Published in final edited form as:

J Am Stat Assoc. 2017 ; 112(520): 1405–1416. doi:10.1080/01621459.2017.1281811.

Conditional Spectral Analysis of Replicated Multiple Time Series with Application to Nocturnal Physiology

Robert T. Krafty,

Associate Professor, Department of Biostatistics, University of Pittsburgh

Ori Rosen,

Professor, Department of Mathematical Sciences, University of Texas at El Paso

David S. Stoffer,

Professor, Department of Statistics, University of Pittsburgh

Daniel J. Buysse, and

Professor, Department of Psychiatry, University of Pittsburgh

Martica H. Hall

Professor, Department of Psychiatry, University of Pittsburgh

Abstract

This article considers the problem of analyzing associations between power spectra of multiple time series and cross-sectional outcomes when data are observed from multiple subjects. The motivating application comes from sleep medicine, where researchers are able to non-invasively record physiological time series signals during sleep. The frequency patterns of these signals, which can be quantified through the power spectrum, contain interpretable information about biological processes. An important problem in sleep research is drawing connections between power spectra of time series signals and clinical characteristics; these connections are key to understanding biological pathways through which sleep affects, and can be treated to improve, health. Such analyses are challenging as they must overcome the complicated structure of a power spectrum from multiple time series as a complex positive-definite matrix-valued function. This article proposes a new approach to such analyses based on a tensor-product spline model of Cholesky components of outcome-dependent power spectra. The approach explicitly models power spectra as nonparametric functions of frequency and outcome while preserving geometric constraints. Formulated in a fully Bayesian framework, a Whittle likelihood based Markov chain Monte Carlo (MCMC) algorithm is developed for automated model fitting and for conducting inference on associations between outcomes and spectral measures. The method is used to analyze data from a study of sleep in older adults and uncovers new insights into how stress and arousal are connected to the amount of time one spends in bed.

Correspondence to: Martica H. Hall.

Code

A Matlab program for implementing the proposed methodology and a file demonstrating its use are available at <https://www.mathworks.com/matlabcentral/fileexchange/61186-mcbspec>.

Keywords

Bayesian Analysis; Coherence; Heart Rate Variability; MCMC; Multivariate Time Series; Sleep; Smoothing Spline; Spectral Analysis; Tensor-Product ANOVA; Whittle Likelihood

1 Introduction

Innovations in data collection and storage have led to an increase in the number of biomedical studies that record multiple time series signals and outcome measures in multiple subjects. For many time series, including common signals such as blood pressure, heart rate and electroencephalography (EEG), frequency patterns that are quantified through the power spectrum contain important information about biological processes. Consequently, studies whose goals are to understand how underlying biological mechanisms are connected to behavioral and clinical outcomes often require an analysis of associations between outcomes and power spectra of multiple time series.

Our motivating application comes from a sleep study whose goal is to better understand the pathways that connect sleep to health and functioning. In the study, heart rate variability (HRV) is recorded in subjects during a night of sleep. HRV is measured through the series of elapsed times between consecutive heart beats, and its power spectrum provides indirect measures of psychological stress and physiological arousal (Hall et al., 2007). Upon awakening, subjects reported subjectively assessed sleep outcomes, such as the amount of time slept during the night, which are associated with many aspects of well-being (Buysse, 2014). Understanding the association between the power spectrum of HRV during different sleep periods (i.e. beginning, middle and end of the night) and self-reported sleep outcomes is essential to understanding how stress connects sleep to health and, consequently, for guiding the use of treatments of poor sleep.

In the biomedical literature, a two-stage approach is typically used to analyze such data. In the first stage, power collapsed within pre-selected frequency bands is estimated individually for each time series (Malik et al., 1996; Hall et al., 2004). A power spectrum is a function of frequency; power collapsed within a frequency band is an integral of the power spectrum over a range of frequencies, which converts the functional parameter into a scalar measure. In the second stage, classical statistical methods, such as ANOVA and linear regression, are used to evaluate associations between these band-collapsed spectral measures and outcomes. Such an approach has three major drawbacks. First, it is highly dependent on the frequency-band collapsed measures selected and there exists a hot debate as to which measures should be considered and/or how they should be interpreted (Burr, 2007). Ideally, an analysis of such data should provide global measures that can be used to understand the entire system while also providing a means to conduct inference on any frequency band-collapsed measure of potential interest. Second, estimated power is treated as if it were not an estimate but the true unknown parameter, leading to inaccurate inference. Finally, band-specific frequency measures are estimated for each time series separately, inhibiting the evaluation of patterns across series. For instance, in our motivating example, this two-stage approach does not

examine how the coherence in HRV between the beginning and end of the night is connected to sleep outcomes.

In the statistics literature, a considerable amount of research has been conducted on methods for analyzing functional variables, a thorough review of which is given by Wang et al. (2016). Included in this body of work are methods for analyzing associations between power spectra and outcomes when one time series is observed per subject (Stoffer et al., 2010; Krafty and Hall, 2013). When one observes multiple time series per subject and interest lies in frequency patterns both within each series and across different series, the problem becomes considerably more challenging. This is the case in our motivating study, where we are interested not only in stress and arousal during particular periods of sleep, but also in their persistence and coherence across periods. While the power spectrum from a single time series is a positive real-valued function of frequency, the power spectrum from multiple time series is a positive-definite Hermitian matrix valued function of frequency. An analysis of associations between power spectra from multiple time series and study outcomes must be able to exhibitly model associations while preserving this positive-definite Hermitian structure.

Efficient nonparametric methods that preserve the positive-definite Hermitian structure of spectral matrices have been developed for the simpler, classical problem of estimating the power spectrum of a multivariate time series from a single subject by modeling Cholesky components of spectral matrices as functions of frequency (Dai and Guo, 2004; Rosen and Stoffer, 2007; Krafty and Collinge, 2013). In this article, we extend this framework to develop a new approach to analyzing data from multiple subjects that models Cholesky components as functions of both frequency and outcome. Rather than being curves as functions of frequency, components of spectral matrices under the proposed model are surfaces. Changes in these surfaces with respect to the outcome provide nonparametric measures of association between outcomes and power spectra. The proposed method is formulated in a fully Bayesian framework; a MCMC algorithm based on the Whittle likelihood, or the asymptotic likelihood derived from the Fourier transform of the data, is developed for model fitting and inference. The method allows one to evaluate the entire outcome-dependent power spectrum and to conduct nonparametric inference on the association between the outcome and any function of the power spectrum.

The rest of the article is organized as follows. Our motivating application, the AgeWise Sleep Study, is discussed in Section 2. A review of spectral analysis in the classical setting, where data are observed from a single subject, is given in Section 3. The proposed methodology for analyzing time series from multiple subjects is presented in Section 4. The proposed method is used to analyze data from the motivating application in Section 5 and some final remarks are offered in Section 6.

2 The AgeWise Sleep Study

An estimated 43% of older adults report problems initiating or maintaining sleep (Foley et al., 1995). Poor sleep in older adults has been linked to depression, heart disease, obesity, arthritis, diabetes and stroke (Foley et al., 2004). With medical and scientific advances leading to an increase in the world's elderly population, the consequences of poor sleep in

older adults pose a major public health concern. The AgeWise study is a NIH-funded Program Project conducted at the University of Pittsburgh that seeks a better understanding of causes, effects, and treatments of poor sleep in older adults. Towards this goal, we consider $N = 108$ men and women between 69–89 years of age who were observed during a night of in-home sleep. Two types of data were collected in each subject. First, subjects were observed during the night through ambulatory polysomnography (PSG), or the continuous collection of electrophysiological changes that occur during sleep. Second, upon awakening, subjects completed the Pittsburgh Sleep Diary (Monk et al., 1994) to record self-reported sleep outcomes during the night.

As previously discussed, HRV is the series of elapsed times between heart beats. It is of interest to researchers, as it reflects neurological control of the heart, and through this capacity, its power spectrum provides indirect measures of stress reactivity and arousal. The PSG used in the study included an electrocardiograph (ECG) to monitor heart activity. The ECG was used to locate the timing of heart beats, which were then differenced, detrended, cubic spline interpolated, and resampled at 1 Hz to compute HRV series throughout the night.

During the night, the body cycles through two types of sleep: rapid eye movement (REM) and non-rapid eye movement (NREM) sleep. In NREM sleep, which contains deepsleep, the parasympathetic branch of the autonomic nervous system that is responsible for unconscious actions and stimulates the body to “rest-and-digest” dominates the sympathetic branch, which drives the “flight-or-fight” response. Parasympathetic nervous system activity during NREM is hypothesized to be responsible for many of the rejuvenating properties of sleep (Siegel, 2005). However, physiological activity during NREM sleep is not constant; in good sleepers, the amount of parasympathetic activity during NREM increases throughout the night (Hall et al., 2004). To enable an analysis that can evaluate autonomic nervous system activity and its changes during the night, we consider 3 HRV time series per subject (at the beginning, middle, and end of the night) by extracting the first 5 minutes of HRV from the first three periods of NREM sleep. Data from two subjects are displayed in Figure 1.

