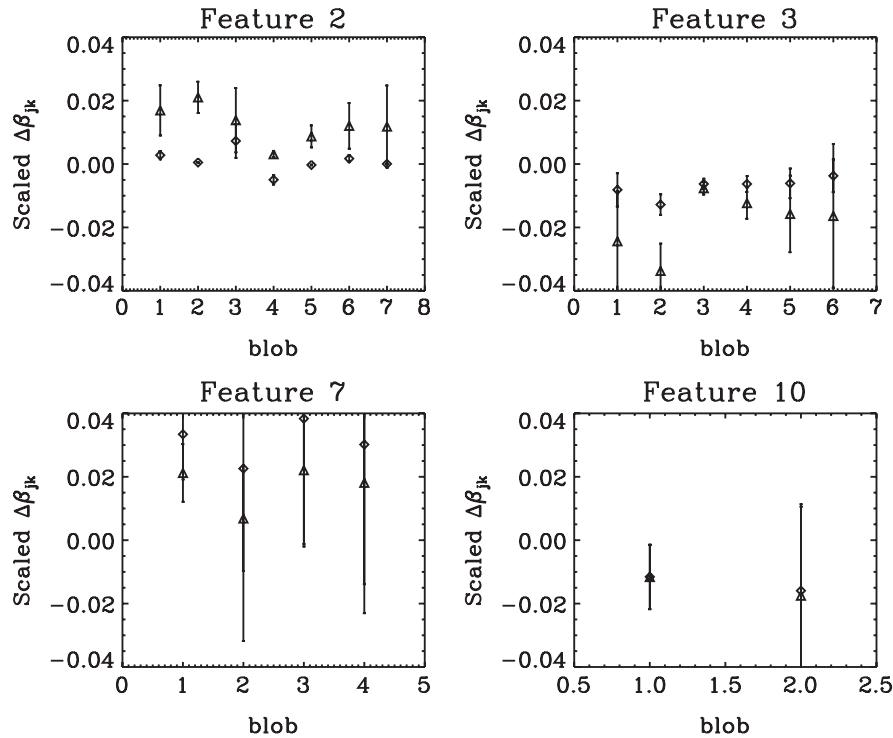


FIG. 23. WMAP channel frequency dependence of the features highlighted by our pipeline. The W-band-normalized difference in pixel-averaged needlet coefficients between the Q and W bands (triangles) and V and W bands (diamonds) are plotted for each blob making up a feature, as highlighted by the full suite of needlet transforms.

The analysis pipeline we have developed takes a two-pronged approach, applied in parallel. The first uses heuristic techniques to test for the presence of features specific to bubble collisions. The second is a fully Bayesian algorithm for the general problem of non-Gaussian source detection, implemented as a patchwise approximation to the full-sky model selection and parameter estimation problem. The data set is segmented in a completely automated way, allowing us to avoid *a posteriori* selection effects associated with choosing the most “interesting” features on the CMB sky by hand. The algorithm is tested and thresholds at each step are calibrated using extensive simulations and then frozen before ever looking at the data, to follow as much as possible the philosophy of a blind analysis. Candidate collisions are identified from an input temperature map based on the response to a suite of needlet transforms (calibrated using simulations with and without bubble collisions), and grouped into blobs. These blobs are scrutinized for circular temperature discontinuities using an edge-detection algorithm. The quantitative significance of an edge is characterized using the CHT. The blobs are also used to construct an approximation to the full-sky Bayesian parameter estimation and model selection problem for bubble collisions. The posterior probability distribution over the expectation value for the number of detectable collisions, \bar{N}_s , is then obtained. This allows us to quantify which of the two models—a theory which predicts on average \bar{N}_s bubble-collision signatures

described by temperature modulations of the form given in Eq. (4), or else the standard model (specified by $\bar{N}_s = 0$) with CMB plus realistic noise and beam effects—better explains the data.

Applying our analysis pipeline to simulations, we have found that a circular temperature discontinuity at the causal boundary is a clear signature of bubble collisions.¹² Although our analysis can identify collisions without temperature discontinuities, their presence greatly increases our ability to make a conclusive detection. Both the edge-detection and Bayesian model selection steps have the ability to identify a causal boundary in the patches of the sky that are highlighted as candidate collisions by the blob-detection step of our analysis pipeline. We have found no evidence for circular temperature discontinuities in the WMAP 7-year data using either method. Based on our analysis of simulations, this allows us to rule out the presence of collisions in the exclusion region of Fig. 17. For collisions larger than $\theta_{\text{crit}} \geq 10^\circ$, this corresponds to $10^5 |z_{\text{crit}}| \lesssim 3\text{--}6$ for the amplitude of the circular temperature discontinuity defined in Eq. (4). For collisions on smaller scales, the CHT step loses sensitivity due to the proliferation of degree-scale blobs in the background CMB fluctuations.

