

Mind the gaps: improved methods for the detection of periodicities in unevenly sampled data

Andrés Gúrpide[★] and Matthew Middleton

School of Physics & Astronomy, University of Southampton, Southampton SO17 1BJ, UK

Accepted 2025 January 24. Received 2025 January 4; in original form 2024 October 15

ABSTRACT

The detection of periodic signals in irregularly sampled time series is a problem commonly encountered in astronomy. Traditional tools used for periodic searches, such as the periodogram, have poorly defined statistical properties under irregular sampling, which complicate inferring the underlying aperiodic variability used for hypothesis testing. The problem is exacerbated in the presence of stochastic variability, which can be easily mistaken for genuine periodic behaviour, particularly in the case of poorly sampled light curves. Here, we present a method based on Gaussian Processes (GPs) modelling for period searches and characterization, specifically developed to overcome these problems. We argue that in cases of irregularly sampled time series, GPs offer an appealing alternative to traditional periodograms because the known distribution of the data (correlated Gaussian) allows a well-defined likelihood to be constructed. We exploit this property and draw from existing statistical methods to perform traditional likelihood ratio tests for an additional (quasi-)periodic component, using the aperiodic variability inferred from the data as the null hypothesis. Inferring the noise from the data allows the method to be fully generalizable, with the only condition that the data can be described as a Gaussian process. We demonstrate the method by applying it to a variety of objects showing varying levels of noise and data quality. Limitations of the method are discussed, and a package implementing the proposed methodology is made publicly available.

Key words: accretion, accretion discs – methods: data analysis – methods: statistical – stars: black holes – stars: neutron – galaxies: active.

1 INTRODUCTION

The identification of periodic/quasi-periodic signals in irregularly sampled time series is a common problem in astronomy due to the predominance of interrupted or intermittent observing. While uninterrupted, regularly sampled observations may be achieved for periods of up to around a day at most, for time-scales extending to the tens or hundreds of days, this becomes impractical, particularly for faint sources, which require the most sensitive of instruments. Such ‘long’ time-scales are, however, of great interest in the study of many phenomena, such as superorbital periods in X-ray binaries (XRBs; e.g. Kotze & Charles 2012; Vasilopoulos et al. 2020) and binary supermassive black hole (SMBH) signals (e.g. Graham et al. 2015).

Arguably the most widely used technique to search for periodicities in time series is the periodogram, which involves calculating the modulus-squared of the discrete Fourier transform. Lomb (1976) and Scargle (1982) extended the periodogram to the case of irregularly sampled time series, a technique known today as the Lomb–Scargle periodogram (see VanderPlas 2018, for a review of this technique). Periodic signals appear as peaks or ‘outliers’ in power, allowing the frequency of the repeating signal to be estimated. If the distribution of

the powers in the absence of a signal is known, the chance probability of generating such an outlier can then be calculated, providing an estimate of the significance of the candidate period.

In the absence of source variability other than the repeating signal, i.e. when the sole source of additional variance in the light curve is due to Poisson noise, the problem is somewhat straightforward. In regularly sampled time series, the problem can be tackled analytically as the powers in the periodogram are independent and follow a χ^2 distribution with 2 degrees of freedom (χ^2_2 ; e.g. van der Klis 1988). In the case of irregularly sampled data, the powers are no longer independent, but the problem can be tackled easily by randomizing the time series (Frescura, Engelbrecht & Frank 2008; VanderPlas 2018).

Searching for periods is made considerably harder when systems show *intrinsic* aperiodic or stochastic (i.e. non deterministic) variability, as is universally observed in both accreting systems (Vaughan et al. 2003) and stars (Bowman & Dorn-Wallenstein 2022). These types of source have steep power spectral densities (PSDs), commonly referred to as ‘red noise’. Failing to account for this background noise tends to overestimate the significance of peaks in the periodogram (Vaughan 2005). For this reason, sources showing stochastic variability are more prone to misidentified periods, exacerbated in the case of uneven sampling (e.g. Vaughan et al. 2016). While this problem was tackled by Israel & Stella (1996) and Vaughan (2005, 2010) in the case of regularly sampled time series,

[★] E-mail: a.gurpide-lasheras@soton.ac.uk

there is as yet no standard procedure for the case of irregularly sampled time series.

In this paper, we present a recipe aimed at detecting periodicities in irregularly sampled time series, with particular focus on cases where the systems under study show additional aperiodic variability (as it is the case in e.g. active galactic nuclei (AGNs); González-Martín & Vaughan 2012), although the method is completely generalizable. We argue that in such cases, Gaussian Process (GP) modelling offers a clear advantage over traditional (Lomb–Scargle) periodograms, because the likelihood of the data is known. This allows us to constrain the underlying, aperiodic variability using GP modelling, and use well-established statistical techniques to determine the candidate period significance and a goodness-of-fit.

This paper is structured as follows: in Section 2, we review the case of period-detection in regularly sampled time series and describe our proposed methodology for the irregularly sampled case. In Section 3, we demonstrate the method, applying it to both simulated and real data. Finally in Section 4, we discuss the advantages of the methodology over more traditional Fourier-based techniques and outline certain limitations and caveats of the proposed methodology.

2 SEARCHING FOR A PERIOD

2.1 Regularly sampled time series

The standard methodology to test for the presence of a narrow peak associated with a periodic/quasi-periodic signal in a periodogram generally involves estimating the broadband noise (or continuum) and using that estimate – and its uncertainties – as the null hypothesis (e.g. Israel & Stella 1996; Vaughan 2005, 2010, see also Gierliński et al. 2008; Pasham et al. 2019; Ashton & Middleton 2021 for an example of the application of such a methodology). In the case of regularly sampled time series, the powers in the periodogram can be considered independent and their distribution is well known (scattered as a χ^2_2 around the underlying PSD; e.g. Klis 1988). Knowing the distribution of powers allows the construction of a well defined likelihood function (e.g. Stella et al. 1994) (commonly known as the Whittle likelihood; see Vaughan 2010, and references therein). This allows forward-fitting of the periodogram¹ and models to be rejected based on the data alone, as it is customarily done when fitting using χ^2 statistics.

Knowledge of the likelihood not only allows models to be rejected based on the data alone, but also to test for the presence of additional components (e.g. quasi-period oscillations; QPOs) by performing a likelihood ratio test (LRT; Protassov et al. 2002). In particular, Vaughan (2010) proposed to follow Protassov et al. (2002) and perform an LRT using the Whittle likelihood function from fits to the periodogram:

$$T_{\text{LRT}} = -2 \ln \frac{L_0}{L_1}. \quad (1)$$

Here, L_0 and L_1 are the maximum of the likelihood functions for the null hypothesis and the alternative model, respectively. Subsequently, one would simulate periodograms drawn from the null hypothesis (and its uncertainties, see e.g. Ashton & Middleton 2021) and derive the same quantity for each of the simulated periodograms. As stated earlier, as the distribution of powers is known in the case of evenly sampled data (scattered as χ^2_2 around the PSD) one can avoid the additional step of simulating light curves (so long as aliasing and

red noise leakage effects are not important). Finally, a comparison of the $observed T_{\text{LRT}}$ against the reference distribution derived from the simulated data sets allows the probability of rejecting the null-hypothesis model to be assessed (through the derived p -value), thereby providing an estimate of the significance of the putative signal. Note that, because the null hypothesis is derived from the data, the method makes no assumptions about the underlying noise, and is completely generalizable. We seek to replicate this process in the case of unevenly sampled data.

2.2 Irregularly sampled time series

In the case of irregular sampling, there is no straightforward way to model the broad-band noise as for the regularly sampled case. If we were to replicate the standard procedure outlined above using the Lomb–Scargle periodogram, we would encounter a variety of problems. First, the powers in the Lomb–Scargle periodogram are known to *not* be statistically independent (e.g. Lomb 1976) and their distribution is therefore unknown and dependent on the underlying (also unknown) PSD. Secondly, the irregular sampling implies that there is no well-defined set of frequencies over which to evaluate the periodogram (e.g. Frescura et al. 2008). Finally, as the Nyquist frequency is ill-defined or nonexistent (VanderPlas 2018), aliasing effects are exacerbated. The combination of these problems typically precludes forward-fitting of the (Lomb–Scargle) periodogram or at the very least, forward-fitting will lead to biased estimates.

To illustrate this, we simulated 1000 light curves using the method proposed by Timmer & Koenig (1995), initially with $N = 1,000$ evenly sampled datapoints using an input PSD where the power ($S(f)$) follows a power law $S(f) \propto f^{-\beta}$ with $\beta = 1$, and 1.8, respectively. The light curves were initially simulated to be ten times longer to introduce red-noise leakage effects and then truncated into the aforementioned length. We then randomly removed 50 datapoints from each light curve and computed Lomb–Scargle periodograms from the resulting light curves. We then fit the periodograms in log space with a linear function (i.e. assuming the powers follow a χ^2_2 ; Vaughan 2005) and retrieve the best-fitting slope (β) in each case. We then progressively removed a further 50 datapoints, until 500 datapoints had been removed (but always keeping the last and the first datapoint to maintain the same light curve length), recording the mean best-fitting β for the ensemble of the 1000 light curves. The mean best-fitting β as a function of number of datapoints removed is shown in Fig. 1 for both β values.

Fig. 1 shows that, for $\beta = 1$, when $\gtrsim 200$ datapoints have been removed, the best-fitting β is underestimated by ~ 20 per cent, illustrating the inherent issues in fitting to the Lomb–Scargle periodogram (see also O’Sullivan & Aigrain 2024). The case of $\beta = 1.8$ shows the bias is more dramatic for steeper PSDs. The situation becomes even worse in a real-case scenario, where the frequency grid will be unknown (here we can at least assume the frequency grid given by the initially evenly sampled light curves) and there will be no way to know whether the fit is an adequate description of the data. Note also that the biases will affect any parameter fitted, including the normalization of the power law, which we have not shown here.

One way to circumvent the above problems is to rely on Monte Carlo simulations of light curves, as pioneered by Done et al. (1992) and later refined by Uttley, McHardy & Papadakis (2002). This approach attempts to find the power spectral model that, when convolved with the observing window, best matches the (Lomb–Scargle) periodogram of the real data, so that all distorting effects are taken into account. However, such methods still run into some problems, particularly when dealing with irregularly sampled time

¹Notably if aliasing and red noise leakage effects are negligible.




Figure 1. Mean best-fitting β for an ensemble of 1000 Lomb–Scargle periodograms of light curves generated with a PSD following a power law $S(f) \sim f^{-\beta}$ with $\beta = 1$ (blue solid line) and 1.8 (orange solid line). The periodograms were fitted with a linear function in log–log space (i.e. assuming the powers follow a χ^2_2 as for the regularly sampled case) as we progressively removed datapoints. The best-fitting β quickly deviates from the input β value (dashed horizontal lines), showing how the Lomb–Scargle periodogram becomes a biased estimator as the sampling regularity decreases.

series. Once again, the unknown distribution of the powers and their lack of independence implies that the choice of the fit-statistic will not be straightforward. One could aim to rebin the periodogram, hoping that enough samples will converge to Gaussianity and independence, but again the lack of a well-defined frequency-grid complicates this approach (whilst the number of averages needed to reach Gaussianity is unclear and depends on the underlying, unknown PSD; Ingram & Done 2011). In addition, if binning cannot be avoided, any periodic component and its structure due to the observing window will be smeared out, which will affect the estimate of the continuum. For that reason, tests relying on simulated Lomb–Scargle periodogram peaks often have to excise the frequency of the candidate period in order to determine the putative underlying noise (e.g. Pasham et al. 2024), thereby making a priori assumption about the presence of any QPO. This is because if the feature is real and not removed, the broadband continuum estimate used for the simulations will be biased, often towards steeper indexes if the QPO is at the low-frequency end. This overestimate of the amount of aperiodic variability will therefore underestimate the significance of the periodic component.

The method of Monte Carlo simulations also becomes quickly computationally expensive as it relies heavily on Monte Carlo simulations for estimation of the best-fitting parameters, and additional simulations are often needed to obtain parameter uncertainties (e.g. Mueller & Madejski 2009; Markowitz 2010). An appealing aspect of such a method, however, is that one can obtain the goodness-of-fit through the simulated light curves, using them to derive the empirical distribution of fit-statistic from which the goodness-of-fit (or ‘rejection confidence’ as per Uttley et al. 2002) can be derived. We note that, as pointed out by Mueller & Madejski (2009), one should re-fit the simulated light curves in the same manner as for the observed data set in order to obtain the empirical distribution of the fit statistic, which would again dramatically increase the computational time. We return to this point in Section 2.6.

An alternative approach to the above is to use time-domain fitting methods such as GP, where the irregular sampling and measurement (heteroskedastic) errors are fully accounted for and

are less susceptible to the distorting effects inherent in a Fourier-domain approach (e.g. Kelly, Sobolewska & Siemiginowska 2011). The covariance functions, or ‘kernels’, when the data is stationary (when they depend only on the $\Delta t_{ij} = |t_i - t_j|$ interval between any two datapoints), describe the autocorrelation function, which can be Fourier-transformed to obtain the PSD (Rasmussen & Williams 2006). Therefore GP modelling offers an equally flexible but frequency-distortion-free access to the PSD, while maximizing data usage by making full account of the measurement uncertainties and avoiding binning. Moreover, complications arising from the unknown distribution of powers in the case of the Lomb–Scargle periodogram are avoided.

