

Deviations from normal distributions in artificial and real time series: a false positive prescription

Paul J. Morris¹,  ¹★ Nachiketa Chakraborty^{2,3}  and Garret Cotter¹

¹*Oxford Astrophysics, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*

²*Max-Planck-Institut für Kernphysik, Saupfercheckweg 1, D-69117 Heidelberg, Germany*

³*Data Assimilation Research Centre, University of Reading, Whiteknights Rd, Reading RG6 6BB, UK*

Accepted 2019 August 9. Received 2019 August 6; in original form 2019 June 12

ABSTRACT

Time-series analysis allows for the determination of the Power Spectral Density (PSD) and Probability Density Function (PDF) for astrophysical sources. The former of these illustrates the distribution of power at various time-scales, typically taking a power-law form, while the latter characterizes the distribution of the underlying stochastic physical processes, with Gaussian and lognormal functional forms both physically motivated. In this paper, we use artificial time series generated using the prescription of Timmer & Koenig to investigate connections between the PDF and PSD. PDFs calculated for these artificial light curves are less likely to be well described by a Gaussian functional form for steep ($\Gamma \gtrsim 1$) PSD indices due to weak non-stationarity. Using the *Fermi* LAT monthly light curve of the blazar PKS2155-304 as an example, we prescribe and calculate a false positive rate that indicates how likely the PDF is to be attributed an incorrect functional form. Here, we generate large numbers of artificial light curves with intrinsically normally distributed PDFs and with statistical properties consistent with observations. These are used to evaluate the probabilities that either Gaussian or lognormal functional forms better describe the PDF. We use this prescription to show that PKS2155-304 requires a high prior probability of having a normally distributed PDF, $P(G) \geq 0.82$, for the calculated PDF to prefer a Gaussian functional form over a lognormal. We present possible choices of prior and evaluate the probability that PKS2155-304 has a lognormally distributed PDF for each.

Key words: accretion, accretion discs – methods: data analysis – galaxies: jets.

1 INTRODUCTION

Time-series analysis is crucial to study periodic and transient phenomena in astrophysics. Rapid X-ray variability is a common feature in the time series of compact accreting objects such as Black Hole X-ray Binaries (BHXRBs) (Shakura & Sunyaev 1973) and Active Galactic Nuclei (AGNs) (Urry & Padovani 1995), which also display rapid variability in high-energy γ -rays (e.g. Nakagawa & Mori 2013). The physical origin of this variability is likely related to underlying physical processes occurring in the object, which can leave an imprint in the time-series properties (Sinha et al. 2018).

One common property of astrophysical time series used to examine variability in astrophysical time series is the Power Spectral Density (PSD), which quantifies the amount of power in given frequencies sampled by the time series. This is typically estimated by utilizing the method of Bartlett (1948), where the time series is divided into M non-overlapping sections for which the discrete

Fourier transform (DFT) is calculated, and the PSD is taken as the mean of the DFTs calculated for each individual section. The functional form of the PSD is often approximately that of a power law, which may include a characteristic break separating regions of different spectral indices (e.g. Uttley, McHardy & Papadakis 2002).

Another property of a time series is the Probability Density Function (PDF), which can be obtained by forming a histogram of the flux values of an individual light curve. A PDF effectively quantifies the probability of a particular source being observed at a given flux value. If a time series is represented by a model containing an additive sequence of independent random variables, then it is said to be linear. Here, the time-series y -values are expected to be normally distributed via the central limit theorem, where y represents a desired measured quantity such as flux. Accordingly, the PDF is well described by a Gaussian functional form. If the functional form of the PDF deviates from a Gaussian, this may be indicative of underlying physical processes occurring in the source of interest. One such distribution known to occur in nature is the lognormal distribution, which can occur in the case of the random constituent elements in a time series, elements being multiplicative

* E-mail: paul.morris@physics.ox.ac.uk

rather than additive (e.g. Uttley, McHardy & Vaughan 2005; Heyde 2010).

PDFs have increasingly been computed for astrophysical sources (Uttley et al. 2005), with lognormality ubiquitous in BHXRBs (Gleissner et al. 2004; Heil, Vaughan & Uttley 2012) and also inferred from PDFs calculated by flux binning several blazar light curves first in X-rays (Giebels & Degrange 2009) and optical (Smith et al. 2018) and later in GeV γ -rays (Kushwaha et al. 2017) and across the electromagnetic spectrum (e.g. Chevalier et al. 2015; Sinha et al. 2016). Though these characteristics have been speculated as applying to all accretion-powered compact objects exhibiting rapid variability (e.g. Uttley et al. 2005), the physical origin, and whether such behaviour generally applies to the blazar population, is currently unknown. The physical motivation for lognormally distributed time series in BHXRBs arises from accretion disc flicker models (e.g. Lyubarskii 1997; Kotov, Churazov & Gilfanov 2001; King et al. 2004; Arévalo & Uttley 2006). Here, fluctuations of the accretion rate propagate inwards through the disc causing the modulation of faster fluctuations at smaller radii. Fluctuations in the X-ray flux are proportional to variations in the accretion rate, resulting in the source PDF being accurately described by a lognormal function. Observations of long-term lags in optical emission relative to X-rays are consistent with this model as optical and X-ray-emitting regions of the accretion disc are modulated together (Breedt et al. 2009). It has been further speculated that lognormal PDFs from AGN jets could be indicative of a disc–jet connection (McHardy 2008).

In this paper, we use simulations of artificial time series to investigate the fundamental relationship between the PSD and PDF, both of which are important properties of astrophysical sources. We begin by investigating the impact on a steepening PSD spectral index on the functional form of the PDF, before looking at how changing the minimum and maximum frequencies sampled in the light curve affects the PDF. We then use these results to motivate a prescription for a false positive fraction, which quantifies the probability of incorrectly measuring the wrong PDF from a source in which the intrinsic PDF has a different functional form. Specifically, we consider the case where the true source PDF can only be either Gaussian or lognormal, and generate many artificial time series for each of these known PDF functional forms. We base our light-curve time sampling on *Fermi* Large Area Telescope (LAT) monthly light curves for our artificial blazar time series and incorporate the measured PDF spectral index into our false positive calculation. We conclude by computing the false positive posterior probability for determining the correct PDF functional form for the blazar PKS2155-304 as an example.

2 TIME-SERIES SIMULATIONS

Observed time series are necessarily discrete, which leads to difficulties computing the PSD. These include red-noise leakage (Harris 1978), the effects of which can be mitigated by simulating a longer light curve (≥ 10 times the desired time-series length) such that power is contained in frequencies lower than the minimum sampled by the time series (Uttley et al. 2002). The finite time resolution can also lead to aliasing effects (e.g. Kirchner 2005). This occurs when frequencies higher than the sampling frequency are present in the time series and add additional power to lower frequency elements and is prevented by Nyquist sampling where the data are sampled at least twice the desired maximum. When producing artificial time series, the effects of aliasing can be avoided by generating time series from a known PSD, in which the

resultant light curves are sampled at the desired Nyquist frequency (Uttley et al. 2002). Artificial, user-defined, time series alleviate further issues that may affect data time series such as uneven sampling or variable signal-to-noise levels. Additionally, time-series simulations can be used to generate an ensemble of ‘realistic’ light curves, i.e. those matching the cadence of observations and with comparable statistical moments. This ensemble can be used to provide statistical tests including computing the significance of a test; an example of this is found in Romoli et al. (2018). Presently, there are two main prescriptions used to simulate artificial light curves, which are outlined in the following subsections.

2.1 PSD functional forms

In the GeV γ -ray regime, the BL Lac and Flat Spectrum Radio Quasar (FSRQ) subclasses of AGNs, which are those that have their relativistic jets pointed along a sightline close to that of the Earth, have PSDs well described by a single power law with average spectral indices of 1.7 and 1.5, respectively (Abdo et al. 2010) (see also Nakagawa & Mori 2013). Accordingly, one method of generating artificial time series relies on the definition of a power-law PSD, $\mathcal{P}_{\text{PL}}(\nu)$, which quantifies the power per unit frequency at frequencies ν . This is defined as

$$\mathcal{P}_{\text{PL}}(\nu) = A\nu^{-\Gamma} \quad (1)$$

for some normalization A and power-law index Γ . Alternatively, many compact object sources have been shown to have a broken power-law (BPL) PSD, which has been calculated using X-ray observations of BHXRBs (e.g. Wijnands & van der Klis 1999; Pottschmidt et al. 2003; Belloni et al. 2005) and AGNs (e.g. Marshall 2015), which can also have BPL PSDs in the optical (Smith et al. 2018). These PSDs typically begin with a white-noise index of $\Gamma = 0$ that transitions to pink or red noise at a characteristic break frequency, which may be related to a characteristic cooling time-scale of the emitting population (e.g. Ishibashi & Courvoisier 2012).