¹²The observational detection of a circular temperature discontinuity is so unlikely to arise spuriously that it provides conclusive evidence of a detection.

The posterior evaluated using the WMAP 7-year data is maximized at $\bar{N}_s = 0$, and constrains $\bar{N}_s < 1.6$ at 68% confidence. We therefore conclude that this data set does not favor the bubble-collision hypothesis for any value of \bar{N}_s . In light of this null detection, comparing with the simulated bubble collisions, we can constrain the central amplitude of the temperature modulation caused by the collision [defined in Eq. (4)] to be $z_0 \lesssim 1 \times 10^{-4}$ over the range of scales $\theta_{\text{crit}} \gtrsim 5^\circ$ we have simulated. If the collision is described by a single super-Hubble wavelength mode confined to a disk on the sky, from Eq. (12) we can use these bounds (with the largest collision size we have simulated $\theta_{\text{crit}} = 25^\circ$) to constrain $\Omega_k^{1/2} \Phi(0) \lesssim 7 \times 10^{-4}$ [where Ω_k is the present component in curvature and $\Phi(0)$ is the initial magnitude of the Newtonian potential caused by the collision]. More generally, Eq. (13) bounds the nucleation rate of bubbles in our parent vacuum, provided gravitational waves and negative curvature are observed with future experiments.

Although we have obtained a null result, our analysis pipeline has identified four features in the WMAP 7-year data that have Bayesian evidence ratios that are significantly larger than expected for false detections from an end-to-end simulation of the WMAP experiment. Two of these features (features 2 and 3) have been noted in previous literature. Feature 2 corresponds to the WMAP Cold Spot [65] (see Ref. [66] for a review of its properties), and feature 3 was identified using standard needlets in Ref. [48]. All four features are far from the Galactic cut of the KQ75 7-year mask (see Fig. 21), and none appear to be responses to the point source components of the mask (see Fig. 22). We have confirmed that the signal in each case is not strongly dependent on the frequency band used (see Fig. 23), providing evidence that these features are not due to astrophysical foregrounds. A number of analyses, most recently the redshift analysis of Ref. [67], suggest that the Cold Spot is primordial and not associated with the integrated Sachs-Wolfe effect of a large void. Further studies of the other three features would be needed to confirm that they are truly primordial.

Our ability to detect bubble collisions will improve greatly with data from the Planck satellite. Decreased instrumental noise will enlarge the exclusion and sensitivity regions in parameter space for the needlet step of the analysis, as evidenced by our ability to detect more simulated collisions in low-noise regions of the WMAP data. The threefold increase in resolution will greatly improve our ability to detect circular edges. In addition, the polarization data from Planck will be of sufficient resolution to look for complementary signatures of bubble collisions [26,68]. Such an analysis should be able to confirm if the features we have identified are in fact bubble collisions.

It is also important to determine if other theories predicting azimuthally symmetric features in the CMB [60,69–71] are better fits to the data. The blob and

edge-detection steps in our analysis pipeline are sensitive to a variety of possible signatures, and given a model, the Bayesian model comparison step could be easily tailored to accommodate different forms of the temperature modulation. Because our pipeline is automated, we can compare the evidence ratios obtained for different models to decide which is a better fit, without recourse to *a posteriori* choices of which features to analyze.

In conclusion, we have presented a powerful algorithm for analyzing CMB data for signatures of bubble collisions. Applying this pipeline to the WMAP 7-year data, we have constrained the possible parameter space of bubble collisions, as well as identifying interesting candidate signatures in the data for further investigation. Future data from the Planck experiment will allow us to greatly improve on these results. If confirmed, the presence of bubble collisions in the CMB would provide extraordinary insight into the origins of our Universe.

ACKNOWLEDGMENTS

We are very grateful to Eiichiro Komatsu and the WMAP Science Team for supplying the end-to-end WMAP simulations used in our null tests. S.M.F. is supported by the Perren Fund. M.C.J. acknowledges support from the Moore Foundation. H.V.P. is supported by Marie Curie Grant No. MIRG-CT-2007-203314 from the European Commission, and by STFC and the Leverhulme Trust. M.C.J. and H.V.P. thank the Aspen Center for Physics, where this project was initiated, for hospitality. H.V.P. and D.J.M. acknowledge the hospitality of the Statistical Frontiers of Astrophysics workshop at IPMU, Tokyo. S.M.F. thanks Filipe Abdalla, Ingo Waldmann, and Michael Hirsch for useful discussions, and Jenny Feeney for proofreading the manuscript. M.C.J. thanks Rebecca Danos for discussions regarding the edge-detection algorithm. H.V.P. thanks Andrew Pontzen for interesting discussions regarding the general problem of CMB anomaly hunting. We acknowledge use of the HEALPIX package and the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. A preprint version of this paper presented only evidence ratios confined to patches. We thank an anonymous referee who encouraged us to develop this algorithm into a full-sky formalism. This calculation is now presented in Appendix and incorporated into our analysis pipeline and results.