Beyond the computational demand, which generally scales as N^3 (although Foreman-Mackey et al. (2017) showed that the computational time can be reduced to scale as $\sim J^2 N$ – where J is the number of model components – for a restricted set of kernels), a more general drawback of GP modelling compared to traditional methods is that there is no measure of goodness-of-fit. As a result, models cannot be rejected solely on the basis of the data, and only model comparison (e.g. using an information criterion) is possible. In addition, although QPO-searches have been performed using GP (e.g. Covino et al. 2020; Covino, Tobar & Treves 2022; Hübner et al. 2022; Zhang, Yang & Dai 2023), establishing the significance of such signals remains challenging. In particular, it is important to quantify the chance probability of generating a fit improvement (or any other metric such as the Bayes factor) when including a QPO/periodic component (hereafter we will refer to this simply as the ‘signal’) given the specific sampling, priors, fitting technique, and other factors involved in the analysis. This is particularly important if such methods are to be extended to include non-stationary kernels, where the time window becomes a parameter of the model (Hübner et al. 2022). In such cases, one needs to account for the additional sets of free-trials or model-flexibility introduced in allowing signals to be transient.

2.3 The method

Our procedure can be considered equivalent to the LRT approach proposed by Vaughan (2010), but adapted to deal with irregularly sampled data. First, to circumvent the issues related to use of the Lomb–Scargle periodogram, we obtain the likelihood directly from the GP modelling in the time domain. We then make a comparison between a continuum-only model (the null hypothesis) to a more complex model that includes the signal, obtaining a fit improvement (quantified through $L_1 - L_0 = \Delta L$ or T_{LRT}). Next, from the posteriors of the null hypothesis modelling, we draw kernel parameter samples and then use the PSD of these kernels to generate a number of simulated light curves via inverse-Fourier transform (the full methodology employed to simulate the light curves is described in Appendix A). We finally perform the same GP modelling on the synthetic light curves to derive the reference distribution for the LRT. While throughout this work we employ uninformative (uniform) priors, Bayesian priors could easily be incorporated, provided the same priors are also used when fitting the simulated data sets.

Intuitively, this method can be understood as follows: if the improvement in fit statistic provided by the added model component (the putative signal) is due to random noise fluctuations in the original data set (i.e. the signal is spurious), the fit improvement obtained in the simulated data sets (which were simulated using the model *without* the additional model component) will be of the same order as that of the real data set. If the signal is real, then the improvement




Figure 2. PSDs of the *celerite* models used in this work. All PSDs are shown with the same integrated variance. The DRW results in a bending power law in Fourier space (dashed orange line), whereas the exponentially decaying sinusoid gives a Lorentzian, which is shown for different values of coherence, Q , in blue solid lines. For $Q \lesssim 3/2$, the Lorentzian becomes broad, mimicking a bending power law (see Belloni, Psaltis & Klis 2002). The PSD of the SHO for the special case of $Q = 1/\sqrt{2}$ and $Q = 1/2$ (which yields the Matérn-3/2 kernel approximation) are shown as a dashed-dotted green line and a solid purple line, respectively. The vertical dashed line indicates the central frequency of the Lorentzian (ω_0) and ω_{bend} for the DRW, Matérn-3/2 and $\text{SHO}_{Q=1/\sqrt{2}}$ kernels. Note that for the Matérn-3/2 $\rho = \sqrt{3}/\omega_{\text{bend}}$.

in fit statistic provided by the additional model component will be generally larger than any of the values obtained in the simulated data sets. Such a procedure not only inherently accounts for the number of free trials – as long as the parameter ranges/priors are kept the same as for the original data set – but also for the fact that some signals may have a more complex profile than a simple peak in a Lomb–Scargle periodogram (as is often assumed in significance testing). Another advantage of relying on fit improvements is that it removes the need to make any prior assumptions about the presence of the QPO in the periodogram, which, as stated in Section 2.2, is often the case when relying on periodogram peaks.





Figure 3. Example of a simulated light curve, generated to test the sensitivity of our method to false negatives ($N = 250\Delta t \approx 4$ days). (Left) Light curve generated using a Lorentzian + DRW, with bending time-scale of 60 d, and period of 100 d. (Right) Corresponding Lomb–Scargle periodogram.

2.4 Kernel functions

There are naturally a range of possible models one could potentially use to describe the underlying noise and the signal, and we refer to Rasmussen & Williams (2006) for some examples. In this paper, we use the *celerite* kernels proposed by Foreman-Mackey et al. (2017) for GP modelling – which we note bear many similarities to CARMA (Kelly et al. 2014) – to reduce the computational burden, but note our method is generalizable to any choice of kernels. These kernels can then be combined through additions or multiplications to achieve more complex covariance matrices. However, as shown by Foreman-Mackey et al. (2017), any multiplication of *celerite* kernels can always be reformulated as an addition under a new parameter set. Therefore we only explore additions of the kernels described below.

The simplest choice of *celerite* kernel for modelling aperiodic variability is the Damped Random Walk (DRW), whose kernel is simply a decaying exponential:

$$k(\Delta t_{ij}) = \sigma^2 \exp(-\omega_{\text{bend}} \Delta t_{ij}) \quad (2)$$

the PSD of which is a bending power law (Fig. 2):

$$S(\omega) = \sqrt{\frac{2}{\pi}} \frac{\sigma^2}{\omega_{\text{bend}}} \frac{1}{1 + \left(\frac{\omega}{\omega_{\text{bend}}}\right)^2} \quad (3)$$

with an index of -2 for $\omega \gg \omega_{\text{bend}}$ bending smoothly to a flat ($S(\omega) \sim \omega^0$) power law around ω_{bend} . In the above, σ^2 is the variance of the process.

A further possible kernel (as proposed by Foreman-Mackey et al. 2017) is the stochastically driven damped harmonic oscillator (SHO), which can model both aperiodic and periodic variability. For the full details of this kernel we refer the reader to Foreman-Mackey et al. (2017); in this work we consider two special cases of this kernel used to model aperiodic noise. The first one is commonly used to model (aperiodic) granular noise in stars:

$$k(\Delta t_{ij}) = S_N \omega_{\text{bend}} e^{-\frac{1}{\sqrt{2}} \omega_{\text{bend}} \Delta t_{ij}} \cos\left(\frac{\omega_{\text{bend}} \Delta t_{ij}}{\sqrt{2}} - \frac{\pi}{4}\right) \quad (4)$$

with a PSD of the form:

$$S(\omega) = \sqrt{\frac{2}{\pi}} \frac{S_N}{(1 + \omega/\omega_{\text{bend}})^4} \quad (5)$$




Figure 4. Distribution of p -values (with colours indicating the density of values) obtained from the application of our PPP method to 100 simulated light curves using a Lorentzian (QPO) + DRW with varying cadence (left) and baseline (right). The mean p -values for each cadence/baseline strategy are shown as short, horizontal lines. The horizontal grey dashed line shows the 99 per cent significance detection level ($p = 0.01$).

here S_N scales the variance of the noise process ($\sigma^2 = \frac{1}{\sqrt{2}} S_N \omega_{\text{bend}}$). The PSD of this kernel is similar to the DRW but here the power law has a steeper index of -4 for $\omega \gg \omega_{\text{bend}}$ (Fig. 2). Hereafter, we refer to this model as $\text{SHO}_{Q=1/\sqrt{2}}$ as this kernel is obtained for the special case of $Q = 1/\sqrt{2}$ within the more general SHO (for more details we refer the reader to Foreman-Mackey et al. 2017).

The second special case of the SHO we consider is an approximation to the Matérn-3/2 kernel², which using `celerite` kernels can be approximated setting $Q = 1/2$ in the SHO:

$$k(\Delta t_{ij}) = \sigma^2 \left(1 + \frac{\sqrt{3}}{\rho} \Delta t_{ij} \right) e^{-\frac{\sqrt{3}\Delta t_{ij}}{\rho}}, \quad (6)$$

where ρ sets the characteristic time-scale in a similar fashion to the DRW. The PSD of this function is only slightly dissimilar to the $\text{SHO}_{Q=1/\sqrt{2}}$ kernel as can be seen in Fig. 2. Hereafter, we refer to this kernel as Matérn-3/2 for simplicity.

Finally, we also considered a ‘Jitter’ or white-noise kernel to model uncorrelated aperiodic variability, parameterized only by its variance:

$$k(\Delta t_{ij}) = \sigma^2 \delta_{ij}, \quad (7)$$

where δ_{ij} is the Kronecker delta, indicating this term simply adds a diagonal term to the covariance matrix. This kernel can be interpreted in two ways. The first is that the uncertainties on the data are underestimated; in this case σ^2 provides the constant, missing contribution to the noise; the second is as an extra white noise term to capture some random variations (e.g. instrumental effects) not captured by the main model. Here, we consider it as an independent model to describe cases where the data does not support the use of a different kernel (signalling that white noise as the null hypothesis might be justified).

²In practice we have found the parameter controlling the approximation in `celerite` to have very small effect on the results and was fixed to the arbitrary small value of 10^{-7} . While preparing this manuscript we have learned that an exact state-representation of the Matérn-3/2 has now been derived in Jordán, Eyheramendy & Buchner (2021).

Table 1. Summary of the analysis carried out to test the sensitivity of our method to false negatives. 100 light curves were simulated per baseline/cadence combination from a Lorentzian + DRW PSD and we examined whether we could detect the additional Lorentzian (QPO) component over the DRW.

N^a	$\langle \Delta t \rangle^b$ d	T^c d	$\langle p \rangle^d$	$n_{p<0.01}^e$
100	10	1000	0.12	65
250	4	1000	0.12	62
500	2	1000	0.11	68
1000	1	1000	0.10	68
300	1	300	0.24	21
400	1	400	0.21	39
500	1	500	0.17	44

Notes.

^a Number of datapoints of the generated light curves.

^b Mean cadence.

^c Observing baseline.

^d Mean retrieved PPP value of 100 light curves for the presence of the QPO (Lorentzian) component.

^e Number of light curves for which the Lorentzian was significantly detected ($p < 0.01$).

For the periodic component, we have employed only a single exponentially decaying sinusoid:

$$\begin{aligned} k(\Delta t_{ij}) &= \sigma^2 \exp(-b\Delta t_{ij}) \cos(\omega_0 \Delta t_{ij}) \\ &= \sigma^2 \exp\left(-\frac{\omega_0}{2Q} \Delta t_{ij}\right) \cos(\omega_0 \Delta t_{ij}), \end{aligned} \quad (8)$$

where $b = \omega_0/2Q$ following the nomenclature of Foreman-Mackey et al. (2017). The resulting PSD takes the form of a Lorentzian (Fig. 2):

$$S(\omega) = \frac{1}{\sqrt{2\pi}} \frac{\sigma^2 b}{b^2 + (\omega - \omega_0)^2} = \sqrt{\frac{2}{\pi}} \frac{\sigma^2 Q \omega_0}{\omega_0^2 + 4Q^2(\omega - \omega_0)^2}, \quad (9)$$

which is a phenomenological model commonly used to model QPOs in X-ray binaries (e.g. Belloni et al. 2002; Vaughan & Uttley 2005),




Figure 5. Example of simulated light curves generated to test the sensitivity of our method to false positives. (Top left) Light curve generated using a DRW with bending time-scale of 65 d ($N = 300$, cadence roughly every 1 d). (Top right) Corresponding Lomb–Scargle periodogram. (Bottom left) As before, but initially with $N = 1000$ and then including two gaps of 45 d, one of 60 d and another of 100 d. (Bottom right) Corresponding Lomb–Scargle periodogram.

and is flexible enough to capture strict periodicities (where the coherence is extremely high). The Lorentzian has three parameters: the period of the oscillation $P = 2\pi/\omega_0$, the coherence or quality factor Q^3 , which sets how stable the oscillation amplitude is over time, or how peaked the Lorentzian is, and σ^2 , which is again the variance of the oscillation. Note that unlike periodogram modelling, which is agnostic to the underlying mechanism broadening the QPO, our periodic model here can only capture variations in amplitude.

The presence of Poisson (or Gaussian) noise can be included in the usual manner, by adding in quadrature the observation uncertainties to the covariance matrix (Rasmussen & Williams 2006; Foreman-Mackey et al. 2017). Here, we take the mean function to simply be the mean of the light curve. This choice also helps to limit the number of variable parameters, but again, our method can be generalized to include any mean function.

³Note that our definition is consistent with Belloni et al. (2002) but differs by a factor 2 compared to other works (e.g. Vaughan & Uttley 2005).

2.5 Model selection

Since we are performing a test for an additional component, our models will always be of the form *underlying noise + periodic component(s)*. Before testing for the presence of a signal, it is important to select a suitable null-hypothesis that captures the underlying, stochastic noise. Information criteria (IC), which penalize more complex models if the increase in fit statistic is not deemed ‘worthy’ of the extra parameters, are commonly used for model selection. If priors are included, model selection can be performed using the Bayes factor. Given that we have used non-informative priors throughout this work, we perform model selection using the Akaike Information Criterion (AIC; Akaike 1998):

$$\text{AIC} = 2k - 2 \ln L_*, \quad (10)$$

where models with higher AIC values are considered to have excessive complexity with respect to the quality of the data. Here, k is the number of model parameters and L_* for a particular model is the maximum of the likelihood function. The AIC is only correct asymptotically (i.e. for large sample sizes) but it can be corrected for

Table 2. Summary of the analysis carried out to test the sensitivity of our method to false positives. We tested for the presence of an additional Lorentzian (QPO) component in 50 light curves, simulated from a DRW PSD per baseline/cadence combination.

