Spectral and Cross-Spectral Analysis - a Tutorial for Psychologists and Social Scientists

M.J. Vowels

University of Surrey

L.M. Vowels

University of Southampton

N.D. Wood

University of Kentucky

Author Note

Matthew J. Vowels, Centre for Computer Vision, Speech and Signal Processing, University of Surrey; Laura M. Vowels, Department of Psychology, University of Southampton; Nathan D. Wood, Department of Family Sciences, University of Kentucky

Correspondence concerning this article should be addressed to Laura M. Vowels, School of Psychology, University of Southampton, Highfield Campus, SO17 1BJ, UK. E-mail: l.vowels@soton.ac.uk

Abstract

Social scientists have become increasingly interested in using intensive longitudinal methods to study social phenomena that change over time. Many of these phenomena are expected to exhibit cycling fluctuations (e.g., sleep, mood, sexual desire). However, researchers typically employ analytical methods which are unable to model such patterns. We present spectral and cross-spectral analysis as means to address this limitation. Spectral analysis provides a means to interrogate time series from a different, frequency domain perspective, and to understand how the time series may be decomposed into their constituent periodic components. Cross-spectral extends this to dyadic data and allows for synchrony and time offsets to be identified. The techniques are commonly used in the physical and engineering sciences, and we discuss how to apply these popular analytical techniques to the social sciences while also demonstrating how to undertake estimations of significance and effect size. In this tutorial we begin by introducing spectral and cross-spectral analysis, before demonstrating its application to simulated univariate and bivariate individual- and group-level data. We employ cross-power spectral density techniques to understand synchrony between the individual time series in a dyadic time series, and circular statistics and polar plots to understand phase offsets between constituent periodic components. Finally, we present a means to undertake non-parameteric boostrapping in order to estimate the significance, and derive a proxy for effect size. A Jupyter Notebook (Python 3.6) is provided as supplementary material to aid researchers who intend to apply these techniques.

Keywords: spectral analysis, cross-spectral analysis, dyadic, time series analysis

Spectral and Cross-Spectral Analysis - a Tutorial for Psychologists and Social Scientists

Introduction

Social scientists are often interested in dynamic systems and social phenomena that change in highly complex ways over time (Bringmann et al., 2016; Nesselroade & Ram, 2004; Vallacher & Nowak, 1997; Wang et al., 2012). Despite this, social scientists have traditionally relied on cross-sectional data that focus exclusively on variation between, not within, individuals and which cannot capture the complexity and the dynamic nature of the system (Molenaar, 2004). In recent years, there has been an increase in intensive longitudinal studies in which a relatively large number of time points are collected to assess the phenomenon of interest (Bolger et al., 2003; Mehl & Conner, 2012; Scollon et al., 2003). These studies enable researchers to study within-person processes as they unfold over time, in order to determine when, how, and why people change (Bolger & Laurenceau, 2013; Bringmann et al., 2016; Ferrer & Nesselroade, 2003; Molenaar, 2004; Molenaar & Campbell, 2009; Nesselroade & Molenaar, 2010; Nesselroade & Ram, 2004). However, despite a surge in the collection of intensive longitudinal data, the majority of the studies still employ overly reductionist analyses that fail to "let the data speak" (van der Laan & Rose, 2011). At best, some models may constitute simple functional forms such as quadratic, cubic, or exponential, and a linear trend to describe a growth pattern (Bringmann et al., 2016), despite the fact that the characteristics of interest may exhibit more complex relationships that may not be uncovered with these techniques (Vallacher & Nowak, 1997). Furthermore, given psychologists are interested in understanding how individuals change over time (Molenaar, 2004), being able to estimate within-person variation in their time series instead of simply pooled averages is important.

Some authors have pushed for more comprehensive mathematical approaches to model societal phenomena (Gottman et al., 2003). Several analytical methods that enable the study of complex time-varying phenomena have been proposed. These include the use

of autoregressive (Bringmann et al., 2016), dynamic structural equation (Asparouhov et al., 2017), hierarchical Bayesian (Driver & Voelkle, 2018), and penalized spline models (Won Suk et al., 2019). However, there are several limitations with these analytical methods, such as requiring prior knowledge of when the time series is likely to change, not accounting for multiple rates of change that may be ongoing at the same time, and not addressing cyclical patterns in fluctuation over time.

Stroud and Booth (2003) define a periodic signal as one whose "function values repeat at regular intervals of the independent variable" (p.183). Many phenomena cycle and oscillate, and are misspecified by a straight line or a polynomial. This is illustrated in Figure 1, where the Ordinary Least Squares fit indicates a significant downward trend, despite the slope of this trend varying depending on the start and end points of the sample period. In other words, if we were to extend the x-axis of the graph in Figure 1 and run the OLS once again, the trend would change direction. Because of this, even when researchers do not have access to enough data to uncover cyclicity in their time series, considering potential theoretical fluctuation in the phenomenon of interest is important. In the case where researchers have intensive longitudinal data, being able to examine potential periodic components in the data is advantageous.

There are a number of analytical methods that can be used to identify periodically fluctuating phenomena over time, notably, latent differential equations (Hu et al., 2014), spectral analysis (Oppenheim & Schafer, 1999), and dynamic factor models which operate within the structural equation modeling framework (Molenaar, 1985, 1987; Molenaar et al., 1992). One of the models using latent differential equations includes the damped oscillator model (Chow et al., 2005; Deboeck et al., 2008; Hu et al., 2014). The technique can be used to model variables that tend to oscillate around some typical value, or equilibrium. In the model, the interplay between frequency and damping determines how quickly the phenomenon returns to equilibrium (Nesselroade & Boker, 1994). When there is no damping in the model, in other words when a phenomenon fluctuates around a mean

without returning to an equilibrium, the integral solution of the damped oscillator model is equal to the sinusoidal model of the spectral analysis (Chow et al., 2005). Thus, in the simple case, these two methods provide equivalent results and the use of spectral analysis may be preferred. Spectral analysis has two key strengths above the damped oscillator model. First, spectral analysis estimates all possible frequencies simultaneously whereas the oscillator model only extracts the most dominant frequency. Second, when researchers are interested in understanding synchrony, or shared fluctuation in the frequency domain, across two time series, cross-spectral analysis can estimate the phase information easily. In contrast, researchers would need to run separate damped oscillator models for each time series and then compute the cross-phase (Chow et al., 2005).

The use of spectral analysis dates at least as far back as the Babylonian times (around 4000 years ago) for making astronomical predictions (Oppenheim & Schafer, 1999). Spectral analysis is currently widely and frequently used by scientists of physics, meteorology, marine science, and others (Abramovich et al., 2000; Chatfield, 2005). However, while there are many phenomena that fluctuate periodically over time that may be of interest to social scientists (e.g., mood, stress, sexual desire, calendar events, monthly payments, sleeping, and eating), there has been relatively little use of spectral analysis in the literature. Exceptions include identification of weekly periodicity in characteristics of mood and personality (Campbell et al., 1991; Larsen, 1987; Larsen et al., 2009; Larsen & Kasimatis, 1990) and fluctuation in sexual desire (Vowels et al., 2018). Dynamic factor models (Molenaar, 1985, 1987) can also be conducted in the frequency domain which require a spectral decomposition of the original time series into the frequency domain as the first step in the analysis process. The damped oscillator model, which is one model using latent differential equations, in its simple form also provides results equivalent to spectral analysis. Additionally, spectral analysis is often used especially in economics to de-seasonalize the data as a data preparation step prior to conducting the main analysis (e.g., see Geweke, 1978; Sims, 1974). In other words, the regularly fluctuating components

are removed from a time series in order to detect a clearer trend.

In addition to identifying periodic change over time in one variable, spectral analysis can also be used to identify shared fluctuation or temporal coordination across two or more variables (cross-spectral analysis). There are various societal phenomena that have been suggested to exhibit such temporal coordination. These include the rhythm and frequency of gaze between mothers and their newborn infants (Gottman, 1979; Gottman & Ringland, 1981; Lester et al., 1985; R. M. Warner et al., 1987); coregulation between romantic partners' physiology and affect (Butner et al., 2007; Gottman et al., 2003; Gottman et al., 2002; Helm et al., 2012; Liu et al., 2013; Papp et al., 2013; Saxbe & Repetti, 2010; Steele & Ferrer, 2011); and interpersonal synchrony between strangers in peer conversations, cooperative games, and group activities (Guastello et al., 2006; Kleinspehn-Ammerlahn et al., 2011; Muller & Lindenberger, 2011; Wiltermuth & Heath, 2009). Only a small number of studies have employed cross-spectral analysis to study synchrony (Gottman & Ringland, 1981; Lester et al., 1985; Liu & Molenaar, 2016; Sadler et al., 2009). For example, in a study of 30 couples, Gottman and Ringland (1981) used cross-spectral analysis and found coherence between partners' affect ratings and their physiological responses. Cross-spectral analysis has also been used to investigate the decreased coherence in brainwave activity in patients with schizophrenia (Bert et al., 2010; Yeragani et al., 2006).

While there are a handful of existing studies in the social sciences that have used spectral and cross-spectral analyses, there is currently very little guidance on how to undertake them (for exceptions, see Gottman, 1979; R. Warner, 1998). Therefore, the purpose of this tutorial is to provide both a theoretical as well as a practical introduction to spectral and cross-spectral analyses. Our hope is that readers can get an intuitive understanding of the theory behind spectral analysis as well as how to conduct analyses in the frequency domain. There are various ways to approach spectral analysis, but we have chosen to use one of the most well-known approaches that derives from the work of Fourier.

This approach is extremely well established in the physical sciences, and readers are encouraged to explore the detailed expositions by many authors (Cooley et al., 1967; Cooley & Tukey, 1965; Jenkins & Priestley, 1957; Jenkins & Watts, 1968; Oppenheim & Schafer, 1999; Percival & Walden, 1998). What follows is an overview of the concepts which we consider to be relevant to those working in the social sciences, as well as practical instructions on how to conduct spectral analysis with estimations of significance and effect size. Furthermore, we also describe the use of circular statistics and polar plots that can be used to visualize and make inferences about the phase offset between different cycling components in cross-spectral analysis. We have also provided a Python notebook that readers can use to test and replicate different elements of the analyses.[1]

Spectral Analysis via Fourier Series

The Foundations

Trigonometry Review

Similarly to how Principal Components Analysis (PCA) might decompose data into a linear sum of orthogonal components, spectral analysis is concerned with the decomposition of data (in our case, we consider time series data) into a linear sum of principal frequency components. These components are orthogonal sinusoidal basis functions. While the basis function for a linear trend is the first power of some term e.g. 't', the basis functions for spectral analysis are sinusoidal functions of time t:

$$f(t) = A\sin(2\pi ft) \tag{1}$$

Where $2\pi f$ constitutes the angular frequency of the wave and is often denoted as ω; f is the frequency of the wave (in cycles per unit time); and A is the peak (i.e. maximum) amplitude of the wave. The interval between the start and end of each repetition or cycle is

[1] Materials can be found at: https://github.com/matthewvowels1/SpectralAnalysisTutorial.

called the 'period', T. Equation 1 assumes continuous time sampling (i.e. we theoretically have data for every possible time point that could exist), but in practice we usually acquire N data from discrete points in time $n = [0, 1, ..., N-1]$, where the square brackets highlight that n takes on discrete values. In this article we may switch between continuous and discrete time where convenient, but whenever we do, we will use t and n, and round and square brackets, respectively. For instance, a discrete time-sampled version of Equation 2 would be:

$$f[n] = A\sin(\omega n) \tag{2}$$

Figure 2 illustrates a simulated wave from a hypothetical study where the data was collected each day (i.e., a sampling period of 1) for 60 days (i.e., the total number of days in the study). In order to simulate this wave of frequency f, with a unit amplitude e.g. $A = 1$, the following process was followed. Begin by treating the sampling period as ranging from 0 to $(1 - 1/N)$ in equally spaced, discrete steps, where the number of steps is determined by N (e.g., $N = 60$). So long as the time index ranges between 0 and $(1 - 1/N)$, the frequency f is simply the number of cycles in that period (e.g., $f = 2$). Then, dividing the $[0, (1 - 1/N)]$ interval into N equally spaced steps, $f[n]$ may then be generated according to the relationship in Equation 2 with $\omega = 2 \times \pi \times 2$. The period of the wave is simply the reciprocal of the frequency $N_t = f^{-1} = 1/2$ (i.e., one cycle lasts half of our sampling time). If the sampling time is two months, and $N = 60$, then we can re-scale the x-axis from $[0, (1 - 1/N)]$ to [0,59], assuming that we want daily samples for two months, and we recover Figure 2 using $f[n] = 1 \times \sin(2 \times \pi \times 2 \times n)$.

This wave can be 'phase shifted' to the left or to the right by respectively adding or subtracting a phase quantity to the term inside the brackets. If this phase shift is given by a quantity ϕ, then Equation 2 can be extended to: $f[n] = A\sin(2\pi f n \pm \phi)$. Two sine (or cosine) waves are indistinguishable if they have the same frequency and amplitude and have a phase shift that is $\phi = 2d\pi$ radians where $d \in \mathbb{Z}$ (i.e., d is an integer). The summation of

two waves sharing the same frequency and amplitude but having a phase shift $\phi = \pi$ radians is 0. In this case, the waves are said to be in inverse polarity and sum to zero.