APPENDIX: STATISTICAL FORMALISM

1. Posterior

In this Appendix, we discuss how Bayesian parameter estimation and model selection for theories which predict

localized sources can be approximated by a patchwise analysis. Consider astronomical observations covering solid angle $\Omega_{\text{obs}} = 4\pi f_{\text{sky}}$ that are of sufficient depth/resolution to identify sources with a particular range of properties (which can then be deemed “detectable”). Given a theory that predicts an expectation value of \bar{N}_s sources over the whole sky, we want to know both what constraints the available data place on \bar{N}_s and whether the data favor a model which predicts one value of \bar{N}_s over another. All the relevant information is encoded in the posterior distribution $\Pr(\bar{N}_s|\mathbf{d}, f_{\text{sky}})$, where \mathbf{d} are the pixelized flux or temperature measurements (and, optionally, any statistics derived from them). Bayes’s theorem allows the posterior to be written as

$$\Pr(\bar{N}_s|\mathbf{d}, f_{\text{sky}}) = \frac{\Pr(\bar{N}_s) \Pr(\mathbf{d}|\bar{N}_s, f_{\text{sky}})}{\Pr(\mathbf{d}|f_{\text{sky}})}, \quad (\text{A1})$$

where $\Pr(\bar{N}_s)$ is the prior distribution on \bar{N}_s , $\Pr(\mathbf{d}|\bar{N}_s, f_{\text{sky}})$ is the likelihood of getting the observed data given the area of observation and the expected number of sources, and $\Pr(\mathbf{d}|f_{\text{sky}})$ ensures that the posterior is normalized over \bar{N}_s . Constraints on \bar{N}_s can be drawn directly from this normalized posterior; the relative probability of models predicting different values of \bar{N}_s can be found by picking out the posterior at two values of \bar{N}_s :

$$\frac{\Pr(\bar{N}_{s,1}|\mathbf{d}, f_{\text{sky}})}{\Pr(\bar{N}_{s,2}|\mathbf{d}, f_{\text{sky}})} = \frac{\Pr(\bar{N}_{s,1}) \Pr(\mathbf{d}|\bar{N}_{s,1}, f_{\text{sky}})}{\Pr(\bar{N}_{s,2}) \Pr(\mathbf{d}|\bar{N}_{s,2}, f_{\text{sky}})}. \quad (\text{A2})$$

In the absence of a prescriptive theory, it is useful to emphasize the role of the data, which can be done by adopting a flat prior on \bar{N}_s . Further assuming that the data will give an upper limit on \bar{N}_s , it is possible to adopt an improper uniform prior $\Pr(\bar{N}_s) = \Theta(\bar{N}_s)$ without any high- \bar{N}_s cutoff. The resultant posterior has the form

$$\Pr(\mathbf{d}|N_s, f_{\text{sky}}) = \int d\mathbf{m}_1 \cdots d\mathbf{m}_{N_s} \prod_{s=1}^{N_s} \Pr(\mathbf{m}_s) \frac{1}{(2\pi)^{N_{\text{px}}/2} |\mathbf{C}|} e^{-[\mathbf{d}-\mathbf{t}(\mathbf{m}_1)\cdots\mathbf{t}(\mathbf{m}_{N_s})]\mathbf{C}^{-1}[\mathbf{d}-\mathbf{t}(\mathbf{m}_1)\cdots\mathbf{t}(\mathbf{m}_{N_s})]^T/2}, \quad (\text{A6})$$

where $\mathbf{t}(\mathbf{m})$ is the data template that would result from a source whose position and profile/scale are defined by the model parameters \mathbf{m} , $\Pr(\mathbf{m}_s)$ is the prior distribution of source parameters for detectable sources, and \mathbf{C} is the pixel-pixel covariance matrix of the nonsource noise (which could include contributions that are considered signal in other contexts, such as the CMB).

Evaluating the full sum in Eq. (A5) is not always practical or even possible. In addition, the evaluation of individual terms in this sum will be computationally limited by the size of the covariance matrix \mathbf{C} and the cost of performing the integral over model parameters for each template. However, it is possible to circumvent these problems and estimate the posterior Eq. (A5) if one knows in advance some of the properties of the integrand in Eq. (A6).

$$\Pr(\bar{N}_s|\mathbf{d}, f_{\text{sky}}) \propto \Theta(\bar{N}_s) \Pr(\mathbf{d}|\bar{N}_s, f_{\text{sky}}), \quad (\text{A3})$$

up to a normalization constant that depends on the data and f_{sky} but not on \bar{N}_s .

