

Why and when to expect Gaussian error distributions in epoch of reionization 21-cm power spectrum measurements

Michael J. Wilensky¹,¹★ Jordan Brown² and Bryna J. Hazelton^{3,4}

¹*Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester M13 9PL, UK*

²*Group in Logic and the Methodology of Science, University of California Berkeley, Berkeley, CA 94720, USA*

³*Department of Physics, University of Washington, Seattle, WA 98195, USA*

⁴*eScience Institute, University of Washington, Seattle, WA 98195, USA*

Accepted 2023 March 11. Received 2023 January 25; in original form 2022 November 24

ABSTRACT

We explore error distributions in Epoch of Reionization 21-cm power spectrum estimators using a combination of mathematical analysis and numerical simulations. We provide closed-form solutions for the error distributions of individual bins in 3d-power spectra for two estimators currently in use in the field, which we designate as ‘straight-square’ and ‘cross-multiply’ estimators. We then demonstrate when the corresponding spherically binned power spectra should (and should not) have Gaussian error distributions, which requires appealing to non-standard statements of the central limit theorem. This has important implications for how upper limits are reported, as well as how cosmological inferences are performed based on power spectrum measurements. Specifically, assuming a Gaussian error distribution can over- or underestimate the upper limit depending on the type of estimator and produces overly compact likelihood functions for the power spectrum.

Key words: methods: data analysis – analytical – numerical – statistical – cosmology: observations – dark ages – reionization – first stars.

1 INTRODUCTION

The Epoch of Reionization (EoR), a period when the Hydrogen content of the universe went from being predominantly neutral to being predominantly ionized, can in principle be mapped directly using the 21-cm line from neutral Hydrogen (Furlanetto, Oh & Briggs 2006; Morales & Wyithe 2010). Since the 21-cm line arises from a forbidden transition in neutral Hydrogen, it acts as a direct tracer of the progress of cosmic reionization. Until imaging sensitivity can be achieved, observational experiments focus on constraining the power spectrum of this cosmic reionization signal i.e. the Fourier dual of the two-point correlation function of the map. This is done using radio interferometers, a few of which have produced steadily deepening upper limits, such as the Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010), the Giant Metrewave Radio Telescope (GMRT; Paciga et al. 2013), the Murchison Widefield Array (MWA; Tingay et al. 2013; Wayth et al. 2018), the LOw Frequency ARray (LOFAR van Haarlem et al. 2013), and the Hydrogen Epoch of Reionization Array (HERA; DeBoer et al. 2017). The layout of the interferometric baselines determines the sensitivity of the array to transverse spatial (angular) modes, while its frequency structure determines line-of-sight spatial characteristics when searching for a redshifted line such as the cosmic reionization signal. Current power spectrum estimators generally fall into two camps: delay spectra, which Fourier transform interferometric visibilities for each baseline along

the frequency axis, and imaging power spectra, which grid the visibilities before Fourier transforming (Parsons et al. 2012; Liu, Parsons & Trott 2014; Morales et al. 2019). Either process generates a power spectrum in three-dimensional \mathbf{k} -space. The quantity most readily compared to theoretical predictions is the spherically averaged power spectrum (i.e. averaging the power spectrum in spherical shells in the aforementioned \mathbf{k} -space). Since the target signal is extremely faint relative to other sources of radio emission, astrophysical or otherwise, current experiments generally focus on systematic mitigation and subsequent placement of an upper limit on the strength of the power spectrum signal on various spatial scales.

As upper limits on the cosmological 21-cm power spectrum signal from the Epoch of Reionization push closer to detection, it will be increasingly important to understand the exact statistical properties of the measurement. Though estimators vary, the power spectrum is inherently a quadratic statistic, and this can make tracking the exact sampling distributions of the estimators complicated. As such, there exist a number of noise characterization techniques meant to provide an estimate of the size of thermal fluctuations in the measurement, often without explicit reference to these sampling distributions.

For a thorough exploration of noise estimation techniques in delay spectrum analysis, see Tan et al. (2021). Such noise estimation techniques apply directly to power spectrum estimates formed by PAPER and HERA. These techniques include bootstrapping, visibility differencing, analytic noise power spectrum estimation (Liu & Shaw 2020), and covariance matrix propagation. Additionally, the LOFAR EoR power spectrum pipeline has deployed a combination of propagation of

* E-mail: michael.wilensky@manchester.ac.uk

Stokes V information and Gaussian Process Regression (GPR; Mertens, Ghosh & Koopmans 2018; Kern & Liu 2021) to inform noise estimates (Patil et al. 2017; Mertens et al. 2020). Other end-to-end noise propagation techniques are deployed in MWA power spectrum estimates that have overlapped with previously mentioned techniques here (Trott et al. 2016; Barry et al. 2019a).

While these noise estimates are all physically grounded and generally indicate the size of statistical uncertainties, a main point of Tan et al. (2021) is that the exact statistical question answered by each estimator varies between them. For instance, it has become a standard to reference the ‘ 2σ ’ upper limit on the reionization signal, but if the sampling distribution is non-Gaussian, then a 2σ upper limit made with reference to some estimated covariance will be different than one made by, say, a 97.7 per cent one-tailed confidence interval estimated by Monte Carlo simulation. An exact value of the upper limit calculated from the sampling distribution corresponding to the power spectrum estimator would be the clearest item to interpret. In some examples, it is explicitly assumed that the sampling distribution is Gaussian, e.g. Li et al. (2019). In Tan et al. (2021), the sampling distribution is checked somewhat exhaustively for approximate Gaussianity, and an exact sampling distribution is supplied for their estimator. This is explained by reference to the well-known central limit theorem (CLT). We remark that there are several variations of the CLT, not well-known to the community at large, and it is in fact the more nuanced variations of the CLT that are required to justify Gaussian error distributions in EoR 21-cm power spectrum measurements. The purpose of this work is to present and explore these more advanced variations in order to delineate when and why the assumption of Gaussianity is mathematically justified.

The paper is laid out as follows. In Section 2, we formally present the variations of the CLT that are relevant to our formal arguments, as well as the relevant mathematical background to understand them. In Section 3, we refresh the reader on complex circular Gaussian random vectors. In Section 4–5, we discuss the sampling distributions of the thermal noise¹ in 3-d power spectra and explore the effects of incoherent averaging (i.e. averaging of power spectra, rather than visibilities), assuming all samples are independent. We broadly classify these into ‘straight-square’ estimators, e.g. Patil et al. (2017) and Mertens et al. (2020), and ‘cross-multiply’ estimators, such as in HERA Collaboration (2022, 2023) and Barry et al. (2019b). We analytically calculate sampling distributions for both and then explore how each type of distribution can (but might not) converge to a Gaussian via different averaging schemes. In the latter case, we reproduce the averaging schemes used on actual data. In Section 6, we explore the effects of combining statistically dependent samples. We summarize our results and draw our conclusions in Section 7.

¹Strictly speaking, the noise term in the interferometric visibilities receives contribution both from thermal fluctuations in the instrument, statistical fluctuations in the emission of non-thermal sources, and other effects. For high-redshift 21-cm observations, the noise fluctuations are dominated by non-thermal galactic synchrotron radiation. However, it is a historical convention to refer to the collective noise fluctuations as ‘thermal noise’ despite non-thermal contributions. The connection is that, at each frequency, the brightness of a non-thermal source is identified with a blackbody of corresponding temperature so that its contribution to the system temperature (called its brightness temperature) predicts the appropriate noise fluctuations as per the radiometer equation. See e.g. Bennett et al. (2003) and de Oliveira-Costa et al. (2008).

2 CENTRAL LIMIT THEOREM

In this section, we review the classical version of the CLT and give generalizations applicable to the error distributions in 21-cm measurement, as the hypotheses of the classical theorem are not satisfied by the sums of error terms arising in current 21-cm measurement processing pipelines. CLTs are about convergence of an infinite sequence and therefore do not apply directly to any finite sum of random variables. Nonetheless, they provide a theoretical justification for modelling such sums as having a Gaussian distribution in many circumstances. We give a precise statement about finite sums from Shevtsova (2010) after a short exposition regarding the infinite limit.

The classical CLT states that the finite sums of an infinite sequence of independent and identically distributed random variables converge in distribution to a Gaussian. The conclusion is not necessarily true if either the assumption of independence or the assumption of identical distribution is dropped. A sequence of random variables, (X_1, X_2, X_3, \dots) , is identically distributed if, for any X_i and X_j in the sequence, the cumulative distribution function (CDF) of X_i is equal to the cumulative distribution function of X_j , i.e. $F_{X_i} = F_{X_j}$. A sequence is independent if, for any finite set I of natural numbers and finite sequence $(a_i)_{i \in I}$, we have $\mathbb{P}(i \in I \rightarrow X_i > a_i) = \prod_{i \in I} \mathbb{P}(X_i > a_i)$, where \mathbb{P} denotes probability of an event. If the variables have probability density functions, this is equivalent to the pdf of the joint distribution being given pointwise by the product of the pdfs of the individual variables. Note that uncorrelated variables may still be dependent; this is the case for X and $|X|$ whenever X has a pdf that is an even function. Additionally, sequences of pairwise independent random variables may not be independent sequences. The hypotheses for the classical central limit theorem are, taken together, quite stringent. We give statements from the exposition in Billingsley (2012).

THEOREM 1 (Billingsley (2012) Theorem 27.1). *Suppose that $\{X_n\}$ is an independent sequence of random variables having the same distribution with mean c and finite positive variance σ^2 . If $S_n = X_1 + \dots + X_n$, then $\frac{S_n - nc}{\sigma\sqrt{n}} \rightarrow N(0, 1)$.*

Here, $N(0, 1)$ indicates a standard normal random variable. In an extreme case of dependency where all variables X_i in the sequence are equal, $\frac{S_n - nc}{\sigma\sqrt{n}}$ is a (zero-mean) translation of X_i and does not converge to a Gaussian if X_i is not itself Gaussian. Most pipelines combine non-independent samples in some form or another, e.g. when computing spherical power spectra from 3d-power spectra, or when averaging cross power spectra over baseline pairs that share a baseline (Tan et al. 2021). This leads to dependency between noise measurements, so this hypothesis is not strictly satisfied.

In addition, the random noise variables are not identically distributed. We see in Sections 4–5 that the relevant variables are Exponential or Laplacian, as they arise from a squaring or product of circular complex Gaussians, and that their scale parameters vary. If successively smaller scale parameters appear in the sum, the sum may never look Gaussian, with approximately exponential tails. However, in many cases, one does obtain Gaussian distributions as in Sections 4–6. The case of independent, but non-identically distributed random variables is characterized by the Lyapunov CLT:

THEOREM 2 (Billingsley (2012) Theorem 27.3). *Suppose that $\{X_n\}$ is an independent sequence of random variables, each with mean zero and finite variance σ_n^2 , and there is some $\delta > 0$, such that, with*

$s_n^2 = \sum_{i=1}^n \sigma_i^2$, $\lim_{n \rightarrow \infty} \sum_{i=1}^n s_n^{-(2+\delta)} \mathbb{E}[|X_i|^{2+\delta}] = 0$. Then $S_n/s_n \rightarrow \mathcal{N}(0, 1)$.

