

# Randomization Inference of Periodicity in Unequally Spaced Time Series with Application to Exoplanet Detection

Panos Toulis

University of Chicago  
Booth School of Business  
[ptoulis@chicagobooth.edu](mailto:ptoulis@chicagobooth.edu)

Jacob Bean

University of Chicago  
Department of Astronomy and Astrophysics  
[jbean@astro.uchicago.edu](mailto:jbean@astro.uchicago.edu)

## Abstract

The estimation of periodicity is a fundamental task in many scientific areas of study. Existing methods rely on theoretical assumptions that the observation times have equal or i.i.d. spacings, and that common estimators, such as the periodogram peak, are consistent and asymptotically normal. In practice, however, these assumptions are unrealistic as observation times usually exhibit deterministic patterns —e.g., the nightly observation cycle in astronomy— that imprint nuisance periodicities in the data. These nuisance signals also affect the finite-sample distribution of estimators, which can substantially deviate from normality. Here, we propose a set identification method, fusing ideas from randomization inference and partial identification. In particular, we develop a sharp test for any periodicity value, and then invert the test to build a confidence set. This approach is appropriate here because the construction of confidence sets does not rely on assumptions of regular or well-behaved asymptotics. Notably, our inference is valid in finite samples when our method is fully implemented, while it can be asymptotically valid under an approximate implementation designed to ease computation. Empirically, we validate our method in exoplanet detection using radial velocity data. In this context, our method correctly identifies the periodicity of the confirmed exoplanets in our sample. For some other, yet unconfirmed detections, we show that the statistical evidence is weak, which illustrates the failure of traditional statistical techniques. Last but not least, our method offers a constructive way to resolve these identification issues via improved observation designs. In exoplanet detection, these designs suggest meaningful improvements in identifying periodicity even when a moderate amount of randomization is introduced in scheduling radial velocity measurements.

**Keywords.** Periodicity, Unequally Spaced Time Series, Set Identification, Randomization Inference, Exoplanet Detection.

# 1 Introduction

The identification of periodicity in time series data is a fundamental task across scientific fields. Applications include the classical study of periodic mass extinctions (Raup, 1986), gene expression cycles (Cho et al., 1998), circadian rhythms in biology (Costa et al., 2013), and seasonal oscillations in epidemiology (Murray and Morse, 2011). Typically, observations times in these applications are unequally spaced, which introduces identification issues that standard statistical methods, be it frequentist or Bayesian, cannot handle. When the hidden true periodicity cannot be point-identified, it is an open question how statistical inference should proceed, and whether it is possible to anticipate such identification issues through better observation designs.

Our work in this paper is motivated by the analysis of radial velocity data in exoplanet detection. The idea behind this method is to look for periodic variations in velocity data from a star. When these variations cannot be explained by known signals (e.g., rotation or magnetic activity of the star), they are attributed to an (unseen) exoplanet orbiting around the star. In recent years, the deployment of an extensive network of high-resolution, Earth-based observatories has led to astonishing exoplanet detections, including “51 Pegasi b” (Mayor and Queloz, 1995) —awarded the 2019 Nobel Prize in Physics— and even a potential discovery in our immediate stellar neighborhood of Proxima Centauri (Anglada-Escudé et al., 2016).

The statistical methodology behind these exoplanet detections proceeds in two main steps. The first step is to test for any sign of periodicity. If this test succeeds, the second step is to estimate the hidden periodicity. To test periodicity, statistical methods rely on the so-called *periodogram* of the data, which essentially quantifies the power of the Fourier spectrum of the data across different periods (Schuster, 1898). For instance, the classical method of Fisher (1929) relies on a statistic defined as the ratio of the periodogram peak over its average value across all periods, which has a known distribution under a harmonic model with i.i.d. normal errors. This method has since been improved in terms of power (Siegel, 1980; Bølviken, 1983; Chiu, 1989), and has been extended to more general hypotheses (Juditsky et al., 2015) and sparse alternatives (Cai et al., 2016). For adaptations of these tests to astronomical applications see also (Baluev, 2008, 2013; Delisle et al., 2020; Linnell Nemeč and Nemeč, 1985), and the concept of “false alarm probability”. In general, testing for periodicity presents no big challenges as current tests are robust and their properties are well-understood.

However, the problem of actually estimating the hidden periodicity is more subtle. For one, standard parametric results assume away identification issues, typically by assuming that the true periodicity lies in a constrained space allowing point-identification with an increasing number of observations (Walker, 1971; Hannan, 1971, 1973; Quinn and Fernandes, 1991; Quinn and Hannan, 2001); this is also related to the problem of aliasing and the Nyquist limit in statistical Fourier analysis (Bloomfield, 2004, Section 2.5). Furthermore, standard methods that rely on Fourier analysis also require equally spaced observation times. In practice, however, observations are usually unequally spaced for practical reasons, such as the difficulty in gathering fossils to study biological extinctions or the logistics of scheduling astronomical observations with large telescopes.

To address the problem of unequally spaced observation times, [Lomb \(1976\)](#) and [Scargle \(1982\)](#) independently developed a modification of Schuster’s periodogram. Even though the Lomb–Scargle periodogram was motivated by Fourier analysis, it has an intuitive statistical interpretation as it is equivalent to fitting the data with a harmonic model through least squares. As such, the Lomb–Scargle periodogram currently occupies a central place in exoplanet detection methodology ([VanderPlas, 2018](#), Section 1.1), where the focus is on assessing whether the periodogram peak is statistically significant, typically via extreme value theory.

A common flaw in current methodology, however, is that a significance result on the periodogram peak is interpreted to mean that the true unknown period is also “somewhere near” the peak. The implicit underlying assumption is, of course, that the finite-sample distribution of the periodogram peak is near the true period, and also normal in shape. This is unrealistic, however, because the likelihood in the harmonic models employed in practice is highly non-smooth and multimodal with respect to the underlying period. As such, without strong, unrealistic assumptions on the observation times, the periodogram peak is not “well-behaved” neither in finite samples nor asymptotically, in the sense that it may deviate significantly from normality, and may even be inconsistent. See also ([Hall and Li, 2006](#)) for similar discussions on the challenges of estimating periodicity through periodograms. Technically, one fundamental problem is that the separability conditions required for consistency of M-estimators ([Van der Vaart, 2000](#), Theorem 5.7), such as the peak of a periodogram, generally do not hold under arbitrary observation designs. Below, we present some examples to illustrate these issues and motivate our work.

**Example 1** (Synthetic data). Let  $t_i = i + 0.05U_i$ ,  $i = 1, \dots, 100$ , be the observation times for data  $y_i = 1.5 \cos(2\pi t_i/\sqrt{2}) + \varepsilon_i$ , where  $U_i \sim \text{Unif}[-1, 1]$  and  $\varepsilon_i \sim N(0, 1)$  i.i.d. The true periodicity is  $\theta^* = \sqrt{2} \approx 1.414$ . The data and its generalized Lomb–Scargle periodogram ([Zechmeister and Kürster, 2009](#)) are shown in [Figure 1](#) (left). In our sample, the peak is at  $\hat{\theta}_n = 3.417$  and is distinct from the peak around the true period. There are also multiple peaks associated with the true period harmonics. Using Fisher’s test we can reject the null of no periodicity ( $p$ -value is 0).

Inference on  $\theta^*$ , however, is challenging as we need to know the sampling distribution of  $\hat{\theta}_n$ . As argued above, normality of  $\hat{\theta}_n$  —as it is commonly assumed in practice— is not plausible without strong assumptions on the observation times, especially in finite samples. Indeed, in this example, we see that the simulated sampling distribution of  $\hat{\theta}_n$  is far from normal, and instead looks discrete with three distinct modes, as shown in [Figure 1](#) (right). The empirical density over 1,000 simulations is shown below (rounded to two significant digits):

$\hat{\theta}_n$	0.37	0.59	0.77	1.41	1.42	3.40	3.41	3.42	3.43
$\text{pr}(\hat{\theta}_n \theta_0)$	0.4%	15.3%	0.2%	55.6%	10.9%	1.6%	7.6%	7.8%	0.6%

We see that the sampling distribution of our estimator is non-regular. This complicates inference as it is unclear, for example, how to construct confidence intervals for the true period. In this paper, we propose a method that can transform such sampling distributions into valid confidence sets that can still cover the true period with a target coverage probability.

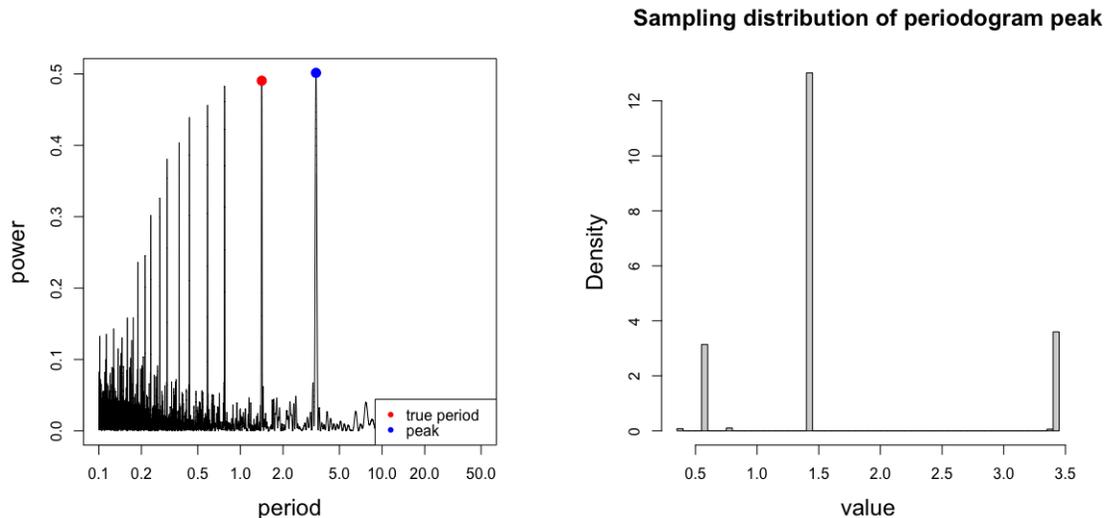


Figure 1: *Left:* Periodogram from one dataset from model of Example 1. The dataset was seeded to illustrate the identification issue. *Right:* Sampling distribution of the periodogram peak from the same model over 1,000 replications. We see that the periodogram peak is not normal, which complicates inference.

**Example 2** (Extinction cycles). In a series of influential articles, [Raup and Sepkoski \(1986\)](#) and [Raup \(1986\)](#) analyzed data from biological extinction events, and attempted to estimate their periodicity. Their approach did not use the periodogram, but it modeled the event times at the most prominent extinction events directly for a best-matching period. Significance was determined through a resampling scheme where all the recorded extinction times (not just the prominent ones) were permuted. Through this method, [Raup and Sepkoski \(1986\)](#) determined that there exists a biological extinction cycle of roughly 26 million years. However, this result may be an artifact of the small sample size (only 40 observations), and the unequal spacing of observation times. Indeed, [Stigler and Wagner \(1987\)](#) showed that the same signal at 26 million years appears even in simulated data where the observation times are fixed but the actual data are random. We analyze the extinction data in Section 4.1 using the method proposed in this paper, and present a statistically principled complement to the heuristic analysis in ([Stigler and Wagner, 1987](#)).

**Example 3** (Planet around  $\alpha$  Centauri B?). In a widely-publicized article, [Dumusque et al. \(2012\)](#) reported the detection of an Earth-mass planet around the star  $\alpha$  Centauri B at a period of 3.24 days. This detection was startling because  $\alpha$  Centauri is the closest stellar system to our own system, implying that habitable exoplanets may be ubiquitous. However, the detection has not yet been confirmed, and there have been some early objections to it ([Hatzes, 2013](#)). Notably, [Rajpaul et al. \(2015b\)](#) followed a methodology similar to ([Stigler and Wagner, 1987](#)) discussed in Example 2, and showed that the 3.24-day signal also appears as a local peak in various simulations with quasi-random data and fixed observation times. In this paper, we analyze these data in Section 5 and show that the statistical evidence is weak in the sense that the observed data could be explained by many other candidate signals. Moreover, in Section 6.1 we discuss how the observation design could be improved to identify the 3.24-day signal, assuming that the signal is true.

