



# Improving the Lomb–Scargle Periodogram with the Thomson Multitaper

Aaron Springford<sup>1</sup> , Gwendolyn M. Eadie<sup>2,3</sup> , and David J. Thomson<sup>4</sup>

<sup>1</sup> Statistician, Toronto, ON, Canada; [stats@aaronspringford.com](mailto:stats@aaronspringford.com)

<sup>2</sup> Department of Astronomy, University of Toronto, Toronto, ON, Canada; [gwen.eadie@utoronto.ca](mailto:gwen.eadie@utoronto.ca)

<sup>3</sup> Department of Statistical Sciences, University of Toronto, Toronto, ON, Canada

<sup>4</sup> Department of Mathematics & Statistics, Queen's University, Kingston, ON, Canada; [djt@queensu.ca](mailto:djt@queensu.ca)

Received 2019 June 20; revised 2020 March 8; accepted 2020 March 11; published 2020 April 13

## Abstract

A common approach for characterizing the properties of time-series data that are evenly sampled in time is to estimate the power spectrum of the data using the periodogram. The periodogram as an estimator of the spectrum is (1) statistically inconsistent (i.e., its variance does not go to zero as infinite data are collected), (2) biased for finite samples, and (3) suffers from spectral leakage. In astronomy, time-series data are often unevenly sampled in time, and it is popular to use the Lomb–Scargle (LS) periodogram to estimate the spectrum. Unfortunately, from a statistical standpoint, the LS periodogram suffers from the same issues as the classical periodogram and has even worse spectral leakage. Here, we present an improvement on the LS periodogram by combining it with the Thomson multitaper approach. The multitaper spectral estimator is well established in the statistics and engineering literature for evenly sampled time series. It is attractive because it directly trades off bias and variance for frequency resolution, and is fast to compute: compared to an untapered spectral estimator, the multitaper adds no more than a couple of seconds for a time series with a million data points on a current desktop computer. Here, we describe an estimator that combines the multitaper with the LS periodogram. We show examples in which this new approach has improved properties compared to traditional approaches in the case of unevenly sampled time series. Finally, we demonstrate an application of the method to astronomy with an application to Kepler data.

*Unified Astronomy Thesaurus concepts:* [Time series analysis \(1916\)](#); [Interdisciplinary astronomy \(804\)](#); [Astrostatistics \(1882\)](#); [Algorithms \(1883\)](#); [Astrostatistics techniques \(1886\)](#); [Astrostatistics strategies \(1885\)](#); [Irregular cadence \(1953\)](#); [Asteroseismology \(73\)](#); [Stellar oscillations \(1617\)](#)

## 1. Introduction

The analysis of time-series data is becoming increasingly common in astrophysics research, and will continue to grow in prevalence as data from telescopes such as the National Science Foundation Vera C. Rubin Telescope (LSST Science Collaboration et al. 2009) comes online in the 2020s. Data from spacecraft and ground-based telescopes such as Kepler (Koch et al. 2010), the Transiting Exoplanet Survey Satellite (TESS; Ricker et al. 2015), and the Zwicky Transient Facility (ZTF; Bellm et al. 2018; Graham et al. 2019) have and are improving our understanding of exoplanets and stars, but we are limited by issues such as unevenly sampled data, measurement uncertainties, and the statistical methods we employ. This paper is most concerned with the latter, and seeks to provide an improved statistical method for identifying and estimating periodic signals in time-series data.

Various methods exist for analyzing data sampled in time, but a common approach is called spectral analysis. Spectral analysis is used to estimate and characterize the process underlying time-series data by expressing power as a function of frequency, and obtaining an estimate of the frequency spectrum or power spectrum. When performing spectral analysis, we obtain from the data a statistical estimate of the true power spectrum—we will never know to true power spectrum, but can hope to estimate it well using reliable methods.

Most spectrum estimators (e.g., the periodogram) depend heavily on the assumption that the data are regularly sampled in time. In astronomy, we do not always have the luxury of well-controlled observation times, and so time-series data are almost always irregularly sampled to some degree. Thus, the

Lomb–Scargle (LS) periodogram (Lomb 1976; Scargle 1982), which accounts for unevenly sampled data, is a popular method for estimating the frequency spectra of astronomical time-series data.

The application of the LS estimator has become very popular in astronomy, despite its many drawbacks (see VanderPlas 2018 for a thorough discussion about LS). One method for spectral analysis that is widely used in other areas of science but less so in astronomy is the multitaper statistic invented by Thomson (1982).<sup>5</sup> One reason for this may be that the Thomson multitaper statistic, like many time-series analysis methods, relies on evenly sampled data. Closely related fields of helioseismology, solar physics, and geophysics have seen the use of the multitaper statistic (e.g., Thomson et al. 1996; Komm et al. 1999; Thomson & Vernon 2015, 2016).

In this paper, we propose the multitaper-Lomb–Scargle (MTLS) statistic: a method for time-series analysis that combines the multitaper approach with the LS technique. The method is a rapid computation of the multitaper spectral statistic (Thomson 1982) under irregular sampling. We foresee this method as an easily implemented improvement to the most common approach to spectrum estimation when times are known but irregular.

We aim to make this paper accessible to those relatively new to time-series analysis and/or statistical notation for time-series analysis. If you are already well versed in both, then you may want to skip ahead to Section 4.

The paper is organized as follows.

<sup>5</sup> This is the mathematically rigorous paper—for a less intense summary, see Thomson (1990). Further statistical properties of the spectrum estimate are described in Thomson & Haley (2014).

**Section 2:** a short introduction to mathematical notation for time-series analysis and a description of important assumptions and concepts such as stationarity, the autocovariance function, and the spectral density function.

**Section 3:** a review of the classical periodogram and its statistical issues, which necessitate the use of tapers.

**Section 4:** description of the discrete prolate spheroidal sequences (DPSS) as tapers, which are used in the multitaper spectral estimator.

**Section 5:** a well-known example that demonstrates the superior estimate of the frequency power spectrum provided by the multitaper statistic over the periodogram.

**Section 6:** issues with interpolating missing data in a time series.

**Section 7:** the MTLs statistic; the LS periodogram is briefly reviewed (Section 7.1) and then combined with the multitaper spectral statistic (Section 7.2).

**Section 8:** the MTLs is used in a preliminary analysis of time-series data from a red giant (RG) star as measured by Kepler.

Much of the work presented in this paper is part of Springford (2017), which also contains more details pertaining to time-series analysis with irregular and latent times (i.e., when there is a gap for time).

## 2. Notation for Time-series Analysis

In time-series analysis, we consider a sequence of random variables  $X_t$  collected at the set of times  $\mathcal{D}$ ,

$$\{X_t: t \in \mathcal{D}\},$$

where we assume the set of times  $\mathcal{D}$  are strictly increasing. Next, assume that the random variables have a joint distribution with finite mean

$$\mu_t = E[X_t], \quad t \in \mathcal{D}$$

and finite variance/covariance

$$\text{cov}(X_t, X_r) = E[(X_t - \mu_t)(X_r - \mu_r)], \quad t, r \in \mathcal{D}.$$

Here,  $t$  and  $r$  are two different times from  $\mathcal{D}$ , and  $E[\ ]$  denotes the expected value of whatever is inside the square brackets.<sup>6</sup>

A time series is a realization  $\mathbf{x}$  drawn from  $\mathbf{X}$ , at a set of times  $\mathcal{D}$ . For example,  $\mathbf{x}$  might be the measurements of the flux of a star collected at a set of times  $\mathcal{D}$ .

A time series is regularly (or evenly) sampled if  $t_i - t_{i-1} = c$  for all  $i$ , where  $c$  is a constant. Regularly sampled time series are the most common in other scientific disciplines, particularly when the observation times can be controlled by the observer.

Of primary interest in the analysis of time series is the covariance of the random variables  $\{X_t\}$ . This differentiates time-series analysis from much of statistics in which random variables are often assumed to be independent or conditionally independent (spatial statistics has a similar interest in covariance structure). It also presents a problem for inference because typically we only obtain a single time series—that is, a single sample—making it impossible to characterize the covariance without simplifying assumptions about the covariance structure. Easily the most common assumption is that of *stationarity*, of which there are two flavors:

1. strong stationarity assumes that any finite collection of random variables  $\{X_{t_1}, \dots, X_{t_m}\}$  has the same joint distribution as  $\{X_{t_1+s}, \dots, X_{t_m+s}\}$ , where  $s \in (-\infty, \infty)$  if the times are continuous, and  $s \in \mathbb{Z}$  for integer-valued discrete times ( $\mathbb{Z}$  is the set of integers); and
2. weak or second-order stationarity assumes that  $E[X_t] = \mu$  for every  $t$  and that the covariance  $\text{cov}(X_t, X_{t+s})$  does not depend on  $t$ , assuming that the variance of each  $X_t$  is finite.

If  $\{X_t\}$  follow a multivariate Gaussian distribution of finite dimension, weak stationarity implies strong stationarity because the multivariate Gaussian is uniquely determined by its mean and covariance. In practice, it is impossible to verify these stationarity assumptions, but some natural systems seem more likely than others to be non-stationary. For example, pulsars often exhibit obvious, predictable, and regular variation in flux over several years. On the other hand, the flux measured from a supernova changes quite a lot depending on when you start measuring!