The goal of our analysis is to understand how the power spectrum of HRV over the three periods of NREM are connected to self-reported sleep. We focus on one particular self-reported sleep measure derived from the Pittsburgh Sleep Diary: time in bed (TIB). TIB is defined as the elapsed time between attempted sleep and final waking. It serves as an upper bound for the amount of time spent asleep during the night, which has been linked to heart disease, hypertension, impaired neurobehavioral performance and mortality (Buysse, 2014). The reported TIB from our sample has a mean of 477.99 minutes and a standard deviation of 71.32 minutes. The resulting data for analysis consist of three epochs of HRV, one during each of the first three periods of NREM sleep, and self-reported TIB from each subject.

3 Methodological Background: Spectral Domain Analysis

Before introducing our proposed model for the spectral analysis of multiple time series from multiple subjects in Section 4, in this section we present background on spectral analysis in

the classical setting, where data are observed from a single subject, for both univariate and multivariate time series.

3.1 Univariate Time Series

3.1.1 Population Parameters—Spectral domain analysis focuses on the cyclic behavior of time series data. An alternate approach is time domain analysis wherein the relationship between the data at different time lags is the focus. For stationary time series, the main time domain tool is the covariance between a current value of the series, say X_t , and the value of the series h time units prior, say X_{t-h} . The autocovariance is a function of lag, and is given by

$$\gamma(h) = \text{Cov}(X_t, X_{t-h}), \quad h=0, \pm 1, \pm 2, \dots$$

If $\gamma(h)$ is absolutely summable (which it is for ARMA models, for example), then there is a duality between the power spectrum, given by

$$f(\omega) = \sum_{h=-\infty}^{\infty} \gamma(h) \exp(-2\pi i \omega h), \quad \omega \in \mathbb{R},$$

and the autocovariance function, namely,

$$\gamma(h) = \int_{-1/2}^{1/2} f(\omega) \exp(2\pi i \omega h) d\omega, \quad h=0, \pm 1, \pm 2, \dots, \quad (1)$$

as the inverse transform of the power spectrum. The relationship is the same as that of a characteristic function and a probability density. Consequently, the information contained in the power spectrum is equivalent to the information contained in the autocovariance function. If we are concerned with lagged behavior, working with $\gamma(h)$ is more informative; if we are concerned with cyclic behavior, as is the case of HRV where cyclical behavior provides interpretable physiological information, working with $f(\omega)$ is more informative.

The power spectrum is nonnegative and we assume that it is positive, so that $f(\omega) > 0$ for all ω . In addition to being positive, f has two other restrictions as a function of frequency. By the nature of the Fourier transform, it is periodic such that $f(\omega) = f(\omega + 2\pi)$, and it is a Hermitian function, or an even function, where $f(\omega) = f(-\omega)$. Consequently, $f(\omega)$ is usually displayed only for $\omega \in [0, 1/2]$.

Putting $h = 0$ in (1) yields

$$\gamma(0) = \text{Var}(X_t) = \int_{-1/2}^{1/2} f(\omega) d\omega,$$

which expresses the total variance of the time series as the integrated power spectrum. In particular, we may think of $f(\omega) d\omega$ as the approximate variance in the data attributed to frequencies in a small band of width $d\omega$ around ω . It is common to view spectral analysis as an analysis of variance (ANOVA) of time series data with respect to frequency. In fact, the power spectrum is a density of variance rather than of probability.

3.1.2 Estimation—The nonparametric estimation of f from an epoch of length n , $X_1, \dots, X_n$, can begin by considering the discrete Fourier transform (DFT)

$$Y_m = n^{-1/2} \sum_{t=1}^n X_t \exp(-2\pi i \omega_m t),$$

where $\omega_m = m/n$ are the Fourier frequencies. When n is large, Y_m are approximately independent mean-zero complex normal random variables with variances $f(\omega_m)$ for $m = 1, \dots, M$, $M = \lfloor (n-1)/2 \rfloor$, and $Y_m = \bar{Y}_{n-m}$ (Shumway and Stoffer, 2011, Appendix C). Consequently, the periodogram $|Y_m|^2$ provides approximately unbiased but noisy estimates of $f(\omega_m)$. Consistent estimates can be obtained by smoothing the periodogram across frequency using tools such as local averaging (Shumway and Stoffer, 2011, Chapter 4.5), splines (Pawitan and O'Sullivan, 1994), and wavelets (Moulin, 1994).

Our estimation approach for multiple time series from multiple subjects, which we develop in Section 4, is based on Bayesian splines. To motivate the development of the new methodology, in this subsection we discuss first a Bayesian smoothing spline model for univariate spectral analysis, then discuss a low-rank approximation. Smoothing spline estimation balances the fit of a function to observed data with a roughness-based measure of regularity. The Bayesian formulation of smoothing splines was first discussed in the case of Gaussian observations by Kimeldorf and Wahba (1970) and Wahba (1978). Under the Bayesian formulation, the likelihood provides a measure of fit to observed data, and regularity is imposed through a mean-zero Gaussian prior on the functional parameter, which induces a prior for the roughness of the function.

For spectrum estimation, the large sample distribution of Y_m provides the Whittle likelihood (Whittle, 1953, 1954)

$$L(Y|f) \approx \prod_{m=1}^M f^{-1}(\omega_m) \exp\{-f^{-1}(\omega_m)|Y_m|^2\}.$$

We adopt generic notation throughout this article where Y will denote all DFT data. Since f is a positive function, $\log f$ is modeled rather than f itself to avoid constraints. Although general measures of regularity can be considered, we focus on measuring the roughness of a function through its integrated squared second derivative,

$$\mathcal{P}(\log f) = \int_0^{1/2} \{[\log f]''(\omega)\}^2 d\omega.$$

The specification of the prior distribution begins by decomposing $\log f$ into a linear part (which is in the null space of $\mathcal{P}$) and a nonlinear part. To define the prior distribution for the nonlinear part, consider the reproducing kernel of the seminorm defined by $\mathcal{P}$

$$J(\omega_i, \omega_j) = \int_0^{1/2} (\omega_i - \nu)_+ (\omega_j - \nu)_+ d\nu,$$

where $(\nu)_+ = \max(\nu, 0)$ (Gu, 2013, Section 2.3.1). For the Bayesian smoothing spline model, the prior distribution for the log-spectrum can be formulated as

$$\log f(\omega) = a_1 + a_2 \omega + \sum_{j=1}^M J(\omega, \omega_j) z_j$$

where $\mathbf{z} = (z_1, \dots, z_M)' \sim N(\mathbf{0}, \tau^2 \mathbf{J}^{-1})$ is independent of $\mathbf{a} = (a_1, a_2)' \sim N(0, \sigma_\alpha^2 \mathbf{I}_2)$, $\mathbf{J} = \{J(\omega_i, \omega_j)\}$ is the $M \times M$ matrix of J evaluated at the Fourier frequencies and $\mathbf{I}_2$ is the 2×2 identity matrix. The reproducing property of the kernel J provides a simple form for the roughness of the log-spectrum, $\mathcal{P}(\log f) = \mathbf{z}' \mathbf{J} \mathbf{z}$ (Gu, 2013, Chapter 2), from which it can be seen that the prior distribution on the coefficients $\mathbf{z}$ induces a prior on the roughness of the spectrum

where $\mathcal{P}(\log f) \sim \tau^2 \chi_M^2$ and χ_M^2 denotes a chi-squared random variable on M degrees of freedom. The smoothing parameter $\tau^2 > 0$ balances the smoothness of the estimator to its fit to the data such that, conditional on τ^2 , the Bayes estimate (posterior median) approaches a linear function as $\tau^2 \rightarrow 0$ and interpolates $|Y_m|^2$ as $\tau^2 \rightarrow \infty$.

Two approaches may be taken for estimation and inference under the Bayesian model: empirical Bayes and fully Bayes. In the empirical Bayes approach, a data driven method, such as generalized cross-validation (GCV) or generalized maximum likelihood (GML), is used to select the smoothing parameter τ^2 . The log-spectrum is then estimated from the posterior distribution conditional on τ^2 , and its median is equivalent to the frequentist smoothing spline obtained by minimizing a penalized Whittle likelihood as $\sigma_\alpha^2 \rightarrow \infty$ (Gu, 1992; Qin and Wang, 2008). In the fully Bayesian approach, τ^2 is treated as a random variable with a noninformative prior, and inference is conducted averaged over the posterior distribution of τ^2 (Speckman and Sun, 2003; Crainiceanu et al., 2005). Our proposed methodology will adopt the later approach, as discussed in Section 4.