We emphasize that neither bootstrap- nor Bayesian-based methods have so far addressed this type of identification issues. On the one hand, bootstrap procedures require random observation times and convergence in distribution for  $\hat{\theta}_n$ , which requires strong, unrealistic assumptions on the design. Moreover, the small-sample performance of bootstrap procedures with time series is theoretically not well understood. This is particularly important in exoplanet detection as we usually have to work with only a few hundred observations. Bayesian methods, on the other hand, have the problem of prior specification. For example, even through they are usually referred to as “non-informative”, uniform priors give preference to parameter regions that not only have high likelihood but are also wide. In effect, this sweeps the identification problem “under the rug” as narrow parameter regions that cause identification problems (such as the region around  $\theta = 0.59$  in Example 1) are washed away; see also (Hall and Yin, 2003, Section 1) for similar arguments that invoke number theory, and (Gelman and Yao, 2020, Section 3) for a more general discussion on how uniform priors can unintentionally lead to sharp inference. In a Bayesian context, this issue also spills over to model selection through Bayes factors in the sense that model decisions based on Bayes factors may strongly depend on features that are esoteric to the selected models (e.g., size or structure of the parameter space); see also (Gelman and Yao, 2020, Section 6) for details.

## 1.1 Our contribution

In this paper, we propose a set identification method to infer hidden periodicity. The main idea is to develop a sharp test for a null hypothesis on a given periodicity value,  $H_0 : \theta_0 = \theta$ , and then invert the test to build a confidence set,  $\Theta_{1-\alpha}$ . Similar to a confidence interval, a confidence set is constructed so that it contains the true parameter with a pre-specified coverage probability:

$$\text{pr}(\theta^* \in \Theta_{1-\alpha}) \geq 1 - \alpha. \tag{1}$$

In contrast to confidence intervals, confidence sets are not necessarily contiguous. The use of confidence sets is therefore more suitable here since it does not require asymptotic normality, or even consistency, of periodogram peaks. Furthermore, the “topology” of  $\Theta_{1-\alpha}$  can reveal identification issues based on, say, how wide or narrow the confidence set is.

We propose two constructions for  $\Theta_{1-\alpha}$ . The first construction is idealized and yields a finite-sample valid procedure, where (1) holds exactly in the sample (Theorem 1). This construction is computationally challenging, in practice, and so we also propose an alternative construction, where (1) holds asymptotically (Theorem 3). This approximate construction may be conservative but in empirical applications we show this not to be a major concern. Our approach borrows ideas from randomization inference in statistics (Lehmann and Romano, 2006; Rosenbaum et al., 2010), where inference is based on structural assumptions (e.g., symmetries) on the data, and is therefore robust (Toulis, 2019). Our approach is also related to partial identification in econometrics (Manski, 2003; Romano and Shaikh, 2008; Tamer, 2010), where test inversion is used to build confidence sets for unknown parameters that cannot be point-identified.

While several heuristic techniques have been used to diagnose identification issues in periodicity estimation (see Examples 1 and 2), our paper is the first to develop a valid, principled statistical procedure for inference when such issues are present. The proposed procedure is described in detail in Section 3.1, and can in fact work with many different choices of the statistic (other than periodogram peaks), and can be parallelized for computational efficiency. Moreover, our procedure does not require normality or i.i.d. errors, and can seamlessly work with equally or unequally spaced data. We note that while some important non-parametric methods with unequally spaced data have recently been proposed, they still require observation designs with equal spacings (Sun et al., 2012; Vogt and Linton, 2014) or i.i.d. spacings between observation times (Hall et al., 2000; Hall, 2008; Hall and Yin, 2003), and emphasize asymptotic consistency over finite-sample performance. That said, our procedure could leverage these methods in the choice of the test statistic. This could combine the best of both worlds, namely the finite-sample validity and robustness from our randomization-based procedure, and the flexibility and power from nonparametric methods. See Remark 3.2 and Section 6.2 for more discussion.

Empirically, we illustrate our method in several examples from exoplanet detection using radial velocity data. In this context, we show that our method gives sharp inference on the confirmed exoplanets in our sample, namely 51 Pegasi b (Mayor and Queloz, 1995, “51Pegb”) and Gliese 436 b (Butler et al., 2004, “GJ436b”), showing that our inference is not conservative. For other, yet unconfirmed, exoplanets in our sample our method gives more mixed results. For example, we show that the statistical evidence is notably weak for a candidate discovery around  $\alpha$  Centauri B (Dumusque et al., 2012) in the sense that the confidence set,  $\Theta_{1-\alpha}$ , contains a large range of signals, some possibly non-planetary, including nuisance 1-day signals that are usually associated with the nightly cycle of astronomical observations. Testing another widely-publicized discovery in Proxima Centauri (Anglada-Escudé et al., 2016), we find that the discovery seems to be robust apart from a nuisance signal at 0.916 days that cannot be rejected at the 1% level. Importantly, our method suggests ways to improve the observation designs for sharper inference. Specifically, we quantify how much additional randomness is needed in the observation times or, alternatively, how many more observations are needed to precisely identify a candidate signal that is suspected to be true. Following this approach, we propose designs that could address the aforementioned identification issues, and also help with future scientific discoveries.

## 2 Problem Setup and Notation

We observe data  $Y^n = \{y_1, \dots, y_n\} \in \mathbb{R}^n$  at observation times  $T^n = \{t_1, t_2, \dots, t_n\}$ . For example, in exoplanet detection, the data could represent measured radial velocity of a star at a time  $t$  expressed as a barycentric Julian date. We decompose the data model as  $y_i = y^p(t_i) + \varepsilon(t_i)$ , where  $y^p(t)$  is an unobserved periodic, non-stochastic function with unknown period  $\theta^*$ , such that

$$y^p(t + \theta^*) = y^p(t), \quad (t \in \mathbb{R}).$$

The terms  $\varepsilon(t_i)$  correspond to the unobserved errors. We use  $Y^{n,P} = (y^P(t_1), \dots, y^P(t_n))$  and  $\varepsilon^n = (\varepsilon(t_1), \dots, \varepsilon(t_n))$  to denote the (unobserved) vector values of these two components.

Our main goal is to do inference on  $\theta^*$ , the true unknown period. We will need the following additional assumptions. First, we require a consistent estimator of  $y^P$ , which could either be parametric or nonparametric. For simplicity, in the following section we start with the simple but popular trigonometric regression model for  $y^P$ . Later, in Section 6.2, we also discuss a nonparametric approach. Second, we assume that the errors,  $\varepsilon(t)$ , form a stochastic process satisfying the following invariance property. For any observation times  $T^n = \{t_1, \dots, t_n\}$ , with  $n$  finite, there exists a known algebraic group  $\mathcal{G}^n$  of  $n \times n$  matrices such that

$$g \cdot \varepsilon^n \stackrel{d}{=} \varepsilon^n \mid T^n, \quad (g \in \mathcal{G}^n). \quad (\text{A1})$$

This formulation follows the framework of randomization tests (Lehmann and Romano, 2006, Chapter 15), where testing is based on structural assumptions rather than analytical. To keep things simple, for the rest of this paper we assume that  $\mathcal{G}^n = \mathbb{D}^n$ , the set of  $n \times n$  diagonal matrices with  $\pm 1$  in the diagonal. This definition implies that the regression errors,  $\varepsilon(t)$ , are assumed to have a sign symmetric distribution. As such, our inference method not only covers Gaussian errors, which are frequently assumed in practice, but also extends to more general model classes where the error distribution is unspecified, or even varies with  $t$ , as long as the symmetry in (A1) holds.

**Remark 2.1** (Generality of (A1)). Assumption (A1) can also allow inference with complex error structures. Suppose, for instance, that  $\varepsilon(t)$  is a Gaussian process (GP) with covariance matrix  $C$ . Then, we could define  $\mathcal{G}^n = \{g' = C^{1/2}gC^{-1/2} : g \in \mathbb{D}^n\}$ . Of course, this requires  $C$  to be known, or be consistently estimated. To accomplish that, one idea would be to combine our approach with recent GP-based methods on radial velocity data (Rajpaul et al., 2015a; Damasso and Del Sordo, 2017; Faria et al., 2016, 2020). We leave this for future work. ■

### 3 Main Method

Here, we follow a parametric approach to model the periodic component,  $y^P$ . A nonparametric approach is discussed later, in Section 6.2. All proofs can be found in the Supplement.

We start with the simple harmonic model

$$y^P(t) \equiv y^P(t; \theta^*, \psi^*) = \psi_1^* + \psi_2^* \cos(2\pi t/\theta^*) + \psi_3^* \sin(2\pi t/\theta^*), \quad (\text{A2})$$

where  $\theta^* \in \Theta \subseteq \mathbb{R}^+$  is the true unknown period, and  $\psi^* = (\psi_1^*, \psi_2^*, \psi_3^*) \in \Psi \subseteq \mathbb{R}^3$  are unknown nuisance parameters, where  $\Psi$  is convex and compact. We note that we use this model for simplicity. For instance, we could work with, say, a Keplerian model that includes additional parameters to describe various patterns of motion (e.g., eccentricity) of a body relative to another (Wright and Howard, 2009). All that is required for our theory to work is for the model of choice to be consistently estimated conditional on the true period,  $\theta^*$ .

To measure the fit of  $y^P$  we use the simple squared loss,

$$L(\theta, \psi | Y^n, T^n) = \sum_{i=1}^n [y_i - y^P(t_i | \theta, \psi)]^2 / \sigma_i^2,$$

where the precisions  $\sigma_i^2$  are known from instrument specifications, as is typical in exoplanet detection. Note that the loss function (i.e., negative log-likelihood under Gaussian errors) is non-convex and non-smooth with respect to  $\theta$ , and so standard statistical inference cannot be applied.

To proceed, let  $L_{0n}(Y^n, T^n) = \sum_{i=1}^n (y_i - \bar{y})^2 / \sigma_i^2$  denote the baseline loss, where  $\bar{y} = (1/n) \sum_{i=1}^n y_i$ . Then, the *oracle* generalized periodogram,  $A_n : \Theta \rightarrow \mathbb{R}^+$  can be defined as follows:

$$A_n(\theta | Y^n, T^n, \psi) = \frac{L_{0n}(Y^n, T^n) - L(\theta, \psi | Y^n, T^n)}{L_{0n}(Y^n, T^n)}. \quad (2)$$

In practice,  $\psi$  is not the target of inference, and so we use the profile likelihood

$$\hat{A}_n(\theta | Y^n, T^n) \equiv A_n(\theta | Y^n, T^n, \hat{\psi}_\theta), \text{ where } \hat{\psi}_\theta = \arg \min_{\psi \in \Psi} L(\theta, \psi | Y^n, T^n). \quad (3)$$

$\hat{A}_n$  is known as the generalized periodogram, and is easy to calculate as  $\hat{\psi}_\theta$  is known in closed form (Zechmeister and Kürster, 2009). Our first goal is to test a global null hypothesis of the form:

$$H_0^{\text{full}} : \theta^* = \theta_0, \psi^* = \psi_0. \quad (4)$$

This hypothesis is not particularly interesting because, ultimately, we are only interested in the unknown period  $\theta^*$ . However, testing for  $H_0^{\text{full}}$  will serve as a good starting point for testing and inference on  $\theta^*$ . Of course, to test any hypothesis we need a test statistic,  $s_n(Y^n, T^n) \in \mathbb{R}$ . Many choices are possible here, and so we defer a detailed discussion about this choice to Section 3.1. Here, we simply define

$$s_n(Y^n, T^n) = \max_{\theta \in \Theta} A_n(\theta | Y^n, T^n, \psi_0) - A_n(\theta_0 | Y^n, T^n, \psi_0).$$

The idea behind this choice of the test statistic is that  $s_n$  will tend to take small values when  $\theta^* = \theta_0$ , as we expect the periodogram to peak around  $\theta_0$  under  $H_0^{\text{full}}$ ; when  $H_0^{\text{full}}$  is not true we expect  $s_n$  to take large values as the periodogram will peak in periods different from  $\theta_0$ , so that the null can be rejected.