Assuming weak stationarity, we are interested in the autocovariance function

$$R(s) = \text{cov}(X_t, X_{t+s}) = E[(X_t - \mu)(X_{t+s} - \mu)]. \quad (1)$$

We will assume for now that we have  $\mathcal{D} = 0, 1, 2, \dots, N-1$ , where  $N$  is the length of the time series. These time values will make the mathematics simpler compared to regular Julian dates, for example.

Under mild regularity conditions, the spectral density function of frequency  $f$  can be represented as

$$S(f) = \sum_{s=-\infty}^{\infty} R(s) \exp(-2\pi ifs), \quad (2)$$

which provides another approach to statistical inference via the spectral density or spectrum of the time series. The autocovariance and the spectrum are Fourier transform pairs, so that given the spectrum, the autocovariance can be represented as

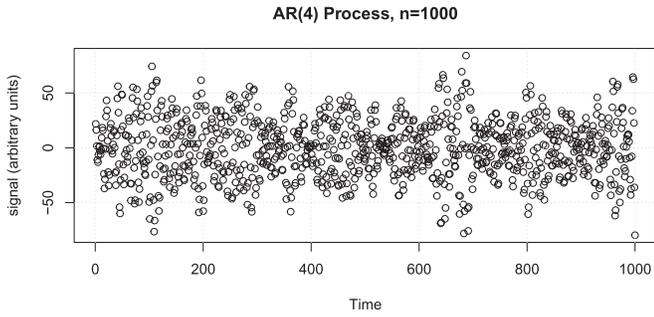
$$R(s) = \int_{-1/2}^{1/2} S(f) \exp(2\pi ifs) df. \quad (3)$$

This is the Einstein–Wiener–Khinchine theorem (Einstein 1914; Wiener 1930; Khinchine 1934).

The spectrum has the following important properties:

Property	Descriptive Meaning
$S(f) \geq 0$	The spectrum is non-negative
$S(f) = S(f+1)$	The spectrum is periodic and has period equal to one
$S(f) = S(-f)$	The spectrum is an even function for real-valued $\mathbf{x}$

<sup>6</sup> Also known as the expectation in physics and astronomy literature.



**Figure 1.** A single, evenly sampled time series generated from an AR(4) process with coefficients  $\varphi_1 = 2.7607$ ,  $\varphi_2 = -3.8106$ ,  $\varphi_3 = 2.6535$ , and  $\varphi_4 = -0.9238$ .

The equivalence of the autocovariance  $R(s)$  and the spectrum  $S(f)$  means that either can be used for inference. However, the spectrum has an arguably simpler interpretation and provides the opportunity to discover something unexpected (Tukey & Hamming 1949). The spectrum is precisely tuned to detect periodicities.

Another strong point in favor of the spectrum compared to the autocovariance is that at suitably spaced frequencies, the distribution of spectrum estimates are or can be made approximately independent. This independence simplifies interpretation and inference considerably.

### 3. Statistical Properties of the Periodogram

A naïve estimator of the spectrum, known as the periodogram (Schuster 1898), is

$$\hat{S}_P(f) = \frac{1}{N} \left| \sum_{t=0}^{N-1} x_t \exp(-2\pi ift) \right|^2. \quad (4)$$

(Recall that  $N$  is the length of the time series.) The estimator in Equation (4) assumes we have a time series  $x$ , a realization of  $X$ , observed at a set of known and discrete times. It also assumes that the observation times are regularly spaced with  $\mathcal{D} = 0, 1, 2, \dots, N - 1$ .

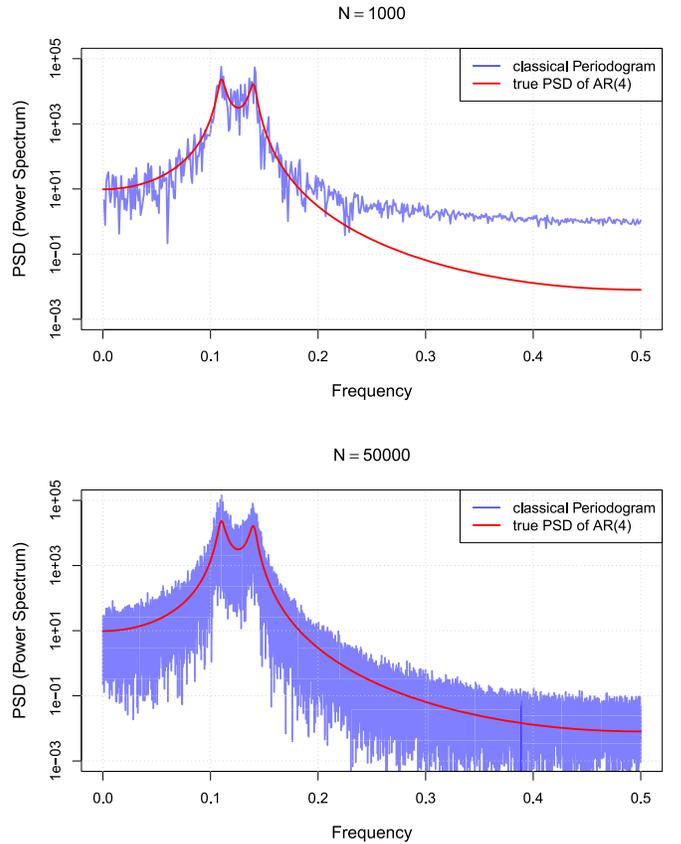
The periodogram estimator of the spectrum,  $\hat{S}_P(f)$ , has some problems. It is inconsistent: the variance of the estimator does not tend to zero as  $N$  goes to infinity (first proved by Rayleigh 1903). In fact, the variance of the estimator does not decrease as the sample size increases.

The spectral window for the periodogram is Fejér’s kernel  $\mathcal{F}(f) = \frac{1}{N} \frac{\sin^2(N\pi f)}{\sin^2(\pi f)}$ . This kernel has side lobes that contribute to the spectrum estimate at frequencies away from a given frequency, an effect known as spectral leakage. Spectral leakage can cause the periodogram to be badly biased (Thomson 1982).

Both the inconsistency and bias of the periodogram estimator can be displayed through an example first presented in Box & Jenkins (1970) and later in Percival & Walden (1993). The example begins by considering a time series generated by an autoregressive process,

$$X_t = c + \sum_{i=1}^p \varphi_i X_{t-i} + \epsilon_t, \quad (5)$$

of order  $p = 4$ , with coefficients  $\varphi_1 = 2.7607$ ,  $\varphi_2 = -3.8106$ ,  $\varphi_3 = 2.6535$ , and  $\varphi_4 = -0.9238$ . This autoregressive process (denoted AR(4)) is used to generate the time series shown in



**Figure 2.** Power spectrum estimate from the classical periodogram (blue curves) and the true power spectral density (red dashed curve) for an AR(4) process time series of length  $N = 1000$  (top) shown in Figure 1. The bottom figure is the same but for  $N = 50,000$ . Note that the classical periodogram estimate of the power spectrum is biased at high frequency for finite  $N$ , and is inconsistent.

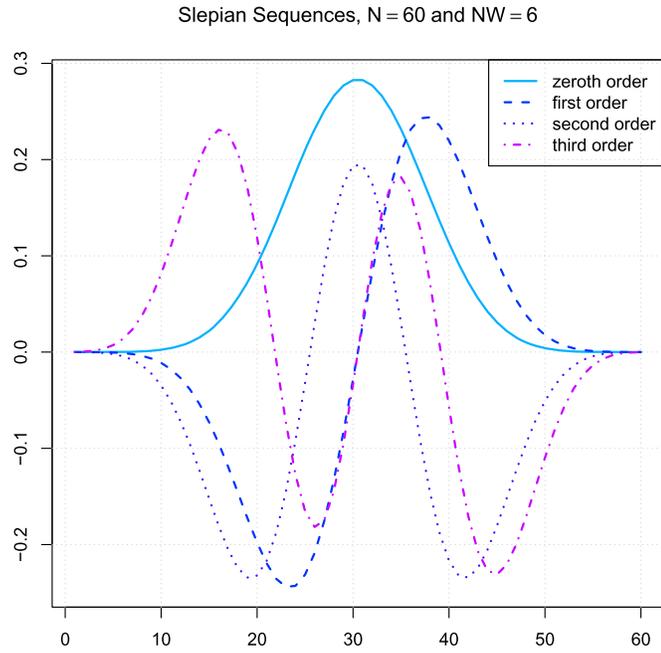
Figure 1. The evenly sampled time series is then used with the classical periodogram estimator to estimate the spectral density (or power spectrum).

Figure 2 shows the power spectrum as estimated by the periodogram (blue curve) for the AR(4) time series of length  $N = 1000$  (top) and  $N = 50,000$  (bottom), and compares it to the true power spectrum that is known analytically (the red curve). Note that the classical periodogram estimator is biased at high frequency for the time series of length  $N = 1000$  (top). The spectrum estimate for the time series of length  $N = 50,000$  (bottom) is less biased, but the power spectrum remains statistically inconsistent—the variance does not go to zero as  $N$  increases.

These statistical issues with the periodogram are well known and so estimators that taper the data are preferred (Brillinger 1981). The tapered estimator of the spectrum, known as a direct estimate, is

$$\hat{S}_D(f) = \left| \sum_{t=0}^{N-1} h_t x_t \exp(-2\pi ift) \right|^2 \quad (6)$$

where  $h_t$  is the data taper, which notably has the property that  $\sum_{t=0}^{N-1} h_t^2 = 1$ . Examining Equation (6), the taper reweights the time-series values, typically down-weighting values at the beginning and end of the time series. In the frequency domain, this results in a spectral window with lower side lobes than



**Figure 3.** Slepian sequences of order  $k = 0$  to  $k = 3$ , assuming a time-series length of  $N = 60$  and a frequency bandwidth of  $2W = 1.2$ .