N^a	$< \Delta t >^b$ d	T^c d	p_{uniform}^d
100	10	1000	0.35
250	4	1000	0.56
500	2	1000	0.16
1000	1	1000	0.35
300	1	300	0.91
400	1	400	0.30
500	1	500	0.98

Notes.

^a Number of datapoints of the generated light curves.

^b Mean cadence.

^c Observing baseline.

^d p -values for the distribution of the 50 retrieved PPPs following a uniform distribution between 0 and 1, as expected when the null hypothesis is true.

finite sample sizes as shown in Hurvich & Tsai (1989):

$$\text{AICc} = \text{AIC} + \frac{2k(k+1)}{N-k-1}. \quad (11)$$

In order to find the best-fitting model, we have implemented a small iterative routine in which we start by testing each of the single-kernel models (Jitter, DRW, Matérn-3/2, $\text{SHO}_{Q=1/\sqrt{2}}$, and Lorentzian) on the data. From these five fits, we selected the one yielding the lowest AICc_{min} and those within $\Delta\text{AICc} = 2$ from AICc_{min} . Subsequently, we tested each of the selected kernels in combination with any of the other five. From this second stage, we again retained the lowest overall model and those within the above ΔAICc , and repeated the process until adding an extra component no longer resulted in a decrease in the AICc.

We have compared the results of this routine with brute-forcing testing all possible model combinations and while we have found this routine yields the correct model in most instances, this was not the case in two of the objects tested here. As an example, there might be instances where a combination of a Lorentzian + DRW might

Table 3. AICc, ΔAICc and p -values for the standardized residuals following a Gaussian distribution for the different models tested on the combined EPIC data of the Seyfert galaxy NGC 1365. ΔAICc refers to the increment in AICc with respect to the first model listed in the Table. Models with lower AICc values are not shown for clarity.

Model	AICc	ΔAICc	p -value
Lorentzian + DRW + $\text{SHO}_{Q=1/\sqrt{2}}$	92.4	0.0	0.08
Lorentzian + Matérn-3/2 + Matérn-3/2	92.7	0.4	0.01
Lorentzian + Matérn-3/2 + $\text{SHO}_{Q=1/\sqrt{2}}$	92.9	0.5	0.45
Lorentzian + Matérn-3/2 + DRW	93.2	0.8	0.01
Matérn-3/2 + $\text{SHO}_{Q=1/\sqrt{2}}$	95.7	3.4	0.18
Matérn-3/2 + DRW	95.8	3.4	0.06
$2 \times \text{Matérn-3/2}$	96.0	3.6	0.19
DRW + $\text{SHO}_{Q=1/\sqrt{2}}$	96.8	4.5	0.06
Lorentzian + Matérn-3/2	98.7	6.3	0.70
$2 \times \text{SHO}_{Q=1/\sqrt{2}}$	99.3	6.9	0.27
Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$	100.6	8.2	0.02
Lorentzian + DRW	101.3	9.0	0.0
Matérn-3/2	103.2	10.9	0.83
DRW	103.8	11.5	0.00
Lorentzian + $2 \times \text{DRW}$	105.3	12.9	0.00
$2 \times \text{DRW}$	107.4	15.0	0.00
$3 \times \text{DRW}$	111.6	19.2	0.00
$\text{SHO}_{Q=1/\sqrt{2}}$	114.7	22.3	0.44
Lorentzian + Jitter	121.3	28.9	0.00

be a better overall combination than a Lorentzian + Matérn-3/2, even if in isolation a Matérn-3/2 may be preferred over a DRW and Lorentzian. Nevertheless, we have found the routine useful in performing a preliminary triage and establishing the number of components required. Therefore after the model minimizing AICc was found, we have tested alternations keeping the same model components to refine the final selection. We leave developing a more refined search process when focusing on large-scale survey searches, where we will calibrate the method against specific data sets.

Once the model (or combination of model components) has been selected for the alternative and null hypothesis models, the posteriors derived from the null hypothesis (the noise-only model) can be used




Figure 6. (Left) Combined EPIC 0.3–10 keV light curve of the Seyfert galaxy NGC1365. (Right) Corresponding Lomb–Scargle periodogram (oversampled by a factor 5). The pink dashed line shows the mean periodogram of 10 000 light curves simulated from the posteriors of the Matérn-3/2 + DRW kernel (best-fitting model), with the shaded areas showing the 16 and 84 per cent percentiles of the distribution. The vertical black arrow shows the putative QPO reported by Yan et al. (2024). The dashed blue line shows the power spectrum of the observing window.




Figure 7. GP modelling results of the combined *XMM-Newton* EPIC data of the Seyfert galaxy NGC 1365 (Fig. 6). (Top left) Best-fitting DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ (solid orange line) and its 1σ uncertainties (shaded areas). The bottom panel shows the standardized residuals of the model. (Top right) PSDs derived from the *celerite* modelling (absolute rms normalization), showing the DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ (null) and the Lorentzian + DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ (alternative) models. The solid and shaded areas show the median and 1σ uncertainties derived from the posteriors. The dashed horizontal line shows the approximate Poisson level ($2\tilde{\Delta}t < \sigma_{\text{err}}^2$ where $\tilde{\Delta}t$ and σ_{err}^2 are the median sampling and the mean square error, respectively). (Bottom left) ACF of the standardized residuals. The shaded areas indicate the 95 percent confidence level expected for white noise. (Bottom right) Posteriors of the best-fitting DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ model. The contours indicate the 2D, 1σ and 2σ confidence levels (39 and 86 percent, respectively) and the dashed lines on the marginalized histograms indicate the 32, 50, and 84 percent percentiles (median $\pm 1\sigma$). The MCMC run for approximately 64 000 steps until convergence, from which we discarded the first 10 000 as burn-in.

to calibrate the reference LRT distribution as proposed by Protassov et al. (2002) and derive the posterior predictive p -value (PPP). In doing so, we are able to map ΔAICc changes to p -values.

Note it may not always be possible to establish a unique pair of alternative and null hypothesis models when the differences between two models are small (typically $\Delta\text{AICc} \lesssim 2$). In such instances hypothesis testing may be repeated using the various competing models in order to assess the robustness of the results to the choice of null hypothesis. This situation is akin to the regularly sampled case (e.g. Alston et al. 2014) and practical examples will be discussed in Sections 3.2.1 and 3.2.3. Note that we refer here to differences in ΔAICc between two alternative models, as small differences in ΔAICc between null and alternative models are precisely the type of situation our method is designed to address.

2.6 Goodness-of-fit

The goodness-of-fit is one of the main statistical quantities lacking in GP modelling. As opposed to the commonly used χ^2 statistic (whose value can be mapped to a p -value, indicating the likelihood that the data was generated by the model), the maximum of the GP likelihood L_* alone tells us nothing about whether the model is a good description of the data or whether the data can be described by a GP. Regarding the latter, there may be concern the light curves of accreting compact objects cannot be described using GPs, because the fluxes are observed to follow a lognormal probability density function (PDF), suggesting a multiplicative process generates the variability (Uttley, McHardy & Vaughan 2005), which we should not be able to describe using GPs. In the Appendix (Section B)

we discuss this aspect, and show through simulations that, despite the light curves having a lognormal distribution, GPs are able to recover the underlying process generating the variability (the PSD), indicating that their applicability might be broader than originally thought.

Despite the likelihood telling us nothing about whether the model is an appropriate description of the data, there are still several diagnostics that can be employed to test whether the model describes the data (or whether the data can be described by a GP). Following Kelly et al. (2014), we derive two diagnostics for model testing. First, we assess whether the standardized residuals follow a standard normal distribution ($\mu = 0$, $\sigma = 1$) by performing a Kolmogorov–Smirnov (KS) test. In practice, if the data cannot be described by a GP, the standardized residuals will be narrower than a standard normal distribution. This indicates the GP is overcompensating by assuming the entire variability in the time series is just random noise. In other words, the GP will indicate excess variance with respect to the data. On the other hand, residuals broader than a standard normal distribution indicate a deficiency in the chosen model. The other quantity we derive verifies there are no remaining trends by computing the autocorrelation function (ACF) of the standardized residuals. Any deviations from white noise will indicate the GP has not captured the full variability present in the light curve.

As described above (Section 2.2), in cases where the underlying distribution of the fit statistic is unknown, the reference distribution can be built empirically using numerical simulations. This general method (e.g. Waller et al. 2003; Kaastra 2017) involves simulating realistic data sets from the best-fitting model parameters that yielded L_* , applying the same fitting procedures and retrieving the reference distribution of L_{sim} values to compare to L_* . While a similar approach was also suggested by Kelly et al. (2011), who proposed to simulate light curves from the best-fit-derived PSD and compare their periodograms to the periodogram of the data, here we avoid the Fourier domain entirely by fitting the simulated light curves in time domain too.

This approach can be understood as follows: if the best-fitting parameters are truly representative of the data, then the simulated data sets (light curves in this case) will yield values of L close to L_* when fitted, and so L_* will sit roughly at the median of the L_{sim} distribution. If the best-fitting parameters are *not* representative of the data, then the value of L_* will be an outlier in the distribution of L_{sim} , i.e. L_* will in general be much lower than each of the L_{sim} values from the synthetic data sets; the model can then be statistically rejected (typically $p \lesssim 0.05$). If the data is overfitted, then the value of L_* will be towards the higher end of the L_{sim} distribution, implying the model has captured the data beyond the statistical noise, which is injected into the simulations (the model is deemed ‘too good’, which may also occur where the errors have been overestimated).

Finally, we note that, as opposed to more traditional χ^2 -fitting, where more complex models always lead to lower χ^2 , in GP modelling this is not necessarily the case. As opposed to χ^2 , where the likelihood depends exclusively on the fit residuals, in GP modelling, the likelihood depends on the residuals *and* a term depending on the kernel (or model) through the determinant of the covariance matrix. Therefore the best fit is determined from a trade-off between the residuals and the part of the likelihood that depends on the model alone. This makes it possible for less complex models to actually have more flexibility than models involving more hyperparameters, yielding better fits even if the complexity of the model is reduced.




Figure 8. Reference LRT distribution generated from simulated light curves from the posteriors of the DRW + SHO $Q=1/\sqrt{2}$ model (null hypothesis). The solid orange line shows a fit to the distribution using a log-normal. The T_{LRT} observed in the data is shown as a dashed black line.

2.7 Recipe

As a summary of the above, we outline the proposed steps of our method:

- (i) Chose a periodic kernel (or set of kernels) and a set of models for the underlying noise.
- (ii) Fit the models (and combinations of) to the data and rank them using one of the widely used IC (e.g. AICc and BIC).
- (iii) Ensure that the model with the lowest IC provides a good fit (e.g. via standardized residuals and their ACF or deriving the reference distribution for L_*).
- (iv) Compare the maximum of the likelihood function L_1 of the best-fitting signal + underlying noise model (the alternative model) to the maximum of the likelihood function L_0 of the model *without* the signal (the null-hypothesis) and retrieve the fit improvement, quantified as T_{LRT} .
- (v) Use the posteriors of the null model (the stochastic noise-only model) to generate synthetic data sets.
- (vi) Fit the synthetic data sets with the alternative and null-hypothesis models, derive the reference distribution for the LRT, and obtain the PPP by locating T_{LRT} in the distribution.
- (vii) Based on the significance of the signal, decide whether the component should be added to the null hypothesis (i.e. whether the signal is present in the data).

A PYTHON package that implements the proposed methodology has been made available at https://github.com/andresgur/mind_the_gaps and was employed throughout this work.

3 APPLICATION

We initially applied the recipe above to simulated data to explore the sensitivity of our method to variations in cadence and observing baseline. We present two sets of tests to examine the robustness of our approach to false negatives (failure to detect a signal) and false positives (identification of spurious signals).

3.1 Application to simulated data

In order to examine the sensitivity of our method to false negatives, we start by generating light curves (using the method explained in




Figure 9. NGC 7793 P13 *Swift*-UVOT light curve, with the segment considered for analysis indicated with a vertical dashed line and an arrow. (Right) Lomb-Scargle periodogram of the *U* band light-curve segment indicated in the left-hand panel. The black vertical arrows indicate harmonics (P/n) at $n = 2$ and $n = 7$ from the fundamental at ~ 64 d. The dashed blue line shows the power spectrum of the observing window.

Appendix A) with a (quasi)periodic component (Lorentzian) and red noise (a DRW) to mimic the case of a QPO identified in a stochastically varying light curve (e.g. Graham et al. 2015).

We assumed a period of 100 d for the QPO with a coherence $Q = 200$, a bending time-scale of 60 d for the DRW and that the QPO and the DRW contribute equally to a total variance of 6.7×10^{-4} (ct s^{-1})². We assumed a mean count rate of 0.1 ct s^{-1} for the source, background contribution of 1 percent and 2 ks exposure for all observations (these values were motivated by the faintest sources *Swift*-XRT is capable of monitoring).