Let  $s_{\text{obs}} = s_n(Y^n, T^n)$  denote the observed value of the test statistic in the sample. To construct the null distribution of  $s_n$ , we simulate data as follows:

$$Y^{n,(i)} = Y_0^{n,P} + G^{(i)} \cdot (Y^n - Y_0^{n,P}),$$

where  $Y_0^{n,P} = [y^P(t_1; \theta_0, \psi_0), \dots, y^P(t_n; \theta_0, \psi_0)]$  are the periodic component values, which are known under the null hypothesis;  $G^{(i)} \sim \text{Unif}(\mathcal{G}^n)$  is a random error transformation from  $\mathcal{G}^n$ .

Then, a  $p$ -value for  $H_0^{\text{full}}$  is

$$\text{pval}(\theta_0, \psi_0) = E\{s_n(Y^{n,(i)}, T^n) \geq s_{\text{obs}} \mid Y^n, T^n\}, \quad (5)$$

where the randomness is with respect to  $G^{(i)}$  while  $Y^n, T^n$  are fixed. The following result shows that this  $p$ -value is valid in finite samples for testing  $H_0^{\text{full}}$ .

**Theorem 1.** *Suppose that Assumptions (A1) and (A2) hold. Then, the  $p$ -value in (5) is exact in finite samples under  $H_0^{\text{full}}$ , that is, for any finite  $n > 0$ ,*

$$\text{pr} \left\{ \text{pval}(\theta_0, \psi_0) \leq \alpha \mid H_0^{\text{full}} \right\} = \alpha.$$

In most settings,  $\psi^*$  is a nuisance parameter, and so we only want to test for  $\theta^*$ :

$$H_0 : \theta^* = \theta_0. \quad (6)$$

In this context, the result in Theorem 1 can still be used, in principle, to build a confidence set for  $\theta^*$  using the following construction:

$$\Theta_{1-\alpha} = \left\{ \theta \in \Theta : \max_{\psi \in \Psi} \text{pval}(\theta, \psi) > \alpha \right\}. \quad (7)$$

**Theorem 2.** *Suppose that Assumptions (A1) and (A2) hold. Then,  $\Theta_{1-\alpha}$  is a finite-sample valid  $100(1 - \alpha)\%$  confidence set for  $\theta^*$ ; i.e., for any finite  $n > 0$ ,*

$$\text{pr}(\theta^* \in \Theta_{1-\alpha} \mid H_0) \geq 1 - \alpha.$$

While  $\Theta_{1-\alpha}$  is a valid confidence set, its construction is challenging due to the optimization problem in (7) that requires maximizing the  $p$ -value over the entire nuisance parameter space,  $\Psi$ . This is especially challenging when  $\Psi$  is multidimensional. To address this computational issue, a feasible procedure is to estimate  $\psi^*$  under the null hypothesis  $H_0$  —for example, through the least-squares estimate  $\hat{\psi}_{\theta_0}$ — and then plug this estimate into (7). We note that this approach works because the estimator  $\hat{\psi}_{\theta_0}$  is consistent for  $\psi^*$  under the standard assumptions of least squares since, under fixed  $\theta$ , the loss function in the harmonic model is smooth and convex with respect to  $\Psi$ . In contrast, estimating  $\theta^*$  in a similar fashion (e.g., through least squares) does not work, precisely because of the identification issues discussed in Section 1.

To test  $H_0$  in a computationally manageable way, we therefore have to simply re-define the test components using the generalized periodogram (3) as follows:

$$s_n(Y^n, T^n) = \max_{\theta \in \Theta} \hat{A}_n(\theta \mid Y^n, T^n) - \hat{A}_n(\theta_0 \mid Y^n, T^n), \quad (8)$$

$$\hat{Y}_0^{\text{P}} = [y^{\text{P}}(t_1; \theta_0, \hat{\psi}_{\theta_0}), \dots, y^{\text{P}}(t_n; \theta_0, \hat{\psi}_{\theta_0})], \quad (9)$$

$$Y^{n,(i)} = \hat{Y}_0^{\text{P}} + G^{(i)} \cdot [Y^n - \hat{Y}_0^{\text{P}}]. \quad (10)$$

The  $p$ -value for  $H_0$  is then defined as

$$\widehat{\text{pval}}(\theta_0) = E\{s_n(Y^{n,(i)}, T^n) \geq s_{\text{obs}} \mid Y^n, T^n\}, \quad (11)$$

where, again, the randomness is with respect to  $G^{(i)}$  while the data  $(T^n, Y^n)$  are fixed. The following construction for the confidence set of  $\theta^*$  is valid asymptotically:

$$\hat{\Theta}_{1-\alpha} = \left\{ \theta \in \Theta : \widehat{\text{pval}}(\theta) > \alpha \right\}. \quad (12)$$

**Theorem 3.** *Suppose that Assumptions (A1) and (A2) hold, and that  $\hat{\psi}_{\theta_0} \xrightarrow{P} \psi^*$  under  $H_0$ . Then,  $\hat{\Theta}_{1-\alpha}$  is an asymptotically valid  $100(1 - \alpha)\%$  confidence set for  $\theta^*$ ; i.e., as  $n$  increases*

$$\text{pr}(\theta^* \in \hat{\Theta}_{1-\alpha} \mid H_0) \geq 1 - \alpha + o_P(1).$$

### 3.1 Concrete procedure and practical considerations

In summary, we propose the following procedure to estimate the true periodicity.

#### PROCEDURE 1

1. Choose a grid of possible period values,  $\Theta$  that contains  $\theta^*$  w.p. 1. Set  $\hat{\Theta}_{1-\alpha} \leftarrow \emptyset$ .
2. Obtain data  $(Y^n, T^n)$ , possibly after removing known stellar signals, e.g., rotational periods, magnetic cycles, etc. (Feigelson and Babu, 2012, Section 11.3.4).
3. **For** all  $\theta_0 \in \Theta$  **do**:
  - (i) Estimate the nuisance parameters,  $\hat{\psi}_{\theta_0} = \arg \min_{\psi \in \Psi} L(\theta_0, \psi \mid Y^n, T^n)$ , through weighted least squares.
  - (ii) Calculate the observed value,  $s_{\text{obs}}$ , of the test statistic defined in (8).
  - (iii) With fixed  $T^n$ , sample new data,  $Y^{n,(i)}$ , where  $i = 1, \dots, R$  for some fixed  $R$ , according to definitions (9) and (10).
  - (iv) Using the samples from 3(iii), calculate the  $p$ -value (11), and if it exceeds  $\alpha$  then include  $\theta_0$  in the confidence set; i.e., set  $\hat{\Theta}_{1-\alpha} \leftarrow \hat{\Theta}_{1-\alpha} \cup \{\theta_0\}$  if  $\widehat{\text{pval}}(\theta_0) > \alpha$ .
4. **Return**  $\hat{\Theta}_{1-\alpha}$  as the  $100(1 - \alpha)\%$  confidence set of  $\theta^*$ .

In words, Procedure 1 considers all possible candidate period signals in  $\Theta$ . For every candidate signal,  $\theta_0$ , the procedure, first, estimates the harmonic model assuming that the true period is  $\theta_0$ , and then it randomly flips the signs of the residuals to generate new synthetic data (i.e., assumes symmetric distribution for the errors). The observed value of the test statistic is then compared against the test statistic values in the synthetic data, and the resulting  $p$ -value determines whether  $\theta_0$  should be in the confidence set or not.

We end this section with a few remarks on some practical aspects of implementing Procedure 1.

**Remark 3.1** (Choice of  $\Theta$ ). The choice of  $\Theta$  can be crucial as the periodogram is generally sensitive to the average spacing in  $\Theta$ . In general, it is a good idea to choose spacings that are smaller than the expected width of the major periodogram peaks, and also induce a spacing in  $\Theta$  that is uniform in the log scale (VanderPlas, 2018, Section 6.1). Another useful technique is to keep refining  $\Theta$  until the shape of the observed periodogram and its major peaks are stable. ■

**Remark 3.2** (Choice of test statistic,  $s_n$ ). Procedure 1 is valid for any choice of the test statistic. Our choice of test statistic in (8) performs well under several controlled simulations (see Section 4), but other options are certainly possible. For instance, Fisher’s classical statistic to test periodicity can be defined as  $s_n = \max_{\theta \in \Theta} \hat{A}_n(\theta) / \bar{A}_n$ , where  $\bar{A}_n = |\Theta|^{-1} \sum_{\theta} A_n(\theta)$  is the average periodogram value; here, we suppressed the dependence on  $Y^n, T^n$  to ease notation. Improvements using a trimmed mean in place of  $\bar{A}_n$  have also been suggested (Bølviken, 1983; Siegel, 1980; Damsleth and Spjøtvoll, 1982), including sparse alternatives (Cai et al., 2016). See also (McSweeney, 2006) for some numerical comparisons. ■

**Remark 3.3** (Power). The power of Procedure 1 depends on how sensitive the chosen test statistic,  $s_n$ , is in detecting violations of the null hypothesis. In general, statistics that lead to better estimators of  $\theta^*$  will also lead to more powerful tests, and thus more narrow confidence sets. In this paper, we don’t explore the power of our randomization tests theoretically to maintain the focus on our application—we refer the reader to (Toulis, 2019) and (Dobriban, 2021) for such theoretical investigations on the power of randomization tests, including those applied in a regression context like ours. Some optimal results in testing periodicity are known but require strong assumptions on the errors. For instance, Fisher’s classical statistic (described above) is optimal only under i.i.d. normal errors in the harmonic model. Empirically, our choice of test statistic in (8) appears to be powerful compared to alternatives—see Section 4 and Section 5 for details. ■

**Remark 3.4** (Computation). The complexity of Procedure 1 can be described, *prima facie*, as  $O(|\Theta|^2 \cdot R \cdot C)$ , where  $C$  is the average computation time for a single weighted least-squares fit. For some ballpark numbers, we can take  $|\Theta| = 10^4$ ,  $R = 10^3$ , and  $C = 50\mu\text{s}$  for an analysis on a conventional laptop of a time series with 200 observation times, using an implementation in R. The full computation requires a total of 1,388 hrs. of wall clock time (approx. 58 days). However, several reductions of computation time are possible. First, Procedure 1 can be fully parallelized since step 3 can be executed independently for each  $\theta_0 \in \Theta$ . With a moderate-sized cluster of 100 computing nodes, the wall clock time thus drops to 14 hrs. Moreover, again in step 3, there is no need to consider all values in  $\Theta$  but only a proportion of the periodogram peaks. In our applications, for example, we only consider those peaks of the periodogram that are at least 20% as high as the highest peak. This leads to a complexity  $O(\gamma|\Theta|^2 \cdot R \cdot C)$  with  $\gamma$  ranging between 0.1%-3% in our applications. As such, the required computation drops dramatically to roughly 30 mins. of wall clock time. Indeed, in our application, we are able to ramp up our samples to  $R = 100,000$  and still finish all exoplanet analyses within a couple of hours using a cluster of 400 computing nodes. Further significant gains are possible if the method is implemented in a more efficient language than R, such as C++ or Python. ■

## 4 Numerical Examples

### 4.1 Extinction data

Here, we return to Example 2 on estimating periodicity in biological extinction events. The original data are shown in the left subplot of Figure 2 below. Raup and Sepkoski (1986) argued that there is a 26 million year periodicity based on fitting a simple periodic model on the largest extinction events. Formally, for any given  $\theta \in \Theta$ , they fit a model  $t_i = t_0 + k_\theta \theta$  for those timepoints that correspond to the eight largest peaks; here  $t_0, k_\theta$  are free parameters and  $k_\theta$  is an integer. Using this model, Raup and Sepkoski (1986) calculated which  $\theta$  yields the smallest least squares fit, and then used a permutation-based test to assess significance.

The periodogram of the data is shown in the right subplot of Figure 2. Indeed, using a set  $\Theta$  ranging between 12.5-62.5 million years, as suggested by Raup and Sepkoski (1986), we find the largest periodogram peak to be around 26 million years. However, we also see additional peaks around 30 million years and 35 million years. Using Procedure 1, none of these values can be rejected as the true periodicity ( $p$ -values are 0.56 and 0.6, respectively).<sup>1</sup> Moreover, a similar analysis with a larger range for  $\Theta$  (not shown in the figure) reveals multiple peaks —at periods less than 0.5 million years— that are even larger than the peaks shown in Figure 2, and are accepted in the confidence set. We conclude that there are significant identification issues with this dataset, and so the discovery cannot be confirmed. These issues were also highlighted by Stigler and Wagner (1987), albeit using a heuristic approach.

