Optimal estimation of the binned mask-free power spectrum, bispectrum, and trispectrum on the full sky: Scalar edition

Oliver H. E. Philcox [*]

Center for Theoretical Physics, Department of Physics, Columbia University,
New York, New York 10027, USA
and Simons Society of Fellows, Simons Foundation, New York, New York 10010, USA

We derive optimal estimators for the two-, three-, and four-point correlators of statistically isotropic scalar fields defined on the sphere, such as the cosmic microwave background temperature fluctuations, allowing for arbitrary (linear) masking and inpainting schemes. In each case, we give the optimal unwindowed estimator (obtained via a maximum-likelihood prescription, with an associated Fisher deconvolution matrix), and an idealized form, and pay close attention to their efficient computation. For the trispectrum, we include both parity-even and parity-odd contributions, as allowed by symmetry. The estimators can include arbitrary weighting of the data (and remain unbiased), but are shown to be optimal in the limit of inverse-covariance weighting and Gaussian statistics. The normalization of the estimators is computed via Monte Carlo methods, with the rate-limiting steps (involving spherical harmonic transforms) scaling linearly with the number of bins. An accompanying code package, POLYBIN, implements these estimators in PYTHON, and we demonstrate the estimators' efficacy via a suite of validation tests.

I. INTRODUCTION

From statistical chemistry to cosmology, the physical sciences abound with examples of random fields. In many instances, the particular realization of the field (i.e., its microstate) is not of physical relevance but the distribution from which it is drawn (i.e., its macrostate) contains valuable information. For example, the precise positions of atoms in a solid are rarely of use, but their distribution encodes the physical properties of the material; likewise, whilst we care not about individual galaxy positions, their ensemble statistics can teach us about the primordial Universe. To understand such systems, we therefore require robust ways of characterizing the statistical properties of random fields.

Perhaps the most well-known statistics are the *correlation functions*. These encode the correlations between the (continuous or discrete) field at different points in space, and, if the system is sufficiently large, can be estimated via spatial averaging. Mathematically, an N-point correlator of some field a takes the form (see, for example, [1])

$$\zeta^{(N)}(\mathbf{x}_1, \ldots, \mathbf{x}_N) = \langle a(\mathbf{x}_1) \cdots a(\mathbf{x}_N) \rangle, \qquad (1)$$

where $\{\mathbf{x}_i\}$ are some positions of relevance, and $\langle \cdot \rangle$ represents an ensemble average over realizations of a. The simplest statistic is the two-point function (or, in Fourier-space, the power spectrum); this has been used to

characterize phenomena as diverse as cell biology, quantum chemistry, and astrophysics. Though this is sufficient for some applications, the rich phenomenology of nature often leads us to consider also the higher-order correlators (see, for example, [2–8]).

In this work, we will focus on random fields defined on the two-sphere. We will further specialize to *statistically isotropic* phenomena, i.e., those whose correlators are invariant under global rotations. Such fields can be conveniently described by working in a spherical harmonic basis; efficient measurement of the harmonic-space correlators, or *polyspectra*, will be the subject of the next thirty pages. Geophysics and cosmology provide a number of examples of statistically isotropic fields: for example, the set of galaxies at some fixed distance from Earth naturally lies on a spherical shell, and, by the Copernican principle, there are no special directions on large scales (see, for example, [9,10]). In the cosmological case, polyspectra with $N > 2$ are of particular relevance, since they are predicted to vanish in the simplest models of inflation (barring a number of secondary effects arising at late times) [11,12], thus we will focus primarily on the extraterrestrial case in this work. By measuring the cosmic bispectrum, trispectrum, and beyond, we can thus probe primordial physics, which operates at energy scales vastly in excess of those encountered on Earth.

Measuring correlators beyond the power spectrum is, in general, a difficult task, with naïve estimation of an N-point correlator scaling exponentially with N. For this reason,

[*]ohep2@cantab.ac.uk

most cosmological analyses have opted not to measure the full statistic, but to constrain a small number of parameters, corresponding to the amplitudes of specific physical templates (see e.g., [13–21]). Whilst this is useful for some studies, it does not facilitate general exploration of the statistic, nor probing physical effects whose forms are not *a priori* known, see e.g., [22]. To this end, it is useful to also measure the full spectra, projected onto some set of ℓ-space bins [23,24], or via some modal decomposition (which parametrize the space using smooth mode functions rather than discrete bins) [25–30].

A second complication arises from observational effects. Usually, one cannot measure the field at all points on the sphere, thus the observed field is modulated by some *mask* (also known as a *window function*), W, depending on the galactic plane, experimental limitations, bright stars, noise-dominated regions, etc. Conventional polyspectrum estimators measure the correlators (known as *pseudospectra*) of the masked field Wa rather than the true correlators of a, which complicates their interpretation, particularly given that W is rarely isotropic. To robustly compare measurement and model, one must either deconvolve the measurement (also taking into account spatially-varying noise), or convolve the theory [31,32]. Both of these are complex for statistics beyond the power spectrum [30], leading to a number of analyses ignoring the window effects, with potentially dangerous consequences.

The above discussion motivates the development of robust and efficient estimators for binned polyspectra on the two-sphere. To this end, we will build on a variety of tools developed throughout the last thirty years. First, efficient estimation of (windowed) polyspectra has been considered in a number of works, including [31–35] for the power spectrum, [13,14,16,19,22–25,30,36–39] for the bispectrum and [26–29,40–47] for the trispectrum, primarily for the amplitudes of specific separable shapes. Other works have considered modal approaches to measuring the bispectrum (often for the purpose of estimating specific nonseparable templates) and higher-order statistics, and some works have considered binned polyspectra directly [22,37]. Here, we opt to use bins rather than modal decompositions for general interpretability (given that we are not concerned with individual models); the latter may provide a more efficient compressed basis in practice however, though we caution that the various modes are not independent, even in ideal scenarios, and, furthermore, it may be nontrivial to project the theory models onto the relevant basis.

Self-consistent treatment of the mask in higher-order polyspectra is a novel feature of this work; to achieve this, we will use maximum-likelihood prescriptions, whereupon one first writes down the likelihood for the observed field (which depends both on W and the statistical properties of a), then maximizes analytically to find an optimal estimator for the statistic of interest. Such estimators are unbiased (i.e., their mean is not affected by the window) and avoid

the need to include the mask in the theory model. This approach has been previously considered for the power spectrum [32,48–53], as well as the two- and three-point statistics of three-dimensional fields [33,54–58], but, to our knowledge, ours is the first such treatment for higher-order spectra on the sphere. In this work, we will pay particular attention to trispectra, which have been rarely measured directly. Unlike the lower-order correlators, these can be decomposed into two pieces, which are even and odd under parity transformations: the latter has not been previously measured in two-dimensional cosmology, and, in our accompanying work [59], we will use it to test the recent claims of parity-violation in large scale structure [60–63]. Finally, we release public code alongside this manuscript which implements all the above estimators (both in full generality, and a simplified form); we envisage that this will facilitate robust analysis of general higher-order correlators in cosmology and beyond.

The remainder of this work is as follows. In Sec. II we set out our definitions for the binned polyspectra, before giving a general discussion of ideal estimators in Sec. III. In Secs. IV, V and VI, we derive estimators for the binned power spectrum, bispectrum, and trispectrum, giving both the idealized form and the optimal unwindowed estimator in each case. Finally, we verify the estimators numerically in Sec. VII before concluding in Sec. VIII. To guide the reader through this (necessarily dense) paper, we indicate key equations with boxes, and summarize the relevant estimators at the end of each section. Each estimator is implemented in the public code POLYBIN: an extensive tutorial can be found on GitHub.[1]

II. IDEAL BINNED POLYSPECTRA

We begin by defining our conventions for the fields and polyspectra used in this work, and present a number of results used in the remainder of this work. In general, we will work with scalar fields defined on the two-sphere, such as the atmospheric pressure on Earth or the CMB temperature fluctuations. A general zero-mean signal, labeled $a(\hat{\mathbf{n}})$ can be expanded in spherical harmonics thus

$$a(\hat{\mathbf{n}}) \equiv \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) \Leftrightarrow a_{\ell m} = \int_{\mathbb{S}^2} d\hat{\mathbf{n}} a(\hat{\mathbf{n}}) Y_{\ell m}^*(\hat{\mathbf{n}}),$$

(2)

where we will keep the summation limits and integration domains implicit henceforth. In many cases, we perform noisy observations of this signal, yielding the observed field, $\tilde{a}$, defined as

$$\tilde{a}(\hat{\mathbf{n}}) \equiv W(\hat{\mathbf{n}})a(\hat{\mathbf{n}}) + n(\hat{\mathbf{n}}),$$

(3)

---

[1]GitHub.com/oliverphilcox/PolyBin.

where $W(\hat{\mathbf{n}})$ is some deterministic window (or mask), defining how various regions of the sphere are observed and $n(\hat{\mathbf{n}})$ is a stochastic noise field. In general, we will denote masked fields with a tilde. Note that we assume W and a to be uncorrelated, such that $\langle Wa \rangle = \langle W \rangle \langle a \rangle = 0$; violation of this assumption will significantly complicate the estimators [e.g., [64]].

A. Ideal correlators

The power spectrum, C, of the signal field can be written

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} \rangle \equiv C_{m_1 m_2}^{\ell_1 \ell_2} \to (-1)^{m_1} \delta_{\ell_1 \ell_2}^{\mathrm{K}} \delta_{m_1 (-m_2)}^{\mathrm{K}} C_\ell, \quad (4)$$

where the angle brackets indicate an average over statistical realizations of the signal, and we have statistical isotropy and homogeneity to obtain the second expression. Whilst this assumption is usually valid for the underlying signal, a, realistic noise and masks are often anisotropic, thus the diagonal approximation cannot be used.

Similarly, the bispectrum, B, takes the form

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle \equiv B_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \to \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}, \quad (5)$$

where the RHS holds under isotropic and homogeneous assumptions, as before, and we have defined the *reduced bispectrum* $b_{\ell_1 \ell_2 \ell_3}$. This is symmetric under any permutation of indices, and requires $|\ell_1 - \ell_2| \le \ell_3 \le \ell_1 + \ell_2$. (5) involves the Gaunt function, defined as the average over three spherical harmonics

$$\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \equiv \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi}}$$
$$\times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}$$
$$\equiv \int d\hat{\mathbf{n}} Y_{\ell_1 m_1}(\hat{\mathbf{n}}) Y_{\ell_2 m_2}(\hat{\mathbf{n}}) Y_{\ell_3 m_3}(\hat{\mathbf{n}}), \quad (6)$$

where the 3×2 matrices are Wigner $3j$ symbols. For an isotropic real scalar field, a, the bispectrum is parity-even, and thus vanishes unless $\ell_1 + \ell_2 + \ell_3$ is even (which is enforced by the Gaunt integral).[2]

Finally, we can define a trispectrum, T, of a via

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} a_{\ell_4 m_4} \rangle_c \equiv T_{m_1 m_2 m_3 m_4}^{\ell_1 \ell_2 \ell_3 \ell_4}. \quad (7)$$

where we take only the *connected* part of the correlator. In this case, the rotationally invariant decomposition is less straightforward, since the reduced trispectrum cannot be fully described by four ℓ-modes, rather we must introduce

---

[2]For anisotropic signals, such as galactic dust, nonzero parity-odd bispectra can exist. To compute these, one can use a modified definition of the reduced bispectrum, as discussed in [22,37].

also a diagonal element, L. As discussed in [27], we can introduce the nonredundant function $T_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L)$, via

$$T_{m_1 m_2 m_3 m_4}^{\ell_1 \ell_2 \ell_3 \ell_4} \to \sum_{L=0}^{\infty} \sum_{M=-L}^{L} (-1)^M \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix}$$
$$\times \begin{pmatrix} \ell_3 & \ell_4 & L \\ m_3 & m_4 & M \end{pmatrix} T_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L), \quad (8)$$

summing over the diagonal, L, and its azimuthal component, and noting that $M = m_1 + m_2 = -m_3 - m_4$. This has a number of nontrivial symmetries, in particular

$$T_{\ell_3 \ell_4}^{\ell_2 \ell_1}(L) = (-1)^{\ell_1 + \ell_2 + L} T_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L), \quad T_{\ell_1 \ell_2}^{\ell_3 \ell_4}(L) = T_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L). \quad (9)$$

For our purpose, it will be useful to introduce a new trispectrum, $t_{\ell_3 \ell_4}^{\ell_1 \ell_2}$, via the symmetric definition (similar to [37] for the bispectrum)

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} a_{\ell_4 m_4} \rangle_c$$
$$\equiv T_{m_1 m_2 m_3 m_4}^{\ell_1 \ell_2 \ell_3 \ell_4} \to \sum_{L=0}^{\infty} \sum_{M=-L}^{L} (-1)^M w_{\ell_1 \ell_2 m_1 m_2}^{L(-M)} w_{\ell_3 \ell_4 m_3 m_4}^{LM} t_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L)$$
$$+ 23 \text{ perms}, \quad (10)$$

summing over twenty-four permutations of $\{\ell_1, \ell_2, \ell_3, \ell_4\}$. This involves a new weighting function, akin to the Gaunt function

$$w_{\ell_1 \ell_2 m_1 m_2}^{LM} \equiv \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2L + 1)}{4\pi}}$$
$$\times \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & M \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}$$
$$\equiv \int d\hat{\mathbf{n}} {}_{+1}Y_{\ell_1 m_1}(\hat{\mathbf{n}}) {}_{+1}Y_{\ell_2 m_2}(\hat{\mathbf{n}}) {}_{-2}Y_{LM}(\hat{\mathbf{n}}), \quad (11)$$

which we have written as an integral over three spin-weighted spherical harmonics, ${}_s Y_{\ell m}(\hat{\mathbf{n}})$, in the final line. Note that this is symmetric under interchange of (ℓ_1, m_1) and (ℓ_2, m_2). The reason for the spin-weighting adopted in w will be explained in Sec. VI A 1.

The reduced trispectrum obeys the symmetries

$$t_{\ell_3 \ell_4}^{\ell_2 \ell_1}(L) = t_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L), \quad t_{\ell_1 \ell_2}^{\ell_3 \ell_4}(L) = t_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L), \quad (12)$$

to fully specify the trispectrum, we thus require only values with $\ell_1 \le \ell_2$, $\ell_3 \le \ell_4$, $\ell_3 \le \ell_1$ and, if $\ell_1 = \ell_3$, $\ell_2 \le \ell_4$. The diagonal, L, satisfies the triangle conditions $|\ell_1 - \ell_2| \le L \le \ell_1 + \ell_2$ and $|\ell_3 - \ell_4| \le L \le \ell_3 + \ell_4$, due to the $3j$ symbols. Finally, under conjugation and parity-inversion, the trispectrum satisfies

$$[t^{\ell_1\ell_2}_{\ell_3\ell_4}(L)]^* = (-1)^{\ell_1+\ell_2+\ell_3+\ell_4} t^{\ell_1\ell_2}_{\ell_3\ell_4}(L), \quad \mathbb{P}[[t^{\ell_1\ell_2}_{\ell_3\ell_4}(L)]$$
$$= (-1)^{\ell_1+\ell_2+\ell_3+\ell_4} t^{\ell_1\ell_2}_{\ell_3\ell_4}(L), \qquad (13)$$

respectively; as such, trispectra with even (odd) $\ell_1 + \ell_2 + \ell_3 + \ell_4$ are parity-even (parity-odd) and purely real (imaginary).

B. Binning

In this work, we will focus on computing polyspectra in some set of ℓ-bins, whose formulation we now turn to. An alternative approach would be to project the spectra onto some type of modal decomposition [26–28,30,37]. Our approach has the benefit that the output spectra can be directly compared to theory, in the limit of narrow bins (see [31] for techniques going beyond this limit). To include binning, we introduce the (arbitrary) binning function $\Theta_\ell(b)$, which is usually defined to be unity if ℓ is in bin b and zero else. For the (isotropic and homogeneous) power spectrum, this leads to the definition

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} \rangle \approx (-1)^{m_1} \delta^K_{\ell_1\ell_2} \delta^K_{m_1(-m_2)} \sum_b \Theta_{\ell_1}(b) C(b), \quad (14)$$

where $C(b)$ are the binned quantities we shall construct estimators for.[3]

The binned bispectrum, $b(\mathbf{b})$, is similarly defined

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle \approx \mathcal{G}^{\ell_1\ell_2\ell_3}_{m_1 m_2 m_3} \sum_{b_1 b_2 b_3} \frac{b(\mathbf{b})}{\Delta_3(\mathbf{b})}$$
$$\times [\Theta_{\ell_1}(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3) + 5 \text{ perms}], \qquad (15)$$

where $\mathbf{b} \equiv \{b_1, b_2, b_3\}$ and the permutations are in $\{\ell_1, \ell_2, \ell_3\}$. To avoid double counting, we restrict the summation to $b_1 \leq b_2 \leq b_3$ (giving rise to the permutations), and introduce a permutation factor

$$\Delta_3(\mathbf{b}) \equiv \begin{cases} 6 & b_1 = b_2 = b_3 \\ 2 & b_1 = b_2 \neq b_3 \text{ or } b_1 \neq b_2 = b_3 \\ 1 & \text{else.} \end{cases} \qquad (16)$$

This ensures that, in the limit of thin bins, $b(\mathbf{b})$ is equal to the bispectrum evaluated at the bin centers. Note that we can incorporate squeezed triangles into this formalism by allowing a larger ℓ_{max} for ℓ_2 and ℓ_3 (equivalently b_2 and b_3) than ℓ_1.

---

[3]Strictly $C(b)$ is related to a sum over the ℓ-bins rather than being the value at the bin-center. The approximation of (14) is valid for slowly varying C_ℓ and suitably narrow bins. Since our focus here is principally on higher-point functions, this is appropriate here.

Finally, we can define the binned trispectrum, $t(\mathbf{b}, B)$, by analogy with (10) as

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} a_{\ell_4 m_4} \rangle_c$$
$$\approx \sum_{LM} (-1)^M w^{L(-M)}_{\ell_1\ell_2 m_1 m_2} w^{LM}_{\ell_3\ell_4 m_3 m_4} \sum_{\mathbf{b}, B} \frac{t(\mathbf{b}, B)}{\Delta_4(\mathbf{b})} \Theta_L(B)$$
$$\times [\Theta_{\ell_1}(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3)\Theta_{\ell_4}(b_4) + 7 \text{ perms}]$$
$$+ (2 \leftrightarrow 3) + (2 \leftrightarrow 4), \qquad (17)$$

summing over the following permutations of $\{1, 2, 3, 4\} \equiv \{b_1, b_2, b_3, b_4\}$:

$$\{1, 2, 3, 4\}, \{1, 2, 4, 3\}, \{2, 1, 3, 4\}, \{2, 1, 4, 3\},$$
$$\{3, 4, 1, 2\}, \{3, 4, 2, 1\}, \{4, 3, 1, 2\}, \{4, 3, 2, 1\}, \qquad (18)$$

due to the symmetry properties given in (12). Here, the trispectrum is defined for all $\mathbf{b} \equiv \{b_1, b_2, b_3, b_4\}$ encompassing the external $\{\ell_1, \ell_2, \ell_3, \ell_4\}$ legs and B describing the diagonal L. We sum over all $b_1 \leq b_2$, $b_3 \leq b_4$, $b_1 \leq b_3$ and, if $b_3 = b_1$, $b_2 \geq b_4$, i.e., over all independent bins, noting additionally that $\{b_1, b_2, B\}$ and $\{b_3, b_4, B\}$ must satisfy triangle conditions.[4] If we wish to include squeezed and doubly-squeezed tetrahedra, we simply extend the ℓ ranges to use a larger ℓ_{max} for ℓ_2, ℓ_4 and L (due to the triangle conditions), or, if our focus is collapsed tetrahedra, we can use a smaller ℓ_{max} for L. To avoid double counting, we introduce the degeneracy factor

$$\Delta_4(\mathbf{b}) = \begin{cases} 8 & b_1 = b_2 = b_3 = b_4 \\ 4 & b_1 = b_2 \text{ and } b_3 = b_4 \\ 2 & b_1 = b_2 \text{ or } b_3 = b_4 \\ 2 & b_1 = b_3 \text{ and } b_2 = b_4 \\ 1 & \text{else,} \end{cases} \qquad (19)$$

which is simply the number of distinct appearances of each term in the above sum over permutations. Finally, we note that we can impose that the trispectrum is parity-even (parity-odd) by adding a factor $[1 \pm (-1)^{\ell_{1234}}]/2$ to (17).

III. MASKED DATA AND MINIMUM-VARIANCE ESTIMATORS

We now turn to the problem of estimating the binned polyspectrum coefficients discussed in Sec. II from the observed masked data. For this purpose, we will first consider the nonideal correlators, for which it is useful to work in map-space, rather than harmonic space. These results may then be used to derive optimal estimators, by maximizing an weakly non-Gaussian likelihood.

---

[4]For the parity-odd trispectrum, we can additionally drop bins satisfying $b_1 = b_3$ and $b_2 = b_4$, which evaluate to zero.

A. Nonideal correlators

When working with observational data, we may no longer assume rotational symmetry, since the weights, W (encoding the response of the map to the underlying field a) and the noise, n, can be inhomogeneous. In this case, the results of Sec. II do not apply. Working in map space, the two-point correlator of the observed field can be instead written as

$$\tilde{\mathsf{C}}^{ij} \equiv \langle \tilde{a}(\hat{\mathbf{n}}^i)\tilde{a}(\hat{\mathbf{n}}^j)\rangle, \tag{20}$$

where i, j index points on the sky (such as HEALPix pixels).[5] From the definition of the masked field (3), this can be written in terms of the unmasked two-point function, $\mathsf{C}^{ij} \equiv \langle a(\hat{\mathbf{n}}^i)a(\hat{\mathbf{n}}^j)\rangle$ and the noise $\mathsf{N}^{ij} \equiv \langle n(\hat{\mathbf{n}}^i)n(\hat{\mathbf{n}}^j)\rangle$

$$\begin{aligned}
\tilde{\mathsf{C}}^{ij} &= W(\hat{\mathbf{n}}^i)\mathsf{C}^{ij}W(\hat{\mathbf{n}}^j) + \mathsf{N}^{ij} \\
&= \sum_{\ell m} B_\ell^2 C_\ell [W(\hat{\mathbf{n}}^i)Y_{\ell m}(\hat{\mathbf{n}}^i)][W(\hat{\mathbf{n}}^j)Y_{\ell m}^*(\hat{\mathbf{n}}^j)] + \mathsf{N}^{ij},
\end{aligned} \tag{21}$$

where we have expanded the true correlator, which *is* rotationally invariant, in terms of (4) and additionally included an (isotropic) beam B_ℓ. When working with discrete data (such as maps in HEALPix format), we can additionally include a pixel window function in B_ℓ (via $B_\ell \to w_\ell B_\ell$ for window w_ℓ), to remove the leading dependence on N_{side}.