Fejér’s kernel, reducing spectral leakage in the estimate (see also Blackman & Tukey 1958; Blackman 1959).

There are many kinds of data tapers  $h_t$  (sometimes called windows) that can be used to mitigate the spectral leakage problem (Harris 1978). The family of cosine tapers are popular (of which the Hann window is a special case), as is the modified cosine taper or Hamming window.<sup>7</sup>

In the next section, we describe the particular set of tapers used by the Thomson (1982) multitaper statistic and why they are an optimal choice.

#### 4. Slepian Tapers

Thomson’s multitaper statistic makes use of not one taper, but multiple tapers. Specifically, it uses a special set of sequences called the DPSS (colloquially referred to as the Slepian sequences in reference to Slepian 1978).

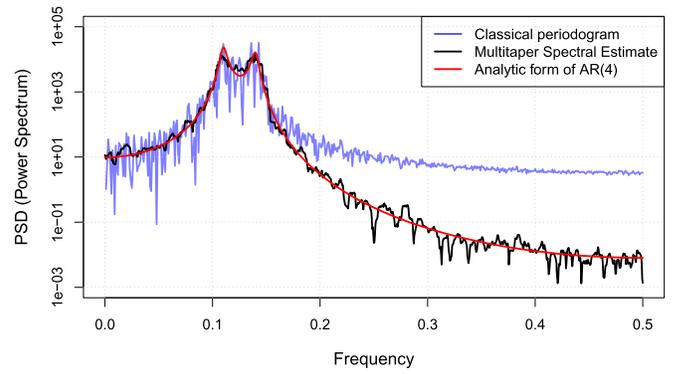
The Slepian sequences are usually denoted as  $v_t^{(k)}(N, W)$ , where  $k$  is the order of the sequence. They are sequences of length  $N$  (i.e., the length of the time series) that have the smallest possible spectral leakage outside of a defined frequency band of width  $2W$ . As an example, the Slepian sequences of order  $k = 0$  through  $k = 3$  for a time series of length  $N = 60$  and frequency band of width  $2W = 1.2$  are shown in Figure 3.

The higher-order DPSS can also have nearly minimal spectral leakage. In fact, the first  $K \lesssim 2NW$  have nearly minimal spectral leakage, so that for any  $k \lesssim 2NW$

$$\hat{S}_k^{(mt)}(f) = \left| \sum_{t=0}^{N-1} v_t^{(k)} x_t \exp(-2\pi ift) \right|^2 \quad (7)$$

will provide an estimate that is characteristically similar (note here that we have replaced Slepian’s notation with an equivalent shorthand:  $v_t^{(k)}$ ). Moreover, the estimates  $\hat{S}_k^{(mt)}(f)$  are approximately uncorrelated because the DPSS tapers are

<sup>7</sup> Note the similarity between the names, which often leads to confusion; the Hann window is often called the “Hanning” window!



**Figure 4.** Comparison of the power spectrum estimate made by the classical periodogram (blue curve) to that made by the multitaper spectral estimator (black curve). The true power spectrum for the AR(4) process is shown as a red dashed curve.

orthogonal (Thomson 1982). This suggests averaging the  $\{\hat{S}_k^{(mt)}(f)\}$  to not only control spectral leakage, but also to reduce the variance of the combined estimate.

The result is the multitaper spectrum estimate

$$\hat{S}^{(mt)}(f) = \frac{1}{K} \sum_{k=0}^{K-1} \hat{S}_k^{(mt)}(f), \quad (8)$$

where  $K$  is the maximum order of the DPSS taper chosen in forming the estimate. An adaptively weighted average can be used to further improve the in-band concentration (Thomson 1982). It is also common to play it safe with respect to spectral leakage and use  $K = \lfloor 2NW \rfloor - 1$  or fewer DPSS tapers (Thomson 1982).

By using multiple tapers, the Thomson multitaper spectrum estimate reduces bias, reduces variance, and reduces spectral leakage. In the next section, we continue the example from Percival & Walden (1993) that helped bolster the application of the multitaper estimator in other disciplines.

#### 5. A Motivating Example for the Multitaper Spectral Estimator

Continuing with the example in Section 3, we now apply the multitaper spectrum estimator to the same time series (AR(4) process) presented in Figure 1.

Figure 4 shows the true spectrum for the AR(4) process (red curve), the spectral estimate from the periodogram (blue curve), and the spectral estimate from the Thomson multitaper approach (black curve; this example appears in Percival & Walden 1993). The Thomson multitaper estimator clearly reduces bias and variance, and accurately predicts the true power spectrum. Note that one way to reduce the variance of the periodogram is to apply a smoothing filter after computing the estimate. We hope that this example makes obvious the shortcoming of this approach: although the variance of the estimate would be reduced, the bias would remain.

The Thomson multitaper estimator is meant for regularly sampled time-series data, and Stoica & Sundin (1999) proved that when the spectrum is relatively smooth, the multitaper estimate is equivalent to the maximum likelihood. In the rest of this paper, we introduce a way to combine the multitaper estimator with the LS periodogram in order to make this approach feasible and useful for irregularly sampled time-series data.

## 6. Irregular Times

There exists a version of the multitaper for irregular time series (Bronze 1988). The approach involves computing generalized DPSS for each frequency band of interest by solving a generalized matrix eigenvalue problem. By definition, the Bronze solution achieves optimal in-band concentration, just like Thomson (1982) in the case of regularly sampled times. However, the estimator is extremely computationally expensive and numerically difficult to compute. These qualities make it ill-suited to problems in which a quick and reliable solution is necessary for Markov Chain Monte Carlo (MCMC) sampling or in which there are thousands of time series to be analyzed (e.g., forthcoming data from the Legacy Survey of Space and Time; LSST Science Collaboration et al. 2009).

Still others have developed solutions for otherwise regular series with gaps (e.g., Fodor & Stark 2000; Smith-Boughner & Constable 2012). Most recently, Chave (2019) introduces a generalization of the multitaper for gappy time-series data that are otherwise regularly sampled. However, this method is not intended for data with arbitrary sampling times, and when applied to time series with many small gaps does not work well (A. Chave 2020, personal communication).

Approaches that interpolate the autocovariance matrix using slotted estimators based on discretization of time lags are also possible (Stoica & Sandgren 2006), but again would seem to be more suited to series that are otherwise regular or comprised of a small set of discrete time increments. A review of various parametric, non-parametric, and heuristic approaches is available in Babu & Stoica (2010).

Given the computational advantage afforded by the spectral estimators under regular sampling, a natural question is whether we can use methods of spectrum estimation that are designed for regularly sampled time series in cases where the times are irregularly sampled.

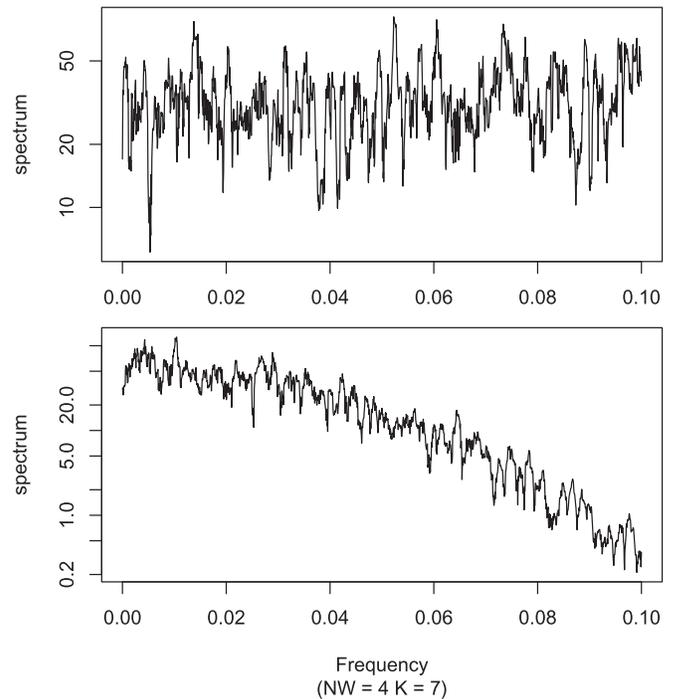
A seemingly popular approach is to interpolate the irregularly sampled time series to a mesh of regular times, then proceed with spectrum estimation as though the interpolated time series were what was originally observed. However, the consequences of the interpolation on the spectrum estimate and the inferred characteristics of the underlying spectrum should be considered. For example, Lepage (2009) demonstrated that spurious peaks in the spectrum can occur as a result of interpolation of an irregularly sampled time series to regular spacing.