The functional form of a BPL PSD is defined as

$$\mathcal{P}_{\text{BPL}}(\nu) = \begin{cases} A\nu^{-\Gamma_1} & \nu \leq \nu_b \\ A\nu_b^{(\Gamma_2-\Gamma_1)}\nu^{-\Gamma_2} & \nu > \nu_b \end{cases} \quad (2)$$

transitioning in spectral index at a given break frequency, ν_b . For observed time series, the PSD index often flattens towards the low-frequency end of the power spectrum due to finite observation length and sampling of the time series (Ishibashi & Courvoisier 2012).

2.2 Timmer–Koenig simulations

A popular prescription for the generation of simulated time series is the method of Timmer and König (Timmer & Koenig 1995, hereafter TK95). TK95 simulations essentially generate artificial light curves with power-law noise that are easily extendable to other user-defined forms and have Gaussian PDFs. The method starts assuming a power spectral shape describing a time series with a general power-law-type noise such as white (0.0), pink (1.0), or red (2.0) noise. Using this, real and imaginary parts of Fourier amplitudes are drawn from a Gaussian distribution with a normalization such that the variance is that of the observed (or user-defined) light curve. These Fourier amplitudes are used to reconstruct the time series, which is now a realization of the underlying distribution of the physical process driving variability that we wish to probe. As stated in TK95, this ensures that we have simulated light curves that preserve the observed variability

properties to 1st order (as mean and variance are matched with observations). In doing so, as usual the length or duration of the simulated light curve is a factor of ≥ 10 longer than the observed light curves to avoid red-noise leakage or loss of power in the longer time-scales.

The discrete inverse Fourier transform (DiFT) of this artificial PSD yields a simulated data set with a PSD consistent with the desired user-defined PSD. This method can also be used to produce artificial light curves with the same PSD as a given data set by performing a DFT and randomizing the amplitudes and phases in Fourier space before calculating the time series via a DiFT. The **TK95** method is popular because of its simplicity, and the artificial time series it produces, $y_{\text{TK95}}(t)$, have PDFs with fluxes that are distributed normally. This property also allows lognormally distributed time series to be produced by exponentiating the output normally distributed time series, i.e. $y_{\text{LN}(t)} = \exp(y_{\text{TK95}}(t))$ (Uttley et al. 2005). Unless otherwise explicitly stated, artificial time series used in this work have been generated using an input power-law PSD as defined in equation (1) and using the method outlined below in Section 2.2.

2.3 Emmanoulopoulos simulations

To generate a time series with a known PDF and PSD, the more sophisticated Emmanoulopoulos (Emmanoulopoulos, McHardy & Papadakis 2013, hereafter **EMP13**) method can be used. This method combines features from **TK95** and that of Schreiber & Schmitz (1996). It works by initially generating two time series of equal length: one from the **TK95** method, $x_{\text{TK95}}(t)$, and the second is drawn directly from the desired PDF, $x_{\text{PDF}, i}(t)$. The DFTs of both of these time series are taken, with the amplitudes from the **TK95** DFT combined with the phases from the PSD obtained from the DFT of $x_{\text{PDF}, i}(t)$ to give a modified PSD. The DiFT of this gives an adjusted time series, $x_{\text{adjusted}, i}(t)$. Elements in $x_{\text{adjusted}, i}(t)$ are then replaced by those from $x_{\text{PDF}, i}(t)$ according to the ranking of the former to give a new time series, with the desired PDF, $x_{\text{PDF}, i+1}(t)$. If $x_{\text{PDF}, i+1}(t) \neq x_{\text{PDF}, i}(t)$, then $x_{\text{PDF}, i+1}(t)$ replaces $x_{\text{PDF}, i}(t)$ until convergence of the two is found. This results in an artificial time series with both the desired PDF and PSD. This method is more computationally intensive than the **TK95** method because of the iterative element.

3 TESTS FOR NORMALITY

We aimed to initially establish the likelihood of simulated **TK95** time series being consistent with having normally distributed fluxes. A metric to assess the normality of a time series was therefore required, with three popular methods considered.

Commonly in astronomy literature, the PDFs are estimated from the light curves (time series) as a histogram with fixed uniform binning to obtain a PDF in a histogram form (e.g. Kushwaha et al. 2017; Shah et al. 2018). The best fit to this PDF is then assessed using a variety of standard models such as Gaussian and lognormal distributions, with the goodness of fit evaluated with a standard statistical test such as the reduced chi squared statistic, $\chi^2_{\mathcal{N}}$ (e.g. Hughes & Hase 2010), where \mathcal{N} is the number of degrees of freedom. The model with the best $\chi^2_{\mathcal{N}}$ value, i.e. $\chi^2_{\mathcal{N}} \approx 1$, is considered to best fit the histogram. This method is further discussed in Section 8. However, because the histogram has a dependence on the binning algorithm we decided to include additional statistical tests.

A common test used to assess the likelihood of a time series being normally distributed in astronomy literature (e.g. Kushwaha

et al. 2017) is the Anderson–Darling (AD) test (Anderson & Darling 1952). The AD test can be used to determine how likely a sample is to come from a specified distribution; however, Monte Carlo simulations have indicated that when assessing if a sample is consistent with being normally distributed, the AD test has less statistical power than a Shapiro–Wilk test (Mohd Razali & Yap 2011; Yap & Sim 2011). We therefore instead use the Shapiro–Wilk (SW) (Shapiro & Wilk 1965) as our test for time series normality. This test can also be used to test for lognormality by taking the natural logarithm of each time series and testing for normality in log space. Furthermore, the SW test is known to be reliably applicable up to relatively large data sample sizes of ≤ 5000 (Rahman & Govindarajulu 1997), thus is suitable for assessing whether the comparatively shorter time-series simulations here are normally distributed.

An important property of the SW test considered in this work is the p -value. This gives the probability of obtaining the data under the condition that the null hypothesis is true. The null hypothesis for an SW test is that the sample is normally distributed. We use the p -value to illustrate the average consistency of our **TK95** simulations to be able to produce normally distributed time series.

4 ANALYTICAL MOTIVATION FOR A TRANSITION AT $\Gamma = 1$

It has been noted that properties of **TK95** simulated time series vary as a function of PSD index. Vaughan et al. (2003) show that the distribution of variances of these time series transitions from a Gaussian for $\Gamma = 0$ and approaches a χ^2 distribution as the index becomes steeper. The interpretation of this is that the steepening PSD spectral index progressively diminishes the effect of high-frequency components, effectively reducing the total number of degrees of freedom in each simulated time series (see also Uttley et al. 2005). More recent work by Alston (2019) has shown that the exponential flux distributions at steep PSD indices can deviate from the expected lognormal distribution.

Analytically, this can be explained by first obtaining the total power in a time series by integrating the PSD (equation 1) with respect to frequency,

$$\int_{\nu_{\min}}^{\nu_{\max}} \mathcal{P}(\nu) d\nu = \int_{\nu_{\min}}^{\nu_{\max}} A \nu^{-\Gamma} d\nu = \frac{A \nu^{1-\Gamma}}{1-\Gamma} \Big|_{\nu_{\min}}^{\nu_{\max}}. \quad (3)$$

It can be seen from equation (3) that the pink noise case when $\Gamma = 1$ is unique as the denominator causes both integration limits to diverge. This means that any finite time series, which by definition has discrete frequencies, contains less power at either end of the PSD than an ideal pink noise power spectrum.

When $\Gamma < 1$, $(1 - \Gamma) > 0$. From equation (3), this causes the high-frequency terms to be significant and approach ∞ as $\nu_{\max} \rightarrow \infty$. This means that the PSD for a finitely sampled time series will contain less power for $\Gamma < 1$ than a pink noise PSD at higher energies. In contrast, the low frequency limit will reduce to zero as $\nu_{\min} \rightarrow 0$. This is important because it shows that when $\Gamma < 1$, the high-frequency PSD components are able to provide a significant contribution to the total power in the time series.

Conversely, if $\Gamma > 1$, $(1 - \Gamma) < 0$. Equation (3) shows that as $\nu_{\max} \rightarrow \infty$, the contribution from high-frequency terms approaches zero, so the relative power in them diminishes at steeper PSD indices and is finite for steep PSD indices. Here, the low frequency limit approaches ∞ as $\nu_{\min} \rightarrow 0$. It is therefore for $\Gamma \gtrsim 1$ that the PSD and time series become dominated by increasingly low-frequency components that contain a substantial fraction of the total power. If

the PSD index is steep enough, the time series will be dominated by a small number of low-frequency components, which we show later can yield time series that depart significantly from normally distributed PDFs.

Therefore, $\Gamma \approx 1$ is roughly where the contribution to the PSD from the highest sampled frequencies becomes subdominant relative to their low-frequency counterparts. Their influence on the behaviour of the artificial light curves begins to diminish, progressively reducing rapid scale variability for high PSD indices. Recent work by Alston (2019) has shown that mismeasurement of the PSD can affect the functional form of the PDF. We undertake simulations that look into the connection between the PDF functional form and PSD spectral index and quantify the impact of this effect on any deviation of simulated light curves from normally distributed fluxes in the following sections.