In order to shift a single wave, such as the one in Figure 2, by a specific quantity, the relationship between the shift and the period of the wave must be known. We cannot simply apply a delay to a series comprising multiple waves of different frequencies, because the delay would affect all waves by the same quantity regardless of their respective frequencies. For example, for a shift of 15 days to a wave with a total period of 30 days, the shift in radians is $\phi = 2\pi(15/30)$ (i.e., shift is a fraction of the total 2π radians). The result in Figure 3 is given by $f[n] = 1\sin(2 \times \pi \times 2 \times n - 2\pi(15/30))$. Note that the 'dots have been joined' in this figure, to give the illusion of continuous time sampling.

Fourier and the Discrete and Fast Fourier Transforms

Spectral analysis is principally grounded in Fourier's theorem, which posits that *any* periodic function of time can be represented as a sum of sine and cosine functions of varying amplitudes and frequencies. In other words, the Fourier series allows us to represent the original time series as a sum of sinusoidal components and associated spectral parameters (where the parameters are amplitude, frequency, and phase). These parameters can be plotted on a spectrum. Indeed, the aim of spectral analysis is to 'work backwards' to identify the parameters of the spectral components that could have resulted in the generation of the observed time series.

A mathematical description of an infinite Fourier series for a signal is presented by Stroud and Booth (2003, p. 183) as:

$$f(t) = \frac{a_0}{2} + \sum_{k=1}^{\infty}(a_k \cos(kt) + b_k \sin(kt)) \tag{3}$$

In Equation 3, a_k and b_k represent the coefficients and are, by definition, the amplitudes of the Fourier series basis functions which have frequency k.

A square wave is an example of a time series with dramatic changes and abrupt

discontinuities. Figure 4 illustrates how a finite sum of continuous sine waves can be used to approximate this wave. Unfortunately, we can never recreate a square wave perfectly without an infinite sum of sine waves. This is because abrupt discontinuities present undefined / infinite gradients, which only an infinite number of continuous functions can represent. However, in the social sciences we are more likely to be concerned with identifying the periodicity of the fluctuation, than we are to be concerned with perfect time-domain simulation or reconstruction of the time series' discontinuities. Given that periodic fluctuation in discrete or binary variables (e.g. yes/no) may be well represented using a square wave, it is important to understand that spectral analysis can identify fluctuations in such discontinuous time series.

In the discrete time scenario, and finite sample setting, whereby the time series is represented as a finite set of datapoints, the number of Fourier series coefficients is often set to be equal to the number of time points. These coefficients are estimated using the principles of the Discrete Fourier Transform (DFT), which undertakes a lossless, bijective, mathematical transformation between the discrete time and discrete frequency domains. The DFT, and its inverse, DFT^{-1} are defined as follows:

$$X[k] = \text{DFT}\{x[n]\} = \sum_{n=0}^{N-1} x[n] e^{-j\frac{2\pi}{N}kn} \tag{4}$$

and

$$x[n] = \text{DFT}^{-1}\{X[k]\} = \frac{1}{K}\sum_{k=0}^{K-1} X[k] e^{j\frac{2\pi}{N}kn} \tag{5}$$

such that

$$\frac{\text{DFT}^{-1}\{X[k]\}}{\text{DFT}^{-1}\{\text{DFT}\{x[n]\}\}} = 1 \tag{6}$$

for $k = \{0, 1, 2, ..(K-1)\}$.

where $x[n]$ is some function of discrete time (i.e., a time series sampled at regular discrete intervals, such as would be the case for a daily diary study). K may (as mentioned above) be equal to N, and describes the maximum number of Fourier coefficients, and $X[k]$

represents the discrete frequency domain representation of the original discrete time domain signal.[2] Once again, we reiterate that this is a lossless, bijective transformation such that one can move back-and-forth between the frequency and time domains without loss of information. In order to compute the K coefficients, we rely on efficient algorithmic implementations of the DFT, most commonly, the Fast Fourier Transform (FFT).

Implementing the Fast Fourier Transform

Overview and Research Questions

This section presents some examples of simple, univariate time series and their corresponding Fast Fourier Transform (FFT) spectra. We also cover a number of important practical considerations, such as the number of required time points, the highest recoverable frequency, how we extract the phase information, and how to improve the interpretability of the spectra using windowing.

Taking the FFT of a single univariate time series allows us to decompose it into its constituent frequency components, and to understand the relative amplitudes and phase shifts for each of these components. For example, a time series may derive from some self-report measure of mood, where the ratings might be expected to contain periodicity due to socioeconomic factors (such as pay-day), or biological factors (such as the circaseptan rhythm, or day of the week). The FFT can either be used to investigate the nature of any periodicity in mood, or it can be used to remove it (de-seasonalize).

An Example Transformation

Taking the two-month, univariate time series in Figure 2, which comprised a sine wave with an amplitude of 1, and a frequency of two cycles in two months, a Fast Fourier

[2] While the introduction of the exponential term may initially seem surprising, readers are encouraged to recall Euler's formula linking trigonometric functions with the exponential: $e^{j\phi} = \cos\phi + j\sin\phi$, where $j = \sqrt{-1}$. As such, the DFT works to establish a measure of cross-covariance between the time series and a series of sinusoidal components represented in exponential form.

Transform (FFT) spectrum of this time series is shown in Figure 5. Note that a peak is identified at a frequency of one cycle per month, although the x-axis could be re-scaled to be in terms of any desired unit of time.

There are a few notable aspects to this result. First, it is reassuring that we have perfectly recovered the ground-truth results for both the amplitude and the frequency, however, in practice, this will not be the case. Second, the frequency axis stops at 15 cycles per month - this relates to what is known as the Nyquist limit. Third, there are only 31 frequency components, even though in the previous section, K was suggested to usually be equal to the number of time samples $N = 60$. Fourth, there is no indication of the phase of the components. We will now discuss these four items in turn.

Perfect Recovery

The FFT results in Figure 5 demonstrate a perfect/unambiguous recovery of the sine wave component from the hypothetical 60 day study. This has occurred because - contrary to what will usually be the case in practice - the original time series had a duration which was an integer multiple (i.e. a whole number) of cycles or wave periods. Effectively, the wave wraps around such that the last sample in the series, functionally precedes the first sample in the series. The DFT thereby assumes the observed wave is a finite sample of an infinitely long, repeating version of the same wave. A further implication of this assumption is that the wave under observation is stationary, i.e. its parameters and therefore its components do not vary over the course of the sampling period. There are tools, such as spectrograms, that provide three dimensional representations for how periodicity varies over time, but these are beyond the scope of this paper.

In reality, it is more common that the sampling period is *not* an integer multiple of the wave period in length. An example is shown in Figure 6. Note how the wave is in a different part of the cycle when data collection stops as compared to the beginning of the study.

Indeed, if this wave is wrapped around such that its last sample precedes the first, an abrupt discontinuity is introduced, and the DFT has no choice but to model this discontinuity as if it repeats periodically. The FFT is shown in Figure 7 suggesting that frequencies on either side of the actual frequency (i.e., side lobes) may be present thereby blurring the identifiability of the principal frequency, and affecting the estimation of the amplitude.

In practice, the further effect of additive noise may also obscure clean results. Figures 8 and 9 depict example time series and corresponding FFT for a signal with Gaussian distributed additive noise. This noise is now present in the FFT result as randomly distributed energy at all frequencies (i.e., white noise).

Figure 8 and the subsequent FFT in Figure 9 are starting to resemble data that researchers may actually collect. We will discuss techniques, such as windowing, to address these challenges later in the paper.

The Nyquist Limit

The second aspect to discuss relates to the highest reported frequency component, which in the examples was a frequency of 15 cycles per month. This frequency corresponds with what is known as the Nyquist limit, which is equal to half the sampling frequency. For example, for 60 samples of daily diary data, the highest frequency that can be identified, and thus the Nyquist limit, is one cycle per every two days (because the sampling occurs once a day), which would imply 30 cycles in total across that 60 day time span. Since the signal is a function of time, and is being represented by functions with regular and stable periodicity, the data must also be sampled at regular intervals. Any component of periodicity that exceeds these limitations will be *aliased* back into the signal and will cause spurious results (Hamilton, 1994).

Unfortunately, researchers may often have little control over the possibility of aliasing, because, with human participants, the choice of sampling frequency is limited

practically. As such, there will be a certain probability that the data will contain aliased artifacts. In such cases, deviations from regular sampling can actually alleviate the regularity associated with bands of aliasing frequencies, so long as these deviations are not severe enough to affect the estimation of the sub-Nyquist frequency components.[3] Therefore, if a researcher wishes to investigate a phenomenon that they suspect to cycle at a maximum frequency of once per week, then a sample rate of at least twice per week is required. Note that the Nyquist limit is exactly that: A limit. In practice, the higher the sampling rate, the better.

Half the Number of Components to Time Samples.

Given that the highest frequency component corresponds with half the sampling frequency, and given also that the number of frequency components is usually set to be equal to the number of time points, why are there only 31 frequency points in Figure 5, when there were $N = 60$ time samples? The DFT (and hence the FFT) actually yield what are known as *even spectra*, which means the spectrum of positive frequencies are mirrored around the y-axis into negative frequencies. Negative frequencies are not semantically meaningful, and so it is usual practice to discard the first $(K/2) - 1$ components, and present only the positive frequencies (including 0). The mathematical operations may be written as follows: $|FFT_{pos}[k']| = 2 \times |FFT[k']/N|$ for $k' = 0, ..., (N/2) + 1$. The vertical lines indicate a magnitude operation, the factor of two exists to compensate for discarding half the spectrum, and the divisor of N exists to normalize the result according to the number of samples in the sampling period. Note that the factor of two multiplication should only be applied to components that are actually mirrored (i.e., not to the zeroth component, and neither to the nyquist component; MathWorks, 2018). The raw output of the FFT is actually $K = N$ complex duals, usually given in $\hat{C}_k = a_k + jb_k$ form, where $j = \sqrt{-1}$ and C_k is the kth frequency component. In order to take the magnitude, the L2

[3] We owe this idea to Dr. Isaac Washburn (personal correspondence).

norm is taken according to:

$$|C_k| = \sqrt{(a_k^2 + b_k^2)} \tag{7}$$

Where is the Phase Parameter?

The phase parameter is embedded in the raw output of the FFT, which was discarded in order to plot semantically meaningful spectra. However, the phase of any of the K frequency components (or, more usually, the $(K/2) + 1$ positive frequency components) can be trivially calculated according to:

$$\angle \hat{C}_k = \arctan \frac{b_k}{a_k} \tag{8}$$

This equation yields a value in radians. Note that phase *wraps* around 2π radians, or 360 degrees, such that $0 = 2\pi = 4\pi...$ radians. For this reason, circular/polar statistics may be useful when comparing phase values and computing descriptive statistics such as means and standard deviations.

Phase information is crucial if a model of the original time series is to be constructed from the spectral decomposition. For example, if one wishes to simulate a canonical time series using the spectral decomposition, then the phase information is critical in maintaining the temporal relationships between the time series' constituent frequency components. However, and as Gottman and Ringland (1981) note, phase information is available for all frequencies, regardless of whether those frequencies are significant and meaningful in themselves. Therefore, phase information can only be meaningfully interpreted for frequencies which are significantly high in amplitude (i.e. not simply due to noise and random fluctuation).

Increasing the Complexity and Windowing

A wave of marginally increased complexity is given by $f[n] = \sin(4\pi n) + \sin(16\pi n) + 0.5\sin(32\pi n) + 0.4\epsilon[n]$ where $\epsilon \sim \mathcal{N}(0, 1)$, $n = [0, 1/N, ..., 1 - (1/N)]$ and $N = 60$. In words, this wave contains cycles at once per month, four times per month, and twice a week, and also contains additive Gaussian noise. This time series is shown in Figure 10. From a visual inspection it may be possible to vaguely identify the presence of cycles, but the parameters of these components would be almost impossible to identify without spectral analysis. The corresponding FFT spectrum is shown in Figure 11, and the four highest peaks in this spectrum correspond exactly with the four constituent waves in our time series. Note that, this is an idealized demonstration because the length of the sampling period was chosen to be an integer multiple of all constituent components.