We performed two types of tests using our input PSD model. First, to test the sensitivity of our method to changes in the sampling frequency, we generated light curves with a length of approximately $T = 1000$ d, with a sampling rate Δt drawn from a Gaussian distribution of mean = 1, 2, 4, and 10 d, and a standard deviation of 0.2 d, such that the light curves had 1000, 500, 250, and 100 datapoints, respectively, and a realistic, irregular observing cadence. Fig. 3 shows a test light curve ($N = 250$ and $\Delta t \approx 4$ d) and its corresponding (Lomb–Scargle) periodogram. Secondly, to test the effects of having a shorter baseline, we fixed Δt (mean and standard deviation of 1 and 0.2 d, respectively) but progressively reduced the number of datapoints to generate light curves of shorter duration. In addition to the 1000 d light curve, we also simulated light curves spanning approximately $T = 500, 400$, and 300 d, respectively.

For a given cadence/baseline combination, we simulated 100 light curves and carried out the PPP method described in Section 2.3, i.e. 100 light curves were fitted with the DRW and the DRW + Lorentzian models and the LRT reference distribution was built using 2000 simulations from the DRW posteriors. We chose to simulate 100 light curves as a trade off between computational time and having roughly a representative sample for each cadence/baseline. Similarly, the rather low number of 2000 simulations was set by computational constraints.

Fig. 4 shows the distribution of the retrieved p -values for the 100 light curves for the case of varying cadence (left panel) and varying baseline (right panel). Table 1 shows the mean retrieved p -values and the number of significant ($p \lesssim 0.01$) detections per cadence/baseline combination.

From Fig. 4, we can see that, despite the rather low count rates of the light curves, we are able to recover the period in more

Table 4. As per Table 3 but now showing the AICc, ΔAICc , and p -values for the standardized residuals following a Gaussian distribution for the different models tested against the *Swift*-UVOT data of P13.

Model	AICc	ΔAICc	p -value
3×Lorentzian + Jitter	−556.0	0.0	0.32
3×Lorentzian + SHO $Q=1/\sqrt{2}$	−555.9	0.2	0.71
3×Lorentzian + Matérn-3/2	−554.1	1.9	0.27
3×Lorentzian + DRW	−553.6	2.4	0.08
2×Lorentzian + Jitter	−552.2	3.8	0.43
2×Lorentzian + DRW	−549.5	6.5	0.35
2×Lorentzian + SHO $Q=1/\sqrt{2}$	−549.0	7.0	0.38
2×Lorentzian + Jitter + SHO $Q=1/\sqrt{2}$	−548.4	7.6	0.008
2×Lorentzian + Jitter + DRW	−548.1	7.9	0.02
2×Lorentzian + Jitter + Matérn-3/2	−547.7	8.3	0.02
2×Lorentzian + Matérn-3/2	−540.1	15.9	0.18
2×Lorentzian	−531.5	24.5	0.52
Lorentzian + Jitter	−529.2	26.8	0.44
Lorentzian + DRW	−526.5	29.5	0.365
Lorentzian + SHO $Q=1/\sqrt{2}$	−525.9	30.1	0.45
SHO $Q=1/\sqrt{2}$ + 2×DRW	−522.6	33.5	0.44
2×SHO $Q=1/\sqrt{2}$ + DRW	−516.5	39.5	0.78
DRW	−495.5	60.5	0.178
Jitter	−495.1	60.9	0.05
Matérn-3/2	−494.3	61.8	0.28
SHO $Q=1/\sqrt{2}$	−493.2	62.8	0.10
2×DRW	−492.9	63.1	0.11
Lorentzian	−492.6	63.4	0.69
SHO $Q=1/\sqrt{2}$	−492.1	63.9	0.21

than half of the instances (in $\geq 62/100$) as long as enough cycles are observed. In particular, there is little improvement in detection rates in the light curves with fixed baseline (left panel). This is partially due to the fact that the parameter space (namely P and ω_{bend}) is accommodated with the sampling (as the lowest cadence in the light curve sets the minimum allowed P and bend time-scale) but it suggests the number of cycles might be the most important metric when attempting to detect a periodicity. This simple result is consistent with the requirement identified by Vaughan et al. (2016) when looking for periodicities in stochastically varying systems.




Figure 10. GP modelling results of the *Swift*-UVOT light-curve segment (Fig. 9) of the pulsating ULX NGC 7793 P13. (Top left) Best-fitting $3 \times \text{Lorentzian} + \text{Jitter}$ model. (Top right) PSD of the best-fitting model. The contribution of the Lorentzians to the PSD (blue solid line) are shown as a orange, green, and red dashed lines. (Bottom left) ACF of the standardized residuals of the best-fitting model. (Bottom right) Posteriors of the periods of the three Lorentzians and the Jitter amplitude (other parameters omitted for clarity). The MCMC sampling run for approximately 49 000 steps until convergence, from which we discarded the first 22 000 as burn-in. Symbols as per Fig. 7.

Interestingly, our results show that light curves spread over time with lower-cadence sampling prove more advantageous for detecting periodicities compared to shorter light curves with a higher number of datapoints (at least as long the time-scale of the period is much longer than the cadence). For instance, in the case of a $\Delta t \approx 2$ d and $N = 500$ (left panel), we are able to recover the period in 68/100 instances, whereas for $\Delta t \approx 1$ d and the same number of datapoints (right panel), only in 44/100 instances we are able to recover the period (Table 1). Our results suggest that, in the presence of stochastic variability, it will be hard to reliably confirm periodicities in a light curve covering five or less cycles of the putative period, in agreement with Vaughan et al. (2016).

Next we turn to examine the robustness of our method to false positives (i.e. misidentifying aperiodic variability as periodic), one of the aspects that motivated us to devise better methods for period detection. To this end, we performed a second series of simulations using a simple DRW with Δt again drawn from the same Gaussian distributions. Here, we use a mean count rate of 35 ct s^{-1} , a

background contribution of 300 ct s^{-1} and a variance of $36 (\text{ct s}^{-1})^2$ and a break at 65 d for the DRW. These parameters are similar to those observed in the *Transiting Exoplanet Survey Satellite* (*TESS*) light curves of Blazars (see Section 3.2.4) and were chosen to generate light curves which *appear* periodic. Fig. 5 (top panel) shows an example light curve ($N = 300$, median $\Delta t = 1$ d) with the corresponding periodogram; clearly naive inspection of the periodogram may lead to the conclusion that some genuine periodicity is present in the light curve.

We ran our PPP method again as above, this time creating 50 light curves per cadence/baseline combination and simulating 2000 light curves from the DRW posteriors, and comparing the fit improvements when adding a Lorentzian. Under the absence of the signal, the distribution of retrieved PPP values from the T_{LRT} tests is expected to be uniformly distributed from 0 to 1. We therefore tested using a KS test whether the 50 retrieved PPP values per cadence/baseline combination followed the aforementioned uniform distribution. Table 2 shows all p_{uniform} -values are consistent with the




Figure 11. (Left) *Swift*-XRT light curve of NGC 7793 P13 with the segment considered for analysis lying between the two vertical, dashed black lines. (Right) Lomb–Scargle periodogram of the segment indicated in the right-hand panel. Symbols as per Fig. 9. The black vertical arrows show the fundamental period frequency ($P \sim 65.6$ d) and an harmonic at $n = 2$.

expected uniform distribution, regardless of the observing strategy, indicating false positives are unlikely, at least for the cadences explored here.

Finally, to test our ability to avoid false-positives when gaps are introduced, we have repeated the method above, keeping the same PSD and using $N = 1000$, samples taken roughly at 1-d intervals but then adding three gaps: one of 45 d, one of 60 d, and another of 100 d. Fig. 5 (bottom panels) shows an example light curve and corresponding periodogram. We then ran the PPP method for 50 sample light curves and tested whether the recovered PPPs followed a uniform distribution as expected for the absence of a signal, finding a p -value of 0.37. This indicates that, at least for the cadence/variability time-scales explored here, our method can robustly avoid false-positives as long as the noise is well described. We will present an exhaustive exploration of period detectability under different combinations of observing strategies in a forthcoming publication.

3.2 Application to real data

Our method has been developed for instances where irregular sampling hampers obtaining the PSD in a straightforward manner. Such a scenario is routinely encountered in many studies of AGNs (e.g. Jiang et al. 2022) and other accreting systems.

Nevertheless, to show our method is not restricted to irregularly sampled time series, we first apply our method to a recent claim of a QPO in *XMM-Newton* data of a Seyfert galaxy (Section 3.2.1). We then explore claims of periodicities in *Swift* data (Gehrels et al. 2004) (both UVOT and XRT) of a ULX (Section 3.2.2), a QPO in an AGN in RXTE data (Section 3.2.3) and finally revisit a recent claim of a QPO in the *TESS* (Ricker et al. 2014) light curve of a Blazar (Section 3.2.4). The results of the analysis are then discussed in Section 3.3. The choice of priors and the procedure used to derive the best-fitting parameters and their posteriors is described in Appendix C and unless stated otherwise, we perform 10 000 simulations to derive the LRT reference distribution in the calculation of the PPP value.

Table 5. As per Table 3 but now showing the AICc, Δ AICc and p -values for the standardized residuals following a Gaussian distribution for the different models tested against the *Swift*-XRT data of the pulsating ULX NGC7793 P13.

Model	AICc	Δ AICc	p -value
$3 \times \text{Lorentzian} + \text{Matérn-3/2} + \text{Jitter}$	-1673.4	0.0	0.001
$3 \times \text{Lorentzian} + \text{DRW} + \text{Jitter}$	-1669.2	4.2	0.001
$2 \times \text{Lorentzian} + \text{Matérn-3/2} + \text{Jitter}$	-1664.5	8.9	0.003
$2 \times \text{Lorentzian} + \text{Matérn-3/2} + \text{DRW} + \text{Jitter}$	-1662.2	11.5	0.006
$2 \times \text{Lorentzian} + \text{DRW} + \text{Jitter}$	-1661.2	12.2	0.003
$2 \times \text{Lorentzian} + \text{DRW} + \text{Jitter} + \text{SHO}_{Q=1/\sqrt{2}}$	-1659.9	13.5	0.02
$2 \times \text{Lorentzian} + 2 \times \text{DRW} + \text{Jitter}$	-1657.1	16.6	0.01
$\text{Lorentzian} + \text{DRW} + \text{Matérn-3/2} + \text{Jitter}$	-1656.1	17.3	0.003
$\text{Lorentzian} + \text{DRW} + \text{Jitter}$	-1649.4	23.9	0.003
$\text{Lorentzian} + 2 \times \text{DRW}$	-1647.4	25.9	0.004
$\text{Lorentzian} + \text{Matérn-3/2} + \text{Jitter}$	-1647.1	26.3	0.00
$3 \times \text{Lorentzian} + 2 \times \text{DRW}$	-1644.8	28.5	0.00
$\text{Lorentzian} + \text{DRW} + \text{Matérn-3/2}$	-1641.8	31.6	0.02
$3 \times \text{Lorentzian} + \text{Matérn-3/2}$	-1641.6	31.8	0.00
$2 \times \text{Lorentzian} + 2 \times \text{Matérn-3/2}$	-1641.5	32.3	0.00
$2 \times \text{Lorentzian} + 2 \times \text{DRW}$	-1640.8	32.5	0.001
$3 \times \text{Lorentzian} + \text{SHO}_{Q=1/\sqrt{2}}$	-1639.0	34.4	0.00
$3 \times \text{Lorentzian} + 3 \times \text{DRW}$	-1638.7	34.6	0.00
$3 \times \text{Lorentzian} + 2 \times \text{Matérn-3/2}$	-1637.4	36.0	0.00
$3 \times \text{Lorentzian} + \text{DRW}$	-1636.9	36.8	0.00
$2 \times \text{Lorentzian} + \text{DRW} + \text{Matérn-3/2}$	-1636.7	37.1	0.00
$3 \times \text{Lorentzian} + 3 \times \text{Matérn-3/2}$	-1632.3	41.4	0.00
$2 \times \text{Lorentzian} + \text{DRW}$	-1631.1	42.7	0.00
$\text{Lorentzian} + \text{DRW}$	-1627.7	46.0	0.00
$\text{Lorentzian} + \text{Jitter}$	-1623.3	50.1	0.00
$\text{Matérn-3/2} + \text{Jitter}$	-1619.3	54.5	0.00
$\text{Lorentzian} + \text{SHO}_{Q=1/\sqrt{2}} + \text{Jitter}$	-1619.2	54.2	0.00
$3 \times \text{Lorentzian} + \text{SHO}_{Q=1/\sqrt{2}} + \text{Jitter}$	-1618.7	55.1	0.00
$\text{Lorentzian} + \text{Matérn-3/2}$	-1617.1	56.3	0.00
$\text{Lorentzian} + \text{SHO}_{Q=1/\sqrt{2}}$	-1611.6	61.8	0.002
$\text{DRW} + \text{Jitter}$	-1607.9	65.9	0.001
Lorentzian	-1602.3	71.4	0.00
DRW	-1594.2	79.5	0.001
Matérn-3/2	-1590.5	83.2	0.00
$\text{SHO}_{Q=1/\sqrt{2}}$	-1588.4	85.4	0.00




Figure 12. GP modelling results of the *Swift*-XRT data of the pulsating ULX NGC 7793 P13. (Top Left) Best-fitting $3 \times$ Lorentzian + Matérn-3/2 + Jitter model to the *Swift*-XRT 0.3–10 keV light-curve segment of NGC 7793 P13 shown in Fig. 12. (Top right) PSD of the best-fitting model. (Bottom left) ACF of the standardized residuals of the best-fitting model. (Bottom right) Posteriors of the periods of the three Lorentzians and the Matérn-3/2 (the other parameters are omitted for clarity). The MCMC run for approximately 620 000 steps until convergence and about 124 000 were discarded for the burn-in. Symbols as per Fig. 7.