5 METHOD AND RESULTS

Initially, we wished to demonstrate the analytic result presented in Section 4, and test the relative contribution of the high-frequency PSD components towards the properties of the PDFs for artificial time series generated with different PSD spectral indices. We decided to use a BPL PSD as defined in equation (2) to restrict the power in the higher frequencies above the break frequency, ν_b . As we analytically expected any transition to occur around $\Gamma = 1$, we defined our BPL PSD to transition from white to red noise such that $\Gamma_1 = 0$ and $\Gamma_2 = 2$ in equation (2). We varied the value of ν_b logarithmically uniformly throughout the entire frequency range of the PSD. For each value of ν_b , 10 000 light curves were simulated using the TK95 method outlined in Section 2.2, which was chosen over the EMP method as the latter requires significantly more computation time as a consequence of the iterative element. To mirror *Fermi* LAT blazar time series, these simulated data had their timing properties taken assuming monthly time bins over the ≈ 10 yr total time observations have been made using the LAT. Our artificial time series were therefore set up to have 128 elements with bin size equal to 1 month. In this test, we assessed the normality of simulated time series with a p -value returned from an SW test, and the lognormality by testing normality in log space with the same test. The null hypothesis is that data are normally distributed.

Fig. 1 shows that the p -value is close to zero for $\nu_b \approx \nu_{\min}$ when the PSD is approximately pure red noise and increases before converging to ~ 0.5 as $\nu_b \rightarrow \nu_{\max}$ (approaching white noise), thus demonstrating that a transition occurs between $\Gamma = 0$ and 2. Intermediate points shift ν_b and the transition point between the two. Fig. 1 also shows that the highest frequencies can significantly contribute towards a time series with a white to red noise PSD if the break frequency is greater than the mean logarithmic frequency of the PSD. It also infers that the characteristic shape of the PDF is not consistent with being Gaussian in functional form for PSDs with a substantial red-noise spectral component. This is because a low ν_b effectively makes the PSD close to pure red noise, lowering the relative contribution from the high-frequency components and reducing the total number of degrees of freedom, as discussed in Vaughan et al. (2003). The figure also indicates greater average consistency with the null hypothesis being true for Gaussian PDFs over lognormal for $\nu_b \gtrsim 10^{-7}$ Hz, yet it should be noted that in this frequency range the null hypothesis cannot be rejected to a significant confidence level in either case as the p -value is above the 2σ threshold of 0.05. The lower average p -value for the SW test for a normally distributed PDF in log space relative to in linear space indicates that the y -values of the artificial time series are inconsistent




Figure 1. Figure showing the evolution of the p -value returned from an SW test (see Section 3) for an ensemble of 10 000 TK95 simulated time series generated with a broken power-law PSD transitioning from white to red noise spectral indices at a variable frequency, ν_b . The figure shows the dependence of the mean p -value on ν_b . Towards the left-hand side of the plot, the PSD spectral break is at low frequencies, increasingly reducing the PSD to a pure red-noise power law. Towards the right-hand side, a white-noise PSD is recovered as the simulated time series have fluxes consistent with being normally distributed. The red triangles indicate tests for normality in linear space, while the blue circles represent tests for normality in logarithmic space.


with being normally distributed in log space for a greater proportion of realizations than those with PDFs that are inconsistent with being normally distributed in linear space.

Secondly, it was decided to investigate the effect that changing the PSD index has on the likelihood of obtaining normally distributed artificial time series. In this test, for the sake of simplicity, it was assumed that simple power law as in equation (1) accurately quantifies the PSD, as appears to be the case for blazars (e.g. Abdo et al. 2010; Nakagawa & Mori 2013). The PSD index for blazars is typically in the range $\Gamma \approx 1$ –2 (e.g. Abdo et al. 2010; Nakagawa & Mori 2013; Sobolewska et al. 2014), so this was included in our range that spanned from white noise ($\Gamma = 0$) to $\Gamma = 3$.

For each PSD index, 10 000 light curves were generated using the TK95 method outlined in Section 2.2. We find that the results show high consistency for a sample size of $\gtrsim 100$ artificial time series generated at each value of Γ . The statistical properties of each time series were evaluated in accordance to the statistical tests outlined in Section 3, with a p -value returned for the SW test. Our main findings may be summarized as follows:

- (i) As expected, a transition occurs for $\Gamma > 1$, with the average p -value decreasing for steeper PSD indices, for TK95 simulated time series. For $\Gamma > 1$, there is a sharp decline in the average p -value, thus these artificial time series are more likely to be inconsistent with the null hypothesis of having normally distributed fluxes. This is shown in Fig. 2 with the mean p -value from SW tests plotted against Γ .
- (ii) The functional form of the curve described above is approximately that of a sigmoid, and is discussed in Section 6.
- (iii) At steeper PSD indices, especially above $\Gamma > 2$, the time series PDF is, on average, not obviously consistent with a physically motivated PDF functional form such as a Gaussian or lognormal.

(a) Figure demonstrating the effect of varying the PSD index on the p -value for a Shapiro-Wilk test assuming a power law noise PSD. The artificial TK95 generated light curves in this calculation were 128 months long and binned in one month time intervals. It can be seen that the p -value sharply decreases at around $\Gamma = 1$, reducing the number of artificial time series which are consistent with having normally distributed time series. Above $\Gamma \approx 2$, the null hypothesis can be rejected to $> 2\sigma$. The mean p -values are calculated using 10,000 simulations for each PSD index. The fiducial fit representing the blue solid line in the top panel is described by Eqn. 4.



(b) Sample time series, PSDs and PDFs for a sample of power law spectral indices used in the top panel. It can be seen that the light curves become increasingly dominated by the low frequency components, which can lead to the PDF deviating from a Gaussian functional form. The PDFs here have been generated using 10 uniform bins in linear space.

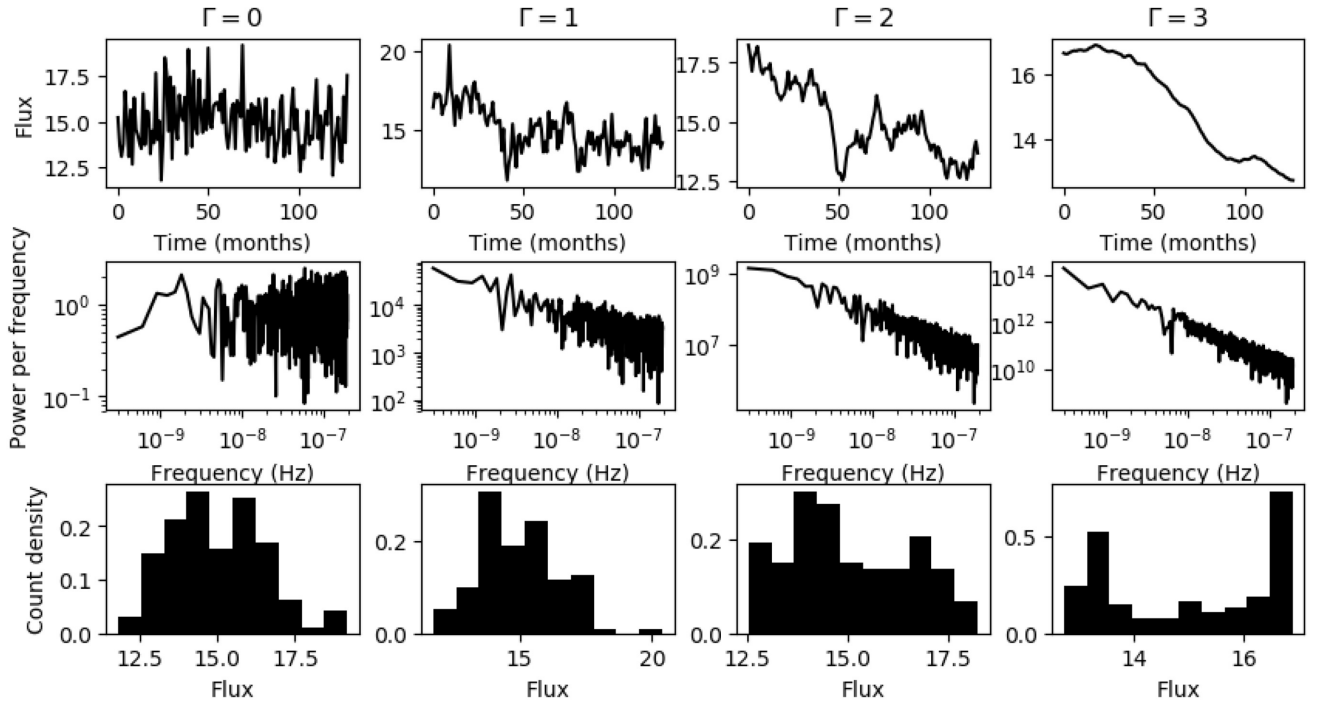


Figure 2. (a) Artificial light curves generated from the TK95 method show reduced p -values for steeper PSD indices, implying greater deviations from normality. (b) Demonstrates that this is because the time series becomes increasingly dominated by low-frequency components. These time series were generated with fixed mean flux values of 15, standard deviations of 1.5, and PSD normalizations of $A = 1$.