There are also effective forms of the Lyapunov CLT, i.e. forms of the theorem that bound the rate of convergence. We state the sharpest bound we know of on the difference between the CDFs and a Gaussian CDF.²

THEOREM 3 (Shevtsova (2010) Theorem 1). *Suppose that $\{X_n\}$ is an independent sequence of random variables, each with mean zero, $E[X_n^2] = \sigma_n^2$, and $E[|X_n|^3] = \beta_n$, with σ_n, β_n finite and $\sum_{i=1}^n \sigma_i^2 = 1$. Writing F_n for the CDF of $\sum_{i=1}^n X_i$ and Φ for the CDF of $\mathcal{N}(0, 1)$, we have*

$$\sup_{x \in \mathbf{R}} |F_n(x) - \Phi(x)| \leq 0.56 \sum_{i=0}^n \beta_i.$$

The quantity on the left-hand side measures the distance of the CDF of the n th partial sum from a Gaussian CDF. The right-hand side is typically easy to compute. If the left-hand side goes to 0 in the infinite limit, then the sequence converges in distribution to a standard normal random variable. Therefore, the right-hand side tells us how close a partial sum is to convergence at any given n .

Unfortunately, the random sequences at issue for the 21-cm error bars are not always independent and the Lyapunov theorem does not strictly apply. However, some estimators may be constructed to satisfy a weak form of independence, introduced in Hoeffding & Robbins (1948), known as m -dependence for a fixed natural number m . Essentially, if the statistical dependence only has a fixed finite range among the samples and the Lyapunov condition is satisfied, then there is convergence to Gaussianity. A sequence $(X_i)_{i=0}^{\infty}$ of random variables is said to be m -dependent if, whenever $s - r > m$, (X_0, X_1, \dots, X_r) and $(X_s, X_{s+1}, \dots, X_n)$ are independent vector-valued random variables. That is, if we write $\mathbf{Y}_0 = (X_0, X_1, \dots, X_r)$ and $\mathbf{Y}_1 = (X_s, X_{s+1}, \dots, X_n)$, then, for any rectangles $R_0 = (a_0, b_0) \times \dots \times (a_r, b_r)$ and $R_1 = (a_s, b_s) \times \dots \times (a_n, b_n)$, we have $\mathbb{P}(\mathbf{Y}_0 \in R_0 \& \mathbf{Y}_1 \in R_1) = \mathbb{P}(\mathbf{Y}_0 \in R_0) \mathbb{P}(\mathbf{Y}_1 \in R_1)$.

If (X_i) is an m -dependent sequence of random variables, we define, for each natural number i , $A_i = \mathbb{E}[X_{i+m}^2] + 2 \sum_{j=1}^m \mathbb{E}[X_{i+m-j} X_{i+m}]$.³

THEOREM 4 (Hoeffding & Robbins (1948) Theorem 1). *If (X_i) is an m -dependent sequence of random variables, each with zero mean, and there is a real number M such that, for all i , $\mathbb{E}[|X_i|^3] \leq M$, and $\lim_{p \rightarrow \infty} p^{-1} \sum_{h=1}^p A_{i+h}$ exists uniformly for all i , with limit A , then the distribution function of $n^{-1/2} A^{-1} \sum_1^n X_i$ converges pointwise to the distribution function of $\mathcal{N}(0, 1)$.*

Hoeffding and Robbins's theorem shows that local dependency does not interfere with the CLT, regardless of whether the random variables are identically distributed. It therefore theoretically justifies the assumption of approximate Gaussianity for the finite sums of random variables that arise in error estimates.

²This theorem makes use of the *supremum*, which for a bounded set of real numbers is the lowest upper bound of the set. For example, the supremum of the real open interval $(-1, 1)$ is 1.

³Note that, if the sequence is m -dependent and $m \leq n$, then the sequence is n -dependent, and we may obtain different A_i depending on whether we consider the sequence as m -dependent or n -dependent; having fixed an m for which we consider (X_i) to be m -dependent at the beginning, the definition of the A_i 's is unambiguous.

3 COMPLEX CIRCULAR GAUSSIAN RANDOM VECTORS

In this section, we quickly review some concepts of circular Gaussian random vectors and explain their relevance to the problem. We draw mainly from Gallager (2013), where detailed proofs of the following claims are held.

The visibility noise of a given baseline at each frequency and time is complex circular Gaussian (i.e. the real and imaginary components are independent and identically distributed Gaussian random variables; Thompson, Moran & Swenson 2017), and is independent of the noise of every other baseline, frequency, and time. This means the noise samples of all the visibilities may be considered together as jointly circular Gaussian. We explore particular estimators in more detail in the following sections, but all of them involve applying linear transformations to the visibilities to produce processed data whose noise-like component is a circular complex Gaussian random vector. These are then squared or cross-multiplied in a particular manner to produce the power spectrum estimator in question. Thus, if we can produce the probability density function of the square or cross-multiplication of two independent circular complex Gaussian random variables and determine how such product distributions that are generally non-Gaussian behave under averaging, then we can see under what conditions a Gaussian error distribution is valid to assume.

Suppose the complex variables $(X_1, X_2, X_3, \dots, X_N)$ are jointly Gaussian. Then we refer to the vector \mathbf{X} whose j th component is X_j as a Gaussian random vector. If

$$\begin{aligned} E[\mathbf{X}] &= \boldsymbol{\mu} \\ E[(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^\dagger] &= \mathbf{C} \\ E[(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T] &= \boldsymbol{\Gamma}, \end{aligned} \quad (1)$$

where \mathbf{X}^T and \mathbf{X}^\dagger are, respectively, the transpose and conjugate transpose of \mathbf{X} , then \mathbf{X} is a complex Gaussian random vector variable with mean, $\boldsymbol{\mu}$, covariance matrix, \mathbf{C} , and pseudo-covariance matrix, $\boldsymbol{\Gamma}$. We write

$$\mathbf{X} \sim \mathcal{CN}(\boldsymbol{\mu}, \mathbf{C}, \boldsymbol{\Gamma}). \quad (2)$$

We say that \mathbf{X} has circular symmetry if $\mathbf{X}e^{i\phi}$ is distributed identically to \mathbf{X} . That is, we can multiply the vector by an overall complex phase and the distribution is unaffected. Inspecting equation (1), we see that it is necessary that

$$\begin{aligned} \boldsymbol{\mu} &= 0 \\ \boldsymbol{\Gamma} &= 0. \end{aligned} \quad (3)$$

For a Gaussian random vector, this is also sufficient for circular symmetry, since the three expectations in equation (1) define all the statistical properties of the distribution. From now on, we say that if $\mathbf{X} \sim \mathcal{CN}(0, \mathbf{C}, 0) \equiv \mathcal{CN}(0, \mathbf{C})$, then \mathbf{X} is circular complex Gaussian. In this case, its probability density function, $f_{\mathbf{X}}(\mathbf{x}|\mathbf{C})$, is written as

follows:⁴

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{\exp[-\mathbf{x}^\dagger \mathbf{C}^{-1} \mathbf{x}]}{\pi^n \det \mathbf{C}}, \quad (4)$$

where n is the number of components of \mathbf{X} .

The most important property for circular complex Gaussian vectors for our purposes is that circular symmetry is maintained under linear transformations. That is, if

$$\mathbf{Y} = \mathbf{A}\mathbf{X}, \quad (5)$$

where \mathbf{A} is some constant matrix such that $\mathbf{A}\mathbf{C}\mathbf{A}^\dagger$ is invertible, then

$$\mathbf{Y} \sim \mathcal{CN}(0, \mathbf{A}\mathbf{C}\mathbf{A}^\dagger). \quad (6)$$

Note that if $\mathbf{A}\mathbf{C}\mathbf{A}^\dagger$ is singular, then a dimensional reduction procedure is required to write the density in the form above, cf. Appendix A in Byrne et al. (2021). Additionally, all linear functionals of circular complex Gaussian vectors are circular complex Gaussian variables, which is to say that for any (non-zero) vector \mathbf{b} , the scalar $\mathbf{b}^T \mathbf{X}$ is a circular complex random variable.

4 ‘STRAIGHT-SQUARE’ POWER SPECTRUM ESTIMATORS

We begin by describing power spectrum estimators that involve squaring the linearly transformed visibilities, such as that used in LOFAR upper limits. Denoting the linearly transformed visibilities as $\tilde{v}(\mathbf{k}, t)$, we may suppose that each voxel has a power of the form

$$\begin{aligned} P(\mathbf{k}, t) &= |\tilde{v}(\mathbf{k}, t)|^2 = |s(\mathbf{k}, t) + n(\mathbf{k}, t)|^2 \\ &= |s(\mathbf{k}, t)|^2 + 2\Re(s^*(\mathbf{k}, t)n(\mathbf{k}, t)) + |n(\mathbf{k}, t)|^2, \end{aligned} \quad (7)$$

where $s(\mathbf{k}, t)$ and $n(\mathbf{k}, t)$ are the signal and noise components of the linearly transformed visibilities, respectively, for the wave mode \mathbf{k} , at sidereal time t . Another data product for which this is relevant are the ‘difference’ cubes formed by Error Propagated Power Spectrum with Interleaved Observed Noise (ε PPSILON), which uses time-differenced data to form power-spectrum-esque objects (Barry et al. 2019a). These cubes are used both as a part of the ε PPSILON power spectrum estimator as well as by themselves as a noise metric, since they are sourced only by those entities that remain after time-differencing at the instrument cadence (typically thermal noise fluctuations and sometimes RFI sources, e.g. Wilensky et al. (2019)). In most cases, it is fair to only consider the quadratic noise term in these difference cubes. ε PPSILON uses a generalization to the Lomb–Scargle periodogram (VanderPlas 2018) for power spectrum estimation, but its implementation can indeed be seen as first performing a linear transformation to the data, and then taking a sum of quadratic terms.

There are three terms in equation (7), two of which contain contributions from the signal, one of which is a signal-noise cross term. It raises the question as to whether we should include this cross term into the description of the error statistics, since it clearly receives contribution from the noise. We note that this

⁴This form of the pdf may look unsightly to those who are only familiar with real-valued normal vectors, for which there is a factor of 2 in the exponent (as well as one associated with each factor of π) and a square root of the determinant. If one writes this circular Gaussian pdf in terms of the real and imaginary components of \mathbf{X} , so that it is instead describing a real-valued normal random vector of twice the length, then one finds these factors and operations in the appropriate locations (after an appropriate renaming of the covariance matrix as well). See Gallager (2013).

term is significant whenever the strength of the signal and noise terms are comparable, otherwise one of the two quadratic terms dominates. Without assuming some explicit model of the signal and the averaging scheme, it appears difficult to treat these terms to any satisfactory level of generality.

All 21-cm power spectrum upper limits involve incoherent averaging (averaging after squaring or cross-multiplying, rather than *coherent* averaging which takes place before), since at the very least samples at different \mathbf{k} -modes on the same spherical shell must be averaged together. Some power spectrum estimators also involve incoherently averaging different sidereal times. The incoherent averaging scheme is tied intimately to the treatment of the signal-noise cross term. For example, if one fixes a set of \mathbf{k} -bins, only ever observes one field on the sky (whose size is determined by the instrumental beam), the astrophysical foreground sources in that field are perfectly subtracted (or equivalently, no \mathbf{k} bins affected by foregrounds are included), and no other systematics are present, then there is only one signal realization that ever enters the incoherent averaging scheme. Since we assume linearity in going from visibilities to $\tilde{v}(\mathbf{k}, t)$, we necessarily assume $n(\mathbf{k}, t)$ is a circular Gaussian random vector. An incoherent average (in this case, the spherical average, since we have essentially assumed only one sidereal time) of the cross terms then amounts to the real component of a weighted average of a subset of the complex components of this circular Gaussian random vector, and in this case the error distribution is potentially tractable since this term will have Gaussian distribution by the properties exposed in Section 3. However, one would still need to understand the signal variation between different \mathbf{k} -modes to obtain the appropriate sampling distribution. If more than one field is combined incoherently, then there are multiple realizations to contend with in the incoherent average, and then one must understand the nature of cosmic variance. Without a reliable model for variations in the signal from field to field or mode to mode, one could not propose a realistic accounting for the cross term.