In general \bar{N}_s is not directly measurable, even for perfect data, because the number of sources present in the observable sky, N_s , is the realization of a Poisson-like process (of mean $f_{\text{sky}}\bar{N}_s$). The possibility that N_s is itself subject to some uncertainty (e.g., due to noisy data or confusion problems) can be incorporated by marginalizing over N_s to give

$$\begin{aligned} \Pr(\mathbf{d}|\bar{N}_s, f_{\text{sky}}) &= \sum_{N_s=0}^{\infty} \Pr(N_s|\bar{N}_s, f_{\text{sky}}) \Pr(\mathbf{d}|N_s, f_{\text{sky}}) \\ &= \sum_{N_s=0}^{\infty} \frac{(f_{\text{sky}}\bar{N}_s)^{N_s} e^{-f_{\text{sky}}\bar{N}_s}}{N_s!} \Pr(\mathbf{d}|N_s, f_{\text{sky}}), \end{aligned} \quad (\text{A4})$$

where the second formula explicitly assumes that the number of observable sources is drawn from a Poisson process. Inserting this second expression into Eq. (A3) then gives

$$\Pr(\bar{N}_s|\mathbf{d}, f_{\text{sky}}) \propto \Theta(\bar{N}_s) e^{-f_{\text{sky}}\bar{N}_s} \sum_{N_s=0}^{\infty} \frac{(f_{\text{sky}}\bar{N}_s)^{N_s}}{N_s!} \Pr(\mathbf{d}|N_s, f_{\text{sky}}). \quad (\text{A5})$$

The form of the likelihood $\Pr(\mathbf{d}|N_s, f_{\text{sky}})$ is treated largely in abstract here, with the specific details of the likelihood calculation for the bubble-collision hypothesis given WMAP 7-year data described in Sec. V C. Assuming the measurements take the form of flux/counts at different positions on the sky (as for a CMB experiment) and that they are subject to (possibly correlated) Gaussian noise, the likelihood would have the form

To see how this works, assume that one has located a set of N_b blobs on the sky that are candidate sources. Segment the sky into $N_r = N_b + 1$ regions consisting of those containing blobs and the rest of the sky. Given N_b , we can now evaluate Eq. (A5) term by term. The likelihood in the first term, for $N_s = 0$, is simply given by

$$\Pr(\mathbf{d}|0, f_{\text{sky}}) = \frac{1}{(2\pi)^{N_{\text{px}}/2} |\mathbf{C}|} e^{-\mathbf{d}\mathbf{C}^{-1}\mathbf{d}^T/2}, \quad (\text{A7})$$

which is the likelihood for the null hypothesis with no sources. Moving on to the $N_s = 1$ term, we first expand the integral over source positions to cover each of the N_r regions:

$$\begin{aligned} \Pr(\mathbf{d}|1, f_{\text{sky}}) &= \sum_{r=1}^{N_r} \int_{\text{region } r} d\mathbf{m} \Pr(\mathbf{m}) \\ &\times \frac{1}{(2\pi)^{N_{\text{px}}/2} |\mathbf{C}|} e^{-[\mathbf{d}-\mathbf{t}(\mathbf{m})]\mathbf{C}^{-1}[\mathbf{d}-\mathbf{t}(\mathbf{m})]^T/2}. \end{aligned} \quad (\text{A8})$$

We now assume that the blobs containing candidate sources include all of the significant contributions to the integral, and replace N_r in the sum by N_b . This will give us a *lower bound* on the likelihood, even if a number of actual sources are not contained within the blobs defined by the candidate sources. We further assume that sources do not overlap. If the covariance matrix is small enough to invert, we could stop here. However, in cases where the covariance matrix is too large to feasibly invert (as is the case for the WMAP 7 year data), we can make one further approximation:

$$\begin{aligned} &\int_{\text{region } b} d\mathbf{m} \Pr(\mathbf{m}) \frac{1}{(2\pi)^{N_{\text{px}}/2} |\mathbf{C}|} e^{-[\mathbf{d}-\mathbf{t}(\mathbf{m})]\mathbf{C}^{-1}[\mathbf{d}-\mathbf{t}(\mathbf{m})]^T/2} \\ &\simeq \int_{\text{region } b} d\mathbf{m} \Pr(\mathbf{m}) \prod_{r=1}^{N_r} L_r(\mathbf{m}), \end{aligned} \quad (\text{A9})$$

where the product is over N_r disjoint regions on the sky and

$$\Pr(\mathbf{d}|N_s, f_{\text{sky}}) = \begin{cases} 0, & \text{if } N_s > N_b, \\ \sum_{b_1, b_2, \dots, b_{N_s}=1}^{N_b} \left[\prod_{s=1}^{N_s} \rho_{b_s} \prod_{i,j=1}^{N_s} (1 - \delta_{s_i, s_j}) \prod_{r=1}^{N_r} L_r(\mathbf{0}) \right], & \text{if } N_s \leq N_b, \end{cases} \quad (\text{A13})$$

where the combinatorics require some explanation. If there are fewer blobs on the sky than proposed sources then the likelihood is very small: By assumption, the likelihood evaluated outside of a blob is small. If there are at least as many blobs as proposed sources, then the likelihood takes the form of a sum that includes every possible association of the N_s sources with the N_b blobs, provided that no two sources are matched to the same blob. Hence the multiple sum generates all possible combinations of source-blob associations and the product over evidence ratios gives the relevant weightings; the product over delta functions removes the terms in which any two sources are attached to the same blob.