As discussed in the section entitled Perfect Recovery, some ambiguity exists in the FFT spectrum of time series which contain a non-integer number of constituent component cycles. The ambiguity manifested itself as side lobes that form either side of main frequency lobes in the spectrum (Wickramarachi, 2003). In order to mitigate these issues, windowing functions may be used to smoothly attenuate the amplitude of the time series towards the start and end of the sampling period, thereby removing discontinuities. Without this attenuation, the time series is essentially being truncated or windowed by a *box-car* function, or a rectangular window:

$$w[n] = \begin{cases} 1 & 0 \leq n \leq N \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

There are various options for windowing, and each of them have various advantages and disadvantages relating to side lobe height and side lobe width trade-offs. A discussion on these trade-offs is beyond the scope of this paper, but interested readers are pointed to Wickramarachi (2003), Brandt (2011), and Oppenheim and Schafer (1999). The Hann

window is considered in this work, and an example windowed sine wave, with the corresponding Hann windowing function, is shown in Figure 12, and corresponding FFT results are shown in Figure 13. It can be seen that the higher frequency component parameter estimation is relatively unaffected by the rectangular truncation, and the application of windowing has actually broadened the side lobes to its detriment. However, at the lower frequency, the truncation is more noticeable, resulting in side lobes that extend a long way either side of the main lobe. Here, the application of windowing reduces these side lobes, and improves amplitude estimation. Note that, pursuant to windowing, an amplitude correction factor must be applied to compensate for the reduction in energy concomitant with windowing. This correction factor is calculated according to Equation 10 (Brandt, 2011). Researchers are encouraged to explore windowing as part of the data analysis.

$$a_{cf} = \frac{N}{\sum_i^N w[n]} \tag{10}$$

Cross-Spectral Analysis

Overview and Research Questions

While it is of some benefit to be able to decompose univariate time series into their constituent frequency components, it is also interesting to be able to identify commonalities, or shared rates of fluctuation, across a bivariate, or dyadic time series. Cross-spectral analysis provides a means to quantify these shared rates of fluctuation, as a form of correlation in the frequency domain. For example, for two individuals living together, cross-spectral analysis provides a means to understand whether their self-report ratings for mood fluctuate at any shared frequencies. Further, most software packages that provide cross-spectral estimation also provide an estimation for the relative phase shift between the shared components. This allows researchers to understand the lead/lag relationship between individuals. It might be, for instance, that while the individuals have

fluctuations in mood that both cycle once per week, each of the individuals' mood ratings peak at different times to one another.

Cross Power Spectral Density

Cross-spectral analysis provides a means to establish the cross-covariance/correlation between the spectra. One common method is the Cross Power Spectral Density (CPSD) method, which first requires computation of the power spectra for each univarite time series, and then undertakes a form a cross-covariance in the frequency domain. The power spectral densities, which are equivalent to periodograms (MathWorks, 2018), are trivially derived from the FFT as follows:

$$P_{xx} = \frac{1}{f_s}|X[k]|^2 \tag{11}$$

The power spectral density can therefore be considered to be the normalized magnitude squared spectrum. Consider the bivariate time series depicted in Figure 14, and the corresponding FFT spectra shown in Figure 15.

The CPSD between two signals x and y, which is designated in this work by P_{xy}, is usually computed using Welch's Overlapped Segmented and Averaged (WOSA) method (Welch, 1967), rather than by simply taking the complex conjugate product of the two as $P_{xy} = X[k]Y^*[k]$, where the asterisk denotes the complex conjugate. For WOSA method, the time signals must first be divided into 'Q' segments, each of length 'S', with overlap between segments of 'R'. Second, the segments are windowed (e.g., using the Hann window), and a Fourier transform is undertaken for each segment yielding X_q and Y_q. Third, P_{xx}, P_{yy} and P_{xy} are calculated for each segment yielding: $I^q_{xx} = |X_q|^2$, $I^q_{yy} = |Y_q|^2$, and $I^q_{xy} = X_q Y^*_q$, so that finally, P_{xy} may be calculated according to the averages of these terms as follows (Haykin & Van Veen, 2007):

$$P_{xy} = \frac{1}{Q}\sum_{q=1}^{Q} I^q_{xy} \tag{12}$$

If $0 \leq P_{xy} \leq 1$ is required, P_{xy} can be squared and normalized by the WOSA estimates of P_{xx} and P_{yy}, which yields what is known as the magnitude squared coherence C_{xy} as follows:

$$C_{xy} = \frac{|\frac{1}{Q}\sum_{q=1}^{Q} I_{xy}^{q}|^2}{\frac{1}{Q}\sum_{q=1}^{Q} I_{xx}^{q} \frac{1}{Q}\sum_{q=1}^{Q} I_{yy}^{q}} \tag{13}$$

This process increases the signal to noise ratio by a factor proportional to the degree of overlap between the windows. This is because averaging with overlap in the presence of independently and identically distributed (i.i.d.) noise reduces the variance associated with the i.i.d. noise in proportion to the 'true' signal.

The CPSD of the bivariate/dyadic time series in Figure 14 is shown in Figure 16. It can be seen that, in spite the large frequency component at eight cycles per month in one of the time series, the CPSD correctly has a value close to zero. If the two time series represented fluctuations in the mood of partners in a relationship, the CPSD would be highlighting that the individuals share common rates of periodic fluctuation.

Dyadic Phase

Note that there is no phase information in the CPSD plot - by taking the power spectra, the phase information is destroyed. However, it is common for implementations of CPSD to additionally provide the phase discrepancy between the time series in the dyad, for each frequency component. It is possible that two time series share the same frequency components (and therefore register a high CPSD at these frequencies) but for the components to be shifted in phase with respect to each other. For example, two individuals who receive their paychecks monthly, but at different times in the month, will likely both exhibit high CPSD at a rate of once per month. However useful it may be to know about this commonality in their spectra, it would also be useful to know whether they cycle *in synchrony*. In the paycheck example, they do not - one individual may have a lowest balance for the month on the same day that the other has just received their paycheck.

The next section explores the use of FFT and CPSD across groups of individuals, and the notion of phase will be revisited.

Group Analyses

Averaging Spectra and Cross-Spectra

Thus far, the univariate and bivariate time series spectral analyses presented would be performed at the person- or dyad-level of analysis respectively. This section concerns the use of spectral analysis over groups of time series (e.g., a number of couples from each of which the researcher has collected daily diary data for mood). In order to demonstrate the process, synthetic data may be generated and analysed according to the specification described below. For real-world problems, simply substitute the data generation steps for data acquisition steps.

The data simulated for the purposes of this article were for M heterosexual dyads (i.e., men and women), 60 days worth of data, time series are generated such that two types of couples exist: the first type has men and women sharing a once-per-month cycle but for a phase difference of 15 days. The women are also arbitrarily given some additional rates of fluctuation not shared with the men. The second group are the same as the first, but for the fact that the women's once-per-month cycle is offset by eight days, instead of 15. As such, we should expect bimodality in the resulting phase distribution. All time series contain additive Gaussian noise.

The time series ought first to be individually de-trended and person mean centered (Gottman & Ringland, 1981; Larsen, 1987; Lester et al., 1985). By person mean centering the time series, the average value across each of the individuals' time series becomes zero. De-trending can be achieved by fitting a first or second order generalized linear model to each individual time series, and subtracting, at each time point, the corresponding prediction from the model. At this point, windowing and window amplitude correction may be applied to mitigate truncation problems. Next, univariate FFTs are computed for

each of the individuals' time series one at a time, and then averaged across individual, per frequency, per group. For example, each of the individuals' spectra will have corresponding amplitude parameters. Each of these amplitudes can be averaged over the group, yielding an average spectrum. Here a group might correspond with a stratification according to gender. This would then allow one to compare the average spectrum for a man with the average spectrum for a woman. Note that averaging must be performed *after* the magnitude of the FFTs has been computed (i.e., after the L2 norm has been computed according to Equation 7). Computing the magnitude removes the phase information from each individual spectrum. Without taking the magnitude, an average in the complex domain may result in constructive and destructive interference (i.e., waves with phase shifts may either sum to reinforce each other, or to cancel each other).

In order to undertake a cross-spectral analysis over multiple dyads, CPSDs may be calculated for each couple individually. Then, similarly to the average univariate FFT analysis, an average may be performed on the magnitude CPSDs across couples, per frequency.

Polar Plots

While the CPSDs are being calculated for each dyad, it may be worthwhile storing the corresponding phase information for each frequency. Note that, because phase wraps around 2π radians, we need to use *polar* histograms, rather than regular histograms. An example is shown in Figure 17. Here, counts extend radially, and the phase discrepancy is given in degrees around the circumference of the polar plot. The polar histogram exhibits bimodality, and at this specific frequency (which happens to be once per month), couples have phase discrepancies of 180 degrees (which is equivalent to 15 days) and 90 degrees (which is equivalent to eight days). These results therefore corroborate our synthetic data specification described above. As mentioned, the phase information can only be meaningfully interpreted at specific, significant frequencies (see below for a description of

how to establish the statistical significance of a frequency component). Finally, the fact that phase wraps also means that conventional descriptive statistics are meaningless, and we encourage researchers to utilize circular equivalents for means and standard deviations when characterizing such distributions.

Statistics

Significance Testing

For noisy, real-world time series, it may be useful to undertake a form of null-hypothesis testing with spectral analysis. In order to avoid potential pitfalls associated with parametric assumptions, and also to enable null-hypothesis testing both on individual as well as groups of time series, it is recommended to undertake a form of bootstrapping.

Our null hypothesis might be stated as follows: no periodic cyclicity exists in the time series. In order to test this, an FFT for i.i.d. noise needs to be compared against an FFT for our time series in question. At this point, it is instructive to illustrate what noise (i.e., a time series devoid of periodic cyclicity) looks like. Consider the twenty Gaussian, i.i.d. time series shown in Figure 18(a). Due to the nature of their i.i.d. sampling over time, there is no auto-correlation structure, besides what might occur spuriously for an arbitrary sample over an arbitrary duration. The corresponding FFTs for these time series are shown in Figure 18(b). Due to the finite number of time points ($N = 60$) there are spurious indications of cyclicity. However, taking an average across the magnitude spectra yields the result shown in Figure 19. Even for 20 time series, it can be seen that the average of the FFTs for i.i.d. time series quickly converges to a flat line, representing the average energy.

It is important to note that the observed time series and the generated i.i.d. noise (which represent samples from our null distribution) must have the same energy. The *same energy* criterion is essential, and thankfully easy to satisfy. The order of the samples in the original time series can be permuted a large number of times (e.g., 1,000 times), which constitutes a form of sampling without replacement. The permutations are devoid of

auto-correlation, providing a set of surrogates against which to compare our original time series. By taking FFTs of the original time series as well as the 1,000 permuted surrogates, we can establish the count for which the amplitude of each component in the FFT of the original signal falls above the corresponding amplitude in the FFTs of the permuted signals. By dividing this count by the number of comparisons, and multiplying by 100, we arrive at a percentage chance that the true amplitude of the wave is at least as strong as the one indicated by the data and analysis, under the assumption that the null effect is true (i.e., it fulfils the same role as the p-value). The procedure also yields bootstrapped confidence intervals.

Figure 20 illustrates the process in terms of variance before and after permutation. On the left, the total variance of $f(t)$ may be attributable to either additive noise $\epsilon(t)$, or to sinusoidal components (in the Figure, only one component is assumed). On the right of the Figure, the variance of the signal post-permutation remains same as the total variance of $f(t)$, but now none of that variance is attributable to sinusoidal components. Note that Figure 20 is not intended to suggest that the variance may necessarily be additive, only that some proportion of the variance is attributable either to the noise or to the sinusoidal component(s). In the section below, we relate the energy and power of a time series to its variance in order to derive a proxy for effect size.

The process described above is appropriate for an individual time series, but is also easily extended to the group scenario: the comparison for a group becomes one of the average FFT across the group, with 1,000 equivalent average FFTs for permuted versions of the time series in that group. Consider Figure 21. This figure may depict either the FFT (bold line) of an individual's time series, and FFTs of the permuted version of this individual's time series (faded lines); or it may equally depict a group average FFT (bold line), and average FFTs of the all permuted time series in the group (faded lines). In either case, the process is the same: To count how many times the amplitudes of the un-permuted time series FFT fall above those of the permuted time series FFTs.

For the example shown in Figure 21, the count results are shown in Figure 22. Here, the x-axis corresponds with the index of the frequency component, and the y-axis corresponds with the percentage frequency with which the amplitude of each component falls above the amplitude of the corresponding component in the permuted time series FFTs.

Once the significance of each spectral component has been established (and an adjustment for the alpha level is recommended in the case of multiple frequency hypothesis testing), a reduced time series model may be built using the amplitudes and phase information for each of the significant components. For example, if the frequencies of $f_1 = 3.5$ and $f_2 = 4.5$ are significant, with amplitudes $A_1 = 0.5$ and $A_2 = 0.2$ and phase shifts of $\phi_1 = 0.2\pi$ and $\phi_2 = 0.1\pi$ respectively, this information can be used to generate a model: $f'[n] = A_1 \sin(2\pi f_1 n + \phi_1) + A_2 \sin(2\pi f_2 n + \phi_2)$. This forms a simplified model comprising only the statistically significant components of the original time series which may be used for forecasting. Alternatively, it may be used to de-seasonalize the original data, i.e. by computing $f[n] - f'[n] = f''[n]$ where $f[n]$ is the original time series, $f'[n]$ is the time series model comprising only significant frequency components, and $f''[n]$ is the de-seasonalized result. An example of before-and-after de-seasonalization is show in Figure 23. In this time series, two dominating cycling components obfuscate a clear quadratic trend. Spectral analysis thereby allows the researcher to decouple the trend from the oscillating components to more clearly model other important, underlying patterns (Sims, 1974).