Consider the following highly simplified example. Suppose we have a white noise process that follows

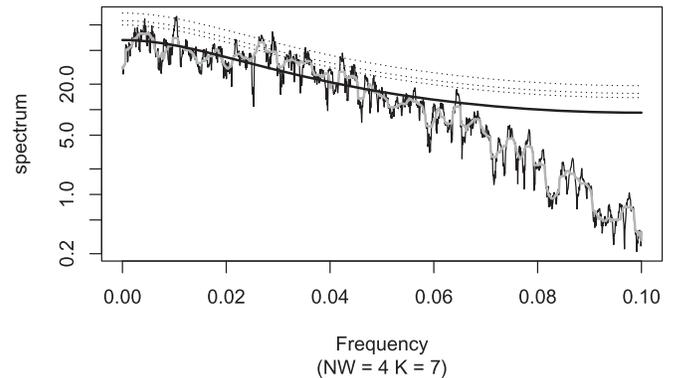
$$X_t \sim \mathcal{N}(0, 7^2)$$

for any  $t$ , where  $\mathcal{N}$  represents a normal distribution. We draw a realization from this process  $\mathbf{x}$  at observation times that are irregular, and in addition at observation times on a regular mesh for comparison. Next, we use linear interpolation to interpolate the times for the irregularly sampled time series. Figure 5 compares the multitaper spectrum estimate obtained by linear interpolation of the irregularly sampled series to the spectrum estimate that would have been obtained had the series been regularly sampled.

The linear interpolation has a noticeable effect on the overall shape of the spectrum. Although the underlying process is white, and should have resulted in a flat spectrum estimate, interpolation has reddened the estimate: high frequencies have



**Figure 5.** Comparison of multitaper spectrum estimates for a white noise process sampled at a time interval of five (top panel), and the same realization sampled irregularly and interpolated to the same time interval of five (bottom panel), showing the effect of interpolation.



**Figure 6.** Application of the Mann & Lees (1996) method to the interpolated white-noise series spectrum seen in the bottom panel of Figure 5. The light gray is the locally smoothed spectrum estimate used to fit the autoregressive spectrum in solid black. The dotted lines are 90th, 95th, and 99th percentiles of the background spectrum. Estimated spectrum values exceeding these percentiles would be interpreted as significant signals.

a lower associated spectrum estimate, and low frequencies have a higher associated spectrum estimate.

Here we might like to know whether the estimate of the power spectrum in the bottom panel of Figure 5 would lead to a false detection of signals (as was seen in Lepage 2009). Mann & Lees (1996) describe a popular method for estimation of background noise and signal detection in climatic time series. The method consists of partitioning the spectral representation of a series into a background autoregressive spectrum and a set of signals. The signals are then tested for significance by comparison to a null distribution. Figure 6 shows the application of the Mann & Lees (1996) method to the interpolated spectrum estimate. Although it appears that the background spectrum does not do a good job of describing the

spectrum estimate at higher frequencies, several low to moderate frequency signals are detected by the procedure. In cases where the method is applied in “black box” fashion, false detection could easily become a common problem.

Given that interpolation can have misleading effects on estimates of the spectrum, it seems that computational efficiency alone does not justify interpolation to regular times. Thus, there is motivation to develop spectral estimators that use the unevenly sampled data directly without interpolation.

The LS periodogram is widely used and considered to be a solution to the spectrum estimation problem for irregular sampling. However, like the periodogram, it has statistical issues, namely:

1. it is inconsistent,
2. it has spectral leakage that is often worse than the classical periodogram, and
3. it has spectral leakage that depends on the pattern of observation times and that is frequency specific.

In Scargle (1982), the author is concerned primarily with the detection and identification of a single periodic component from a background noise process that is of no interest. Because of this narrow focus, Scargle argues that inconsistency and spectral leakage are not of particular concern, but that the estimator could be improved by applying a window function to reduce spectral leakage if desired.

Here, we propose using interpolated DPSS as windows for the LS periodogram, and averaging the resulting spectral statistics in the same way as for the regular sampling multitaper. Although not theoretically optimal in the sense of Bronez (1988), the statistic can be made fast to compute using a nonequispaced fast Fourier transform (nfft; Press & Rybicki 1989).

The properties of the multitaper spectral likelihood for the evenly sampled case do not necessarily transfer to the irregular sampling case. In particular, the spectral window under irregular sampling varies as a function of the actual sample times and differs as a function of frequency (Scargle 1982). In the next section, we describe an approximate multitaper spectral statistic based on the nfft, the LS periodogram, and the DPSS tapers.

## 7. The Multitaper LS Statistic

In this section, we review the LS periodogram and describe how it can be computed using a fast algorithm. Next, we define a multitaper version of the LS periodogram and demonstrate that it has spectral windows that exhibit reduced spectral leakage.

### 7.1. Fast Computation of the LS Periodogram

The LS periodogram estimator (Lomb 1976; Scargle 1982) is

$$P_{LS}(f) = \frac{1}{2} \frac{\left(\sum_{i=1}^N x_i \cos \omega(t_i - \tau_{LS})\right)^2}{\sum_{i=1}^N \cos^2 \omega(t_i - \tau_{LS})} + \frac{1}{2} \frac{\left(\sum_{i=1}^N x_i \sin \omega(t_i - \tau_{LS})\right)^2}{\sum_{i=1}^N \sin^2 \omega(t_i - \tau_{LS})} \quad (9)$$

where  $\tau_{LS}$  is defined by

$$\tan 2\omega\tau_{LS} = \frac{\sum_{i=1}^N \sin 2\omega t_i}{\sum_{i=1}^N \cos 2\omega t_i}, \quad (10)$$

$N$  is the length of the time series,  $\omega = 2\pi f$  is the angular frequency,  $x_i$  is the  $i$ th observation, and  $t_i$  is the sampling time of the  $i$ th observation.

The LS estimator has the following properties.

1. The estimator is invariant to time translation. That is, if the  $t_i$  are replaced with  $t_i - t^*$  where  $t^*$  is a constant, the estimator is unchanged.
2. The estimator has the same distribution for a Gaussian white-noise process as the regularly sampled periodogram.
3. The estimator is equivalent to a least-squares fit of a sinusoid at a given frequency  $f$  to the data  $\mathbf{x}$  in the time domain.

Leroy (2012) points out that the nfft can be used to compute the LS periodogram using the decomposition of Press & Rybicki (1989). Following Press & Rybicki, the LS periodogram can be represented using the following sums:

$$s_h = \sum_{i=1}^N x_i \sin \omega t_i \quad c_h = \sum_{i=1}^N x_i \cos \omega t_i \quad (11)$$

$$s_2 = \sum_{i=1}^N \sin 2\omega t_i \quad c_2 = \sum_{i=1}^N \cos 2\omega t_i. \quad (12)$$

If the sums in Equations (11) and (12) can be computed rapidly, then the LS periodogram can be computed rapidly as well using the following terms:

$$A_1 = c_h \cos \omega \tau_{LS} + s_h \sin \omega \tau_{LS} \quad (13)$$

$$A_2 = \frac{N}{2} + \frac{1}{2} c_2 \cos 2\omega \tau_{LS} + \frac{1}{2} s_2 \sin 2\omega \tau_{LS} \quad (14)$$

$$B_1 = s_h \cos \omega \tau_{LS} - c_h \sin \omega \tau_{LS} \quad (15)$$

$$B_2 = \frac{N}{2} - \frac{1}{2} c_2 \cos 2\omega \tau_{LS} - \frac{1}{2} s_2 \sin 2\omega \tau_{LS}, \quad (16)$$

where  $\cos 2\omega \tau_{LS}$  and  $\sin 2\omega \tau_{LS}$  can be computed from Equation (12), and  $\cos \omega \tau_{LS}$  and  $\sin \omega \tau_{LS}$  can be computed using half-angle formulae from  $\cos 2\omega \tau_{LS}$  and  $\sin 2\omega \tau_{LS}$ . The LS periodogram is thus expressed as

$$P_{LS}(f) = \frac{1}{2} \left( \frac{A_1^2}{A_2} + \frac{B_1^2}{B_2} \right), \quad (17)$$

To compute the sums in Equations (11) and (12), we use the nfft library (Keiner et al. 2009), which provides a fast Fourier transform for irregularly sampled data. The nfft library defines the nonequispaced discrete Fourier transform (NDFT) as

$$x_{t_j^*} = \sum_{k=-L/2}^{L/2-1} \tilde{x}_k e^{-i2\pi t_j^* k}, \quad (18)$$

where the  $\tilde{x}$  are input complex Fourier coefficients, and  $x_{t_j^*}$  are the time series as observed at times  $t_j^*$ , standardized so that the times are in the interval  $[-1/2, 1/2)$ . Writing the NDFT in

matrix notation, we have

$$\mathbf{x} = \mathbf{A}\tilde{\mathbf{x}} \quad (19)$$

$$A_{jk} = e^{-i2\pi kt_j^*}. \quad (20)$$

(Note that  $k$  here represents the index over which the summation occurs in Equation (18), not the  $k$ th taper.)

To make the nfft library useful for time-series analysis, we need to consider the  $\mathbf{x}$  as an input and the Fourier representation as an output. The adjoint (conjugate transpose) of the matrix  $\mathbf{A}$  multiplied by the time-series vector  $\mathbf{x}$  is

$$\tilde{\mathbf{z}} = \mathbf{A}^H \mathbf{x}, \quad (21)$$

which is

$$\tilde{z}_k = \sum_{j=1}^N x_{t_j^*} e^{i2\pi kt_j^*}. \quad (22)$$

Thus, the adjoint NDFT performs the necessary Fourier transform of the input time series  $\mathbf{x}$ , except that it seems to return the complex conjugate of the usual transform. In fact, because of the centering of the times to the interval  $[-1/2, 1/2]$ , there is a further rotation of  $k\pi$  due to the  $-1/2$  offset. This additional rotation requires a further transformation of  $\tilde{z}_k$ ,

$$\tilde{z}_k^* = \begin{cases} \overline{\tilde{z}_k} & k \text{ is even} \\ -\tilde{z}_k & k \text{ is odd.} \end{cases} \quad (23)$$

This subtlety is not mentioned by Leroy (2012).