Inserting the likelihood given in Eq. (A13) into Eq. (A5) yields the unnormalized posterior on \bar{N}_s :

$$\begin{aligned} \Pr(\bar{N}_s | \mathbf{d}, f_{\text{sky}}) &\propto \Theta(\bar{N}_s) e^{-f_{\text{sky}} \bar{N}_s} \sum_{N_s=0}^{N_b} \frac{(f_{\text{sky}} \bar{N}_s)^{N_s}}{N_s!} \\ &\times \sum_{b_1, b_2, \dots, b_{N_s}=1}^{N_b} \left[\prod_{s=1}^{N_s} \rho_{b_s} \prod_{i,j=1}^{N_s} (1 - \delta_{s_i, s_j}) \right], \end{aligned} \quad (\text{A14})$$

under the assumption that

$$L_r(\mathbf{m}) = \frac{1}{(2\pi)^{N_{\text{px}}/2} |\mathbf{C}_r|} e^{-[\mathbf{d}_r - \mathbf{t}_r(\mathbf{m})]\mathbf{C}_r^{-1}[\mathbf{d}_r - \mathbf{t}_r(\mathbf{m})]^T/2} \quad (\text{A10})$$

is the contribution to the likelihood of data in region r , defined in terms of the covariance of the pixels in this region, \mathbf{C}_r , the data in this region, \mathbf{d}_r , and the source template in this region, $\mathbf{t}_r(\mathbf{m})$. This is exact for a diagonal covariance, but is only approximate in the case where there are off-diagonal elements. Using the assumption that the integral has a significant contribution only inside the blobs, in the rest of the sky we can replace $\mathbf{t} = \mathbf{0}$. The one-source model likelihood then becomes

$$\Pr(\mathbf{d}|1, f_{\text{sky}}) \simeq \sum_{b=1}^{N_b} \prod_{r=1}^{N_r} L_r(\mathbf{0}) \rho_b, \quad (\text{A11})$$

where

$$\rho_b = \frac{\int_{\text{region } b} d\mathbf{m} \Pr(\mathbf{m}) L_b(\mathbf{m})}{L_b(\mathbf{0})}. \quad (\text{A12})$$

This is the evidence ratio for a single source template centered in region b .

For a general number of blobs and sources, the model likelihood is

$$\Pr(\mathbf{d}|0, f_{\text{sky}}) = \prod_{r=1}^{N_r} L_r(\mathbf{0}), \quad (\text{A15})$$

in which case regions that do not contain a blob are irrelevant for determining the posterior. Equation (A14) is the main result of this calculation, from which all following results can be derived. In the limit of a single isolated observation Eq. (A14) reproduces the Bayesian source detection formalism developed in [72,73].

2. Special cases

a. Perfect data

With infinite, perfect data the number of sources on the sky would be directly determined by counting the N_b blobs in the data and so $\Pr(\mathbf{d}|N_s, f_{\text{sky}}) = \delta_{N_s, N_b}$. The posterior in Eq. (A14) would become

$$\Pr(\bar{N}_s | N_b, f_{\text{sky}}) \propto \Theta(\bar{N}_s) \bar{N}_s^{N_b} e^{-f_{\text{sky}} \bar{N}_s}, \quad (\text{A16})$$

the standard result for constraining a rate variable from a single measurement, modified slightly to account for the fact that the constraint on \bar{N}_s is weakened if $f_{\text{sky}} \ll 1$. In the even more particular case that no blobs were detected in perfect data, the posterior would be

$$\Pr(\bar{N}_s|0, f_{\text{sky}}) = \Theta(\bar{N}_s) f_{\text{sky}} e^{-f_{\text{sky}} \bar{N}_s}. \quad (\text{A17})$$

If a single blob was detected unequivocally then the posterior would be

$$\Pr(\bar{N}_s|1, f_{\text{sky}}) = \Theta(\bar{N}_s) f_{\text{sky}} (f_{\text{sky}} \bar{N}_s) e^{-f_{\text{sky}} \bar{N}_s}. \quad (\text{A18})$$

If two blobs were detected unequivocally then the posterior would be

$$\Pr(\bar{N}_s|2, f_{\text{sky}}) = \Theta(\bar{N}_s) \frac{f_{\text{sky}}}{2} (f_{\text{sky}} \bar{N}_s)^2 e^{-f_{\text{sky}} \bar{N}_s}. \quad (\text{A19})$$

b. No blobs

If there are no identified blobs then $N_b = 0$ and there is no evidence for any sources at all. This is really a weaker constraint than the above situation if the data are perfect, but in the approximation used here the final result is the same. In Eq. (A14) the first sum is truncated at the first term and so

$$\Pr(\bar{N}_s|0, \mathbf{d}, f_{\text{sky}}) = \Theta(\bar{N}_s) f_{\text{sky}} e^{-f_{\text{sky}} \bar{N}_s}, \quad (\text{A20})$$

matching Eq. (A17).