Finally, consider what happens if, instead of counting how many times our original time series FFT falls *above* the permuted time series FFTs, we count how many times our original time series FFT falls *below* the permuted time series FFTs.[4] In this case, instead of identifying periodicity, we identify a persistence, or a reluctance to change across frequency. This alternative may also be of interest to researchers, who wish to explore phenomena

[4] We owe this idea to personal correspondence with Prof. Michael Renfro.

that exhibit a form of time-domain inertia.

Effect Sizes

As well as identifying significant frequency components, it may also be of interest to be able to quantify the importance or strength of these components, much akin to a measure of effect size in conventional statistics. Despite the apparent in-applicability of electronic engineering notions such as electrical power and energy to social science data, these concepts are useful when building analogies for the metrics in the two fields.

Start by assuming the following model: $f(t) = A\sin(\omega t) + \epsilon(t)$. In words, our model comprises a sinusoidal function of time, with some additive i.i.d. noise e.g. $\epsilon(t) \sim \mathcal{N}(0, 1)$. Variance is a quantity usually associated with random variables. Given that a deterministic sinusoidal function of time $\sin(\omega t)$ is not a random variable, its variance might be fairly considered to be zero, and its expectation simply equal to itself i.e. $\mathbb{E}[\sin(\omega t)] = \sin(\omega t)$. The expectation and variance of $f(t)$ are therefore equal to the expectation and variance of the additive noise alone, ϵ, which are zero and one respectively.

However, by deviating from these formalities and for the sake of this exposition, one can compute the variance of a time series comprising a deterministic sine wave, as if it were a random variable, according to the empirical expression for variance, (i.e., the average squared difference of the sample values from the mean). This assumes person mean centering and de-trending, which, as described above, is recommended as a data pre-processing step. In electronic engineering, this *empirical variance* quantity is actually known as the electrical power of a signal. The square root of the empirical variance quantity is the standard deviation, and this is actually the same as what is known in electronic engineering as the root mean square (RMS) value. The RMS value is often used to characterize the strength of an a.c. signal, and gives the equivalent level of a d.c. signal that would dissipate the same amount of electrical power. The squaring operation in Root Mean Square is required as *rectification*, and thereby makes all negative values positive,

yielding a non-zero mean. One can easily switch between RMS and the peak amplitude of a sine wave by multiplying or dividing by the square root of two, respectively. If one wants to compute the *energy* of a time series, the variance (or power) can be multiplied by the length of the signal. For example, the energy of a time series of duration $N = 60$ days that has a variance/power equal to one is 60. Here, the units have been omitted because they depend on the original measure used (for example, in electronics, the power might be measured in Watts, and the energy in the above calculation would be measured in Watt-Days).

The FFT diagrams presented thus far are known as amplitude spectra, as they represent the amplitudes of the frequency components in the Fourier decomposition. As described already, these amplitude spectra can easily be converted into power spectra (or periodograms) by taking the raw FFT output and taking the squared magnitude before normalizing by the length of the signal. The power spectrum therefore provides the variance (or power) of each of the constituent frequency components. We recommend that these per-frequency variance values be used as proxies for effect size.

Proportion of Variance Explained

As described above, the power spectrum provides a means to decompose the time series into the individual contributions of variance/power from all frequency components, and we proposed that these per-frequency variance quantities be treated as measures of effect size. Taking inspiration from the Signal to Noise Ratio (SNR) in electronic engineering (Proakis & Manolakis, 1996), in order to make these measures comparable across studies we propose a ratio which we call the Signal to Variance Ratio (SVR). Essentially, by summing across every power spectrum component, the energy of the wave is computed (similarly to how multiplying the time series variance by the number of time points also gives the energy). The ratio of a particular power component V_k to the total sum of all components $\sum_k^K V_k$ therefore provides a measure of the relative contribution of

each component:

$$\mathrm{SVR}_k = \frac{V_k}{\sum_k^K V_k} \tag{14}$$

Whether or not the numerator and denominator are both multiplied by the number of components to normalize them according to power or energy is irrelevant as the ratio will remain the same. The SVR can be multiplied by 100 to provide an estimation in terms of a percent contribution.

Given the spectral analogies with conventional statistical concepts such as effect size, proportion of variance explained, and null hypothesis testing, one is equipped with the relevant tools to additionally undertake statistical inference with confidence intervals, and statistical power analyses.

Discussion

The purpose of this tutorial was to introduce the theory and application of spectral and cross-spectral analysis, which allow researchers to identify cycling components in time series data. There are a limited number of examples of its application in the social sciences (Campbell et al., 1991; Gottman, 1979; Gottman & Ringland, 1981; Larsen, 1987; Larsen et al., 2009; Larsen & Kasimatis, 1990; Lester et al., 1985; Vowels et al., 2018; R. M. Warner et al., 1987) and, to the best of our knowledge, there are no primers or practical guidelines for how to apply it in social sciences. This paper provides a much needed introduction to spectral analysis and our hope is that it will help to encourage researchers to consider potential periodicity and synchrony in their data. In this tutorial we have: Extended the traditional spectral and cross-spectral methodologies to enable a non-parametric, bootstrapped estimation of the significance of univariate and bivariate frequency components; described how power spectra may be interpreted to derive a proxy for effect size; discussed why and how to undertake windowing; and illustrated how the use of circular statistics enables us to investigate the phase offset between two signals.

There are a number of strengths to spectral analysis which are worth highlighting. First, the FFT and DFT are bijective transforms, not models, which means that they retain all of the information present in the original time domain representation. The methods provide researchers with a new perspective for interrogating data and identifying patterns which might otherwise be visually unidentifiable. Once the frequency domain spectra have been derived, researchers can then identify significant spectral components and associated effect sizes in order to build time series models. Second, the bootstrapping technique that we presented as a means to estimate significance is non-parameteric, and therefore does not rely on potentially restrictive assumptions. This bootstrapping technique can also be applied to estimate the significance of the spectral components in *individual*, as well as groups of, time series. It thereby provides clinicians with a means to build time series models on a case by case basis. Third, the two advantages above are equally applicable to cross-spectral analysis, which additionally provides a means to understand synchrony in dyadic data. Fourth, spectral analysis can also be useful as a descriptive tool as it provides a convenient way of extracting cycles from individual's data (Chow et al., 2005). Finally, our proposal for the use of circular statistics enables researchers to interrogate data that wraps, which in our case was phase information, but could be equally useful for calendar data where, for example, the first day of the year should be considered to be one day away from the last day of the year.

While this paper provides an introduction to spectral and cross-spectral analysis, there are many potential extensions that have not been discussed in detail. For example, we briefly referred to spectrograms (Rabiner & Schafer, 1978) which provide an estimation of the spectral components as they evolve over time. Spectrograms, for example, can be used to model dampening effects similar to the damped oscillator model (Chow et al., 2005). However, spectrograms work by undertaking multiple FFTs over a sliding window, and the requirements for adequate sampling frequency and sampling duration are challenging to fulfil when collecting typical social sciences data. Another extension involves the use of

spectral parameters/coefficients from an FFT as features in other analyses. For example, the magnitude parameters can be used as predictors in ordinary least squares regression, dynamic factor models, or features in machine learning algorithms such as random forests, deep neural networks, and computer vision algorithms (Breiman, 2001; Goodfellow et al., 2016; Molenaar, 1987; Szeliski, 2010). Different elements from the spectral decomposition (e.g., amplitude, phase, couple CPSD, significant frequencies) can also be used as outcomes in other models. Finally, we also alluded to the use of spectral analysis for time series decomposition (i.e., for decomposing a time series into its trend and seasonal components) (Brockwell & Davis, 2002). Time series decomposition is important in removing the influence of seasonality on the estimation of the trend and overall growth pattern. Without such de-seasonalization, the estimation of the trend may be meaningless, and become dependent on the start and end points of the sampling duration (e.g., see Figure 1).

There are also a number of different flavors of spectral analysis, which may yield different results. One example is wavelet analysis, which is more robust in analyzing signals that have spectral compositions that change over time (Abramovich et al., 2000; Nason & Savchev, 2014). Other options include variants of Spectral Proper Orthogonal Decomposition, which may be more computationally complex, but provide researchers with additional avenues for identifying patterns/commonality within signals (Paul & Verma, 2016; Sieber et al., 2016; Towne et al., 2017).

As part of the work, we have provided a Jupyter Notebook (written in Python 3.6) that explains how to conduct spectral analysis step by step. This notebook contains many of the simulation examples used in this tutorial. It is worth noting that we have also applied these techniques to real social sciences data, and a case study can be found in Vowels et al., 2018. Spectral analysis may also be undertaken in R (for an example, see Wearing, 2010). Note that some functions (e.g., windowing options) described in this paper and in the Jupyter Notebook may not be available as ready functions. Many of the standard statistical programs such as SPSS and SAS also have their own spectral analysis

functions that may be sufficient for some analyses.

Limitations and Considerations

There are several limitations for spectral analysis: First, there is a trade-off between time domain and frequency domain resolution/accuracy (Jenkins & Watts, 1968; MathWorks, 2018; Stowell & Plumbley, 2014; Watkinson, 2001). This trade-off between time domain accuracy and frequency domain accuracy is analogous to the Heisenberg uncertainty principle: Accuracy cannot be achieved simultaneously in both domains. In other words, the longer and less temporally specific the time domain sample is, the higher the accuracy in the frequency domain (assuming the frequency content does not change over the sampling duration). Conversely, the shorter and more temporally specific the time domain sample is, the less precise the frequency domain representation is. However, for signals that are stationary, in that the characteristics of the periodic components do not significantly vary over time, high time domain accuracy is not important, and it makes more sense to maximize the frequency domain accuracy. For example, in order to identify cycling components in individuals' mood over the course of a week, researchers may attain more accurate estimations of the frequencies of the components the longer the period of study is. In this case, researchers should study the individuals over the course of as many weeks as is feasible. In doing so, the assumption is that the frequency components remain stationary over the course of that period. If the researchers suspect that the frequency content changes over the course of the data collection period, they may utilize time windowing approaches and compute a spectrogram.

Second, the signal must be sampled at a frequency that is at least twice the highest frequency of interest (Jenkins & Watts, 1968; Watkinson, 2001). In other words, the highest frequency that can be recovered is a frequency half the sampling frequency, and this is known as the Nyquist limit. Relating this consideration to the example above, in seeking to identify mood fluctuations that occur once daily, sampling must occur at least twice a

day. Furthermore, the sampling must be undertaken regularly and consistently such that the time period in between samples does not vary dramatically. In reality, if a researcher is collecting daily reports from participants, they may complete the report at different times each day. Discussion of how to deal with these instances can be found elsewhere (Voelkle & Oud, 2013) Additionally, there are situations in which the signal may not always cycle at the same frequencies. For example, mood which is thought to cycle weekly, may be impacted by longer weekends or holidays and the calendar months vary in their length. These will add variation to the data making its spectral decomposition more ambiguous.

Third, there should ideally be no missing data points (Cowpertwait & Metcalfe, 2009), although Broersen (2006) and Christmas (2013) discuss potential ways around this common problem. Our suggestion is that researchers interpolate the time series (e.g., with polynomial splines; Won Suk et al., 2019) and subsequently re-sample it at regular intervals. Fourth, because spectral analysis is not a modelling technique but rather a transform, it is not possible to model measurement error. When researchers have access to multivariate time series data, they may wish to take their analysis further by using dynamic factor models rather than stopping after the frequency domain transformation (Molenaar, 1987). Finally, spectral analysis cannot be used to identify causal influence (i.e., which variable influences the other variable in a bivariate cross-spectral analysis). Similarly, it cannot tell which variable precedes the other - the phase information is ambiguous in this regard, and only tells us the relative offset between components and not which component comes first. Bivariate time-series analysis may be a better choice in cases where a researcher believes autocorrelation may be an issue, but wishes to test for cross-correlation and interaction between time series (Gottman & Ringland, 1981).

Many of the usual statistical assumption checks may not be relevant to whether spectral analysis can be used on a particular set of time series data. For instance, a test of normality is not relevant - this is trivially demonstrated by considering the distribution of a sine wave, which resembles a Beta$(0.5, 0.5)$ distribution, and the time series samples

therefore need not be sampled from a parametric distribution. Furthermore, tests for outliers which exhibit large deviations from the mean of a group's time series are also not relevant, because such mean offsets are either removed as part of individual de-trending, or constitute the zeroeth frequency component (i.e., the d.c. offset). Whether or not the individuals exhibiting such outlying characteristics are problematic for other reasons (e.g. measurement error) is, similarly, an orthogonal issue. Researchers are, however, encouraged to evaluate the distribution of spectral frequency components *after* a spectral analysis has been undertaken, in order to identify individuals with outlying spectral characteristics (e.g. individuals with outlying component amplitudes, or individuals with outlying spectral compositionality). Our bootstrapped approach to significance testing makes no parametric assumptions about the distribution of spectral amplitudes, but this is not to say that these distributions should not be interrogated for outliers or other distributional anomalies. Finally, interested researchers are pointed to work by Percival and Walden, 1998 who discuss some of the more detailed assumptions associated with utilization of a particular estimation algorithm for univariate spectral analysis.