The utility of (21) is that the windowed correlator is expressed in terms of the quantity we wish to estimate, C_ℓ (neglecting binning for now). This is further simplified if one assumes the same window for data and noise: in this case, the masks simply pre and postmultiply the correlator. In general, the map-level correlator is an $N_{\text{pix}} \times N_{\text{pix}}$ matrix which is difficult if not impossible to compute explicitly; however, its action on a map can be straightforwardly defined. We will discuss this in Sec. IV C.

The nonideal three- and four-point correlators take similar forms,

$$\begin{aligned}
\tilde{\mathsf{B}}^{ijk} &\equiv \langle \tilde{a}(\hat{\mathbf{n}}^i)\tilde{a}(\hat{\mathbf{n}}^j)\tilde{a}(\hat{\mathbf{n}}^k)\rangle, \\
\tilde{\mathsf{T}}^{ijkl} &\equiv \langle \tilde{a}(\hat{\mathbf{n}}^i)\tilde{a}(\hat{\mathbf{n}}^j)\tilde{a}(\hat{\mathbf{n}}^k)\tilde{a}(\hat{\mathbf{n}}^l)\rangle.
\end{aligned} \tag{22}$$

As before, these can be straightforwardly written in terms of the map-space ideal correlators (B^{ijk} and T^{ijkl}), and, via the results of Sec. II A, the binned quantities we wish to measure. Furthermore, whilst they are high-dimensional objects (rank 3- and 4-tensors), their action on maps will prove straightforward, due to the internal symmetries in the unwindowed correlators.

---

[5]Throughout this work, we will used serif fonts to denote map-space correlators, e.g., B, *sans serif* to denote harmonic correlators, e.g., B, and lower case for reduced spectra, e.g., b. Latin indices $i, j, k, \ldots$ will always denote summation over spatial points. Tildes are added to indicate masked fields.

B. Optimal estimators

Optimal estimators for the binned polyspectra can be derived by maximizing the likelihood of the observed map $\tilde{a}(\hat{\mathbf{n}})$. Assuming weak non-Gaussianity, this is given by an Edgeworth expansion in terms of the map-space correlators [e.g., [65]],

$$\begin{aligned}
L[\tilde{a}] \propto \exp\left[-\frac{1}{2}h_i\tilde{\mathsf{C}}^{ij}h_j\right] \\
\times \left\{1 + \frac{1}{3!}\tilde{\mathsf{B}}^{ijk}\mathcal{H}_{ijk} + \frac{1}{4!}\tilde{\mathsf{T}}^{ijkl}\mathcal{H}_{ijkl} + \cdots\right\},
\end{aligned} \tag{23}$$

where we sum over all repeated indices, such that $\alpha_i\beta^i \equiv \int d\hat{\mathbf{n}}\alpha(\hat{\mathbf{n}})\beta(\hat{\mathbf{n}})$. Here, we have defined the Wiener-filtered map, $h(\hat{\mathbf{n}}) \equiv [\tilde{\mathsf{C}}^{-1}\tilde{a}](\hat{\mathbf{n}})$ (recalling that $\tilde{\mathsf{C}}$ contains both signal and noise, and (for now) assuming it to be invertible), as well as the (map-space) Hermite polynomials

$$\begin{aligned}
\mathcal{H}_{ijk} &\equiv h_ih_jh_k - (h_i\tilde{\mathsf{C}}_{jk}^{-1} + 2 \text{ perms}) \\
\mathcal{H}_{ijkl} &\equiv h_ih_jh_kh_l - (h_ih_j\tilde{\mathsf{C}}_{kl}^{-1} + 5 \text{ perms}) \\
&\quad + (\tilde{\mathsf{C}}_{ij}^{-1}\tilde{\mathsf{C}}_{kl}^{-1} + 2 \text{ perms}).
\end{aligned} \tag{24}$$

In (23), all cosmology dependence appears through the map-space correlators, $\tilde{\mathsf{C}}$, $\tilde{\mathsf{B}}$, $\tilde{\mathsf{T}}$, which can be related to their binned coefficients, $C(b)$, $b(\mathbf{b})$, $t(\mathbf{b}, L)$ using the results of Sec. II. To see this, let us consider some binned quantity $x(\mathbf{b})$ arising only in the $N > 2$-point correlator, $\tilde{\mathsf{X}}^{i_1\cdots i_N}$. An estimator for $x(\mathbf{b})$ can be obtained by maximizing $\log L[\tilde{a}](x)$, yielding

$$\hat{x}(\mathbf{b}) \propto \frac{1}{N!}\frac{\partial \tilde{\mathsf{X}}^{i_1\cdots i_N}}{\partial x(\mathbf{b})}\mathcal{H}_{i_1\cdots i_N}, \tag{25}$$

working in the limit of small x. The normalization factor (which is, in general, a matrix), can be derived by requiring that the estimator is unbiased, i.e., $\mathbb{E}[\hat{x}(\mathbf{b})] = x(\mathbf{b})$ for expectation operator $\mathbb{E}$. We will refer to its inverse as the *Fisher matrix*, defined as

$$\mathcal{F}_N(\mathbf{b}, \mathbf{b}') = \frac{1}{N!}\frac{\partial \tilde{\mathsf{X}}^{i_1\cdots i_N}}{\partial x(\mathbf{b})}\tilde{\mathsf{C}}_{i_1j_1}^{-1}\cdots\tilde{\mathsf{C}}_{i_Nj_N}^{-1}\frac{\partial \tilde{\mathsf{X}}^{j_1\cdots j_N}}{\partial x(\mathbf{b}')}, \tag{26}$$

where we note that all disconnected terms cancel in the expectation of the Hermite tensor $\mathcal{H}_{i_1\cdots i_N}$.

The estimator has the following properties:

(i) *Unbiased*: This follows from the definition of the Fisher matrix, which ensures $\mathbb{E}[\hat{x}(\mathbf{b})] = x(\mathbf{b})$. One caveat should be noted. This assumes that the polyspectrum is completely characterized by the set of measured bins $\{x(\mathbf{b})\}$ (which generically include noise contributions), i.e., that $\tilde{\mathsf{X}}^{i_1\cdots i_N} = \sum_{\mathbf{b}} x(\mathbf{b})(\partial\tilde{\mathsf{X}}^{i_1\cdots i_N}/\partial x(\mathbf{b}))$. Violations of this can occur if there is contribution from modes outside

the binning range or unaccounted-for leakage between polyspectra of different parity. For ideal polyspectrum estimators with $N \leq 3$, the Fisher matrix is diagonal, so such effects may be neglected; in the general case, one can ameliorate this by constructing the estimator using a slightly larger range of bins than desired in the output data-product, and, if necessary, accounting for leakage between states of different parity (cf. Sec VI B).

(ii) *Window-Free*: Since we consistently include window functions in the $\tilde{X}$ definitions and take derivatives with respect to the *true* binned correlators, $x(\mathbf{b})$, the measurements are not biased by the window function, i.e., the estimators take into account the response of the map to the underlying signal. This lies in contrast to the pseudospectra measured by direct polyspectrum estimators, and occurs due to the appearance of the mask in the estimator, and the form of the normalization matrix.

(iii) *Optimal*: Since (25) has been derived using maximum-likelihood techniques, it is optimal in the small-correlator limit. As such, the covariance is given by $\mathcal{F}_N^{-1}$.[6]

(iv) *Efficiency* As shown below, the estimator numerators scale at most linearly with the total number of bins in the statistic, N_{bin}. When computed using Monte Carlo methods, the rate-determining step in computation of the Fisher matrix is also linear in the number of bins.

Whilst the above estimators have significant formal utility, in practice, it will prove useful to consider estimators with a more general choice of weighting, due to the difficulties obtaining accurate noise covariances, N^{ij}, and inverting the covariance $\tilde{\mathsf{C}}^{ij}$. In the below, we will consider a more general choice of weighting to enable efficient computation, defining $h(\hat{\mathbf{n}}) = \mathsf{S}^{-1}\tilde{a}$ instead of $\tilde{\mathsf{C}}^{-1}\tilde{a}$, where S^{-1}, is some weighting matrix that is not required to be symmetric or invertible. For example, one may wish to project out contaminated areas of a map, which, *a priori*, has a uniform weight; in this case, S^{-1} would excise regions of the map, and thus be noninvertible, whilst W could be set to the identity operator. Arbitrary linear operations (such as linear inpainting, cf. [35]) can also be included in S^{-1}, as appropriate to the task in question.

Replacing $\tilde{\mathsf{C}}^{-1}$ with S^{-1} in (25) and (26) we obtain an estimator which is always unbiased, and minimum variance in the limit of $\mathsf{S}^{-1} \to \tilde{\mathsf{C}}^{-1}$. In the Gaussian regime, the covariance for general (not necessarily invertible) S^{-1} is given by

$$\text{cov}_N(\mathbf{b}, \mathbf{b}') = \mathcal{F}_N^{-1}[\mathsf{S}^{-1}]\mathcal{F}_N[\mathsf{S}^{-1}\tilde{\mathsf{C}}\mathsf{S}^{-\text{T}}]\mathcal{F}_N^{-\text{T}}[\mathsf{S}^{-1}], \quad (27)$$

where $\mathcal{F}_N[\mathsf{M}]$ as the Fisher matrix with weighting M (which may be asymmetric); if S^{-1} is invertible, the degree of suboptimality is second order in $(\mathsf{S} - \tilde{\mathsf{C}})$ [36], and, if not, an optimal estimator is obtained if $\mathsf{S}^{-1} = \mathsf{S}^{-1}\tilde{\mathsf{C}}\mathsf{S}^{-\text{T}}$. In the remainder of this work, we consider how such estimators can be efficiently computed.

IV. OPTIMAL POWER SPECTRUM ESTIMATION

With the above formalism in place, we may now proceed to derive the optimal estimator for the binned full-sky power spectra, $C(b)$, analogous to [32,33,66]. This is derived in a slightly different manner to the bispectrum and trispectrum estimators discussed below (though ends up taking the same functional form), since the two-point function appears both in the Gaussian likelihood of (23), and in the Wiener filtering. Our estimators can be shown to be equivalent to those of the MASTER formalism in certain limits [31] (in particular a uniform mask W, and a weighting scheme of the form $[\mathsf{S}^{-1}v](\hat{\mathbf{n}}) = s(\hat{\mathbf{n}})v(\hat{\mathbf{n}})$, for some filter s, possibly with incomplete support).

Taking derivatives with respect to $C(b)$ (dropping all non-Gaussian correlators) gives

$$\frac{\partial \log L[\tilde{a}]}{\partial C(b)} = \frac{1}{2}\frac{\partial \tilde{\mathsf{C}}^{ij}}{\partial C(b)}h_ih_j - \frac{1}{2}\text{Tr}\left[\tilde{\mathsf{C}}^{-1}\frac{\partial \tilde{\mathsf{C}}}{\partial C(b)}\right]$$

$$\frac{\partial^2 \log L[\tilde{a}]}{\partial C(b)\partial C(b')} = -\left[\frac{\partial \tilde{\mathsf{C}}}{\partial C(b)}\tilde{\mathsf{C}}^{-1}\frac{\partial \tilde{\mathsf{C}}}{\partial C(b)}\right]^{ij}h_ih_j$$

$$+ \frac{1}{2}\text{Tr}\left[\tilde{\mathsf{C}}^{-1}\frac{\partial \tilde{\mathsf{C}}}{\partial C(b)}\tilde{\mathsf{C}}^{-1}\frac{\partial \tilde{\mathsf{C}}}{\partial C(b')}\right], \quad (28)$$

noting that $\tilde{\mathsf{C}}_{,\alpha}^{-1} = -\tilde{\mathsf{C}}^{-1}\tilde{\mathsf{C}}_{,\alpha}\tilde{\mathsf{C}}^{-1}$. To derive the optimal estimator, we expand the likelihood to second order around some fiducial spectrum $\bar{C}(b)$, and maximize with respect to the true spectrum $C(b)$, yielding

$$\hat{C}(b) = \bar{C}(b) + \frac{1}{2}\sum_{b'}\mathcal{F}_{2,\text{opt}}^{-1}(b, b')\frac{\partial \tilde{\mathsf{C}}^{ij}}{\partial C(b')}[h_ih_j - \tilde{\mathsf{C}}_{ij}^{-1}],$$

$$(29)$$

defining

$$\mathcal{F}_{2,\text{opt}}(b, b') = \frac{1}{2}\text{Tr}\left[\tilde{\mathsf{C}}^{-1}\frac{\partial \tilde{\mathsf{C}}}{\partial C(b)}\tilde{\mathsf{C}}^{-1}\frac{\partial \tilde{\mathsf{C}}}{\partial C(b')}\right], \quad (30)$$

and evaluating all quantities at the fiducial spectrum $\bar{C}$. Our final estimator is formed by replacing $\tilde{\mathsf{C}}^{-1}$ by some generic weighting S^{-1} (which, as above, need not be symmetric or invertible), and noting that the first and last

---

[6]This is strictly true only for real-valued $x(\mathbf{b})$. Later, we will find that the trispectrum contains an imaginary piece, for which the covariance is $-\mathcal{F}_4^{-1}$. As long as we always take the imaginary part of any such quantities, the above logic holds.

terms in (29) differ only by the noise correlation function, N^{ij}. This gives

$$
\begin{aligned}
\{\hat{C}(b) = \frac{1}{2}\sum_{b'}\mathcal{F}_2^{-1}(b,b')&\left[\frac{\partial\tilde{C}^{ij}}{\partial C(b')}h_ih_j\right.\\
&\left.-\mathrm{Tr}\left(\frac{\partial\tilde{C}}{\partial C(b')}S^{-1}NS^{-T}\right)\right],\\
\mathcal{F}_2(b,b') = \frac{1}{2}\mathrm{Tr}&\left[S^{-T}\frac{\partial\tilde{C}}{\partial C(b)}S^{-1}\frac{\partial\tilde{C}}{\partial C(b')}\right],
\end{aligned}
\tag{31}
$$

where $h \equiv S^{-1}\tilde{a}$ and S^{-T} is the transpose of S^{-1}. The second term in $\hat{C}(b)$ subtracts off the estimator bias induced by noise in the data; for the remainder of this work, we will absorb this into $\hat{C}(b)$, and drop this term. Notably, $\mathcal{F}_2(b,b')$ is only symmetric if $S^{-1} = S^{-T}$ (which does not hold if the filtering projects out modes). This is the minimum variance estimator in the limit of a Gaussian likelihood, $\bar{C}(b) \to C^{\mathrm{true}}(b)$, and $S^{-1} \to \tilde{C}^{-1}$ (or $S^{-1} \to S^{-1}\tilde{C}S^{-1}$ in general).

A. Idealized form

Let us consider (31) in the idealized limit, i.e., without a mask or beam and assuming isotropic noise (absorbed into the binned spectrum). Here, the relevant correlator derivative becomes

$$
\frac{\partial C^{ij}}{\partial C(b)} = \sum_{\ell m}\Theta_\ell(b)Y_{\ell m}(\hat{\mathbf{n}}^i)Y_{\ell m}^*(\hat{\mathbf{n}}^j),
\tag{32}
$$

moving to harmonic-space and using the binned power spectrum definition (14). The Wiener-filtered field can be simply written as $h_{\ell m} = S_\ell^{-1}a_{\ell m}$ (assuming S to be diagonal in harmonic space), thus, following a little algebra, we obtain the estimator

$$
\begin{aligned}
\hat{C}_{\mathrm{ideal}}(b) &= \frac{1}{2}\mathcal{F}_{2,\mathrm{ideal}}^{-1}(b)\sum_{\ell m}\Theta_\ell(b)\frac{|a_{\ell m}|^2}{S_\ell^2},\\
\mathcal{F}_{2,\mathrm{ideal}}(b) &= \frac{1}{2}\sum_\ell\Theta_\ell(b)\frac{2\ell+1}{S_\ell^2}.
\end{aligned}
\tag{33}
$$

This is just the conventional binned power spectrum estimator (summing over all ℓ, m allowed by the binning function), albeit including a Wiener-filter weighting. In this case, the normalization is diagonal in the bins, i.e., each measurement is independent. Computation of (33) requires one harmonic transform (to define $a_{\ell m}$), then a simple sum over $\{\ell, m\}$; the latter process scales as $\mathcal{O}(N_{\mathrm{bin}})$ for N_{bin} bins in $\{b\}$.

B. General form

In the presence of a mask, we instead simplify (31) by inserting the following two-point function derivative,

$$
\frac{\partial\tilde{C}^{ij}}{\partial C(b)} = \sum_{\ell m}B_\ell^2\Theta_\ell(b)[W(\hat{\mathbf{n}}^i)Y_{\ell m}(\hat{\mathbf{n}}^i)][W(\hat{\mathbf{n}}^j)Y_{\ell m}^*(\hat{\mathbf{n}}^j)],
\tag{34}
$$

cf. (21). The estimator numerator can be written

$$
\hat{C}(b) \propto \frac{1}{2}\sum_{\ell m}B_\ell^2\Theta_\ell(b)[Wh]_{\ell m}[Wh]_{\ell m}^*,
\tag{35}
$$

where $[Wh]_{\ell m}$ is the harmonic-space representation of $W(\hat{\mathbf{n}})h(\hat{\mathbf{n}}) \equiv W(\hat{\mathbf{n}})[S^{-1}\tilde{a}](\hat{\mathbf{n}})$. This is straightforward to compute by a direct harmonic-space sum, and scales as $\mathcal{O}(N_{\mathrm{bin}})$, as before (with only one invocation of S^{-1} required).

The Fisher matrix of (31) is more difficult to compute due to the trace, and formally requires $\mathcal{O}(N_{\mathrm{pix}}^2)$ operation. One option to compute it is to note that, if the estimator is optimal, it is equal to the covariance of the $\tilde{C}(b)$ numerator. As such, one could compute this quantity for a number of Monte Carlo simulations and form the covariance empirically. However, this is slow to converge (since we require the inverse Fisher matrix), and only exact if S^{-1} is the true inverse covariance.

To form a practically implementable estimator, we can instead rewrite the Fisher matrix as an expectation over some Gaussian random field (GRF) u, as in [36,54,55]. This corresponds to writing

$$
\mathcal{F}_2(b,b') = \frac{1}{2}\left\langle\left(\frac{\partial\tilde{C}}{\partial C(b)}S^{-1}u\right)^{\mathrm{T}}S^{-1}\left(\frac{\partial\tilde{C}}{\partial C(b')}U^{-1}u\right)\right\rangle_u
\tag{36}
$$

or, more simply,

$$
\mathcal{F}_2(b,b') = \frac{1}{2}\langle Q_2^{\mathrm{T}}[S^{-1}u](b)\times WS^{-1}W\times Q_2[U^{-1}u](b')\rangle_u,
\tag{37}
$$

where U is the (arbitrary, but invertible) covariance of u, and given that $\langle U^{-1}uu^{\mathrm{T}}\rangle_u$ is just the identity matrix, is equivalent to the Fisher matrix given in (31). In the second line, we have defined a filtered map, Q_2,

$$
\begin{aligned}
Q_2[x](\hat{\mathbf{n}}^i;b) &\equiv B\cdot\frac{\partial C^{ij}}{\partial C(b)}[Wx](\hat{\mathbf{n}}^j) \Leftrightarrow Q_{2,\ell m}[x](b)\\
&= B_\ell^2\Theta_\ell(b)[Wx]_{\ell m},
\end{aligned}
\tag{38}
$$

where the second (harmonic-space) definition holds due to (34), and the first derivative includes the (optionally pixel-windowed) beam (unlike in Sec. IV A). Similar definitions will be used for the higher-order polyspectra. Thus, to form the Fisher matrix, one must compute the Q_2 filters on a set of N_{GRF} GRFs, u, transform them to map space, and combine them with a $WS^{-1}W$ weighting. Each realization

is symmetric only if $S^{-1} = U^{-1}$ and $S^{-1} = S^{-T}$; the average over realizations requires only the latter condition. This can be straightforwardly achieved using repeated spherical harmonic transforms. Notably, it does not require any $\mathcal{O}(N_{\mathrm{pix}}^2)$ operations to implement. In practice, we will find that $\mathcal{O}(100)$ GRFs is sufficient for the above calculation, and adds an extra variance to the power spectrum at the $\mathcal{O}(1 + N_{\mathrm{GRF}}^{-1})$ level. This is much more efficient than the Monte Carlo covariance estimation method discussed above. Computationally it involves two processes: computation of N_{bin} WQ_2 and $S^{-1}WQ_2$ maps, and a summation over all such pairs. The first operation scales as $\mathcal{O}(N_{\mathrm{bin}})$ (involving N_{bin} S^{-1} operations), whilst the second is quadratic in N_{bin}. Usually, the first dominates (since each term requires several harmonic transforms, depending on the form of S^{-1}), resulting in a linear computational scaling (in the limit of large memory).

C. Choice of S^{-1} and U

To implement the general estimator described above, we must specify: (a) the mask, W, (b) the weighting matrix S^{-1}, which downweights low signal-to-noise or troublesome regions, and (c) the Fisher GRF covariance U. The first is usually the simplest; this is the linear response of the data d to the underlying field a, which, for a full-sky map (or an inpainted version thereof), is often unity. For cut-sky data, this indicates which regions are observed and is often a binary mask, albeit with some additional smoothing. A variety of additional phenomena can be included here, such as pixel weights and discreteness effects.

The optimal choice for the weighting operator S^{-1} is the inverse data covariance $\tilde{C}^{-1}$. In realistic scenarios, this is neither diagonal in map nor harmonic space, and is thus difficult to invert (though may be possible via approaches such as [67], which supplement the map with additional uncorrelated noise). A simpler choice may be use a diagonal-in-ℓ approximation of the covariance to define S^{-1} (equal to $1/S_\ell$ in harmonic-space). Often, one may wish to downweight or remove specific regions of the map in the analysis before applying such a weight; this can be achieved by first removing areas of the map with some projection matrix Π, e.g.,

$$[S^{-1}v](\hat{n}) = \sum_{\ell m} Y_{\ell m}(\hat{n}) \frac{[\Pi v]_{\ell m}}{S_\ell}, \qquad (39)$$

for arbitrary map v, where $[x]_{\ell m}$ is the harmonic transform of x. If Π is not of full-rank, S^{-1} is not invertible.[7] One may

additionally wish to 'inpaint' the map, by filling in small holes with the mean of the surrounding pixels [35]. Since this is a linear operation it can be included within S^{-1}, and fully accounted for in the normalization (noting that the below estimators require only the action of S^{-1} on maps v, rather than the explicit form of S).