3.2.1 A high-frequency QPO in the Seyfert NGC 1365

Using *XMM-Newton* data and employing techniques such as the Lomb–Scargle periodogram, Yan et al. (2024) recently reported the detection (significance of 3.6σ) of a high-frequency (~ 4566 s) QPO in the Seyfert galaxy NGC1365. We used obsid 0 205 590 301 where Yan et al. (2024) reported the detection of the QPO and reanalysed the EPIC-pn and MOS data using tasks `eproc` and `emproc` in SAS version 20.0.0. We filtered the light curves for particle flaring by first extracting background 10–12 keV light curves and then inspected these visually to set a threshold count-rate to reject times of high-background flaring. We applied the standard quality filters and selected `PATTERN` ≤ 4 events for pn and `PATTERN` ≤ 12 events for the MOS cameras. We used `eregionanalyse`, with the input source coordinates, to select a suitable source region. The circular region as determined by the task contained a fainter source near to the target in some instances, so to avoid contamination we reduced the radius to ~ 55 arcsec, but keeping the same centroid position. A slightly larger circular region on the same chip, away from the

readout region and as close as possible to the source region, was selected for background light curve extraction. The final light curve was corrected for effects including losses due to vignetting, chip gaps, and bad pixels using `epicclcorr`. Following Yan et al. (2024), the three light curves were binned into 200s and their net count rates were combined into a final light curve. Because the asynchronicity of the three instruments can introduce spurious variability (Barnard et al. 2007), we ensured the start and end times were the same for the three detectors and inspected the individual and combined light curves visually.

Fig. 6 shows the 0.3–10 keV combined EPIC light curve of NGC 1365, which comprises 289 datapoints and a duration of 57 800 s. The right-hand panel of Fig. 6 shows the corresponding periodogram, with an arrow at ~ 0.05 d indicating the claimed QPO by Yan et al. (2024).

Table 3 lists the models tested to the data, ranked by AICc value. We can see that the best-fitting model comprises a Lorentzian (describing the putative QPO) and $\text{DRW} + \text{SHO}_{Q=1/\sqrt{2}}$ kernels to describe the underlying noise. Compared to a $\text{DRW} + \text{SHO}_{Q=1/\sqrt{2}}$




Figure 13. (Left) RXTE light curve of the AGN NGC 4945. The segment where Smith et al. (2020) reported the significance of the QPO to be strongest is highlighted with a dashed line and an arrow. (Right) Lomb–Scargle periodogram of the segment indicated in the right-hand panel. The vertical black arrow shows the QPO reported by Smith et al. (2020). The dashed blue line shows the power spectrum of the observing window. The dashed purple line shows the model derived from periodograms of 10 000 light curves generated from the posteriors of the $\text{SHO}_{Q=1/\sqrt{2}}$ model parameters (best-fitting model), with the shaded areas indicating the 16 and 84 per cent percentiles of the distribution. See text for details.

Table 6. As per Table 3 but now showing the AICc, ΔAICc , and p -values for the standardized residuals following a Gaussian distribution for the different models tested against the RXTE data of the AGN NGC4945. The left and right values are for the analysis of the full light curve and the segment shown in Fig. 13, respectively.

Full light curve				Smith et al. (2020)			
Model	AICc	ΔAICc	p -value	Model	AICc	ΔAICc	p -value
DRW	246.1	0.00	0.00	Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$	97.9	0.0	0.22
DRW + Jitter	247.2	1.1	0.00	Lorentzian + Matérn-3/2	98.3	0.4	0.07
DRW + $\text{SHO}_{Q=1/\sqrt{2}}$	248.5	2.4	0.00	Lorentzian + DRW	99.2	1.3	0.03
DRW + Matérn-3/2	248.8	2.7	0.00	Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$ + Jitter	99.5	1.6	0.22
Lorentzian + DRW	248.8	2.7	0.00	2×Lorentzian	101.0	3.1	0.42
2×DRW	249.5	3.4	0.00	Lorentzian + DRW + Jitter	101.6	3.7	0.03
Lorentzian + 2×DRW	250.2	4.1	0.00	Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$ + DRW	102.0	4.1	0.12
Lorentzian + Jitter	251.6	5.5	0.03	Lorentzian + 2× $\text{SHO}_{Q=1/\sqrt{2}}$	102.4	4.4	0.08
Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$	253.7	7.6	0.00	Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$ + Matérn-3/2	102.9	5.0	0.16
Lorentzian + Matérn-3/2	261.6	11.4	0.00	Lorentzian + Matérn-3/2 + DRW	102.9	5.0	0.12
Matérn-3/2	255.2	9.1	0.002	Matérn-3/2	103.8	5.9	0.14
$\text{SHO}_{Q=1/\sqrt{2}}$	264.9	18.8	0.02	DRW	104.5	6.6	0.03
Lorentzian	267.9	21.8	0.03	Lorentzian	104.6	6.7	0.22
Jitter	310.4	64.3	0.09	$\text{SHO}_{Q=1/\sqrt{2}}$	104.9	7.0	0.33

only model, the addition of the Lorentzian represents a $\Delta\text{AICc} = 4.5$ fit improvement. Fig. 7 shows the best-fitting DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ model, its PSD, the ACF of the standardized residuals and the posteriors. Both models provide an adequate description of the data whilst the ACF (bottom left panel) shows that the variability is approximately well captured by the DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ model. Therefore, using the posteriors of the DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ model, we tested whether the addition of the Lorentzian was supported by the data.

Fig. 8 shows the reference LRT distribution derived from light-curve simulations generated from the posteriors of the DRW + $\text{SHO}_{Q=1/\sqrt{2}}$ model. As can be seen from the figure, the addition of the Lorentzian (the QPO component) is significant only at the ~ 91 per cent level ($\sim 1.7\sigma$).

As stated in Section 2.5, owing to the relatively small difference in fit improvement ($\Delta\text{AICc} = 0.4$) with respect to the Lorentzian + 2×Matérn-3/2, we have repeated the significance calculation

with the posteriors of the this other model too. We have found the significance of ~ 86 per cent, in line with the lower $\Delta\text{AICc} = 3.6$ provided by this model with respect to the null hypothesis. Therefore we do not support the presence of a QPO in this light curve of NGC 1365.

3.2.2 The pulsating ULX NGC 7793 P13

Since the discovery of its ~ 63 d period (Motch et al. 2014), the pulsating neutron star ULX NGC 7793 P13 (Fürst et al. 2016; Israel et al. 2017, P13 hereafter) has been intensively monitored by *Swift*. Being among the brightest ULXs in the optical bands with a V magnitude of around ~ 20.2 (Motch et al. 2014), it is one of only a small number of ULXs where the long-term variability can be studied by both the *Swift*-UVOT and *Swift*-XRT. The irregular sampling of the monitoring of this source has revealed two closely but significantly different periods: an ~ 64 -d period in the U band




Figure 14. GP modelling results of the RXTE light-curve segment shown in Fig. 13 of the AGN NGC 4945. (Top left) Best-fitting Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$ model. (Top right) PSDs of the null hypothesis and alternative models. (Bottom left) ACF of the standardized residuals of the Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$ model. (Bottom right) Posterior parameters for the Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$ model. The MCMC sampler run for approximately 135 000 steps until convergence and about 35 000 steps were discarded for the burn-in. Symbols as per Fig. 7.

and an ~ 65 -d period in the X-rays (Hu et al. 2017; Fürst et al. 2018). An advantage of using GP for period searching is that uncertainties are well-defined as we can marginalize over the noise parameters. Therefore we can assess both the significance of the claimed periodicities and also the difference between them.

The UVOT data were kindly provided by Khan & Middleton (2023), to which we refer the reader for the data reduction details. The *U* band contained the largest amount of observations (Fig. 9; 260 observations compared to $\lesssim 20$ in other bands); we therefore analysed only this band. While an advantage of GP modelling is that more data, regardless of the gaps, should lead to tighter constraints, here the few additional and largely spaced datapoints at the beginning of the monitoring increase the computational cost dramatically for a small gain in accuracy, particularly in our false-alarm probability calculation. Therefore we only considered the data after MJD 57 500 where the monitoring is denser (Fig. 9).

Table 4 lists the models tested in our fit to the data, ranked by AICc value. Part of the modelling was guided by a visual inspection of the Lomb–Scargle periodogram of the light curve segment, which

we show in Fig. 9. We can see the main peak at $P \sim 64$ d and some harmonics at 32 d ($P/2$) and ~ 9 d ($P/7$), indicating the periodicity – if real – is not a pure sinusoid.

The harmonics are also reflected in the GP modelling: we can see from Table 4 that the preferred model consists of three Lorentzians + a Jitter component for the underlying noise. From Table 4 we can also see that this model is preferred over one where the underlying noise is instead described by a DRW ($\Delta\text{AICc} = 2.4$), suggesting that white noise is the statistically preferred null hypothesis.

The standardized residuals of the best-fitting model are fully consistent with a Gaussian distribution, indicating the variability is well-described by a GP, whereas the overall variability is also well-captured, as indicated by the ACF of the standardized residuals (Fig. 10 bottom left panel).

Having established white noise to be a good representation for the underlying noise, we proceed to test for the presence of the Lorentzian(s) components in a hierarchical manner. First, we test for the first Lorentzian over the Jitter-only model, using the posteriors of the Jitter model. If significant, we subsequently use the posteriors

of the Lorentzian + Jitter to test for an additional Lorentzian until the new added Lorentzian is no longer significant.

For the first Lorentzian, we found that none of the simulations showed a T_{LRT} as high as that observed in the data. Fitting the LRT distribution with a lognormal, we estimate the period to be highly significant (99.999 per cent or $\sim 6\sigma$). For the second harmonic at ≈ 32 d, we find the significance to be ≈ 99.1 per cent, while for the third component the significance is 95.8 per cent.

From our best fit (Fig. 10), we obtained $P = 63.9 \pm 0.4$ d, with a coherence $Q = 220^{+727}_{-146}$, indicating the period amplitude is stable over this time period.

We now examine the *Swift*-XRT data and the claimed ~ 65 -d period. The full *Swift*-XRT light curve is shown in Fig. 11 and was extracted using the online tools (Evans et al. 2007, 2009) keeping all snapshots with a detection significance of $\geq 2\sigma$. Modelling the full light curve would add additional complexity due to potential deviations from stationarity and more complex fine-tuning of the mean function. Additionally, the few largely spaced datapoints would again add little gain in constraining power at the expense of significant computational time. Hence, we analysed the indicated segment in Fig. 11 where the variability appears stationary and the monitoring is densest. This segment lasts 984.3 d with a mean observing cadence of 3.1 d.

Table 5 shows the ΔAICc with the various models tested along with their p -values for the standardized residuals following a standard normal distribution. The best-fitting model is a combination of three Lorentzians + a Matérn-3/2, where a Jitter term is needed as there is additional white-noise variability (jittering) that cannot be captured by any of the kernels. Indeed, we find that none of the models adequately describes the data based on the p -values of the residuals following a standard normal distribution. This may not be surprising as the distribution of count-rates is itself non Gaussian ($p = 0.008$ for rejecting a Gaussian distribution based on a KS test). Nevertheless, from Fig. 12, we can see that the failure to describe the data is mostly due to a few datapoints strongly deviating from the model. This is clearly seen in the ACF (Fig. 12 bottom left panel), which confirms the lack of trends in the standardized residuals. Thus, while the model may not capture the full complexity of the data, we can at least ascertain that the variability is well represented by the combination of the three Lorentzians + Matérn-32 + Jitter. The PSD of this models is shown in Fig. 12 (top right panel).

As with the UV data, we tested for the addition of the Lorentzians to the noise model in a hierarchical manner. The first Lorentzian with $P \sim 65$ d is found to be significant at the 99.99 per cent level ($> 3\sigma$), the addition of a second is significant at the ≈ 99.2 per cent level and the third Lorentzian is marginally significant, at the ~ 91.5 per cent level.

From Fig. 12 our final estimate for the period is $P = 65.6 \pm 0.6$ d. As for the UV data, the high coherence ($Q \gtrsim 100$) suggests the periodicity is stable throughout the segment.

3.2.3 NGC 4945

Smith, Robles & Perlman (2020) claimed an ~ 42 -d QPO in the irregularly sampled *RXTE* data of the Type 2 Seyfert, NGC 4945 with a significance of 10.2σ .⁴

⁴Note that the authors also quote a false-alarm probability of 2.87 per cent, which corresponds to $\sim 2.2\sigma$ only.