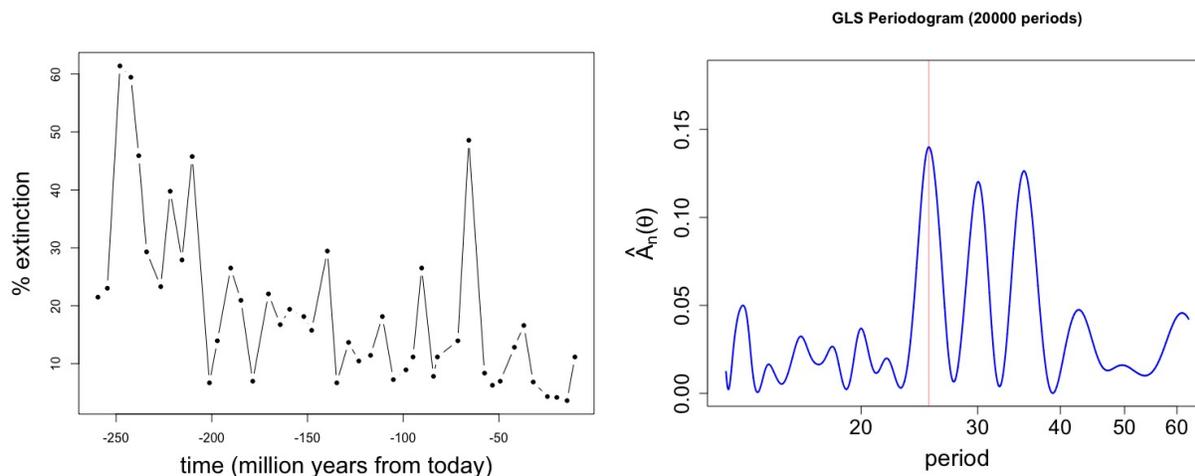


Figure 2: *Left*: Original extinction data of Raup and Sepkoski (1986). The data include 40 observations of extinction events measured as % of estimated biomass lost during the event. *Right*: The periodogram of the extinction data in the range 12.5-62.5 million years, as suggested by Raup and Sepkoski (1986).

<sup>1</sup>All data and code for the numerical examples of this section, including the exoplanet detection applications of the next sections, are provided with the Supplementary Code. An online repository with R code to replicate all results and figures in the following sections can be found at <https://github.com/ptoulis/ri-exoplanet-detection>. A Python implementation can be found at <https://github.com/afeinstein20/periodic-ri>.

## 4.2 Simulated data with realistic observation designs

Here, we revisit Example 1 but with a more realistic design for the observation times. Specifically, we build upon the setup of Hall and Li (2006, Section 3.2) and define the data model:

$$y_i = 1 - \cos(2\pi t_i/\theta^*) + N(0, \sigma^2), \quad (i = 1, \dots, n). \quad (13)$$

Here, the design of observation times relies on a density  $f_1$  on the interval  $[0, 1]$  that represents a full 24-hour day, with  $[0, 0.5]$  denoting day time and  $[0.5, 1]$  denoting night time ( $t = .75$  is midnight, and so on). Adapting the setup of Hall and Li (2006), we consider three choices for  $f_1$ :

- (i)  $f_1(t) = 1(.5 \leq t \leq 1)$ ;
- (ii)  $f_1(t) = \max\{0, -\pi \sin(2\pi t)\}1(0 \leq t \leq 1)$ ;
- (iii)  $f_1(t) = 1(|t - 0.75| \leq 0.0208)$ ;

The full observation design is then defined as:

$$f(T^n) \propto \prod_{i=1}^n \left( \sum_{j=0}^{n-1} f_1(t_i - j\theta_{\text{obs}}^*) \right). \quad (14)$$

We set  $\theta^* = \sqrt{2}$  as the true signal period with  $\sigma^2 = 1.5^2$ , and  $\theta_{\text{obs}}^* = 1$  for the observation period. Thus, the particular form of  $f(T^n)$  in (14) represents a density that has daily periodicity in the interval  $[0, n]$ , which represents  $n$  days of observations. The definitions of  $f_1$  are meant to represent certain recognizable classes of observation designs. Specifically, in design (i) the observation times  $t_i$  are sampled uniformly during night time. In design (ii), we define a more realistic scenario where observations happen only at night time, but are more probable during midnight. Finally, in design (iii) the observations are uniformly sampled within one hour window around midnight.

Figure 3 shows the sampling distribution of  $\hat{\theta}_n = \arg \max_{\theta \in \Theta} \hat{A}_n(\theta \mid Y^n, T^n, \hat{\psi}_\theta)$ , the periodogram peak, where  $(T^n, Y^n)$  are sampled, respectively, according to (14) and (13), and  $\hat{A}_n$  is the generalized Lomb–Scargle periodogram in (3). As discussed in Section 1, we see that the sampling distribution is far from normal across all designs, and is better described as an irregular, discrete distribution on  $\Theta$ . The amount of irregularity depends on the observation design. Under Design (i), for example,  $\hat{\theta}_n$  has converged to a point distribution around  $\theta^*$  at  $n = 500$  data points, but there is significant irregularity even at  $n = 100$ . Under Design (iii), the situation is much worse since the distribution of  $\hat{\theta}_n$  is irregular even at  $n = 500$  with three distinct modes similar to the mode around  $\theta^*$ . This example illustrates one of the key motivations of this paper, which is that standard asymptotic normal theory is inappropriate for inference of periodicity in time series due to the irregularity of periodic models.

Figure 4 shows how to address this problem using the identification method proposed in this paper. Specifically, following Example 1, we seed a dataset  $(T^n, Y^n)$  such that estimation is problematic in the sense that the observed periodogram peak is not at the true value,  $\theta^*$  (see Figure 4,

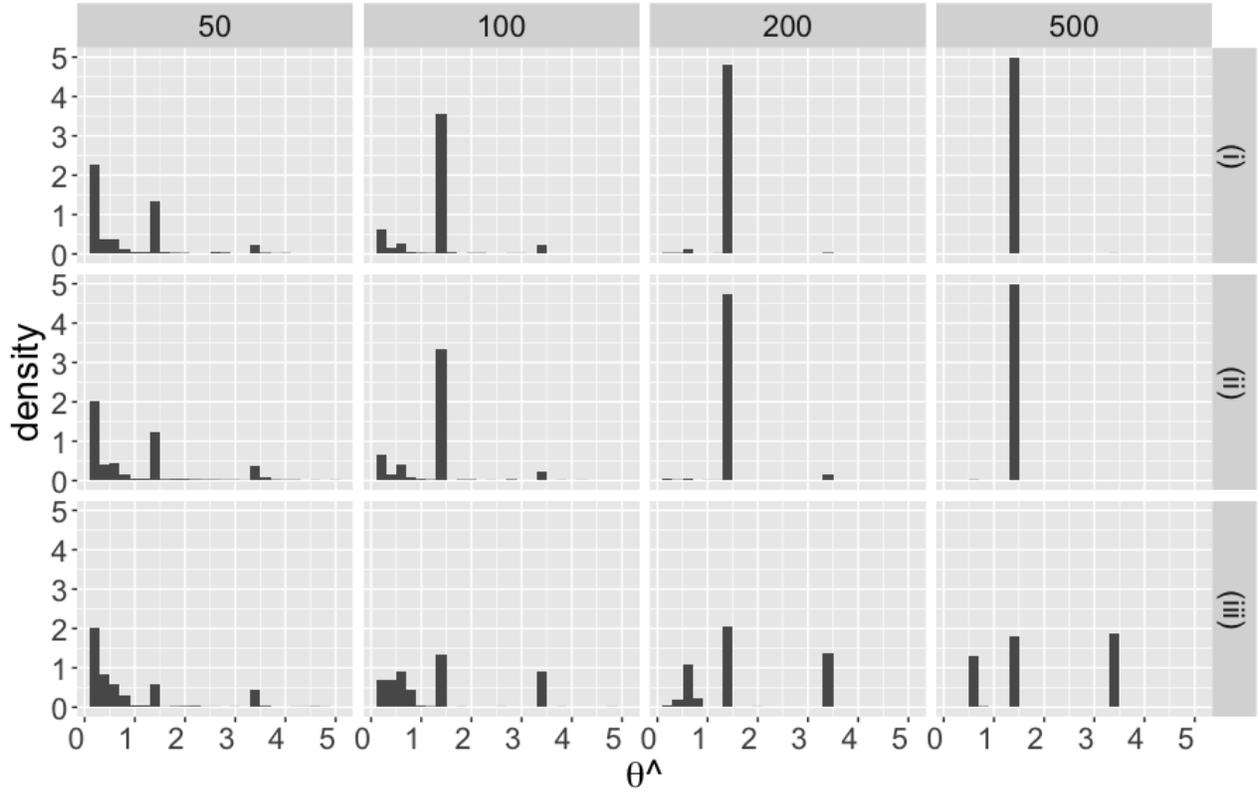


Figure 3: Sampling distribution of periodogram peak (approximated via 1,000 samples) in the model described by (14) and (13).

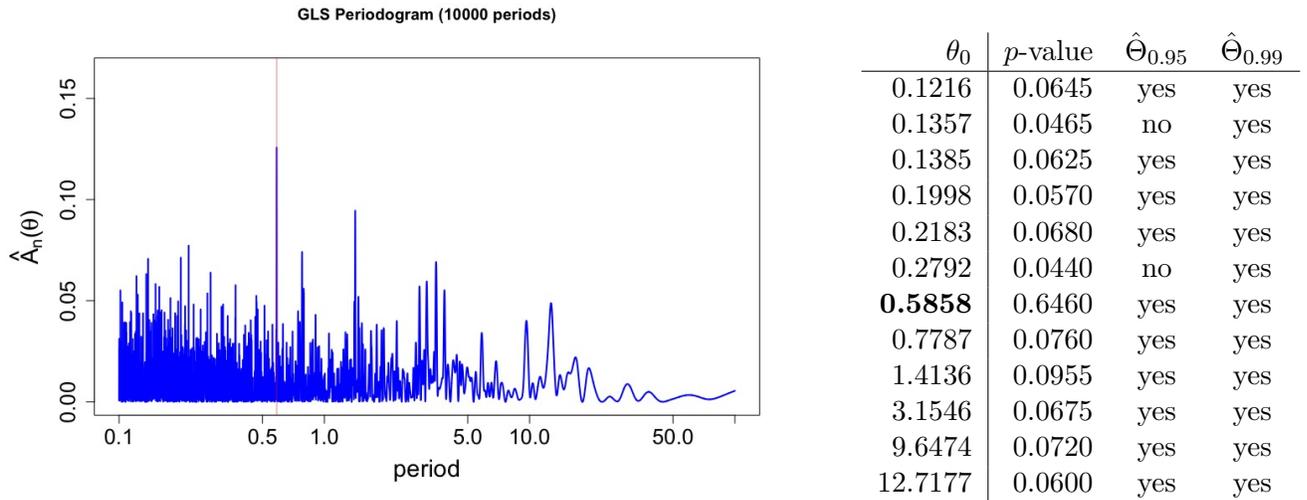


Figure 4: *Left*: Periodogram for a problematic dataset sampled according to (14) and (13), where the peak  $\hat{\theta}_n = 0.586$  is not at the true period  $\theta^* = 1.414$ . *Right*: Inference of periodicity based on Procedure 1. The table shows the  $p$ -values for the hypothesis  $H_0 : \theta^* = \theta_0$  for values of  $\theta_0$  that correspond to the highest peaks of the periodogram shown on the left.

left) —in this dataset  $\hat{\theta}_n = 0.586$ . The table on the right of Figure 4 shows the  $p$ -values for testing  $H_0 : \theta^* = \theta_0$  obtained from Procedure 1. For simplicity, we only show results for the 10 highest peaks in the observed periodogram. We see that all these values cannot be rejected at the 1% level and so they are included in the 99% confidence set,  $\hat{\Theta}_{0.99}$  (see column “ $\hat{\Theta}_{0.99}$ ” in the table). By looking at the range of values included in this confidence set, we conclude that there are severe identification issues in the data. This was not immediately evident from the periodogram as the peak in Figure 4 is fairly distinct and far exceeds the 1% false alarm probability threshold (not shown above). We also note that, beyond the mode at  $\theta_0 = 0.586$ , our proposed method puts the highest confidence at the value  $\theta_0 = 1.4136$  ( $p$ -value = 0.096), which is close to the true signal.