## 7.2. MTLs: Multitapering the LS Periodogram

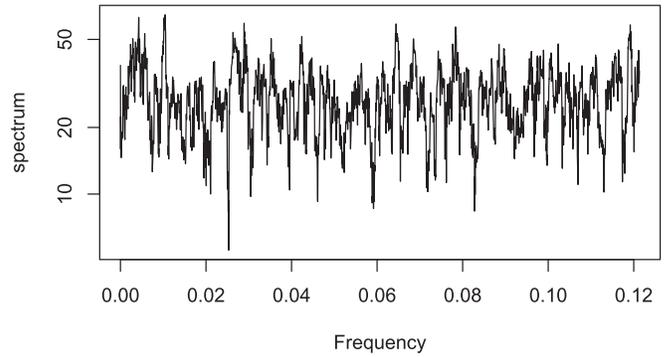
Scargle (1982) suggests that the spectral leakage in the irregular sampling case can be improved through data tapering. The Thomson multitaper limits the spectral leakage to a negligible amount in the case of regularly sampled times, so that multitaper spectral statistics spaced far enough apart in frequency are nearly uncorrelated. Here we generalize this approach to unevenly sampled times.

Our approach for multitapering using the DPSS for irregular time series is as follows.

1. We compute the Slepians at a mesh corresponding to an average sampling interval  $\overline{\Delta t} = T/N$ , where  $T$  is  $t_N - t_1$  and  $N$  is the number of times in the series. The DPSS on the regular mesh are denoted as before:  $v_t^{(k)}(N, W)$ , where  $k$  is the order of the sequence.
2. The taper weights at intermediate points corresponding to the irregular times are obtained by interpolation using a cubic spline. We denote these interpolated weights as  $v_t^{*(k)}(N, W)$ .
3. The interpolated weights for each DPSS order are renormalized so that  $\sum_{t=t_1}^{t_N} (v_t^{*(k)}(N, W))^2 = 1$  for each  $k$ .

The approximation in step 2 is best when the set of times is close to regular—for example, there are no long gaps—which we assume is true for many astronomical data sets (e.g., from the Kepler or TESS spacecraft). The resulting multitaper statistics are then used as in the regular sampling case described previously.

The MTLs is a compromise between the optimal quadratic estimator of Bronez (1988), which is difficult and costly to compute, and the untapered LS periodogram, which can be made fast to compute but does not achieve minimum spectral leakage out of band. Because the spectral leakage performance



**Figure 7.** MTLs estimate for a white noise process sampled irregularly. The spectrum estimate is comparable in shape—that is, flat—to the estimate that would have been obtained under regular sampling (Figure 5, top). Moreover, the average spectrum value and variation in the estimate are similar.

of the MTLs depends on the observation times  $\mathcal{D}$  as well as the frequency band, it is crucial that the spectral pseudowindows (see Section 7.3) be examined across a range of frequencies for a given set of irregular times for which the statistic is to be computed.

We now revisit the earlier example of an irregularly sampled white-noise process, interpolated to regular sampling (bottom of Figures 5 and 6 in Section 6). We saw that interpolation to regular sampling induced an overall shape to the spectrum estimate and could lead to false signal detection. As shown in Figure 7, the MTLs spectrum estimate is able to overcome the irregular sampling without introducing any shape or other spurious features.

## 7.3. Checking Pseudowindows

Scargle (1982) defines and examines pseudowindows and shows that in some cases, spectral leakage can be significant even at remote frequencies. This suggests that checking pseudowindows is good practice in all cases of spectrum estimation under irregular sampling, whether using the MTLs, LS, the method of Bronez, or any other.

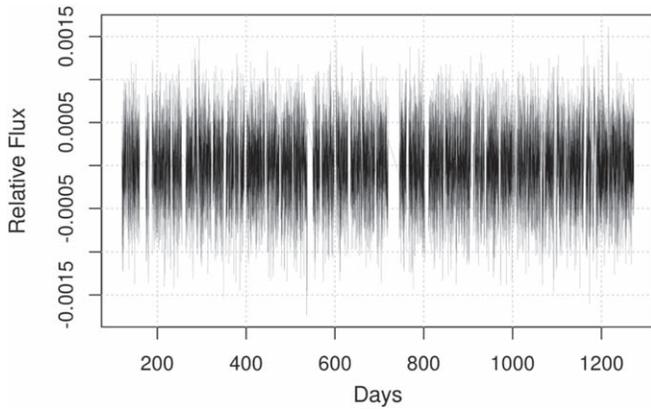
A pseudowindow is the frequency-domain response of the estimator to a pure sinusoidal input at a particular frequency. In order to check pseudowindows, it suffices to create an artificial signal at a given frequency, sample it at the same times as the original irregular series, and examine the resulting spectral statistics. This reveals the pattern of spectral leakage out of the band defined by the Slepians.

When using the MTLs for irregular times, the DPSS are interpolated, and the resulting spectral leakage depends on both the particular set of irregular times and the frequency band. Although the spectral leakage properties of the MTLs should be close to optimal for nearly regular sampling, it is a good idea to check the pseudowindows for a range of bands to avoid errors in the interpretation of the MTLs statistics.

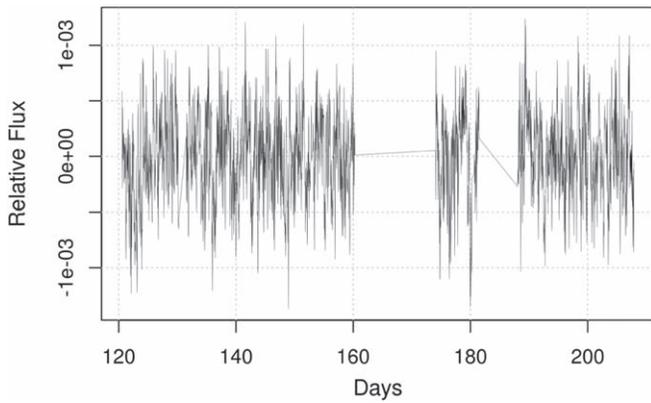
## 8. Application to Kepler Data

It is hoped that the MTLs spectral statistic can be used as an approximate direct estimator of the spectrum in cases where the observation times are known but irregular. In particular, we believe this estimator may be usefully applied to study time-varying processes in astronomical data.

Here, we perform a preliminary analysis of Kepler time-series data from an RG star, Kepler-91. Figure 8 shows the time



**Figure 8.** Relative flux data for the RG star Kepler-91.



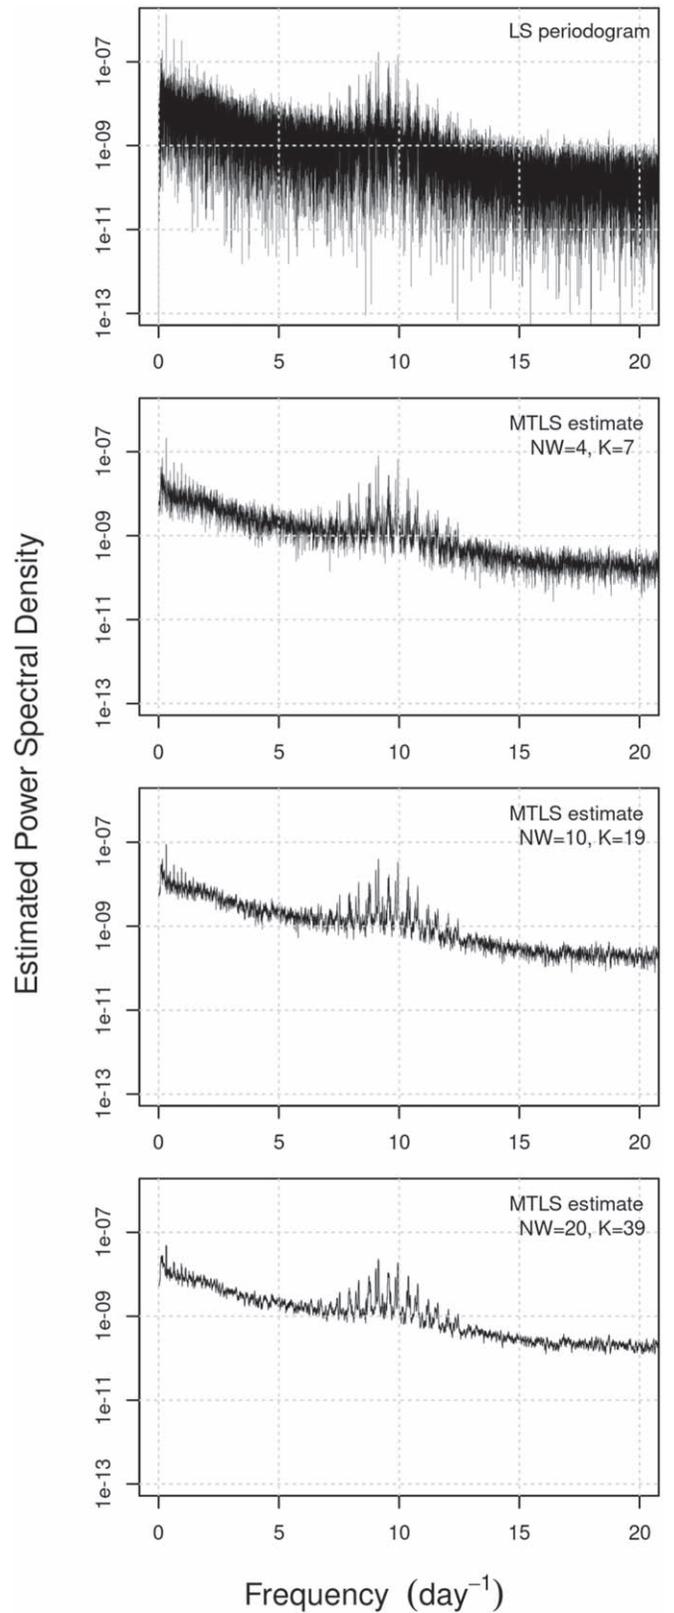
**Figure 9.** Zoom-in to the first 3000 data points shown in Figure 8.