The presence of systematic effects tends to exacerbate the difficulty of the cross term. While astrophysical sources may behave similarly to the cosmological signal depending on the stationarity of the sources in the fields of interest, RFI is generally non-stationary and has poorly understood statistics at faint intensities. Excess power from flagging (Offringa, Mertens & Koopmans 2019; Ewall-Wice et al. 2021; Wilensky, Hazelton & Morales 2022) may also vary field-to-field and night-to-night, producing non-trivial statistics if not ameliorated, etc. Since a general treatment of systematics is beyond the scope of this work and we must assume something particular about the cosmological signal in order to characterize its effect in the cross term, we will ignore it in this analysis and focus ourselves on the noise-only term. This means this analysis is mainly applicable in the regime where the noise dominates the signal. Note that in the simplest case (only one field, no systematics other than potentially foregrounds), the Gaussianity of the distribution is unaffected by the addition of the cross term, though its mean and variance will be affected, which has important implications regarding reported confidence intervals, etc.

In general, there may be statistical dependency between different \mathbf{k} -modes or redundant measurements of the same \mathbf{k} -mode, and so the sampling distribution of the cylindrical or spherical power spectrum noise contamination is not necessarily straightforward to compute. In general, if \mathbf{X} is a circular Gaussian random vector (with arbitrary covariance), then the variable

$$\mathbf{Y}' = \mathbf{X}^\dagger \mathbf{B}\mathbf{X}, \quad (8)$$

for some complex matrix \mathbf{B} , has real and imaginary components that are a special case of ‘Generalized χ^2 ’ random variables. The pdf can be calculated in closed form by calculating the characteristic function and then performing an inverse Fourier transform (Imhof 1961), which generally produces unwieldy expressions, or through more recently discovered numerical methods (Das & Geisler 2021). We present probability density functions for special cases relevant to 21-cm power spectrum estimation in this paper. A well-known result is that the modulus square of a circular Gaussian random variable is an exponential random variable i.e. one whose pdf is of exponential form:

$$f_X(x) = \lambda e^{-\lambda x}. \quad (9)$$

We refer to λ as the ‘scale parameter’ of the exponential random variable. This relation can be gleaned by considering the complex Gaussian variable as a 2-d real Gaussian and inspecting the density in polar coordinates. Since it is instructive, and correlations are often short-ranging, we derive the sampling distribution for a sum of independent exponential random variables, and also show that in some instances, even sums of independent exponential random variables do not approach Gaussianity in the infinite limit.

If $\mathbf{B} = \mathbf{I}$ and each component of \mathbf{X} is independent of every other component, but the variance of each component is not necessarily identical, then the variable Y ,

$$Y \equiv \mathbf{X}^\dagger \mathbf{X} \quad (10)$$

is called a ‘hypoexponential’ random variable. A special case of the hypoexponential random variable is the Erlang random variable, which occurs when all the variances of \mathbf{X} are equal. The scale parameters for each of the terms in the sum, denoted λ_j , are equal to $1/\sigma_j^2$, where σ_j^2 is the total variance j th vector component of \mathbf{X} (sum of variance of real and imaginary components). Due to the convolution theorem, the characteristic function of Y is the product of the characteristic functions of each exponential term:

$$E[e^{itY}] = \prod_{j=1}^N \frac{\lambda_j}{\lambda_j - it}, \quad (11)$$

where N is the dimension of \mathbf{X} . Note that the variances do not all have to be distinct in this treatment. To compute its probability density function, we can perform an inverse Fourier transform:

$$f_Y(y) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} e^{-ity} \prod_{j=1}^N \frac{\lambda_j}{\lambda_j - it} dt. \quad (12)$$

We can compute this integral using contour integration methods. The integrand is a meromorphic function i.e. it is complex differentiable everywhere except for a countable number of poles. It has one pole per distinct variance, and the order of each pole is given by the multiplicity of that variance. All of the poles are located in the lower half of the complex plane. When $y < 0$, we can use a semi-circular contour that encloses the upper-half plane and Jordan’s lemma (Byron Jr & Fuller 1970), showing that the result is 0. This is just an expression of positive-definiteness. When $y \geq 0$, we enclose the lower-half plane (and thus the residues) with a semi-circular contour. The residues are readily calculated using Cauchy’s integral formula (Byron Jr & Fuller 1970). The final result in the fully general case is

$$f_Y(y) = \sum_{j=1}^n \sum_{m=0}^{m_j-1} \frac{\lambda_j^{m_j}}{(m_j - m - 1)! m!} y^{m_j - m - 1} e^{-\lambda_j y} \Psi_{m_j}(-\lambda_j), \quad (13)$$

where m_j is the multiplicity of the j th variance and Ψ_{m_j} is the function

$$\Psi_{m_j}(t) = \frac{d^{m_j}}{dt^{m_j}} \left[\prod_{l \neq j} \left(\frac{\lambda_l}{\lambda_l + t} \right)^{m_l} \right]. \quad (14)$$

This is the probability density function for the positive-definite noise bias in a straight-square estimator when all the \mathbf{k} modes are independent, up to scaling factors from the weights (and the fact that this represents a sum, not an average), which just rescale the λ_j ’s.

The hypoexponential random variable is also a special case of random variable that arises in the study of finite-state continuous-time Markov chains, called a phase-type random variable (Breuer & Baum 2005). A phase-type random variable describes the sampling distribution for a finite-state Markov process to reach an absorbing state. These processes are frequently studied in queuing theory, telecommunications theory, and biology (Nelson, Randolph 1995; Breuer & Baum 2005). The hypoexponential distribution is the absorption time sampling distribution for a Markov process where all states are attainable and there are no cycles. Its distribution functions can be efficiently computed using the following theorem.

THEOREM 5 (Breuer & Baum (2005) Theorem 9.2, Example 9.4). *The cumulative distribution function of the sum of independent, exponentially distributed random variables with scale parameters $\lambda_1, \dots, \lambda_n$ is equal to $1 - \boldsymbol{\alpha}^T e^{\boldsymbol{\theta} x} \mathbf{1}$, and the pdf of the sum is*

$$-\boldsymbol{\alpha}^T e^{\boldsymbol{\theta} x} \boldsymbol{\theta} \mathbf{1}, \text{ where } \boldsymbol{\alpha} \in \mathbf{R}^n, \boldsymbol{\alpha} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \dots \\ 0 \end{bmatrix}, \mathbf{I} \in \mathbf{R}^n \text{ is } \mathbf{1} = \begin{bmatrix} 1 \\ 1 \\ \dots \\ 1 \end{bmatrix},$$

$$\text{and } \boldsymbol{\theta} = \begin{bmatrix} -\lambda_1 & \lambda_1 & 0 & \dots & \dots & 0 \\ 0 & -\lambda_2 & \lambda_2 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \dots & -\lambda_{n-1} & \lambda_{n-1} \\ 0 & 0 & \dots & \dots & 0 & -\lambda_n \end{bmatrix} \text{ is an } n \times n \text{ matrix. Here}$$

$$e^{\mathbf{A}} = \sum_{i=0}^{\infty} \frac{\mathbf{A}^i}{i!} \text{ is the matrix exponential.}^5$$

In the next subsection, we use this form of hypoexponential random variable to efficiently compute the distribution functions of some pedagogical hypoexponential examples with a high degree of numerical stability. In particular, we expose some distinct cases in which the hypoexponential random variable is and is not approximately Gaussian.

4.1 Numerical Exploration

Of particular note for hypoexponential random variables is that there are some sets of λ_j ’s for which the variable is not even approximately Gaussian in the infinite limit. We can understand this by examining the skew of the hypoexponential random variable. Since the second and third central moments are cumulants (they add for independent random variables), we can formulate the skew of the hypoexponential as

$$E \left[\left(\frac{Y - \mu}{\sigma} \right)^3 \right] = \frac{\sum_{j=1}^N \frac{2}{\lambda_j^3}}{\left(\sum_{j=1}^N \frac{1}{\lambda_j^2} \right)^{3/2}}. \quad (15)$$

⁵The quantity $\boldsymbol{\alpha}^T e^{\boldsymbol{\theta} x} \mathbf{1}$ is independent of the ordering of the scale parameters $\lambda_1, \dots, \lambda_n$, though $e^{\boldsymbol{\theta} x}$ is not. This is most clearly understood through the connection between hypoexponential distributions and absorption-time distributions for Markov chains explored in detail in Breuer & Baum (2005).




Figure 1. PDF plot for sums of 200 independent exponential random variables (one with linear and one with semilog axes), with scale parameters λ_n for each sequence as indicated by the legend, compared to a Gaussian distribution. All variables are translated to have mean 0 and scaled to have variance 1. When $\lambda_n = n$, the sum retains its skew as computed in equation (16).

Since a Gaussian random variable has 0 skew, and we can construct sums in the form above that provide a non-zero skew in the infinite limit, we can say that some sums of exponential random variables do not converge to a Gaussian even in the infinite limit. To begin listing a few examples, consider when all the λ_j 's are identical. We then obtain a skew of $2/\sqrt{N}$, which vanishes in the infinite limit. If the λ_j 's are the reciprocals of the positive integers, then we also see that the ratio of sums in equation (15) vanishes in the infinite limit. In contrast, if the λ_j 's are the positive integers, then the ratio of sums converges to a non-zero value. In particular, it converges to the value

$$\lim_{N \rightarrow \infty} \frac{\sum_{j=1}^N \frac{2}{\lambda_j^3}}{\left(\sum_{j=1}^N \frac{1}{\lambda_j^2}\right)^{3/2}} = \frac{2\zeta(3)}{\zeta(2)^{3/2}}, \quad (16)$$

where $\zeta(x)$ is the Riemann zeta function. We show a figure demonstrating these sums in Fig. 1, where we plot the probability density functions of these choices of hypoexponential random variables.

Note that this is not identical to the Lyapunov condition from Section 2, so we cannot necessarily say that when the skew vanishes in the limit that the sum of exponentials converges to Gaussianity. However, it is clear from Fig. 1 that some finite sums of exponential random variables are approximately Gaussian.