If we wish to use optimal weights, however, some progress can be made using conjugate gradient descent methods. In this framework, we compute the action of the inverse matrix, $\tilde{C}^{-1}$, on some map $\alpha(\hat{n})$, using only applications of the uninverted matrix on maps $\beta(\hat{n})$. These can be computed thus (for arbitrary β),

$$
\begin{aligned}
[\tilde{C}\beta](\hat{n}^i) &= W(\hat{n}^i)C^{ij}[W\beta](\hat{n}^j) + [N\beta](\hat{n}^i) \\
&= W(\hat{n}^i)[B_\ell^2 C_\ell[W\beta]_{\ell m}](\hat{n}^i) + [N\beta](\hat{n}^i), \qquad (40)
\end{aligned}
$$

where we have written $\tilde{C}$ in harmonic space to obtain the second quantity via (21), and (as above) denoted forward- and backward harmonic transforms by $[\cdots]_{\ell m}$ and $[\cdots](\hat{n})$ respectively. To apply $\tilde{C}$ to a map, our approach is thus: (1) multiply the map by W in map-space; (2) transform to harmonic-space and multiply by $B_\ell^2 C_\ell$, (3) transform to map-space and multiply by $W(\hat{n})$; (4) add on the noise term (which is straightforward if it is diagonal in map space). Given the above mapping, and an appropriate preconditioner (such as the diagonal harmonic-space covariance), we may form $\tilde{C}^{-1}\alpha$ iteratively for any given map α. We caution that conjugate-gradient-descent inversion is computationally expensive procedure (involving two harmonic transforms per iteration), and we require at least $(N_{\mathrm{bins}} + 2)$ applications of the inverse map for the full estimator (one for h, one for $S^{-1}u$, and one for each bin in the Fisher matrix). In practice, therefore, we will primarily use a diagonal approximation in this work. Notably, this does not bias any results, but will lead to a slight loss of optimality.

Finally, some care is warranted regarding our choice of the Monte Carlo maps u, and their covariance U. Ideally, we require maps that are easy to simulate, i.e., GRFs. This also simplifies the interpretation, since no higher-point correlators need to be removed (which will be relevant for the bispectrum and beyond). We further require the maps to have a precisely known and simply invertible covariance. The convergence of the Monte Carlo estimates is fastest if U^{-1} is close to the weighting matrix S^{-1}; in this work, we will fix $U^{-1} = S^{-1}$, assuming a diagonal form for both.

V. OPTIMAL BISPECTRUM ESTIMATION

We now turn to the window-free bispectrum. As shown in Sec. III, the general estimator for the binned bispectrum takes the form

---

[7]Notably, there is a degeneracy between the mask W and the weighting S^{-1}: multiplying the data by some invertible function $f(\hat{n})$ sends $W \to fW$, resulting in the same estimator if $S^{-1} \to f^{-1}S^{-1}f^{-1}$. Note that this also affects the noise correlator also.

$$\hat{b}(\mathbf{b}) = \frac{1}{3!}\sum_{\mathbf{b}'}\mathcal{F}_3^{-1}(\mathbf{b},\mathbf{b}')\frac{\partial\tilde{\mathsf{B}}^{ijk}}{\partial b(\mathbf{b})}$$
$$\times[h_ih_jh_k - (h_i\langle h_jh_k\rangle + 2\text{ perms})],$$

$$\mathcal{F}_3(\mathbf{b},\mathbf{b}') = \frac{1}{3!}\frac{\partial\tilde{\mathsf{B}}^{ijk}}{\partial b(\mathbf{b})}\mathsf{S}_{il}^{-1}\mathsf{S}_{jm}^{-1}\mathsf{S}_{kn}^{-1}\frac{\partial\tilde{\mathsf{B}}^{lmn}}{\partial b(\mathbf{b}')} \tag{41}$$

for $h \equiv \mathsf{S}^{-1}\tilde{a}$, which is optimal in the limit of weak non-Gaussianity, and $\mathsf{S}^{-1} \to \tilde{\mathsf{C}}^{-1}$. In the above, we sum over all bins with $b_1 \le b_2 \le b_3$, and note that the numerator contains both a three- and one-field term. The latter does not affect the mean, but ensures optimality on large scales, and takes a slightly different form from the Hermite tensor definitions (24) since we have introduced a generic weighting S^{-1}. This bears similarities to the estimators of [e.g., [13,14,19,23,24,26,36]], but now includes full treatment of masks and weights, and incorporates arbitrary ℓ-space binning. We discuss its practical implementation below.

A. Idealized form

Assuming rotational invariance, a unit beam, and a trivial mask, (41) can be simplified by first rewriting the numerator in harmonic space

$$\hat{b}_{\text{ideal}}(\mathbf{b}) \propto \frac{1}{3!}\sum_{\ell_im_i}\frac{\partial B^{\ell_1\ell_2\ell_3}_{(-m_1)(-m_2)(-m_3)}}{\partial b(\mathbf{b})}$$
$$\times[h_{\ell_1m_1}h_{\ell_2m_2}h_{\ell_3m_3} - (h_{\ell_1m_1}\langle h_{\ell_2m_2}h_{\ell_3m_3}\rangle$$
$$+ 2\text{ perms})], \tag{42}$$

summing over all ℓ_i and m_i with $i \in \{1,2,3\}$. Inserting the relation between the ideal harmonic-space bispectrum and the binned form $b(\mathbf{b})$ given in (15) yields

$$\hat{b}_{\text{ideal}}(\mathbf{b}) \propto \frac{1}{6\Delta_3(\mathbf{b})}\sum_{\ell_im_i}\mathcal{G}^{\ell_1\ell_2\ell_3}_{m_1m_2m_3}$$
$$\times[\Theta_{\ell_1}(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3) + 5\text{ perms}]$$
$$\times[h_{\ell_1m_1}h_{\ell_2m_2}h_{\ell_3m_3} - (h_{\ell_1m_1}\langle h_{\ell_2m_2}h_{\ell_3m_3}\rangle$$
$$+ 2\text{ perms})], \tag{43}$$

where $\mathcal{G}$ is the Gaunt factor. Assuming a diagonal choice for S^{-1}, the one-field terms require $\ell_2 = \ell_3$, and thus, by the triangle conditions $\ell_1 = 0$ (or permutations thereof); in the ideal limit, they may thus be dropped.

To implement (43) one could perform the ℓ, m summation directly, though this has $\mathcal{O}(\ell_{\max}^6)$ complexity. A more efficient scheme (first considered in [13]), is to rewrite the Gaunt factor as a spherical harmonic integral using (6); this separates the three (ℓ_i, m_i) terms, yielding

$$\hat{b}_{\text{ideal}}(\mathbf{b}) \propto \frac{1}{\Delta_3(\mathbf{b})}\sum_{\ell_im_i}\int d\hat{\mathbf{n}}\, Y_{\ell_1m_1}(\hat{\mathbf{n}})Y_{\ell_2m_2}(\hat{\mathbf{n}})Y_{\ell_3m_3}(\hat{\mathbf{n}})$$
$$\Theta_{\ell_1}(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3)h_{\ell_1m_1}h_{\ell_2m_2}h_{\ell_3m_3} \tag{44}$$

additionally absorbing the permutation symmetries and dropping a factor of $(-1)^{\ell_1+\ell_2+\ell_3}$, noting that statistically isotropic temperature correlators must be are parity-even. Defining fields

$$H_b^{\text{ideal}}[x](\hat{\mathbf{n}}) = \sum_{\ell m}\Theta_\ell(b)\frac{x_{\ell m}}{S_\ell}Y_{\ell m}(\hat{\mathbf{n}}) \Leftrightarrow H_{b,\ell m}^{\text{ideal}}[x]$$
$$= \Theta_\ell(b)\frac{x_{\ell m}}{S_\ell}, \tag{45}$$

this can be written

$$\hat{b}_{\text{ideal}}(\mathbf{b}) \propto \frac{1}{\Delta_3(\mathbf{b})}\int d\hat{\mathbf{n}}\, H_{b_1}^{\text{ideal}}[a](\hat{\mathbf{n}})H_{b_2}^{\text{ideal}}[a](\hat{\mathbf{n}})H_{b_3}^{\text{ideal}}[a](\hat{\mathbf{n}}). \tag{46}$$

This is efficient to compute, requiring just one harmonic transform per choice of b, and a summation for each choice of $\mathbf{b}$; as such, the leading scaling is $\mathcal{O}(N_\ell)$, recalling that N_ℓ is the number of one-dimensional ℓ-bins.

For the Fisher matrix, working in harmonic-space for diagonal S^{-1}, we can write

$$\mathcal{F}_{3,\text{ideal}}(\mathbf{b},\mathbf{b}') = \frac{1}{6}\sum_{\ell_im_i}\frac{\partial B^{\ell_1\ell_2\ell_3}_{(-m_1)(-m_2)(-m_3)}}{\partial b(\mathbf{b})}S_{\ell_1}^{-1}S_{\ell_2}^{-1}S_{\ell_3}^{-1}$$
$$\times\frac{\partial B^{\ell_1\ell_2\ell_3}_{m_1m_2m_3}}{\partial b(\mathbf{b}')}. \tag{47}$$

Inserting the binned definition, we will have a sum over two Gaunt factors, which evaluates to

$$\sum_{m_1m_2m_3}[\mathcal{G}^{\ell_1\ell_2\ell_3}_{m_1m_2m_3}]^2 = \frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi}$$
$$\times\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ 0 & 0 & 0\end{pmatrix}^2, \tag{48}$$

and a sum over permutations of binning functions, which evaluates to

$$\Theta_{\ell_1}(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3)[\Theta_{\ell_1}(b_1')\Theta_{\ell_2}(b_2')\Theta_{\ell_3}(b_3')$$
$$+ 5\text{ perms}] = \Delta_3(\mathbf{b})\delta^{\text{K}}_{\mathbf{bb}'}, \tag{49}$$

recalling that bins are ordered and nonoverlapping, such that $\Theta_\ell(b)\Theta_\ell(b') = \delta^{\text{K}}_{bb'}$. Just as for the power spectrum, the Fisher matrix is diagonal in $\mathbf{b}$, and can be evaluate as a triple sum over ℓ_i, which has $\mathcal{O}(\ell_{\max}^3)$ complexity.

Collecting results the ideal bispectrum estimator becomes

$$\hat{b}_{\text{ideal}}(\mathbf{b}) = \frac{1}{\Delta_3(\mathbf{b})} \mathcal{F}_{3,\text{ideal}}^{-1}(\mathbf{b}) \int d\hat{\mathbf{n}} H_{b_1}^{\text{ideal}}(\hat{\mathbf{n}}) H_{b_2}^{\text{ideal}}(\hat{\mathbf{n}}) H_{b_3}^{\text{ideal}}(\hat{\mathbf{n}})$$

$$\mathcal{F}_{3,\text{ideal}}(\mathbf{b}) = \frac{1}{\Delta_3(\mathbf{b})} \sum_{\ell_1\ell_2\ell_3} \frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi} \frac{\Theta_{\ell_1}(b_1)}{S_{\ell_1}} \frac{\Theta_{\ell_2}(b_2)}{S_{\ell_2}} \frac{\Theta_{\ell_3}(b_3)}{S_{\ell_3}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (50)$$

where H^{ideal} is defined in (45) and the Fisher matrix is equal to the estimator variance if $S_\ell = C_\ell$.

B. General form

1. Numerator

The general estimator can be derived in a similar manner to the ideal case. Working in harmonic space, the numerator is akin to (42), but includes window functions (due to the $\tilde{B}$ correlator), '

$$\hat{b}(\mathbf{b}) \propto \frac{1}{3!} \sum_{\ell_i m_i} \frac{\partial B_{(-m_1)(-m_2)(-m_3)}^{\ell_1\ell_2\ell_3}}{\partial b(\mathbf{b})} [[Wh]_{\ell_1 m_1}[Wh]_{\ell_2 m_2}[Wh]_{\ell_3 m_3}$$

$$- ([Wh]_{\ell_1 m_1}\langle[Wh]_{\ell_2 m_2}[Wh]_{\ell_3 m_3}\rangle + 2 \text{ perms})], \quad (51)$$

where the derivative includes the beam, $B_{\ell_1}B_{\ell_2}B_{\ell_3}$. Inserting the bispectrum derivative and rewriting the Gaunt factor as an integral, the three-field term takes a similar form to before,

$$\hat{b}^{(3)}(\mathbf{b}) \propto \frac{1}{\Delta_3(\mathbf{b})} \int d\hat{\mathbf{n}} H_{b_1}[a](\hat{\mathbf{n}}) H_{b_2}[a](\hat{\mathbf{n}}) H_{b_3}[a](\hat{\mathbf{n}}), \quad (52)$$

where H is now defined as

$$H_b[x](\hat{\mathbf{n}}) = \sum_{\ell m}[WS^{-1}x]_{\ell m} B_\ell \Theta_\ell(b) Y_{\ell m}(\hat{\mathbf{n}})$$

$$\Leftrightarrow H_{b,\ell m}[x] = [WS^{-1}x]_{\ell m} B_\ell \Theta_\ell(b). \quad (53)$$

In the presence of a mask, the one-field term is nontrivial, but can be computed via a Monte Carlo average. Defining a set of fields $\{\alpha\}$ with covariance $\tilde{C}_\alpha$, we can write

$$\hat{b}^{(1)}(\mathbf{b}) = -\frac{1}{\Delta_3(\mathbf{b})} \int d\hat{\mathbf{n}} H_{b_1}[a](\hat{\mathbf{n}}) \langle H_{b_2}[\alpha](\hat{\mathbf{n}}) H_{b_3}[\alpha](\hat{\mathbf{n}})\rangle_\alpha$$

$$+ 2 \text{ perms}, \quad (54)$$

where the average is taken over the random fields. For the estimator to be optimal, $\tilde{C}_\alpha$ should be equal to the data covariance $\tilde{C}$; however, given that $\langle a \rangle = 0$, the estimator does not become biased if this condition is not satisfied. This is in contrast with the trispectrum estimators of Sec. VI, which require accurate random simulations to remove the disconnected contributions. This has the same computational scalings as the ideal numerator (linear in N_ℓ), but with runtime additionally proportional to the number of MC simulations, N_{MC}.

2. Fisher matrix

In the nonideal case, the Fisher matrix is difficult to compute analytically. As for the power spectrum (Sec. IV B), we can use a Monte Carlo procedure, first writing the covariance in real-space,

$$\mathcal{F}_3(\mathbf{b}, \mathbf{b}') = \frac{1}{6} \frac{\partial \tilde{B}^{ijk}}{\partial b(\mathbf{b})} S_{il}^{-1} S_{jm}^{-1} S_{kn}^{-1} \frac{\partial \tilde{B}^{lmn}}{\partial b(\mathbf{b}')}$$

$$= \frac{1}{12} \frac{\partial \tilde{B}^{ijk}}{\partial b(\mathbf{b})} S_{il}^{-1} S_{jj'}^{-1} S_{kk'}^{-1} [U^{j'm'}U^{k'n'}$$

$$+ U^{j'n'}U^{k'm'}]U_{m'm}^{-1} U_{n'n}^{-1} \frac{\partial \tilde{B}^{lmn}}{\partial b(\mathbf{b}')}, \quad (55)$$

inserting two copies of the identity matrix in the second line, for arbitrary invertible matrix U. As for the power spectrum, this is symmetric only if $S^{-1} = S^{-T}$. The Fisher matrix can be evaluated by introducing a set of GRFs $\{u\}$ with covariance U, noting that the quantity inside the square brackets is equal to $\langle u^{j'}u^{k'}u^{m'}u^{n'}\rangle - \langle u^{j'}u^{k'}\rangle\langle u^{m'}u^{n'}\rangle$. In this case, the Fisher matrix becomes

$$\mathcal{F}_3(\mathbf{b}, \mathbf{b}') = \frac{1}{12} \int d\hat{\mathbf{n}} \, d\hat{\mathbf{n}}' \langle Q_3[S^{-1}u, S^{-1}u](\hat{\mathbf{n}}; \mathbf{b})[WS^{-1}W](\hat{\mathbf{n}}, \hat{\mathbf{n}}')Q_3[U^{-1}u, U^{-1}u](\hat{\mathbf{n}}'; \mathbf{b}')\rangle_u$$

$$- \frac{1}{12} \int d\hat{\mathbf{n}} \, d\hat{\mathbf{n}}' \langle Q_3[S^{-1}u, S^{-1}u](\hat{\mathbf{n}}; \mathbf{b})\rangle_u[WS^{-1}W](\hat{\mathbf{n}}, \hat{\mathbf{n}}')\langle Q_3[U^{-1}u, U^{-1}u](\hat{\mathbf{n}}'; \mathbf{b}')\rangle_u, \quad (56)$$

where we have introduced the map [analogous to (38) for the power spectrum],

$$Q_3[x, y](\hat{\mathbf{n}}^i; \mathbf{b}) \equiv \frac{\partial B^{ijk}}{\partial b(\mathbf{b})}[Wx]_j[Wx]_k. \quad (57)$$

Inserting the bispectrum derivative and converting to harmonic space, we find

$$Q_{3,\ell m}[x,y](\mathbf{b}) = \frac{1}{\Delta_3(\mathbf{b})} \sum_{\ell_2 \ell_3 m_2 m_3} \mathcal{G}^{\ell \ell_2 \ell_3}_{m m_2 m_3} [Wx]^*_{\ell_2 m_2} [Wy]^*_{\ell_3 m_3} B_\ell B_{\ell_2} B_{\ell_3} [\Theta_\ell(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3) + 5 \text{ perms.}]$$

$$= \frac{2}{\Delta_3(\mathbf{b})} B_\ell \Theta_\ell(b_1) \int d\hat{\mathbf{n}} Y^*_{\ell m}(\hat{\mathbf{n}}) H[x](\hat{\mathbf{n}};b_2) H[y](\hat{\mathbf{n}};b_3) + 2 \text{ perms.}, \quad (58)$$

where we inserted the integral form of the Gaunt factor (6) in the second line, and used the H maps defined in (53). This is straightforwardly evaluated as a harmonic transform.

Whilst possible to implement (56) is somewhat unwieldy, since it requires the average of a map, $Q_3(\hat{\mathbf{n}})$, over a set of random fields, which is expensive to store (though [54] took this approach). Instead, one may proceed by introducing *two* (uncorrelated) sets of random fields $\{u_1\}$ and $\{u_2\}$ with covariance U, as in [40]. These can be combined to form the following symmetric combination:

$$\alpha(\langle u_1^{j'} u_1^{k'} u_1^{m'} u_1^{n'} \rangle + \langle u_2^{j'} u_2^{k'} u_2^{m'} u_2^{n'} \rangle) + \beta(\langle u_1^{j'} u_1^{k'} u_2^{m'} u_2^{n'} \rangle + \langle u_2^{j'} u_2^{k'} u_1^{m'} u_1^{n'} \rangle); \quad (59)$$

this is equal to the combination of U covariances appearing in (55) if $\alpha = -\beta = 1/2$.[8] Defining

$$F_3^{ab,cd}(\mathbf{b},\mathbf{b}') = \frac{1}{12} \int d\hat{\mathbf{n}} \, d\hat{\mathbf{n}}' \langle Q_3[\mathsf{S}^{-1}u_a, \mathsf{S}^{-1}u_b](\hat{\mathbf{n}};\mathbf{b})[W\mathsf{S}^{-1}W](\hat{\mathbf{n}},\hat{\mathbf{n}}')Q_3[\mathsf{U}^{-1}u_c,\mathsf{U}^{-1}u_d](\hat{\mathbf{n}}';\mathbf{b}')\rangle_{u_a,u_b,u_c,u_d}, \quad (60)$$

we can write

$$\mathcal{F}_3(\mathbf{b},\mathbf{b}') = \frac{1}{2}(F_3^{11,11}(\mathbf{b},\mathbf{b}') + F_3^{22,22}(\mathbf{b},\mathbf{b}')) - \frac{1}{2}(F_3^{11,22}(\mathbf{b},\mathbf{b}') + F_3^{22,11}(\mathbf{b},\mathbf{b}')), \quad (61)$$

which makes efficient use of the Monte Carlo simulations. Computation requires N_ℓ H maps to be computed, which are combined into N_{bin} $Q_{\ell m}$ maps, involving $\mathcal{O}(N_{\text{MC}}N_\ell^2)$ harmonic transforms. These are then combined via map-space summation, yielding an estimator for the Fisher matrix that is again linear in N_{bin} (in the large-memory limit), and proportional to the number of Fisher simulations, N_{fish} (which are analyzed independently). Note that there is no scaling with ℓ_{max}, except for that incurred by the choice of HEALPix N_{side}.

In summary, the optimal window-free bispectrum estimator is given by

$$\hat{b}(\mathbf{b}) = \sum_{\mathbf{b}'} \frac{\mathcal{F}_3^{-1}(\mathbf{b},\mathbf{b}')}{\Delta_3(\mathbf{b}')} \int d\hat{\mathbf{n}} [H_{b_1'}[a](\hat{\mathbf{n}})H_{b_2'}[a](\hat{\mathbf{n}})H_{b_3'}[a](\hat{\mathbf{n}}) - (\langle H_{b_1'}[\alpha](\hat{\mathbf{n}})H_{b_2'}[\alpha](\hat{\mathbf{n}})\rangle_\alpha H_{b_3'}[a](\hat{\mathbf{n}}) + 2 \text{ perms})], \quad (62)$$

with the Fisher matrix defined in (61) subject to the Q definitions of (58). This is straightforward to implement and optimal in the limit of $\mathsf{S}^{-1} \to \tilde{\mathsf{C}}^{-1}$ and weak non-Gaussianity.

VI. OPTIMAL TRISPECTRUM ESTIMATION

Finally, let us consider optimal estimation of the full-sky trispectrum. Unlike for lower-order statistics, this has rarely been considered previously (though see [27,40,41] for notable examples) and the impact of masks has not been carefully assessed. Furthermore, as noted in Sec. II, the trispectrum contains both a parity-even and a parity-odd part; the estimators below are the first to measure the latter part.

As discussed in Sec. III, the general trispectrum estimator takes the form

$$\hat{t}(\mathbf{b},B) = \frac{1}{4!} \sum_{\mathbf{b}'} \mathcal{F}_4^{-1}(\mathbf{b},B;\mathbf{b}',B') \frac{\partial \tilde{T}^{ijkl}}{\partial t(\mathbf{b},B')} [h_i h_j h_k h_l - (h_i h_j \langle h_k h_l \rangle + 5 \text{ perms}) + (\langle h_i h_j \rangle \langle h_k h_l \rangle + 2 \text{ perms})]$$

$$\mathcal{F}_4(\mathbf{b},B;\mathbf{b}',B') = \frac{1}{4!} \frac{\partial \tilde{T}^{ijkl}}{\partial t(\mathbf{b},B)} \mathsf{S}^{-1}_{im} \mathsf{S}^{-1}_{jn} \mathsf{S}^{-1}_{ko} \mathsf{S}^{-1}_{lp} \frac{\partial \tilde{T}^{mnop}}{\partial t(\mathbf{b}',B')}, \quad (63)$$

---

[8]For full generality, we could include a third set of terms of the form $\langle u_1^{j'} u_2^{k'} u_1^{m'} u_2^{n'} \rangle$ and permutations thereof. The inclusion of these may lead to a slight reduction in the number of Monte Carlo simulations required, but we neglect them for simplicity here.

for $h \equiv S^{-1}\tilde{a}$. As noted in Sec. II B, the bins satisfy $b_1 \leq b_2$, $b_3 \leq b_4$, $b_1 \leq b_3$, and, if $b_1 = b_3$, $b_2 \leq b_4$, as well as a diagonal L, binned in some bin B (satisfying triangle conditions on $\{b_1, b_2, B\}$ and $\{b_3, b_4, B\}$). Similarly to before, this is optimal in the limit of vanishing non-Gaussianity and $S^{-1} \rightarrow \tilde{C}^{-1}$.