series of the relative flux measurements measured by the Kepler spacecraft (preprocessing done by T. Bedding 2020, private communication), and Figure 9 shows detail by zooming in on the first 3000 observations. The star has a transiting planet, Kepler-91b, with an orbital period of approximately 6.25 days (Batalha et al. 2013; Lillo-Box et al. 2014), in addition to other periodic variations in flux indicative of stellar oscillations. This system has been extensively studied previously (e.g., Lillo-Box et al. 2014; Barclay et al. 2015; Placek et al. 2015; Budding et al. 2016).

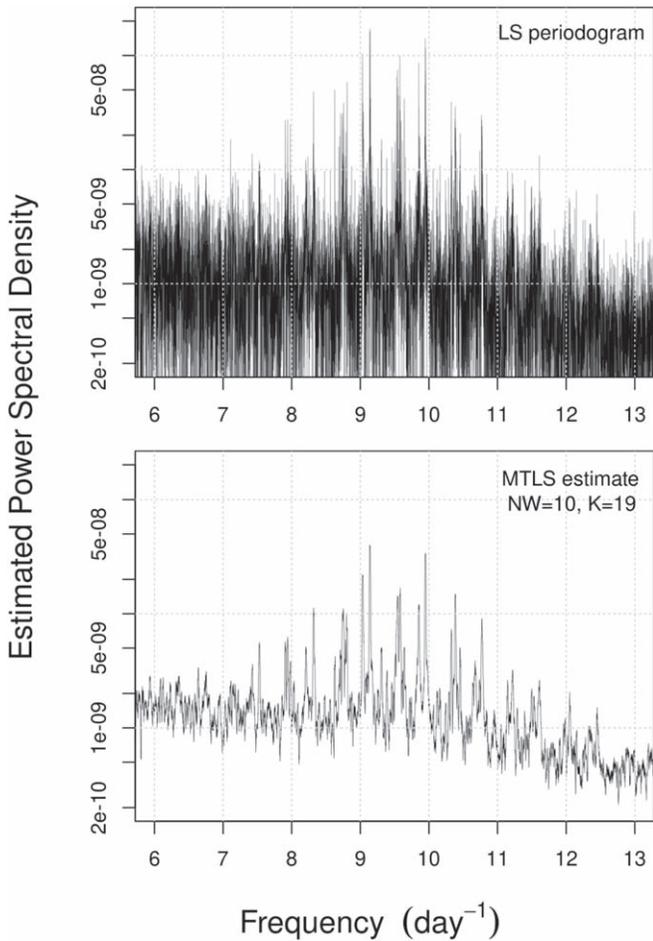
By using measurements of total flux over time from the Kepler spacecraft and estimating the power spectrum, asteroseismologists can unravel stellar properties (e.g., helium core burning in RG stars; Bedding et al. 2011). This inference is based on the peak mode and the spacing of adjacent line components in the spectrum of the time series of the flux, which relies on accurate identification of signals from noise in the spectrum estimate.

The data in Figure 8 are sampled irregularly due to the orbit of the spacecraft around the Sun and correction to the heliocentric barycenter. There are also some longer time gaps in the series (Figure 8). For this preliminary analysis, the series was analyzed in its entirety, across these longer gaps.

We obtain estimates of the power spectrum of the relative flux using MTLs and three different frequency bandwidth values (corresponding to  $NW = 4, 10, \text{ or } 20$ ), and also using the LS periodogram for comparison. Figure 10 shows the LS periodogram and the MTLs estimate for increasing numbers of



**Figure 10.** Four spectrum estimates using the Kepler-91 flux data from Figure 8. The top panel is the usual LS periodogram. The second panel is the MTLs spectrum estimate with  $NW = 4$  and  $K = 7$  tapers. The third panel is the MTLs spectrum estimate with  $NW = 10$  and  $K = 19$  tapers. The bottom panel is the MTLs spectrum estimate with  $NW = 20$  and  $K = 39$  tapers. The progression shows the trade-off between frequency resolution and estimator variance.



**Figure 11.** Zoom-in of the top and third panel of Figure 10 to illustrate the reduced variance in the MTLS-estimated power spectrum.

tapers  $k$ . The increasing number of tapers with increasing bandwidth has a noticeable effect on the variance of the estimate. In the extreme case of the LS periodogram (top), signals appear to be less obvious than for even the narrow-bandwidth MTLS (second panel).

Comparing the three MTLS spectrum estimates, the one with a time-bandwidth parameter of 10 appears to retain the relative amplitude of any signals present while providing a good separation relative to the noise in the estimate. Figure 11 shows a zoom-in of the top and third panel in Figure 10 for comparison.

Although some work has been done on the choice of bandwidth parameter under some assumptions about the underlying true spectrum (Haley & Anitescu 2017), there is no universally applicable method for determining the optimal choice for this parameter.

The MTLS estimate of the spectrum appears to confirm previous estimates of the orbital period of the transiting planet Kepler-91b (Figure 12(b)). In particular, there is a peak in the spectrum estimate at a period that matches previous estimates (i.e., the planet orbital period of 6.25 days; Batalha et al. 2013; Lillo-Box et al. 2014), as well as a series of harmonics, which is to be expected due to the shape of the light curve attenuation as the planet passes in front of the star. There are also a number of apparent signals or line components in the frequency band from 8 to 12 cycles per day, approximately 92–140  $\mu\text{Hz}$  (Figure 12(c)). The high-frequency portion of the MTLS

estimate (seen in Figures 12(a) and (d)) should not be trusted, for reasons explained next.

The apparently large spectrum values at high frequency are potentially due to leakage from elsewhere in the spectrum. To investigate this possibility, we generated synthetic signals across the frequency range and calculated the response of the MTLS statistic for each.

The frequency response shows a pattern of aliasing beyond about 10 cycles per day (Figure 13). This pattern was verified by computing the LS periodogram of the same synthetic signals using an independent computer code (Ruf 1999). The aliasing effect is a result of the particular pattern of times.