Adopting a flat prior on \bar{N}_s , the posterior probability ratio for a model predicting a generic $\bar{N}_s > 0$ versus one predicting no collisions in this case is given by

$$\frac{\Pr(\bar{N}_s|0, \mathbf{d}, f_{\text{sky}})}{\Pr(0|0, \mathbf{d}, f_{\text{sky}})} = e^{-f_{\text{sky}} \bar{N}_s}. \quad (\text{A21})$$

This is always less than 1, and so as expected, a theory which predicts \bar{N}_s sources on the sky is always disfavored when compared to a theory that predicts no sources on the sky.

c. One blob

Probably the most important simple case is where there is a single identified blob, which might represent a first detection of this class of source. Inserting $N_b = 1$ into Eq. (A14), the sum includes the possibilities of either one source on the (observed) sky or no sources; the posterior evaluates to

$$\Pr(\bar{N}_s|1, \mathbf{d}, f_{\text{sky}}) = \Theta(\bar{N}_s) f_{\text{sky}} e^{-f_{\text{sky}} \bar{N}_s} \frac{1 + f_{\text{sky}} \bar{N}_s \rho_b}{1 + \rho_b}. \quad (\text{A22})$$

In the limit in which the data in this region are much better fit by a source then $\rho_b \gg 1$, and the posterior becomes

$$\Pr(\bar{N}_s|1, \mathbf{d}, f_{\text{sky}}) = \Theta(\bar{N}_s) f_{\text{sky}} (f_{\text{sky}} \bar{N}_s) e^{-f_{\text{sky}} \bar{N}_s}, \quad (\text{A23})$$

which matches (A18) above. Conversely, in the limit that the source is a worse fit to the data (possible given that the source has been forced to be detectable), then $\rho_b \ll 1$ and

$$\Pr(\bar{N}_s|1, \mathbf{d}, f_{\text{sky}}) = \Theta(\bar{N}_s) f_{\text{sky}} e^{-f_{\text{sky}} \bar{N}_s}, \quad (\text{A24})$$

matching Eq. (A17) which was obtained under the assumption that there was no blob in the first place.

Adopting again a flat prior on \bar{N}_s , the posterior probability ratio for a model predicting a generic $\bar{N}_s > 0$ versus the no-bubble case is given by

$$\frac{\Pr(\bar{N}_s|1, \mathbf{d}, f_{\text{sky}})}{\Pr(0|1, \mathbf{d}, f_{\text{sky}})} = e^{-f_{\text{sky}} \bar{N}_s} (1 + f_{\text{sky}} \bar{N}_s \rho_b). \quad (\text{A25})$$

Here, it can be seen that two things are necessary to favor the theory with sources given one detection: $\bar{N}_s \sim \mathcal{O}(1)$ and $\rho_b \gg 1$.

d. Two blobs

If two blobs are identified then the sum in (A13) has three terms, for which the likelihoods are

$$\Pr(2, \mathbf{d}|0, f_{\text{sky}}) = \prod_{r=1}^{N_r} L_r(\mathbf{0}), \quad (\text{A26})$$

$$\Pr(2, \mathbf{d}|1, f_{\text{sky}}) = [\rho_{b_1} + \rho_{b_2}] \prod_{r=1}^{N_r} L_r(\mathbf{0}), \quad (\text{A27})$$

and

$$\Pr(2, \mathbf{d}|2, f_{\text{sky}}) = \rho_{b_1} \rho_{b_2} \prod_{r=1}^{N_r} L_r(\mathbf{0}). \quad (\text{A28})$$

Hence the (unnormalized) posterior is

$$\Pr(\bar{N}_s|2, \mathbf{d}, f_{\text{sky}}) \propto \Theta(\bar{N}_s) e^{-f_{\text{sky}} \bar{N}_s} \{1 + f_{\text{sky}} \bar{N}_s [\rho_{b_1} + \rho_{b_2}] + (f_{\text{sky}} \bar{N}_s)^2 \rho_{b_1} \rho_{b_2}\}. \quad (\text{A29})$$

In the limit in which the evidence for both sources is strong (i.e., $\rho_{b_1} \gg 1$ and $\rho_{b_2} \gg 1$) then the third term in the curly braces dominates and

$$\Pr(\bar{N}_s|2, \mathbf{d}, f_{\text{sky}}) = \Theta(\bar{N}_s) \frac{f_{\text{sky}}}{2} (f_{\text{sky}} \bar{N}_s)^2 e^{-f_{\text{sky}} \bar{N}_s}, \quad (\text{A30})$$

which matches the perfect data case with $N_b = 2$, as expected. In the limit where one blob is a false candidate, but the other yields a strong evidence (e.g., $\rho_{b_1} \gg 1$ and $\rho_{b_2} \ll 1$), then we recover the perfect data case with $N_b = 1$.