Estimator (63) contains a four-, two-, and a zero-field term; unlike for the bispectrum, all terms are nontrivial, as they subtract off the mean of the signal. One exception to this is the ideal parity-odd trispectrum: since parity-violation only appears at fourth-order for scalars, the disconnected terms vanish in the ideal limit, making this contribution somewhat easier to estimate. In the below, we will consider estimators for both the parity-even and parity-odd trispectra below, which will be denoted $t_\pm(\mathbf{b}, L)$. We caution that the parity-odd components are purely imaginary, thus their Fisher matrix is negative definite (and equal to the negative of the covariance, if odd- and even-modes are uncorrelated).

A. Idealized form

1. Four-field term

In the ideal limit, the trispectrum numerator can be written in harmonic space as

$$\hat{t}_{\pm,\text{ideal}}(\mathbf{b}, B) \propto \frac{1}{4!} \sum_{\ell_i m_i} \frac{\partial T^{\ell_1 \ell_2 \ell_3 \ell_4}_{(-m_1)(-m_2)(-m_3)(-m_4)}}{\partial t_\pm(\mathbf{b}, B)} [h_{\ell_1 m_1} h_{\ell_2 m_2} h_{\ell_3 m_3} h_{\ell_4 m_4} - (h_{\ell_1 m_1} h_{\ell_2 m_2} \langle h_{\ell_3 m_3} h_{\ell_4 m_4} \rangle + 5 \text{ perms.})$$
$$+ (\langle h_{\ell_1 m_1} h_{\ell_2 m_2} \rangle \langle h_{\ell_3 m_3} h_{\ell_4 m_4} \rangle + 2 \text{ perms.})]. \tag{64}$$

Inserting the explicit definition of the binned trispectra (17), the four-field term can be written

$$\hat{t}^{(4)}_{\pm,\text{ideal}}(\mathbf{b}, B) \propto \frac{1}{\Delta_4(\mathbf{b})} \sum_{\ell_i m_i} (-1)^{\ell_{1234}} \sum_{LM} (-1)^M w^{L(-M)}_{\ell_1 \ell_2 m_1 m_2} w^{LM}_{\ell_3 \ell_4 m_3 m_4} \Theta_L(B) \left[\frac{1 \pm (-1)^{\ell_{1234}}}{2} \right] \Theta_{\ell_1}(b_1) \cdots \Theta_{\ell_4}(b_4) h_{\ell_1 m_1} \cdots h_{\ell_4 m_4}, \tag{65}$$

where we have noted that all 24 permutations are equivalent (due to the symmetry of the four h fields), and explicitly restricted to even or odd $\ell_{1234} \equiv \ell_1 + \ell_2 + \ell_3 + \ell_4$. By expanding the square bracket, this can be split into two coupled pieces:

$$\hat{t}^{(4)}_{\pm,\text{ideal}}(\mathbf{b}, B) \propto \pm \frac{1}{2\Delta_4(\mathbf{b})} \sum_{LM} (-1)^M \Theta_L(B) [A^{\text{ideal}}_{b_1 b_2}(L, -M) A^{\text{ideal}}_{b_3 b_4}(L, M) \pm \bar{A}^{\text{ideal}}_{b_1 b_2}(L, -M) \bar{A}^{\text{ideal}}_{b_3 b_4}(L, M)], \tag{66}$$

subject to the definitions

$$A^{\text{ideal}}_{b_1 b_2}(L, M) = \sum_{\ell_1 \ell_2 m_1 m_2} w^{LM}_{\ell_1 \ell_2 m_1 m_2} \Theta_{\ell_1}(b_1) \Theta_{\ell_2}(b_2) h_{\ell_1 m_1} h_{\ell_2 m_2}$$
$$\bar{A}^{\text{ideal}}_{b_1 b_2}(L, M) = \sum_{\ell_1 \ell_2 m_1 m_2} (-1)^{\ell_1 + \ell_2 + L} w^{LM}_{\ell_1 \ell_2 m_1 m_2} \Theta_{\ell_1}(b_1) \Theta_{\ell_2}(b_2) h_{\ell_1 m_1} h_{\ell_2 m_2}, \tag{67}$$

which are symmetric under $b_1 \leftrightarrow b_2$. The separable form given in (66) is significantly more efficient than a naive estimation using (65), with computation scaling as $\ell^6_{\max}$ instead of $\ell^{10}_{\max}$ for some global maximum scale $\ell_{\max}$ [given that each of (L, M) coefficient involves $\mathcal{O}(\ell^4_{\max})$ terms, and there are $\mathcal{O}(N^2_\ell) = \mathcal{O}(\ell^2_{\max})$ such pieces].

Rather than performing the sum over ℓ_i, m_i explicitly, it is preferred to compute A and $\bar{A}$ by first rewriting the weighting function in terms of spin-weighted spherical harmonics, as in (11). Inserting this relation, we find

$$A^{\text{ideal}}_{b_1 b_2}(L, M) = \int d\hat{\mathbf{n}}_{-2} Y_{LM}(\hat{\mathbf{n}}) \left[\sum_{\ell_1 m_1} h_{\ell_1 m_1} \Theta_{\ell_1}(b_1)_{+1} Y_{\ell_1 m_1}(\hat{\mathbf{n}}) \right] \left[\sum_{\ell_2 m_2} h_{\ell_2 m_2} \Theta_{\ell_2}(b_2)_{+1} Y_{\ell_2 m_2}(\hat{\mathbf{n}}) \right]$$
$$\equiv \int d\hat{\mathbf{n}}_{-2} Y_{LM}(\hat{\mathbf{n}}) H^+_{b_1,\text{ideal}}(\hat{\mathbf{n}}) H^+_{b_2,\text{ideal}}(\hat{\mathbf{n}}), \tag{68}$$

and similarly

$$\bar{A}^{\text{ideal}}_{b_1 b_2}(L, M) = (-1)^L \int d\hat{\mathbf{n}}\, _{-2}Y_{LM}(\hat{\mathbf{n}}) \bar{H}^+_{b_1,\text{ideal}}(\hat{\mathbf{n}}) \bar{H}^+_{b_2,\text{ideal}}(\hat{\mathbf{n}}), \tag{69}$$

defining the spin-weighted fields

$$H^{\pm}_{b,\text{ideal}}(\hat{\mathbf{n}}) = \sum_{\ell m} h_{\ell m} \Theta_{\ell}(b)_{\pm 1}Y_{\ell m}(\hat{\mathbf{n}}), \quad \bar{H}^{\pm}_{b,\text{ideal}} = \sum_{\ell m} (-1)^{\ell} h_{\ell m} \Theta_{\ell}(b)_{\pm 1}Y_{\ell m}(\hat{\mathbf{n}}). \tag{70}$$

The H fields satisfy the following identity

$$[H^+_{b,\text{ideal}}(\hat{\mathbf{n}})]^* = -\sum_{\ell m} h_{\ell m} \Theta_{\ell}(b)_{-1}Y_{\ell m}(\hat{\mathbf{n}}) \equiv -H^-_{b,\text{ideal}}(\hat{\mathbf{n}}), \tag{71}$$

(using properties of the spin-weighted spherical harmonics and assuming $a(\hat{\mathbf{n}})$ to be real), implying that $A^{\text{ideal},*}_{b_1 b_2}(L, M) = (-1)^M \bar{A}^{\text{ideal}}_{b_1 b_2}(L, -M)$. This has the useful implication that

$$\hat{t}^{(4)}_{\pm,\text{ideal}}(\mathbf{b}, B) \propto \pm \frac{1}{\Delta_4(\mathbf{b})} \sum_{LM} \Theta_L(B) \begin{cases} \text{Re}[\bar{A}^{\text{ideal},*}_{b_1 b_2}(L, M) A^{\text{ideal}}_{b_3 b_4}(L, M)] \\ i\text{Im}[\bar{A}^{\text{ideal},*}_{b_1 b_2}(L, M) A^{\text{ideal}}_{b_3 b_4}(L, M)] \end{cases}, \tag{72}$$

which makes clear that parity-even (odd) trispectra are purely real (imaginary). Additionally, it can be used to write the estimator entirely in terms of $M \geq 0$ modes (noting that codes such as HEALPix generally store only these, by symmetry),

$$\hat{t}^{(4)}_{\pm,\text{ideal}}(\mathbf{b}, B) \propto \pm \frac{1}{2\Delta_4(\mathbf{b})} \sum_{L, M \geq 0} (1 + \delta^{\text{K}}_{M>0}) \Theta_L(B) \begin{cases} \text{Re}[\bar{A}^{\text{ideal},*}_{b_1 b_2}(L, M) A^{\text{ideal}}_{b_3 b_4}(L, M) + A^{\text{ideal},*}_{b_1 b_2}(L, M) \bar{A}^{\text{ideal}}_{b_3 b_4}(L, M)] \\ i\text{Im}[\bar{A}^{\text{ideal},*}_{b_1 b_2}(L, M) A^{\text{ideal}}_{b_3 b_4}(L, M) - A^{\text{ideal},*}_{b_1 b_2}(L, M) \bar{A}^{\text{ideal}}_{b_3 b_4}(L, M)], \end{cases} \tag{73}$$

where the factor involving a Kronecker delta gives 2 if $M > 0$ and 1 else.

Utilizing these relations, we can compute the four-point term by first assembling all possible $H^{\pm}_{b,\text{ideal}}(\hat{\mathbf{n}})$ fields (a total of N_{ℓ}), then combining to form each of the $\mathcal{O}(N_{\ell}^2)$ combinations of $A^{\text{ideal}}_{b_1 b_2}(L, M)$ and performing a pairwise sum over harmonics, restricting to the relevant bin in L. In practice, $H^{\pm}_{\text{ideal}}$ can be obtained via spin-weighted harmonic transform, since $\pm H^{\pm}_{\text{ideal}}(\hat{\mathbf{n}})$ is the map-space spin-± 1 conjugate to the harmonic-space spin-± 1 fields $\pm h_{\ell m} \Theta_{\ell}(b)$. Similarly, $\bar{A}^*_{LM}$ and A^*_{LM} are the harmonic-space spin-± 2 conjugates of the spin-± 2 maps $H^+(\hat{\mathbf{n}})H^+(\hat{\mathbf{n}})$ and $H^-(\hat{\mathbf{n}})H^-(\hat{\mathbf{n}})$ respectively. Thus, the computational cost to form the A fields is $\mathcal{O}(N_{\ell}^2)$, whilst that of the summation is $\mathcal{O}(N_{\text{bin}}) = \mathcal{O}(N_{\ell}^4)$, but does not involve harmonic transforms, thus is not likely to be rate limiting.

2. Aside: Spin definitions

It is interesting to consider why the above decomposition is possible. In the definition of [27], the trispectrum coefficients, $T^{\ell_1 \ell_2}_{\ell_3 \ell_4}(L)$ are defined via

$$T^{\ell_1 \ell_2 \ell_3 \ell_4}_{m_1 m_2 m_3 m_4} = \sum_{LM} (-1)^M \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & L \\ m_3 & m_4 & M \end{pmatrix} T^{\ell_1 \ell_2}_{\ell_3 \ell_4}(L) \tag{74}$$

as in (8). With this definition, the trispectrum estimator will involve terms of the form

$$\sum_{m_1 m_2} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & M \end{pmatrix} h_{\ell_1 m_1} h_{\ell_2 m_2} = \begin{pmatrix} \ell_1 & \ell_2 & L \\ -s_1 & -s_2 & s_{12} \end{pmatrix}^{-1} \sqrt{\frac{4\pi}{(2\ell_1 + 1)(2\ell_2 + 1)(2L + 1)}}$$

$$\times \int d\hat{\mathbf{n}}\, _{-s_{12}}Y_{LM}(\hat{\mathbf{n}}) \left[\sum_{m_1} h_{\ell_1 m_1 s_1}Y_{\ell_1 m_1}(\hat{\mathbf{n}})\right] \left[\sum_{m_2} h_{\ell_2 m_2 s_2}Y_{\ell_2 m_2}(\hat{\mathbf{n}})\right]. \tag{75}$$

On the rhs, we have inserted the spin-weighted Gaunt factor definition (11) for a general set of spins $\{s_1, s_2, s_{12}\}$. This allows the m_i summations to be rewritten as an integral (or equivalently, a set of spin-weighted harmonic transforms); given an appropriate definition for the *reduced* trispectrum coefficients, it also allows us to separate the ℓ_i summations.

To perform the above trick, we must carefully choose the spins. In particular, we require the $3j$ symbol to be nonzero for all ℓ of interest. Assuming $\ell_i \geq 2$, $|\ell_1 - \ell_2| \leq L \leq \ell_1 + \ell_2$, and $|s_i| \leq \ell_i$, one might consider $\{s_1, s_2, s_{12}\} = \{0, 0, 0\}, \{\pm 1, \mp 1, 0\}, \{\pm 2, \mp 2, 0\},$ $\{\pm 1, \pm 1, \mp 2\}$. Whilst the former choice matches that used in the bispectrum, it requires even $\ell_1 + \ell_2 + L = 0$, and thus cannot be used for the parity-odd trispectrum. Similarly, the second and third vanish upon symmetrization, thus we here utilize the third, fixing $s_1 = s_2 = -1$ and

$s_{12} = 2.$[9] To this end, we absorb the first line on the rhs of (75) into the trispectrum definition, yielding the reduced trispectrum of (10), and allowing separation of the ℓ_i summations.

3. Two-field term

The two-field term can be obtained by first noting that $\langle h_{\ell m} h_{\ell' m'} \rangle = (-1)^m \delta^K_{\ell \ell'} \delta^K_{m(-m')} C_\ell / S^2_\ell$, assuming uniform weights S_ℓ. As such, the estimator takes the form

$$\hat{t}^{(2)}_{\pm,\text{ideal}}(\mathbf{b}, B) \propto -\frac{1}{\Delta_4(\mathbf{b})} \sum_{\ell_i m_i} (-1)^{\ell_{1234}} \sum_{LM} (-1)^M w^{L(-M)}_{\ell_1 \ell_2 m_1 m_2} w^{LM}_{\ell_3 \ell_4 m_3 m_4} \Theta_L(B)$$

$$\times \left[\frac{1 \pm (-1)^{\ell_{1234}}}{2} \right] \Theta_{\ell_1}(b_1) \cdots \Theta_{\ell_4}(b_4) \left[h_{\ell_1 m_1} h_{\ell_2 m_2} (-1)^{m_3} \delta^K_{\ell_3 \ell_4} \delta^K_{m_3(-m_4)} \frac{C_{\ell_3}}{S^2_{\ell_3}} + 5 \text{ perms} \right]. \tag{76}$$

Due to the Kronecker deltas, the first two permutations contain the term

$$\sum_{m_3 M} (-1)^{m_3 + M} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_3 & L \\ m_3 & -m_3 & M \end{pmatrix} \propto \begin{pmatrix} \ell_1 & \ell_2 & 0 \\ m_1 & m_2 & 0 \end{pmatrix} \delta^K_{L0} \delta^K_{M0}. \tag{77}$$

using properties of Wigner $3j$ symbols [68] and separating out part of the w^{LM} weighting matrices. Since we restrict to $L \geq 2$, this term vanishes always. The other four permutations contain contributions of the form

$$\sum_{m_1 M} (-1)^{m_2} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_1 & L \\ m_3 & -m_1 & M \end{pmatrix} \propto \delta^K_{\ell_2 \ell_3} \delta^K_{m_2(-m_3)}. \tag{78}$$

In both cases, two pairs of momenta are restricted to be equal, thus ℓ_{1234} is even, and any parity-odd contribution to the trispectrum must vanish. For the parity-even part, we find

$$\hat{t}^{(2)}_{+,\text{ideal}}(\mathbf{b}, B) \propto -\frac{1}{\Delta_4(\mathbf{b})} (\delta^K_{b_1 b_4} \delta^K_{b_2 b_3} + \delta^K_{b_1 b_3} \delta^K_{b_2 b_4}) \sum_{\ell_1 \ell_2 L} \frac{(2\ell_1 + 1)(2L + 1)}{4\pi} \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2$$

$$\times (-1)^{\ell_1 + \ell_2 + L} \Theta_L(B) (\Theta_{\ell_1}(b_1) \Theta_{\ell_2}(b_2) + \Theta_{\ell_2}(b_1) \Theta_{\ell_1}(b_2)) \frac{C_{\ell_1}}{S^2_{\ell_1}} \sum_{m_2} |h_{\ell_2 m_2}|^2, \tag{79}$$

involving the empirical power spectrum estimate $\sum_{m_2} |h_{\ell_2 m_2}|^2 / (2\ell_2 + 1)$. This scales as $\mathcal{O}(\ell^3_{\max})$.

4. Zero-field term

The zero-field term may be evaluated using a similar prescription. First, we note that this requires two pairs of ℓ_i to be equal: due to the $1 \pm (-1)^{\ell_{1234}}$ term, the odd-piece must vanish. For the even piece, there are only two nontrivial permutations (due to the above arguments removing the $\ell_1 = \ell_2$, $\ell_3 = \ell_4$ term),

$$\hat{t}^{(0)}_{+,\text{ideal}}(\mathbf{b}, B) \propto \frac{1}{\Delta_4(\mathbf{b})} \sum_{\ell_i m_i} (-1)^{\ell_{1234}} \sum_{LM} (-1)^M w^{L(-M)}_{\ell_1 \ell_2 m_1 m_2} w^{LM}_{\ell_3 \ell_4 m_3 m_4} \Theta_L(B)$$

$$\times \Theta_{\ell_1}(b_1) \cdots \Theta_{\ell_4}(b_4) \left[(-1)^{m_1 + m_2} \frac{C_{\ell_1} C_{\ell_2}}{S^2_{\ell_1} S^2_{\ell_2}} (\delta^K_{\ell_1 \ell_3} \delta^K_{\ell_2 \ell_4} \delta^K_{m_1(-m_3)} \delta^K_{m_2(-m_4)} + \delta^K_{\ell_1 \ell_4} \delta^K_{\ell_2 \ell_3} \delta^K_{m_1(-m_4)} \delta^K_{m_2(-m_3)}) \right]. \tag{80}$$

[9]Other choices are possible; these will lead to reduced trispectra differing by powers of $\sqrt{\ell}$.

To simplify this, we note that

$$\sum_{m_1 m_2 M} (-1)^{m_1+m_2+M} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & L \\ -m_1 & -m_2 & M \end{pmatrix} = (-1)^{\ell_1+\ell_2+L}, \tag{81}$$

thus

$$\hat{t}^{(0)}_{+,\text{ideal}}(\mathbf{b}, L) \propto \frac{1}{\Delta_4(\mathbf{b})} (\delta^{\mathrm{K}}_{b_1 b_4} \delta^{\mathrm{K}}_{b_2 b_3} + \delta^{\mathrm{K}}_{b_1 b_3} \delta^{\mathrm{K}}_{b_2 b_4}) \sum_{\ell_1 \ell_2 L} \frac{(2\ell_1+1)(2\ell_2+1)(2L+1)}{4\pi} \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \Theta_L(B)$$

$$\times (-1)^{\ell_1+\ell_2+L} \Theta_{\ell_1}(b_1) \Theta_{\ell_2}(b_2) \frac{C_{\ell_1}}{S^2_{\ell_1}} \frac{C_{\ell_2}}{S^2_{\ell_2}}, \tag{82}$$

which can be straightforwardly computed in $\mathcal{O}(\ell^3_{\text{max}})$ operations.

5. Normalization

We now turn to the trispectrum Fisher matrix. As we shall find below, this is somewhat more complex than for the power spectrum or bispectrum, since there are off-diagonal correlations even in the ideal case, i.e., bins with different **b** can correlate. This arises due to the degeneracy in the quadrilateral definition: there are two choices of diagonal momentum L for any given tetrahedron. As such, the off-diagonal terms will contribute only when $\{b'_1, b'_2, b'_3, b'_4\}$ is some permutation of $\{b_1, b_2, b_3, b_4\}$.

To compute the normalization, we start from (63) and insert the harmonic-space definitions of the binned trispectrum (17), noting that we can absorb a symmetry factor of 24 since T^{ijkl} is fully symmetric under index exchange. This gives

$$\mathcal{F}^{\text{ideal}}_{4\pm}(\mathbf{b}, B; \mathbf{b}', B') = \frac{1}{\Delta_4(\mathbf{b})\Delta_4(\mathbf{b}')} \sum_{\ell_i m_i} (-1)^{\ell_{1234}} \sum_{LL'MM'} (-1)^{M+M'} \left[\frac{1-(-1)^{\ell_{1234}}}{2}\right]^2 S^{-1}_{\ell_1} \cdots S^{-1}_{\ell_4}$$

$$\times \Theta_L(B)\Theta_{L'}(B')\Theta_{\ell_1}(b_1)\cdots\Theta_{\ell_4}(b_4) w^{L(-M)}_{\ell_1 \ell_2 m_1 m_2} w^{LM}_{\ell_3 \ell_4 m_3 m_4}$$

$$\times \{[w^{L'(-M')}_{\ell_1 \ell_2 m_1 m_2} w^{L'M'}_{\ell_3 \ell_4 m_3 m_4} \Theta_{\ell_1}(b'_1)\cdots\Theta_{\ell_4}(b'_4) + 7 \text{ perms}] + (2 \leftrightarrow 3) + (2 \leftrightarrow 4)\}, \tag{83}$$

where the $(-1)^{\ell_{1234}}$ term comes from switching m_i to $(-m_i)$ in one of the trispectrum derivatives. As before, the binning functions satisfy $\Theta_\ell(b)\Theta_\ell(b') = \delta^{\mathrm{K}}_{bb'}\Theta_\ell(b)$, for contiguous bins; this restricts which bins contribute to the coupling matrix. To proceed it is useful to consider the three permutations in the bottom line separately. The first involves

$$\sum_{m_1 m_2} w^{L(-M)}_{\ell_1 \ell_2 m_1 m_2} w^{L'(-M')}_{\ell_1 \ell_2 m_1 m_2} \propto \sum_{m_1 m_2} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & L' \\ m_1 & m_2 & -M' \end{pmatrix} = \frac{1}{2L+1} \delta^{\mathrm{K}}_{LL'} \delta^{\mathrm{K}}_{(-M)M'}, \tag{84}$$

which implies the matrix is diagonal in L. Similarly, the binning functions yield a factor $\delta^{\mathrm{K}}_{\mathbf{bb}'}\Delta_4(\mathbf{b}')$ (noting the selection rules on **b**), leading to the final contribution,

$$\mathcal{F}^{\text{ideal},(a)}_{4\pm}(\mathbf{b}, B; \mathbf{b}', B') = \pm \frac{\delta^{\mathrm{K}}_{\mathbf{bb}'}\delta^{\mathrm{K}}_{BB'}}{\Delta_4(\mathbf{b})} \sum_{\ell_i L} \left[\frac{1 \pm (-1)^{\ell_{1234}}}{2}\right] \Theta_{\ell_1}(b_1)\cdots\Theta_{\ell_4}(b_4)\Theta_L(B) \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \begin{pmatrix} \ell_3 & \ell_4 & L \\ -1 & -1 & 2 \end{pmatrix}^2$$

$$\times \frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)(2\ell_4+1)(2L+1)}{(4\pi)^2} S^{-1}_{\ell_1} \cdots S^{-1}_{\ell_4}. \tag{85}$$

Notably, this factorizes into a piece involving (ℓ_1, ℓ_2, L) and another involving (ℓ_3, ℓ_4, L): as such, computation cost is $\mathcal{O}(N^3_\ell)$.