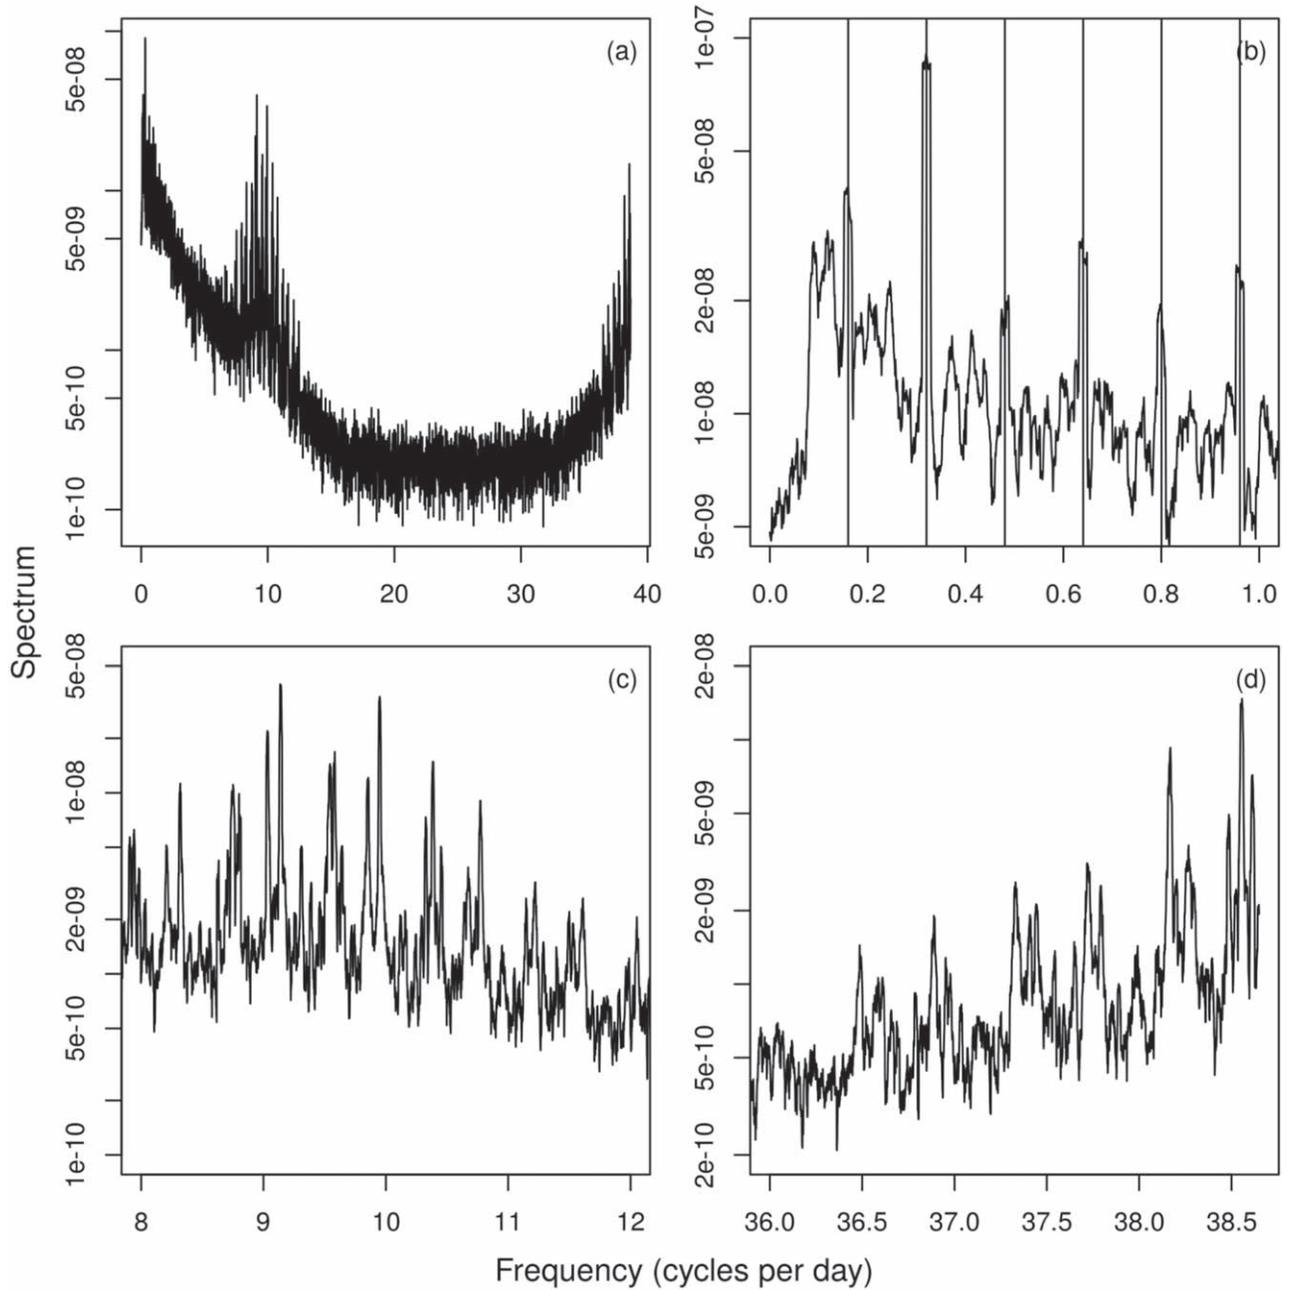
## 9. Discussion

Unlike many branches of statistics in which observations are assumed to be realizations of independent or conditionally independent random variables, time-series analysis is preoccupied with the dependence of sequences of random variables. Unfortunately for time-series analysts, the random variables underlying the observations are fleeting, and it is impossible to obtain repeated measures of a random variable corresponding to a moment in time. Instead, methods of time-series analysis make simplifying assumptions about the nature of the underlying random variables and their covariance structure. The assumption of stationarity, for example, allows measurement to begin at any point in time, because if the assumption holds then the joint distribution of random variables depends only on their relative position in time. The stationarity assumption allows the covariance to be summarized as a function of time separation.

The autocovariance can be obtained from the spectrum if desired, meaning that spectral statistics or spectrum estimates can be used instead of estimates of the autocovariance. Inference in the spectral domain is convenient because many natural processes are approximately periodic, and show up as signals in the spectrum of a time series. This makes non-parametric spectrum estimation of the kind described here an important contributor to the exploratory data analysis toolbox. In cases where little is known beforehand, a good estimate of the spectrum can provide some insight into the underlying processes, while suggesting parametric approaches that could be successful in describing the data. Like other exploratory approaches, analysis might stop at the spectrum estimate, or might proceed to additional modeling.

Moreover, spectral statistics can be made to be approximately independent using data tapers, so that the total likelihood of a set of spectral statistics can be approximated using a product likelihood. This could greatly simplify inference for parametric models of stellar processes that imply a theoretical spectrum—given a model with unknown parameters, the theoretical spectrum could be used to calculate the likelihood for each of the spectral statistics. Because of their approximate independence, the total likelihood is the product of each, and maximum-likelihood estimates or Bayesian posterior distributions for model parameters are easily obtained. A more general example of this approach is described in Springford (2017).

For irregularly sampled time series, the LS spectrum estimate is widely used. Scargle (1982) was motivated by the case of a single periodic component embedded in a white noise process, and for this type of process the LS periodogram has its place. However, the spectral leakage properties of the estimator



**Figure 12.** MTLS spectrum estimate with  $NW = 10$  and  $K = 19$  tapers. Panel (a) is the spectrum estimate across all frequencies. Panel (b) is the low-frequency portion of the spectrum; lines show the fundamental frequency (1/6.25 days) and harmonics of the previously identified transiting planet (Batalha et al. 2013; Lillo-Box et al. 2014). Panel (c) is the band from 8 to 12 cycles per day, showing many apparent signals. Panel (d) is the highest-frequency portion of the spectrum. This portion of the spectrum should be interpreted with caution due to the possibility of leakage or aliasing from distant parts of the spectrum.

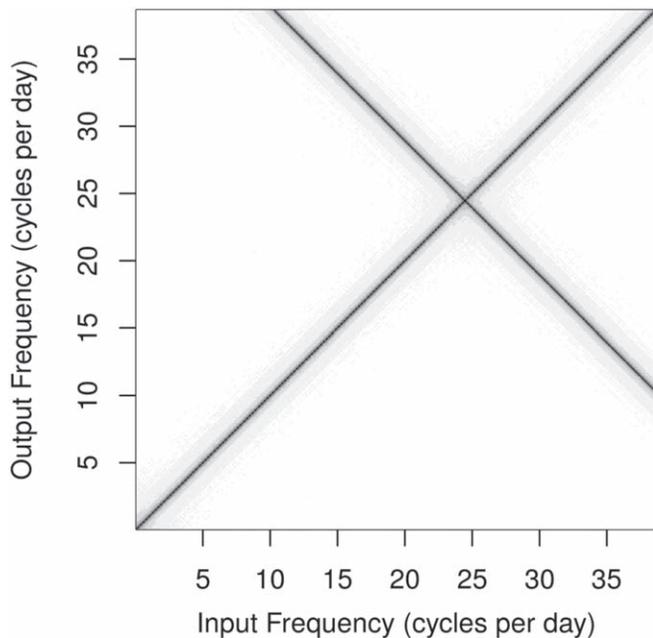
can be extremely poor, and if the underlying spectrum is not of the type envisioned then it is possible to be misled by the result.

Scargle (1982) suggests tapering the time series to control spectral leakage before computing his periodogram, a suggestion that we take to the same conclusion as Thomson (1982) by multitapering using interpolated DPSS tapers.

The generalized LS periodogram (GLS; Zechmeister & Kürster 2009) and the Bayesian GLS (Mortier et al. 2015) are increasingly popular in astronomy. While these methods offer improvements to the standard LS, these improvements are complementary to those offered by tapering, primarily: (1)

measurement of error-derived weights and (2) separate mean values by frequency. Analysts who are interested in employing the GLS or BGLS for these reasons can also benefit by tapering or multitapering.

The MTLS statistic we introduce here matches the Thomson multitaper under regular sampling. Deviations from regular sampling introduce frequency-specific differences that depend on the particular set of observation times. For a given set of times, careful examination of the pseudowindows is recommended. If there appear to be strong signals or a large dynamic range in the spectrum, frequencies with high associated spectral



**Figure 13.** Frequency response of the MTLs for the Kepler-91 time series. Frequencies above about 10 cycles per day show pseudo-aliasing to higher frequencies.

statistics should be examined for their leakage properties because they are liable to contribute to spectral statistics at other frequencies in the spectrum.

Unlike in the case of regular sampling, leakage (or aliasing) to distant frequencies is possible for irregular sampling. In particular, high-frequency parts of the spectrum may not be trustworthy and, if judged to be suspect, should be ignored.

In the example of the RG star data, there appear to be many signals present, and the overall shape of the spectrum at high frequencies seems to be affected by a kind of aliasing. This limits the effective Nyquist frequency to something much smaller than  $1/2\Delta t = N/2T$ .

One possibility is that the aliasing is due to regular gaps in the series. If this is the case, an MTLs spectrum estimate could be computed for each contiguous portion of the series, and averaged as in a non-overlapped Welch estimate (Welch 1967). Although this does not resolve the higher-frequency components of the spectrum estimate, if the aliasing is removed it makes interpretation of the estimate more straightforward. Regardless, the technique of separating time series according to the location of long gaps could be helpful in certain cases where the long gaps make the high-frequency portion of the spectrum estimate unreliable.

There are obvious extensions to the current MTLs as an estimator of the spectrum, based on analogies to the use of the multitaper estimator in the regular sampling case.