The presented smoothing spline model contains a large number of coefficients, which can impede computation and limit practicality. Low rank approximations, such as those obtained by using a subset of the kernel functions $J(\cdot, \omega_j)$ (Gu and Kim, 2002) or another set of basis functions contained in the column space of $\mathbf{J}$ (Wood, 2006), can be used to ease computational burden without sacrificing model fit. Here, we consider the basis formed from the scaled eigenvectors of $\mathbf{J}$, which has been used for power spectrum estimation by Wood et

al. (2002), Rosen et al. (2009) and Rosen et al. (2012). This basis can model smooth functions with a relatively few number of basis functions, provides a diagonal prior covariance structure that aids computation, and maintains the intuitive interpretation of the prior distribution regularizing roughness as measured through $\mathcal{P}$ (Nychka and Cummins, 1996). To formulate this low-rank approximation, we first consider an equivalent formulation of the Bayesian smoothing spline model for $\log f$ at the Fourier frequencies. Let $\log \mathbf{f} = [\log f(\omega_1), \dots, \log f(\omega_M)]'$ be the log-spectrum evaluated at the Fourier frequencies. Further, let $J = V_J D_J V_J'$ be the spectral decomposition of J , $Q_J = V_J D_J^{1/2}$ and $\mathbf{c} = D_J^{1/2} V_J' \mathbf{z}$. Then an equivalent model at the Fourier frequencies is

$$\log \mathbf{f} = L_J \mathbf{a} + Q_J \mathbf{c}.$$

where $L_J = (\mathbf{1}_M \mid \boldsymbol{\omega})$, $\mathbf{1}_M$ is the M -vector of ones, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_M)'$, and $\mathbf{c} \sim N(\mathbf{0}, \tau^2 I_M)$.

The eigenvectors contained in the columns of V_J are in increasing order of roughness and the eigenvalues contained in the diagonal matrix D_J decay rapidly (Eubank, 1999). Smooth functions can be accurately modeled through the first several $n_J < M$ to provide a low-rank approximation. With a slight abuse of notation to avoid the need to introduce further variables, we let Q_J represent the $M \times n_J$ matrix of the first n_J eigenvectors with corresponding coefficients $\mathbf{c} \sim N(\mathbf{0}, \tau^2 I_{n_J})$.

The selection of n_J provides a compromise between low-rank computational feasibility and loss of exibility relative to full M -rank model. An intuitive measure of the loss of exibility is the fraction of the total variance of the covariance matrix J explained by its first n_J eigenvectors (FVE), or the sum of its first n_J eigenvalues divided by the total sum of all of its M eigenvalues. Wood et al. (2002), Rosen et al. (2009) and Rosen et al. (2012) suggest using $n_J = 10$ basis functions, which, in each of their settings considered, equates to an FVE of 97.975%. Our empirical findings support this suggestion, and we recommend selecting n_J to achieve a 97.975% FVE. Under this rule, $n_J = 7$ for $n = [15, 18]$, $n_J = 8$ for $n = [19, 22]$, $n_J = 9$ for $n \in [23, 40]$ and $n_J = 10$ for $n \in [41, 10^4]$.

3.2 Multivariate Time Series

3.2.1 Population Parameters—The ideas presented in the univariate case generalize to the multivariate case wherein we observe a P -dimensional vector-valued time series, say $\{\mathbf{X}_t\}$. Under stationarity, the autocovariance function is a $P \times P$ matrix given by

$$\Gamma(h) = \text{Cov}(\mathbf{X}_t, \mathbf{X}_{t-h}), \quad h = 0, \pm 1, \pm 2, \dots$$

If $\sum_h \|\Gamma(h)\| < \infty$, the spectral density matrix of the series $\mathbf{X}_t$ is given by

$$f(\omega) = \sum_{h=-\infty}^{\infty} \Gamma(h) \exp(-2\pi i \omega h), \quad \omega \in \mathbb{R}$$

and the inverse relationship is

$$\Gamma(h) = \int_{-1/2}^{1/2} f(\omega) \exp(2\pi i \omega h) d\omega, \quad h=0, \pm 1, \pm 2, \dots$$

For each $\omega \in \mathbb{R}$, $f(\omega)$ is a $P \times P$ non-negative definite Hermitian matrix with the diagonal elements, $f_{pp}(\omega)$ for $p = 1, \dots, P$, being the spectra of the individual components, and the off-diagonal elements, $f_{qp}(\omega)$ for $q \neq p = 1, \dots, P$, being the cross-spectra. Throughout this article, we assume that $f(\omega)$ is non-singular for all $\omega \in \mathbb{R}$. As in the univariate case, f is a periodic and Hermitian function of frequency were, for matrix-valued functions, Hermitian as a function of frequency is defined as $f(\omega) = f^*(-\omega)$, and $f^*(\omega)$ is the complex conjugate of $f(\omega)$.

An important example of the application of the cross-spectrum is to the problem of linearly predicting one of the component series, say X_{qt} , from another component, say X_{pt} . A measure of the strength of such a relationship is the squared coherence function defined as

$$\rho_{qp}^2(\omega) = \frac{|f_{qp}(\omega)|^2}{f_{qq}(\omega)f_{pp}(\omega)}.$$

This is analogous to conventional squared correlation between two finite-variance random variables; e.g., $0 \leq \rho_{qp}^2(\omega) \leq 1$. This analogy motivates the interpretation of squared coherence as the squared correlation between two time series at frequency ω . These ideas extend in an obvious way to the concept of multiple coherence and partial coherence functions obtained from the full spectral matrix in much the same way that multiple correlation and partial correlation can be obtained from a covariance matrix. Full details of these results may be found in Shumway and Stoffer (2011, Chapters 4 & 7).

3.2.2 Estimation—In the multivariate setting, let

$$\mathbf{Y}_m = n^{-1/2} \sum_{t=1}^n \mathbf{X}_t \exp(-2\pi i \omega_m t)$$

be the DFTs of the data. In this case, the Whittle likelihood is

$$L(Y|f) \approx \prod_{m=1}^M |f^{-1}(\omega_m)| \exp\{-\mathbf{Y}_m^* f^{-1}(\omega_m) \mathbf{Y}_m\},$$

and the periodogram $\mathbf{Y}_m \mathbf{Y}_m^*$ is an approximately unbiased but noisy estimate of $f(\omega_m)$, from which consistent estimates can be obtained by smoothing.

While in the univariate setting the spectrum is smoothed on the logarithmic scale to preserve positivity, Cholesky components of spectral matrices can be smoothed to preserve positive-definiteness in the multivariate setting (Dai and Guo, 2004; Rosen and Stoffer, 2007; Krafty and Collinge, 2013). The modified Cholesky decomposition assures that, for a spectral matrix $f(\omega)$, there exists a unique $P \times P$ lower triangular complex matrix $\Theta(\omega)$ with ones on the diagonal and a unique $P \times P$ positive diagonal matrix $\Psi(\omega)$ such that

$$f^{-1}(\omega) = \Theta(\omega) \Psi^{-1}(\omega) \Theta^*(\omega).$$

There are P^2 -Cholesky components to estimate: $\Re\{\Theta_{k\ell}\}$ and $\Im\{\Theta_{k\ell}\}$ for $k > \ell = 1, \dots, P-1$, and Ψ_{kk}^{-1} for $k = 1, \dots, P$. Since the diagonal terms $\Psi_{kk}^{-1}(\omega) > 0$, we model $\log \Psi_{kk}^{-1}$.

Letting $\boldsymbol{\theta}_{k\ell} = [\Theta_{k\ell}(\omega_1), \dots, \Theta_{k\ell}(\omega_M)]'$ and $\log \psi_{kk}^{-1} = [\log \Psi_{kk}^{-1}(\omega_1), \dots, \log \Psi_{kk}^{-1}(\omega_M)]'$, we model:

$$\Re\{\boldsymbol{\theta}_{k\ell}\} = L_J \mathbf{a}_{rk\ell} + Q_J \mathbf{c}_{rk\ell}, \quad k > \ell = 1, \dots, P-1 \quad (2)$$

$$\Im\{\boldsymbol{\theta}_{k\ell}\} = L_J \mathbf{a}_{ik\ell} + Q_J \mathbf{c}_{ik\ell}, \quad k > \ell = 1, \dots, P-1 \quad (3)$$

$$\log \psi_{kk}^{-1} = L_J \mathbf{a}_{dkk} + Q_J \mathbf{c}_{dkk}, \quad k = 1, \dots, P, \quad (4)$$

where

$\mathbf{c}_{rk\ell} \sim N(\mathbf{0}, \tau_{rk\ell}^2 I_n)$, $\mathbf{c}_{ik\ell} \sim N(\mathbf{0}, \tau_{ik\ell}^2 I_n)$, $\mathbf{c}_{dkk} \sim N(\mathbf{0}, \tau_{dkk}^2 I_n)$, $\mathbf{a}_{rk\ell} \sim N(\mathbf{0}, \sigma_{\alpha}^2 I_2)$, $\mathbf{a}_{ik\ell} \sim N(\mathbf{0}, \sigma_{\alpha}^2 I_2)$ and $\mathbf{a}_{dkk} \sim N(\mathbf{0}, \sigma_{\alpha}^2 I_2)$. Throughout this article, r , i and d are used to denote coefficients for real components of Θ , imaginary components of Θ and the logarithm of the diagonal components of Ψ^{-1} , respectively.