The other permutations do not require $L = L'$, and thus source a (small) mixing between modes. The second involves the following combination of $3j$ symbols (from the modified Wigner symbols):

$$\sum_{m_1m_2m_3m_4MM'}(-1)^{M+M'}\begin{pmatrix}\ell_1 & \ell_2 & L\\ m_1 & m_2 & -M\end{pmatrix}\begin{pmatrix}\ell_3 & \ell_4 & L\\ m_3 & m_4 & M\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_3 & L'\\ m_1 & m_3 & -M'\end{pmatrix}\begin{pmatrix}\ell_2 & \ell_4 & L'\\ m_2 & m_4 & M'\end{pmatrix}$$

$$= (-1)^{\ell_2+\ell_3}\begin{Bmatrix}L & \ell_1 & \ell_2\\ L' & \ell_4 & \ell_3\end{Bmatrix}, \tag{86}$$

simplifying in terms of a $6j$ symbol in the second line. Similarly, the third has

$$\sum_{m_1m_2m_3m_4MM'}(-1)^{M+M'}\begin{pmatrix}\ell_1 & \ell_2 & L\\ m_1 & m_2 & -M\end{pmatrix}\begin{pmatrix}\ell_3 & \ell_4 & L\\ m_3 & m_4 & M\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_4 & L'\\ m_1 & m_4 & -M'\end{pmatrix}\begin{pmatrix}\ell_3 & \ell_2 & L'\\ m_3 & m_2 & M'\end{pmatrix}$$

$$= (-1)^{L+L'}\begin{Bmatrix}L & \ell_1 & \ell_2\\ L' & \ell_3 & \ell_4\end{Bmatrix}. \tag{87}$$

This leads to the following matrix contributions:

$$\mathcal{F}_{4\pm}^{\text{ideal,(b)}}(\mathbf{b},B;\mathbf{b}',B')$$

$$= \pm\frac{1}{\Delta_4(\mathbf{b})\Delta_4(\mathbf{b}')}[\delta_{b_1b_1'}^{\text{K}}\delta_{b_2b_3'}^{\text{K}}\delta_{b_3b_2'}^{\text{K}}\delta_{b_4b_4'}^{\text{K}} + 7\text{ perms}]\sum_{\ell_iLL'}\left[\frac{1\pm(-1)^{\ell_{1234}}}{2}\right]\Theta_{\ell_1}(b_1)\cdots\Theta_{\ell_4}(b_4)$$

$$\times \Theta_L(B)\Theta_{L'}(B')\frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)(2\ell_4+1)(2L+1)(2L'+1)}{(4\pi)^2}(-1)^{\ell_2+\ell_3}\begin{Bmatrix}L & \ell_1 & \ell_2\\ L' & \ell_4 & \ell_3\end{Bmatrix}$$

$$\times S_{\ell_1}^{-1}\cdots S_{\ell_4}^{-1}\begin{pmatrix}\ell_1 & \ell_2 & L\\ -1 & -1 & 2\end{pmatrix}\begin{pmatrix}\ell_3 & \ell_4 & L\\ -1 & -1 & 2\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_3 & L'\\ -1 & -1 & 2\end{pmatrix}\begin{pmatrix}\ell_2 & \ell_4 & L'\\ -1 & -1 & 2\end{pmatrix}, \tag{88}$$

and

$$\mathcal{F}_{4\pm}^{\text{ideal,(c)}}(\mathbf{b},B;\mathbf{b}',B')$$

$$= \pm\frac{1}{\Delta_4(\mathbf{b})\Delta_4(\mathbf{b}')}[\delta_{b_1b_1'}^{\text{K}}\delta_{b_2b_4'}^{\text{K}}\delta_{b_3b_3'}^{\text{K}}\delta_{b_4b_2'}^{\text{K}} + 7\text{ perms}]\sum_{\ell_iLL'}\left[\frac{1\pm(-1)^{\ell_{1234}}}{2}\right]\Theta_{\ell_1}(b_1)\cdots\Theta_{\ell_4}(b_4)$$

$$\times \Theta_L(B)\Theta_{L'}(B')\frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)(2\ell_4+1)(2L+1)(2L'+1)}{(4\pi)^2}(-1)^{L+L'}\begin{Bmatrix}L & \ell_1 & \ell_2\\ L' & \ell_3 & \ell_4\end{Bmatrix}$$

$$\times S_{\ell_1}^{-1}\cdots S_{\ell_4}^{-1}\begin{pmatrix}\ell_1 & \ell_2 & L\\ -1 & -1 & 2\end{pmatrix}\begin{pmatrix}\ell_3 & \ell_4 & L\\ -1 & -1 & 2\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_4 & L'\\ -1 & -1 & 2\end{pmatrix}\begin{pmatrix}\ell_3 & \ell_2 & L'\\ -1 & -1 & 2\end{pmatrix}. \tag{89}$$

Computation of this scales as $\mathcal{O}(\ell_{\max}^6)$, due to the presence of the Wigner $6j$ symbol.

Combining results, our ideal trispectrum estimators are given by

$$\hat{t}_{+,\text{ideal}}(\mathbf{b},B) = \sum_{\mathbf{b}'B'}\mathcal{F}_{4+}^{\text{ideal},-1}(\mathbf{b},B;\mathbf{b}',B')[\hat{t}_{+,\text{ideal}}^{(4)}(\mathbf{b}',B') + \hat{t}_{+,\text{ideal}}^{(2)}(\mathbf{b}',B') + \hat{t}_{+,\text{ideal}}^{(0)}(\mathbf{b}',B')]$$

$$\hat{t}_{-,\text{ideal}}(\mathbf{b},L) = \sum_{\mathbf{b}'B'}\mathcal{F}_{4-}^{\text{ideal},-1}(\mathbf{b},B;\mathbf{b}',B')\hat{t}_{-,\text{ideal}}^{(4)}(\mathbf{b}',B'), \tag{90}$$

for the parity-even and parity-odd contributions respectively, where the numerators are given in (72), (79) and (82) and the Fisher matrix is a sum of (85), (88) and (89). Note also that there is no correlation between even- and odd-trispectra, since they require even ℓ_{1234} and odd ℓ_{1234} respectively. As before, the Fisher matrix is equal to the estimator variance if $S_\ell = C_\ell$ in the Gaussian limit (or its negative, for the imaginary parity-odd trispectrum).

B. General form

At the final level of complexity we have the binned trispectrum of a masked field. The numerator of this takes a

similar form to the ideal case discussed above, and the Fisher matrix can be computed similarly to that of the bispectrum Sec. V B. However, we note that, in the general case, the two- and zero-field terms in the parity-odd estimator do not vanish, and, at least in principle, there can be nontrivial mixing between odd- and even-parity trispectra induced by the window function. We show how

to account for such effects below, considering each piece of the estimator in turn.

1. Four-field term

Analogously to Sec. VI A 1, the four-field component of the full trispectrum numerator is given by

$$
\hat{t}_{\pm}^{(4)}(\mathbf{b}, B) \propto \frac{1}{24} \sum_{\ell_i m_i} \frac{\partial T^{\ell_1 \cdots \ell_4}_{(-m_1) \cdots (-m_4)}}{\partial t(\mathbf{b}, B)} [Wh]_{\ell_1 m_1} \cdots [Wh]_{\ell_4 m_4}
$$

$$
= \pm \frac{1}{\Delta_4(\mathbf{b})} \sum_{\ell_i m_i} \left[\frac{1 \pm (-1)^{\ell_{1234}}}{2} \right] \sum_{LM} (-1)^M w^{L(-M)}_{\ell_1 \ell_2 m_1 m_2} w^{LM}_{\ell_3 \ell_4 m_3 m_4} B_{\ell_1} B_{\ell_2} B_{\ell_3} B_{\ell_4} \Theta_L(B) \Theta_{\ell_1}(b_1) \cdots \Theta_{\ell_4}(b_4)
$$

$$
\times [Wh]_{\ell_1 m_1} \cdots [Wh]_{\ell_4 m_4}, \tag{91}
$$

inserting the binned trispectrum definition in the second line. This differs only from the ideal case by the replacement $h \to Wh \equiv W\mathsf{S}^{-1}\tilde{a}$, and is similar to the even-parity estimator of [40]. Following similar logic to before, the estimator separates into a more straightforwardly computable form:

$$
\hat{t}_{\pm}^{(4)}(\mathbf{b}, B) \propto \pm \frac{1}{2\Delta_4(\mathbf{b})} \sum_{LM} (-1)^M \Theta_L(B) [A_{b_1 b_2}(L, -M) A_{b_3 b_4}(L, M) \pm \bar{A}_{b_1 b_2}(L, -M) \bar{A}_{b_3 b_4}(L, M)], \tag{92}
$$

which could be expressed as a real or imaginary part as in (72). This uses the (mask-dependent) definitions

$$
A_{b_1 b_2}[x, y](L, M) = \int d\hat{\mathbf{n}}\, _{-2}Y_{LM}(\hat{\mathbf{n}}) H^+_{b_1}[x](\hat{\mathbf{n}}) H^+_{b_2}[y](\hat{\mathbf{n}})
$$

$$
\bar{A}_{b_1 b_2}[x, y](L, M) = (-1)^L \int d\hat{\mathbf{n}}\, _{-2}Y_{LM}(\hat{\mathbf{n}}) \bar{H}^+_{b_1}[x](\hat{\mathbf{n}}) \bar{H}^+_{b_2}[y](\hat{\mathbf{n}})
$$

$$
H^{\pm}_b[x](\hat{\mathbf{n}}) = \sum_{\ell m} [Wx]_{\ell m} B_\ell \Theta_\ell(b)_{\pm 1}Y_{\ell m}(\hat{\mathbf{n}})
$$

$$
\bar{H}^{\pm}_b[x](\hat{\mathbf{n}}) = \sum_{\ell m} (-1)^\ell [Wx]_{\ell m} B_\ell \Theta_\ell(b)_{\pm 1}Y_{\ell m}(\hat{\mathbf{n}}). \tag{93}
$$

These may be computed via weighted spherical harmonic transforms, as discussed in Sec. VI B 1, and differ only by the mask W and the beam B_ℓ. As such, the four-field term can be computed as a set of forward and reverse harmonic transforms, and finally a harmonic space sum in some bin B. As for the ideal case, the computational scaling is $\mathcal{O}(N_\ell^2)$ for the A fields, and $\mathcal{O}(N_{\rm bin})$ for the overall summation.

2. Two-field term

As mentioned above, the two-field term is not guaranteed to vanish in the general parity-odd estimator, nor does it take a simple form in the general parity-even estimator. This is due to multipole mixing induced by the mask: even ℓ_{1234} in the true map does not necessarily correspond to even ℓ_{1234} in the windowed map.[10] For the parity-odd case,

however, the two-field term is likely to be small, assuming a relatively well-behaved window function.

In general, the two-field term is equal to the four-field term but with two of the Monte Carlo fields contracted, i.e., with the replacement $h_i h_j \to \langle h_i h_j \rangle$. As for the bispectrum (Sec. V B), we will compute this by averaging over a set of simulations, $\{\alpha\}$, with covariance $\tilde{\mathsf{C}}_\alpha$.[11] In this case, however, the estimator will be biased if $\tilde{\mathsf{C}}_\alpha$ is not equal to the data covariance $\tilde{\mathsf{C}}$, though, the bias is expected to be small in the parity-odd case, given that the term vanishes in the ideal limit. Furthermore, in the weakly non-Gaussian regime, the disconnected terms are large compared to the connected ones, thus we may require a substantial number of simulations to compute this contribution, to avoid additional sources of variance. We can write the two-field term in the following manner:

---

[10]See [22] for further discussion of this in the context of the parity-odd bispectrum.

[11]Note that there is no requirement for the simulations to have accurate statistics beyond $\tilde{\mathsf{C}}$; this is discussed in Sec. VI B 1.

$$\hat{t}_{\pm}^{(2)}(\mathbf{b}, B) \propto \mp \frac{1}{2\Delta_4(\mathbf{b})} \sum_{LM} (-1)^M \Theta_L(B) \{ A_{b_1 b_2}[h, h](L, -M) \langle A_{b_3 b_4}[\mathsf{S}^{-1}\alpha, \mathsf{S}^{-1}\alpha](L, M) \rangle_{\alpha}$$

$$\pm \bar{A}_{b_1 b_2}[h, h](L, -M) \langle \bar{A}_{b_3 b_4}[\mathsf{S}^{-1}\alpha, \mathsf{S}^{-1}\alpha](L, M) \rangle_{\alpha} \} + 5 \text{ perms}, \tag{94}$$

where the permutations are over positions of the α mocks, arising due to the permutations contained within the binned trispectrum definition. To implement (94), we must compute both $\langle A_{bb'}[\mathsf{S}^{-1}\alpha, \mathsf{S}^{-1}\alpha](L, M) \rangle$ and $\langle A_{bb'}[h, \mathsf{S}^{-1}\alpha](L, -M) A_{b''b'''}[h, \mathsf{S}^{-1}\alpha](L, M) \rangle$; in practice, computation is dominated by the latter, since we must combine the data with each of N_{MC} simulations, with each requiring a harmonic transform per bin pair.

3. Zero-field term

The general zero-field term may be computed analogously, and takes the form

$$\hat{t}_{\pm}^{(0)}(\mathbf{b}, B) \propto \pm \frac{1}{4\Delta_4(\mathbf{b})} \sum_{LM} (-1)^M \Theta_L(B) \{ \langle A_{b_1 b_2}[\mathsf{S}^{-1}\alpha_1, \mathsf{S}^{-1}\alpha_2](L, -M) A_{b_3 b_4}[\mathsf{S}^{-1}\alpha_1, \mathsf{S}^{-1}\alpha_2](L, M) \rangle_{\alpha_1, \alpha_2}$$

$$- \langle \bar{A}_{b_1 b_2}[\mathsf{S}^{-1}\alpha_1, \mathsf{S}^{-1}\alpha_2](L, -M) \bar{A}_{b_3 b_4}[\mathsf{S}^{-1}\alpha_1, \mathsf{S}^{-1}\alpha_2](L, M) \rangle_{\alpha_1, \alpha_2} + 5 \text{ perms} \}, \tag{95}$$

where $\{\alpha_1\}$ and $\{\alpha_2\}$ are two independent sets of simulations with the same covariance, and we sum over their possible locations. If the simulations were Gaussian, one could use only a single set and compute the four-point average via $\langle \alpha^4 \rangle \sim \tilde{\mathsf{C}}_\alpha \tilde{\mathsf{C}}_\alpha$; here, we allow for non-Gaussianities (for example from lensing), thus use only two-point averages.[12]

4. Normalization

Mask-induced multipole mixing can lead to nontrivial leakage between even- and odd-parity trispectra. As such, the general Fisher matrix contains even-even correlations (denoted $\mathcal{F}_{4++}$), odd-odd correlations ($\mathcal{F}_{4--}$) and even-odd correlations ($\mathcal{F}_{4+-}$ and $\mathcal{F}_{4-+}$). Thanks to the optimal estimator formalism, the full trispectrum estimates obtained should be free from this mixing, i.e., the measured parity-odd modes should not contain a parity-even contribution. This is important if one is searching for a signal in the former, and wants to avoid, for example, lensing-based contributions to the latter.

To compute the Fisher matrix, we start from the general relation given in (63), and denote the two parity states by $\lambda, \lambda' \in \{\pm 1\}$,

$$\mathcal{F}_{4\lambda\lambda'}(\mathbf{b}, B; \mathbf{b}', B') = \frac{1}{24} \frac{\partial \tilde{\mathsf{T}}^{ijkl}}{\partial t_\lambda(\mathbf{b}, B)} \mathsf{S}_{im}^{-1} \mathsf{S}_{jn}^{-1} \mathsf{S}_{ko}^{-1} \mathsf{S}_{lp}^{-1}$$

$$\times \frac{\partial \tilde{\mathsf{T}}^{mnop}}{\partial t_{\lambda'}(\mathbf{b}', B')}. \tag{96}$$

As with the bispectrum, this must be significantly simplified to avoid a heinously expensive sum. Whilst one could compute $\mathcal{F}_4$ as the covariance of the unnormalized $\hat{t}$ estimator applied to a set of GRFs, this requires a large number of Monte Carlo simulations to converge and is accurate only in the limit of $\mathsf{S}^{-1} \to \tilde{\mathsf{C}}^{-1}$. Instead (following the bispectrum logic, and [36,40,54]), we can use the following identity:

$$\mathsf{S}_{im}^{-1}[\mathsf{S}_{jn}^{-1}\mathsf{S}_{ko}^{-1}\mathsf{S}_{lp}^{-1} + 5 \text{ perms.}] = \frac{1}{6}\mathsf{S}_{im}^{-1}\mathsf{S}_{jj'}^{-1}\mathsf{S}_{kk'}^{-1}\mathsf{S}_{ll'}^{-1}[\mathsf{U}_{j'n'}\mathsf{U}_{k'o'}\mathsf{U}_{l'p'} + 5 \text{ perms}]\mathsf{U}_{n'n}^{-1}\mathsf{U}_{o'o}^{-1}\mathsf{U}_{p'p}^{-1}$$

$$= \frac{1}{6}\mathsf{S}_{im}^{-1}\mathsf{S}_{jj'}^{-1}\mathsf{S}_{kk'}^{-1}\mathsf{S}_{ll'}^{-1}\mathsf{U}_{n'n}^{-1}\mathsf{U}_{o'o}^{-1}\mathsf{U}_{p'p}^{-1}\langle u_{j'}u_{k'}u_{l'}u_{n'}u_{o'}u_{p'}\rangle_{\mathrm{fc}}, \tag{97}$$

where we have inserted three copies of the unit matrix, for symmetric invertible matrix U and GRFs u, which satisfy $\langle uu^{\mathrm{T}} \rangle = \mathsf{U}$. The correlator has the subscript 'fc' corresponding to 'fully-connected', i.e., we consider only two-point contractions when each of $\{j', k', l'\}$ contracted with one of $\{n', o', p'\}$.

With the decomposition (97), the Fisher matrix can be split into two pieces, connected only by a known matrix, S^{-1}. Explicitly, each takes the form

---

[12]One may also utilize non-Gaussian simulations to remove unwanted trispectra (arising from lensing or noise, for example); this is detailed in [40].

$$\frac{\partial \tilde{\mathsf{T}}^{ijkl}}{\partial t_\lambda(\mathbf{b}, B)} x_j y_k z_l = W(\hat{\mathbf{n}}^i) \sum_{\ell_1 \cdots \ell_4 m_1 \cdots m_4} Y^*_{\ell_1 m_1}(\hat{\mathbf{n}}^i) \frac{\partial T^{\ell_1 \cdots \ell_4}_{(-m_1) \cdots (-m_4)}}{\partial t(\mathbf{b}, B)} [Wx]_{\ell_2 m_2} [Wy]_{\ell_3 m_3} [Wz]_{\ell_4 m_4}$$

$$\equiv W(\hat{\mathbf{n}}^i) Q_{4\lambda}[x, y, z](\hat{\mathbf{n}}^i; \mathbf{b}, B) \tag{98}$$

for some $\{x, y, z\}$, converting the trispectrum to harmonic space, and introducing $Q_{4\pm}$ functions, akin to (57). This function is just a real-space map for each choice of $\mathbf{b}$ and L. Using the above definition, the coupling matrix can be written

$$\mathcal{F}_{4\lambda\lambda'}(\mathbf{b}, B; \mathbf{b}', B') = \frac{1}{144} \int d\hat{\mathbf{n}}\, d\hat{\mathbf{n}}' \langle Q_{4\lambda}[\mathsf{S}^{-1}u, \mathsf{S}^{-1}u, \mathsf{S}^{-1}u](\hat{\mathbf{n}}; \mathbf{b}, B) [W\mathsf{S}^{-1}W](\hat{\mathbf{n}}, \hat{\mathbf{n}}')$$

$$\times Q_{4\lambda'}[\mathsf{U}^{-1}u, \mathsf{U}^{-1}u, \mathsf{U}^{-1}u](\hat{\mathbf{n}}', \mathbf{b}', B') \rangle_{\text{fc}}, \tag{99}$$

which is a Monte Carlo average over realizations of u, akin to (56) for the bispectrum.

To compute the fully-connected correlator, we must subtract off the unwanted correlations. This can be done by introducing multiple sets of GRFs, denoted, $\{u_n\}$, and computing expressions of the form

$$F^{abc,def}_{4\lambda\lambda'} \equiv \frac{1}{144} \int d\hat{\mathbf{n}}\, d\hat{\mathbf{n}}' \langle Q_{4\lambda}[\mathsf{S}^{-1}u_a, \mathsf{S}^{-1}u_b, \mathsf{S}^{-1}u_c](\hat{\mathbf{n}}; \mathbf{b}, B) [W\mathsf{S}^{-1}W](\hat{\mathbf{n}}, \hat{\mathbf{n}}')$$

$$\times Q_{4\lambda'}[\mathsf{U}^{-1}u_d, \mathsf{U}^{-1}u_e, \mathsf{U}^{-1}u_f](\hat{\mathbf{n}}', \mathbf{b}', B') \rangle_{u_a, u_b, u_c, u_d, u_e, u_f}, \tag{100}$$

analogous to those in Sec. V B. Most simply, we could use three such sets, giving $\mathcal{F}_{4\lambda\lambda'}(\mathbf{b}, B; \mathbf{b}', B') = 6F^{123,123}_{4\lambda\lambda'}$, such that only fully connected terms can contribute. As shown in [40], a more efficient way is to instead use two sets of GRFs, and compute the Fisher matrix as

$$\mathcal{F}_{4\lambda\lambda'}(\mathbf{b}, B; \mathbf{b}', B') = \frac{1}{8}[(F^{111,111}_{4\lambda\lambda'} + F^{222,222}_{4\lambda\lambda'}) + 9(F^{112,112}_{4\lambda\lambda'} + F^{122,122}_{4\lambda\lambda'}) - 6(F^{111,122}_{4\lambda\lambda'} + F^{222,112}_{4\lambda\lambda'})], \tag{101}$$

with coefficients chosen to minimize the variance of the $\mathcal{F}_4$ estimate, *i.e.* reduce the number of Monte Carlo simulations required.