Finally, we check whether the finite-sample coverage of  $\hat{\Theta}_{1-\alpha}$  used in Procedure 1 matches the asymptotic nominal coverage, as derived in Theorem 3. For this simulation we consider all aforementioned designs ( $f_1$ ), and also multiple sample sizes ( $n$ ), and observation times restricted to a 6-month observation period. For each setting we calculate how often  $\hat{\Theta}_{.95}$  covers the true period,  $\theta^*$ , over 4,000 replications for each ( $f_1, n$ )-pair. The results are shown in the table below.

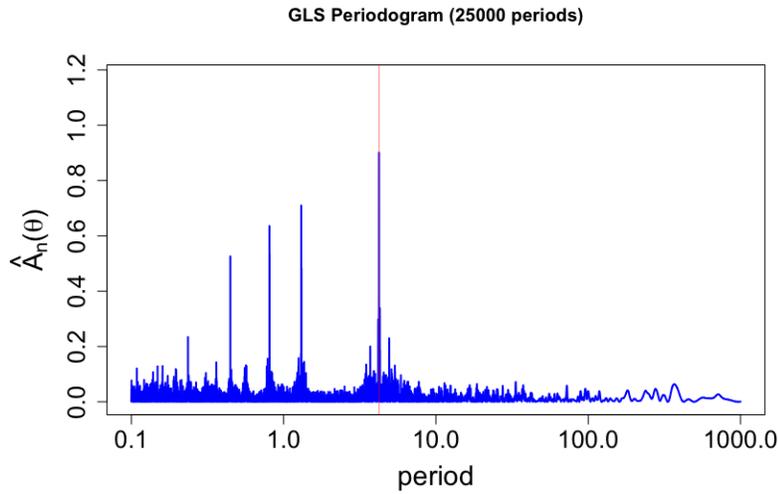
Design, $f_1$	Sample size, $n$				
	25	50	100	250	500
(i)	97.97%	97.32%	96.05%	94.45%	95.27%
(ii)	97.82%	97.65%	95.42%	94.50%	97.50%
(iii)	97.72%	98.85%	96.95%	93.40%	93.60%

We note that coverage gets closer to the nominal value for larger sample sizes, which is consistent with Theorem 3. Moreover, we see that better observation designs generally lead to better coverage.

## 5 Application: Exoplanet Detection

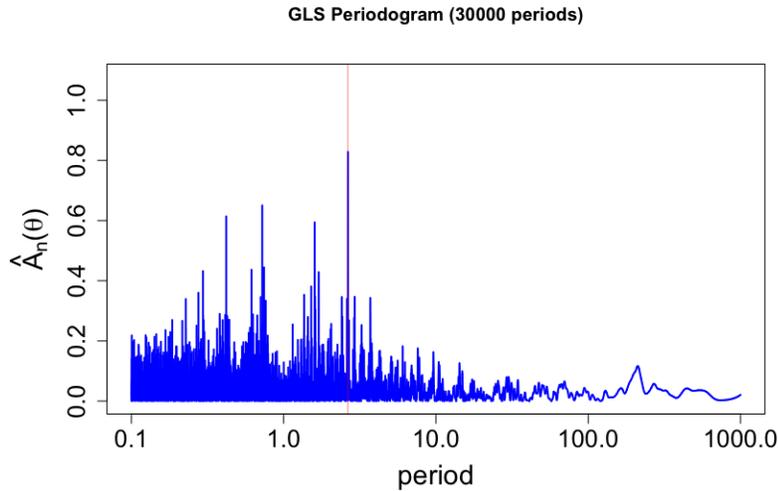
In this section, we use our method to analyze radial velocity data from four known star systems. Specifically, we first analyze data on two exoplanets, namely 51 Pegasi b (Mayor and Queloz, 1995, “51Pegb”) and Gliese 436 b (Butler et al., 2004, “GJ436b”). These exoplanets have been confirmed independently by other means, including direct observations of transit. We also analyze data from two candidate exoplanets potentially orbiting the stars  $\alpha$  Centauri B (Dumusque et al., 2012), and Proxima Centauri (Anglada-Escudé et al., 2016), respectively.

We start our analysis with the confirmed exoplanets in our sample. The periodogram of the stellar radial velocity data for 51Pegb is depicted in Figure 5. This shows a clear peak around  $\hat{\theta}_n = 4.23$  days with additional distinct peaks around 1-day and 0.5-day periods. Our method correctly shows no identifiability issues here. All periods corresponding to the highest peaks are strongly rejected ( $p$ -values numerically zero) except for  $\theta_0 = 4.23$  (observed peak). Next, we analyze data on GJ436b. Compared to 51Pegb, the periodogram for GJ436b shows a peak around  $\hat{\theta}_n = 2.6441$ , but this peak is less distinct than other peaks. Our method again correctly shows that the true signal can be easily identified. As in 51Pegb, all periods corresponding to the highest peaks are strongly rejected ( $p$ -value = 0) except for the observed peak at  $\theta_0 = 2.64$ .



$\theta_0$	$p$ -value	$\hat{\Theta}_{0.95}$	$\hat{\Theta}_{0.99}$
0.3085	0.0000	no	no
0.5662	0.0000	no	no
0.8069	0.0000	no	no
0.8089	0.0000	no	no
0.8295	0.0000	no	no
1.3047	0.0000	no	no
1.3095	0.0000	no	no
3.7033	0.0000	no	no
4.1807	0.0000	no	no
<b>4.2311</b>	1.0000	yes	yes
4.2821	0.0000	no	no
4.9331	0.0000	no	no

Figure 5: *Left*: Periodogram of radial velocity on exoplanet “51Pegb”. Here,  $\Theta = \{0.1, \dots, 1000\}$  is split uniformly in the log-space so that  $|\Theta| = 25,000$ . *Right*: Inference of periodicity of 51Pegb based on Procedure 1. The table shows the  $p$ -values for the hypothesis  $H_0 : \theta^* = \theta_0$  for values of  $\theta_0$  that correspond to high peaks of the periodogram shown on the left. We see that there are no identification issues, as the 4.23-day signal is the only one accepted in the confidence sets.



$\theta_0$	$p$ -value	$\hat{\Theta}_{0.95}$	$\hat{\Theta}_{0.99}$
0.4200	0.0000	no	no
0.6155	0.0000	no	no
0.7067	0.0000	no	no
0.7438	0.0000	no	no
1.3641	0.0000	no	no
1.5187	0.0000	no	no
1.6013	0.0000	no	no
1.6086	0.0000	no	no
1.7008	0.0000	no	no
2.4103	0.0000	no	no
<b>2.6441</b>	1.0000	yes	yes
3.7092	0.0000	no	no

Figure 6: *Left*: Periodogram of radial velocity on exoplanet “GJ436b”. Here,  $\Theta = \{0.1, \dots, 1000\}$  is split uniformly in the log-space so that  $|\Theta| = 30,000$ . *Right*: Inference of periodicity of GJ436b based on Procedure 1. The table shows the  $p$ -values for the hypothesis  $H_0 : \theta^* = \theta_0$  for values of  $\theta_0$  that correspond to high peaks of the periodogram shown on the left. We see that there are no identification issues, as the 2.64-day signal is the only one accepted in the confidence sets.

Our results are mixed for the other, yet unconfirmed exoplanets. For the candidate exoplanet orbiting  $\alpha$  Centauri B (Dumusque et al., 2012), in particular, our analysis —summarized in Figure 7— shows severe identification issues. The periodogram shows a peak around  $\hat{\theta}_n = 3.23$  days but it is not as distinct from the other peaks. Moreover, there is significant variation in the periodogram. Indeed, our method reveals that many periods corresponding to the highest peaks cannot be rejected at the 1% level, and there is even a period ( $\theta_0 = 0.7622$ ) that cannot be rejected at the 5% level. These results indicate that additional statistical evidence is necessary to claim that the particular exoplanet discovery is plausible. As in the case of the extinction data, we note that earlier work has raised concerns about this discovery as well (Rajpaul et al., 2015b, e.g.), albeit using only heuristic methods similar to those suggested by Stigler and Wagner (1987).

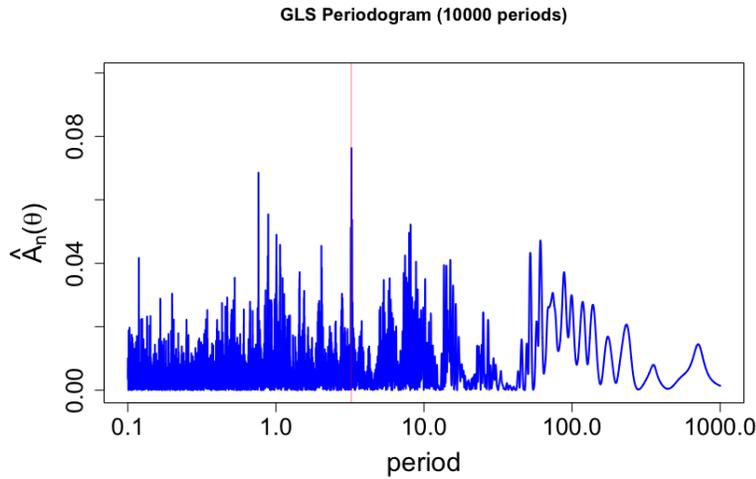
Finally, we discuss the data for a recent discovery of an exoplanet ostensibly orbiting Proxima Centauri (Anglada-Escudé et al., 2016). We use a recent release of 63 observations of Proxima Centauri from the ESPRESSO spectrograph (Mascareño et al., 2020) that has increased precision. The analysis and results are shown in Figure 8. The periodogram shows a peak around  $\hat{\theta}_n = 11.17$  days but also a substantial peak around the 1-day period signal, possibly associated with the observation design. Indeed, our method shows that out of all highest peaks in the periodogram, only  $\theta_0 = 11.1739$  and  $\theta_0 = 0.9164$  cannot be rejected at the 1% level. Overall, while the finding of Anglada-Escudé et al. (2016) seems to be robust, there is still a small chance of observing the observed periodogram pattern under the particular observation times, at least with respect to those aspects of the periodogram that are captured by our test statistic. More statistical evidence is therefore needed to statistically validate this discovery, possibly through the use of a slightly more randomized observation design. We discuss such designs in the following section.

## 6 Discussion and Extensions

### 6.1 Observation design

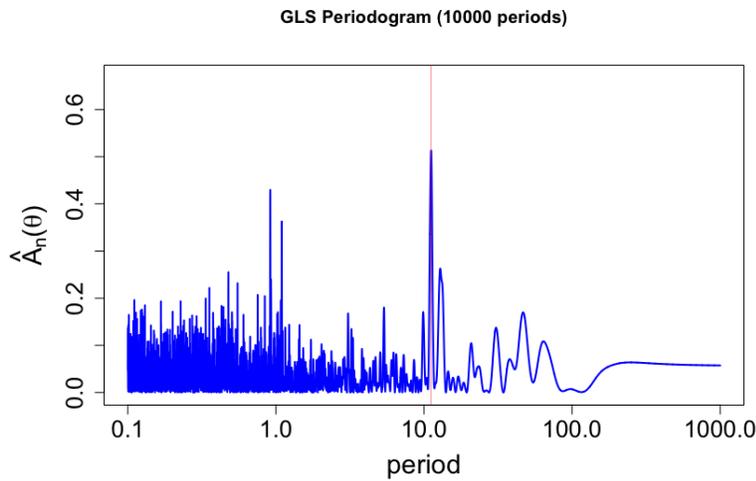
The importance of observation times in identifying a periodic signal is well understood (Feigelson and Babu, 2012; VanderPlas, 2018; Ivezić et al., 2014). It may therefore come as a surprise that there is little systematic work on optimizing observation designs. One exception is perhaps (Loredo, 2003), which discusses a Bayesian approach to sequential experimental design of observations—see also (Wynn, 1984; Kiefer, 1975; Mitchell, 2000; Smucker et al., 2018) for a statistical background. However, this approach also relies on regularity conditions for the underlying data model, and so it does not account for the identification issues discussed in Section 1.