Fig. 2 shows the dependence of the mean p -value on PSD index. By comparison with Fig. 1, it can be seen that the p -value level for white noise in Fig. 1 is equal to the p -value before $\Gamma = 1$ in Fig. 2, thus in this region on average there is little difference to time series generated from white-noise PSDs. Fig. 2 also offers a qualitative

explanation that explains why the PDF functional form becomes consistent with being non-Gaussian for steeper PSD indices.

Panel (b) on this figure shows example time series, PSDs, and PDFs for integer PSD indices from $\Gamma = 0$ to 3. The time series corresponding to a PSD index of $\Gamma = 0$ is effectively white noise,

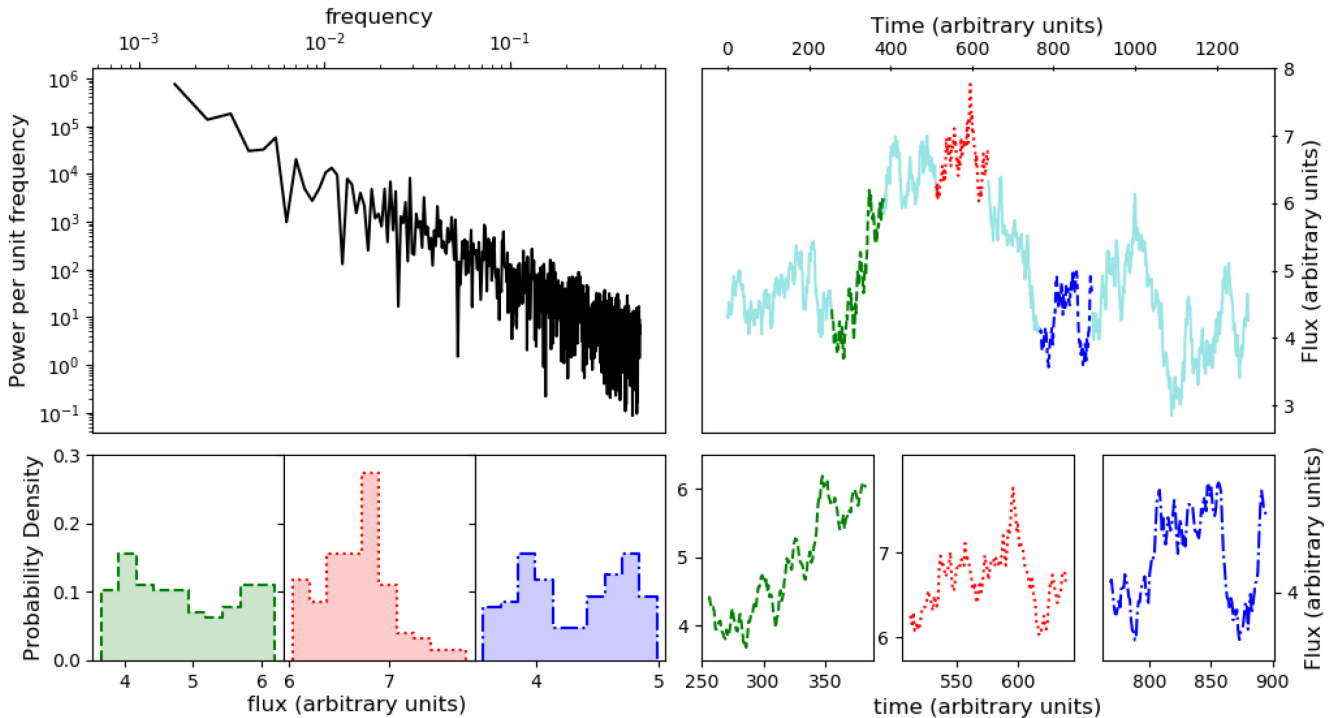


Figure 3. Figure demonstrating the non-stationarity of time series for a power-law PSD with an index $\Gamma = 2$. The top left panel shows the user-defined long-light-curve PSD, while the top right shows the corresponding long artificial light curve generated from it. The coloured sections (green dashed, red dots, and blue dot-dashed) in the long time series, which are representative of [TK95](#) generated light curves, are plotted in the bottom right subplots. Their corresponding PDFs are shown by the colour and border style matching histograms below the PSD. It can be seen that light curves sampled from the same long time series can have PDFs that differ vastly in functional form, indicating a lack of stationarity. This highlights the difficulty associated with correctly measuring the PDF for an object whose light curve is only sampled over a relatively short amount of time.

and the PDF functional form is approximately Gaussian. For $\Gamma = 1$, there is less rapid variability in the time series, although the PDF does not appear to deviate significantly from a normal distribution. In the more extreme example cases of $\Gamma = 2$ and even more so for $\Gamma = 3$, the time series has even less structure and the PDF has almost a double-peaked structure. In the case of $\Gamma = 3$, the time series is almost completely dominated by the large amount of power concentrated at the lowest frequencies. The long wavelengths corresponding to these frequencies largely determine the light-curve properties. The corresponding example PDF for $\Gamma = 3$ is clearly significantly different from having a Gaussian or lognormal functional form. These properties were common for time series generated from power-law PSDs with these spectral indices.

This is further explored in Fig. 3, which demonstrates that short light curves obtained from the same simulated long light curve and therefore have been created from the same PSD may not necessarily be determined to have the same PDF. In this figure, the user-defined PSD for the long light curve takes a power-law form with a red-noise spectral index of $\Gamma = 2$. The figure illustrates that the three PDFs corresponding to the three short time series drawn from the long light curve have different functional forms; the green dashed light curve has a PDF with a nearly flat distribution, the red dotted light curve exhibits a PDF with a single large peak, and the blue dot-dashed time series PDF looks bimodal. This figure highlights the difficulty in correctly measuring a true PDF for a source of which limited time-series data are available. Fig. 3 illustrates that the divergence from normally distributed fluxes shown in Fig. 2 can be attributed to weak non-stationarity, where the variance of individual light-curve segments is not constant in time but instead

varies about a mean value determined by the underlying PSD (Uttley et al. 2005).

6 FUNCTIONAL FORM OF MEAN p -VALUE DEPENDENCE ON PSD INDEX

To quantify the impact that the PSD spectral index has on [TK95](#) time-series simulations, we performed an empirical fit to the data shown in Fig. 2(a). This functional form of this curve is given by that of a modified sigmoid, namely,

$$f(\Gamma) = \frac{\alpha \exp(\beta - \eta\Gamma)}{1 + \exp(\beta - \eta\Gamma)}, \quad (4)$$

where Γ represents the PSD spectral index assuming a PSD defined by equation (1) and the parameters α , β , and η are free. The parameter α effectively normalizes the curve, β represents an x -axis translation, and η quantifies the steepness of the decline in average p -value, which becomes more rapid as η increases.

In the following section, we describe further simulations undertaken to understand the influence of the frequency range of the PSD on the p -value versus Γ sigmoid curve and discuss their impact on the best fit of the functional form described by equation (4).

7 FURTHER TESTS

Having analytically motivated and demonstrated the decline of normally distributed [TK95](#) time series beyond $\Gamma \approx 1$, we decided to undertake further experiments to establish whether the characteristic shape outlined in equation (4) is influenced by other properties of




Figure 4. Figure showing the evolution of the best fit described by equation (4) as in Fig. 2 with respect to the length of the time series. Artificial light curves were generated via the TK95 method and all have intrinsically Gaussian PDFs. They were tested for normality in linear space (solid lines) and log space (dashed lines) via SW tests. It can be seen that, for $\Gamma \leq 1$, the normalization is roughly constant for the former, but the artificial time series PDFs are on average less likely to be lognormal for longer light curves, indicated by the declining SW test p -value associated with the test for normality in log space. Additionally, the deviations from normality occur at a more rapid rate for longer time series. This figure illustrates the effect changing ν_{\min} has on TK95 time series produced from a power-law PSD.

TK95 simulations to test each of them. Specifically, we wished to establish if the position and severity of the break are influenced by the following:

- (i) The total length of each artificial time series. As a longer time series samples lower frequencies, this affects ν_{\min} .
- (ii) The amount of red-noise leakage in the time series.
- (iii) The binning of, or the smallest time-scale associated with, the time series. This changes ν_{\max} .

To look into this, further simulations were carried out. To address the first point, the time binning was kept constant as monthly but the total length of the time series was increased incrementally up to 800 elements, equivalent to ~ 67 yr. The same simulation set-up was taken, with a plot of mean p -value versus PSD index obtained for each time series length, with the mean taken from computing the SW p -value from 10 000 artificial TK95 light curves generated for each PSD index with intrinsically Gaussian PDFs. This plot was obtained for testing the mean p -values of the produced time series both in linear and log space, with a mean best-fitting parameter curve produced from these simulations by fitting equation (4) and varying α , β , and η .