Figure 15. LRT distribution generated from simulated light curves from the posteriors of the $\text{SHO}_{Q=1/\sqrt{2}}$ model (null hypothesis) for NGC 4945. The solid orange line shows a fit to the distribution using a log-normal. The LRT observed in the data is shown as a dashed black line.

Following Smith et al. (2020), we obtained the *RXTE*/PCA data from the University of California archive⁵ (for details regarding data filtering criteria, we refer the reader to their website). Fig. 13 shows a segment of the light curve where the monitoring was densest (cf. fig. 2 in Smith et al. 2020). The full light-curve spans 442 d, with a median cadence of 2.25 d. The authors found the significance of the QPO to be the strongest in the segment towards the end of the light curve after the vertical dashed line in Fig. 13. This segment spans 192 d with a mean cadence of 2.04 d.

We first focus on the analysis of the full light curve. Table 6 shows the ΔAICc for the set of models explored, with the DRW yielding the lowest AICc (with $\Delta\text{AICc} = 2.7$ over the Lorentzian + DRW). This already suggests the data can be explained under a simpler, stochastic model. Performing light-curve simulations from the DRW posteriors, we find a significance of ≈ 39 per cent for the Lorentzian component, indicating the addition of the Lorentzian is not supported by the data. However, we note the residuals in all models are narrower than a standard normal distribution ($\sigma = 0.67$), indicating the variability is not well described by a GP. The downward trend in flux around MJD 53 900 may indicate the process is non-stationary over the time-scales analysed here.

We proceed to focus on the segment indicated to the right-hand side of the vertical dashed line in Fig. 13, where the authors claimed the QPO significance to be highest. Table 6 lists the models tested against the data in this segment. In this case we find potential evidence for a periodic component, as a model including a Lorentzian (the broadband noise modelled with an $\text{SHO}_{Q=1/\sqrt{2}}$) provides the lowest AICc . The standardized residuals and their ACF are shown in Fig. 14. The standardized residuals are compatible with Gaussianity and, therefore, the assumption of a GP is reasonable. We note, however, the ACF indicates there is still slight variability not captured by the model potentially indicating that models outside *celerite* might be more appropriate. Nevertheless, most of the variability is reasonably captured. The central frequency of the Lorentzian is found to be $P = 42^{+2}_{-3}$ d, which matches the periodicity reported by Smith et al. (2020).

⁵<https://cass.ucsd.edu/rxteagn/>

Using the posteriors from the best-fitting $\text{SHO}_{Q=1/\sqrt{2}}$ as the null hypothesis, we obtained the significance of the (quasi)-periodic component. To estimate the background contribution for our light-curve simulations, we assumed the mean source rate was 5 per cent of the background rate. Although in the average spectra the source contributed 10 per cent to the total rate, 5 per cent is both consistent with previous work (Done et al. 2003) and we found the simulated lightcurve errorbars matched more closely the data errorbars. As we are assuming a higher background than in the average spectrum, our simulations will be less likely to generate a spurious signal (they will have increased levels of white noise) and the estimated QPO significance will tend to overestimate the true significance, if anything.

Fig. 15 shows the reference LRT distribution derived from the posteriors of the $\text{SHO}_{Q=1/\sqrt{2}}$ model. The putative periodicity has a significance of ~ 98.7 per cent (i.e. $\approx 2.5\sigma$), which is indeed quite high, but it may not be considered sufficient to claim a detection.

Fig. 13 shows the mean Lomb–Scargle periodogram of 10 000 light curves simulated using the posteriors of the $\text{SHO}_{Q=1/\sqrt{2}}$ model. Using the process outlined in Section 2.6 to map the $L_{\text{max}}^{\text{obs}}$ to a goodness-of-fit, we find a p -value of 0.9 using the $\text{SHO}_{Q=1/\sqrt{2}}$ model, a deviation of $\sim 1.7\sigma$ from the mean, indicating the fit is an acceptable description of the data.

Given that the Matérn-3/2 provides the lowest AICc when used in isolation (Table 6), one could argue it represents the best null-hypothesis. Moreover, as discussed in Section 2.5, based on the low $\Delta\text{AICc} = 0.4$ between the Lorentzian + $\text{SHO}_{Q=1/\sqrt{2}}$ and the Lorentzian + Matérn-3/2 models, it may be argued these two models offer similar levels of goodness of fit (i.e. we cannot distinguish between the two with the data at hand). We therefore repeated the significance calculation with the posteriors from the Matérn-3/2-only model, and found a similar value for the significance (~ 96 per cent). This is consistent with the lower ΔAICc provided by this model when the Lorentzian is added compared to the $\text{SHO}_{Q=1/\sqrt{2}}$ model (Table 6) and indicates our results are not strongly dependent on the continuum choice (so as long as it is representative of the data).

3.2.4 The Blazar B0537-441

Tripathi et al. (2024) recently reported the detection of a QPO of ~ 6.5 d in the Blazar B0537-441, using *TESS* data. We obtained the *TESS* light curves from sectors 32 and 33 (as analysed by these authors), reduced by the Science Processing Operations Center (SPOC; Jenkins et al. 2016) using the python *lightcurve* package. The light curve is extracted using aperture photometry and then corrected with the presearch data conditioning module to remove long-term trends and systematics caused by the spacecraft. The data from sectors 32 and 33 had been processed with pipeline versions *spoc-5.0.21-20210107* and *spoc-5.0.22-20210121*, respectively. For computational reasons we rebinned the light curves to 1 h which still allowed us to analyse the variability present and probe the relevant time-scales, see Fig. 16. The light curve had a duration of ~ 53 and a total of 1164 datapoints, with two ~ 1 -d gaps at MJD ~ 2186 and MJD ~ 2214 due to the satellite’s orbit and another gap at MJD ~ 2202 due to the observing strategy of *TESS*.

Table 7 lists the models tested against this data set. While some of the single-component models provide seemingly better fits according to the AICc, we can see from their p -values that none of these models provide a satisfactory description of the data ($p \lesssim 0.05$). In such models, the standardized residuals are broader ($\sigma \sim 2$) than expected for a standard normal distribution, indicating a deficiency in the fit.

The first model that provides an adequate description of the data ($p = 0.4$) is a combination of a Lorentzian and a DRW. The model that maximizes the likelihood, its standardized residuals, ACF, and posterior parameters is shown in Fig. 17. The best-fitting suggests a quasi-periodicity ($Q = 4^{+3}_{-1}$) with $P = 4.8^{+0.5}_{-0.4}$ d. On the other hand, the bend of the DRW is not well constrained, most likely owing to the relatively short baseline (≈ 54 d) of the data; in our modelling, the DRW mostly acts as a power law with $\beta = -2$. The DRW + Lorentzian model provides an $\Delta\text{AICc} = 66$ with respect to the DRW-only model; we proceeded to test whether the ΔAICc was significant using the posteriors of the DRW-only model. We found the Lorentzian component to be significant at the ~ 99.98 per cent ($\sim 3\sigma$) level (Fig. 18). This is in agreement with the high ΔAICc




Figure 16. (Left) *TESS* light curve of the Blazar, B0537-441 from sectors 32 and 33 (cf. fig. 3b in Tripathi et al. 2024). (Right) Corresponding Lomb–Scargle periodogram (black solid line). The power spectrum of the observing window is shown as per Fig. 9. The purple dashed and dotted green lines show the best-fitting Lorentzian + DRW, and DRW-only models (with the shaded areas showing the 16 and 84 per cent percentiles of 10 000 simulations; see text for details).

Table 7. As per Table 3, showing the AICc, Δ AICc, and p -values for the different models tested against the TESS data of the Blazar, B0537-441.

Model	AICc	Δ AICc	p -value
Matérn-3/2 + SHO $_{Q=1/\sqrt{2}}$	3526.2	0.0	0.0
2×Matérn-3/2	3526.3	0.1	0.0
Matérn-3/2 + DRW	3526.5	0.3	0.0
Lorentzian + SHO $_{Q=1/\sqrt{2}}$	3527.6	1.4	0.0
Lorentzian + Matérn-3/2	3527.7	1.5	0.0
2×Lorentzian + SHO $_{Q=1/\sqrt{2}}$	3527.8	1.3	0.0
2×Lorentzian + Matérn-3/2	3528.4	2.2	0.0
Matérn-3/2	3528.6	3.5	0.0
SHO $_{Q=1/\sqrt{2}}$	3536.9	10.7	0.0
2×Lorentzian + DRW	3560.3	32.4	0.01
Lorentzian + DRW	3573.9	47.7	0.40
Lorentzian + 2×DRW	3578.1	50.5	0.35
2×Lorentzian	3578.3	52.2	0.22
2×Lorentzian + Jitter	3581.1	53.7	0.28
Lorentzian	3608.0	81.8	0.94
Lorentzian + Jitter	3610.0	83.9	0.97
DRW	3639.9	113.8	0.27
2×DRW	3644.0	117.8	0.27

observed between the Lorentzian + DRW and the DRW-only model. Therefore, we deem the addition of the Lorentzian to be supported by the data.

Fig. 16 shows a comparison of the periodogram of the best-fitting model (Lorentzian + DRW), the periodogram of the DRW-only model and the Lomb–Scargle of the data. The Lomb–Scargle periodograms of the best-fitting models were derived by taking the Lomb–Scargle periodogram of 10 000 light curves generated from the posteriors of each of the two models. We can see that the strongest period in periodogram is consistent with our best-fitting period of 4.8 d and that the noise is reasonably captured by the DRW.

3.3 Analysis of the results

As we have shown, the methodology outlined here is particularly suited for the analysis of irregularly sampled time series commonly associated with monitoring of systems such as ULXs or AGNs (e.g. Uttley et al. 2002), but its applicability is not restricted to irregularly-sampled time series (Section 3.2.1). We have shown its application to the short (~ 50 ks) nearly regularly sampled time series that may be obtained with observatories such as *XMM-Newton* or *NICER*, by analysing the QPO recently claimed by Yan et al. (2024) in the Seyfert galaxy NGC 1365. These authors found a significance of about 3.6σ by comparing periodogram peaks of light curves simulated from the continuum-fitted PSD. Using our method we have found instead a much lower significance, of about 1.7σ (or 91 per cent). From their analysis, it is unclear where this discrepancy in the estimate of the significance originates. Yan et al. (2024) report fitting the PSD with a bending power law, and then use this model to produce light-curve simulations to test the significance of the highest peak in the periodogram. However, there is no information regarding the fitting procedure, namely the statistic used to fit the periodogram and whether the appropriateness of the model was taken into account. It is also unclear whether the uncertainties on the model were taken into account in the estimation of the false-alarm probability and how the number of trials were considered. It is likely that a combination of these factors can explain the difference in our results.

We have then applied our method to the ULX in NGC 7793 P13, where slightly dissimilar superorbital periods had been claimed in the sparsely sampled *Swift*-UVOT and XRT light curves (Hu et al. 2017; Fürst et al. 2018). As can be seen from Fig. 12, we have found the X-ray period to be 65.6 ± 0.6 d, which is indeed significantly longer than the period in the UV ($P = 63.9 \pm 0.4$ d; Fig. 10). As stated above, using time domain methods allows to marginalize over the noise components and obtain accurate uncertainties on the parameters of the periodic component. Thus, we can support earlier assertions that the X-ray period is significantly longer (Hu et al. 2017; Fürst et al. 2018) than the optical/UV period. The high coherence $Q \gtrsim 300$ from the Lorentzian components suggest the period amplitude is stable throughout the segment, consistent with the long-term behaviour of the source (Fürst et al. 2021). Regarding the significance of the periodicities, while the third harmonics were marginally significant (at ~ 95 per cent and 91 per cent for the UV and X-ray light curves) given that these constitute the harmonics of the same periodicity, these significances likely underestimate the true significance as one could repeat the analysis tying the periods or considering the combined fit improvement provided by the three Lorentzians altogether, but this is beyond the scope of this work.

We have also examined the putative QPO claimed by Smith et al. (2020) in the AGN NGC 4945. When analysing the whole light curve, we have seen that we could not explain the data under a GP, which could indicate deviations from stationarity. Indeed, when analysing the last portion of the whole light curve (Fig. 13) we have seen the preferred rednoise kernel (SHO $_{Q=1/\sqrt{2}}$) differed from that obtained for the full light curve (DRW), which could support the non-stationarity of the process. Nevertheless, our analysis suggests there is little evidence for periodic variability when analysing the light curve as a whole. In the segment where Smith et al. (2020) reported the significance of the QPO to be the highest (Fig. 13 left panel), we have found the putative periodicity to have a significance of ~ 98.7 per cent (i.e. $\approx 2.5\sigma$), much lower than reported by Smith et al. (2020). The fact that we are able to produce simulations with comparable fit improvements as that observed in the data implies our test is well-calibrated and suggests our significance estimate is more plausible. The most obvious discrepancy is that we have correctly accounted for the presence of rednoise. Instead, Smith et al. (2020) relied on the analytical recipe provided by Horne & Baliunas (1986), which may be appropriate in cases where employing white noise as the null hypothesis—but see Frescura et al. (2008) for caveats on this method.