One way to address these issues is to use the identification method developed in this paper. The idea is simply to synthesize data under alternative designs, and then pick the design that yields the most precise confidence set,  $\hat{\Theta}_{1-\alpha}$ . Conceptually, the main challenge with this approach is how to define an appropriate space for these alternative designs in a way that is both practical and computationally efficient. Given a choice for the design space, we can then simulate synthetic data and choose from those designs that yield narrow confidence sets.



$\theta_0$	$p$ -value	$\hat{\Theta}_{0.95}$	$\hat{\Theta}_{0.99}$
0.7622	0.0705	yes	yes
0.8882	0.0271	no	yes
1.0086	0.0174	no	yes
1.0678	0.0079	no	no
2.0292	0.0122	no	yes
3.2074	0.0163	no	yes
<b>3.2371</b>	1.0000	yes	yes
3.2670	0.0178	no	yes
7.9394	0.0116	no	yes
8.1169	0.0175	no	yes
52.2242	0.0121	no	yes
61.1334	0.0226	no	yes

Figure 7: *Left*: Periodogram of radial velocity on candidate exoplanet orbiting  $\alpha$  Centauri B (Dumusque et al., 2012). Here,  $\Theta = \{0.1, \dots, 1000\}$  is split uniformly in the log-space, so that  $|\Theta| = 10,000$ . *Right*: Inference of periodicity of the exoplanet based on Procedure 1. The table shows the  $p$ -values for the hypothesis  $H_0 : \theta^* = \theta_0$  for values of  $\theta_0$  that correspond to high peaks of the periodogram shown on the left. We see that there are severe identification issues as several signals other than the periodogram peak are accepted in the confidence sets.



$\theta_0$	$p$ -value	$\hat{\Theta}_{0.95}$	$\hat{\Theta}_{0.99}$
0.1106	0.0007	no	no
0.3355	0.0022	no	no
0.3552	0.0055	no	no
0.4778	0.0025	no	no
0.5512	0.0047	no	no
0.7532	0.0052	no	no
0.8412	0.0059	no	no
0.9164	0.0173	no	yes
0.9266	0.0005	no	no
1.0957	0.0080	no	no
<b>11.1739</b>	1.0000	yes	yes
12.8769	0.0006	no	no

Figure 8: *Left*: Periodogram of radial velocity on candidate exoplanet Proxima Centauri b (Anglada-Escudé et al., 2016). Here,  $\Theta = \{0.1, \dots, 1000\}$  split regularly in the log-space, so that  $|\Theta| = 10,000$ . *Right*: Inference of periodicity of the exoplanet based on Procedure 1. The table shows the  $p$ -values for the hypothesis  $H_0 : \theta^* = \theta_0$  for values of  $\theta_0$  that correspond to high peaks of the periodogram shown on the left. We see that there are no severe identification issues, and the detection appears to be robust. There is only one signal (at 0.9164 days) that cannot be rejected at the 1% level, and thus causes concern. An improved observation design could decisively rule out this nuisance signal; see Section 6.1 for more details.

In practice, we usually face a scenario where we suspect the peak  $\hat{\theta}_n$  in the periodogram to be the true signal, but there is also another “nuisance peak”, say  $\theta^{\text{nuis.}}$ . Typically, this nuisance peak appears near the 1-day signal as a by-product of a daily observation design. For example, in the ESPRESSO data discussed in Section 5 and Figure 8, the problem is that even though there is very strong evidence for a signal at  $\hat{\theta}_n = 11.17$  days, there is also a signal at  $\theta^{\text{nuis.}} = 0.9164$  that cannot be rejected at the 1% level ( $p$ -value = 0.017). The observed peak could statistically be an artifact of the daily observation design. The key question is thus:

“What kind of observation designs should we use to accurately identify the signal at  $\hat{\theta}_n$ , assuming this signal to be the true one?”

To address this problem, we propose the following high-level procedure.

1. Select a space  $\mathcal{D}$  for observation designs and a tolerance threshold  $0 < \epsilon < |\hat{\theta}_n - \theta^{\text{nuis.}}|$ .
2. **For** each design  $D$  in the design space, **do**:
  - (a) Assume  $\theta^* = \hat{\theta}_n$  as the true signal, and then calculate  $\hat{\psi}_{\theta^*}$  as defined in (3).
  - (b) **For**  $r = 1, \dots, R$ , **do**:
    - (i) Generate synthetic data  $(\tilde{T}^{n'}, \tilde{Y}^{n'})$  according to the design and  $(\theta^*, \hat{\psi}_{\theta^*})$  from step 2(a), where we may introduce additional (synthetic) observations, such that  $n' \geq n$ .
    - (ii) Calculate  $\hat{\Theta}_{1-\alpha}$  based on synthetic data  $(\tilde{T}^{n'}, \tilde{Y}^{n'})$ .
    - (iii) Calculate how many periods in  $\hat{\Theta}_{1-\alpha}$  are  $\epsilon$ -far from  $\theta^*$ :

$$I_D^r = | \{ \theta \in \hat{\Theta}_{1-\alpha} : |\theta - \theta^*| > \epsilon \} | .$$

3. Select the “simplest” design possible that also yields sharp identification:

$$\hat{D} = \arg \min_{D \in \mathcal{D}} \left\{ |D| : \sum_{r=1}^R I_D^r = 0 \right\}, \quad (15)$$

where  $|D| \in \mathbb{R}$  denotes the complexity of the design (explained below).

In words, the design  $\hat{D}$  in (15) is the simplest possible observation design for which the confidence sets include only values in the  $\epsilon$ -neighborhood of  $\theta^*$ . In this case, we can accurately identify the underlying true signal. To implement the optimal design procedure outlined above, we only have to define the design space  $\mathcal{D}$  along with an appropriate notion of design complexity, and also discuss choices for threshold  $\epsilon$ . Here, for simplicity, we consider two choices for the design space:

- (A) **Fixed observations**,  $n$ . Each design in this space is parameterized by  $\delta > 0$  such that  $\tilde{T}^n = T^n + \delta U^{(n)}$ , where  $U^{(n)}$  is a vector of  $n$  independent  $\text{Unif}([-1, 1])$ , i.e., random variables uniformly distributed on  $[-1, 1]$ . Thus, the alternative observation time  $\tilde{t}_i$  is uniformly distributed  $\pm\delta$  days around the actual observation time  $t_i$ . Given  $\tilde{T}^n$ , we may generate synthetic

radial velocity data,  $\tilde{Y}^n$ , through  $y(\tilde{t}_i) = y^p(\tilde{t}_i; \theta^*, \hat{\psi}_{\theta^*}) + \varepsilon_i$ , where  $y^p$  is the periodic component defined in (A2), and  $\varepsilon_i \sim N(0, \sigma_i^2)$ , i.i.d. The complexity of a design  $D$  with parameter  $\delta$  can simply be defined as  $|D| = \delta$ .

- (B) **Additional observations**,  $n' - n$ . Each design in this space is parameterized by  $n' > 0$  such that  $\tilde{T}^{n'} = (T^n, \tilde{T}^{(n'-n)})$ , where we sample  $(n' - n)$  additional observation times in a hypothetical span of 3 months as follows:  $\tilde{t}_j = \max\{T^n\} + \text{Unif}(\{1, \dots, 90\}) + 0.01 \cdot \text{Unif}([-1, 1])$ , for  $j = 1, \dots, (n' - n)$ . Given  $\tilde{T}^{n'}$ , we can then generate synthetic radial velocity data  $\tilde{Y}^{n'}$  exactly as in Design (A). The complexity of a design  $D$  can simply be defined as  $|D| = n' - n$ .

Intuitively, Design (A) adds some level of randomization (controlled by  $\delta$ ) on top of the actual observation times  $T^n$ . Roughly speaking, this determines the amount of randomization necessary for accurate identification in future observations. When there are identification issues, a large  $\delta$  will generally be necessary for identification, and, conversely, when there are no identification issues a small  $\delta$  (or zero) will be needed. Design (B) adds new observation time points at a pre-specified level of noise. This level is set to  $\pm 0.01$  on the day-scale which amounts to  $2 \times 0.01 \times 24 \times 60 \approx 30$  mins of randomized observation window around  $n' - n$  additional daily observations. Similar to Design (A), when there are identification issues, a large number of additional observations,  $n' - n$ , will be necessary for identification; and, conversely, when there are no identification issues, a small number of additional observations will be needed.

Regarding the hyper-parameters, we note that  $R$  controls the number of synthetic datasets generated per design, and thus large  $R$  leads to more accurate calculation of the optimal design in (15). Practically, since we want to estimate small, near-zero proportion values,  $R$  could be small, e.g.,  $R = 1,000$  or even  $R = 100$ . Moreover, a reasonable value for  $\epsilon$  is in the range  $[0.1, 0.5]$ . Increasing that threshold to larger values could conflate real small-periodicity signals with the common 1-day period signal that is associated with the daily observation designs.

To illustrate these ideas, we apply the above optimal design procedure to the exoplanet detection problems of Section 5. The results are shown in Table 1. We see that for 51 Pegasi b and Gliese 436 b, there is no need for improving the observation times. Indeed, we need neither additional randomness in observation times ( $\delta = 0$ ), nor any additional observations ( $n' - n = 0$ ). In contrast, for the candidate exoplanet around  $\alpha$  Centauri B (Dumusque et al., 2012) we need a substantial improvement to the observation design. Specifically, in order to identify precisely the observed signal at  $\hat{\theta}_n = 3.23$  at the 1% level (assuming it's true), we need a large additional variation of  $\pm 0.18$  days around the actual observation times (i.e.,  $\pm 4.32$  hrs./observation). Alternatively, we need 137 additional observations with a random variation of  $\pm 15$  mins./observation, as specified by Design (B), almost equal to the size of the actual dataset. Finally, the situation is much better for Proxima Centauri (Anglada-Escudé et al., 2016). Specifically, in order to identify the observed signal at  $\hat{\theta}_n = 11.17$  at the 1% level—and so eliminate the only nuisance signal at  $\theta^{\text{nuis.}} = 0.9164$ —we need an additional variation of  $\pm 0.06$  days (i.e.,  $\pm 1.44$  hrs./observation) around the actual observation times. Alternatively, we only need 17 additional observations with a random variation of  $\pm 15$  mins./observation, as specified by Design (B).

Candidate Exoplanet (or Star)	Design (A)		Design (B)
	randomness needed for identification (best $\delta$ )	$\pm$ hrs.	#additional obs. needed for identification (best $n' - n$ )
51 Pegasi b (Mayor and Queloz, 1995)	0	0	0
Gliese 436 b (Butler et al., 2004)	0	0	0
$\alpha$ Centauri B (Dumusque et al., 2012)	0.18	4.32	137
Proxima Centauri (Anglada-Escudé et al., 2016)	0.06	1.44	17

Table 1: Observation designs (A) and (B) to achieve identification in the exoplanet applications of Section 5. Design (A) introduces additional randomness in the observation times, while design (B) introduces additional observations. We see that 51Pegb and GJ436b require no improvement in the observation times as the signal is well-identified. To identify the candidate exoplanet around  $\alpha$  Centauri B, however, we need an additional variation of  $\pm 0.18$  days around the actual observation times (i.e.,  $\pm 4.32$  hrs./observation). Alternatively, we need 137 additional observations with a random variation of  $\pm 15$  mins./observation. For the candidate exoplanet around Proxima Centauri (Anglada-Escudé et al., 2016) we need an additional variation of  $\pm 0.06$  days (i.e.,  $\pm 1.44$  hrs./observation) on the actual observation times. Alternatively, we only need 17 additional observations with a random variation of  $\pm 15$  mins./observation.

In summary, the results in Table 1 illustrate the benefits of the design procedure outlined above in guiding the selection of observation times for future exoplanet detections. In practice, we can envision a two-stage approach where at the first stage the scientist analyzes the stellar signal through the periodogram and identifies the most plausible peak along with other nuisance peaks. In the second stage, the scientist performs an analysis as described here in order to quantify how much variation or how many more samples are needed for identification of the plausible signal at a desired significance level, and proceeds to make these observations.

Of course, one limitation of our discussion here is that we have only considered two simple classes of observations designs. In practice, we need to also take into account the particular logistics of Earth-based observations to determine more realistic observation designs. The optimal design procedure in (15) could then be used in conjunction with such more realistic designs. Overall, this is an intriguing statistical problem, which we suggest for future work.