Many error estimates fundamentally work with the estimated standard deviation of the power spectrum bin (typically reporting a 2σ upper limit based on the estimate of σ). For a Gaussian random variable, particular distances from the origin, measured in standard deviations, correspond uniquely to well-known values of the CDF of a Gaussian random variable. When the error distribution is non-

Gaussian, reporting a significance in terms of a standard deviation no longer corresponds to the same values of the relevant CDF. For any hypoexponential random variable X , there is an exponential random variable Y such that, for large x , $f_Y(x) < f_X(x)$, meaning that a given significance reported in standard deviations will tend to underestimate the size of the corresponding confidence interval if the significance is interpreted according to the correspondence between confidence intervals and standard deviations for a normal distribution. For nearly Gaussian hypoexponential random variables, the significance at which the difference in tails matters is high, i.e. one does not observe the non-Gaussianity except at distances far from the origin. Since many analyses report a 2σ upper limit, we characterize the degree to which this underestimates the actual inferred confidence interval as a figure of merit. If $\Phi(x)$ is the standard normal cumulative distribution function,

$$\Phi(x) = \frac{1}{2} \left(1 + \operatorname{erf}(x/\sqrt{2}) \right), \quad (17)$$

and $\operatorname{erf}(x)$ is the error function,

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt. \quad (18)$$

then we calculate

$$Q \equiv \frac{F^{-1}(\Phi(2)) - \mu}{2\sigma}, \quad (19)$$

where μ is the mean of the variable in question. For a Gaussian random variable, Q is exactly 1. For a single exponential random variable,

$$Q_{\text{exp}} = \frac{-\ln(1 - \Phi(2)) - 1}{2} = 1.39159\dots \quad (20)$$

This is to say that, for a *single* exponential random variable, if one perfectly subtracts the noise bias and quotes a 2σ noise level based on an estimate of the standard deviation in order to represent a typical point at which a measurement might be designated as noise limited, then this will undershoot the actual 97.7 per cent noise level by approximately 40 per cent. Generally, we observe that for any hypoexponential random variable, Q is bounded between 1 and the value above. While this has the consequence that the power at which a measurement can be deemed noise limited is higher, it also means that the reported upper limits on the signal are inaccurate without reference to the true sampling distribution. Since any actual upper limit entails averaging multiple bins together, the amount of undershoot on the noise level is generally less than Q_{exp} for a straight-square estimator and decreases substantially with just a few averages ($N \sim 20$) assuming the noise in each sample is about the same.

The exact inaccuracy of the upper limit calculation imposed by the non-Gaussianities will depend on the reported value as well as the particulars of the contributing voxels. Specifically, the upper limit is derived from the probability distribution of the signal, conditional on the power spectrum measurement. Neglecting signal-noise cross terms, equation (7) reads (omitting function arguments for brevity)

$$P = |s|^2 + |n|^2. \quad (21)$$

Since these two terms are positive definite, $|s|^2 \leq P$. This means that the measured value always acts as a 100 per cent upper limit on the signal. If the noise distribution is treated as Gaussian, then one may inadvertently report an erroneous ' 2σ ' upper limit greater than P in some cases. This also suggests a potentially simpler method for placing upper limits when cross terms can be ignored, i.e. just quoting P as the 100 per cent upper limit. However, this can make comparisons between analyses with very different




Figure 2. The skews and Q -values of sums of independent exponentials with scale parameters given by $10^{-[5n/42]}$ for $n < 42$. Each time an exponential is added with a scale parameter an order of magnitude smaller than the previous scale parameters, the tails revert to looking nearly exponential, as evidenced by the Q -value and skew spikes.

sampling distributions difficult, such as the cross-multiply estimators in Section 5, which have an infinite 100 per cent upper limit.

We remark that Q is only a proxy for the relative Gaussianity of the distribution. One could instead regard Q as a function so that what we have defined as Q is $Q(2)$. Then one may calculate $Q(x)$ for some x greater than 2 if agreement with a Gaussian to some higher level of significance is desired. In other words, if the analyst only has reason to be concerned with a certain level of significance, x , then if $Q(x)$ is close to 1, a Gaussian approximation may suffice for that significance, but this depends on the context in which the error distribution is being used and so one should take care to understand the distribution to the extent that it affects their analysis, e.g. if one is performing inference based on the cosmological 21-cm signal.

In order to intentionally create a pathological example for which the Lyapunov CLT does not apply, and to expose a type of behaviour that can disrupt the Gaussianity of a random variable, we add a sequence of independent, exponentially distributed random variables with scale parameters that decrease exponentially to 0 in discrete steps. We add exponential random variables of a fixed scale parameter, but after some number of additions, we begin adding exponential random variables with a scale parameter a factor of 10 lower than the initial one. We repeat this with smaller and smaller scale parameters (i.e. larger and larger variances), decreasing by a factor of 10 each time. We see in Fig. 2 that, in such cases, the distribution of the sum resembles the distribution of the exponential with the smallest scale parameter included in the sum. When several variables are added with a fixed scale parameter, we observe the sum beginning to converge to a Gaussian distribution, though in this ‘comb’ example the dropping scale parameters prevent it from ever getting very close.

A unifying theme of the examples that do not converge to a Gaussian in the infinite limit is that the distribution of variances among the summed variables spans many scales (though this is not uniformly sufficient, e.g. the green curve in Fig. 1). Translating this to a practical application, analysts should be wary of combining measurements with substantially different noise levels. For example, Fig. 2 demonstrates a scenario in which a sum that is tending towards Gaussianity can be shocked back towards exponentiality whenever

a sample of standard deviation 10 times the typical value is added.⁶ A more thorough examination of the error distribution is warranted when there is a high degree of variability in the uncertainty of co-added samples. We see this theme return through different examples in the next section, which discusses cross-multiply estimators.

However, we also point out that most power spectrum estimators use ‘inverse variance weighting,’ which in a straight-square estimator will weight each summed exponential variable by an estimate of its scale parameter (recall the scale parameter is the inverse variance of the sample to be squared). If this estimate is exact, then the noise-only term will be a sum of identically distributed exponential random variables, collapsing to the simplest statement of the central limit theorem if statistical dependence can be ignored. In general, the estimate will have errors due to imperfect knowledge of the system temperature or whatever other parameters guide the estimate. If the errors are small, then one could appeal to the Lyapunov CLT for intuition. We demonstrate this in the next section. If the errors are large, and in particular if they are much larger for some data than others, then a significantly non-Gaussian result may still be expected despite inverse variance weighting.

5 ‘CROSS-MULTIPLY’ POWER SPECTRUM ESTIMATORS

Some power spectrum estimators, such as ϵ PPSILON, and that deployed in the recent HERA limits (HERA Collaboration 2022, 2023) involve cross-multiplying data that share a signal component but have independent noise components. This results in a different error distribution, notably one that is not positive-definite but rather zero-mean. In the case of HERA, the estimator takes the form (Tan et al. 2021),

$$P(\mathbf{k}, t) = |s(\mathbf{k}, t)|^2 + s(\mathbf{k}, t)n_2^*(\mathbf{k}, t) + s^*(\mathbf{k}, t)n_1(\mathbf{k}, t) + n_1(\mathbf{k}, t)n_2^*(\mathbf{k}, t) \quad (22)$$

and similarly for the estimator in ϵ PPSILON (Barry et al. 2019a), except only involving the real component of this quantity. Here, $n_1(\mathbf{k}, t)$ and $n_2(\mathbf{k}, t)$ are independent circular Gaussian random vectors. As in Section 4, we will ignore the signal-noise cross terms for the reasons expressed there. The noise-only term can be expanded as (omitting the arguments for brevity) follows:

$$n_1 n_2^* = \Re(n_1)\Re(n_2) + \Im(n_1)\Im(n_2) + i(\Im(n_1)\Re(n_2) - \Re(n_2)\Im(n_1)). \quad (23)$$

Since by circularity, the real and imaginary components of either one of the noise vectors are independent, identically distributed, and zero-mean, the components of the product are also independent, identically distributed, and zero-mean. Each component is a sum of two terms that are products of independent, one-dimensional Gaussian random variables. Denoting any one of these terms as W , and writing the standard deviations of n_1 and n_2 as σ_1 and σ_2 , respectively, we calculate its probability density function using the rule for products of independent random variables:

$$f_W(w) = \frac{1}{2\pi\sigma_1\sigma_2} \int_{-\infty}^{\infty} \frac{dx}{|x|} e^{-x^2/2\sigma_1^2} e^{-w^2/2x^2\sigma_2^2}. \quad (24)$$

We are free to rewrite this like so:

$$f_W(w) = \frac{1}{\pi\sigma_1\sigma_2} \int_0^{\infty} \frac{dx}{x} e^{-\frac{|w|}{2\sigma_1^2}(e^{\ln(x^2\sigma_2/|w|\sigma_1)} + e^{\ln(|w|\sigma_1/x^2\sigma_2)})}. \quad (25)$$

⁶To be clear, the final result is independent of the order of summation, but the orderliness of this example helps isolate the relevant phenomenon.

Using the substitution $t = \ln(x^2\sigma_2/|w|\sigma_1)$, we can write

$$f_W(w) = \frac{1}{\pi\sigma_1\sigma_2} \int_0^\infty dt e^{-|w|\cosh t/\sigma_1\sigma_2} = \frac{K_0(|w|/\sigma_1\sigma_2)}{\pi\sigma_1\sigma_2}, \quad (26)$$

where K_0 is a modified Bessel function of the second kind, and the final identity can be found using contour integration methods (Watson, G.N. 1966).

W is a sum of two independent terms with exactly this distribution. The resulting sum is distributed according to the autoconvolution of the above distribution. Rather than compute this convolution directly, we instead evaluate the square of the characteristic function, which is the characteristic function of the autoconvolution by the convolution theorem. From Abramowitz & Stegun (1964), we have

$$K_0(x) = \frac{1}{2} \int_{-\infty}^{\infty} dk \frac{\cos(kx)}{\sqrt{k^2+1}}. \quad (27)$$

Since $\sin(kx)$ is odd and $\sqrt{k^2+1}$ is even, we may write this as follows:

$$\frac{K_0(|x|)}{\pi} = \frac{1}{2\pi} \int_{-\infty}^{\infty} dk \frac{e^{-ikx}}{\sqrt{k^2+1}}, \quad (28)$$

where we have taken advantage of the fact that only the even part of e^{-ikx} contributes to the integral in order to introduce the $|x|$ on the left-hand side. Fourier inversion and the convolution theorem implies then that the characteristic function of W is

$$E[e^{ikw}] = \frac{1}{\sigma_1^2\sigma_2^2k^2+1}, \quad (29)$$

which is precisely the characteristic function of a Laplacian random variable centred on the origin with a scale parameter

$$\lambda = \frac{1}{\sigma_1\sigma_2}. \quad (30)$$

Invertibility of the Fourier transform finally implies that the noise-only terms in the aforementioned cross-multiply estimators have a Laplacian probability density function before any incoherent averaging. Incoherent averaging of independent power spectra will then have noise-only terms whose pdfs are that of sums of independent Laplacian random variables.

5.1 Numerical exploration with simulations

In this section, we explore the properties of sums of Laplacian random variables using various simulations. We show examples for different mixtures of scale parameters. Some mixtures clearly tend towards Gaussianity as they are summed, while others do not. We find that mixtures whose scale parameters are concentrated around some value seem to tend towards Gaussianity.

Recall the figure of merit used in Section 4.1, expressed by equation (19). We again use this figure of merit in an exploration of the relative Gaussianity of a finite sum of Laplacian random variables. For a single Laplacian random variable, we have

$$Q_{\text{laplace}} = -\frac{1}{2\sqrt{2}} \ln\left(2(1-\Phi(2))\right) = 1.09249\dots \quad (31)$$

For a Normal random variable, this is equal to exactly 1. In general, a sum of Laplacian random variables will have a Q between these two values. Note that Q is only a proxy for the Gaussianity of the random variable unless one can prove that the convergence in distribution is uniform. It does, however, allow us to speak to the degree to which citing a variance underestimates thermal errors. The effect on the upper limit in this case is that it is generally underestimated when the

Gaussian approximation is used. The degree of underestimation can vary depending on the value of the measurement. Since the signal-only term is clearly positive definite, the probability distribution of the signal conditional on the measurement must be truncated below 0 (Li et al. 2019, Appendix A). If the measured value is close to 0, then this can shift the mass of the signal distribution in an asymmetric way.