The results of this test are shown in Fig. 4. It can be seen that although there is variation in the curve with respect to the total length of the time series, each curve still indicates that the generated artificial time series show, on average, increasing deviations from being normally distributed at higher PSD indices. However, while increasing the length of the time series has little effect on the shape of the curve for the linear space SW p -value for $\Gamma \leq 1$, the p -value testing normality in log space decreases as a function of time series length, corresponding to a reduction of the parameter α . To some extent, this is a trivial result as we would expect to recover the intrinsic PDF shape more easily when taking more samples as in the case of a longer time series, yet it highlights the need for caution




Figure 5. The effect of red-noise leakage as a function of PSD index. Here, artificial time series each containing 128 elements and with fixed ν_{\max} were generated from power-law PSDs via the TK95 method. The length of the long light curve, N_{long} , from which a shorter light curve of the desired length was extracted, was varied to investigate the effect of red-noise leakage, with the corresponding noisy power-law PSDs extending to lower frequencies for larger N_{long} . As leakage more strongly reduces the power contained in low-frequency components and transferring it to high-frequency components, this effect is more prominent in low- Γ regions where high-frequency components contribute significant power to the time series.

when attempting to determine the intrinsic PDF for a source with a relatively small time series containing few data elements. This arises because the SW test has more power for larger sample sizes, thus returns a result more likely to correctly reject the null hypothesis in the lognormal case.

Fig. 4 was generated from simulations of long light curves 10 times the desired length with noisy power-law PSDs (TK95). Accordingly, the PSD extends to lower frequencies as the length of the light curve increases. It follows that for a PSD index steeper than white noise, the ratio of power contained in the constant maximum sampled frequency (shortest sampled time-scale) relative to the variable minimum frequency decreases as the minimum frequency becomes smaller. This is shown in Fig. 4 by the steeper transition at $\Gamma \approx 1$ (corresponding to an increase in the best-fitting parameter η in equation 4) for longer time series as there is relatively less power contained at high frequencies for longer time series relative to shorter ones as the former have power-law PSDs that extend to a lower minimum frequency. This allows the higher frequency elements to have a greater influence on shorter time series, reflected in the less rapid decline of the mean p -value in Fig. 4 relative to those for longer time series.

The effects of red-noise leakage can be mitigated by simulating time series from PSDs that extend to a lower frequency range and drawing a light curve of the desired length from a longer time series (Vaughan et al. 2003). Therefore, to quantify the effect of this another test was undertaken, this time varying the length of the long light curve, N_{long} , from which the desired short time series (of 128 elements) was extracted. For higher values of N_{long} , the (power-law) PSD extended to a proportionally lower frequency range. Results are shown in Fig. 5. It is immediately obvious that at steep PSD indices ($\Gamma \gtrsim 1.5$) all of the curves show great consistency. This is because at steep PSD indices, the lower frequency components dominate the time series irrespective of leakage. Leakage more strongly reduces the power contained in lower frequency components in




Figure 6. As Fig. 4 but changing the time-scale associated with the smallest bin. Unlike Fig. 4, it can be seen that the bin time has no effect on the overall dependence of the produced time series with respect to normality. This is because although ν_{\max} is changed, the relative power contributions from each frequency element in the PSD remain the same.

the time series, indicated by the divergence of the curves for $\Gamma \gtrsim 1.5$. Generally, the curves have lower mean p -values for larger N_{long} when the effects of leakage are smaller. This is because red-noise leakage more strongly reduces the power contained in low-frequency components, effectively making the PSD power-law index less steep and changing the p -value to be consistent with that expected from a shallower PSD index.

A third test was undertaken, this time varying the minimum sampling time, which changes the maximum frequency in the PDF. The range explored was from 1 h to 1 yr, with common sampling such as daily and monthly time bins investigated. The results of this are shown in Fig. 6. It is clear that changing this parameter has almost no effect on the parametrization curves at all, as they are all close to being exactly superimposed. Changing the highest PDF frequency therefore has a negligible effect on the properties of TK95 time-series simulations. This is because the relative power in each frequency component is the same in this test, i.e. the time series has the same number of elements regardless of the sampling time. This effectively translates the PSD on the y-axis, whereas in the previous test increasing the length of the time series extended the frequency range and light curve length, thus changing the relative distribution of power into frequencies sampled by each artificial time series.

Throughout our tests, if the sample size of the time series is ≥ 30 , there is great consistency in the value of the mean p -value for $\Gamma \lesssim 1$ when the light-curve PDFs are consistent with a Gaussian functional form. This is shown in Figs 2(a), 4, and 6 and is equivalent to stating that α in equation (4) is approximately 0.5. This result is expected as the distribution of p -values should be flat if the null hypothesis cannot be rejected (Rice 2006; Murdoch, Tsai & Adcock 2008), i.e. TK95 time series PDFs are consistent with having normally distributed fluxes for $\Gamma \leq 1$. This result is perhaps more intuitively understood when considering the mean p -value when testing for normality in log space, as depicted in Fig. 4. Here, the p -value decreases as the length of the time series increases, implying that the p -value distribution becomes skewed to favour lower numbers as a greater number of TK95 simulations reject the null hypothesis. In this figure, as the time series length increases so does the statistical




Figure 7. Based on 10 000 TK95 generated artificial times series each with 128 elements (each being 1 month) and intrinsic Gaussian PDFs. It can be seen that for $\Gamma \leq 1$, a Gaussian is, on average, an appropriate functional form for describing the PDF when evaluated with a χ^2_N test. In this Γ range, a lognormal also often provides a good fit to the PDF, though the quality of the fit is on average poorer than the Gaussian. It can be seen that the quality of both fits declines rapidly beyond $\Gamma = 1$, after which neither distribution is likely to accurately describe the PDF. This is as a result of non-stationarity of the PSD as all segments of the long simulated time series have different properties (see also Alston 2019).

power, and a greater proportion of time-series realizations correctly have the null hypothesis, i.e. that they are normally distributed in log space, rejected. Therefore the distribution of p -values becomes skewed as it favours values closer to zero, which is reflected in the lower mean values.

8 FALSE POSITIVE RATE FOR PDFS

It should also be stated that there are numerous ways of estimating the PDF from data, including non-parametric ones (see Wegman 1972, for a summary). Each has its own distinct advantages and disadvantages, but a comprehensive discussion is beyond the scope of this paper. In astronomy literature, the functional form of the PDF is often evaluated by binning the time series and fitting to the resultant histogram (e.g. Guo et al. 2016; Kushwaha et al. 2017; Shah et al. 2018), with the comparison of different functional forms often completed using a χ^2_N test. So far, we have assessed deviations from normality in the PDFs of simulated TK95 series using the SW test as our metric. The most popular, physically motivated functional forms of the PDF found in the literature are Gaussian and lognormal, so we first extend the previous work to establish whether Gaussian TK95 PDF light curves can be erroneously measured as lognormal.

Fig. 7 shows the distributions of χ^2_N values for an ensemble of 10 000 TK95 generated time series. This reflects the result Fig. 2, although this figure shows that the best fit of a lognormal function to the histograms gets progressively worse on average for $\Gamma \gtrsim 1$, much like for the Gaussian fit. This is an important result as it captures the essence of statistical behaviour arising from loss of strict stationarity. Although these simulations deviate from being normally distributed at steep PSD indices, they do not show any preference for the lognormal models either. This is because for steep PSDs, the divergence of power at long time-scales implies divergence of the corresponding variance. And this renders higher moments of the PDF ill-constrained (and varying significantly from




Figure 8. Comparison of Gaussian versus lognormal fits to the PDFs of artificial time series. A lognormal model was deemed a better fit for $\chi^2_{N, \text{LN}} < \chi^2_{N, \text{G}}$, with the fraction of PDFs preferred indicated by the grey portion of the bars. The subplots divide the fits based on the values of $\chi^2_{N, \text{G}}$. Produced by generating 10 000 artificial **TK95** times series each with 128 elements (with a time-step of 1 month between adjacent elements) and intrinsic Gaussian PDFs. For clarity, only indices for which at least 30 convergent fits were obtained are displayed with a bar in the figure. It can be seen that in these simulations a normal distribution typically better describes the shape of the PDF than a lognormal function for low $\chi^2_{N, \text{G}}$ and when $\Gamma < 1$. However, above this index, especially when both fits are poor, the false positive fraction is around 50 per cent, which is indicated by the dashed horizontal line.

realization to realization). As a by-product, this diminishes the chance of falsely obtaining an ‘incorrect result’, which corresponds to identifying distribution as lognormal. This is owing to the binary choice, ‘normal versus lognormal’ provided. The true PDF, of course, could be more complex and also different from both choices. However, extracting this general form, which is well motivated physically, from the data is challenging when multiple, complex processes are at play.