In conclusion, spectral and cross-spectral methods are important for uncovering cycling components in time series data. Ignoring the presence of these components has the potential to render trend estimations meaningless, and an investigation of the spectral components themselves may yield new insight into how many social phenomena fluctuate over time. The techniques presented in this tutorial (and in the accompanying Jupyter Notebook) enable researchers to perform traditional significance and effect size estimation without restrictive parametric assumptions. It is hoped that researchers can extend their analytical arsenal with the techniques, in order to investigate important social phenomena that fluctuate over time.

References

Abramovich, F., Bailey, T. C., & Sapatinas, T. (2000). Wavelet analysis and its statistical applications. *The Statistician*, *49*(1), 1–29. https://doi.org/10.1111/1467-9884.00216

Asparouhov, T., Hamaker, E., & Muthen, B. (2017). Dynamic structural equation models. *Structural Equation Modeling: A Multidisciplinary Journal*, *25*(3), 359–388. https://doi.org/10.1080/10705511.2017.1406803

Bert, S., Park, G., & Liddle, P. F. (2010). Impaired cross-spectral phase coherence in schizophrenia. *Schizophrenia Research*, *117*(2-3), 478–479. https://doi.org/10.1016/j.schres.2010.02.902

Bolger, N., Davis, A., & Rafaeli, E. (2003). Diary methods: Capturing life as it is lived. *Annual Review of Psychology*, *54*(1), 579–616. https://doi.org/10.1146/annurev.psych.54.101601.145030

Bolger, N., & Laurenceau, J. P. (2013). *Intensive longitudinal methods*. New York, The Guilford Press.

Brandt, A. (2011). *Noise and vibration analysis: Signal analysis and experimental procedures*. John Wiley; Sons Ltd.

Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Bringmann, L., Hamaker, E., Vigo, D., Aubert, A., Borsboom, D., & Tuerlinckx, F. (2016). Changing dynamics: Time-varying autoregressive models using generalized additive modelling. *Psychological Methods*, *22*, 409–425. https://doi.org/10.1037/met0000085

Brockwell, P., & Davis, R. (2002). *Introduction to time series and forecasting* (S. F. G. Casella & I. Olkin, Eds.). New York, NY, Springer Texts In Statistics.

Broersen, P. M. T. (2006). Automatic spectral analysis with missing data. *Digital Signal Processing*, *16*(6), 754–766. https://doi.org/10.1016/j.dsp.2006.01.001

Butner, J., Diamond, L., & Hicks, A. (2007). Attachment style and two forms of affect coregulation between romantic partners. *Personal Relationships*, *14*, 431–455. https://doi.org/10.1111/j.1475-6811.2007.00164.x

Campbell, J. D., Chew, B., & Scratchley, L. S. (1991). Cognitive and emotional reactions to daily events: The effects of self-esteem and self-complexity. *Journal of Personality*, *59*(3), 473–505. https://doi.org/10.1111/j.1467-6494.1991.tb00257.x.

Chatfield, C. (2005). *The analysis of time series, an introduction.* (6th). New York, Chapman; Hall.

Chow, S., Ram, N., Boker, S. M., & Clore, G. (2005). Emotion as a thermostat: Representing emotion regulation using a damped oscillator model. *Emotion*, *5*(2), 208–225. https://doi.org/10.1037/1528-3542.5.2.208

Christmas, J. (2013). The effect of missing data on robust bayesian spectral analysis. *IEEE International Workshop on Machine Learning for Signal Processing*. https://doi.org/10.1109/MLSP.2013.6661980

Cooley, J. W., Lewis, P. A., & Welch, P. D. (1967). Historical notes on the fast fourier transform. *IEEE Transactions on Audio and Electroacoustics*, *15*(2), 76–79. https://doi.org/10.1109/TAU.1967.1161903

Cooley, J. W., & Tukey, J. W. (1965). An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, *19*, 297–301. https://doi.org/10.1090/S0025-5718-1965-0178586-1

Cowpertwait, P. S. P., & Metcalfe, A. V. (2009). *Introductory time series with R*. New York, Spinger.

Deboeck, P., Boker, S. M., & Bergeman, C. (2008). Modeling individual damped linear oscillator processes with differential equations. *Multivariate Behavioral Research*, *43*, 497–523. https://doi.org/10.1080/00273170802490616

Driver, C., & Voelkle, M. (2018). Hierarchical bayesian continuous time dynamic modeling. *Psychological Methods*, *23*(4), 774–799. https://doi.org/10.1037/met0000168

Ferrer, E., & Nesselroade, J. (2003). Modeling active processes in dyadic relations via dynamic factor analysis. *Emotion*, *3*(4), 344–360. https://doi.org/10.1037/1528-3542.3.4.344

Geweke, J. (1978). The temporal and sectoral aggregation of seasonally adjusted time series. *Seasonal Analysis of Economic Time Series*, 411–432.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. Cambridge, Massachusetts, MIT Press.

Gottman, J. M. (1979). Detecting cyclicity in social interaction. *Psychological Bulletin*, *86*(2), 338–348. https://doi.org/10.1037/0033-2909.86.2.338

Gottman, J. M., Murray, J. D., Swanson, C. C., Tyson, R., & Swanson, K. R. (2003). *The mathematics of marriage: Dynamic nonlinear models.* Cambridge, MIT Press.

Gottman, J. M., & Ringland, J. T. (1981). The analysis of dominance and bidirectionality in social development. *Society for Research in Child Development*, *52*, 393–412. https://doi.org/10.2307/1129157

Gottman, J. M., Swanson, C., & Swanson, K. (2002). A general systems theory of marriage: Nonlinear difference equation modeling of marital interaction. *Personality and Social Psychology Review*, *6*, 326–340. https://doi.org/10.1207/S15327957PSPR0604_07

Guastello, S., Pincus, D., & Gunderson, P. (2006). Electrodermal arousal between participants in a conversation: Nonlinear dynamics and linkage effects. *Nonlinear Dynamics, Psychology, and Life Sciences*, *10*, 365–399.

Hamilton, J. D. (1994). *Time series analysis*. New Jersey, Princeton University Press.

Haykin, S., & Van Veen, B. (2007). *Signals and systems*. John Wiley; Sons.

Helm, J., Sbarra, D., & Ferrer, E. (2012). Assessing cross-partner associations in physiological responses via coupled oscillator models. *Emotion*, *12*, 748–762. https://doi.org/10.1037/a0025036

Hu, Y., Boker, S., Neale, M., & Klump, K. (2014). Coupled latent differential equation with moderators: Simulation and application. *Psychological Methods*, *19*(1). https://doi.org/10.1037/a0032476

Jenkins, G. M., & Priestley, M. B. (1957). The spectral analysis of time-series. *Journal of the Royal Statistical Society*, *19*(1), 1–12. https://doi.org/10.1111/j.2517-6161.1957.tb00240.x

Jenkins, G. M., & Watts, D. G. (1968). *Spectral analysis and its applications*. San Francisco, Holden-Dan.

Kleinspehn-Ammerlahn, A., Riediger, M., Schmiedek, F., von Oertzen, T., Li, S.-C., & Lindenberger, U. (2011). Dyadic drumming across the lifespan reveals a zone of proximal development in children. *Developmental Psychology*, *47*, 632–644. https://doi.org/10.1037/a0021818

Larsen, R. J. (1987). The stability of mood variability: A spectral analytic approach to daily mood assessments. *Journal of Personality and Social Psychology*, *52*(6), 1195–1204. https://doi.org/10.1037/0022-3514.52.6.1195

Larsen, R. J., Augustine, A. A., & Prizmic, Z. (2009). A process approach to emotion and personality: Using time as a facet of data. *Cognition and Emotion*, *23*(7), 1407–1426. https://doi.org/10.1080/02699930902851302

Larsen, R. J., & Kasimatis, M. (1990). Individual differences in entrainment of mood to the weekly calendar. *Journal of Personality and Social Psychology*, *58*(1), 164–171. https://doi.org/10.1037//0022-3514.58.1.164

Lester, B., Hoffman, J., & Brazelton, T. B. (1985). The rhythmic structure of mother-infant interaction in term and preterm infants. *Child Development*, *56*, 15–27. https://doi.org/10.2307/1130169

Liu, S., & Molenaar, P. (2016). Testing for granger causality in the frequency domain: A phase resampling method. *Multivariate Behavioral Research*, *51*, 53–66. https://doi.org/10.1080/00273171.2015.1100528

Liu, S., Rovine, M., Klein, L., & Almeida, D. (2013). Synchrony of diurnal cortisol pattern in couples. *Journal of Family Psychology*, *27*, 579–588. https://doi.org/10.1037/a0033735

MathWorks. (2018). Matlab user's guide. *The MathWorks, Inc.*

Mehl, M., & Conner, T. (2012). *Handbook of research methods for studying daily life.* New York, NY, Guilford Press.

Molenaar, P. (1985). A dynamic factor model for the analysis of multivariate time series. *Psychometrika*, *50*, 181–202. https://doi.org/10.1007/BF02294246

Molenaar, P. (1987). Dynamic factor analysis in the frequency domain: Causal modeling of multivariate psychophysiological time series. *Multivariate Behavioral Research*, *22*, 329–353. https://doi.org/10.1207/s15327906mbr2203_5

Molenaar, P. (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement; Interdisciplinary Research and Perspectives*, *2*, 201–218. https://doi.org/10.1207/s15366359mea0204_1

Molenaar, P., & Campbell, C. G. (2009). The new person-specific paradigm in psychology. *Current Directions in Psychological Science*. https://doi.org/10.1111/j.1467-8721.2009.01619.x

Molenaar, P., De Gooijer, J. G., & Schmitz, B. (1992). Dynamic factor analysis of nonstationary multivariate time series. *Psychometrika*, *57*, 333–349.

Muller, V., & Lindenberger, U. (2011). Cardiac and respiratory patterns synchronize between persons during choir singing. *PLoS One*, *6*(9). https://doi.org/10.1371/journal.pone.0024893

Nason, G. P., & Savchev, D. (2014). White noise testing using wavelets. *The ISI's Journal for the Rapid Dissemination of Statistics Research*, *3*(1), 351–362. https://doi.org/10.1002/sta4.69

Nesselroade, J., & Boker, S. M. (1994). Can personality change? In T. Heatherton & J. Weinberger (Eds.). Washington DC, American Psychological Association.

Nesselroade, J., & Molenaar, P. (2010). The handbook of life-span development. In R. M. Lerner & W. F. Overton (Eds.). Hoboken, NJ, Wiley.

Nesselroade, J., & Ram, N. (2004). Studying intraindividual variability: What we have learned that will help us understand lives in context. *Research in Human Development*, *1*(1-2), 9–29. https://doi.org/10.1080/15427609.2004.9683328

Oppenheim, A. V., & Schafer, R. W. (1999). *Discrete-time signal processing*. New Jersey, Pearson International.

Papp, L., Pendy, P., Simon, C., & Adam, E. (2013). Spouses' cortisol associations and moderators: Testing physiological synchrony and connectedness in everyday life. *Family Process*, *52*, 284–298. https://doi.org/10.1111/j.1545-5300.2012.01413.x

Paul, S., & Verma, M. K. (2016). Proper orthogonal decomposition vs. Fourier analysis for extraction of large-scale structures of thermal convection, In *Advances in computation, modeling and control of transitional and turbulent flows*. World Scientific Publishing Company.

Percival, D., & Walden, A. (1998). *Spectral analysis for physical applications: Multitaper and conventional univariate techniques*. Cambridge, Cambridge University Press.

Proakis, J. G., & Manolakis, D. G. (1996). *Digital signal processing: Principles, algorithms, and applications*. New Jersey, Prentice-Hall International, Inc.

Rabiner, L. R., & Schafer, R. W. (1978). *Digital processing of speech signals* (A. Oppenheim, Ed.). Prentice-Hall.

Sadler, P., Ethier, N., Gunn, G., Duong, D., & Woody, E. (2009). Are we on the same wavelength? interpersonal complementarity as shared cyclical patterns during interactions. *Journal of Personality and Social Psychology*, *97*, 1005–1020. https://doi.org/10.1037/a0016232

Saxbe, D., & Repetti, R. (2010). For better or worse? coregulation of couples' cortisol levels and mood states. *Journal of Personality and Social Psychology*, *98*, 92–103. https://doi.org/10.1037/a0016959

Scollon, C., Prieto, C.-K., & Diener, E. (2003). Experience sampling: Promises and pitfalls, strength and weaknesses. *Journal of Happiness Studies*, *4*, 5–34. https://doi.org/10.1023/A:1023605205115

Sieber, M., Oberleithner, K., & Paschereit, C. O. (2016). Spectral proper orthogonal decomposition. *Journal of Fluid Mechanics*, *792*, 798–828. https://doi.org/10.1017/jfm.2016.103

Sims, C. (1974). Seasonality in regression. *Journal of the American Statistical Association*, *69*(347), 618–626. https://doi.org/10.2307/2285991

Steele, J., & Ferrer, E. (2011). Latent differential equation modeling of self-regulatory and coregulatory affective processes. *Multivariate Behavioral Research*, *46*(6), 956–984. https://doi.org/10.1080/00273171.2011.625305

Stowell, D., & Plumbley, M. D. (2014). Large-scale analysis of frequency modulation in birdsong data bases. *Methods in Ecology and Evolution*, *5*, 901–912. https://doi.org/10.1111/2041-210X.12223

Stroud, K. A., & Booth, D. J. (2003). *Advanced engineering mathematics* (4th). New York, Industrial Press.