An analytic formula of the pdf for the case where all the scale parameters are distinct is given in Tan et al. (2021), and a more general one can be derived for when some of the scale parameters are degenerate. These formulas involve sums of products of many numbers (similar to equation (13)), and we found them to be prone to numerical instability when large numbers of variables are summed. We have therefore elected to do most of our exploration using semi-numerical techniques. Using the convolution theorem, we have (for a general continuous random variable)

$$\int_{-x_0}^{x_0} f_X(x)dx = \frac{x_0}{\pi} \int_{-\infty}^{\infty} \text{sinc}(x_0t)\hat{f}_X(t)dt, \quad (32)$$

where \hat{f}_X is the characteristic function of the random variable in question and

$$\text{sinc}(x) = \frac{\sin(x)}{x}. \quad (33)$$

Since our probability density function is always even, we may write⁷

$$\begin{aligned} F_Y(y) &= \frac{1}{2} + \int_0^y f_Y(y')dy' \\ &= \frac{1}{2} + \frac{x_0}{\pi} \int_0^\infty \text{sinc}(x_0t)\hat{f}_Y(t)dt, \end{aligned} \quad (34)$$

if Y is a sum of Laplacian random variables and $y > 0$. To establish Q for some Y of interest, we evaluate equation (34) numerically using Gaussian quadrature for a handful of y values in the interval $[1, 1.1]$ (which contains Q_{Laplace}), and then invert the CDF using a simple linear interpolation scheme. We find the type of interpolation does not change the result meaningfully. All of this is implemented in PYTHON using the NUMPY and SCIPY libraries (Harris et al. 2020; Virtanen et al. 2020).

In Fig. 3, we calculate Q for various mixtures of scale parameters. In each panel, we draw the scale parameters for a particular summed Laplacian from a power law with a minimum cut-off at $\lambda = 1$, and variable index and upper cut-off depending on the panel, i.e.

$$f_\Lambda(\lambda) = \begin{cases} \frac{1-a}{\lambda_u^{1-a}-1} \lambda^{-a}, & 1 < \lambda < \lambda_u \\ 0, & \text{otherwise.} \end{cases} \quad (35)$$

We then calculate Q for each partial sum over the mixture, making no special choice about the order of the scale parameters for each realization. For the case of inverse variance weighting, one should think of this as the distribution of the scale parameters after weighting the cross-multiplied samples, i.e. the distribution of scale parameters generated by errors in the estimates of the data variances. Each panel shows 100 different mixtures for the given power law parameters. We find that when the mixtures are less concentrated, which can be accomplished with shallower power-law index and larger upper cut-off, the distribution of Q at some given number of sums is also less concentrated. In the most diffuse case on display, even after 50 independent additions, Q is more or less uniformly distributed over the possible range.

⁷To see this, it may help to note that since f is real, \hat{f} is also even, as can be gleaned from equation (29).




Figure 3. Ratio of the numerically calculated 97.7 per cent confidence interval divided by twice the standard deviation of sums of Laplacian random variables with different scale parameters (Equation 19). In each panel, the scale parameters for each line are drawn from a power law with index a , upper cut-off λ_u (see annotations of each panel), and a lower cut-off of 1 (Equation 35). We simulate 100 independent groups of 50 scale parameters in each panel, and track the ratio of the confidence interval to twice the standard deviation as we cumulatively average the Laplacian random variables. This ratio is approximately 1.0925 for a Laplacian random variable. As the average proceeds towards Gaussianity, this ratio goes to 1. When scale parameters of very different size enter the average, the averaged variable returns closer to a Laplacian random variable than a Gaussian one. For power laws with large tails (upper panels), this effect is so pronounced that even after averaging 50 random variables, many of the distributions still appear Laplacian. On the other hand, steeper power laws are more concentrated towards the lower cut-off, and generally proceed towards Gaussianity in an orderly fashion (bottom panels). Note that we more often produce something approximately Gaussian than something significantly non-Gaussian.

In Fig. 4, we show the probability density functions of the summed Laplacians in the most extreme case of Fig. 3, where the distribution of Q after 50 additions is nearly uniform. We see that some of the PDFs demonstrate the kurtosis of the Laplacian and some of them appear more Gaussian. We can also see there is significant variation in the rate of tail fall-off on the semilog plot (right panel). In general, this is a rather extreme case of variance distribution for combined

power spectrum measurements, and unlikely to be encountered in practice, since it essentially requires extremely uneven spatial mode or time sampling.

In practice, error distributions are more likely to be Gaussian, e.g. Tan et al. (2021). We can say from this analysis and the highly Gaussian error distributions in Fig. 9 of that work that the power spectrum variance distributions for any spectra that were incoherently




Figure 4. Collection of distributions of independent Laplacian sums after 50 additions, with scale parameters drawn from a power law with index 0.4 and range 1 to 1000 (top right panel of Fig. 3). Since the power-law distribution of lambdas can result in wildly different final variances, each resulting variable has been standardized. We also plot a standard normal distribution for reference. The excess kurtosis is visible in realizations that resemble the original Laplacian distribution even after the summations. These are highly extreme cases that are unlikely to be encountered in practice. More realistic scale parameter distributions tend to produce more Gaussian results.

combined in HERA Collaboration (2022) were probably strongly concentrated. As another practical example, we show histograms of scale parameters that contribute to different spherical k -modes in the ϵ PPSILON estimator from one night of MWA data in Fig. 5. The scale parameters are concentrated over approximately one order of magnitude. When we calculate Q using these scale parameters, we find that the values are extremely close to 1, indicating that there is negligible error in the uncertainty estimate by assuming a Gaussian distribution and quoting twice the standard deviation as the 97.7 per cent confidence interval. Here we have used all of the measured voxels that contribute to a spherical bin, including those that lie in the ‘foreground wedge,’ a region of power spectrum space that is often excluded from spherical power spectrum averages due to the presence of immense foreground contamination (Datta, Bowman & Carilli 2010; Morales et al. 2012; Trott, Wayth & Tingay 2012; Liu et al. 2014), which is unlike what is done in practice with MWA power spectra. Excluding modes with significant foreground contamination will use fewer bins, and therefore generally increase the Q -value in practice. However, this selection will not change the fact that the noise levels are generally homogeneous enough to produce approximately Gaussian sampling distributions, assuming enough independent samples are combined. Furthermore, ϵ PPSILON employs inverse variance weighting, which will homogenize the samples even more except in the unlikely event that the errors in the estimates are larger than the estimates themselves. One would need a model of the uncertainty in these particular estimates to investigate this fully.

Our analysis in this section as well as the previous one leads us to suggest that analysts should be cautious when incoherently combining power spectrum samples with vastly different statistics. We find that if variations of the scale parameters of the (potentially inverse variance weighted) power spectrum samples can be constrained to one order of magnitude, then approximate Gaussianity can be expected assuming more than a handful of samples are combined. About $N > 20$ appears sufficient in Fig. 3, though the exact case should be evaluated by the analyst if there is any question. Including the occasional noisy sample does not totally prevent approximate Gaussianity. For example many of the Q -values are somewhat close to 1 in the bottom-right four panels of Fig. 3 after 50 additions,

though shallower power laws are less reliable. We still urge caution for analysts using inverse variance weighting since estimates of the scale parameters for the noisiest samples may themselves be the estimates with the most error. While it is difficult to speak for every case, and we have yet to portray the effects of statistical dependence, we expect that practical circumstances will produce approximately Gaussian measurements of spherical power spectra more often than not, and so this assumption is usually justified.

6 DEPENDENT RANDOM VARIABLES

Another central component to most basic statements of the CLT are that the data be statistically independent. Complementary to the mathematical experiment depicted in Fig. 1, it is possible to combine statistically dependent but identically distributed data in such a way that Gaussianity is never achieved even in the infinite limit. An example of a potential occurrence of this in 21-cm power spectrum estimation is neatly described in Appendix A of Tan et al. (2021), where it is shown that a persistent skew can be retained with a particular choice of delay spectrum estimator. This is also an interesting example in the sense that the marginal distribution of each summand in the estimator is Laplacian, but due to the statistical dependency of the summands, a skew is produced.

Nevertheless, as discussed in Section 2, there exist statements of the CLT that relax the requirement that all the samples in question be independent. In particular, as long as the statistical dependency only has a finite range among the samples, then Gaussianity can be obtained in the infinite limit. In this section, we explore how correlated thermal noise, which may exist at several stages of the pipeline in different forms depending on the particular estimator, behaves under averages. As an example of how correlations might arise in the thermal noise, consider that all 21-cm power spectrum estimation pipelines make use of a spectral tapering function in the frequency domain before Fourier transforming along this axis. This translates to a convolution in the dual space, meaning that samples which were originally uncorrelated in that space can become correlated. Dependency can also arise in gridded power spectrum estimators, since the information from a given visibility

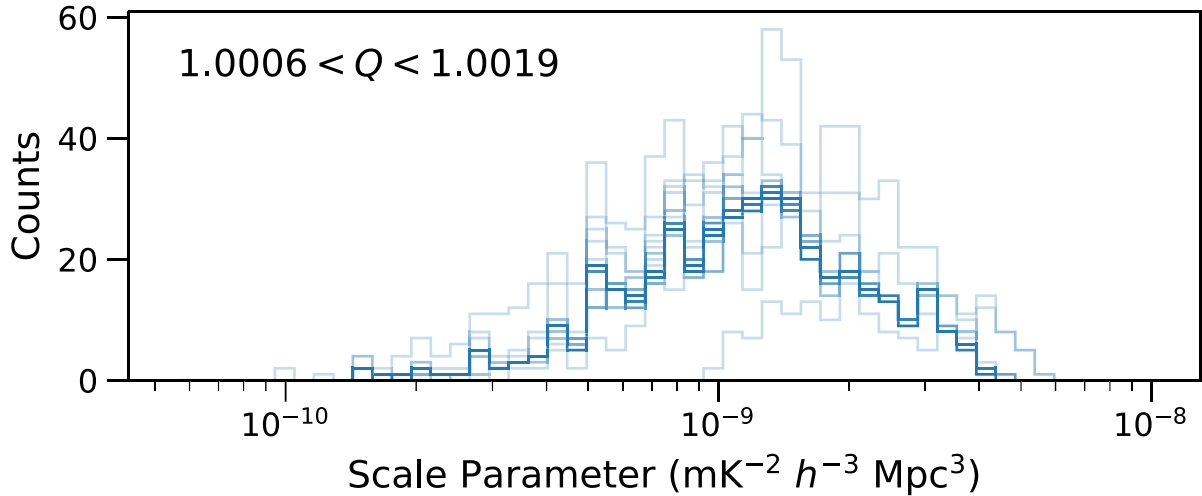


Figure 5. Histograms of scale parameters for 3d power spectrum bins estimated by ϵ PPSILON on one night of data, excluding data that contribute to the two lowest spherical k -modes. Each individual histogram (plotted with an opacity of 0.2 – darker shades arise from overlap of multiple histograms) represents the scale parameters of the 3d bins that contribute to a given spherical wave number. The scale parameters are concentrated over approximately one order of magnitude. Calculating Q from these scale parameters produces highly concentrated values close to 1, suggesting that each resulting error distribution would be highly Gaussian if statistical dependence between nearby voxels were negligible, and even if inverse variance weighting was not applied.

is usually assigned to several uv -pixels, meaning that there will be correlations in the perpendicular wave modes (which we ignored in the construction of Fig. 5). These correlations can produce non-trivial effects on the noise properties of spherically binned power spectra. The range of statistical dependency amongst the samples is highly idiosyncratic to the estimator in question, and a fully general treatment is beyond the scope of this paper. We therefore encourage analysts to understand the range of statistical dependencies in their data set and how they interplay to provide the ultimate sampling distribution of their estimate. We explore a particularly well-behaved and practical example next, in order to demonstrate the possibility of Gaussianity.