Fig. 8 compares the $\chi^2_{N, \text{G}}$ values for Gaussian fits to those for lognormal fits for the same PSD index, for a range of values. The four bands shown are somewhat arbitrary, with the $\chi^2_{N, \text{G}}$ range corresponding to the Gaussian fits only. A lognormal model was deemed a better fit if $\chi^2_{N, \text{LN}} < \chi^2_{N, \text{G}}$. These cases are indicated by the grey portion of each bar, whereas $\chi^2_{N, \text{G}} < \chi^2_{N, \text{LN}}$ are shown in black. Here, the proportion of each bar that is grey is the false positive fraction, i.e. the percentage of time in which a lognormal PDF fit is erroneously preferred over a Gaussian functional form. It can be seen for this example that the incorrect intrinsic PDF is the preferred fit in ≈ 60 per cent of cases, with a false positive unsurprisingly less likely to be determined in the event of a good $\chi^2_{N, \text{G}}$ fit. Fig. 9 shows the number of counts used to normalize each bar in Fig. 8, and also shows the evolution of the $\chi^2_{N, \text{G}}$ distribution. It can be seen that the quality of fit of a Gaussian function to the histogrammed PDF worsens significantly as the PSD index steepens, verifying our result from Section 5. This result is important as it demonstrates the use that simulated time series can have in testing whether an observed time series has a correctly measured PDF.

We have investigated the suitability of **TK95** simulated time series for reproducing normally distributed artificial time series. Before applying these results to real time series, we briefly outline some additional caveats of generating artificial time series.




Figure 9. Based on 10 000 **TK95** simulated times series each with 128 elements (each being 1 month). This figure shows the raw counts in each $\chi^2_{N, \text{G}}$ band, from which the normalized fractions in Fig. 8 were calculated. It can be seen that good fits to the PDF from a Gaussian or lognormal function are more likely to be obtained when the PSD index is low.

One such caveat of generating artificial time series is that they are unable to reproduce the same skewness as observed for real time series. This has been showed by calculating the bicoherence, which quantifies the coupling between variations at different time-scales (and in principle can be used to distinguish between linear, Gaussian processes and non-linear processes, which can produce non-Gaussian PDFs), for real and artificial time series (Maccarone & Coppi 2002; Uttley et al. 2005). Presently, no such method for the generation of artificial time series exists that can include this feature, thus it is not accounted for here.

An additional caveat is that real time series can show strong non-stationarity, which manifest themselves as an additional distortion of the flux PDF. This has been demonstrated by monitoring the X-ray variability of the narrow-line Seyfert 1 (NLS1) galaxy IRAS 13224-3809 (Alston 2019; Alston et al. 2019), where the PSD is strongly non-stationary and the PDF deviates from the expected lognormal functional form. As the computation of the PSD becomes more challenging for smaller data sets, this effect is unlikely to be a major issue for the majority of time series, and can be tested for using the methods described in Vaughan et al. (2003) and Alston et al. (2019).

In the remainder of this paper, we apply our results to real blazar data. The objective is to prescribe a false positive fraction, specifically what fraction of the time one can expect to measure a PDF that is not intrinsically the PDF of the object of interest. When assessing real data, it is difficult to ascertain the true PDF functional form, therefore simulated time series with known PDFs and PSDs are a powerful tool that can be utilized to derive a false positive fraction. In this section, we assume that sources can have PDFs of a Gaussian or lognormal functional form, and ignore other possibilities. To begin with, let us assume that a PDF consistent with a Gaussian functional form is observed. We wish to know what is the probability that the true source PDF is Gaussian given we have measured it to be so, $P(\text{G}|\text{G}_{\text{PDF}})$. Bayes’ theorem (e.g. Sivia & Skilling 2006) tells us this is

$$P(\text{G}|\text{G}_{\text{PDF}}) = \frac{P(\text{G}_{\text{PDF}}|\text{G})P(\text{G})}{P(\text{G}_{\text{PDF}})}, \quad (5)$$

where $P(G)$ is the prior probability that the true PDF is Gaussian. We wish to evaluate the probability of obtaining a Gaussian PDF, $P(G_{PDF})$, which under our assumption that the true PDF must be either Gaussian or lognormal can only be obtained in two ways. These are either from correctly measuring a Gaussian PDF from a source with a true Gaussian PDF, or incorrectly measuring a Gaussian PDF from a source with an intrinsic lognormal PDF. Using $+$ or $-$ to denote either a correct or incorrect PDF measurement, and $P(LN)$ as the probability for a true lognormal PDF allows us to write $P(G_{PDF})$ as

$$P(G_{PDF}) = P(+|G)P(G) + P(-|LN)P(LN), \quad (6)$$

which, when substituted into equation (5), gives

$$P(G|G_{PDF}) = \frac{P(+|G)P(G)}{P(+|G)P(G) + P(-|LN)P(LN)}, \quad (7)$$

which defines the true positive fraction. Similarly, we can calculate the false positive fraction as the probability of measuring a lognormal PDF given the true PDF is Gaussian as

$$P(LN|G_{PDF}) = \frac{P(-|LN)P(LN)}{P(+|G)P(G) + P(-|LN)P(LN)}. \quad (8)$$

Equations (7) and (8) apply to sources with intrinsically Gaussian PDFs. In a similar fashion, we can calculate the true and false positive fractions for intrinsically lognormal sources using

$$P(LN|LN_{PDF}) = \frac{P(+|LN)P(LN)}{P(+|LN)P(LN) + P(-|G)P(G)}, \quad (9)$$

$$P(G|LN_{PDF}) = \frac{P(-|G)P(G)}{P(+|LN)P(LN) + P(-|G)P(G)}. \quad (10)$$

In these formulae, $P(G)$ and $P(LN)$ are the priors, while the variables $P(+|G)$, $P(-|G)$, $P(+|LN)$, and $P(-|LN)$ can be used from the time-series simulation techniques outlined in the previous sections. We demonstrate this in the remainder of the paper, demonstrating the use of this technique with a worked example.

8.1 A worked example

8.1.1 Fermi analysis

For our example object, we used a bright source because if the source in question is often only weakly or undetected, it may skew the shape of the PDF as flux values by definition cannot be negative. In this example, we investigate the blazar PKS2155-304, for which we produce a time series in the energy range 100 MeV–300 GeV by analysing publicly available *Fermi* LAT data.

To produce a time series, we first downloaded a photon file centred on the target source and containing all detected photons within a 15° radius that had been emitting during the first ≈ 10 yr of *Fermi* LAT observations, beginning 2008 August 5 and ending 2019 February 9. For the same interval, we downloaded the spacecraft file that details the relative orientation of the LAT to each source and is necessary for the analysis. The detected photons were then divided into 128 30-d time bins to ensure the source was significantly detected ($TS \geq 25$) in each light curve interval. The data corresponding to each time bin were subject to an unbinned analysis using the P8R2_SOURCE_V6 instrument response functions, accounting for the Galactic and isotropic background photons by incorporating the models `gll_iem_v06.fits` and `iso_P8R2_SOURCE_V6_v06.fits`, of which the normalization of the latter was left as a free parameter.




Figure 10. Monthly time-binned *Fermi* LAT light curve for the blazar PKS2155-304. The energy range is from 100 MeV to 300 GeV. The light curve shows the 128 30-d intervals between 2008 August 5 and 2019 February 9.

These models can be downloaded from the *Fermi* LAT data server.¹ The input model map for our analysis froze all parameters for *Fermi* 3FGL catalogue sources (Acero et al. 2015) $>10^\circ$ from our target of interest as the point-spread function (PSF) of the *Fermi* LAT at 100 MeV is $\approx 3.5^\circ$ and decreases at higher energies (Aharonian et al. 2013), so we do not expect photons from these sources to contaminate data from PKS2155-304. Sources within 10° of PKS2155-304 were modelled with their normalizations as free parameters, freezing other parameters to their 3FGL catalogue values (Acero et al. 2015) as we are not interested in spectral information and only wished to produce a time series. The best fit was found using the NEWMINUIT algorithm (James 1994), which returned flux values and uncertainties based on predicted photon counts associated with each source. The produced time series can be seen in Fig. 10.

8.1.2 Determining the false positive fraction

From the produced time series, both the PDF and PSD can be estimated. To create a PDF, we bin the light curve in accordance to pre-existing practices in the literature, whereby we account for the size of the error bars on the data points (e.g. Kushwaha et al. 2017) and ensure that each bin in the PDF is wider than the mean error bar of data points within that bin. Additionally, we impose the condition that each bin must contain ≥ 5 data points (see Hughes & Hase 2010) to prevent a small number of outliers from skewing the distribution that may erroneously lead to one PDF functional form being erroneously preferred. The resultant PDF is illustrated in Fig. 11, and it can be seen that its functional form is better described by a Gaussian than for a lognormal.