- [1] C. L. Bennett *et al.* (WMAP Collaboration), *Astrophys. J.* **583**, 1 (2003).
- [2] J. A. Tauber *et al.*, *A&A* **520**, A1 (2010).
- [3] L. Susskind, in *Universe or Multiverse* (Cambridge University Press, Cambridge, England, 2003).
- [4] A. Aguirre, in *Beyond the Big Bang* (Springer, New York, 2008).
- [5] M. Bucher, A. S. Goldhaber, and N. Turok, *Phys. Rev. D* **52**, 3314 (1995).
- [6] J. R. Gott, *Nature (London)* **295**, 304 (1982).
- [7] J. Garcia-Bellido, arXiv:hep-ph/9803270.
- [8] S. R. Coleman, *Phys. Rev. D* **15**, 2929 (1977).
- [9] S. R. Coleman and F. De Luccia, *Phys. Rev. D* **21**, 3305 (1980).
- [10] A. H. Guth and E. J. Weinberg, *Phys. Rev. D* **23**, 876 (1981).
- [11] A. H. Guth and E. J. Weinberg, *Nucl. Phys.* **B212**, 321 (1983).
- [12] J. R. Gott and T. S. Statler, *Phys. Lett.* **136B**, 157 (1984).
- [13] J. Garriga, A. H. Guth, and A. Vilenkin, *Phys. Rev. D* **76**, 123512 (2007).
- [14] A. Aguirre, M. C. Johnson, and A. Shomer, *Phys. Rev. D* **76**, 063509 (2007).
- [15] S. W. Hawking, I. G. Moss, and J. M. Stewart, *Phys. Rev. D* **26**, 2681 (1982).
- [16] Z.-C. Wu, *Phys. Rev. D* **28**, 1898 (1983).
- [17] A. Aguirre and M. C. Johnson, *Phys. Rev. D* **77**, 123536 (2008).
- [18] A. Aguirre, M. C. Johnson, and M. Tysanner, *Phys. Rev. D* **79**, 123514 (2009).
- [19] S. Chang, M. Kleban, and T. S. Levi, *J. Cosmol. Astropart. Phys.* **04** (2008) 034.
- [20] S. Chang, M. Kleban, and T. S. Levi, *J. Cosmol. Astropart. Phys.* **04** (2009) 025.
- [21] A. Dahlen, *Phys. Rev. D* **81**, 063501 (2010).
- [22] B. Freivogel, M. Kleban, A. Nicolis, and K. Sigurdson, *J. Cosmol. Astropart. Phys.* **08** (2009) 036.
- [23] R. Easther, J. T. Giblin, Jr., L. Hui, and E. A. Lim, *Phys. Rev. D* **80**, 123519 (2009).
- [24] K. Larjo and T. S. Levi, *J. Cosmol. Astropart. Phys.* **08** (2010) 034.
- [25] J. Zhang and Y.-S. Piao, *Phys. Rev. D* **82**, 043507 (2010).
- [26] B. Czech, M. Kleban, K. Larjo, T. S. Levi, and K. Sigurdson, *J. Cosmol. Astropart. Phys.* **12** (2010) 023.
- [27] A. Aguirre and M. C. Johnson, *Rep. Prog. Phys.* **74**, 074901 (2011).
- [28] N. Kaiser and A. Stebbins, *Nature (London)* **310**, 391 (1984).
- [29] A. S. Lo and E. L. Wright, arXiv:astro-ph/0503120.
- [30] R. J. Danos and R. H. Brandenberger, *Int. J. Mod. Phys. D* **19**, 183 (2010).
- [31] S. Amsel, J. Berger, and R. H. Brandenberger, *J. Cosmol. Astropart. Phys.* **04** (2008) 015.
- [32] N. Jarosik *et al.*, *Astrophys. J. Suppl. Ser.* **192**, 14 (2011).
- [33] S. M. Feeney, M. C. Johnson, D. J. Mortlock, and H. V. Peiris, arXiv:1012.1995.
- [34] C. L. Bennett *et al.*, *Astrophys. J. Suppl. Ser.* **192**, 17 (2011).
- [35] Y. Sekino, S. Shenker, and L. Susskind, *Phys. Rev. D* **81**, 123515 (2010).
- [36] E. Komatsu *et al.*, *Astrophys. J. Suppl. Ser.* **192**, 18 (2011).
- [37] B. Freivogel, M. Kleban, M. Rodriguez Martinez, and L. Susskind, *J. High Energy Phys.* **03** (2006) 039.
- [38] A. De Simone and M. P. Salem, *Phys. Rev. D* **81**, 083527 (2010).
- [39] C. Gordon, W. Hu, D. Huterer, and T. Crawford, *Phys. Rev. D* **72**, 103002 (2005).