6.1 Exploration Using a Fringe Rate Filter

As a toy example, we explore the use of fringe rate filters (Parsons et al. 2016), which have been deployed to some extent in the HERA pipeline (HERA Collaboration 2022, 2023) and may be useful for further improvement of power spectrum upper limits. In short, a fringe rate filter can be implemented as a finite impulse response filter in time. One convolves a particular weighting function with the visibilities as a function of time, where the weighting function is chosen to enhance sensitivity to certain fringe rates (Fourier dual to the time axis of the data). Resulting samples will have some correlation length in time depending on the shape of the filter. Depending on how these correlated samples are used, these correlations can track all the way to the power spectrum estimate. If power spectra from correlated times are averaged incoherently, then depending on the strength of the correlation, the resulting error distribution may appear more or less Gaussian. For example, if one were to average perfectly correlated samples with one another, the sampling distribution would be unchanging at any length of averaging. This is an impractical case, since one would never intentionally generate a sequence of perfectly correlated data to average together; however, it suggests that if one does produce strong correlations, then one may find that the error distribution only slowly migrates towards Gaussianity, if at all.

We explore this possibility by implementing a fringe rate filter on a dynamic spectrum of Gaussian noise that is uncorrelated in

both the frequency and time direction. To be clear, the following implementation of a fringe-rate filter and power spectrum estimation is not an exact replication of what is used by HERA or any other telescope currently. It is a stripped-down version that is meant to expose the effect of averaging correlated data on the 21-cm error distribution. We implement a naive bandpass filter in fringe-rate space that simply multiplies included modes by 1 and multiplies all other modes by 0. We express this in terms of linear algebra as follows. Let \mathbf{n}_i be the noise samples for the i th baseline. Let \mathbf{T} be the operator that performs the timelike Discrete Fourier Transform (DFT). Then let \mathbf{D} be a diagonal matrix that has a value of 1 for all included modes and 0 for all excluded modes. The fringe-rate filtration operator is then

$$\mathbf{F} \equiv \mathbf{T}^{-1} \mathbf{D} \mathbf{T}. \quad (36)$$

When the number of fringe rates modes is an odd number, this operator has a tidy closed form solution. The j kth component of \mathbf{F} is given by

$$F_{jk} = \frac{e^{2\pi i f_c (t_j - t_k)}}{N_t} \frac{\sin [(2N_f + 1)\pi \Delta f (t_j - t_k)]}{\sin [\pi \Delta f (t_j - t_k)]}, \quad (37)$$

where N_t is the total number of times in the samples to be filtered, f_c is the central fringe rate to be included, t_j is the j th integration time in consideration, $2N_f + 1$ is the total number of fringe rate modes being included (i.e. N_f is the number of modes to either side of the central fringe rate being included), and Δf is the fringe-rate spacing produced by the usual discrete Fourier transform conventions. This is sometimes referred to as the Dirichlet kernel and looks like a periodic sinc function (Stein & Shakarchi 2003).

In any case, the fringe-rate filtered noise, $\tilde{\mathbf{n}}_i$, is given by

$$\tilde{\mathbf{n}}_i = \mathbf{F} \mathbf{n}_i. \quad (38)$$

Let \mathbf{N}_i be the corresponding noise covariance matrix. Then the covariance matrix after fringe rate filtering, $\tilde{\mathbf{N}}_i$, is

$$\tilde{\mathbf{N}}_i = E [\mathbf{F} \mathbf{n}_i \mathbf{n}_i^\dagger \mathbf{F}^\dagger] = \mathbf{F} \mathbf{N}_i \mathbf{F}^\dagger. \quad (39)$$

In the special case of uncorrelated Gaussian noise, where \mathbf{N}_i is proportional to the identity (say with diagonal element, σ_i^2), we

have

$$\begin{aligned}\tilde{\mathbf{N}}_i &= \sigma_i^2 \mathbf{F} \mathbf{F}^\dagger \\ &= \sigma_i^2 \mathbf{T}^{-1} \mathbf{D}^2 \mathbf{T} \\ &= \sigma_i^2 \mathbf{T}^{-1} \mathbf{D} \mathbf{T} \\ &= \sigma_i^2 \mathbf{F},\end{aligned}\quad (40)$$

where we have cancelled inverses where appropriate and made use of the fact that \mathbf{D} is Hermitian and $\mathbf{D}^2 = \mathbf{D}$. Given equation (37), we expect oscillating correlations and anticorrelations depending on the number of modes included, even when the original noise covariance is proportional to the identity.

If we want to include the effects of coherent averaging, we may make use of the fact that the covariance is a bilinear function of zero-mean random variables. Namely, if

$$Z = \frac{1}{N} \sum_{i=1}^N z_i \quad (41)$$

and

$$W = \frac{1}{N} \sum_{i=1}^N w_i \quad (42)$$

and furthermore, the covariance of any z_i with any other z_j or w_k is known, then (assuming all variables are zero-mean for notational simplicity)

$$E[ZW] = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N E[z_i w_j], \quad (43)$$

and

$$E[Z^2] = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N E[z_i z_j], \quad (44)$$

and similarly for $E[W^2]$. In other words, we may analytically produce the covariance matrix for coherent averaging of the fringe-rate filtered noise samples simply by summing the appropriate blocks of $\tilde{\mathbf{N}}_i$, and applying the $1/N^2$ factor. In the case of a weighted average, the weights must be tracked through the sums, and the $1/N^2$ factor changes to a quadratic function of the weights, but the overall process is the same. This allows us to short-circuit the simulation process. Rather than simulate independent samples and apply the fringe rate filter, we can instead generate jointly Gaussian samples that are consistent with the analytically propagated covariance matrix.

From here, we can form a (noise) delay spectrum estimate, \mathbf{p}_τ , by performing a delay transform (denoted by the operator \mathbf{P}_τ) on two independently simulated fringe-rate-filtered dynamic spectra and cross-multiplying, in effect simulating the cross-multiplication of the delay spectra of two independent baselines at each (potentially coherently averaged) time bin:

$$\mathbf{p}_\tau \equiv (\mathbf{P}_\tau \tilde{\mathbf{x}}_i)^* \circ \mathbf{P}_\tau \tilde{\mathbf{x}}_j. \quad (45)$$

where \circ is an elementwise-multiply. At this stage, the delay spectrum estimate (which we denote as vector-valued since we have computed one for each time at the specified delay mode) will be a Laplacian random variable at each time. The correlation between different times can be tracked using a relationship between the fourth and second moments of the $\tilde{\mathbf{x}}$. For real jointly Gaussian random vectors, this is known as Isserlis' theorem. It is generalized in Janssen & Stoica (1988) for complex jointly Gaussian random vectors (no circularity required). For white noise and two totally independent baselines, the

answer is⁸

$$E[\mathbf{p}_\tau \mathbf{p}_\tau^\dagger] = N_v^2 \tilde{\mathbf{N}}_i^* \circ \tilde{\mathbf{N}}_j \quad (46)$$

where N_v is the number of frequencies involved in the data.

We are interested at this moment in investigating under what circumstances an incoherent average over the time axis of these estimates produces an approximately Gaussian distribution. To do this, we generate many samples of a circular Gaussian random vector according to the appropriate covariance matrix that represents the fringe-rate filtration and coherent averaging process in several different cases of each. We then form simulated power spectra from these samples according to equation (45), incoherently average them, and histogram the results. We show this in Figs 6 and 7 for different filter choices.

In each simulation, we start with 300 uncorrelated time samples. Since we can analytically propagate the covariance through the coherent averaging and Fourier transform steps, we directly sample from a circular multivariate Gaussian with consistent covariance as if we applied a fringe rate filter and coherently averaged 10 time steps (leaving 30 steps to be incoherently averaged later). In Fig. 6, we set $N_f = 1$, which includes 3 modes out of the potential 150. This is an extremely narrow filter that is meant to represent the most extreme case of correlation structure that might be introduced by such a filter, and is not meant to represent a practical filter. In Fig. 7, we use $N_f = 30$ (about 20 per cent of the possible fringe modes), which is a potentially more realistic (though still quite narrow) setting depending on the goals of the analyst for the baselines in question. In each figure, we show the analytically calculated covariance matrix after fringe-rate filtering, as well as after coherent averaging and cross-multiplication with an independent baseline noise. We observe that for the extremely narrow filter, there is an induced correlation length in the delay spectra over approximately half of the samples in question. For the broader filter, the correlation length is about one sample to either side. The correlations extend around the time boundary due to the DFT implementation, which implicitly assumes a periodic input. These corner correlations could be avoided using an alternate implementation. The most similar alternative would be to apply a Toeplitz matrix made of discrete prolate spheroidal sequences (Slepian 1978; Ewall-Wice et al. 2021), rather than the circulant matrix generated by our DFT implementation. The resulting delay spectra (from pure noise) are incoherently averaged, divided by their analytically calculated standard deviation, and histogrammed in the bottom panel. We also show a counterpart histogram with no fringe-rate filtration and a standard normal pdf.

For the extremely narrow filter (Fig. 6), there is a pronounced discrepancy between the tails of the filtered and unfiltered histograms. While the unfiltered samples are close to Gaussian throughout the range, the filtered samples significantly depart from the Gaussian pdf at about 3σ . Both histograms appear to have exponential tails, meaning that the discrepancy worsens as we consider points further from the origin. While a comparison between the 97.7 per cent interval and 2σ may be relatively close, a figure of merit based further along the tail will inevitably produce strong discrepancies between significance in terms of the actual CDF and the Gaussian equivalent. This may be particularly important for inferences involving the

⁸It is common practice when delay transforming to apply a spectral tapering function, in which case N_v gets replaced with the sum of the squared values of the tapering function at each frequency. There will also then be correlations between different delay modes, which should be accounted for if combining delay modes in a spherical average.