The eigencoefficients (Thomson 1982) are the complex-valued result of taking the discrete Fourier transform of the DPSS-windowed time series. The eigencoefficients are the basis of several extensions to the multitaper method in the regular sampling case. The Thomson F-test (Thomson 1982) is a test for the presence of periodic components that is based on a regression estimate of the eigencoefficients. The variance explained by the regression is compared to the residual variance of the estimate to form the F-test. In principle a similar test could be constructed for the MTLs, but this might

require an adjustment to the method of computation described in Section 7.2. A targeted approach in which frequencies of interest are tested for periodicity might be the most efficient if the fast computing advantage is lost.

Coherence, the frequency-domain analog of correlation, would also be a useful tool in cases where an estimate of the relationship between two or more time series is of interest. Jackknife error estimates for the spectrum estimate, periodic F-test, and coherence are also possible, and would provide insight into the uncertainties and biases associated with each (Thomson 2007).

## 10. Conclusion

We have described a multitapered version of the popular LS estimator of the spectral density of an irregularly sampled time series, which we call the MTLs. The MTLs was shown to have limited spectral leakage, but because the leakage depends on the particular pattern of observation times, examination of the pseudowindows is necessary in any particular application.

As an estimator, the MTLs was applied to a time series of photometry data and displayed the reduced variance and spectral leakage characteristic of the multitaper estimator for regularly spaced data. However, examination of the frequency response (pseudowindows) revealed aliasing and a much smaller effective Nyquist frequency than might have been expected.

The subfield of time-series analysis in astronomy is growing with data sets related to exoplanets, asteroseismology, and transients through Kepler, TESS, and ZTF. The subfield is destined to become even more prevalent as survey projects such as the LSST on the Vera C. Rubin Observatory begin in the 2020s. When scientific questions require the identification of more than a single mode in the power spectrum, the LS estimator is a weak tool because it suffers from spectral leakage, inconsistency, and bias. The MTLs estimator, however, shows promise for applications in astronomy where multiple modes are of interest.

In a future work, we will explore other applications of this method to astronomical data and compare it to other recent advances for estimating the power spectrum. Equally important is the development of an F-test for the MTLs statistic, which determines the probability that a frequency signal in the spectrum estimate originates from a genuine periodic process. The F-test provides a way to test for the presence of real periodic signals in time-series data that may contain evidence of, e.g., exoplanets, asteroseismic modes, and stellar rotation.

This research was funded in part by a postgraduate scholarship to the first author from the Natural Sciences & Engineering Research Council of Canada (held at Queen's University). This research was also supported by the Washington Research Foundation (WRF) and by a Data Science Environments project award from the Gordon and Betty Moore Foundation (Award #2013-10-29) and the Alfred P. Sloan Foundation (Award #3835) to the University of Washington eScience Institute. Support was also made possible by a Moore, Sloan, and WRF Innovation in Data Science Postdoctoral fellowship and a NSERC Postdoctoral Fellowship to the second author, both held at the University of Washington. The second author would like to thank the Institute for Data Intensive Research in Astrophysics & Cosmology (DIRAC) for supporting this research. We thank

T. Bedding for supplying the preprocessed Kepler data and for providing comments that helped improve this manuscript. We also thank D. Riebert for helpful discussions, and an anonymous reviewer whose comments and suggestions significantly improved the manuscript.

*Software:* R Statistical Software Environment (R Development Core Team 2012), including the following packages:

1. *stats* (part of base R)
2. *MASS: Modern Applied Statistics with S* (Venables & Ripley 2002)
3. *multitaper* (Rahim et al. 2014)
4. *ggplot2* (Wickham 2016)
5. *TSA: Time Series Analysis* (Chan & Ripley 2018)

Other software used:

1. *nfft* library (Keiner et al. 2009)

### ORCID iDs

Aaron Springford  <https://orcid.org/0000-0001-7179-9049>

Gwendolyn M. Eadie  <https://orcid.org/0000-0003-3734-8177>

### References

- Babu, P., & Stoica, P. 2010, *Digital Signal Processing*, 20, 359
- Barclay, T., Endl, M., Huber, D., et al. 2015, *ApJ*, 800, 46
- Batalha, N. M., Rowe, J. F., Bryson, S. T., et al. 2013, *ApJS*, 204, 24
- Bedding, T., Mosser, B., Huber, D., et al. 2011, *Natur*, 471, 608
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2018, *PASP*, 131, 018002
- Blackman, R. B., & Tukey, J. W. 1958, *The Bell System Technical Journal*, 37, 185
- Blackman, R. B. R. B. 1959, *The Measurement of Power Spectra from the Point of View of Communications Engineering* (New York: Dover Publications)
- Box, G. E. P., & Jenkins, G. M. 1970, *Time Series Analysis Forecasting and Control* (San Francisco, CA: Holden-Day)
- Brillinger, D. 1981, *ITASS*, 29, 1075
- Bronze, T. 1988, *ITASS*, 36, 1862
- Budding, E., Püsküllü, Ç., Rhodes, M. D., Demircan, O., & Erdem, A. 2016, *Ap&SS*, 361, 17
- Chan, K.-S., & Ripley, B. 2018, *TSA: Time Series Analysis*, <https://CRAN.R-project.org/package=TSA>
- Chave, A. D. 2019, *GeoJI*, 218, 2165
- Einstein, A. 1914, *Archives des sciences physiques et naturelles*, 37, 254
- Fodor, I., & Stark, P. 2000, *ITSP*, 48, 3472
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, 131, 078001
- Haley, C. A. L., & Anitescu, M. 2017, *ISPL*, 24, 1696
- Harris, F. J. 1978, *IEEEP*, 66, 51
- Keiner, J., Kunis, S., & Potts, D. 2009, *ACM Transactions on Mathematical Software (TOMS)*, 36, 19
- Khintchine, A. 1934, *MatAn*, 109, 604
- Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, *ApJL*, 713, L79
- Komm, R. W., Gu, Y., Hill, F., Stark, P. B., & Fodor, I. K. 1999, *ApJ*, 519, 407
- Lepage, K. 2009, PhD thesis, Queen's Univ.
- Leroy, B. 2012, *A&A*, 545, A50
- Lillo-Box, J., Barrado, D., Henning, T., et al. 2014, *A&A*, 568, L1
- Lillo-Box, J., Barrado, D., Moya, A., et al. 2014, *A&A*, 562, A109
- Lomb, N. R. 1976, *Ap&SS*, 39, 447
- LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
- Mann, M. E., & Lees, J. M. 1996, *CiCh*, 33, 409
- Mortier, A., Faria, J. P., Correia, C. M., Santerne, A., & Santos, N. C. 2015, *A&A*, 573, A101
- Percival, D., & Walden, A. 1993, *Spectral Analysis for Physical Applications* (Cambridge: Cambridge Univ. Press)
- Placek, B., Knuth, K. H., Angerhausen, D., & Jenkins, J. M. 2015, *ApJ*, 814, 147
- Press, W. H., & Rybicki, G. B. 1989, *ApJ*, 338, 277
- R Development Core Team 2012, *R: A Language and Environment for Statistical Computing*, <http://www.R-project.org/>
- Rahim, K. J., Burr, W. S., & Thomson, D. J. 2014, PhD thesis, Queen's Univ., <https://CRAN.R-project.org/package=multitaper>
- Rayleigh, L. 1903, *PMag*, 5, 238
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *JATIS*, 1, 014003
- Ruf, T. 1999, *Biological Rhythm Research*, 30, 178
- Scargle, J. D. 1982, *ApJ*, 263, 835
- Schuster, A. 1898, *JGR*, 3, 13
- Slepian, D. 1978, *The Bell System Technical Journal*, 57, 1371
- Smith-Boughner, L., & Constable, C. 2012, *GeoJI*, 190, 1404
- Springford, A. 2017, PhD thesis, Queen's Univ., <http://hdl.handle.net/1974/23784>
- Stoica, P., & Sandgren, N. 2006, *Digital Signal Processing*, 16, 712
- Stoica, P., & Sundin, T. 1999, *CSSP*, 18, 169
- Thomson, D. 1982, *IEEEP*, 70, 1055
- Thomson, D. 2007, *ISPM*, 24, 20
- Thomson, D. J. 1990, *RSPTA*, 330, 601
- Thomson, D. J., & Haley, C. L. 2014, *RSPSA*, 470, 21
- Thomson, D. J., MacLennan, C. G., & Lanzerotti, L. J. 1996, *Natur*, 381, 32
- Thomson, D. J., & Vernon, F. L. 2015, *GeoJI*, 202, 1690
- Thomson, D. J., & Vernon, F. L. 2016, *Proc. IEEE*, 104, 2220
- Tukey, J. W., & Hamming, R. W. 1949, in *The Collected Works of John W. Tukey*, Volume I: Time Series, ed. D. R. Brillinger (London: Chapman and Hall), 129
- VanderPlas, J. T. 2018, *ApJS*, 236, 16
- Venables, W. N., & Ripley, B. D. 2002, *Modern Applied Statistics with S* (4th ed.; New York: Springer)
- Welch, P. 1967, *IEEE Transactions on Audio and Electroacoustics*, 15, 70
- Wickham, H. 2016, *ggplot2: Elegant Graphics for Data Analysis* (New York: Springer)
- Wiener, N. 1930, *Acta Math.*, 55, 117
- Zechmeister, M., & Kürster, M. 2009, *A&A*, 496, 577