For our false positive calculation, we choose to input the measured PSD spectral index of the time series in question into our TK95 simulations. Measuring the PSD index for real data can be achieved by taking a DFT of the time series such that $\mathcal{P} \propto |\mathcal{F}(t)|^2/N$ (e.g. TK95). From this, the PSD spectral index can be estimated in different ways. Or equivalently there are different estimators. Two

¹<https://fermi.gsfc.nasa.gov/ssc/data/access/>




Figure 11. The PDF for PKS2155-304 corresponding to the time series shown in Fig. 10. The vertical error bars were calculated assuming the error is proportional to the number of counts in each bin. The histogram was generated using the condition that the mean error of data in each bin, indicated by the horizontal error bars, could not exceed the bin width. Additionally, each bin was required to have ≥ 5 data points. It can be seen that the histogram functional form is better described by a lognormal distribution as opposed to a Gaussian function.

estimates are made here. First is to compare the chi-square, χ^2_{dist} , of the mean simulated power to that of the observed power χ^2_{obs} , which are given by $\chi^2_{\text{dist},i} = \sum_{v=v_{\min}}^{v=v_{\max}} \frac{[P(v)_{\text{sim},i} - (P(v)_{\text{sim}})]^2}{[(\Delta P(v)_{\text{sim}})]^2}$ for the i -th simulation and $\chi^2_{\text{obs}} = \sum_{v=v_{\min}}^{v=v_{\max}} \frac{[P(v)_{\text{obs}} - (P(v)_{\text{sim}})]^2}{[(\Delta P(v)_{\text{sim}})]^2}$, respectively. The fraction of simulations where $\chi^2_{\text{obs}} < \chi^2_{\text{dist},i}$ is the fraction whose PSD estimate values are within the statistical spread of simulations. For PKS 2155-304, we find that for a range of indices from 0.2 to 2.8 with increments of 0.2, this fraction peaks at 1.2. In an alternate way, we can determine both the central value and uncertainty of the index. Here we calculate the mean power in each frequency bin over range of indices as shown in Fig. 12. Note that since the cadence of the observations as well as the mean and variance of the light curve is imposed on the simulations, this mean power map or ‘mean PSD map’ is not unbiased but carries information on the moments of the observed light curve. From this mean ‘PSD map’, we can compute the best-fitting value, which for PKS2155-304 is $\Gamma \sim 0.9 \pm 0.5$. This best fit is derived from the ‘most probable powers’ in each frequency bin, which are shown in red in Fig. 12. The green points show the periodogram directly computed from the light curve. Fig. 12 shows that the observed light curve is roughly consistent being a realization of a pink-noise process.

Previously, we have shown that TK95 simulations are appropriate for constraining $P(+|G)$, $P(-|G)$, $P(+|LN)$, and $P(-|LN)$ in equations (7)–(10). $P(+|G)$, $P(-|G)$, $P(+|LN)$, and $P(-|LN)$ can be determined in the same way as used to produce Fig. 8. Using the best fits to both the Gaussian and lognormal distributions fitted to the PDF, 10 000 artificial time series are generated for each. Using the results from our calculation of the PSD index, for each artificial time series we draw a value of Γ from a normal distribution characterized with a mean $\Gamma = 0.9$ and $\sigma = 0.5$. Fig. 7 shows that in this regime, time series that are on average normally distributed can be quickly produced via the TK95 method. Similarly, lognormal time series can be obtained by exponentiating a normally distributed




Figure 12. The PSD estimate of *Fermi* light curve of PKS 2155-304 using modified Timmer and Koenig simulations with exponential transform for lognormality is shown. The colour map represents the simulated power in each frequency bin for PSD indices from 0.2 to 2.8 in steps of 0.2. The colour scale represents the fraction of simulations within a power bin for a frequency range. The red represents the ‘most probable power’ in a given frequency bin for every index being equiprobable. The green points represent the periodogram of the observed light curve. The PSD is consistent with pink noise.

time series. One caveat here is that exponentiating the time series can change the time series PSD. Uttley et al. (2005) have investigated this effect (see Appendix B) and find that although slightly more power is present in high-frequency components, the overall effect is small for PSDs lacking sharp features, such as the power-law PSDs used in this example. Alternatively, the EMP13 algorithm may be used; it is more general with user-defined PDF and PSD, which are refined consistently with the data. However, it is naturally more complex and computationally expensive. In this case, we use TK95 as we are in the $\Gamma \gtrsim 1$ regime, so the flux distribution for each artificially generated time series will on average show a high degree of consistency with the true PDF. Each simulated time series has its PDF evaluated in exactly the same way as the real time series, using the same data binning as we are assuming error bars on each of the simulated time-series data points to be correlated with that of the real data (although we do not explicitly calculate them). This allows $P(+|G)$, $P(-|G)$, $P(+|LN)$, and $P(-|LN)$ to be determined.

It remains to determine the prior probability that the PDF we are measuring is Gaussian ($P(G)$) or lognormal [$P(LN)$]. From our assumption that the PDF can only take one of these functional forms, $P(LN) = 1 - P(G)$. Fig. 13 shows the posterior distribution as a function of $P(LN)$, demonstrating the effect that varying $P(LN)$ has on the posterior distributions for a true positive, indicated by the red solid line and defined by equation (9), and for a false positive, which is shown by the blue dashed line and defined in equation (10). In this figure, we highlight two possible choices of prior.

The first of these utilizes the principal of indifference, in which we assume the simplest possible non-informative (or flat) prior where $P(G) = P(LN) = 0.5$. This is marked on Fig. 13 using the black dotted line, and corresponds to values of $P(G|LN_{\text{PDF}}) = 0.163$ and $P(LN|LN_{\text{PDF}}) = 0.837$, indicating a preference for lognormality somewhere between 1σ and 2σ .




Figure 13. Figure showing the variation of the posterior distributions for the PDF of γ -ray time series of PKS2155-304 being lognormally distributed as a function of the prior probability that the source has a lognormal PDF, $P(\text{LN})$. It was assumed that lognormal and normally distributed PDFs were the only options, such that the probability of a Gaussian PDF is $P(\text{G}) = 1 - P(\text{LN})$. It can be seen for our example of PKS2155-304, an intrinsically Gaussian PDF is only preferred if there is a high $[P(\text{G}) \gtrsim 0.82]$ prior probability that the PDF is Gaussian. The PDF and PSD for PKS2155-304 are shown in Figs 11 and 12, respectively. The dotted and double-dashed lines refer to choices of prior, which are discussed in the text.

Secondly, we compute a second example using priors weighted on the probability that our histogram is representative of the true PSD. To do this, we use the Bayesian inference to calculate the probability of obtaining each datum, D_i , assuming the model with parameters, θ_i , is correct. For the entire histogram PDF shown in Fig. 11, this can be expressed as

$$P(D|\theta_i) \propto \prod_i \exp \left[-\frac{(D_i - m_i(\theta_i))^2}{2\sigma_i^2} \right], \quad (11)$$

where $m_i(\theta_i)$ are the model values and σ_i are the size of the uncertainties. It is assumed that measurements of each data point are normally distributed about the mean model value, and N is the number of bins in the PDF. The total likelihood is the product of the likelihoods of each bin value being correct. We apply equation (11) to the normal and lognormal models for the PDF of PKS2155-304 and re-normalize under our assumption that the true PDF must be either Gaussian or lognormal. For PKS2155-304, we obtain $P(\text{G}) = 0.000587$ and $P(\text{LN}) = 0.999413$, where the low-probability terms in $P(\text{G})$ come from the bins at either end of the histogram, which are clearly a poor fit to the Gaussian model in Fig. 11. The result of these much more extreme values on our final results is that the false positive fraction becomes $P(\text{G}|\text{LN}_{\text{PDF}}) = 0.000115$ and the true positive $P(\text{LN}|\text{LN}_{\text{PDF}}) = 0.999885$, with the latter of these $> 3\sigma$. It is clear that this prior leads to a much stronger preference for lognormality, but we stress this is largely because for this prior the probability of obtaining the histogram bins on the tail (the highest flux bins) of the calculated PDF are very small assuming a Gaussian model, and we have not tested other functional forms.

Fig. 13 is a key and informative result as it demonstrates the difficulty associated with being able to conclusively say that an individual source has a particular intrinsic PDF distribution, irrespective of the prior chosen. In our example of PKS2155-304, we have shown that a simple histogram PDF exhibits a preference for a lognormal functional form, yet we cannot say we are confident that is a correct measurement of the true intrinsic PDF to much more than

3σ . Even so, Fig. 13 shows an overall preference of PKS2155-304 to have a lognormally distributed PDF as we are only likely to obtain a Gaussian PDF measurement if our lognormal prior has a value $P(\text{LN}) \lesssim 0.18$. The point where $P(\text{G}|\text{LN}_{\text{PDF}}) = P(\text{LN}|\text{LN}_{\text{PDF}})$ is therefore an important diagnostic for determining a false positive rate for an individual time series. We recommend undertaking artificial light-curve simulations using the method outlined above to determine how likely obtaining the inferred result is.