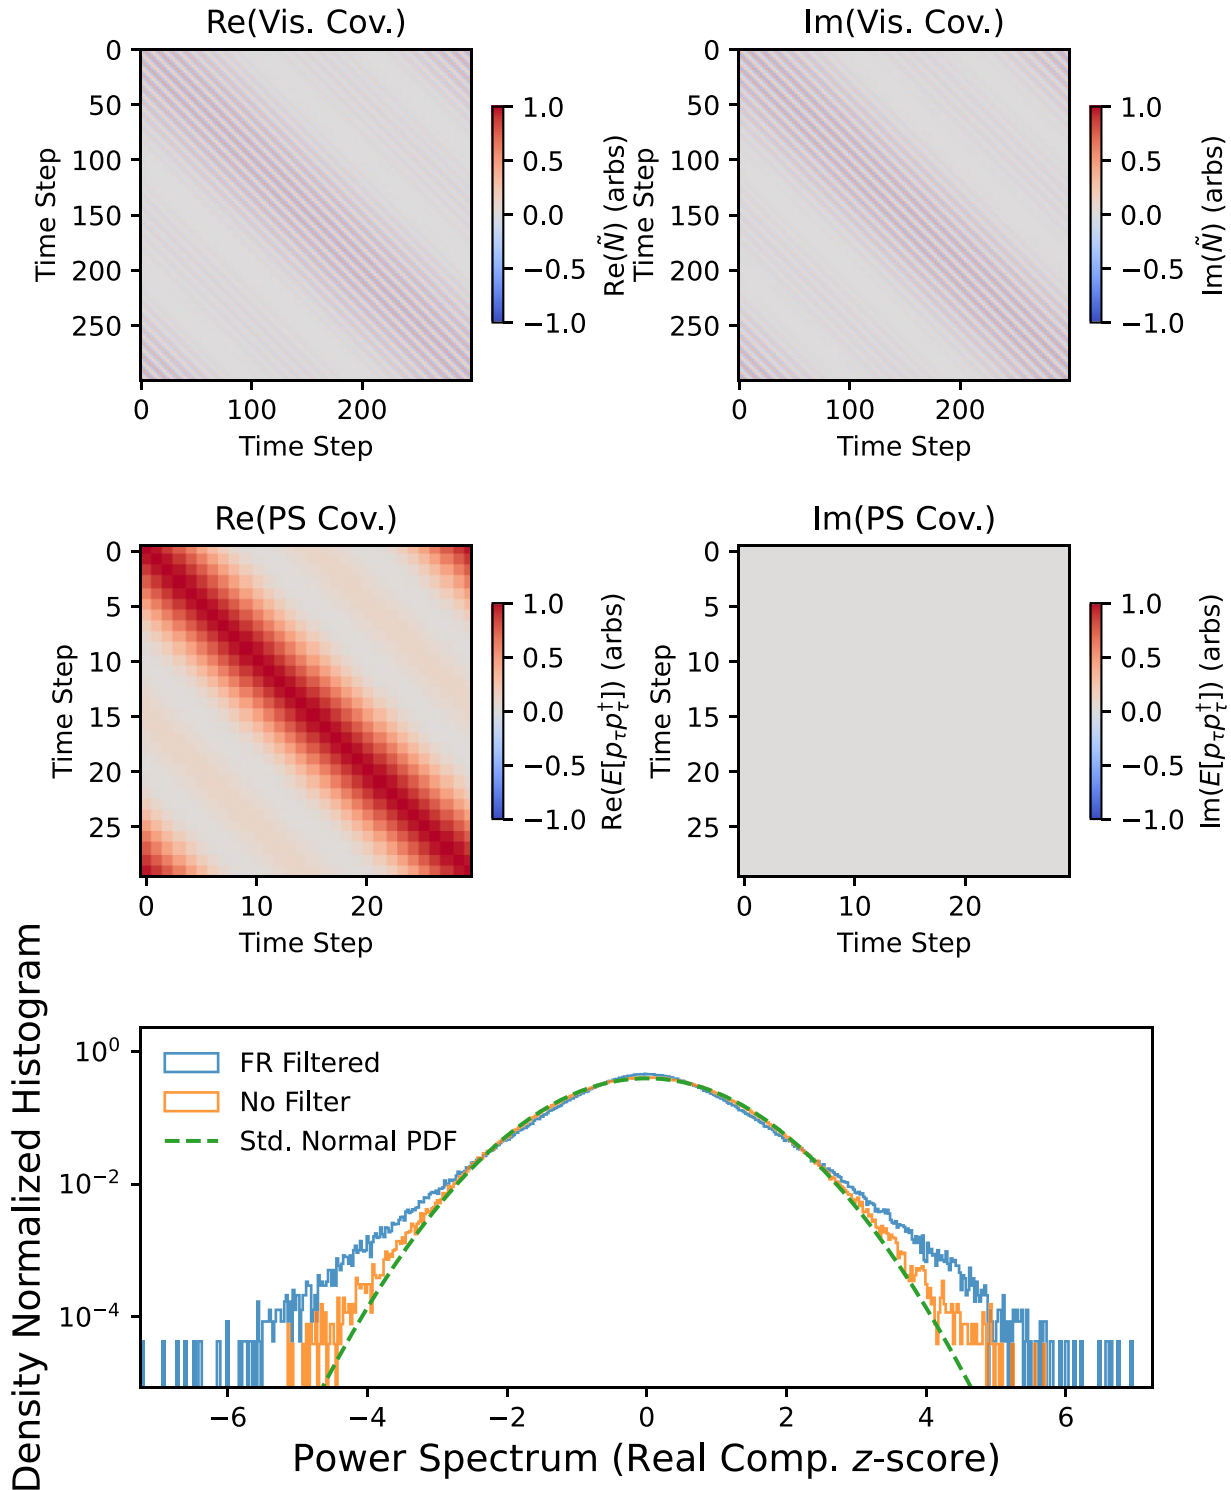


Figure 6. Top: Time-time covariance of independent and identically distributed noise after applying the DFT-based fringe rate filter described in the main text. In this case, only 3 fringe rates out of 150 are included (1 per cent of the available fringe modes), representing a worst case scenario in terms of induced correlation structure. Due to the DFT implementation of the filter, the correlation rises again at long lag, which could be avoided by e.g. convolving with discrete prolate spheroidal sequences (Slepian 1978; Ewall-Wice et al. 2021). Middle: Time-time covariance of the delay spectra after fringe-rate filtering, coherently averaging over 10 time steps, and cross-multiplying totally independent baselines with identical noise distribution. Since the baselines are identical, the imaginary part of the covariance matrix is exactly 0. Bottom: Histogram of simulated delay spectra (real component) with fringe-rate filtering (blue and orange, respectively). We normalize each datum according to its analytic standard deviation and plot a standard normal pdf for reference (dashed green). In this case, where 30 time steps have been incoherently averaged, the distribution of the delay spectra is clearly non-Gaussian when the filter is applied. Conversely, the unfiltered samples agree reasonably well with the standard normal pdf until the more extreme depths of the tails.

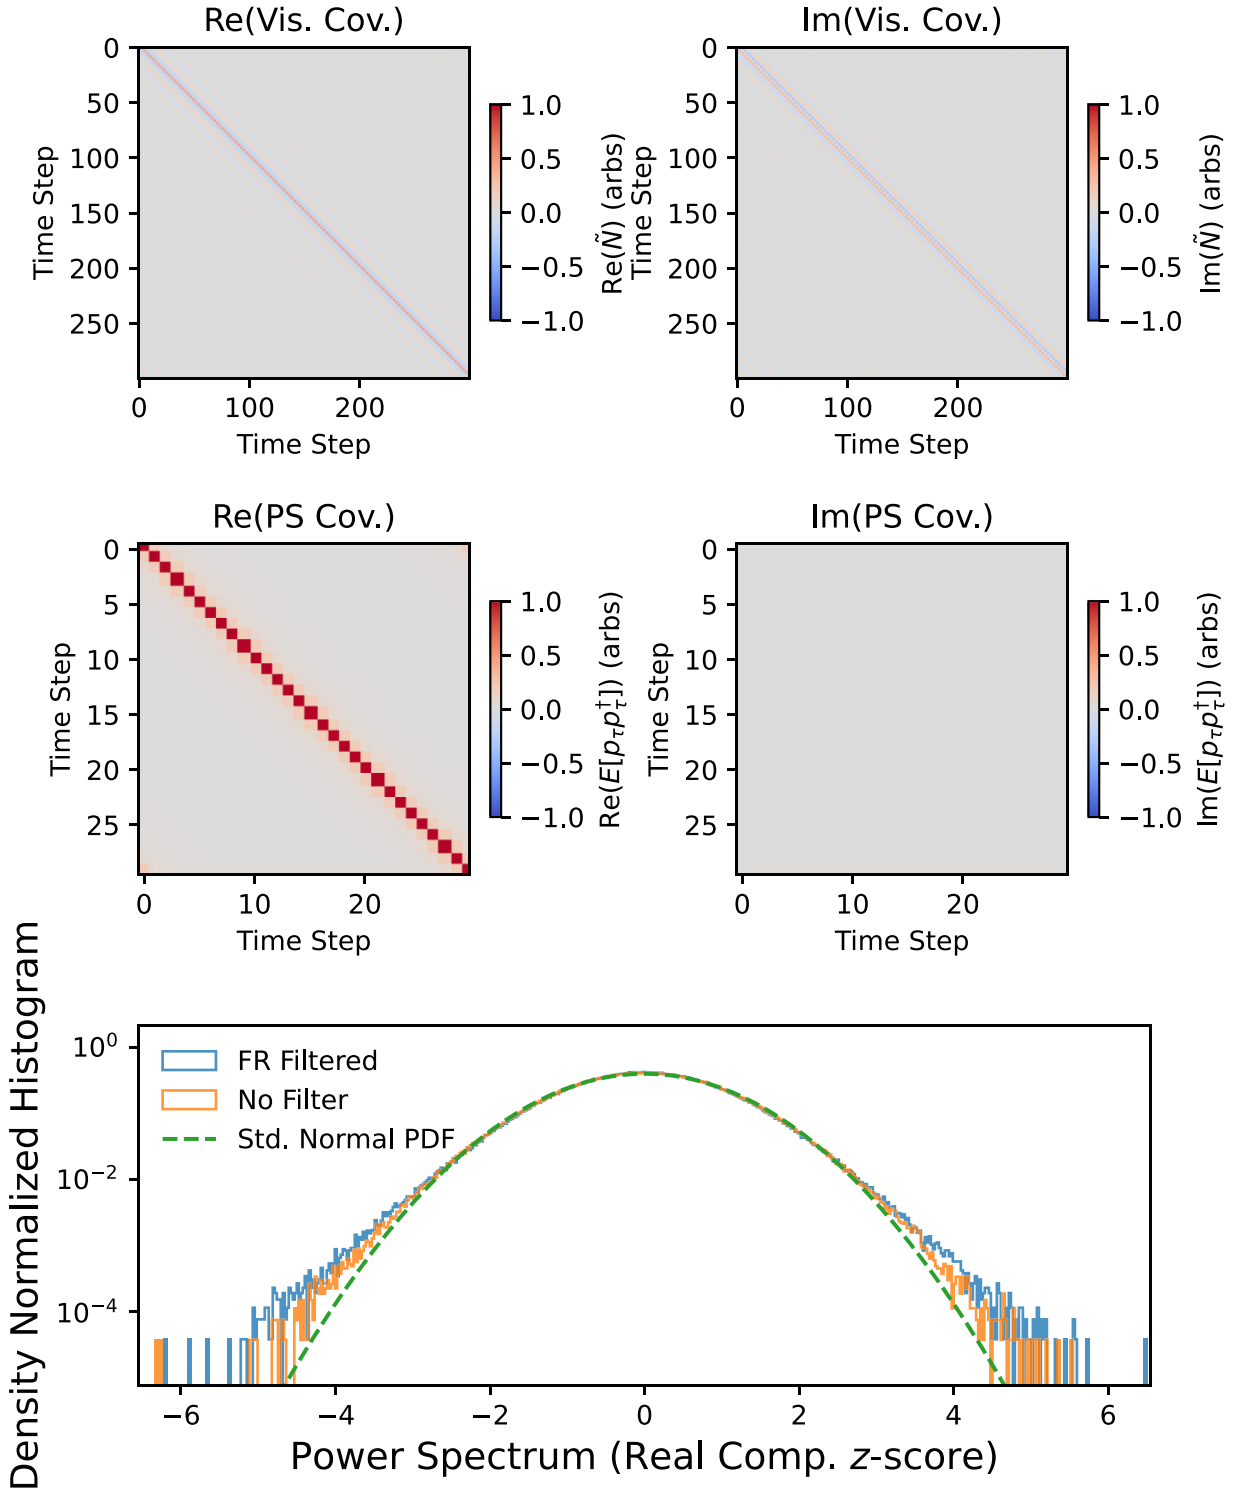


Figure 7. Same as Fig. 6, except including about 60 fringe modes out of 300 (20 percent of the available modes). This noticeably shrinks the amount of correlation in the samples, and the resulting delay spectrum distribution appears nearly identical to the case in which no fringe-rate filtering was performed.

cosmological 21-cm power spectrum signal, where the exact form of the likelihood of a model could drive the inference in a different direction. Fortunately for fringe rate filters, this effect appears to diminish rapidly as the size of the filter grows, as can be gleaned from Fig. 7. However, due to the simplicity of this analysis (single baseline-pair, no spherical averaging, DFT implementation), and the

potential ramifications for cosmological inference, we recommend that analysts assess the effects of potential correlations (or more general statistical dependence e.g. Appendix A of Tan et al. (2021)), in their own context.