We now turn to the computation of $Q_{4\pm}$ maps. First, we insert the explicit trispectrum of (17) into (98), finding

$$Q_{4\pm}[x, y, z](\hat{\mathbf{n}}^i; \mathbf{b}, B) = \pm \frac{1}{\Delta_4(\mathbf{b})} \sum_{\ell_i m_i} Y^*_{\ell_1 m_1}(\hat{\mathbf{n}}^i) \sum_{LM} (-1)^M w^{L(-M)}_{\ell_1 \ell_2 m_1 m_2} w^{LM}_{\ell_3 \ell_4 m_3 m_4} B_{\ell_1} B_{\ell_2} B_{\ell_3} B_{\ell_4} [Wx]_{\ell_2 m_2} [Wy]_{\ell_3 m_3} [Wz]_{\ell_4 m_4} \Theta_L(B)$$

$$\times \left[\frac{1 \pm (-1)^{\ell_{1234}}}{2} \right] [\Theta_{\ell_1}(b_1) \cdots \Theta_{\ell_4}(b_4) + 7 \text{ perms.}] + (2 \leftrightarrow 3) + (2 \leftrightarrow 4). \tag{102}$$

Naïve computation of this expression is highly expensive, due to the large number of coupled ℓ summations. To simplify, we insert the definitions of A and $\bar{A}$ given in (93) and expand the first weighting matrix in terms of spin-weighted spherical harmonics. For the first permutation, this gives

$$Q^{(a)}_{4\pm}[x, y, z](\hat{\mathbf{n}}^i; \mathbf{b}, B) = \pm \frac{1}{2\Delta_4(\mathbf{b})} \sum_{\ell_1 \ell_2 m_1 m_2} Y^*_{\ell_1 m_1}(\hat{\mathbf{n}}^i) \sum_{LM} B_{\ell_1} B_{\ell_2} \Theta_L(B) \Theta_{\ell_1}(b_1) \Theta_{\ell_2}(b_2) [Wx]_{\ell_2 m_2} (-1)^M$$

$$\times \int d\hat{\mathbf{n}}\, {}_{-2}Y_{LM}(\hat{\mathbf{n}})\, {}_{+1}Y_{\ell_1 m_1}(\hat{\mathbf{n}})\, {}_{+1}Y_{\ell_2 m_2}(\hat{\mathbf{n}}) [A_{b_3 b_4}[y, z](L, -M) \pm (-1)^{\ell_1 + \ell_2 + L} \bar{A}_{b_3 b_4}[y, z](L, -M)]. \tag{103}$$

Next, the ℓ_1 and L summations can be written as spin-weighted spherical harmonic transforms. Denoting

$$_s[X]^b(\hat{\mathbf{n}}) = \sum_{\ell m} {}_s Y_{\ell m}(\hat{\mathbf{n}}) \Theta_\ell(b) X_{\ell m}, \tag{104}$$

and using the relation $(-1)^\ell {}_s Y_{\ell m}(\hat{\mathbf{n}}) = {}_{-s} Y_{\ell m}(-\hat{\mathbf{n}})$, we can write the harmonic-space $Q_{4\pm}$ as

$$Q^{(a)}_{4\pm, \ell_1 m_1}[x, y, z](\mathbf{b}, B) = \mp \frac{B_{\ell_1} \Theta_{\ell_1}(b_1)}{2\Delta_4(\mathbf{b})} \left[\int d\hat{\mathbf{n}}\, {}_{-1}Y^*_{\ell_1 m_1}(\hat{\mathbf{n}}) H^+_{b_2}[x](\hat{\mathbf{n}})\, {}_{-2}[\bar{A}^*_{b_3 b_4}[y, z]]^B(\hat{\mathbf{n}}) \right.$$

$$\left. \pm \int d\hat{\mathbf{n}}\, {}_{+1}Y^*_{\ell_1 m_1}(\hat{\mathbf{n}}) H^-_{b_2}[x](\hat{\mathbf{n}})\, {}_{+2}[A^*_{b_3 b_4}[y, z]]^B(\hat{\mathbf{n}}) \right], \tag{105}$$

recalling that $_sY^*_{\ell m} = (-1)^{s+m}{}_{-s}Y_{\ell(-m)}$. Here, the $\hat{\mathbf{n}}$ integral can be evaluated as a further spin-1 spherical harmonic transform. Summing over permutations, we find the final form,

$$Q_{4\pm,\ell m}[x,y,z](\mathbf{b},B) = [Q^{(a)}_{4\pm,\ell m}[x,y,z](\{b_1,b_2,b_3,b_4\},B) + 7 \text{ perms}] + (x \leftrightarrow y) + (x \leftrightarrow z), \quad (106)$$

where the permutations preserve the $\{b_1, b_2\}$ and $\{b_3, b_4\}$ pairs. This involves $\mathcal{O}(N_\ell^4)$ harmonic transforms (since only $\Theta_\ell(b)$ can be separated from the above expression), giving a slightly more favorable scaling than $\mathcal{O}(N_{\text{bin}}) = \mathcal{O}(N_\ell^5)$.

Given $Q_{4\pm,\ell m}$ we can compute $Q_{4\pm}(\hat{\mathbf{n}})$, and thus compute the coupling via (99), multiplying by the filter $W\mathsf{S}^{-1}W$, which requires straightforward transformations between real and harmonic space, with a computational cost of $\mathcal{O}(N_{\text{bin}})$. Computing the above factors is likely the most labor-intensive section of optimal trispectrum estimator, but, thanks to the above simplifications, still scales favorably with the number of Monte Carlo simulations (linearly, and much faster than simply using them to numerically estimate the covariance of the unnormalized estimator [36]), and the number of bins (technically quadratically, but with the rate-limiting pieces (computation of $W\mathsf{S}^{-1}Q_{4\pm}$) scaling linearly).

To summarize, our estimator of the full-sky binned trispectrum is given by

$$\hat{t}_\lambda(\mathbf{b}, B) = \sum_{\mathbf{b}'B'\lambda'} \mathcal{F}^{-1}_{4\lambda\lambda'}(\mathbf{b}, B; \mathbf{b}', B')\{\tau_{\lambda'}[h,h,h,h](\mathbf{b}', B') - 6\langle\tau_{\lambda'}[h,h,\mathsf{S}^{-1}\alpha,\mathsf{S}^{-1}\alpha](\mathbf{b}',B')\rangle_\alpha$$
$$+ 3\langle\tau_{\lambda'}[\mathsf{S}^{-1}\alpha_1,\mathsf{S}^{-1}\alpha_1,\mathsf{S}^{-1}\alpha_2,\mathsf{S}^{-1}\alpha_2](\mathbf{b}',B')\rangle_{\alpha_1,\alpha_2}\}, \quad (107)$$

where the unnormalized estimator can be written explicitly as

$$\tau_\pm[\alpha,\beta,\gamma,\delta](\mathbf{b},B) = \pm\frac{1}{48\Delta_4(\mathbf{b}')}\sum_{LM}(-1)^M\Theta_L(B)\{A_{b_1b_2}[\alpha,\beta](L,-M)A_{b_3b_4}[\gamma,\delta](L,M)$$
$$\pm \bar{A}_{b_1b_2}[\alpha,\beta](L,-M)\bar{A}_{b_3b_4}[\gamma,\delta](L,M)\} + 23 \text{ perms}. \quad (108)$$

where the permutations are over the positions of $\{\alpha, \beta, \gamma, \delta\}$, and α are random fields satisfying $\langle\alpha\alpha^{\mathsf{T}}\rangle = \tilde{\mathsf{C}}$. The A and $\bar{A}$ fields are defined in (93) and the general Fisher matrix is given in (99). If one wishes to ignore the even-odd coupling in the estimator, one just evaluates the above expression fixing $\lambda' = \lambda$.

VII. VALIDATION

In the above sections, we have derived optimal and ideal estimators for the full-sky power spectrum, bispectrum, and (parity-even and odd) trispectrum. To demonstrate their efficacy, we will now consider a variety of tests on synthetic data, both for Gaussian and non-Gaussian maps, optionally including a nontrivial mask. This section makes extensive use of the public POLYBIN code,[13] which implements the above estimators in PYTHON, with harmonic manipulations performed using HEALPix [69]. Spectra can be computed using arbitrary binning schemes, with the option of different binning for squeezed and collapsed configurations. For this purpose, we will specialize to CMB applications, though we note that the tools developed above apply much more generally.

A. Practicalities

To test our estimators, we will primarily use synthetic Gaussian random fields (GRFs) created using HEALPix. These are constructed using the following (statistically isotropic) correlator:

$$\langle a_{\ell m} a_{\ell'm'}\rangle = (-1)^m\delta^{\mathsf{K}}_{\ell\ell'}\delta^{\mathsf{K}}_{m(-m')}[B_\ell^2 C_\ell^{TT} + N_\ell], \quad (109)$$

where C_ℓ^{TT} is the CMB temperature power spectrum predicted by CLASS with the *Planck* best-fit parameters [70], and we set the beam, B_ℓ, to unity. Since we generate and analyze simulations with the same HEALPix N_{side} we do not include a pixel beam. The noise model is given by

$$N_\ell = \Delta_T^2\exp\left(\frac{\ell(\ell+1)\theta_{\text{FWHM}}^2}{8\log 2}\right), \quad (110)$$

where we fix $\Delta_T = 1\mu K$-arcmin and $\theta_{\text{FWHM}} = 5$ arcmin. Though we will usually work with Gaussian fields, we also consider simulations with a synthetic bispectrum injected. These can be obtained following [36], via the transformation on a GRF $a_{\ell m}$

$$a_{\ell m} \rightarrow a_{\ell m} + \frac{1}{6}\mathcal{G}^{\ell\ell_2\ell_3}_{mm_2m_3}b^{\text{theory}}_{\ell\ell_2\ell_3}h^*_{\ell_2m_2}h^*_{\ell_3m_3}, \quad (111)$$

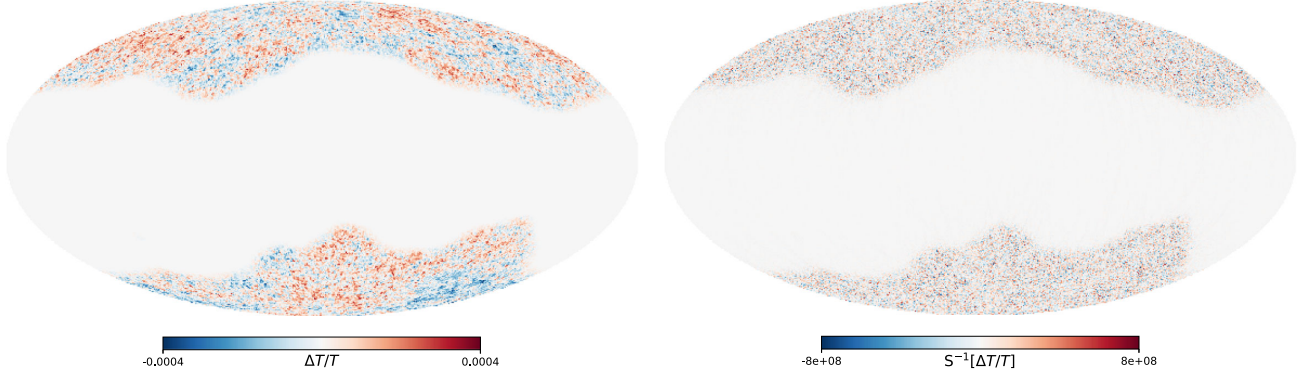[13]GitHub.com/oliverphilcox/PolyBin.

FIG. 1.   Examples of Gaussian full-sky maps used to test the polyspectrum estimators. The left panel shows a dataset constructed with the *Planck* power spectrum and noise parameters and an apodized sky mask. In the right panel, we show the map after application of the quasi-optimal S^{-1} weighting.

where $b_{\ell_1\ell_2\ell_3}^{\text{theory}}$ is the desired bispectrum (see also Eq. (1.3) in [37]) and $h_{\ell m} \equiv [C^{-1}a]_{\ell m}$. Here, we will use the factorized form $b_{\ell_1\ell_2\ell_3} = \prod_{i=1}^{3}\beta_{\ell_i}$, where, for definitiveness, we set $\beta_\ell = 2\exp[(\ell-2)/40]$ (in μK^3 units). Due to the factorization, this can be written

$$a_{\ell m} \to a_{\ell m} + \frac{1}{6}\beta_\ell \int d\hat{\mathbf{n}}\, Y_{\ell m}^*(\hat{\mathbf{n}})[\beta h]^*(\hat{\mathbf{n}})[\beta h]^*(\hat{\mathbf{n}}), \quad (112)$$

using (6) and writing $[\beta h](\hat{\mathbf{n}}) = \sum_{\ell m}\beta_\ell h_{\ell m}Y_{\ell m}(\hat{\mathbf{n}})$, which can be evaluated as a harmonic transform.[14] Finally, we will often consider windowed fields: for this, we utilize a *Planck* 40% Galactic sky mask, with 2° Gaussian apodization, denoted W[15]; this is akin to a (highly anisotropic) window that would be used in a realistic *Planck* analysis, though we pick a somewhat severe example for the sake of demonstration. The full field is given by $d(\hat{\mathbf{n}}) = W(\hat{\mathbf{n}})a(\hat{\mathbf{n}})$, an example of which is shown in Fig. 1.

To form the window-free estimators we require the random fields u and their covariance, U. As noted in Sec. IV C, the estimator is unbiased for any U, however, the Monte Carlo variance can be reduced if U^{-1} is close to the weighting matrix S^{-1}. Here, we will assume a diagonal weighting, such that

$$\langle u_{\ell m}u_{\ell'm'}\rangle = (-1)^m\delta_{\ell\ell'}^{K}\delta_{m(-m')}^{K}[B_\ell^2 C_\ell^{TT} + N_\ell]$$
$$\equiv (-1)^m\delta_{\ell\ell'}^{K}\delta_{m(-m')}^{K}U_\ell, \quad (113)$$

as in (109); if the synthetic data is unwindowed, this matches the true covariance C. With the definition, the action of the U^{-1} weighting on a map β is given by

$$U_{ij}^{-1}\beta^j = \sum_{\ell m}\frac{\beta_{\ell m}}{U_\ell}Y_{\ell m}(\hat{\mathbf{n}}_i), \quad (114)$$

which is straightforwardly computed as a harmonic transform. In practice, one does not deal with continuous maps on the two-sphere, but discrete HEALPix pixels: this affects things only by introducing a factor $A_{\text{pix}} = 4\pi/N_{\text{pix}}$ whenever a summation over pixels is involved.[16] Finally, we must choose a form for the S^{-1} optimality weighting. Here, we use a diagonal approximation (neglecting the window function, except for some rescaling, which cancels), fixing $S^{-1} = U^{-1}$, which we expect to be close to optimal on the scales considered herein. An example of the S^{-1}-filtered data is shown in Fig. 1. As noted in Sec. IV C, an alternative approach would be to omit the window from $\tilde{a}$ (and its correlators), and instead include it as a projection in S^{-1} (*i.e.* treating the true map as the input, and zero-weighting bad regions).[17] This may be more appropriate for real analyses with complex window functions and inpainted maps, and will be the approach used in [59].

B. Power spectrum

We begin by validating the power spectrum estimators of Sec. IV. For this, we apply both the ideal (Sec. IV A) and optimal (Sec. IV B) estimators to a suite of GRF simulations created as described above, optionally including a *Planck* mask. For this test, we focus on comparatively large scales (where the impact of the mask is largest), considering the binned power spectrum in $n_\ell = 100$ linearly spaced bins of width $\Delta\ell = 4$ from $\ell_{\min} = 2$ to $\ell_{\max} = 402$, though we drop the last bin in all cases to mitigate correlations of the extremal bins with their neighbors. Data are constructed using a HEALPix grid of $N_{\text{side}} = 256$, giving $\ell_{\max}^{\text{HEALPIX}} = 767$, far above the scales of interest here. To construct the Fisher

---

[14]We also subtract off the mean of the signal, to ensure that $\langle a_{\ell m}\rangle = 0$.
[15]Available at pla.esac.esa.int/pla.

[16]Here, we neglect discrete pixel weights, which could be included to ensure that U is the exact covariance of u. Assuming that the ℓ range in question is sufficiently small compared to $\ell_{\max}$, this approximation is justified.
[17]See [67] for an alternative approach that allows for invertible covariances.
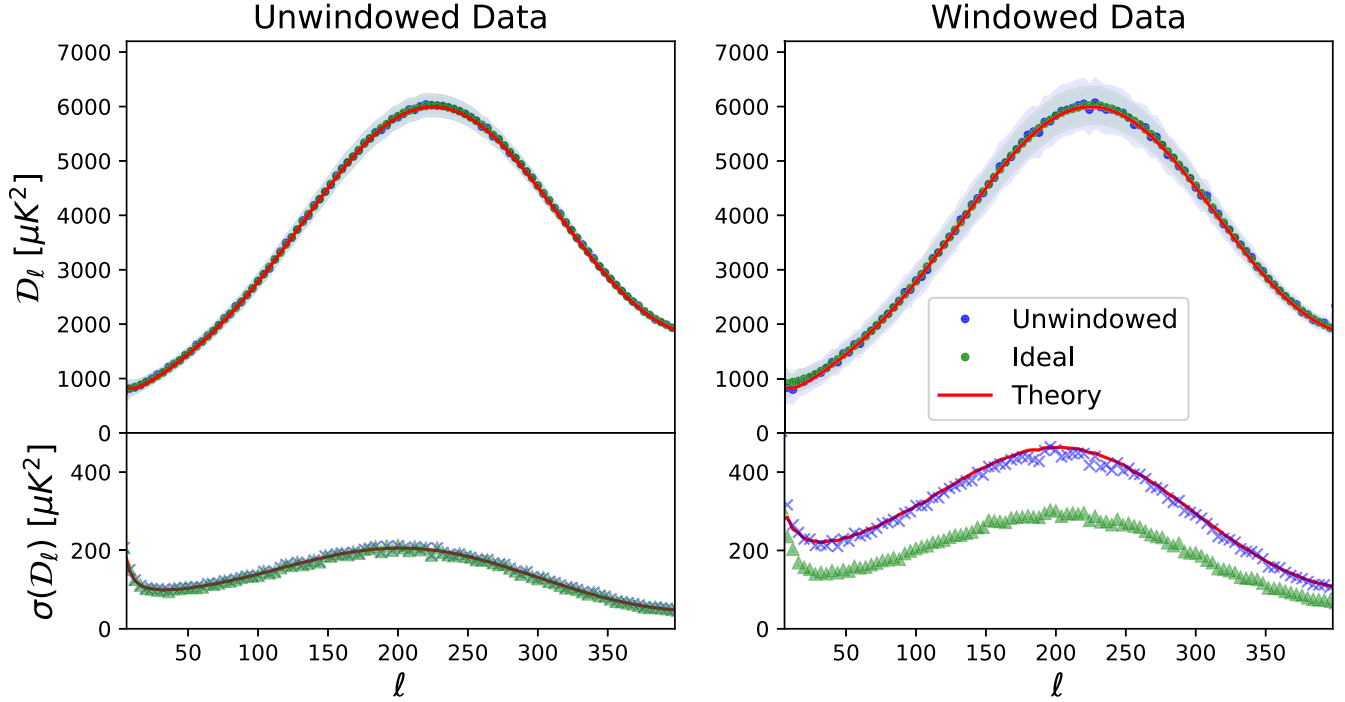
FIG. 2. Comparison of binned power spectrum estimators for unwindowed (left) and windowed (right) data. In each case, we plot the binned power spectrum obtained from the ideal (Sec. IV A, blue) and maximum-likelihood (Sec. IV B, green) estimators, alongside the true theory model obtained from CLASS. The top panels show the raw measurements, in bins of width $\Delta\ell = 4$ [normalized by $\ell(\ell + 1)/(2\pi)$], whilst the bottom panels show the errors. Data is obtained from 1000 Gaussian random field simulations, with the Fisher matrix of the optimal estimators constructed using 100 Monte Carlo realizations, using quasi-optimal weights (cf. Sec. VII A). In all cases the estimators appear unbiased, and the unwindowed variances are almost identical to the inverse Fisher matrix (red lines). For the windowed data, the ideal estimator appears to have lower variance than the optimal estimator: this is due to significant correlations between neighboring bins, as seen in Fig. 3.

matrices required in the optimal estimator, we use $N_{\rm fish} = 100$ simulations. Computation required ≈ 30 CPU-seconds per Fisher realization, and ≈ 0.4 CPU-seconds for each estimator numerator (both for the ideal and optimal approaches); the optimal estimator thus required ≈ 1 CPU-hour of additional time to compute, though we note that this is independent of the number of simulations analyzed.

To compare theory and data, we require some procedure for estimating the binned models from the unbinned spectra $C_\ell^{\rm th}$. An appropriate choice is the following:

$$C^{\rm th}(b) = \left[\sum_\ell \Theta_\ell(b)(2\ell + 1)\frac{C_\ell^{\rm th}}{S_\ell^2}\right] \Big/ \left[\sum_\ell \Theta_\ell(b)(2\ell + 1)\frac{1}{S_\ell^2}\right], \tag{115}$$

derived from considering the expectation of the ideal estimator. This matches the approach of [24] for the bispectrum, but includes our custom weighting S_ℓ (or the diagonal part thereof).

Figure 2 shows the measured binned power spectrum from the unmasked and masked simulations alongside the true injected power spectrum, averaging over 1000 simulations. In both cases, we find excellent agreement

between data and theory, as expected. When the synthetic data does not include a window, the two estimators agree precisely; when a mask is included, the means are consistent, but the variance properties differ. In the latter case, the variance is significantly increased (by a factor of approximately $\langle W^4\rangle/\langle W^2\rangle^2$, due to the reduced area observed), and the ideal estimator seems to considerably outperform the optimal one. This appears paradoxical; however, it occurs since the various bins are correlated in the ideal estimator, but anticorrelated in the optimal approach (with both estimators yielding similar signal-to-noise).