4 Methodology: Replicated Multiple Time Series

The primary question considered in this article is how to assess the association between the power spectrum of P -variate time series of length n , $\{\mathbf{X}_{j1}, \dots, \mathbf{X}_{jn}\}$, and real-valued static variables, U_j , observed from $j = 1, \dots, N$ independent subjects. In the motivating study, there are $N = 108$ participants, U_j is self-reported TIB, and $\mathbf{X}_{jt}$ are time series of HRV during the first $P = 3$ periods of NREM. To address this question, we first introduce a new measure in Section 4.1, the conditional power spectrum, which quantifies associations between power spectra and outcomes. Then, in Section 4.2, we develop a tensor product model for the conditional power spectrum that extends the Bayesian spline model of Cholesky components of a single multivariate time series to account for dependence on both frequency and outcome.

As previously mentioned, there are two approaches to conducting a Bayesian analysis with splines: empirical Bayes and fully Bayesian. Each approach has strengths and weaknesses. In the empirical Bayes approach, smoothing parameters are estimated through a data-driven procedure. Estimates conditional on smoothing parameters can be quickly computed through Fisher's scoring or Newton-Raphson and conditional inference on the modeled functions (Cholesky components in our setting) can be conducted through approximate "Bayesian confidence intervals" (Gu, 1992). In the fully Bayesian approach, smoothing parameters τ^2 are treated as random variables with noninformative priors and MCMC techniques are used to sample from the joint distribution of coefficients and smoothing parameters (Speckman and Sun, 2003; Crainiceanu et al., 2005). The sample simulated from the posterior distribution using MCMC provides a natural means of conducting inference on any function of the spectrum averaged over the distribution of the smoothing parameters, which accounts for uncertainty in the smoothing parameters when conducting inference. As will be illustrated in Section 5, inference on squared coherence, univariate spectra, and integral functions thereof (all of which are nonlinear functions of modeled Cholesky components) are of direct scientific interest. We develop the proposed methodology under a fully Bayesian framework, presenting prior distributions in Section 4.3 and the sampling scheme in Section 4.4.

4.1 Conditional Power Spectrum

Without loss of generality, we formulate the methodology assuming that U_j is scaled to take values within $[0, 1]$. To quantify the association between the power spectrum of the time series $\mathbf{X}_{jt}$ and the outcome U_j , we define the conditional power spectrum

$$f(\omega, u) = \sum_{h=-\infty}^{\infty} \text{Cov}(\mathbf{X}_{jt}, \mathbf{X}_{j,t+h} | U_j = u) e^{-2\pi i \omega h}, \quad \omega \in \mathbb{R}, u \in [0, 1].$$

As with the power spectrum of a single multivariate time series, the spectral matrices $f(\omega, u)$ are positive-definite $P \times P$ Hermitian matrices, and $f(\cdot, u)$ is a periodic and Hermitian function of frequency for fixed u . In a traditional spectral analysis without a cross-sectional variable, spectral measures such as f_{pq} and $\rho_{pq}^2 = |f_{pq}|^2 / (f_{pp}f_{qq})$ are curves as functions of frequency. In the conditional setting, these are surfaces as functions of both frequency and the variable u . How these functions change with respect to u provides information as to how spectral measures are associated with the variable.

4.2 Bayesian Tensor-Product Model of Cholesky Components

As in the classical setting discussed in Section 3.2.2, where a multivariate time series observed from a single subject, to preserve positive definiteness, we model the Cholesky components. Let

$$f^{-1}(\omega, u) = \Theta(\omega, u) \Psi^{-1}(\omega, u) \Theta^*(\omega, u),$$

be the modified Cholesky decomposition of the conditional power spectrum. We use Bayesian tensor product models for the P^2 -unique Cholesky components, which decompose the bivariate functions into products of univariate functions of ω and of u .

Bayesian models for Cholesky components as functions of ω were discussed in Section 3.2.2. Similarly, a low-rank approximate Bayesian smoothing spline model for a function of outcomes at the observed values can be formulated. Since the domain of the outcome values is $[0, 1]$, as opposed the domain of the frequency values $[0, 1/2]$, we consider the kernel

$$H(u_i, u_j) = \int_0^1 (u_i - v)_+ (u_j - v)_+ dv.$$

Letting $\mathbf{u} = (u_1, \dots, u_N)'$, a low-rank model for functions of the outcome evaluated at the observed values is:

$$L_H \mathbf{a} + Q_H \mathbf{b}, \quad (5)$$

where $L_H = (\mathbf{1}_N \mid \mathbf{u})$, Q_H is the $N \times n_H$ matrix of the first n_H columns of $V_H D_H^{1/2}$, $H = V_H D_H^{1/2} V_H'$ is the spectral decomposition of the $N \times N$ matrix $H = \{H(u_i, u_j)\}$, and $\mathbf{a} \sim N(\mathbf{0}, \sigma_a^2 I_2)$ is independent of $\mathbf{b} \sim N(\mathbf{0}, \tau^2 I_{n_H})$.

To write the tensor-product model at the observed frequency-outcome points, concatenate components across frequency and outcome to define the NM -vectors

$$\boldsymbol{\theta}_{k\ell} = [\{\Theta_{k\ell}(\omega_1, u_1), \dots, \Theta_{k\ell}(\omega_M, u_1)\}, \dots, \{\Theta_{k\ell}(\omega_1, u_N), \dots, \Theta_{k\ell}(\omega_M, u_N)\}]'$$

for $k > \ell = 1, \dots, P-1$. Similarly define $\log \psi_{kk}^{-1}$ for $k = 1, \dots, P$. The real and imaginary parts of $\boldsymbol{\theta}_{k\ell}$ and $\log \psi_{kk}^{-1}$ can then be expressed as tensor products of the spline models for functions of frequency (given in Equations (2) – (4)) and outcome (given in Equation (5))

$$\begin{aligned} \Re\{\boldsymbol{\theta}_{k\ell}\} &= \{L_H \otimes L_J\} \mathbf{a}_{rk\ell} + \{Q_H \otimes L_J\} \mathbf{b}_{rk\ell} + \{L_H \otimes Q_J\} \mathbf{c}_{rk\ell} + \{Q_H \otimes Q_J\} \mathbf{d}_{rk\ell} \\ \Im\{\boldsymbol{\theta}_{k\ell}\} &= \{L_H \otimes L_J\} \mathbf{a}_{ik\ell} + \{Q_H \otimes L_J\} \mathbf{b}_{ik\ell} + \{L_H \otimes Q_J\} \mathbf{c}_{ik\ell} + \{Q_H \otimes Q_J\} \mathbf{d}_{ik\ell} \\ \log \psi_{kk}^{-1} &= \{L_H \otimes L_J\} \mathbf{a}_{dkk} + \{Q_H \otimes L_J\} \mathbf{b}_{dkk} + \{L_H \otimes Q_J\} \mathbf{c}_{dkk} + \{Q_H \otimes Q_J\} \mathbf{d}_{dkk}. \end{aligned}$$

This model decomposes conditional Cholesky components into combinations of products of univariate functions of frequency and univariate functions of outcome. The parameters $\mathbf{a}$ are coefficients for functions that are products of linear functions of both ω and u , $\mathbf{b}$ are coefficients for functions that are products of linear functions of ω and nonlinear functions of u , $\mathbf{c}$ are coefficients for functions that are products of nonlinear functions of ω and linear functions of u , and $\mathbf{d}$ are coefficients for functions that are products of nonlinear functions of ω and of u .