Szeliski, R. (2010). *Computer vision: Algorithms and applications*. Spinger.

Towne, A., Schmidt, O. T., & Colonius, T. (2017). Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics*, *847*, 821–867. https://doi.org/10.1017/jfm.2018.283

Vallacher, R. R., & Nowak, A. (1997). The emergence of dynamical social psychology. *Psychological Inquiry*, *8*(2), 73–99. https://doi.org/10.1207/s15327965pli0802_1

van der Laan, M. J., & Rose, S. (2011). *Targeted learning - causal inference for observational and experimental data*. New York, Springer International.

Voelkle, M., & Oud, J. (2013). Continuous time modelling with individually varying time intervals for oscillating and non-oscillating processes. *British Journal of Mathematical and Statistical Psychology*, *66*(1), 103–126. https://doi.org/10.1111/j.2044-8317.2012.02043.x.

Vowels, M. J., Mark, K., Vowels, L. M., & Wood, N. (2018). Using spectral and cross-spectral analysis to identify patterns and synchrony in couples' sexual desire. *PLoS One*, *13*(10), e0205330. https://doi.org/10.1371/journal.pone.0205330

Wang, L., Hamaker, E., & Bergeman, C. (2012). Investigating inter-individual differences in short-term intra-individual variability. *Psychological Methods*, *17*(4), 567–581. https://doi.org/10.1037/a0029317

Warner, R. M., Malloy, D., Schneider, K., Knoth, R., & Wilder, B. (1987). Rhythmic organization of social interaction and observer ratings of positive affect and involvement. *Journal of Nonverbal Behavior*, *11*(2), 57–74. https://doi.org/10.1007/BF00990958

Warner, R. (1998). *Spectral analysis of time series data*. New York, Guilford Press.

Watkinson, J. (2001). *The art of digital audio* (3rd). Oxford, Focul Press.

Wearing, H. (2010). Spectral analysis in R. *https://ms.mcmaster.ca/ bolker/eeid/2010/Ecology/Spectral.pdf*.

Welch, P. D. (1967). The use of the Fast Fourier Transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, *15*, 70–73. https://doi.org/10.1109/TAU.1967.1161901

Wickramarachi, P. (2003). Effects of windowing on the spectral content of a signal. *Sound and Vibration, Data Physics Corporation*, 10–11.

Wiltermuth, S., & Heath, C. (2009). Synchrony and cooperations. *Psychological Science*, *20*, 1–5. https://doi.org/10.1111/j.1467-9280.2008.02253.x

Won Suk, H., West, S., Fine, K., & Grimm, K. (2019). Nonlinear growth curve modelling using penalized spline models: A gentle introduction. *Psychological Methods*, *24*, 269–290. https://doi.org/10.1037/met0000193

Yeragani, V. K., Cashmere, D., Miewald, J., Tancer, M., & Keshavan, M. (2006). Decreased coherence in higher frequency ranges (beta and gamma) between central and frontal EEG in patients with schizophrenia: A preliminary report. *Psychiatry Research*, *141*(1), 53–60. https://doi.org/10.1016/j.psychres.2005.07.016.

Figure 1

Cycles and a Linear Trend

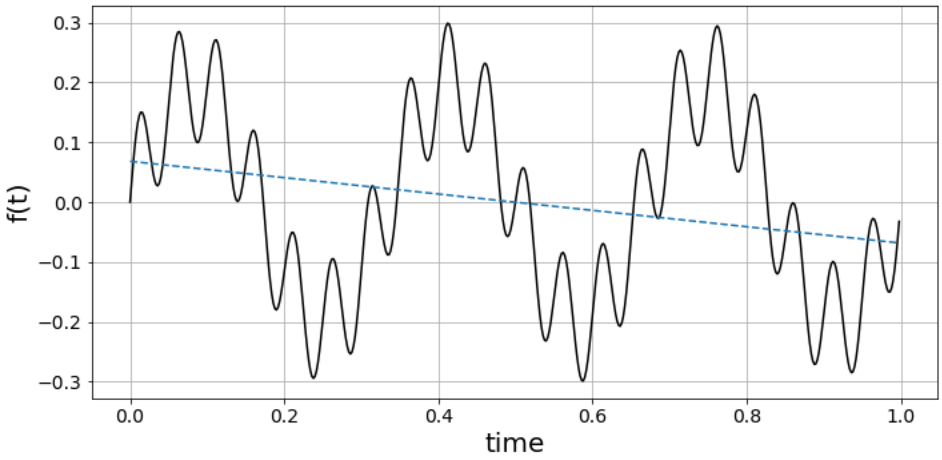

Note. A cycling phenomenon, with no linear trend (solid line), may be badly misrepresented by its corresponding linear regression model.

Figure 2

A Simple Sine Wave

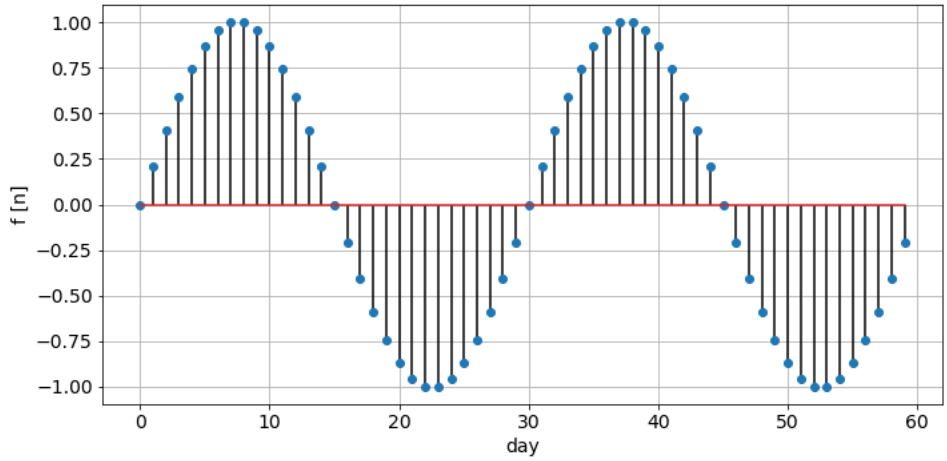

Note. An example of a univariate, discrete-time sampled sinusoidal function for $f[n] = 1 \times \sin(2 \times \pi \times 2 \times n)$ for $n = [0, 1/N, ..., 1 - (1/N)]$ and $N = 60$.

Figure 3

A Phase Shifted Sine Wave

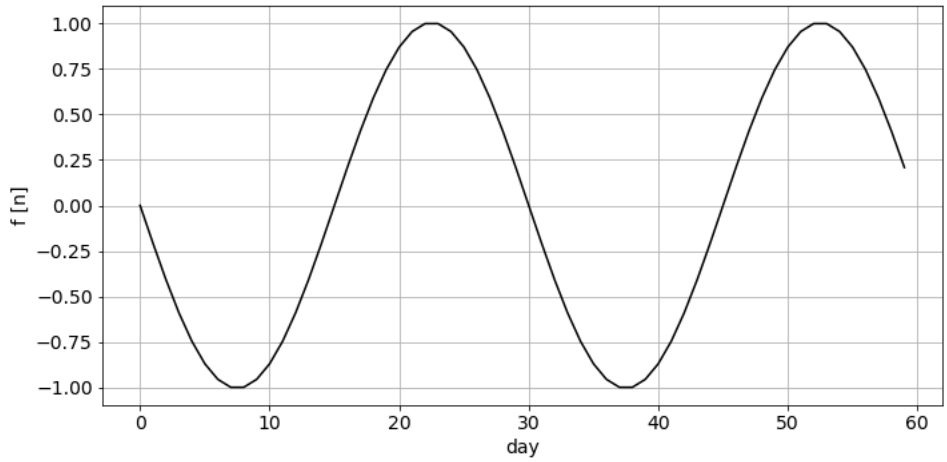

Note. The same sine wave as in Figure 3, shifted by π radians: $f[n] = 1 \sin(2 \times \pi \times 2 \times n - 2\pi(15/30))$.

Figure 4

A Sum of Sine Waves

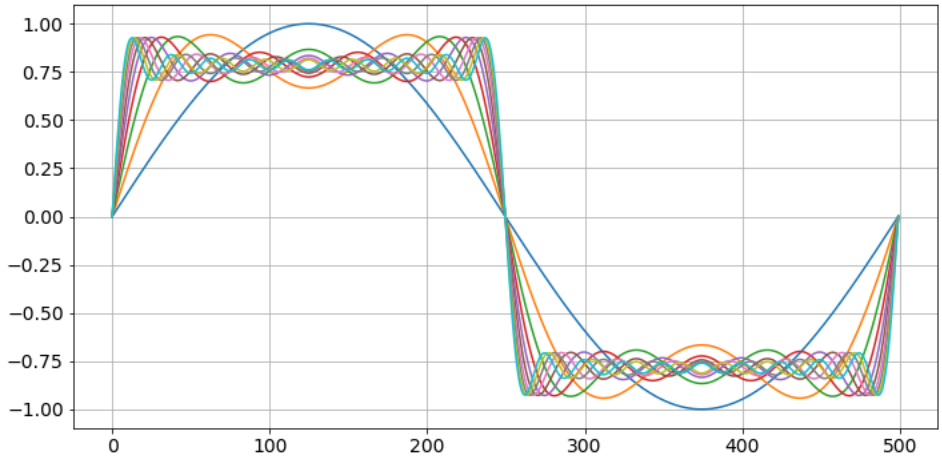

Note. Cumulative sums of sine waves approximating a square/rectangular wave. Each wave includes an additional component.

Figure 5

FFT of the Time Series in Figure 2

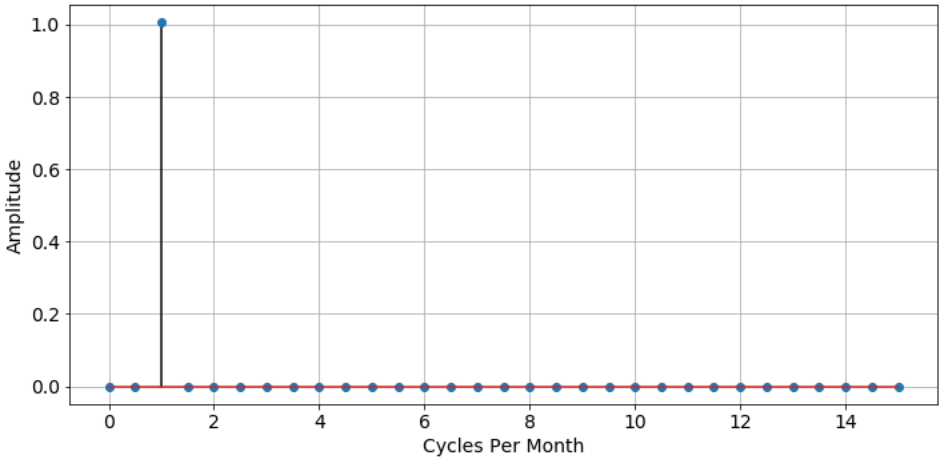

Figure 6

An Example Time Series

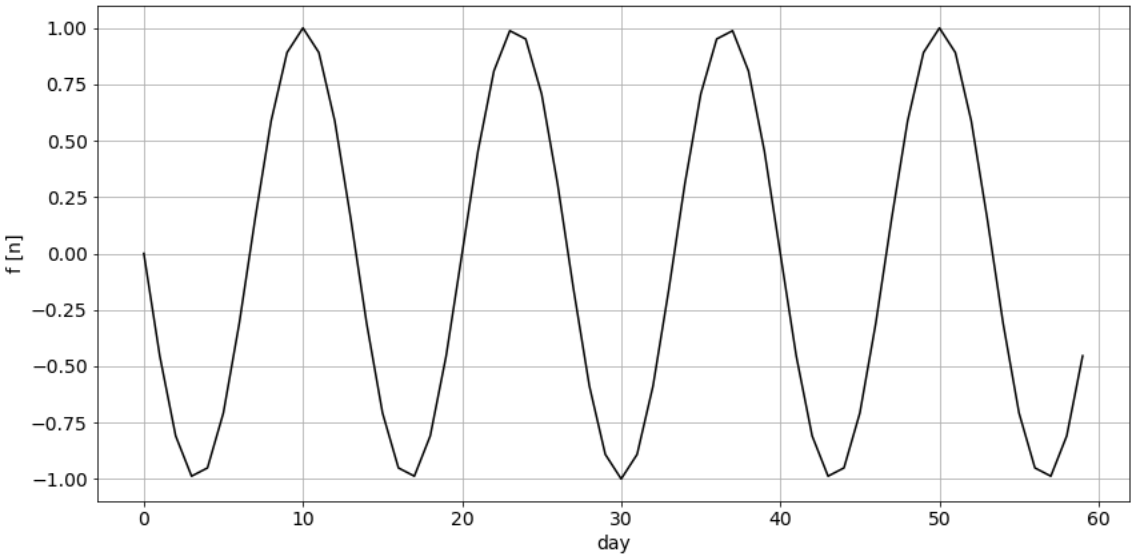

Note. Time series for $f[n] = \sin(2 \times \pi \times 4.5 \times n - 2\pi(15/30))$