We note the significances quoted for NGC 4945 may be considered optimistic, as the selection of this segment seems driven by a ‘posteriori’ analysis of the data (a form of the stopping rule discussed in Vaughan 2010), rather than a data-driven decision (such as to avoid a gap in the light curve, e.g. Section 3.2.2). While this is beyond the scope of this work, one could in principle account for this by simulating light curves using the full length of the monitoring (either Fig. 14 or the entire RXTE history) and then selecting the segment that maximizes the likelihood ratio for each simulation.

Finally, we have examined the QPO claimed in the Blazar 0537-441 by Tripathi et al. (2024) using *TESS* data. While our analysis supports the presence of a QPO-like feature (at the $\sim 3.7\sigma$), the identified period is marginally consistent with the 6.5-d QPO reported by Tripathi et al. (2024), although no uncertainties on the claimed periodicity are provided by Tripathi et al. (2024). Differences in our results may be attributed to the different processing of the data and treatment of the underlying noise (see also Covino et al. 2022, for a similar situation).




Figure 17. GP modelling results of the RXTE data of the *TESS* light curve of the Blazar, B0537-441 shown in Fig. 16. (Top Left) Best-fitting Lorentzian + DRW model. (Top Right) PSD of the null hypothesis and alternative models. (Bottom left) ACF of the standardized residuals of the Lorentzian + DRW model. (Bottom right) Posterior parameters for the Lorentzian + DRW model. The MCMC run for 17 500 steps until convergence, of which 5390 were discarded for burn in. Symbols as per Fig. 7.




Figure 18. As per Fig. 15 but testing the QPO in the *TESS* light curve of the Blazar, B0537-441 using the DRW posteriors as the null hypothesis.

4 DISCUSSION AND CONCLUSIONS

The presence of red noise variability, ubiquitously found in accreting systems, makes the detection of periodicities challenging. On the one hand, most periodicity tests are derived for cases of Gaussian white noise, which makes the problem analytically tractable (e.g. Scargle 1982). On the other hand, the presence of red noise increases the likelihood of producing spurious features in the periodogram, particularly because the scatter in the power is proportional to the power itself (e.g. Vaughan 2005). When the data is unevenly sampled, the problem becomes even more profound as stochastic variability can easily be mistaken for periodic behaviour (cf. Vaughan et al. 2016).

Extrapolation of tests for periodicities against red noise-like variability was presented by Israel & Stella (1996) and Vaughan (2005), who proposed to capture the underlying broad-band noise using either a parametric (restricted to PSDs following a power law; Vaughan 2005) or non-parametric approach (Israel & Stella 1996) and use these estimates and associated uncertainties to derive the probability of obtaining a spurious signal in the periodogram

above a certain level. Vaughan (2010) expanded on previous work to model any arbitrary PSD shape using a Bayesian approach, which allowed for the inclusion of priors. All these techniques concerned the case where the time series is evenly sampled such that (as discussed in Section 2) the periodogram has some well-known statistical properties, which allows a well-defined likelihood (and other statistical tests such as goodness-of-fit) to be defined.

Here, we have provided a method for periodicity searches in the case of unevenly sampled data, where constraining the aperiodic variability is considerably more challenging and where it appears preferable to perform the fitting in the time domain (where the probability distribution is known and is generally Gaussian/Poissonian) using GP modelling. Here, we have exploited the known likelihood with well-established statistical techniques (Protassov et al. 2002) to estimate the significance of a putative (quasi)periodic component. In a similar manner to the regularly sampled case, the noise is inferred from the data, allowing a test for the presence of an additional component (e.g. a QPO) by building an empirical T_{LRT} distribution using the method proposed by Protassov et al. (2002). Given that the method is entirely data-driven, it is completely generalizable to any system/variability and even choice of mean function (which we have not exploited here).

If the PSD is of interest, this quantity can be accessed by Fourier-transforming the best-fitting GP kernel, rather than the data itself, thereby including the data (heteroscedastic) uncertainties in the final estimate. In doing so, frequency-distorting effects arising from irregular sampling are mitigated, while the data usage is maximized. There is additionally no requirement to rebin the data (so as long as there are enough counts for the data to be Gaussian distributed). A similar approach is discussed in Kelly et al. (2014) using CARMA, who also advocates for time-domain fitting. The recipe outlined in this work may equally be used employing CARMA kernels and will suffer from the same limitations we discuss below in Section 4.1. However, there seem to be certain advantages of using *celerite* over CARMA. *celerite* kernels have a more flexible form than CARMA ones (Foreman-Mackey et al. 2017). While the PSD of CARMA kernels are restricted to Lorentzian functions, steeper PSDs may be achieved using a single *celerite* kernel (the $\text{SHO}_{Q=1/\sqrt{2}}$ kernel being an example used here; Fig. 2), which in CARMA may not be straightforward to describe. From a computational point of view, in principle the computational is the same for both *celerite* and *carma* implementations Kelly et al. (2014) and Foreman-Mackey et al. (2017), scaling as $\mathcal{O}(NJ^2)$. However, Foreman-Mackey et al. (2017) showed that in practice *celerite* seems to perform better computationally.

There are several important advantages of using GP over Lomb–Scargle periodograms. As shown in Section 3.2.2, we cannot only access more accurately the underlying noise by performing model selection, but also marginalize over the noise parameters, therefore carrying over the full set of uncertainties into our determination of a candidate period’s frequency. Instead, both model selection and uncertainties are inaccessible when using Lomb–Scargle periodograms.

4.1 Limitations and caveats

Regardless of its power and improvement over traditional approaches, there remain several limitations of our method, arguably the most pressing being the computational time involved. The computational time of the GP modelling itself scales as N^3 , which can become intractable if several models need to be tested or for large data sets. Here, we have chosen to minimize the compute time using *celerite* (where the computational time-scales as

NJ^2) at the expense of flexibility, which may not be much more computationally expensive than the Lomb–Scargle periodogram.⁶ In addition to the model evaluation, there is the computational time required to perform the simulations for hypothesis testing. This problem is partially mitigated because the likelihood allows us to perform initial model selection (in our case through the ΔAICc) and filter out the most prominent cases. Therefore, only in cases with limited signal-to-noise ratio or where the ΔAICc does not provide sufficient indication (e.g. Graham et al. 2015), light-curve simulations may need to be performed, although having to rely on simulations for hypothesis testing equally applies to regularly sampled time series (e.g. Ashton & Middleton 2021).

Another common drawback of GP modelling is how to choose what kernels to test against the data. A straightforward approach to alleviate this problem is to simply stack basis functions until the minimum of the IC is found (Kelly et al. 2014; Foreman-Mackey et al. 2017; Zhang et al. 2023). Secondly, inspection of the standardized residuals can reveal trends indicative of the model not capturing the full variability (as also illustrated by the ACF).

We have also discussed how to identify cases where the GP might not be a good fit for the data (Section 3.2.3) either due to the process not being a GP or due to the assumption of stationarity not being fulfilled. In the former case, it is still unclear whether our method is still valid. Through our simulations (Section A), we have noted that when the light curves are produced using a lognormal PDF, the standardized residuals never show compliance with a standard Gaussian distribution, even if the input model parameters are well captured (see Section B). Thus, preliminary tests indicate that the variability is still well-captured even when the flux distribution is not Gaussian.

Note also that while the assumption of stationarity is another limitation of GPs, the same assumption is inherently made in standard periodograms. In fact, GPs are also more flexible on this regard, as the mean of the time series does not need to be constant. In any instance, in a similar vein as for dynamical periodograms (Kotze & Charles 2012), one could envision splitting the time series into approximately stationary segments and applying an independent GP modelling to each segment. Then the posteriors of a particular parameter (e.g. the period frequency P) could be examined to discern whether a given quantity is varying over the full observation baseline.

Lastly, compared to periodogram fitting, where any functional form may be employed, the fitting process in GP is restricted by the functional form of the kernels. This latter problem may be alleviated at the expense of computational cost, by using kernels outside *celerite* (e.g. Rasmussen & Williams 2006), or decomposing the power spectrum using basis functions as recently demonstrated in Lefkir et al. (2025).

ACKNOWLEDGEMENTS

This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester. This paper includes data collected by the *TESS* mission. Funding for the *TESS* mission is provided by the NASA’s Science Mission Directorate. This work has made use of light curves provided by the University of California, San Diego Center for Astrophysics and Space Sciences, X-ray Group (R.E. Rothschild, A.G. Markowitz, E.S. Rivers, and B.A. McKim), obtained at <http://cass.ucsd.edu/~rxteagn/>. The authors

⁶The fastest implementation in *astropy* scales as $\mathcal{O}(N\log(M))$ where M is the number of frequencies being evaluated.

acknowledge support by STFC through grant ST/V001000/1 and the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton. We would like to thank the referee for their interest in our manuscript and thoughtful comments that helped improve it. A. Gúrpide is also grateful to S. Vaughan, Z. Irving, E. Agol, W. Alston, and M. M. Ward for stimulating discussion on time domain analysis and Gaussian Processes modelling. Software: CORNER (Foreman-Mackey 2016), NIFTY-LS (Garrison et al. 2024), MIND_THE_GAPS (?).

DATA AVAILABILITY

All the data used in this paper is publicly available in the corresponding archives. The code used for this manuscript has also been made publicly available.

REFERENCES

- Akaike H., 1998, in Parzen E., Tanabe K., Kitagawa G. eds, *Selected Papers of Hirotugu Akaike*. Springer, New York, NY, p. 199, http://link.springer.com/10.1007/978-1-4612-1694-0_15
- Alston W. N., Markeviciute J., Kara E., Fabian A. C., Middleton M., 2014, *MNRAS*, 445, L16
- Ashton D. I., Middleton M. J., 2021, *MNRAS*, 501, 5478
- Barnard R., Trudolyubov S., Haswell C. A., Kolb U. C., Osborne J. P., Priedhorsky W. H., 2007, *AIP Conf. Proc.* Vol. 924, *The Multicolored Landscape of Compact Objects and Their Explosive Origins*. Am. Inst. Phys., New York, p. 691
- Belloni T., Psaltis D., Klis M. V. D., 2002, *ApJ*, 572, 392
- Bowman D. M., Dorn-Wallenstein T. Z., 2022, *A&A*, 668, A134
- Breidt E. et al., 2010, *MNRAS*, 403, 605
- Covino S., Landoni M., Sandrinelli A., Treves A., 2020, *ApJ*, 895, 122
- Covino S., Tobar F., Treves A., 2022, *MNRAS*, 513, 2841
- Done C., Madejski G. M., Mushotzky R. F., Turner T. J., Koyama K., Kunieda H., 1992, *ApJ*, 400, 138
- Done C., Madejski G. M., Życki P. T., Greenhill L. J., 2003, *ApJ*, 588, 763
- Emmanoulopoulos D., McHardy I. M., Papadakis I. E., 2013, *MNRAS*, 433, 907
- Evans P. A. et al., 2007, *A&A*, 469, 379
- Evans P. A. et al., 2009, *MNRAS*, 397, 1177
- Foreman-Mackey D., 2016, *J. Open Source Softw.*, 1, 24
- Foreman-Mackey D., Agol E., Ambikasaran S., Angus R., 2017, *AJ*, 154, 220
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Frescura F. A. M., Engelbrecht C. A., Frank B. S., 2008, *MNRAS*, 388, 1693
- Fürst F. et al., 2016, *ApJ*, 831, L14
- Fürst F. et al., 2018, *A&A*, 616, A186
- Fürst F. et al., 2021, *A&A*, 651, A75
- Garrison L. H., Foreman-Mackey D., Shih Y.-H., Barnett A., 2024, *RNAAS*, 8, 250
- Gehrels N. et al., 2004, *ApJ*, 611, 1005
- Gierliński M., Middleton M., Ward M., Done C., 2008, *Nature*, 455, 369
- González-Martín O., Vaughan S., 2012, *A&A*, 544, A80
- Graham M. J. et al., 2015, *Nature*, 518, 74
- Horne J. H., Baliunas S. L., 1986, *ApJ*, 302, 757
- Hu C.-P., Li K. L., Kong A. K. H., Ng C.-Y., Lin L. C.-C., 2017, *ApJ*, 835, L9
- Hübner M., Huppenkothen D., Lasky P. D., Inglis A. R., Ick C., Hogg D. W., 2022, *ApJ*, 936, 17
- Hurvich C. M., Tsai C.-L., 1989, *Biometrika*, 76, 297–307
- Ingram A., Done C., 2011, *MNRAS*, 415, 2323
- Israel G. L. et al., 2017, *MNRAS*, 466, L48
- Israel G. L., Stella L., 1996, *ApJ*, 468, 369
- Jenkins J. M. et al., 2016, in Chiozzi G., Guzman J. C., eds, *Proc. SPIE Conf. Ser.* Vol. 9913, *Software and Cyberinfrastructure for Astronomy IV*. SPIE, Bellingham, p. 1232
- Jiang N. et al., 2022, preprint([arXiv:2201.11633](https://arxiv.org/abs/2201.11633)) <http://arxiv.org/abs/2201.11633>
- Jordán A., Eyheramendy S., Buchner J., 2021, *Res. Notes Am. Astron. Soc.*, 5, 107
- Kastra J. S., 2017, *A&A*, 605, A51
- Kelly B. C., Becker A. C., Sobolewska M., Siemiginowska A., Uttley P., 2014, *ApJ*, 788, 33
- Kelly B. C., Sobolewska M. g., Siemiginowska A., 2011, *ApJ*, 730, 52
- Khan N., Middleton M. J., 2023, *MNRAS*, 524, 4302
- Kotze M. M., Charles P. A., 2012, *MNRAS*, 420, 1575
- Kraft R. P., Burrows D. N., Nousek J. A., 1991, *ApJ*, 374, 344
- Lefkir M., Vaughan S., Huppenkothen D., Uttley P., Anilkumar V., 2025, preprint([arXiv:2501.05886](https://ui.adsabs.harvard.edu/abs/2025arXiv250105886L)) <https://ui.adsabs.harvard.edu/abs/2025arXiv250105886L>
- Lomb N. R., 1976, *Astrophys. Space Sci.*, 39, 447
- Markowitz A., 2010, *ApJ*, 724, 26
- Motch C., Pakull M. W., Soria R., Grisé F., Pietrzyński G., 2014, *Nature*, 514, 198
- Mueller M., Madejski G., 2009, *ApJ*, 700, 243
- O'Sullivan N. K., Aigrain S., 2024, *MNRAS*, 531, 4181
- Pasham D. R. et al., 2019, *Science*, 363, 531
- Pasham D. R. et al., 2024, preprint([arXiv:2402.09689](https://arxiv.org/abs/2402.09689)) <http://arxiv.org/abs/2402.09689>
- Protassov R., van Dyk D. A., Connors A., Kashyap V. L., Siemiginowska A., 2002, *ApJ*, 571, 545
- Rasmussen C. E., Williams C. K. I., 2006, *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning, MIT Press, Cambridge
- Ricker G. R. et al., 2014, in Oschmann J. M., Jr., Clampin M., Fazio G. G., MacEwen H. A., eds, *Proc. SPIE Conf. Vol. 9143, Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave*. SPIE, Bellingham, p. 556
- Scargle J. D., 1982, *ApJ*, 263, 835
- Smith E., Robles R., Perlman E., 2020, *ApJ*, 902, 65
- Stella L., Arlandi E., Tagliaferri G., Israel G. L., 1994, preprint([arXiv:astro-ph/9411050](https://arxiv.org/abs/astro-ph/9411050)) <http://arxiv.org/abs/astro-ph/9411050>
- Timmer J., Koenig M., 1995, *A&A*, 300, 707
- Tripathi A., Smith K. L., Wiita P. J., Wagoner R. V., 2024, *MNRAS*, 528, 6608
- Uttley P., McHardy I. M., Papadakis I. E., 2002, *MNRAS*, 332, 231
- Uttley P., McHardy I. M., Vaughan S., 2005, *MNRAS*, 359, 345
- van der Klis M., 1988, *NATO Science Series C, Mathematical and Physical Sciences*. Advanced Science Institute
- VanderPlas J. T., 2018, *ApJS*, 236, 16
- Vasilopoulos G., Lander S. K., Koliopoulos F., Bailyn C. D., 2020, *MNRAS*, 491, 4949
- Vaughan S., 2005, *A&A*, 431, 391
- Vaughan S., 2010, *MNRAS*, 402, 307
- Vaughan S., Edelson R., Warwick R. S., Uttley P., 2003, *MNRAS*, 345, 1271
- Vaughan S., Uttley P., 2005, *MNRAS*, 362, 235
- Vaughan S., Uttley P., Markowitz A. G., Huppenkothen D., Middleton M. J., Alston W. N., Scargle J. D., Farr W. M., 2016, *MNRAS*, 461, 3145
- Waller L. A., Smith D., Childs J. E., Real L. A., 2003, *Ecol. Model.*, 164, 49
- Yan Y., Zhang P., Liu Q., Chang Z., Liu G., Yan J., Zeng X., 2024, *A&A*, 691, A7
- Zhang H., Yang S., Dai B., 2023, *ApJ*, 946, 52