4.3 Prior Distributions

We define two types of prior distributions: prior distributions on coefficients conditional on smoothing parameters and prior distributions on smoothing parameters. The tensor product model naturally enables the formulation of prior distributions that regularize its components as univariate functions of frequency and outcome. Letting

$\eta_{rkl} = (a'_{rkl}, b'_{rkl}, c'_{rkl}, d'_{rkl})'$, $\eta_{ikl} = (a'_{ikl}, b'_{ikl}, c'_{ikl}, d'_{ikl})'$ and $\eta_{dkk} = (a'_{dkk}, b'_{dkk}, c'_{dkk}, d'_{dkk})'$, conditional on smoothing parameters, we assume the diagonal Gaussian smoothing priors

$$\begin{aligned}\eta_{rkl} &\sim N(0, D_{rkl}) \text{ where } D_{rkl} = \text{diag}(\sigma_\alpha^2 \mathbf{1}'_4, \tau_{\beta rkl}^2 \mathbf{1}'_{2n_H}, \tau_{\gamma rkl}^2 \mathbf{1}'_{2n_J}, \tau_{\delta rkl}^2 \mathbf{1}'_{n_H \times n_J}), \\ \eta_{ikl} &\sim N(0, D_{ikl}) \text{ where } D_{ikl} = \text{diag}(\sigma_\alpha^2 \mathbf{1}'_4, \tau_{\beta ikl}^2 \mathbf{1}'_{2n_H}, \tau_{\gamma ikl}^2 \mathbf{1}'_{2n_J}, \tau_{\delta ikl}^2 \mathbf{1}'_{n_H \times n_J}), \\ \eta_{dkk} &\sim N(0, D_{dkk}) \text{ where } D_{dkk} = \text{diag}(\sigma_\alpha^2 \mathbf{1}'_4, \tau_{\beta dkk}^2 \mathbf{1}'_{2n_H}, \tau_{\gamma dkk}^2 \mathbf{1}'_{2n_J}, \tau_{\delta dkk}^2 \mathbf{1}'_{n_H \times n_J})\end{aligned}$$

where $\mathbf{1}_n$ is the n -vector of ones.

Prior distributions on the smoothing parameters are placed by assuming that $\tau_{\beta rkl}$, $\tau_{\gamma rkl}$, $\tau_{\delta rkl}$, $\tau_{\beta ikl}$, $\tau_{\gamma ikl}$, $\tau_{\delta ikl}$, $k = 1, \dots, P-1$, $\tau_{\beta dkk}$, $\tau_{\gamma dkk}$, $\tau_{\delta dkk}$, $k = 1, \dots, P$, are independent Half-t(ν , G) random variables with pdf $p(x) \propto [1 + (x/G)^2/\nu]^{-(\nu+1)/2}$, $x > 0$, where the hyperparameters ν and G are assumed known (Gelman, 2006). Computationally, it is convenient to utilize the following scale mixture representation (Wand et al., 2012): $(\tau^2 | g) \sim \text{IG}(\nu/2, \nu/g)$, $g \sim \text{IG}(1/2, 1/G^2)$, where $\text{IG}(a, b)$ is the inverse Gamma distribution with pdf $p(x) \propto x^{-(a+1)} \exp(-b/x)$, $x > 0$. The larger the value of G , the less informative the prior, and we set G to a large fixed value. We found analyses to be insensitive to the choice of G , with $G = 10$ and $G = 10^5$ giving indistinguishable results in both simulations and in the analysis of the AgeWise data. The hyperparameter σ_α^2 , which is the prior variance of the coefficients of the linear terms, is assumed to be a known large value. In our computations, $\sigma_\alpha^2 = 10^2$ and $\sigma_\alpha^2 = 10^7$ gave indistinguishable results.

4.4 Whittle Likelihood, Sample Scheme and Inference

Given observed time series, we define the DFT for the j th subject at frequency ω_m as

$$\mathbf{Y}_{jm} = n^{-1/2} \sum_{t=1}^n \mathbf{X}_{jt} \exp(-2\pi i \omega_m t).$$

For large n , conditional on u_j , $\mathbf{Y}_{jm}$ are approximately independent mean-zero complex Gaussian random variables. This provides the conditional Whittle likelihood

$$L(Y|f) \approx \prod_{j=1}^N \prod_{m=1}^M |f^{-1}(\omega_m, u_j)| \exp\{-\mathbf{Y}_{jm}^* f^{-1}(\omega_m, u_j) \mathbf{Y}_{jm}\}.$$

There are $P^2 [(n_H + 2)(n_J + 2) + 3]$ parameters in the model of f : $(n_H + 2)(n_J + 2)$ regression coefficients and 3 smoothing parameters for each of the P^2 Cholesky components. We develop a sampling scheme to sample from the joint posterior distribution of the coefficients $\boldsymbol{\eta}$'s and smoothing parameters $\boldsymbol{\tau}^2$'s conditional on the DFT Y and the observed outcomes $\mathbf{u}$. To aid in developing this sampling scheme, it is advantageous to consider a more compact notation by defining

$$Q = \begin{pmatrix} L_H \otimes L_J & \hat{\mathbf{a}}\tilde{\mathbf{T}}\mathbf{E} & Q_H \otimes L_J \hat{\mathbf{a}}\tilde{\mathbf{T}}\mathbf{E} & L_H \otimes Q_J & \hat{\mathbf{a}}\tilde{\mathbf{T}}\mathbf{E} & Q_H \otimes Q_J \end{pmatrix}$$

so that

$$\Re\{\boldsymbol{\theta}_{k\ell}\} = Q\boldsymbol{\eta}_{rk\ell}, \quad \Im\{\boldsymbol{\theta}_{k\ell}\} = Q\boldsymbol{\eta}_{ik\ell} \quad \text{and} \quad \log\psi_{kk}^{-1} = Q\boldsymbol{\eta}_{dkk}.$$

Each iteration of the sampling scheme consists of three steps. First, the coefficients corresponding to the real and imaginary components of $\boldsymbol{\Theta}$ ($\boldsymbol{\eta}_{rk\ell}$ and $\boldsymbol{\eta}_{ik\ell}$) are sequentially sampled as Gaussian random variables from their conditional posterior distributions conditional on the current values of all other parameters. In the second step, the coefficients corresponding to the diagonal elements of Ψ^{-1} ($\boldsymbol{\eta}_{dkk}$) are drawn. The log of the conditional posterior distribution of these coefficients is given by

$$\log p(\boldsymbol{\eta}_{dkk} | \mathbf{v}_k, D_{dkk}) \stackrel{c}{=} \sum_{j=1}^N \sum_{m=1}^M \left\{ \mathbf{q}'_{jm} \boldsymbol{\eta}_{dkk} - \exp(\mathbf{q}'_{jm} \boldsymbol{\eta}_{dkk}) v_{kjm} \right\} - \frac{1}{2} \boldsymbol{\eta}'_{dkk} D_{dkk}^{-1} \boldsymbol{\eta}_{dkk}, \quad (6)$$

where $\mathbf{q}'_{jm}$ is the row of Q corresponding to the j th subject and m th frequency, $\mathbf{v}_k$ is a vector with components v_{kjm} depending on Y and on other parameters held fixed (its exact form is given in Appendix B), and $\stackrel{c}{=}$ denotes equality up to a constant. Since this is not a known distribution, $\boldsymbol{\eta}_{dkk}$ are drawn in a Metropolis-Hastings step. The last step samples smoothing parameters from their posterior distributions conditional on other parameters. For ease of notation, in what follows we describe the sampling scheme for the case $P = 3$. Further details are given in Appendix B. After initializing all the parameters, the s th iteration, $1 \leq s \leq S$, of the Gibbs sampler consists of the following steps.

1. Sample the coefficients corresponding to $\boldsymbol{\Theta}$:

a. Draw

$$\begin{aligned} (\boldsymbol{\eta}_{r21}^{(s)} | Y, \psi_{11}^{-1(s-1)}, \boldsymbol{\theta}_{31}^{(s-1)}, D_{r21}^{(s-1)}) &\sim N(\boldsymbol{\mu}_{r21}, \sum_{r21}) \\ (\boldsymbol{\eta}_{i21}^{(s)} | Y, \psi_{11}^{-1(s-1)}, \boldsymbol{\theta}_{31}^{(s-1)}, D_{i21}^{(s-1)}) &\sim N(\boldsymbol{\mu}_{i21}, \sum_{i21}) \end{aligned}$$

and update $\boldsymbol{\theta}_{21}^{(s)} = Q\boldsymbol{\eta}_{r21}^{(s)} + iQ\boldsymbol{\eta}_{i21}^{(s)}$.

b. Draw

$$(\boldsymbol{\eta}_{r31}^{(s)} | Y, \boldsymbol{\psi}_{11}^{-1(s-1)}, \boldsymbol{\theta}_{21}^{(s)}, D_{r31}^{(s-1)}) \sim N(\boldsymbol{\mu}_{r31}, \boldsymbol{\Sigma}_{r31})$$

$$(\boldsymbol{\eta}_{i31}^{(s)} | Y, \boldsymbol{\psi}_{11}^{-1(s-1)}, \boldsymbol{\theta}_{21}^{(s)}, D_{i31}^{(s-1)}) \sim N(\boldsymbol{\mu}_{i31}, \boldsymbol{\Sigma}_{i31})$$

and update $\boldsymbol{\theta}_{31}^{(s)} = Q\boldsymbol{\eta}_{r31}^{(s)} + iQ\boldsymbol{\eta}_{i31}^{(s)}$.

c. Draw

$$(\boldsymbol{\eta}_{r32}^{(s)} | Y, \boldsymbol{\psi}_{22}^{-1(s-1)}, D_{r32}^{(s-1)}) \sim N(\boldsymbol{\mu}_{r32}, \boldsymbol{\Sigma}_{r32})$$

$$(\boldsymbol{\eta}_{i32}^{(s)} | Y, \boldsymbol{\psi}_{22}^{-1(s-1)}, D_{i32}^{(s-1)}) \sim N(\boldsymbol{\mu}_{i32}, \boldsymbol{\Sigma}_{i32})$$

and update $\boldsymbol{\theta}_{32}^{(s)} = Q\boldsymbol{\eta}_{r32}^{(s)} + iQ\boldsymbol{\eta}_{i32}^{(s)}$.