Similar to analyses in previous sections, we find the ratio of the 97.7 percent confidence interval to twice the standard deviation (Q -




Figure 8. Ratio of 97.7 percent confidence bound to twice the standard deviation for fringe rate filters of increasing size (solid), where the colours of the line indicate how many time steps were combined incoherently. We also plot non-filtered counterparts for each choice of incoherent averaging length (dashed). Results between the filtered and non-filtered counterparts are similar so long as the filter is sufficiently wide and enough spectra are combined incoherently. This plot suggests that so long as correlation lengths are relatively small in the data (a few samples at most), there is almost no difference between a treatment that takes these correlations into account and one that does not, though measures at further points on the tail may be more discrepant.

value; equation (19) for the filter in different circumstances, shown in Fig. 8. In all cases, we consider 300 time steps, but within each realization we adjust the size of the filter or the coherent averaging length, thus adjusting the number of incoherent averages as well. We then construct delay spectrum samples and statistically estimate Q from them (solid lines). For each choice of incoherent averaging length, we also plot an estimate from unfiltered data but otherwise identical settings. We infer from the plot that this statistic quickly becomes limited by the incoherent averaging length rather than the correlation structure as we increase the breadth of the filter. It also suggests that all but the narrowest filters will produce noise properties that hardly alter the 2σ significance level. We stress again that this statistic does not speak for the entire distribution, and one should investigate as far along the tail as is relevant to one’s analysis. However, the bottom panel of Fig. 7, which corresponds with the rightmost point of the blue line of Fig. 8, suggests that dramatic discrepancies in excess of what is observed in independent identically distributed noise may not appear until very high significance values when correlation lengths are less than a few samples.

7 SUMMARY AND CONCLUSIONS

In this work, we looked in detail at error distributions of common 21-cm power spectrum estimators. We focused on terms that only

contained noise, ignoring cross terms with the cosmological and foreground signals. We proceeded by calculating the exact sampling distribution of the noise-only terms and then investigating the effect of finite sums of random variables. The main guiding principle in this work was to understand when and why Gaussian approximations to the sampling distribution are appropriate. This is often attributed to the CLT, despite the fact that the data in question do not obey the hypotheses of the most commonly known CLT, and despite the fact that all central limit theorems refer specifically to convergence in an infinite limit often with no reference to the effect of (finite) partial sums. Nevertheless, we found that in all but the most extreme violations of the typical hypotheses (namely independent and identically distributed samples), the finite sums indeed produced approximately Gaussian sampling distributions. This behaviour is in congruence with a variety of CLTs that relax the classical hypotheses.

We broke power spectrum estimators into two groups: those involving direct squaring of the data in Fourier space and those that cross-multiply independent samples. For a single sample, these have exponential and Laplacian sampling distributions, respectively. We then addressed the problem of understanding the summation of independent but non-identically distributed variables of the given type. We found that as long as the data are roughly homoschedastic (e.g. standard deviations spanning one order of magnitude), that

one can typically expect an approximately Gaussian distribution by averaging a few dozen independent samples together. We also examined the violation of the independence assumption for the case of cross-multiplication estimators. We did this by implementing a fringe-rate filter on white noise and observing its effect on a single-baseline delay spectrum. While the treatment was highly simplistic, we expect that the lesson applies fairly generally, which is that approximate Gaussianity can still be expected as long as correlation lengths in the samples are less than a few samples and sufficiently many samples are averaged. However, long correlation lengths can produce distinctly non-Gaussian distributions. The detailed tracking of correlations and more general statistical dependencies among power spectrum samples is complicated and can appear in multiple stages of any given pipeline. We therefore recommend analysts investigate their own cases thoroughly.

Other than being an interesting mathematical exercise, this type of work is important for the scientific quality of upper limits on the cosmic reionization signal. Understanding the sampling distribution of one's estimator allows one to more accurately represent confidence intervals for the signal, and also allows one to write down more exact Bayesian likelihood functions. This ultimately impacts the quality of the statistical inference underlying our understanding of the astrophysical processes driving reionization.

ACKNOWLEDGEMENTS

This result is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 948764). This research was supported by NSF grant #1613855. We gratefully acknowledge Phil Bull and Jianrong Tan for extremely helpful comments and discussions.

8 DATA AVAILABILITY

No new data were generated as a part of this work. The MWA data used are publicly available using the MWA All-Sky Virtual Observatory (ASVO).⁹

REFERENCES

- Abramowitz M., Stegun I. A., 1964, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, ninth Dover printing, tenth GPO printing edn. Dover, New York
- Barry N., Beardsley A. P., Byrne R., Hazelton B., Morales M. F., Pober J. C., Sullivan I., 2019a, *Publ. Astron. Soc. Aust.*, 36, e026
- Barry N. et al., 2019b, *ApJ*, 884, 1
- Bennett C. L. et al., 2003, *ApJS*, 148, 97
- Billingsley P., 2012, *Probability and Measure*. Wiley Series in Probability and Statistics, Wiley, New York
- Breuer L., Baum D., 2005, *An Introduction to Queueing Theory and Matrix-Analytic Methods*. Springer, Berlin
- Byrne R., Morales M. F., Hazelton B. J., Wilensky M., 2021, *MNRAS*, 503, 2457

- Byron F. W., Jr, Fuller R. W., 1970, *Mathematics of Classical and Quantum Physics*. Dover Publications, Inc., New York
- Das A., Geisler W. S., 2021, *J. Vis.*, 21, 1
- Datta A., Bowman J. D., Carilli C. L., 2010, *ApJ*, 724, 526
- DeBoer D. R. et al., 2017, *Publications of the Astronomical Society of the Pacific*, 129, 045001
- Ewall-Wice A. et al., 2021, *MNRAS*, 500, 5195
- Furlanetto S. R., Oh S. P., Briggs F. H., 2006, *Physics Reports*, 433, 181
- Gallager R., 2013, *Stochastic Processes: Theory for Applications*. Cambridge University of Press, Cambridge
- HERA Collaboration, 2022, *ApJ*, 925, 221
- HERA Collaboration, 2023, *ApJ*, 945, 124
- Harris C. R. et al., 2020, *Nature*, 585, 357
- Hoeffding W., Robbins H., 1948, *Duke Math. J.*, 15, 773
- Imhof J. P., 1961, *Biometrika*, 48, 419
- Janssen P., Stoica P., 1988, *IEEE Trans. Autom. Control*, 33, 867
- Kern N. S., Liu A., 2021, *MNRAS*, 501, 1463
- Li W. et al., 2019, *ApJ*, 887, 141
- Liu A., Shaw J. R., 2020, *Publ. Astron. Soc. Pac.*, 132, 062001
- Liu A., Parsons A. R., Trott C. M., 2014, *Phys. Rev. D*, 90, 023018
- Mertens F. G., Ghosh A., Koopmans L. V. E., 2018, *MNRAS*, 478, 3640
- Mertens F. G. et al., 2020, *MNRAS*, 493, 1662
- Morales M. F., Wyithe J. S. B., 2010, *Annu. Rev. Astron. Astrophys.*, 48, 127
- Morales M. F., Hazelton B., Sullivan I., Beardsley A., 2012, *ApJ*, 752, 137
- Morales M. F., Beardsley A., Pober J., Barry N., Hazelton B., Jacobs D., Sullivan I., 2019, *MNRAS*, 483, 2207
- Randolph Nelson, 1995, *Probability, Stochastic Processes, and Queueing Theory: The Mathematics of Computer Performance Modeling*. Springer-Verlag, Berlin
- Offringa A. R., Mertens F., Koopmans L. V. E., 2019, *MNRAS*, 484, 2866
- Paciga G. et al., 2013, *MNRAS*, 433, 639
- Parsons A. R. et al., 2010, *AJ*, 139, 1468
- Parsons A. R., Pober J. C., Aguirre J. E., Carilli C. L., Jacobs D. C., Moore D. F., 2012, *ApJ*, 756, 165
- Parsons A. R., Liu A., Ali Z. S., Cheng C., 2016, *ApJ*, 820, 51
- Patil A. H. et al., 2017, *ApJ*, 838, 65
- Shevtsova I., 2010, *Dokl. Math.*, 82, 862
- Slepian D., 1978, *Bell Syst. Tech. J.*, 57, 1371
- Stein E., Shakarchi R., 2003, *Fourier Analysis: An Introduction*. Princeton University Press, Princeton, NJ <https://books.google.com/books?id=CJngEACAAJ>
- Tan J. et al., 2021, *ApJS*, 255, 26
- Thompson A. R., Moran J. M., Swenson George W. J., 2017, *Interferometry and Synthesis in Radio Astronomy*, 3rd Edition. Springer, Berlin
- Tingay S. J. et al., 2013, *PASA*, 30, e007
- Trott C. M., Wayth R. B., Tingay S. J., 2012, *ApJ*, 757, 101
- Trott C. M. et al., 2016, *ApJ*, 818, 139
- VanderPlas J. T., 2018, *ApJS*, 236, 16
- Virtanen P. et al., 2020, *Nat. Meth.*, 17, 261
- Watson G. N., 1966, *A Treatise on the Theory of Bessel Functions*. Cambridge University Press, Cambridge
- Wayth R. B. et al., 2018, *Publ. Astron. Soc. Aust.*, 35, 33
- Wilensky M. J., Morales M. F., Hazelton B. J., Barry N., Byrne R., Roy S., 2019, *Publ. Astron. Soc. Pac.*, 131, 114507
- Wilensky M. J., Hazelton B. J., Morales M. F., 2022, *MNRAS*, 510, 5023
- de Oliveira-Costa A., Tegmark M., Gaensler B. M., Jonas J., Landecker T. L., Reich P., 2008, *MNRAS*, 388, 247
- van Haarlem M. P. et al., 2013, *Astron. Astrophys.*, 556, A2

⁹<https://asvo.mwatelescope.org>

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.