Figure 7

FFT of the Time Series in Figure 6

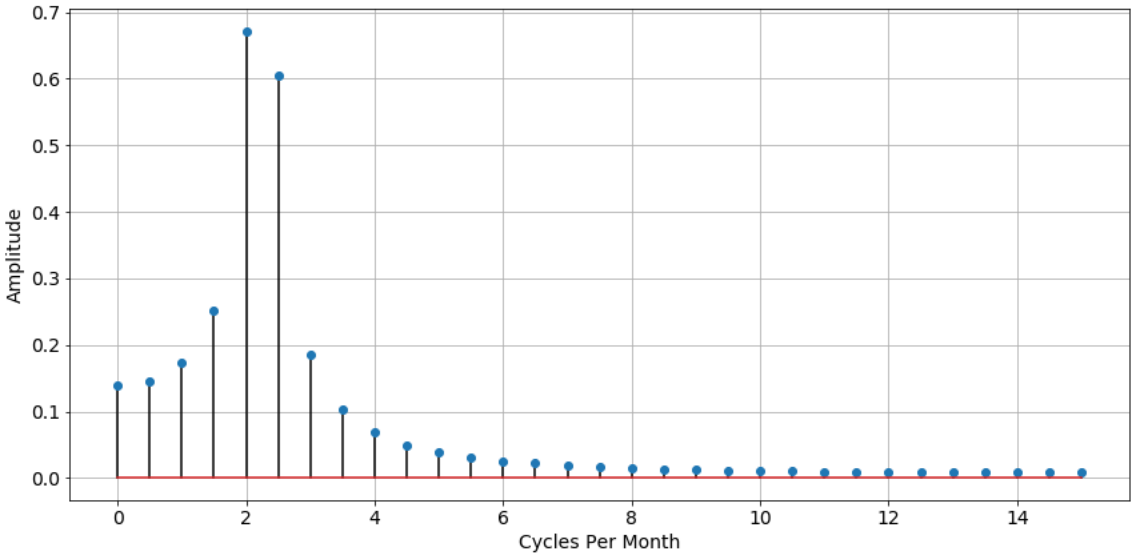

Note. As can be seen from the graphs for FFTs, there is a potentially non-zero amplitude zero frequency component included. Whilst a cycle cannot technically be cycling at a rate of zero cycles per unit time, this component actually represents the d.c. offset of the wave, where d.c. stands for *direct current* and is terminology borrowed from electronics. In contrast, a.c. (i.e, alternating current) represents cycling, or oscillating components, but a.c. is rarely (if ever) used in spectral analysis. If researchers person mean-center their time series, this offset will not exist.

Figure 8

An Example Time Series

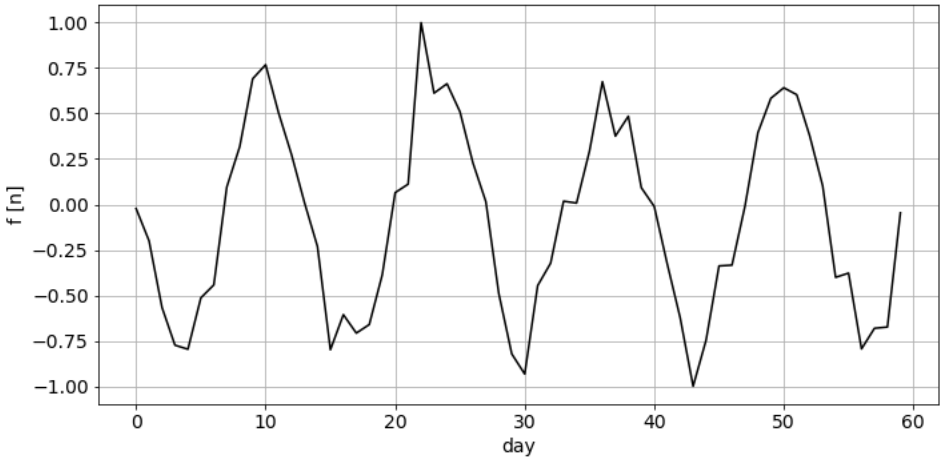

Note. Time series for $f[n] = \sin(2 \times \pi \times 4.5 \times n - 2\pi(15/30)) + \epsilon[n]$ where $\epsilon \sim \mathcal{N}(0, 1)$.

Figure 9

FFT of an Example Time Series with Noise

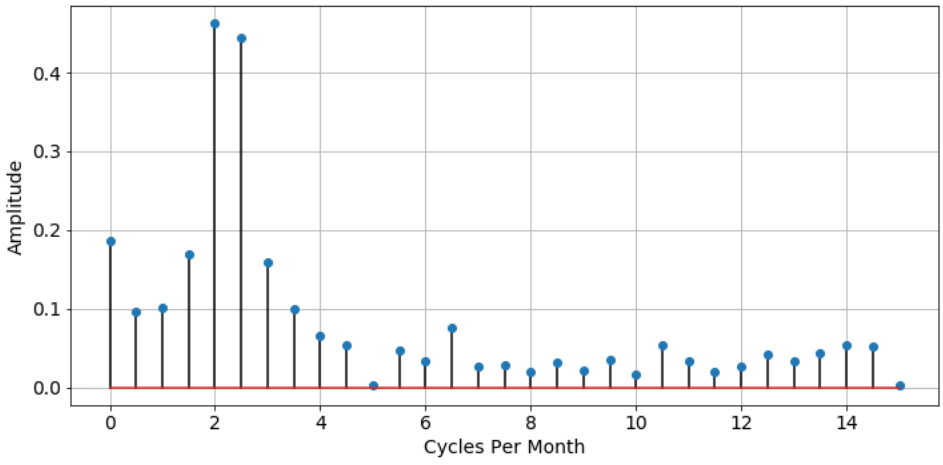

Note. FFT of the time series in Figure 8.

Figure 10

Example Time Series with Increased Complexity

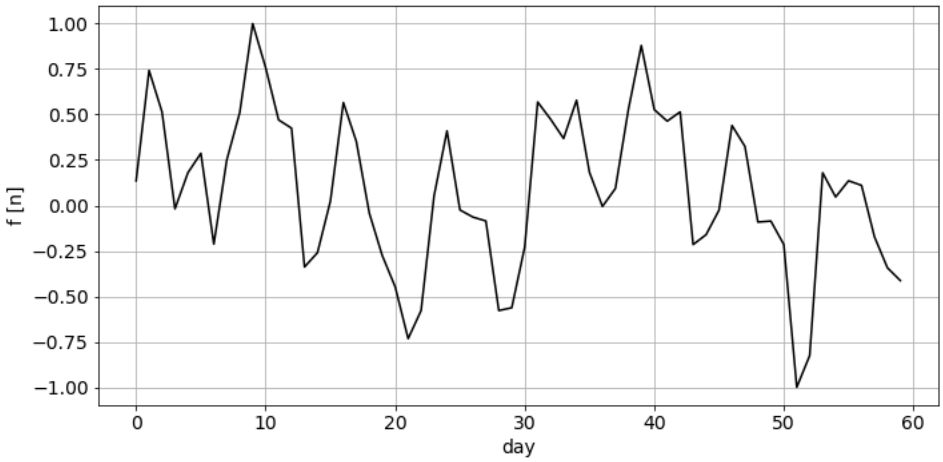

Note. Time series for $f[n] = \sin(4\pi n) + \sin(16\pi n) + 0.5\sin(32\pi n) + 0.4\epsilon[n]$ where $\epsilon \sim \mathcal{N}(0, 1)$, $n = [0, 1/N, ..., 1 - (1/N)]$ and $N = 60$.