The exact forms of the conditional means and covariances, $\boldsymbol{\mu}_{ck\ell}$ and $\boldsymbol{\Sigma}_{ck\ell}$ $c = r, i$, are given in Appendix B.

2. Sample coefficients corresponding to Ψ^{-1} :

```

for k = 1, 2, 3 do
  (a) Draw  $\boldsymbol{\eta}_{dkk}^{(s)} \sim t_\nu(\hat{\boldsymbol{\eta}}_{dkk}, \hat{\boldsymbol{\Sigma}}_{dkk})$ , where  $\hat{\boldsymbol{\eta}}_{dkk}$  is the maximizer of (6) and  $\hat{\boldsymbol{\Sigma}}_{dkk}$  is
    the inverse of the observed information matrix at  $\hat{\boldsymbol{\eta}}_{dkk}$ .
  (b) Compute
      
$$r^{(s)} = \frac{p(\boldsymbol{\eta}_{dkk}^{(s)} | \mathbf{v}_k, D_{dkk}) f_T(\boldsymbol{\eta}_{dkk}^{(s-1)})}{p(\boldsymbol{\eta}_{dkk}^{(s-1)} | \mathbf{v}_k, D_{dkk}) f_T(\boldsymbol{\eta}_{dkk}^{(s)})},$$

      where  $f_T$  is the density of the  $t_\nu(\hat{\boldsymbol{\eta}}_{dkk}, \hat{\boldsymbol{\Sigma}}_{dkk})$  distribution.
  (c) With probability  $\min(1, r^{(s)})$  accept  $\boldsymbol{\eta}_{dkk}^{(s)}$ , otherwise  $\boldsymbol{\eta}_{dkk}^{(s)} = \boldsymbol{\eta}_{dkk}^{(s-1)}$ .
  (d) Update  $\boldsymbol{\psi}_{kk}^{-1(s)} = \exp(Q\boldsymbol{\eta}_{dkk}^{(s)})$ .
end

```

3. Sample smoothing parameters:

```

for l = 1, 2 do
  for k = 2, 3 do
    Draw
      
$$\tau_{\beta rkl}^{2(s)} \sim IG((n_b + \nu)/2, \mathbf{b}_{rkl}^{(s)} \mathbf{b}_{rkl}^{(s)}/2 + \nu/g_{\beta rkl}^{(s-1)})$$

      
$$\tau_{\gamma rkl}^{2(s)} \sim IG((n_c + \nu)/2, \mathbf{c}_{rkl}^{(s)} \mathbf{c}_{rkl}^{(s)}/2 + \nu/g_{\gamma rkl}^{(s-1)})$$

      
$$\tau_{\delta rkl}^{2(s)} \sim IG((n_d + \nu)/2, \mathbf{d}_{rkl}^{(s)} \mathbf{d}_{rkl}^{(s)}/2 + \nu/g_{\delta rkl}^{(s-1)})$$

      
$$g_{\beta rkl}^{(s)} \sim IG((\nu + 1)/2, \nu/\tau_{\beta rkl}^{2(s)} + 1/G^2)$$

      
$$g_{\gamma rkl}^{(s)} \sim IG((\nu + 1)/2, \nu/\tau_{\gamma rkl}^{2(s)} + 1/G^2)$$

      
$$g_{\delta rkl}^{(s)} \sim IG((\nu + 1)/2, \nu/\tau_{\delta rkl}^{2(s)} + 1/G^2).$$

  end
end

```

The smoothing parameters for the imaginary and diagonal components are similarly drawn from inverse gamma distributions.

Point Estimates and Credible Intervals: The sample generated via MCMC methods provides a means of obtaining point estimates and credible intervals for any function of the

spectrum averaged over the distribution of smoothing parameters through the sample mean and percentiles of the empirical distribution of the function evaluated at each iteration of the sampling algorithm. For instance, a measure of interest in the analysis of HRV is the log-spectrum from the p th period of NREM, $\log f_{pp}$. Consider

$$f^{(s)}(\omega_m, u_j) = \{\Theta^{(s)}(\omega_m, u_j) \Psi^{-1(s)}(\omega_m, u_j) \Theta^{*(s)}(\omega_m, u_j)\}^{-1}$$

as the estimated spectral matrix at the s th iteration corresponding to u_j and ω_m with p th diagonal element $f_{pp}^{(s)}(\omega_m, u_j)$. The matrix $\Theta^{(s)}(\omega_m, u_j)$ has $k\ell$ th element

$$\theta_{k\ell jm}^{(s)} = \mathbf{q}_{jm}' \boldsymbol{\eta}_{r k \ell}^{(s)} + i \mathbf{q}_{jm}' \boldsymbol{\eta}_{i k \ell}^{(s)}, \quad k > \ell = 1, \dots, P-1,$$

and $\Psi^{-1(s)}(\omega_m, u_j)$ has $k\ell$ th element $\exp(\mathbf{q}_{kk}' \boldsymbol{\eta}_{d k k}^{(s)})$. If S iterations of the sampling algorithm are run with a burn-in of S_0 , then an estimate of $\log f_{pp}(\omega_m, u_j)$ can be computed as the mean of the values $\{\log f_{pp}^{(s)}(\omega_m, u_\ell); S_0 \leq s \leq S\}$, and a 95% credible interval computed as their 2.5 and 97.5 empirical percentiles.

Scientific interest also lies in measures collapsed across frequency. For example, as will be discussed in the following section, in the analysis of HRV, collapsed power within the high-frequency band (HF) between 0.15–0.40 Hz

$$f_p^{HF}(u) = \int_{.15}^{.40} f_{pp}(\omega, u) d\omega \quad (7)$$

provides a measure of autonomic nervous system activity during the p th NREM period among people with a TIB of u . Letting $f_p^{HF(s)}(u_\ell) = W^{-1} \sum_{.15 \leq \omega_m \leq .40} f_{pp}^{(s)}(\omega_m, u_\ell)$ where W is the number of Fourier frequencies within the HF band, an estimate of $f_p^{HF}(u_\ell)$ is given by the mean of the values $\{f_p^{HF(s)}(u_\ell); S_0 \leq s \leq S\}$, and a 95% credible interval is given by their 2.5 and 97.5 empirical percentiles.

5 Application to the AgeWise Study

We used the proposed methodology to analyze the association between TIB and the power spectrum of the first three periods of NREM from $N = 108$ AgeWise subjects, as described in Section 2. The method was fit using $n_H = n_J = 10$ basis functions, with hyperparameters $G = \sigma_\alpha^2 = 10^5$, and for $S = 3000$ iterations of the MCMC algorithm with a burn-in of $S_0 = 500$. Note that, in this example, there are a total of $P^2 (n_H + 2) (n_J + 2) = 1296$ coefficients and $3P^2 = 27$ smoothing parameters. The average run time per iteration was 5.46 seconds with a standard deviation of 0.24 seconds using the program that is available on the journal's

website in Matlab 2016b and macOS Sierra v10.12.1 on a 2.9 GHz Intel Core i7 processor with 16 GB RAM.

Although all desired analyses are obtained from one MCMC chain, to aid the biological and clinical discussion of the results, we present the analysis in three stages. First, in Section 5.1 we examine the estimated period-specific spectra and squared coherences as frequency-outcome surfaces. In the subsequent two stages, we explore power collapsed within certain frequency bands as functions of TIB: first for power within periods in Section 5.2, then for coherence between periods in Section 5.3. The results from these analyses provide new insights into biological underpinnings of spending too little or too much time in bed. In particular, our analysis suggests that (i) short TIB is connected to elevated stress and arousal within-periods of NREM towards the end of the night and (ii) long TIB is associated with a persistence in arousal in the beginning of the night.

5.1 Analysis of the Conditional Spectrum

Point estimates of the within-period conditional log-spectral surfaces, $\log \{f_{pp}(\omega, u)\}$, and of the cross-period conditional logit squared coherence surfaces,

$$\text{logit} \{\rho_{pq}^2(\omega, u)\} = \log [\rho_{pq}^2(\omega, u) / \{1 - \rho_{pq}^2(\omega, u)\}],$$

are displayed in Figure 2. These estimates are plotted on the logarithmic and logistic scales, respectively, to aid visualization. The conditional spectra at each of the first three periods of NREM and the squared coherence between NREM 1 and 2 display different characteristics within low frequencies that are less than 0.15 Hz compared to higher frequencies between 0.15–0.40 Hz.