In summary, our false positive prescription may be outlined as follows:

- (i) Use the observed time series to compute a PSD and PDF, fitting the desired functional forms to each and evaluating them with a statistical test such as a reduced χ^2 .
- (ii) Determine the priors $P(\text{G})$ and $P(\text{LN})$. In our example, we have computed results for a flat prior and a histogram-weighted prior.
- (iii) Produce N artificial time series from each distribution (Gaussian and lognormal) used to best fit the data PDF, including the measured PSD spectral index. We recommend using TK95 simulations to produce normally and lognormally distributed time series (normally distributed time series can be exponentiated to create lognormal time series), although this approach is only able to reliably produce time series with the correct PDF functional form when the PSD index is $\Gamma \lesssim 1$.
- (iv) As the true PDF functional form is known for the artificial time series, this will give $P(+|\text{G})$, $P(-|\text{G})$, $P(+|\text{LN})$ and $P(-|\text{LN})$, where $+$ and $-$ symbolize correct or incorrect measurement of inherently Gaussian (G) or lognormal (LN) PDFs.
- (v) Depending on whether a Gaussian or lognormal PDF is measured from the time series, evaluate either equations (7) and (8) or equations (9) and (10) as appropriate to obtain positive and false positive fractions.

9 SUMMARY AND CONCLUSIONS

We have presented time-series simulations based on the method of TK95, and investigated the relationship between the input PSD and properties of the output PDF. We initially began by assuming a simple power-law PSD and investigated how this affected the average p -value returned from an SW statistical test, which tests for normality of the time series. We have used these tests to create a false positive prescription rate for evaluating the likelihood of the correct measurement of the PDF functional form for astrophysical sources. Our results may be summarized as follows:

- (i) PDFs of Timmer & Koenig (TK95) time-series simulations are more likely to deviate from having normally distributed fluxes as the PSD index steepens. This is analytically expected because there is progressively less power in the higher frequency components relative to the lower frequencies. Beyond PSD index $\Gamma \approx 1$, artificial time series corresponding to a power-law PSD are increasingly dominated by contributions from these lower frequencies, many of which have wavelengths longer than the desired observation time span.
- (ii) The mean p -value from an SW test on TK95 simulated time series is roughly constant for PSD indices $\lesssim 1$, but sharply begins to decrease above this index, eventually reaching ~ 0 , implying artificial power-law time series with $\Gamma > 1$ on average reject the null hypothesis that the data are normally distributed.
- (iii) The characteristic fiducial fit to this curve is given by a sigmoid function, specifically

$f(\Gamma) = \alpha \exp(\beta - \eta\Gamma) / [1 + \exp(\beta - \eta\Gamma)]$, where Γ is the PSD spectral index and α , β , and η are free parameters.

(iv) In general, it is difficult to distinguish between lognormal and Gaussian PDF functional forms for time series data, even more so for relatively small ($N \lesssim 100$) data sets. We therefore recommend that the calculation of these for individual sources is accompanied by a false positive rate, the parametrizations of which are deduced from generating large numbers of artificial time series with known PDF distributions.

(v) As an example, we show that the *Fermi* LAT γ -ray light curve of the blazar PKS2155-304 shows evidence of lognormality to 83.66 per cent when using a flat prior (assuming that an intrinsic Gaussian or lognormal PDF or equally likely). This increases to 99.99 per cent when weighting the priors using a Bayesian inference to evaluate the probability that the data points are correct assuming each model is.

ACKNOWLEDGEMENTS

PJM acknowledges support from a Hintze Scholarship. NC acknowledges kind support by the AvH foundation and MPIK during development of this project. GC acknowledges support from STFC grants ST/N000919/1 and ST/M00757X/1 and from Exeter College, Oxford. This work was supported by the Oxford Hintze Centre for Astrophysical Surveys, which is funded through generous support from the Hintze Family Charitable Foundation.

REFERENCES

- Abdo A. A. et al., 2010, *ApJ*, 722, 520
 Acero F. et al., 2015, *ApJS*, 218, 23
 Aharonian F., Bergström L., Dermer C., Walter R., Türler M., 2013, *Astrophysics at Very High Energies: Saas-Fee Advanced Course 40*. Swiss Society for Astrophysics and Astronomy. Springer, Berlin Heidelberg
 Alston W. N., 2019, *MNRAS*, 485, 260
 Alston W. N. et al., 2019, *MNRAS*, 482, 2088
 Anderson T. W., Darling D. A., 1952, *Ann. Math. Stat.*, 23, 193
 Arévalo P., Uttley P., 2006, *MNRAS*, 367, 801
 Bartlett M. S., 1948, *Nature*, 161, 686
 Belloni T., Homan J., Casella P., van der Klis M., Nespoli E., Lewin W. H. G., Miller J. M., Méndez M., 2005, *A&A*, 440, 207
 Breedt E. et al., 2009, *MNRAS*, 394, 427
 Chevalier J., Kastendieck M. A., Rieger F. M., Maurin G., Lenain J.-P., Lamanna G., 2015, *Proc. Sci., Long Term Variability of the Blazar Pks 2155-304*. SISSA, Trieste, PoS#829
 Emmanoulopoulos D., McHardy I. M., Papadakis I. E., 2013, *MNRAS*, 433, 907
 Giebels B., Degrange B., 2009, *A&A*, 503, 797
 Gleissner T., Wilms J., Pottschmidt K., Uttley P., Nowak M. A., Staubert R., 2004, *A&A*, 414, 1091
 Guo Y. C., Hu S. M., Li Y. T., Chen X., 2016, *MNRAS*, 460, 1790
 Harris F. J., 1978, *Proc. IEEE*, 66, 51
 Heil L. M., Vaughan S., Uttley P., 2012, *MNRAS*, 422, 2620
 Heyde C. C., 2010, *On a Property of the Lognormal Distribution*. Springer, New York, NY, p. 16
 Hughes I., Hase T. P. A., 2010, *Measurements and their Uncertainties*. Oxford Univ. Press, Oxford
 Ishibashi W., Courvoisier T. J. L., 2012, *A&A*, 540, L2
 James F., 1994, *Comput. Phys. Commun.*, 79, 111
 King A. R., Pringle J. E., West R. G., Livio M., 2004, *MNRAS*, 348, 111
 Kirchner J. W., 2005, *Phys. Rev. E*, 71, 066110
 Kotov O., Churazov E., Gilfanov M., 2001, *MNRAS*, 327, 799
 Kushwaha P., Sinha A., Misra R., Singh K. P., de Gouveia Dal Pino E. M., 2017, *ApJ*, 849, 138
 Lyubarskii Y. E., 1997, *MNRAS*, 292, 679
 Maccarone T. J., Coppi P. S., 2002, *MNRAS*, 336, 817
 McHardy I., 2008, *Blazar Variability across the Electromagnetic Spectrum*. p. 14
 Marshall K., 2015, *ApJ*, 810, 52
 Mohd Razali N., Yap B., 2011, *J. Stat. Model. Analytics*, 2, 21
 Murdoch D. J., Tsai Y.-L., Adcock J., 2008, *Am. Stat.*, 62, 242
 Nakagawa K., Mori M., 2013, *ApJ*, 773, 177
 Pottschmidt K. et al., 2003, *A&A*, 407, 1039
 Rahman M. M., Govindarajulu Z., 1997, *J. Appl. Stat.*, 24, 219
 Rice J., 2006, *Mathematical Statistics and Data Analysis*. Thomson Brooks/Cole, Belmont, CA, USA
 Romoli C., Chakraborty N., Dorner D., Taylor A., Blank M., 2018, *Galaxies*, 6, 135
 Schreiber T., Schmitz A., 1996, *Phys. Rev. Lett.*, 77, 635
 Shah Z., Mankuzhiyil N., Sinha A., Misra R., Sahayanathan S., Iqbal N., 2018, *Res. Astron. Astrophys.*, 18, 141
 Shakura N. I., Sunyaev R. A., 1973, *A&A*, 500, 33
 Shapiro S. S., Wilk M. B., 1965, *Biometrika*, 52, 591
 Sinha A. et al., 2016, *A&A*, 591, A83
 Sinha A., Khatoon R., Misra R., Sahayanathan S., Mandal S., Gogoi R., Bhatt N., 2018, *MNRAS*, 480, L116
 Sivia D., Skilling J., 2006, *Data Analysis: A Bayesian Tutorial*. Oxford Sci. Publ., OUP Oxford
 Smith K. L., Mushotzky R. F., Boyd P. T., Malkan M., Howell S. B., Gelino D. M., 2018, *ApJ*, 857, 141
 Sobolewska M. A., Siemiginowska A., Kelly B. C., Nalewajko K., 2014, *ApJ*, 786, 143
 Timmer J., Koenig M., 1995, *A&A*, 300, 707
 Urry C. M., Padovani P., 1995, *PASP*, 107, 803
 Uttley P., McHardy I. M., Papadakis I. E., 2002, *MNRAS*, 332, 231
 Uttley P., McHardy I. M., Vaughan S., 2005, *MNRAS*, 359, 345
 Vaughan S., Edelson R., Warwick R. S., Uttley P., 2003, *MNRAS*, 345, 1271
 Wegman E. J., 1972, *Technometrics*, 14, 533
 Wijnands R., van der Klis M., 1999, *ApJ*, 514, 939
 Yap B., Sim C., 2011, *J. Stat. Comput. Simul.*, 81, 2141

This paper has been typeset from a \LaTeX file prepared by the author.