- [40] J. T. Giblin, Jr., L. Hui, E. A. Lim, and I.-S. Yang, *Phys. Rev. D* **82**, 045019 (2010).
- [41] M. C. Johnson and I.-S. Yang, *Phys. Rev. D* **82**, 065023 (2010).
- [42] M. S. Turner, *Phys. Rev. D* **44**, 3737 (1991).
- [43] A. L. Erickcek, S. M. Carroll, and M. Kamionkowski, *Phys. Rev. D* **78**, 083012 (2008).
- [44] J. P. Zibin and D. Scott, *Phys. Rev. D* **78**, 123529 (2008).
- [45] J. Garcia-Bellido, A. R. Liddle, D. H. Lyth, and D. Wands, *Phys. Rev. D* **52**, 6750 (1995).
- [46] D. Marinucci *et al.*, *Mon. Not. R. Astron. Soc.* **383**, 539 (2007).
- [47] D. Pietrobon, A. Balbi, and D. Marinucci, *Phys. Rev. D* **74**, 043524 (2006).
- [48] D. Pietrobon *et al.*, *Phys. Rev. D* **78**, 103504 (2008).
- [49] P. Baldi, G. Kerkycharian, D. Marinucci, and D. Picard, arXiv:math/0606599.
- [50] F. Guilloux, G. Fay, and J.-F. Cardoso, arXiv:0706.2598.
- [51] S. Scodeller *et al.*, *Astrophys. J.* **733**, 121 (2011).
- [52] J. Canny, *IEEE Trans. Pattern Anal. Mach. Intell. PAMI-8*, 679 (1986).
- [53] C. Kimme, D. Ballard, and J. Sklansky, *Commun. ACM* **18**, 120 (1975).
- [54] F. Feroz, M. P. Hobson, and M. Bridges, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [55] D. Larson *et al.*, *Astrophys. J. Suppl. Ser.* **192**, 16 (2011).
- [56] K. M. Gorski *et al.*, *Astrophys. J.* **622**, 759 (2005).
- [57] A. de Oliveira-Costa and M. Tegmark, *Phys. Rev. D* **74**, 023005 (2006).
- [58] N. E. Groeneboom, L. Ackerman, I. K. Wehus, and H. K. Eriksen, *Astrophys. J.* **722**, 452 (2010).
- [59] J. Hoftuft *et al.*, *Astrophys. J.* **699**, 985 (2009).
- [60] M. Cruz, E. Martínez-González, P. Vielva, J. M. Diego, M. Hobson, and N. Turok, *Mon. Not. R. Astron. Soc.* **390**, 913 (2008).
- [61] A. Pontzen and H. V. Peiris, *Phys. Rev. D* **81**, 103008 (2010).
- [62] R. T. Cox, *Am. J. Phys.* **14**, 1 (1946).
- [63] J. Skilling, in *Proceedings of the 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, AIP Conf. Proc. No. 735 (AIP, New York, 2004), pp. 395–405.
- [64] A. Lewis and S. Bridle, *Phys. Rev. D* **66**, 103511 (2002).
- [65] M. Cruz, E. Martinez-Gonzalez, P. Vielva, and L. Cayon, *Mon. Not. R. Astron. Soc.* **356**, 29 (2005).
- [66] M. Cruz, E. Martinez-Gonzalez, and P. Vielva, *Highlights of Spanish Astrophysics V*, Astrophysics and Space Science Proceedings (Springer-Verlag, Berlin, 2010), p. 275.

- [67] M. N. Bremer, J. Silk, L. J. M. Davies, and M. D. Lehnert, [arXiv:1004.1178](https://arxiv.org/abs/1004.1178) [Mon. Not. R. Astron. Soc. (to be published)].
- [68] C. Dvorkin, H. V. Peiris, and W. Hu, *Phys. Rev. D* **77**, 063008 (2008).
- [69] N. J. Cornish, D. N. Spergel, and G. D. Starkman, *Classical Quantum Gravity* **15**, 2657 (1998).
- [70] N. Afshordi, A. Slosar, and Y. Wang, *J. Cosmol. Astropart. Phys.* **01** (2011) 019.
- [71] E. D. Kovetz, A. Ben-David, and N. Itzhaki, *Astrophys. J.* **724**, 374 (2010).
- [72] M. P. Hobson and C. McLachlan, *Mon. Not. R. Astron. Soc.* **338**, 765 (2003).
- [73] M. P. Hobson, G. Rocha, and R. S. Savage, in *Bayesian Methods in Cosmology*, edited by M. P. Hobson, A. H. Jaffe, A. R. Liddle, P. Mukeherjee, and D. Parkinson (Cambridge University Press, Cambridge, UK, 2010), p. 167.