From a biological perspective, these results are not surprising and produce interpretable measures. As was discussed in Section 2, the autonomic nervous system is classically divided into two branches: the parasympathetic branch that is responsible for activities related to resting and digestion and the sympathetic branch that is responsible for the flight-or-fight response. Researchers have shown that power within the high frequency band (HF) within 0.15–0.40 Hz provides a measure of parasympathetic nervous system activity and that power within the low frequency band (LF) between 0.04–0.15 Hz is a measure of the combined modulation of both the sympathetic and parasympathetic nervous systems. Consequently, the ratio of power from low frequencies versus high frequencies (LF/HF) can be interpreted as a measure of sympathetic modulation relative to parasympathetic modulation. Blunted HF and elevated LF/HF power are often interpreted as indirect measures of physiological arousal and psychological stress (Hall et al., 2004, 2007). To obtain inference on associations between these measures and TIB, in the next two subsections we examine power and coherence collapsed within these bands as functions of TIB.

5.2 Analysis of Within-Period Power

We consider two collapsed measures of within-period power. In addition to HF previously defined in Equation (7), we also consider LF/HF as

$$f_p^{LF/HF}(u) = \left\{ \int_{.04}^{.15} f_{pp}(\omega, u) d\omega \right\} / \left\{ \int_{.15}^{.40} f_{pp}(\omega, u) d\omega \right\}.$$

Estimates and 95% pointwise credible intervals for these two measures as functions of TIB are displayed in Figure 3 for each period.

HF power is relatively constant across TIB during NREM 1, while participants with a TIB of less than 400 minutes have decreased HF power during NREM 2 and 3 compared to those who spend more time in bed. Further, those who have an exceedingly small TIB display increased LF/HF power during NREM sleep compared to those who spend more TIB, especially during NREM 3. These characteristics are indicative of heightened physiological arousal and psychological stress.

Sleeping less than 7 hours per night has been shown to be associated with a multitude of negative health effects, including increased mortality (Buysse, 2014). The results of our analysis provide a potential pathway through which short sleep, which is inherently bounded by TIB, is connected to well-being: through increased stress and arousal towards the end of the night.

5.3 Analysis of Cross-Period Coherence

To investigate connections between cross-period coherence and TIB, we consider conditional HF band-squared coherence

$$\rho_{pq}^{2,HF}(u) = \left| \int_{.15}^{.40} f_{pq}(\omega, u) d\omega \right|^2 / \{ f_p^{HF}(u) f_q^{HF}(u) \}$$

and display estimates on the logit scale, $\text{logit}(\rho_{pq}^{2,HF}) = \log[\rho_{pq}^{2,HF}/(1-\rho_{pq}^{2,HF})]$, in the top panel of Figure 4. To better understand how changes in TIB are associated with HF coherence, we also examine first derivatives,

$$D_{pq}^{HF}(u) = d[\rho_{pq}^{2,HF}(u)]/du,$$

whose estimates are displayed in the bottom panel of Figure 4. We find that the derivative of HF coherence between NREM 1 and 2 is positive for TIB greater than 500 minutes. This indicates that excessive increases in the amount of time spent in bed are associated with increased coherence in parasympathetic activity in the beginning of the night.

The relationship between excessive TIB and ill-health led Youngstedt and Kripke (2004) to propose modest sleep restrictions to increase quality of life and survival, especially for older

adults, who tend to spend more time in bed as compared to younger adults. However, these restrictions must be used with great care as they can potentially lead to negative health effects (Reynolds III et al., 2010; Reynold et al., 2014). Our results demonstrate that excessive TIB is associated with a coherence in parasympathetic activity in the beginning of the night that is not present in moderate and short TIB. A possible explanation for this relationship is that extensive TIB can cause an increase in the amount of time spent awake while in bed, or lead to fragmented sleep. The roles of and relationships between physiological activity during different sleep cycles could change as sleep becomes more fragmented. These findings provide some of the first potential insights into the biological pathway through which excessive TIB can be connected to negative health, which can potentially be used to inform optimal sleep restriction strategies in older adults.

6 Final Remarks

This article introduces a novel approach to analyzing associations between multiple time series and cross-sectional outcomes when data are observed from multiple subjects. A new measure of association, the conditional power spectrum, is introduced and its Cholesky components are modeled as bivariate functions of frequency and cross-sectional outcome. A MCMC algorithm is developed for model fitting allowing for inference on any function of the power spectrum. The method was motivated by a sleep study and uncovered connections between excessive time in bed and heightened arousal and stress that could not have been uncovered through traditional methods.

We conclude this section by discussing three extensions to the proposed methodology. First, the model is formulated to investigate the association between power spectra and a single cross-sectional variable. The model could easily be extended through higher-order tensor product models to include multiple variables, such as the amount of time it takes to fall asleep and the number of awakenings during the night. However, such a model would provide inference on the effect of these variables on the power spectrum conditional on the other variables, complicating interpretation when these variables are highly correlated. Future work will explore an interpretable canonical correlation type dimension reduction of a collection of correlated variables and multivariate spectral matrices, which can be viewed as a multivariate extension of Krafty and Hall (2013). Second, our application focused on HRV, due to the insights that it provides into autonomic nervous system activity. One could also explore the spectral analysis of other PSG channels, as well as the simultaneous coupling of channels. However, each channel of the PSG is sampled at a different rate. The second extension will develop conditional spectral analysis of time series with different sampling rates. Finally, since we were motivated by the analysis of HRV during epochs within NREM that are approximately stationary, we focused on stationary time series. For more highly sampled signals such as EEG, this assumption is not valid. A conditional time-frequency analysis for signals that are locally stationary will also be explored.

Acknowledgments

This work was supported by NIH grants R01GM113243, P01AG020677, R01HL104607, G12MD007592 and RR024153, NSF grants DMS1506882, 2G12MD007592 and DMS1512188, and NSA grant H98230-12-1-0246.

The authors thank the Editor, Associate Editor and two referees for comments and insights that greatly improved the manuscript.

References

- Burr RL. Interpretation of normalized spectral heart rate variability indices in sleep research: A critical review. *Sleep*. 2007; 30:913–919. [PubMed: 17682663]
- Buyse DJ. Sleep health: Can we define it? Does it matter? *Sleep*. 2014; 37:9–17. [PubMed: 24470692]
- Crainiceanu CM, Ruppert D, Wand MP. Bayesian analysis of penalized spline regression using WinBUGS. *Journal of Statistical Software*. 2005; 14:1–24.
- Dai M, Guo W. Multivariate spectral analysis using Cholesky decomposition. *Biometrika*. 2004; 91:629–643.
- Eubank, RL. *Nonparametric Regression and Spline Smoothing*. 2. New York: Marcel Dekker; 1999.
- Foley DJ, Ancoli-Israel S, Britz P, Walsh J. Sleep disturbances and chronic disease in older adults: Results of the 2003 National Sleep Foundation Sleep in America Survey. *Journal of Psychosomatic Research*. 2004; 56:497–502. [PubMed: 15172205]
- Foley DJ, Monjan AA, Brown SL, Simonsick EM, Wallace RB, Blazer DG. Sleep complaints among elderly persons: An epidemiologic study of three communities. *Sleep*. 1995; 18:425–432. [PubMed: 7481413]
- Gelman A. Prior distributions for variance parameters (Comment on article by Browne and Draper). *Bayesian Analysis*. 2006; 1:515–534.
- Gu C. Cross-validating non-Gaussian data. *Journal of Computational and Graphical Statistics*. 1992; 1:169–179.
- Gu, C. *Smoothing Spline ANOVA Models*. 2. New York: Springer-Verlag; 2013.
- Gu C, Kim Y-J. Penalized likelihood regression: general formulation and efficient approximation. *Canadian Journal of Statistics*. 2002; 30:619–628.
- Hall M, Thayer J, Germain A, Moul D, Vasko R, Puhl M, Miewald J, Buysse D. Psychological stress is associated with heightened physiological arousal during NREM sleep in primary insomnia. *Behavioral Sleep Medicine*. 2007; 5:178–193. [PubMed: 17680730]
- Hall M, Vasko R, Buysse D, Ombao H, Chen Q, Cashmere J, Kupfer D, Thayer J. Acute stress affects heart rate variability during sleep. *Psychosomatic Medicine*. 2004; 66:56–62. [PubMed: 14747638]
- Kimeldorf GS, Wahba G. A correspondence between Bayesian estimation on stochastic processes and smoothing by splines. *Annals of Mathematical Statistics*. 1970; 41:495–502.
- Krafty RT, Collinge WO. Penalized multivariate Whittle likelihood for power spectrum estimation. *Biometrika*. 2013; 100:447–458.
- Krafty RT, Hall M. Canonical correlation analysis between time series and static outcomes, with application to the spectral analysis of heart rate variability. *Annals of Applied Statistics*. 2013; 7:570–587. [PubMed: 24851143]
- Loader CR. Bandwidth selection: classical or plug-in? *Annals of Statistics*. 1999; 27:415–438.
- Malik M, Bigger JT, Camm AJ, Kleiger RE, Malliani A, Moss AJ, Schwartz PJ. Heart rate variability - standards of measurement, physiological interpretation, and clinical use. *Circulation*. 1996; 93:1043–1065. [PubMed: 8598068]
- Monk T, Reynolds CF, Kupfer DJ, Buysse DJ, Coble PA, Hayes AJ, MacHen MA, Petrie SR, Ritenour AM. The Pittsburgh Sleep Diary. *Journal of Sleep Research*. 1994; 3:111–120.
- Moulin P. Wavelet thresholding techniques for power spectrum estimation. *IEEE Transactions on Signal Processing*. 1994; 42:3126–3136.
- Nychka D, Cummins D. Comment on paper by Eilers and Marx. *Statistical Science*. 1996; 11:104–105.
- Pawitan Y, O'Sullivan F. Nonparametric spectral density estimation using penalized Whittle likelihood. *Journal of the American Statistical Association*. 1994; 89:600–610.
- Qin L, Wang Y. Nonparametric spectral analysis with applications to seizure characterization using EEG time series. *Annals of Applied Statistics*. 2008; 2:1432–1451.