In Fig. 3, we plot the correlation matrices for the two estimators applied to the windowed data-set (noting that the unwindowed case is trivially diagonal). If the optimal estimator is, as the name would suggest, optimal, its covariance should be equal to the inverse of the Fisher matrix, $\mathcal{F}_2$. From Fig. 3 and the lower part of Fig. 2, this is exactly what is observed on all scales, implying that our choice of weighting, S^{-1}, is appropriate.[18] As noted

---

[18]In practice, we find little dependence of the power spectrum measurements on the choice of weighting scheme, which occurs since the data is mostly uncorrelated and the ℓ-bins are narrow [cf. [55]].
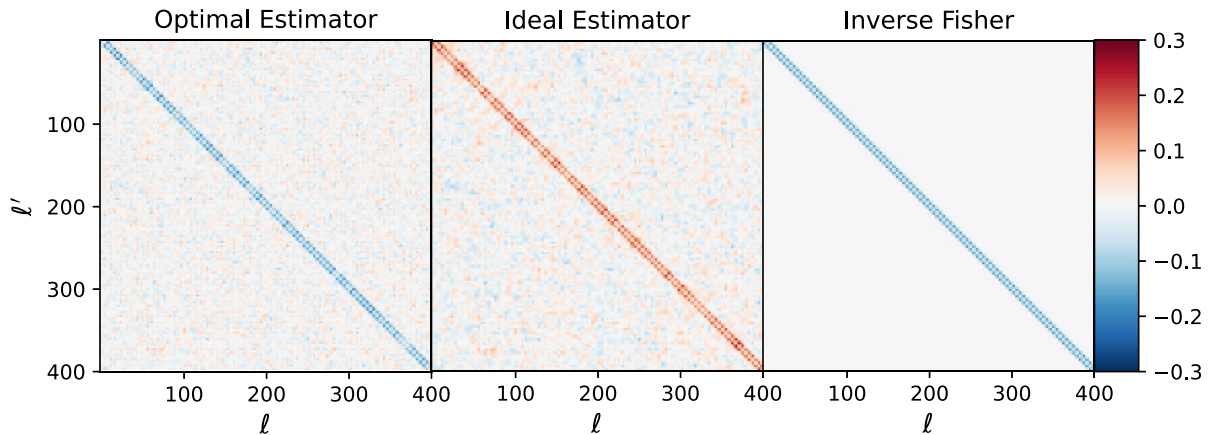
FIG. 3.   Correlation matrices for the ideal and unwindowed power spectrum measurements plotted in Fig. 2 (right panel). In all cases, the underlying data contains a *Planck* sky mask, which, for the conventional (ideal) estimator, gives a clear correlation between neighboring bins. In the optimal estimator, we see an anticorrelation, which is closely matched by the inverse Fisher matrix, as expected. The correlation matrix is defined by $\mathcal{C}_{ij}/\sqrt{\mathcal{C}_{ii}\mathcal{C}_{jj}}$ for covariance $\mathcal{C}_{ij}$, and we subtract off the leading diagonal for clarity.

above, we observe different correlation properties for the optimal and ideal estimators, with a positive correlation between neighboring bins seen in the latter case. One feature of the optimal prescription is that we can naturally form a quantity of unit variance, with no cross-correlations, $\sum_{b'}\mathcal{F}_2^{1/2}(b,b')\hat{C}(b')$, e.g., [51]; we have verified that the optimal estimator correlation matrix of this object shows no obvious departures from the identity matrix beyond that expected from noise fluctuations. Finally, we consider the dependence on the number of Monte Carlo simulations used to define the Fisher matrix (the limiting step in the estimator). Reducing to just ten realizations ($N_{\text{fish}} = 10$) changes the power spectrum predictions by at most 0.25σ, thus we conclude that the above choice of $N_{\text{fish}} = 100$ is both sufficient and conservative.

C. Bispectrum

Next, we turn to the three-point function. Here, we will consider two scenarios: (1) pure GRFs without a window function (to test optimality), and (2) simulations with an injected bispectrum and a mask (to test bias). Due to the higher dimensionality of the three point function, we consider broader (linear) bins, using $\ell_{\min} = 2$, $\Delta\ell = 10$, and $n_\ell = 15$.[19] To avoid edge effects, we will drop any bin containing the largest ℓ values; this reduces the total number of elements in the data vector from 372 to 308. Given the lower $\ell_{\max}$ used in this test, we fix $N_{\text{side}} = 128$, which significantly reduces runtime. Here, we require 240 CPU-seconds to compute the Fisher matrix using a single pair of GRF realizations, and 50 CPU-seconds to compute the estimator numerators. The runtime is dominated by the

100 Monte Carlo simulations (used to compute the one-field term, cf. Sec. V B), but greatly reduced if one analyzes multiple datasets in series (since maps relating to the Monte Carlo simulations do not need to be recomputed). The ideal estimator (which does not include a one-field term) requires only 0.2 CPU-seconds, though with another 35 CPU-seconds to compute the (diagonal) normalization. In this case, one should bin the theory model in the following manner (cf. [24]):

$$
\begin{aligned}
b^{\text{th}}(\mathbf{b}) \propto &\sum_{\ell_{123}} \Theta_{\ell_1}(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3) \\
&\times \frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi} \\
&\times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}^2 \frac{b^{\text{th}}_{\ell_1\ell_2\ell_3}}{S_{\ell_1}S_{\ell_2}S_{\ell_3}},
\end{aligned} \tag{116}
$$

where the normalization factor takes the same form but without $b^{\text{th}}_{\ell_1\ell_2\ell_3}$.

Figure 4 shows our measurements of the reduced bispectrum. In the absence of a signal, we recover null detections (as expected), and find a similar (though not identical) variance between the optimal and ideal estimators, with the optimal estimator performing somewhat better on large scales. When a signal is included, we find unbiased results from both estimators, and, as before, note that the variance of the optimal estimator lies very close to the inverse Fisher matrix (and somewhat higher than the ideal estimator variance, due to bin anticorrelations). This again indicates that the optimal estimator is close to minimum variance.[20]

---

[19]In practice, it may be preferable to use nonlinearly spaced bins, such that the signal-to-noise is more evenly distributed across bins.

[20]Note that this is not guaranteed in this case even if $\mathsf{S}^{-1} = \tilde{\mathsf{C}}^{-1}$, since the field is non-Gaussian, thus the covariance strictly contains a piece proportional to B^2.
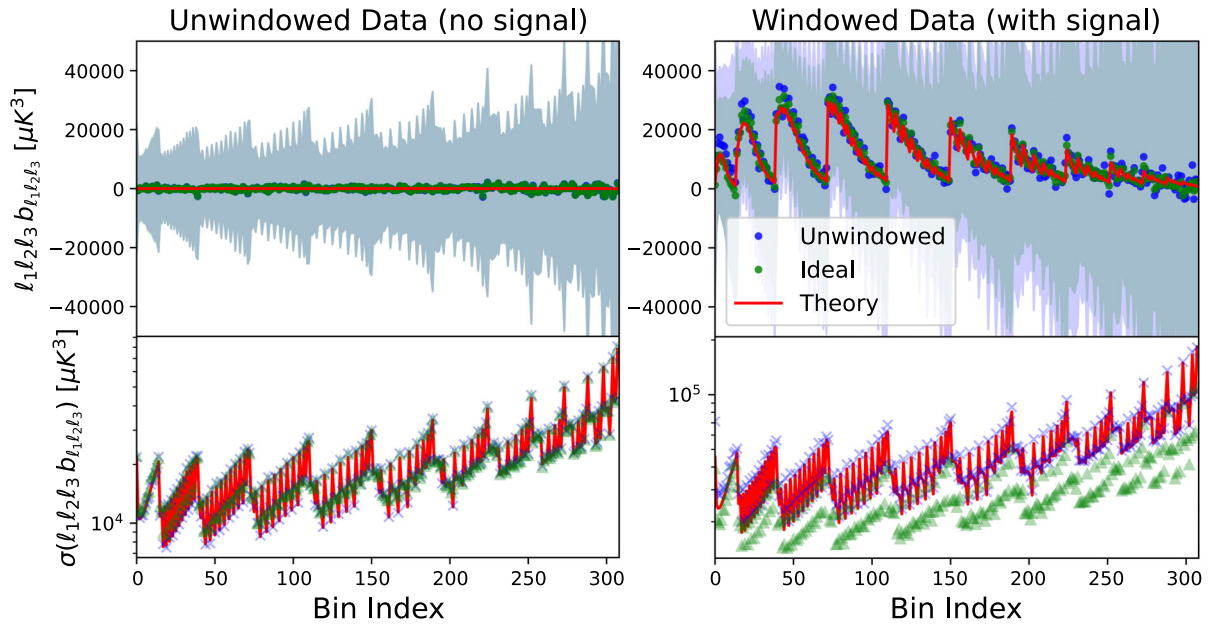
FIG. 4. Comparison of binned bispectrum estimators for unwindowed zero-signal (left) and windowed nonzero signal (right) data. The measurements are akin to those in Fig. 2, but use broader bins with $\Delta\ell = 10$ and $\ell \in [2, 142]$, and we weight the data by $\ell_1\ell_2\ell_3$, averaging over 1000 simulations. Here, we plot all bispectrum bins $\{b_1, b_2, b_3\}$, which satisfying the triangle conditions (at the bin centers) and $b_1 \leq b_2 \leq b_3$. These are collapsed into one dimension for visualization, starting from the lowest ℓ bins on the lhs, and sequentially updating b_3, b_2, then b_1. We see that the estimator is unbiased in both cases, and that the variance of the optimal estimator matches its theoretical prediction (red lines), implying that it is close to minimum variance. The corresponding correlation matrix is shown in Fig. 5.

The correlation matrices shown in Fig. 5 confirm the above results. Here, the mask induces nontrivial correlations between the various bins (the size of which depend on the ratio of the bin width and the characteristic scale of the mask), particularly those with $\mathbf{b}' = \mathbf{b} \pm \{1, 0, 0\}$, or some permutation thereof, extending up to high ℓ (large scales). This structure is well captured by the Fisher matrix, and differs significantly from the (generally positive) correlations of the ideal estimator. In particular, the low-ℓ region

shows strong correlations between a variety of bins, up to $\sim 20\%$. These may be difficult to model, and are not found in the optimal estimator, due to its particular choice of weighting scheme. As before, we find that the Fisher matrix is well-converged; reducing to $N_{\mathrm{fish}} = 10$ biases the bispectrum measurements by at most 0.2σ.

Finally, it is interesting to consider the impact of the linear term in the bispectrum estimator, i.e., that proportional to $\langle\alpha\alpha\rangle a$. As noted above, this term does not
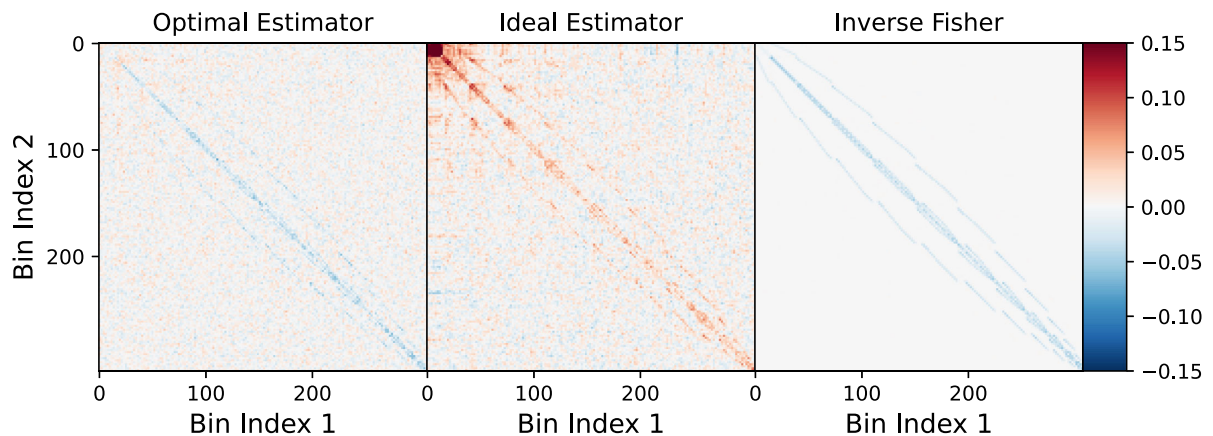

FIG. 5. As Fig. 3, but showing the correlation of the reduced bispectrum measurements displayed in Fig. 4. As before, the covariance of the optimal estimator matches the inverse Fisher matrix to high accuracy; in this case, the correlation structure is more complex due to window-function induced mode coupling. The ideal estimator shows strong correlations between bins, particularly at the top left, corresponding to the lowest bins. These would need to be modeled in any analysis.
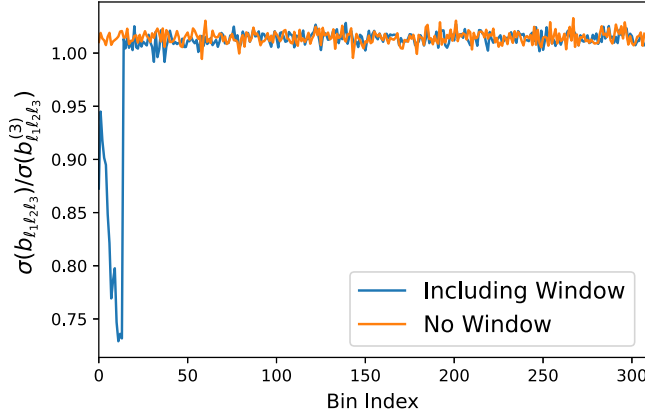
FIG. 6. Impact of the linear term in the bispectrum estimators. We plot the ratio of the errorbars between bispectra estimated including and excluding the one-field term [cf. Eq. (62)], for realizations with (blue) and without (orange) a mask. Significant differences are observed for the first few bins of the windowed data; these correspond to configurations including modes in the lowest ℓ bin.

contribute to the mean of the signal, but can have nontrivial impacts on its covariance. In Fig. 6 we compare the estimator variance both with and without the one-field term. For the first fifteen or so bins, we find a notable reduction in the error bar from including the linear term for windowed data, up to $\approx 25\%$, but essentially no change for the other bispectrum components, nor for unwindowed data. Noting that the affected bins are the only ones to contain the lowest ℓ modes (here with $\ell_1 \in [2, 12]$ and free ℓ_2, ℓ_3), we conclude that the linear term significantly reduces the measurement uncertainty on large scales if a window is present. This matches previous results, see e.g., [30,36] (noting that many nonideal estimators include such a term) and is important if one wishes to constrain large-scale signals, such as those of primordial non-Gaussianity.

D. Trispectrum

Finally, we validate the trispectrum estimator. As discussed in Sec. VI the trispectrum contains two contributions (of even- and odd-parity); here we will measure both simultaneously, and, for the sake of plotting, work with the imaginary part of t_-. Generating realizations with injected non-Gaussianity is nontrivial, especially for the parity-odd terms (though see [40]), though for the parity-even terms, one may consider using lensed simulations, which include a known four-point function. However, to verify the estimators it is sufficient to check that (a) before subtraction of the disconnected terms, the parity-even estimator recovers the Gaussian expectation (i.e., that of the form C_ℓ^2), (b) after subtraction, the estimator is consistent with zero when applied to Gaussian realizations, (c) the estimator variance matches the Fisher prediction. For the parity-odd case, it is usually sufficient to restrict to comparatively large

scales, since (if the underlying theory is statistically isotropic), any parity-violating trispectrum must vanish in the small-scale regime.[21]

As for the bispectrum, the trispectrum is a high-dimensional object, containing $\mathcal{O}(n_\ell^5)$ elements (for $n_\ell \ell$-bins). To keep the computation tractable, we will consider the following binning parameters (using linear bins for simplicity, noting that other choices may be more efficient): $\ell_{\min} = 2$, $\Delta\ell = 20$, $n_\ell = 6$, and drop the largest ℓ-bin to avoid edge effects. We again work at $N_{\rm side} = 128$, which is appropriate for these large-scale modes. In total, we estimate 455 even-parity and 386 odd-parity configurations, which reduces to 249 and 222 when removing the final ℓ bin.[22] Here, we apply our estimator to the 1000 GRF simulations described above, with the Fisher matrix computed using $N_{\rm fish} = 100$ realizations. We also utilize 100 GRFs to compute the disconnected two- and zero-field terms in the optimal estimator. Each Fisher realization requires 40 CPU-minutes to analyze, with the data piece taking $\approx$10 CPU-minutes per simulation, again dominated by the Monte Carlo computations. The ideal numerator is significantly faster, since it does not involve Monte Carlo simulations, and requires only $\approx$1 CPU-second per iteration. We caution however, that the ideal normalization is nontrivial for trispectra, due to its off-diagonal correlators and $6j$ symbols. For the binning parameters discussed above, the ideal Fisher matrix required 24 CPU-hours to compute (after removing the largest ℓ bin); this scales as $\mathcal{O}(\ell_{\max}^6)$, which is prohibitive for large $\ell_{\max}$ (unlike the optimal schemes).

To compare theory and observations for the trispectrum, we should bin the underlying reduced trispectrum $t_{\ell_3\ell_4}^{\ell_1\ell_2,\rm th}(L)$ in the following manner:

$$
t^{\rm th}(\mathbf{b}, B)
$$
$$
\propto \sum_{\ell_{1234}L} \Theta_{\ell_1}(b_1)\Theta_{\ell_2}(b_2)\Theta_{\ell_3}(b_3)\Theta_{\ell_4}(b_4)
$$
$$
\times \frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)(2\ell_4 + 1)(2L + 1)}{(4\pi)^2}
$$
$$
\times \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \begin{pmatrix} \ell_3 & \ell_4 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \frac{t_{\ell_3\ell_4}^{\ell_1\ell_2,\rm th}(L)}{S_{\ell_1}S_{\ell_2}S_{\ell_3}S_{\ell_4}},
$$
$$(117)$$

---

[21]This occurs since, at high-ℓ, the trispectrum is approximately plane-parallel. On $\mathbb{R}^2$, a parity flip is equivalent to a rotation in $\mathbb{R}^3$, and thus trivial if the theory is invariant under rotations. This strictly requires *all* the ℓ modes to be small: in practice, one may wish to include parity-odd modes in the squeezed configuration (depending on the physical models of interest). These can be included by allowing for larger ℓ_2, ℓ_4 and (by the triangle conditions) L.

[22]We recall that the parity-odd estimator vanishes if $b_1 = b_3$ and $b_2 = b_4$, unlike the parity-even case.
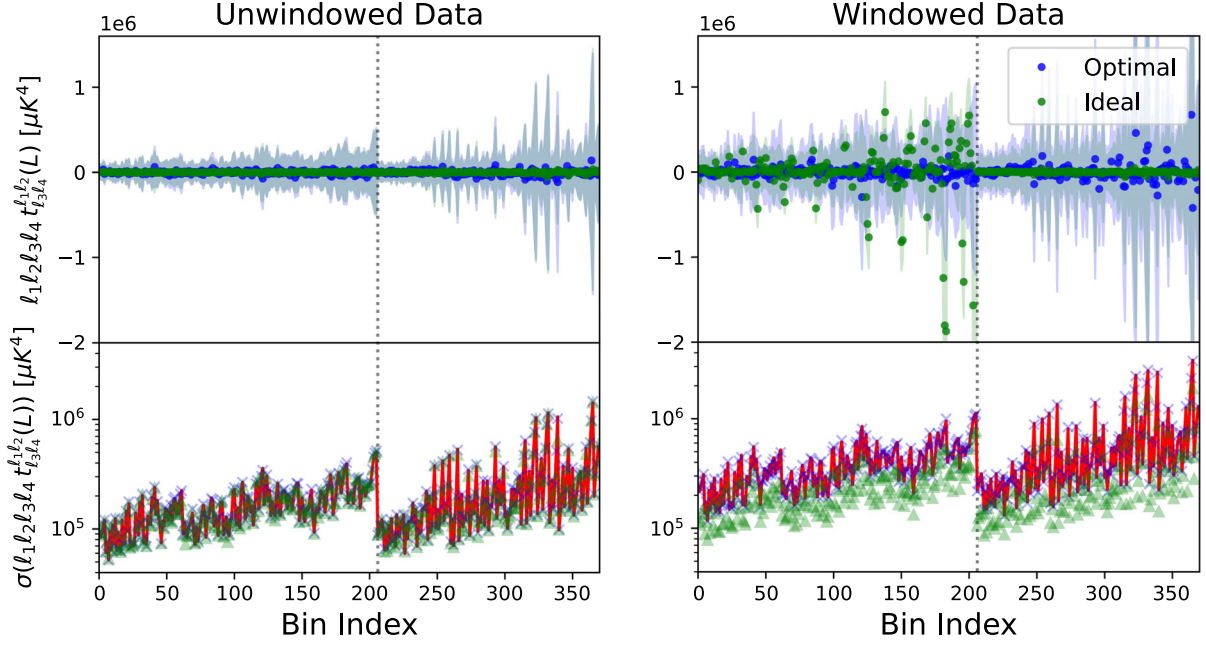
FIG. 7. Comparison of binned trispectrum estimators applied to 1000 unwindowed (left) and windowed (right) Gaussian simulations. The measurements are similar to those of Fig. 4, except with the binning $\Delta\ell = 20$, $\ell \in [2, 102]$ and weighting the data by $\ell_1\ell_2\ell_3\ell_4$. Both parity-even and parity-odd measurements are plotted, demarcated by the vertical dotted lines, condensing all allowed trispectra in bins $\{b_1, b_2, b_3, b_4, B\}$ into a single dimension for visualization. As before, the characteristic ℓ values in the bin gradually increase in size from the left to the right. The mean of both the ideal and optimal estimators appears consistent with zero, and, for the optimal case, the variance matches the inverse Fisher matrix (red lines), as expected. Correlation matrices for the windowed field are shown in Fig. 8.

with an appropriate normalization factor. This is again derived from the expectation of the idealized estimator, but we drop a $6j$ term (which mixes different L and L' modes), which is subdominant, and prevents efficient factorization in $\{\ell_1, \ell_2, L\}$ and $\{\ell_3, \ell_4, L\}$.

In Fig. 7 we show the trispectrum measurements extracted from the GRF realizations. Though detailed interpretation of this plot is hampered by the statistic's high dimensionality, it is clear that both estimators return amplitudes consistent with zero (though there may be some outliers in the ideal windowed scenario, due mask-induced effects). This indicates that the subtraction of the disconnected terms is working as expected. Furthermore, the variances of the optimal estimator are consistent with those predicted by the inverse Fisher matrix, for both the parity-even and parity-odd components (with a ratio of 1.021 ± 0.003). This is shown further in Fig. 8, where we observe that the complex correlation structure of the Fisher matrix matches the covariance of the simulated realizations, implying that the estimator is close to optimal, and that we have used sufficient number of simulations to compute the disconnected terms. When the mask is included, the variance of the estimator increases significantly (roughly by a factor of $\langle W^8 \rangle / \langle W^2 \rangle^4$), and the correlation structure changes, seen particularly in the low-ℓ modes. Unlike the power spectrum and bispectrum

estimators, the covariance of the unmasked fields is non-diagonal; this is in accordance with the discussion of Sec. VI, and is due to the labeling degeneracy, where the diagonal of the quadrilateral $ABCD$ can be placed between sides A and C or B and D. Finally, we note that $N_{\text{fish}} = 100$ is sufficient for Fisher matrix convergent, as before; reducing to N_{fish} gives a (stochastic) bias of at most 0.04σ.