APPENDIX A: LIGHT-CURVE SIMULATIONS

In order to simulate light curves from the kernel PSDs, we have used the method devised by Timmer & Koenig (1995) and Emmanoulopoulos, McHardy & Papadakis (2013). The method proposed by Emmanoulopoulos et al. (2013), as opposed to the method of Timmer & Koenig (1995), which by construction generates

Gaussian distributed data, can generate light curves with any flux PDF and PSD model. Therefore, along with having more realistic light curves matching more closely the real data and its uncertainties, the issue of negative fluxes is also avoided.

As stated above, we used the PSD from the GP kernel as input PSD for the method. For the PDF, we used either a Gaussian distribution (in which case we used the Timmer & Koenig (1995) algorithm), in cases where the observed data was consistent with being Gaussian distributed (as determined using a KS test) or log-normal distribution if this was not the case (this was only the case for the XRT data of P13; Section 3.2.2) – in which case we reverted to the Emmanoulopoulos et al. (2013) algorithm. In any instance, in practice we have found the PDF used to simulate the light curves did not affect the results. The mean of the distribution was set as for the observed data and the variance was determined by integrating the PSD kernel in frequency space from $1/T$, where T was the duration of the light curve, to a pseudo Nyquist frequency defined as $1/2 \text{ min}(\Delta t)$ where $\text{min}(\Delta t)$ indicates the smallest exposure time in the light curve. In this manner we obtained the *intrinsic* variance that generated the light curve prior to resampling, as opposed to the observed variance after resampling. The light curves were initially generated on a regular grid with a sampling $\text{min}(\Delta t)$ in order to introduce aliasing effects, and a few times longer (typically 5–20 depending on the light curve) than the real light-curve length to introduce red noise leakage. We then drew a random segment matching the duration of the real monitoring and re-sampled it with the same exposure times and cadence as the real observations. We finally added Poisson noise and estimated realistic uncertainties taking into account the background rates and exposure times for each individual snapshot. For the *Swift*-XRT, as for the real light curves, in cases where the simulated source counts dropped below 15, we used instead the posterior probability function derived by Kraft, Burrows & Nousek (1991), which is more suited for the low-count regime and prevents having negative counts.

APPENDIX B: GP MODELLING OF LOGNORMAL LIGHT CURVES

It is commonly observed that all accreting systems show a lognormal flux distribution, which translates into the universally observed linear relationship between the square root of their variance (the rms)

and their mean flux (the so-called linear ‘rms–flux’ relation; Uttley et al. 2005). The implication is that the process generating the flux variations must be multiplicative. A pertinent question to ask is therefore whether the light curves of accreting systems can be modelled as a GP, or at the very least, how the retrieved parameters are affected by the lognormality of the fluxes.

The skewness of a lognormal distribution with mean μ and variance σ^2 and with Gaussian parameters μ_L and σ_L is given by

$$\gamma = (e^{\sigma_L^2} + 2)\sqrt{e^{\sigma_L^2} - 1}, \quad (\text{B1})$$

where $\sigma_L^2 = \ln(1 + \frac{\sigma^2}{\mu^2}) = \ln(1 + F_{\text{var}}^2)$ and where F_{var} is the fractional rms variability amplitude ($F_{\text{var}} = \sqrt{\sigma^2/\mu^2}$; Vaughan et al. 2003). This implies that $\gamma = (F_{\text{var}}^2 + 3)F_{\text{var}}$ and so for low F_{var} , the lognormal tends to be symmetric and resembles a Gaussian distribution, but as F_{var} increases, the lognormal distribution becomes more skewed and deviates more strongly from Gaussianity (see also Uttley et al. 2005). This is shown in Fig. B1. Naively, we then may expect that GPs might be able to recover the input parameters more readily when the F_{var} is low. Similarly, Gaussian-like light curves will show no dependence (or a flat) rms–flux relationship, and as F_{var} increases the rms will show a linear dependence with flux (see also Uttley et al. 2005).

In order to inspect any biases introduced by modelling lognormal light curves by a GP, we have generated light curves possessing a lognormal flux distribution using the method proposed by Emmanoulopoulos et al. (2013; see Appendix A). The light curves were generated 10^6 s long, sampled every 10 s and with exposure times of 1 s, roughly matching the light curve of Cygnus X–1 presented by Uttley et al. (2005). The generative PSD was a DRW, where the bending time-scale was set to ~ 930 s to ensure it could be well-detected by the choice of sampling. The variance was adjusted to produce light curves with a varying degree of F_{var} while the mean count rate was fixed to 5000 ct s^{-1} . In particular, we have tested whether we could recover the input PSD parameters (ω_{bend} and variance σ^2) using GP modelling of lognormal light curves having $F_{\text{var}} = 0.1, 0.2, 0.4$, and 0.6 . The light curves were produced free of Poisson noise (and the uncertainties were set to zero in the fitting process) as we are only interested in examining any biases introduced by the lognormality of the fluxes.




Figure B1. (Left) Example of the PDF of the lognormal light curves as a function of F_{var} . As F_{var} increases, the lognormal deviates more strongly from Gaussianity. (Right) RMS–flux relationship of light curves simulated having a lognormal distribution. These were averaged over the ensemble of the 1000 light curves, by averaging the mean and rms calculated using 5000 s segments.




Figure B2. Best-fitting ω_{bend} and variance σ^2 of an ensemble of 5000 light curves generated with varying flux PDFs (as indicated in the legend) and F_{var} . As can be seen, the recovered parameters are in agreement with the input values, regardless of the PDF used to generated the light curves or the F_{var} .




Figure B3. Best-fitting ω_{bend} and variance σ^2 of an ensemble of 5000 light curves generated from a lognormal PDF and varying F_{var} , but now including Poisson noise and uncertainties in the fitting. The blue colour and orange colour shows the results for varying mean ($\mu = 1000$ and 5000 ct s^{-1} , respectively). As can be seen, the bias in the recovered parameters increases with F_{var} , but even at the highest F_{var} the bias remains small. This suggest GPs have broad applicability.

As can be seen from Fig. B1, the generated light curves naturally follow the observed linear rms–flux relationship. In particular, we can see that for the lowest γ (or equivalently, F_{var}), the relationship is flat, as expected for a Gaussian distribution. As F_{var} increases, we see the linear rms–flux relationship is recovered.

Fig. B2 shows histograms of the recovered ω_{bend} and σ^2 for an ensemble of 5000 light curves. As can be seen, we do not observe any deviation from the input parameters in the recovered parameters, regardless of F_{var} , despite the light curves following the universal linear rms–flux relationship.

As a further test, we now incorporate Poisson noise and take into account the (Poissonian) uncertainties in the fitting process. We run this test for the lognormal light curves only, as for high F_{var} ($\gtrsim 0.3$)

the Gaussian distributed light curves produce negative counts due to the distribution not being strictly positively defined. Fig. B3 shows the histogram of the recovered parameters for an ensemble of 5000 light curves with a lognormal distribution, varying the F_{var} and for two different mean values of $\mu = 1000$ and 5000 ct s^{-1} , respectively. As expected, deviations from the input parameters are stronger as F_{var} increases. For $\mu = 1000 \text{ ct s}^{-1}$ and the largest F_{var} values, deviations are at most of the order of ~ 7 per cent, affecting more strongly ω_{bend} . However, we can see that for the higher mean count-rate case ($\mu = 5000 \text{ ct s}^{-1}$), even at the highest F_{var} of 0.6, biases remain below the order of ~ 2 per cent. This suggests that most of the biases we see for $\mu = 1000 \text{ ct s}^{-1}$ are due to Poisson statistics, and that lognormality of the flux has little impact on the recovered

parameters. Moreover, since F_{var} values higher than $\gtrsim 0.5$ are rarely observed in AGNs or X-ray binaries (e.g. Breed et al. 2010), this experiment suggests that there is broad applicability of GPs for the recovery of the variability processes in accreting sources.

APPENDIX C: MCMC SAMPLING

Here, we describe the process for the derivation of the best-fitting parameters and their posteriors. These were found by first minimizing the negative log likelihood function using the L-BFGS algorithm. We then applied a small nudge to the best-fitting parameters and used MCMC methods to sample the posterior running 32 independent chains (or walkers) using the `emcee` python library (Foreman-Mackey et al. 2013). More specifically, after the fitting process the walkers were distributed around the best-fitting parameters by drawing from a Gaussian with mean equal to the best-fitting values and standard deviation equal to 10 per cent of their values.

We adopted fairly uninformative (uniform) priors. Limits on the frequencies of the period and the aperiodic kernel time-scales were set based on data constraints. The shortest time-scale was set by a pseudo Nyquist frequency ($f_{\text{nyq}} = 1/2 < \Delta t >$ where $< \Delta t >$ was

the mean cadence of the light curve). The maximum allowed time-scale was T for the aperiodic kernels and $T/2$ for the periodic kernel, with T being the light-curve length. We further imposed $Q \gtrsim 3/2$; (Fig. 2) for the Lorentzian component to force this component to always represent a periodic signal and avoid degeneracy with the aperiodic kernels. The upper bound of Q was effectively unconstrained to allow for cases where the amplitude of the oscillation is not seen to decay.

In order to ensure convergence, the MCMC sampler was run until (a) the number of steps reached 100 times the integrated autocorrelation time (τ), which was estimated on the fly every 800 samples, and (b) τ changed less than 1 per cent compared to the previous estimate. We then discarded the first $30 \times \tau$ number of samples (the burn in) and thinned the chains by $\tau/2$ to build the posterior probability density function. We additionally inspected the chains for stationarity and compared the variances within each chain to the variance between chains following Vaughan (2010) (and references therein).

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