## 6.2 Non-parametric approach

The method developed in Section 3 is parametric and employs the harmonic model for the periodic component (A2). While this can be extended to include other periodic models (e.g., Keplerian model) it is interesting to consider non-parametric approaches as well. While nonparametric approaches have been considered in the literature before (Feigelson and Babu, 2012; Ivezić et al., 2014; Hall et al., 2000; Hall and Li, 2006; Hall, 2008), their asymptotics require strong assumptions on the observation design (typically i.i.d observation times), which is unrealistic. As such, non-parametric methods face the the same identification problems tied to irregular observation times as parametric methods. For example, the identification of a 26-million year biological extinction cycle in Example 2 used a nonparametric randomization test. However, as discused in Section 4.1, the period cannot be identified from this dataset.

Here, we discuss a nonparametric variant of our set identification method. As before, we assume that the signal is decomposed as  $y_i = y^p(t_i) + \varepsilon(t_i)$ , but make no parametric assumption about  $y^p$ . Before we define and test our main hypothesis, we need some notation. For any two time points  $t, t'$  and period  $\theta \in \Theta$ , let  $t \equiv t' \pmod{\theta}$  denote that  $t - t' = k\theta$  for some integer  $k$ . Also, let  $\mathcal{S}_n$  denote the symmetric group of permutations of  $n$ -element vectors. For some fixed  $\theta$ , define:

$$\Pi(T^n; \theta) = \{\pi \in \mathcal{S}_n : t_{\pi(i)} \equiv t_i \pmod{\theta}, i = 1, \dots, n\}.$$

In words,  $\Pi(T^n; \theta)$  is the set of permutations of  $\{1, \dots, n\}$  such that any time  $t_i$  is mapped only to an observation time that is equivalent to  $t_i$  modulo  $\theta$ . It is straightforward to see that  $\Pi(T^n; \theta)$  is a subgroup of  $\mathcal{S}_n$  since the mod operator defines an equivalence relationship. Let us assume that for any  $T^n$  and period  $\theta \in \Theta$ , the errors satisfy the invariance Assumption (A1) with  $\mathcal{G}^n = \Pi(T^n; \theta)$ . This assumption is generally mild. It requires that the data are exchangeable modulo  $\theta$ , and actually allows certain correlated structures for the errors  $\varepsilon(t)$ .

We wish to test the following nonparametric null hypothesis of periodicity  $\theta_0$ :

$$H_0^{\text{np}} : y^p(t') = y^p(t), \text{ for all } t', t \text{ such that } t' \equiv t \pmod{\theta_0}.$$

To test  $H_0^{\text{np}}$  we can adapt Procedure 1 as follows.

#### PROCEDURE 2

1. **For** all  $r = 1, \dots, R$  **do**:

(i) Sample  $\pi \sim \text{Unif}(\Pi(T^n; \theta_0))$ .

(ii) Generate synthetic outcome data  $Y^{n,(r)} = \pi \cdot Y^n$  obtained by permuting the data  $Y^n$  according to  $\pi$  while the observation times,  $T^n$ , are fixed.

2. Using the samples from 2(ii), calculate the  $p$ -value, say  $\text{pval}(\theta_0)$ , as in (11), and reject if the  $p$ -value is less than  $\alpha$ .

In words, to test  $H_0^{\text{np}}$ , Procedure 2, first, picks a permutation of the observation times in  $T^n$  such that the permuted times are equivalent to the corresponding original periods modulo  $\theta_0$ . Then, it uses this permutation to generate a new outcome vector  $Y^{n,(r)}$ . In the last step, the procedure recalculates the test statistic using data  $(T^n, Y^{n,(r)})$ , which can then be used in the calculation of the  $p$ -value.

**Theorem 4.** *Suppose that Assumptions (A1) and (A2) hold with  $\mathcal{G}^n = \Pi(T^n; \theta_0)$ . Then, the  $p$ -value from Procedure 2 is exact under  $H_0^{\text{np}}$  conditionally on the observation times, that is,*

$$\text{pr} \{ \text{pval}(\theta_0) \leq \alpha \mid H_0^{\text{np}}, T^n \} = \alpha.$$

*It follows that the  $p$ -value is also exact unconditionally.*

According to Theorem 4, the test of Procedure 2 is exact without requiring a parametric model for the periodic component,  $y^P$ . This is possible because under  $H_0^{\text{np}}$  and the invariance assumption on the errors, any two observations  $y(t), y(t')$  are exchangeable when  $t, t'$  are equivalent modulo  $\theta_0$ .

Despite the strength of this result, one practical difficulty with Procedure 2 is that it requires us to find several pairs of distinct time points in  $T^n$  that are periodically equivalent modulo  $\theta$ , for all  $\theta \in \Theta$ . This may be unrealistic in practical applications, including exoplanet detection, because the sample sizes can be small. Technically, the problem is that  $\Pi(T^n; \theta)$  may be a singleton group (including only the identity permutation) for many values of  $\theta$ , so that Procedure 2 leads to an exact by “empty” test, and thus to confidence sets that are too conservative. To address this problem, it may be tempting to expand the definition of  $\Pi(T^n; \theta)$  to include permutations between  $t, t'$  that are  $\epsilon$ -away from being periodically equivalent modulo  $\theta$ . Under this approach, however,  $\Pi(T^n; \theta)$  is not a group anymore. An alternative approach would be to use the nonparametric estimators of  $\theta^*$  developed by Hall et al. (2000); Hall and Li (2006); Hall (2008) together with a variation of Procedure 1 or Procedure 2. Both these procedures do not require regularity conditions on the observation times but only a consistent estimator for the periodic component,  $y^P$ . We leave these directions for future work.

## 7 Conclusion

We developed a method of set identification for hidden periodicity in unequally spaced time series. Our approach is more appropriate than standard methods of statistical inference because common estimators, such as the periodogram peak, are not well-behaved and may even be inconsistent without strong (and practically unrealistic) assumptions on the observation design. We illustrated empirically our method in examples from exoplanet detection using radial velocity data. In the confirmed exoplanets of our sample, our method can identify the hidden periodicity precisely. We also find that the statistical evidence in a recent candidate around  $\alpha$  Centauri B is particularly weak, as several nuisance signals cannot be rejected even at the 5% level. In another popular discovery around Proxima Centauri, the evidence appears to be robust but a nuisance signal due to the daily observational cycle can still not be rejected at the 1% level. Our method suggests ways to improve the observation designs, either by introducing some randomness in the observation times, or by just making more observations. These improved designs could resolve the aforementioned identification issues, and also help in efficient scheduling of measurements for future discoveries.

## 8 Acknowledgements

We would like to thank Yang Chen, Eric Feigelson, Vinay Kashyap, Xiao-Li Meng, Aneta Siemiginowska, Katja Smetanina, Michael Vogt, and the participants at the Statistics Department Seminar at UC Berkeley and the “Topics in Astrostatistics” Seminar at Harvard University for useful suggestions and feedback.

## References

- Anglada-Escudé, G., Amado, P. J., Barnes, J., Berdiñas, Z. M., Butler, R. P., Coleman, G. A., de La Cueva, I., Dreizler, S., Endl, M., Giesers, B., et al. (2016). A terrestrial planet candidate in a temperate orbit around proxima centauri. *Nature*, 536(7617):437–440.
- Baluev, R. V. (2008). Assessing the statistical significance of periodogram peaks. *Monthly Notices of the Royal Astronomical Society*, 385(3):1279–1285.
- Baluev, R. V. (2013). Detecting multiple periodicities in observational data with the multifrequency periodogram—i. analytic assessment of the statistical significance. *Monthly Notices of the Royal Astronomical Society*, 436(1):807–818.
- Bloomfield, P. (2004). *Fourier analysis of time series: an introduction*. John Wiley & Sons.
- Bølviken, E. (1983). New tests of significance in periodogram analysis. *Scandinavian Journal of Statistics*, pages 1–9.
- Butler, R. P., Vogt, S. S., Marcy, G. W., Fischer, D. A., Wright, J. T., Henry, G. W., Laughlin, G., and Lissauer, J. J. (2004). A neptune-mass planet orbiting the nearby m dwarf gj 436. *The Astrophysical Journal*, 617(1):580.
- Cai, T. T., Eldar, Y. C., Li, X., et al. (2016). Global testing against sparse alternatives in time-frequency analysis. *Annals of Statistics*, 44(4):1438–1466.
- Chiu, S.-T. (1989). Detecting periodic components in a white gaussian time series. *Journal of the Royal Statistical Society: Series B (Methodological)*, 51(2):249–259.
- Cho, R. J., Campbell, M. J., Winzeler, E. A., Steinmetz, L., Conway, A., Wodicka, L., Wolfsberg, T. G., Gabrielian, A. E., Landsman, D., Lockhart, D. J., et al. (1998). A genome-wide transcriptional analysis of the mitotic cell cycle. *Molecular cell*, 2(1):65–73.
- Costa, M. J., Finkenstädt, B., Roche, V., Lévi, F., Gould, P. D., Foreman, J., Halliday, K., Hall, A., and Rand, D. A. (2013). Inference on periodicity of circadian time series. *Biostatistics*, 14(4):792–806.
- Damasso, M. and Del Sordo, F. (2017). Proxima centauri reloaded: Unravelling the stellar noise in radial velocities. *Astronomy & Astrophysics*, 599:A126.
- Damsleth, E. and Spjøtvoll, E. (1982). Estimation of trigonometric components in time series. *Journal of the American Statistical Association*, 77(378):381–387.
- Delisle, J.-B., Hara, N., and Ségransan, D. (2020). Efficient modeling of correlated noise. i. statistical significance of periodogram peaks. *Astronomy & Astrophysics*.
- Dobriban, E. (2021). Consistency of invariance-based randomization tests. *arXiv preprint arXiv:2104.12260*.
- Dumusque, X., Pepe, F., Lovis, C., Ségransan, D., Sahlmann, J., Benz, W., Bouchy, F., Mayor, M., Queloz, D., Santos, N., et al. (2012). An earth-mass planet orbiting  $\alpha$  centauri b. *Nature*, 491(7423):207–211.
- Faria, J., Adibekyan, V., Amazo-Gómez, E., Barros, S., Camacho, J., Demangeon, O., Figueira, P., Mortier, A., Oshagh, M., Pepe, F., et al. (2020). Decoding the radial velocity variations of hd 41248 with espresso. *Astronomy & Astrophysics*, 635:A13.

- Faria, J., Haywood, R., Brewer, B., Figueira, P., Oshagh, M., Santerne, A., and Santos, N. (2016). Uncovering the planets and stellar activity of corot-7 using only radial velocities. *Astronomy & Astrophysics*, 588:A31.
- Feigelson, E. D. and Babu, G. J. (2012). *Modern statistical methods for astronomy: with R applications*. Cambridge University Press.
- Fisher, R. A. (1929). Tests of significance in harmonic analysis. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 125(796):54–59.
- Gelman, A. and Yao, Y. (2020). Holes in bayesian statistics. *Journal of Physics G: Nuclear and Particle Physics*, 48(1):014002.
- Hall, P. (2008). Nonparametric methods for estimating periodic functions, with applications in astronomy. In *COMPSTAT 2008*, pages 3–18. Springer.
- Hall, P. and Li, M. (2006). Using the periodogram to estimate period in nonparametric regression. *Biometrika*, 93(2):411–424.
- Hall, P., Reimann, J., and Rice, J. (2000). Nonparametric estimation of a periodic function. *Biometrika*, 87(3):545–557.
- Hall, P. and Yin, J. (2003). Nonparametric methods for deconvolving multiperiodic functions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(4):869–886.
- Hannan, E. (1971). Non-linear time series regression. *Journal of Applied Probability*, 8(4):767–780.
- Hannan, E. J. (1973). The estimation of frequency. *Journal of Applied probability*, pages 510–519.
- Hatzes, A. P. (2013). The radial velocity detection of earth-mass planets in the presence of activity noise: the case of  $\alpha$  centauri bb. *The Astrophysical Journal*, 770(2):133.
- Ivezić, Ž., Connolly, A. J., VanderPlas, J. T., and Gray, A. (2014). *Statistics, data mining, and machine learning in astronomy: a practical Python guide for the analysis of survey data*, volume 1. Princeton University Press.
- Juditsky, A., Nemirovski, A., et al. (2015). On detecting harmonic oscillations. *Bernoulli*, 21(2):1134–1165.
- Kiefer, J. (1975). Optimal design: Variation in structure and performance under change of criterion. *Biometrika*, 62(2):277–288.
- Lehmann, E. L. and Romano, J. P. (2006). *Testing statistical hypotheses*. Springer Science & Business Media.
- Linnell Nemeč, A. and Nemeč, J. M. (1985). A test of significance for periods derived using phase-dispersion-minimization techniques. *The Astronomical Journal*, 90:2317–2320.
- Lomb, N. R. (1976). Least-squares frequency analysis of unequally spaced data. *Astrophysics and space science*, 39(2):447–462.
- Loredo, T. (2003). Bayesian adaptive exploration in a nutshell. *Statistical Problems in Particle Physics, Astrophysics, and Cosmology*, 1:162.
- Manski, C. F. (2003). *Partial identification of probability distributions*. Springer Science & Business Media.