To check the parity-even estimator in more detail it is useful to examine the disconnected terms. In the ideal Gaussian limit, $\mathbb{E}[t_4] = -(1/2)\mathbb{E}[t_2] = \mathbb{E}[t_0]$, and the unnormalized estimators satisfy

$$\mathbb{E}[t_0(\mathbf{b}, L)] \propto \frac{1}{\Delta_4(\mathbf{b})} \sum_{\ell_1\ell_2} \frac{(2\ell_1 + 1)(2\ell_2 + 1)(2L + 1)}{4\pi}$$
$$\times \begin{pmatrix} \ell_1 & \ell_2 & L \\ -1 & -1 & 2 \end{pmatrix}^2 \frac{(-1)^{\ell_1 + \ell_2 + L}}{C_{\ell_1} C_{\ell_2}}$$
$$\times (\delta^{\text{K}}_{b_1 b_3} \delta^{\text{K}}_{b_2 b_4} + \delta^{\text{K}}_{b_1 b_4} \delta^{\text{K}}_{b_2 b_3}). \quad (118)$$

In Fig. 9, we plot the various terms entering the trispectrum numerators, for both the unwindowed and windowed estimators. In the former case, we observe excellent agreement between the disconnected pieces and (118), whilst for the latter (for which theoretical predictions are nontrivial),
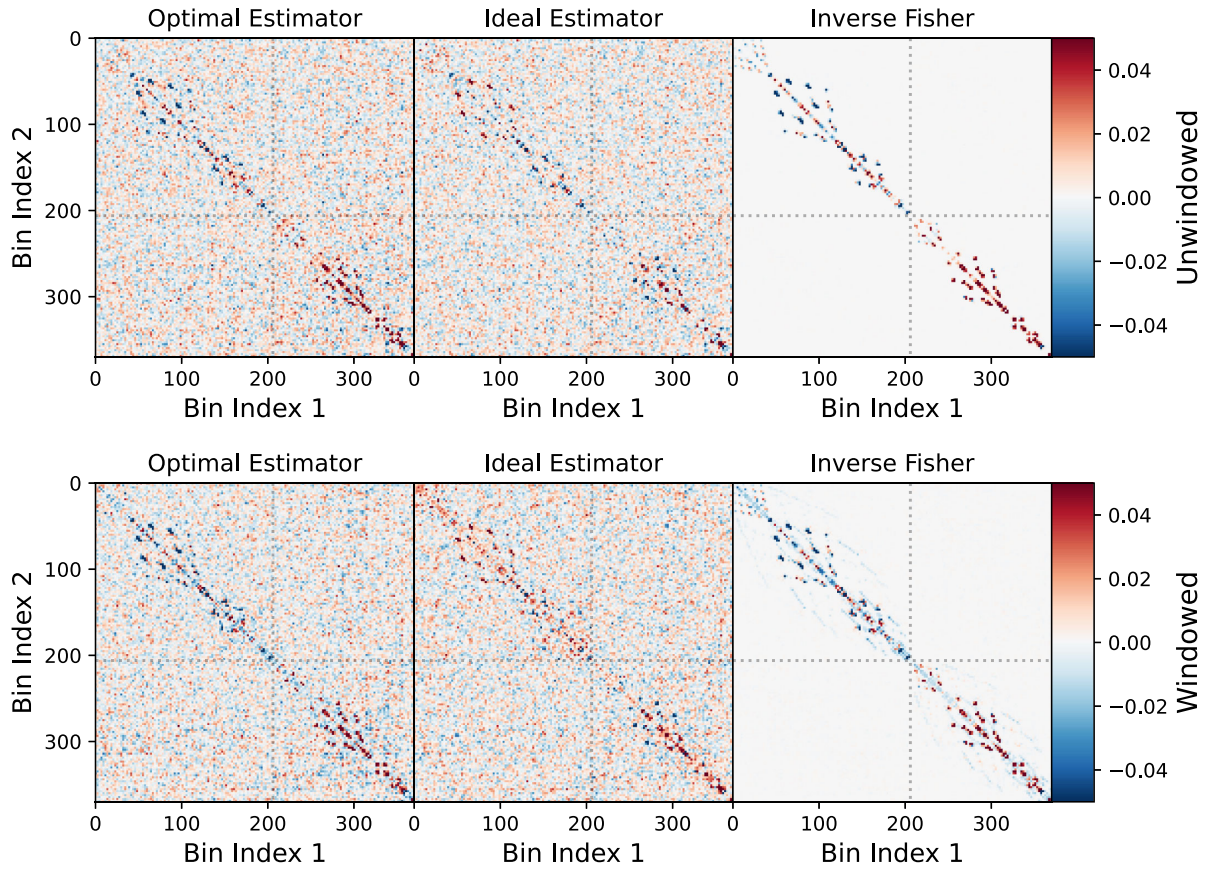
FIG. 8. As Fig. 5, but showing the correlation of the reduced trispectrum measurements displayed in Fig. 7. We utilize the same binning strategy as before, with the parity-even modes shown in the top left and the parity-odd in the bottom right of each matrix. We show results both for unwindowed (top) and windowed (bottom) data, noting that the correlation structure is nontrivial in both cases, due to degeneracies within the trispectrum definition. Regardless of the mask, the inverse Fisher matrix closely matches the optimal estimator covariance, indicating that the estimator is approaching the maximum likelihood solution. The mask is seen to induce non-negligible correlations on small scales, though we do not find significant mixing between the parity-even and parity-odd trispectra.
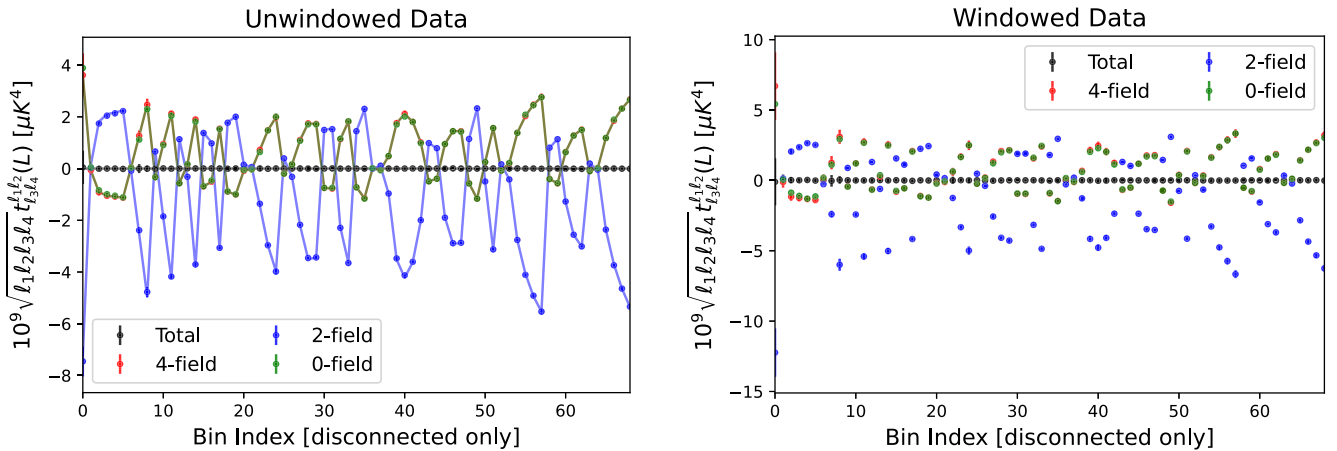
FIG. 9. Comparison of the optimal trispectrum estimators, specializing to the bins with a nontrivial contribution to the disconnected (two- and zero-field) terms. The left and right panels show results for unwindowed and windowed data respectively, with colors discriminating the various terms. In each case, we plot the relevant term in the trispectrum numerator, normalized by the unwindowed Fisher matrix for visualization. For the left panel, we plot also the theoretical prediction, as given in (118). The disconnected terms closely match their expectations and, as expected, cancel when combined to yield a zero detection of the connected trispectrum. In most cases, the error bars are too small to discern.
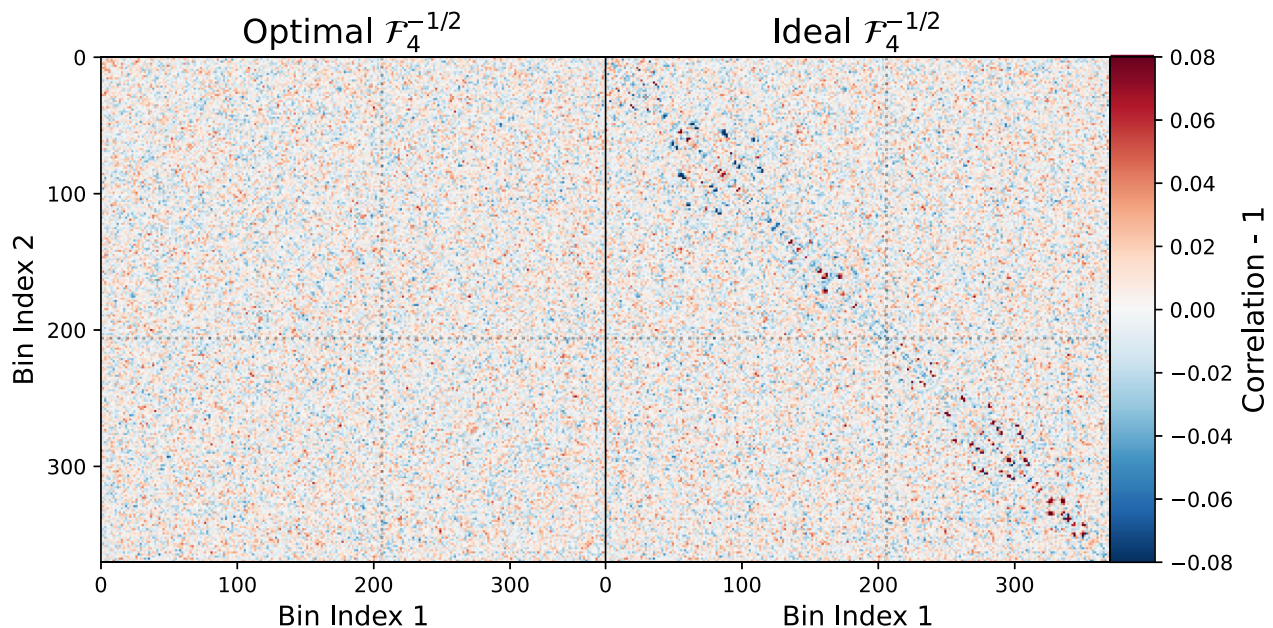
FIG. 10.   Correlation matrix of the windowed trispectrum dataset, weighted by the Cholesky factorization of $\mathcal{F}_4$. The left and right panels show the results using $\mathcal{F}_4$ matrices obtained from optimal and ideal estimators, subtracting the leading diagonal in each case. If the estimator is ideal, this matrix should be consistent with unity. Here, we find good results for the optimal estimators (with a variance of 0.98 ± 0.04) but clear structure in the ideal case (with a variance of 5.8 ± 1.4). This again indicates that the optimal estimators are close to minimum variance, and sources a useful projection of the data.

we see similar behavior as a function of scale. Crucially, whilst the disconnected terms themselves are large, their sum is negligible; this indicates that the estimators are performing as expected, and do not yield a false detection.

Our final consistency check is shown in Fig. 10. Here, we consider the quantity $\hat{t}^{\mathrm{uncorr}} \equiv \mathcal{F}_4^{-\mathrm{T}/2}\hat{t}$, which, as noted in [51], follows a unit Gaussian distribution if the estimator is optimal. In this limit, its covariance would be simply be equal to the identity matrix. From the figure, we find that, using the Fisher matrix obtained from the optimal estimator, the covariance of $\hat{t}^{\mathrm{uncorr}}$ (on the masked dataset) is indistinguishable from a unit normal, and thus the estimator is close to optimal. If one instead uses the 'ideal' Fisher matrix $\mathcal{F}_4$ (which can be obtained without Monte Carlo methods, albeit with large computational costs), we find clear structure to the correlation matrix, due to the impact of the window functions on the Fisher matrix. This decomposition also provides a useful projection scheme; for realistic scenarios (including non-Gaussian effects such as CMB lensing), the various bins of $\hat{t}^{\mathrm{uncorr}}$ are expected to remain almost uncorrelated.

VIII. CONCLUSIONS

Through the measurement and interpretation of random processes, we can understand the physics of a wide variety of phenomena. Correlation functions, or polyspectra, are a key tool with which to do this, allowing for the rich statistics of a stochastic field to be expressed in terms of

low-dimensional functions. In this work, we consider the measurement of such quantities for fields on the two-sphere, relevant to a range of disciplines including cosmology and geophysics. In particular, we derive estimators for the two-, three- and four-point correlators (power spectra, bispectra, and trispectra, respectively), and discuss how they may be efficiently applied to isotropic data. We consider two classes of estimators: 'ideal' and 'optimal'. The first match standard definitions in the literature, and are derived under ideal assumptions, *i.e.* assuming isotropic noise without masks. In contrast, our optimal estimators defined by maximizing the theoretical likelihood for the masked data (including beams), which yields a number of useful properties. These include:

(i) *Optimality*: Assuming that field is close to Gaussian, the variance of the optimal estimators takes its minimum value. Strictly, this is true only if the data is optimally weighted: we have additionally considered close-to-optimal weighting schemes that come close to saturating this bound in realistic scenarios.

(ii) *Bias*: The optimal estimators are unbiased, such that their expectation is equal to the true underlying statistic, regardless of the survey mask and (isotropic) beam. This allows the measurements to be directly compared to data, unlike for the ideal estimators, for which the window should be included in the theory model, requiring a complex convolution.

(iii) *Separability*: Since we specialize to binned polyspectra, the estimators can be efficiently computed through a set of spherical harmonic transforms.

The accompanying Fisher matrices may be estimated via Monte Carlo methods, which are shown to converge quickly.

(iv) *Computational Efficiency*: Computation of the estimator numerators involves sets of harmonic transforms scaling as N_ℓ (for the power spectrum and bispectrum) and N_ℓ^2 (for the trispectrum), as well as a summation scaling as $N_{\rm bin}$, for $N_\ell\ell$-bins and $N_{\rm bin}$ total bins. Similarly, the rate limiting step of the optimal Fisher matrix estimator has the scaling $\mathcal{O}(N_{\rm bin}N_{\rm fish})$, unlike naïve $\mathcal{O}(N_{\rm bin}^2)$ expectations, utilizing $N_{\rm fish} \sim 10$–100 simulations.

To facilitate general use, we have implemented the above estimators in a publicly available PYTHON package, which has been extensively tested in Sec. VII. These could be used for a number of applications, including general (model-independent) non-Gaussian analyses of the cosmic microwave background (CMB) or cosmic shear. A particularly exciting prospect concerns the parity-odd trispectrum. Utilizing these estimators, we robustly measure the statistic, taking into account subtleties such as the leakage of disconnected terms and parity-even modes, and thus place the first CMB-derived constraints on scalar parity-violation in the Universe. This will be discussed in [59]. Naturally, many other applications are possible.

ACKNOWLEDGMENTS

[1] P. J. E. Peebles, *The Large-Scale Structure of the Universe* (Princeton University Press, Princeton, NJ, 1980).

[2] O. H. E. Philcox and S. Torquato, Phys. Rev. X **13**, 011038 (2023).

[3] S. Garrett-Roe and P. Hamm, J. Chem. Phys. **128**, 104507 (2008).

[4] O. H. E. Philcox, J. Hou, and Z. Slepian, arXiv:2108.01670.

[5] J. G. Berryman, J. Comput. Phys. **75**, 86 (1988).

[6] V. S. Dotsenko, Mod. Phys. Lett. A **06**, 3601 (1991).

[7] K. Hwang, B. Schmittmann, and R. K. P. Zia, Phys. Rev. E **48**, 800 (1993).

[8] F. Šanda and S. Mukamel, Phys. Rev. E **72**, 031108 (2005).

[9] *Analysis of Geophysical Potential Fields*, Advances in Exploration Geophysics Vol. 5, edited by P. S. Naidu and M. Mathew (Elsevier, New York, 1998), pp. 75–143.

[10] J. A. Peacock, *Cosmological Physics* (Cambridge University Press, Cambridge, UK, 1999).

[11] A. H. Guth, Phys. Rev. D **23**, 347 (1981).

[12] A. D. Linde, Phys. Lett. **108B**, 389 (1982).

[13] E. Komatsu, D. N. Spergel, and B. D. Wandelt, Astrophys. J. **634**, 14 (2005).

[14] E. Komatsu, B. D. Wandelt, D. N. Spergel, A. J. Banday, and K. M. Gorski, Astrophys. J. **566**, 19 (2002).

[15] Y. Akrami *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A10 (2020).

[16] A. J. Duivenvoorden, P. D. Meerburg, and K. Freese, Phys. Rev. D **102**, 023521 (2020).

[17] Y. Akrami *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A9 (2020).

[18] A. F. Heavens, Mon. Not. R. Astron. Soc. **299**, 805 (1998).

[19] A. P. S. Yadav, E. Komatsu, B. D. Wandelt, M. Liguori, F. K. Hansen, and S. Matarrese, Astrophys. J. **678**, 578 (2008).

[20] D. Munshi, A. Heavens, A. Cooray, J. Smidt, P. Coles, and P. Serra, Mon. Not. R. Astron. Soc. **412**, 1993 (2011).

[21] D. Munshi and A. Heavens, Mon. Not. R. Astron. Soc. **401**, 2406 (2010).

[22] W. R. Coulton and D. N. Spergel, J. Cosmol. Astropart. Phys. 10 (2019) 056.

[23] M. Bucher, B. Van Tent, and C. S. Carvalho, Mon. Not. R. Astron. Soc. **407**, 2193 (2010).

[24] M. Bucher, B. Racine, and B. van Tent, J. Cosmol. Astropart. Phys. 05 (2016) 055.

[25] J. R. Fergusson and E. P. S. Shellard, Phys. Rev. D **80**, 043510 (2009).

[26] J. R. Fergusson, D. M. Regan, and E. P. S. Shellard, arXiv:1012.6039.

[27] D. M. Regan, E. P. S. Shellard, and J. R. Fergusson, Phys. Rev. D **82**, 023520 (2010).

[28] J. R. Fergusson and E. P. S. Shellard, arXiv:1105.2791.

[29] J. R. Fergusson, Phys. Rev. D **90**, 043533 (2014).

[30] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, J. Cosmol. Astropart. Phys. 12 (2012) 032.

[31] E. Hivon, K. M. Gorski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, Astrophys. J. **567**, 2 (2002).

[32] M. Tegmark, Phys. Rev. D **55**, 5895 (1997).

[33] M. Tegmark, A. N. Taylor, and A. F. Heavens, Astrophys. J. **480**, 22 (1997).

[34] M. S. Madhavacheril, K. M. Smith, B. D. Sherwin, and S. Naess, J. Cosmol. Astropart. Phys. 05 (2021) 028.

[35] H. F. Gruetjen, J. R. Fergusson, M. Liguori, and E. P. S. Shellard, Phys. Rev. D **95,** 043532 (2017).

[36] K. M. Smith and M. Zaldarriaga, Mon. Not. R. Astron. Soc. **417,** 2 (2011).

[37] M. Shiraishi, M. Liguori, and J. R. Fergusson, J. Cosmol. Astropart. Phys. 05 (2014) 008.

[38] M. G. Santos, A. Heavens, A. Balbi, J. Borrill, P. G. Ferreira, S. Hanany, A. H. Jaffe, A. T. Lee, B. Rabii, P. L. Richards, G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu, Mon. Not. R. Astron. Soc. **341,** 623 (2003).

[39] A. Gangui and J. Martin, Phys. Rev. D **62,** 103004 (2000).

[40] K. M. Smith, L. Senatore, and M. Zaldarriaga, arXiv:1502 .00635.

[41] D. Munshi, A. Heavens, A. Cooray, J. Smidt, P. Coles, and P. Serra, Mon. Not. R. Astron. Soc. **412,** 1993 (2011).

[42] K. Marzouk, A. Lewis, and J. Carron, J. Cosmol. Astropart. Phys. 08 (2022) 015.

[43] J. Smidt, A. Amblard, C. T. Byrnes, A. Cooray, A. Heavens, and D. Munshi, Phys. Rev. D **81,** 123007 (2010).

[44] M. Kamionkowski, T. L. Smith, and A. Heavens, Phys. Rev. D **83,** 023007 (2011).

[45] S. Mizuno and K. Koyama, J. Cosmol. Astropart. Phys. 10 (2010) 002.

[46] T. L. Smith and M. Kamionkowski, Phys. Rev. D **86,** 063009 (2012).

[47] A. Troja, S. Donzelli, D. Maino, and D. Marinucci, IAU Symp. **306,** 48 (2014).

[48] J. Chen et al., Astrophys. J. Suppl. Ser. **257,** 27 (2021).

[49] J. R. Bond, A. H. Jaffe, and L. Knox, Phys. Rev. D **57,** 2117 (1998).

[50] S. P. Oh, D. N. Spergel, and G. Hinshaw, Astrophys. J. **510,** 551 (1999).

[51] A. J. S. Hamilton and M. Tegmark, Mon. Not. R. Astron. Soc. **312,** 285 (2000).

[52] A. J. S. Hamilton, Lect. Notes Phys. **665,** 415 (2008).

[53] A. J. S. Hamilton, Lect. Notes Phys. **665,** 433 (2008).

[54] O. H. E. Philcox, Phys. Rev. D **104,** 123529 (2021).

[55] O. H. E. Philcox, Phys. Rev. D **103,** 103504 (2021).

[56] M. Tegmark, A. J. S. Hamilton, M. A. Strauss, M. S. Vogeley, and A. S. Szalay, Astrophys. J. **499,** 555 (1998).

[57] O. H. E. Philcox and Z. Slepian, Proc. Natl. Acad. Sci. U.S. A. **119,** e2111366119 (2022).

[58] F. Beutler and P. McDonald, J. Cosmol. Astropart. Phys. 11 (2021) 031.

[59] O. H. E. Philcox, arXiv:2303.12106.

[60] O. H. E. Philcox, Phys. Rev. D **106,** 063501 (2022).

[61] J. Hou, Z. Slepian, and R. N. Cahn, arXiv:2206.03625.

[62] G. Cabass, M. M. Ivanov, and O. H. E. Philcox, Phys. Rev. D **107,** 023523 (2023).

[63] R. N. Cahn, Z. Slepian, and J. Hou, Phys. Rev. Lett. **130,** 201002 (2023).

[64] M. Lembo, G. Fabbian, J. Carron, and A. Lewis, Phys. Rev. D **106,** 023525 (2022).

[65] E. Sellentin, A. H. Jaffe, and A. F. Heavens, arXiv:1709 .03452.

[66] J. Borrill, Phys. Rev. D **59,** 027302 (1998).

[67] M. Millea et al., Astrophys. J. **922,** 259 (2021).

[68] NIST Digital Library of Mathematical Functions (DLMF), edited by f. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, and B. V. Saunders, https://dlmf.nist.gov/.

[69] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, Astrophys. J. **622,** 759 (2005).

[70] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo et al. (Planck Collaboration), Astron. Astrophys. **641,** A6 (2020).