Figure 11

The FFT for the Time Series Shown in Figure 10

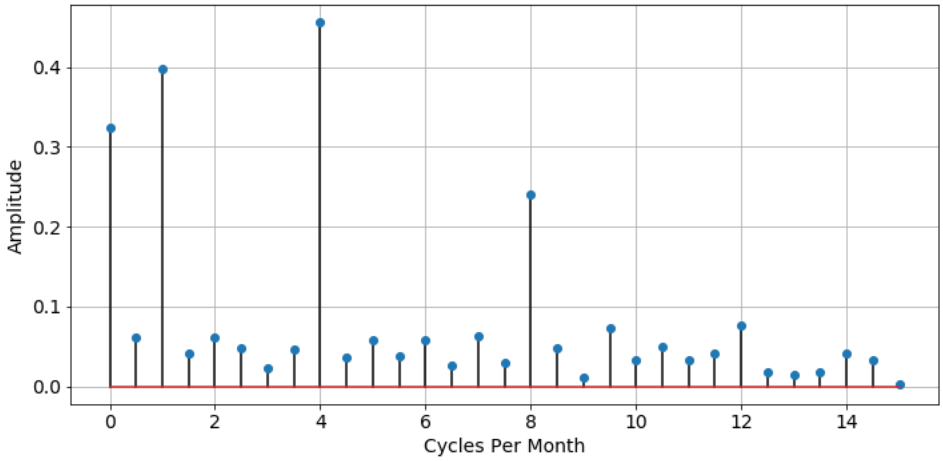

Figure 12

Example Hann-Windowed Time Series

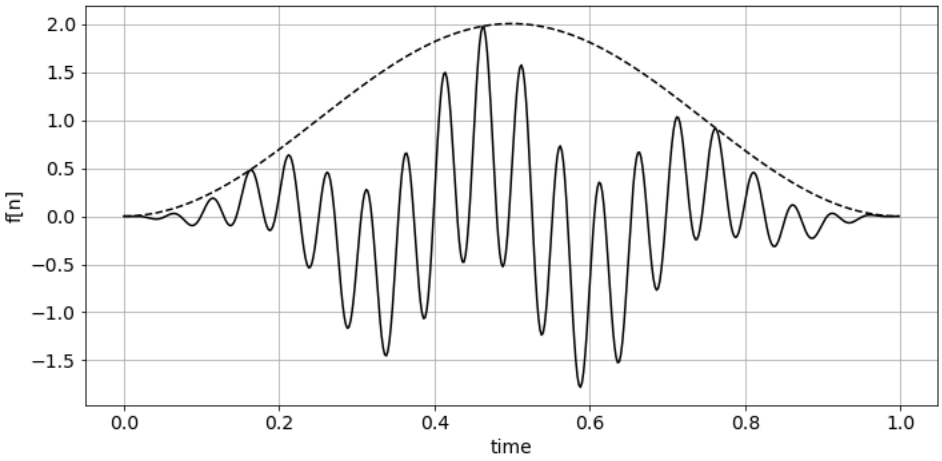

Figure 13

FFTs for Example Hann-Windowed and Un-Windowed Time Series

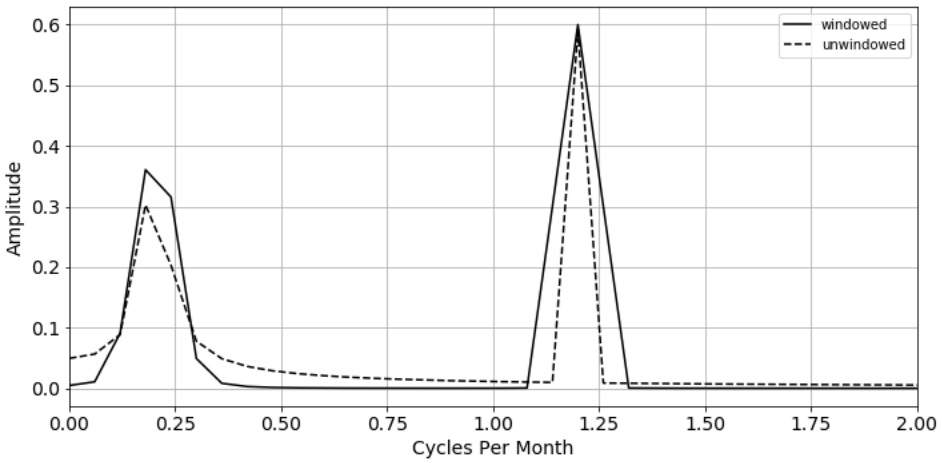

Figure 14

Bivariate Time Series

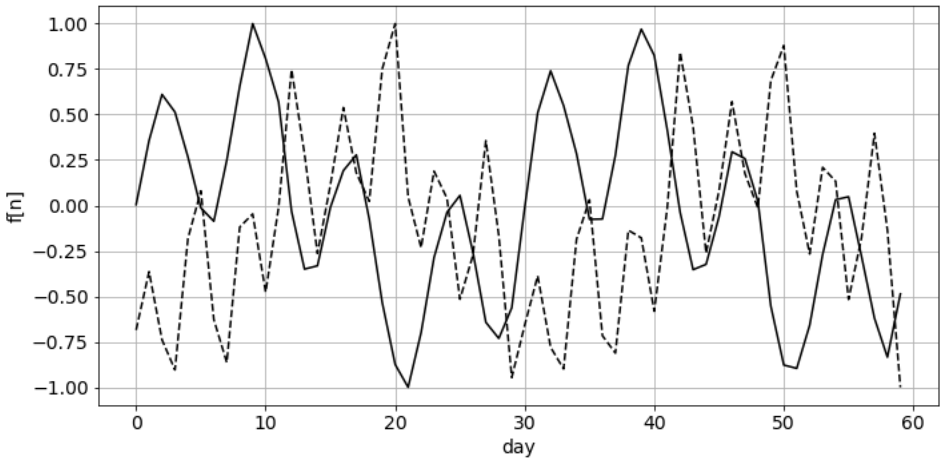

Figure 15

FFTs of the Bivariate Time series Shown in Figure 14

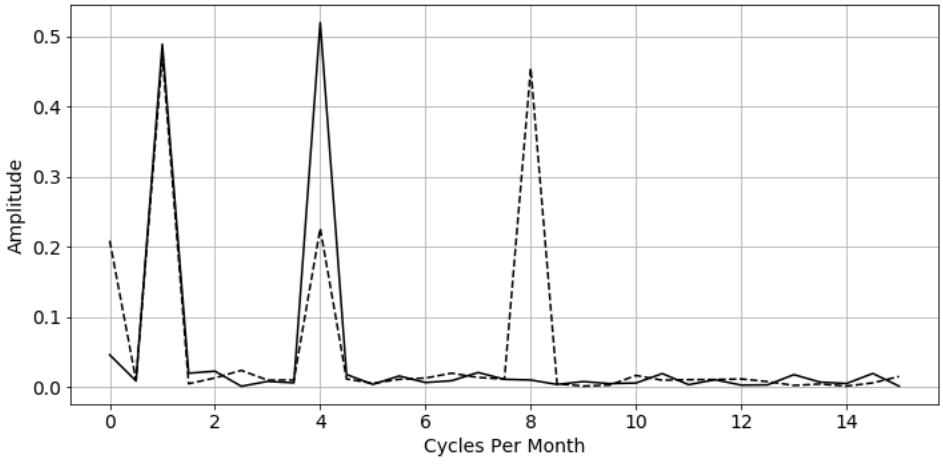

Figure 16

CPSD of the Dyadic Time Series in Figure 14

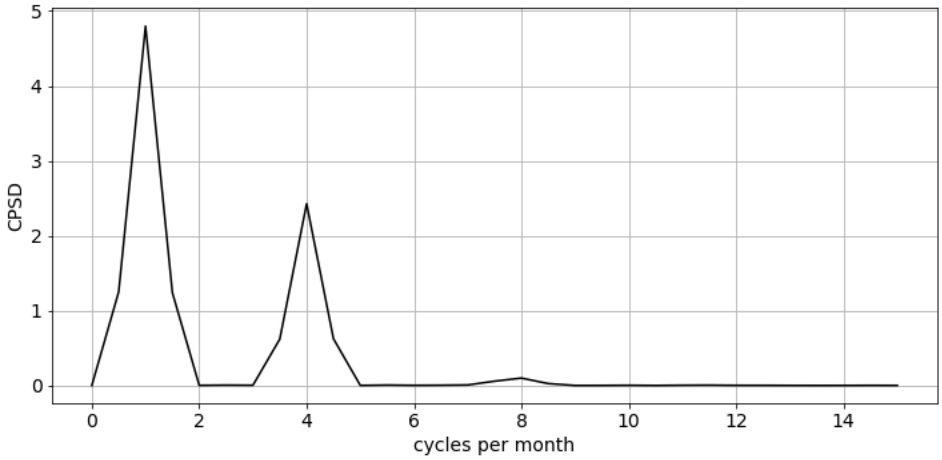

Figure 17

CPSD Phase Histogram

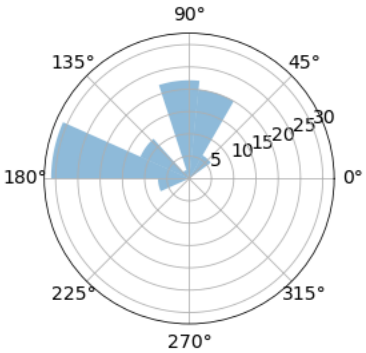

Note. An example CPSD phase polar histogram for a specific frequency component.

Figure 18

Twenty Time Series (Left) and Corresponding FFTs (Right) for Gaussian Noise

(a) *Time Series*

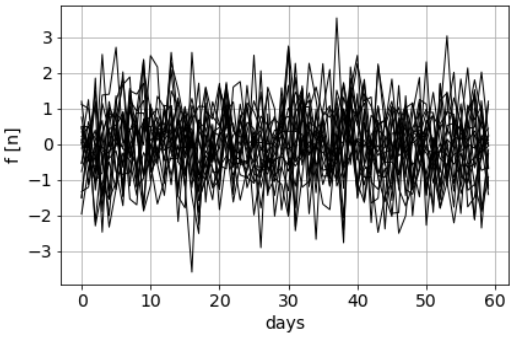

(b) *FFTs*

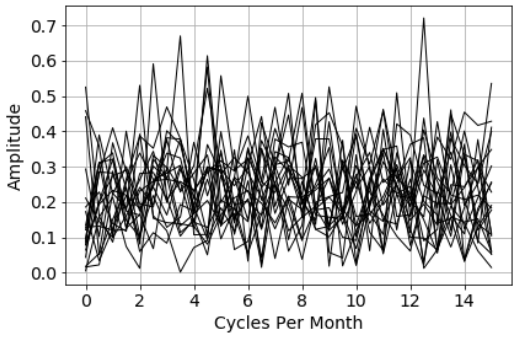

Figure 19

Average FFT for Gaussian Noise

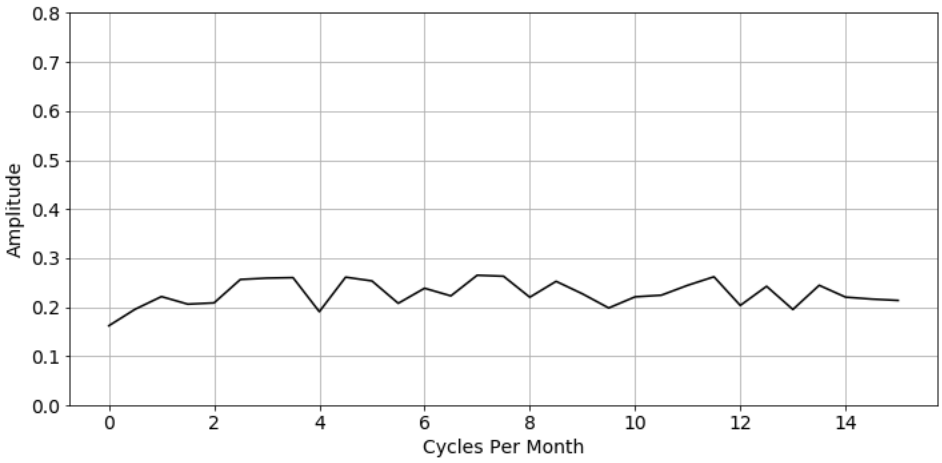

Note. Average FFT for twenty Gaussian noise time series.

Figure 20

Variance Remains Constant Before and After Permutation

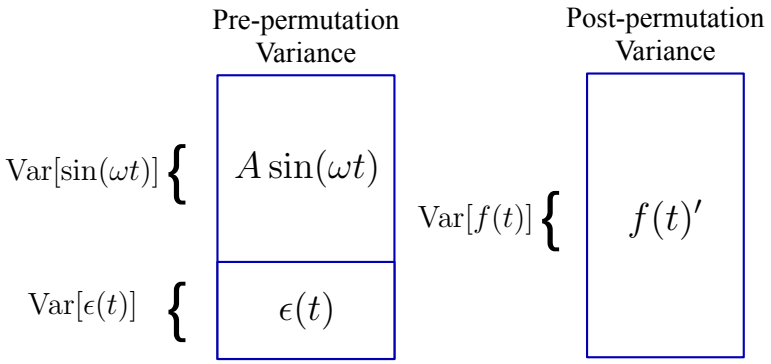

Figure 21

Example FFT and FFT of Permuted Time Series.

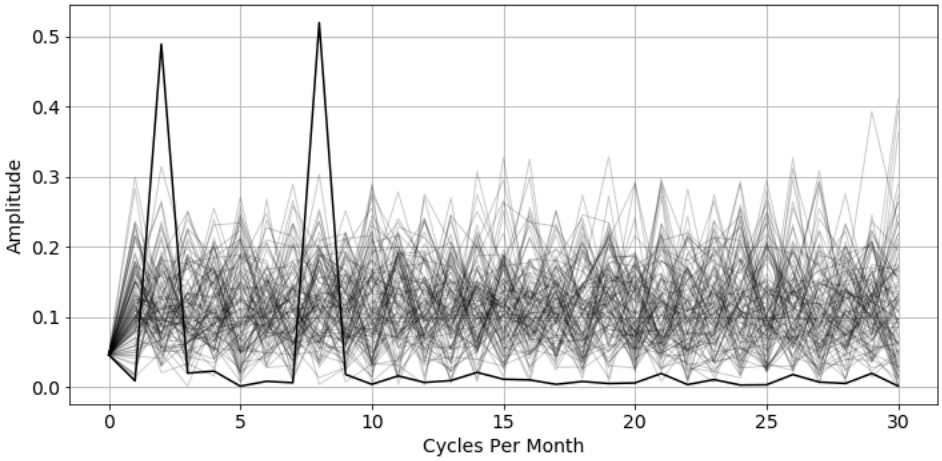

Figure 22

Significant Component Counts

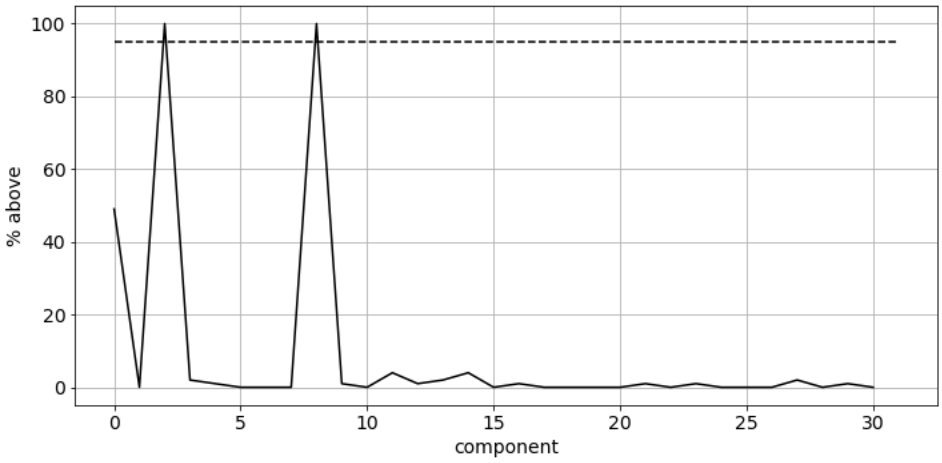

Note. For each frequency component, this graph shows the percentage frequency that the amplitude of this component falls above the corresponding amplitude components in the randomized/permuted time series FFTs. Dashed line corresponds with the 95% confidence bound.

Figure 23

De-Seasonalizing Time Series

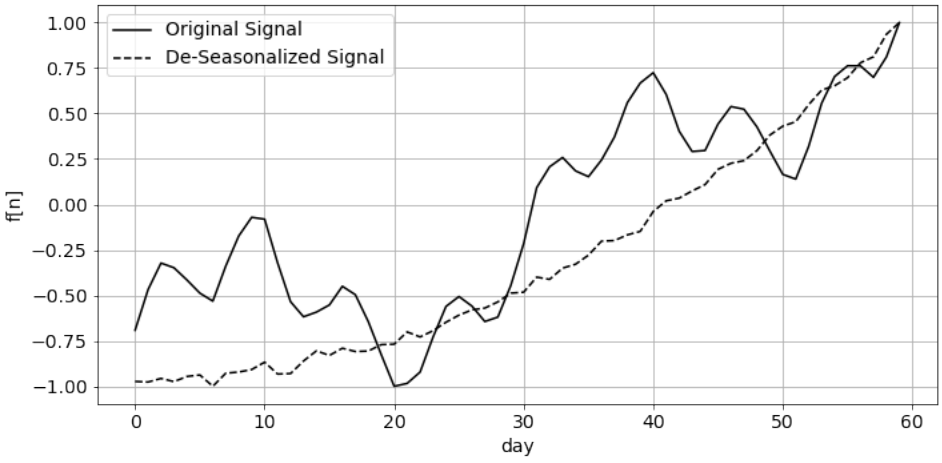

Note. A time series composed from a quadratic trend and two significant frequency components (solid line) and the same time series after de-seasonalizing the time series by subtracting the cycling components.