- Mascareño, A. S., Faria, J., Figueira, P., Lovis, C., Damasso, M., Hernández, J. G., Rebolo, R., Cristiani, S., Pepe, F., Santos, N., et al. (2020). Revisiting proxima with espresso. *Astronomy & Astrophysics*, 639:A77.
- Mayor, M. and Queloz, D. (1995). A jupiter-mass companion to a solar-type star. *Nature*, 378(6555):355–359.
- McSweeney, L. A. (2006). Comparison of periodogram tests. *Journal of Statistical Computation and Simulation*, 76(4):357–369.
- Mitchell, T. J. (2000). An algorithm for the construction of “d-optimal” experimental designs. *Technometrics*, 42(1):48–54.
- Murray, E. J. and Morse, S. S. (2011). Seasonal oscillation of human infection with influenza a/h5n1 in egypt and indonesia. *PloS one*, 6(9):e24042.
- Quinn, B. and Fernandes, J. (1991). A fast efficient technique for the estimation of frequency. *Biometrika*, 78(3):489–497.
- Quinn, B. G. and Hannan, E. J. (2001). *The estimation and tracking of frequency*. Number 9. Cambridge University Press.
- Rajpaul, V., Aigrain, S., Osborne, M. A., Reece, S., and Roberts, S. (2015a). A gaussian process framework for modelling stellar activity signals in radial velocity data. *Monthly Notices of the Royal Astronomical Society*, 452(3):2269–2291.
- Rajpaul, V., Aigrain, S., and Roberts, S. (2015b). Ghost in the time series: no planet for alpha cen b. *Monthly Notices of the Royal Astronomical Society: Letters*, 456(1):L6–L10.
- Raup, D. M. (1986). Biological extinction in earth history. *Science*, 231(4745):1528–1533.
- Raup, D. M. and Sepkoski, J. J. (1986). Periodic extinction of families and genera. *Science*, 231(4740):833–836.
- Romano, J. P. and Shaikh, A. M. (2008). Inference for identifiable parameters in partially identified econometric models. *Journal of Statistical Planning and Inference*, 138(9):2786–2807.
- Rosenbaum, P. R. et al. (2010). *Design of observational studies*, volume 10. Springer.
- Scargle, J. D. (1982). Studies in astronomical time series analysis. ii-statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal*, 263:835–853.
- Schuster, A. (1898). On the investigation of hidden periodicities with application to a supposed 26 day period of meteorological phenomena. *Terrestrial Magnetism*, 3(1):13–41.
- Siegel, A. F. (1980). Testing for periodicity in a time series. *Journal of the American Statistical Association*, 75(370):345–348.
- Smucker, B., Krzywinski, M., and Altman, N. (2018). Optimal experimental design. *Nature methods*, 15(8):559–560.
- Stigler, S. M. and Wagner, M. J. (1987). A substantial bias in nonparametric tests for periodicity in geophysical data. *Science*, 238(4829):940–945.

- Sun, Y., Hart, J. D., and Genton, M. G. (2012). Nonparametric inference for periodic sequences. *Technometrics*, 54(1):83–96.
- Tamer, E. (2010). Partial identification in econometrics. *Annu. Rev. Econ.*, 2(1):167–195.
- Toulis, P. (2019). Life after bootstrap: Residual randomization inference in regression models. *arXiv preprint arXiv:1908.04218*.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- VanderPlas, J. T. (2018). Understanding the lomb–scargle periodogram. *The Astrophysical Journal Supplement Series*, 236(1):16.
- Vogt, M. and Linton, O. (2014). Nonparametric estimation of a periodic sequence in the presence of a smooth trend. *Biometrika*, 101(1):121–140.
- Walker, A. (1971). On the estimation of a harmonic component in a time series with stationary independent residuals. *Biometrika*, 58(1):21–36.
- Wright, J. and Howard, A. (2009). Efficient fitting of multiplanet keplerian models to radial velocity and astrometry data. *The Astrophysical Journal Supplement Series*, 182(1):205.
- Wynn, H. P. (1984). Jack kiefer’s contributions to experimental design. *The Annals of Statistics*, 12(2):416–423.
- Zechmeister, M. and Kürster, M. (2009). The generalised lomb-scargle periodogram—a new formalism for the floating-mean and keplerian periodograms. *Astronomy & Astrophysics*, 496(2):577–584.

## A. Proof of Theorem 1

$$g \cdot \varepsilon^n \stackrel{d}{=} \varepsilon^n \mid T^n \quad (g \in \mathcal{G}^n). \quad (\text{A1})$$

$$y^p(t; \theta^*, \psi^*) = \psi_1^* + \psi_2^* \cos(2\pi t/\theta^*) + \psi_3^* \sin(2\pi t/\theta^*). \quad (\text{A2})$$

The definition of the  $p$ -value for  $H_0^{\text{full}}$  is as follows:

$$\text{pval}(\theta_0, \psi_0) = E\{s_n(Y^{n,(i)}, T^n) \geq s_{\text{obs}} \mid T^n\},$$

**Theorem 1.** *Suppose that Assumptions (A1) and (A2) hold. Then, the  $p$ -value in (5) is exact in finite samples under  $H_0^{\text{full}}$ , that is, for any finite  $n > 0$ ,*

$$\text{pr}\{\text{pval}(\theta_0, \psi_0) \leq \alpha \mid H_0^{\text{full}}\} = \alpha.$$

*Proof.* Recall that  $Y^{n,(i)} = Y^{n,p} + G^{(i)}\varepsilon^n$  where  $G^{(i)} \sim \text{Unif}(\mathcal{G}^n)$  and  $\varepsilon^n = Y^n - Y^{n,p}$ . Under  $H_0^{\text{full}}$  and Assumptions (A1)-(A2) we obtain  $Y^{n,(i)} \stackrel{d}{=} Y^n \mid T^n$ . It follows that

$$s_n(Y^{n,(i)}, T^n) \stackrel{d}{=} s_{\text{obs}} \mid T^n$$

and so the  $p$ -value in (5) is finite-sample valid, that is,  $E[1\{\text{pval}(\theta_0, \psi_0) \leq \alpha\} \mid H_0^{\text{full}}, T^n] = \alpha$ . It follows that the  $p$ -value is also valid *unconditionally*:

$$E[1\{\text{pval}(\theta_0, \psi_0) \leq \alpha\} \mid H_0^{\text{full}}] = E[E[1\{\text{pval}(\theta_0, \psi_0) \leq \alpha\} \mid H_0^{\text{full}}, T^n] \mid T^n] = \alpha.$$

□

## B. Proof of Theorem 2

$$H_0 : \theta^* = \theta_0.$$

$$\Theta_{1-\alpha} = \{\theta \in \Theta : \max_{\psi \in \Psi} \text{pval}(\theta, \psi) > \alpha\}.$$

**Theorem 2.** *Suppose that Assumptions (A1) and (A2) hold. Then,  $\Theta_{1-\alpha}$  is a finite-sample valid  $100(1-\alpha)\%$  confidence set for  $\theta^*$ ; i.e., for any finite  $n > 0$ ,*

$$\text{pr}(\theta^* \in \Theta_{1-\alpha} \mid H_0) \geq 1 - \alpha.$$

*Proof.* Suppose that  $H_0$  in (6) holds., Then,

$$\text{pr}(\theta_0 \notin \Theta_{1-\alpha} \mid H_0) = E[1\{\max_{\psi} \text{pval}(\theta^*, \psi) \leq \alpha\}] \leq E[1\{\text{pval}(\theta^*, \psi^*) \leq \alpha\}] = \alpha.$$

The first equality follows by definition of  $H_0$  in (6); the second follows from the indicator identity  $1(a \leq b) \leq 1(c \leq b)$  when  $a \geq c$ ; and the third follows from Theorem 1. Since  $\theta^* \in \Theta$  w.p. 1 it follows that  $\text{pr}(\theta^* \in \Theta_{1-\alpha}) \geq 1 - \alpha$ . □

## Proof of Theorem 3

**Theorem 3.** *Suppose that Assumptions (A1) and (A2) hold, and that  $\hat{\psi}_{\theta_0} \xrightarrow{P} \psi^*$  under  $H_0$ . Then,  $\hat{\Theta}_{1-\alpha}$  is an asymptotically valid  $100(1-\alpha)\%$  confidence set for  $\theta^*$ ; i.e., as  $n$  increases*

$$\text{pr}(\theta^* \in \hat{\Theta}_{1-\alpha} \mid H_0) \geq 1 - \alpha + o_P(1).$$

*Proof.* Following a similar reasoning as above:

$$\begin{aligned} \text{pr}(\theta_0 \notin \hat{\Theta}_{1-\alpha} \mid H_0) &= \text{pr}\{\text{pval}(\theta^*, \hat{\psi}_{\theta^*}) \leq \alpha \mid H_0\} && \text{(i)} \\ &= \text{pr}\{\text{pval}(\theta^*, \psi^*) + \underbrace{\text{pval}(\theta^*, \hat{\psi}_{\theta^*}) - \text{pval}(\theta^*, \psi^*)}_{\Delta} \leq \alpha \mid H_0\} \\ &\leq \text{pr}\{\text{pval}(\theta^*, \psi^*) \leq \alpha + |\Delta| \mid H_0\} && \text{(ii)} \\ &\leq \text{pr}(|\Delta| > \epsilon \mid H_0) + \text{pr}\{\text{pval}(\theta^*, \psi^*) \leq \alpha + \epsilon \mid H_0\} && \text{(iii)} \\ &= \alpha + \epsilon + P(|\Delta| > \epsilon \mid H_0). && \text{(iv)} \\ &\leq \alpha + \min_{\epsilon} \{\epsilon + P(|\Delta| > \epsilon \mid H_0)\} && \text{(v)} \\ &= \alpha + o_P(1). && \text{(vi)} \end{aligned}$$

In these derivations, line (i) follows from definition of the confidence set and  $H_0$ ; line (ii) follows from the indicator identity  $1(A + B \leq a) \leq 1(A \leq a + |B|)$  for any scalar random variables  $A, B$  and  $a \in \mathbb{R}$ ; line (iii) follows from the indicator identity  $1(A \leq a + B) \leq 1(B > \epsilon) + 1(A \leq a + \epsilon)$  for any  $\epsilon < 0$ , and  $A, B, a$  as above; line (iv) follows from Theorem 1; line (v) gives the tightest bound exploiting the fact that  $\epsilon > 0$  is arbitrary. Finally, line (vi) follows from consistency of  $\hat{\psi}_{\theta^*}$ .  $\square$

## Proof of Theorem 4

**Theorem 4.** *Suppose that Assumptions (A1) and (A2) hold with  $\mathcal{G}^n = \Pi(T^n; \theta_0)$ . Then, the  $p$ -value from Procedure 2 is exact under  $H_0^{\text{np}}$  conditionally on the observation times, that is,*

$$\text{pr}\{\text{pval}(\theta_0) \leq \alpha \mid H_0^{\text{np}}, T^n\} = \alpha.$$

*It follows that the  $p$ -value is also exact unconditionally.*

*Proof.* The proof follows the same steps as Theorem 1 as we only have to derive that  $s_n \stackrel{d}{=} s_{\text{obs}} \mid T_n$ . The only difference is that this equivalence is implied immediately from the assumed invariance based on  $\Pi(T^n; \theta)$  rather than relying on the correct periodic component  $y^p$  as in Theorem 1.  $\square$