- Reynold AM, Bowles ER, Saxena A, Fayad R, Youngstedt SD. Negative effects of time in bed extensions: A pilot study. *Journal of Sleep Medicine and Disorders*. 2014; 1:1002. [PubMed: 25374963]
- Reynolds CF III, Serody L, Okun ML, Hall MH, Houck PR, Patrick S, Maurer J, Bensasi S, Mazumdar S, Bell B, Nebes RD, Miller MD, Dew MA, Nofzinger EA. Protecting sleep, promoting health in later life: A randomized clinical trial. *Psychosomatic Medicine*. 2010; 72:178–186. [PubMed: 19995884]
- Rosen O, Stoffer D. Automatic estimation of multivariate spectra via smoothing splines. *Biometrika*. 2007; 94:335–345.
- Rosen O, Stoffer DS, Wood S. Local Spectral Analysis via a Bayesian Mixture of Smoothing Splines. *Journal of the American Statistical Association*. 2009; 104:249–262.
- Rosen O, Wood S, Stoffer D. AdaptSPEC: Adaptive Spectral Estimation for Nonstationary Time Series. *Journal of the American Statistical Association*. 2012; 107:1575–1589.
- Shumway, RH., Stoffer, DS. *Time Series Analysis and Its Applications: With R Examples*. 3. New York: Springer; 2011.
- Siegel JM. Clues to the functions of mammalian sleep. *Nature*. 2005; 437:1264–1271. [PubMed: 16251951]
- Speckman PL, Sun D. Fully Bayesian spline smoothing and intrinsic autoregressive priors. *Biometrika*. 2003; 90:289–302.
- Stoffer DS, Han S, Qin L, Guo W. Smoothing spline ANOPOW. *Journal of Statistical Planning and Inference*. 2010; 140:3789–3796.
- Wahba G. Improper priors, spline smoothing and the problem of guarding against model errors in regression. *Journal of the Royal Statistical Society, Series B*. 1978; 40:364–372.
- Wand M, Ormerod J, Padoan S, Frühworth R. Mean field variational Bayes for elaborate distributions. *Bayesian Analysis*. 2012; 7:847–900.
- Wang J-L, Chiou J-M, Müller H-G. Review of Functional Data Analysis. *Annual Review of Statistics and Its Application*. 2016; 3 In Press.
- Whittle P. Estimation and information in stationary time series. *Arkiv for Matematik*. 1953; 2:423–434.
- Whittle P. Some recent contributions to the theory of stationary processes. *A Study in the Analysis of Stationary Time Series*. 1954; 2:196–228.
- Wood SA, Jiang W, Tanner M. Bayesian mixture of splines for spatially adaptive nonparametric regression. *Biometrika*. 2002; 89:513–528.
- Wood, SN. *Generalized Additive Models: An Introduction with R*. Boca Raton: CRC Press; 2006.
- Youngstedt SD, Kripke DF. Long sleep and mortality: Rationale for sleep restriction. *Sleep Medicine Review*. 2004; 8:159–174.

Appendices

A Simulation Study

To illustrate the proposed model and to investigate its empirical properties, we consider the $P=3$ dimensional second order moving average model (MA(2))

$$\mathbf{X}_{jt} = \boldsymbol{\varepsilon}_{jt} + \Theta_1 \boldsymbol{\varepsilon}_{j,t-1} + \Theta_2 \boldsymbol{\varepsilon}_{j,t-2}, \quad j=1, \dots, N, \quad t=1, \dots, n,$$

where $\Theta_1 = -I$, $\Theta_2 = 0.6I$, I is the 3×3 identity matrix, and $\boldsymbol{\varepsilon}_{jt}$ are independent $N[\mathbf{0}, \Omega(u_j)]$ random variables with

$$\Omega(u) = \sigma^2(u) \begin{bmatrix} 1 & \rho(u) & \rho(u) \\ \rho(u) & 1 & \rho(u) \\ \rho(u) & \rho(u) & 1 \end{bmatrix},$$

$\sigma^2(u) = (2 - u)^2$ and $\rho(u) = 0.6 + 0.25 \cos(\pi u)$. Two simulated epochs of length $n = 300$, one with $u_j = 0.04$ and one with $u_j = 1$, are displayed in Figure A.1. The conditional spectrum is given by

$$f(\omega, u) = \Theta(\omega) \Omega(u) \Theta(\omega)^*,$$

where $\Theta(\omega) = I + \Theta_1 \exp(-2\pi i \omega) + \Theta_2 \exp(-4\pi i \omega)$. The log-spectra, $\log[f_{pq}(\omega, u)]$, and their estimates under the proposed procedure from a random sample of $N = 25$ independent epochs of length $n = 300$ are displayed in Figure A.2. Plots of the logit-squared coherence, $\text{logit}[\rho_{pq}^2(\omega, u)]$, and their estimates are displayed in Figure A.3. The band-collapsed measures $f_p^{HF}(u)$, $f_p^{LF/HF}(u)$ and $\rho_{pq}^{2,HF}(u)$, along with their estimates and 95% credible intervals, are displayed in Figure A.4.

Figure A.1.

Simulated conditional MA(2) epochs of length $n = 300$ with $u_j = 0.04$ and $u_j = 1$.

Figure A.2.

Conditional log-spectra from the MA(2) model (top panels) and estimated conditional log-spectra from a random sample of time series of length $n = 300$ from $N = 25$ subjects (lower panels).

We simulated 100 random samples of conditional MA(2) time series of length n from N subjects with $u_j = j/N$ for the four combinations of $n = 300, 500$ and $N = 25, 50$. The estimation procedure was run using $n_J = 10$, $n_H = 5$, and for 2000 iterations of the MCMC algorithm with burn-in of 500 iterations. To investigate the sensitivity of the proposed estimation procedure with respect to hyperparameters, the procedure was run twice: for $G = 10^5$ and for $G = 10^{10}$. Table A.1 reports the mean and standard deviation of run times per iteration for each setting using the program that is available on the journal's website in Matlab 2016b and macOS Sierra v10.12.1 on a 2.9 GHz Intel Core i7 processor with 16 GB RAM.

Figure A.3.

Conditional squared coherence from the MA(2) model (top panels) and estimated conditional squared coherence from a random sample of time series of length $n = 300$ from $N = 25$ subjects (lower panels).

Table A.1

Mean (*standard deviation*) run time per iteration of the sampling algorithm in seconds.

	$n = 300$	$n = 500$
$N = 25$	0.35 (0.02)	0.92 (0.03)
$N = 50$	1.29 (0.03)	3.40 (0.28)

To investigate the performance of the proposed procedure for conducting inference on band-collapsed measures as functions of outcome, we computed pointwise 95% credible intervals for the nine band-collapsed curves

f_1^{HF} , f_2^{HF} , f_3^{HF} , $f_1^{LF/HF}$, $f_2^{LF/HF}$, $f_3^{LF/HF}$, $\rho_{12}^{2,HF}$, $\rho_{23}^{2,HF}$, and $\rho_{13}^{2,HF}$. The mean and standard deviation of pointwise coverage probabilities integrated across u are given in Table A.2. The integrated coverage was near the nominal 95% level for each component, ranging between 94.3%–97.4%. Coverage probabilities under different tuning parameters were indistinguishable.

Figure A.4. HF band power (top panels), LF/HF band power (middle panels), and HF band squared coherence (bottom panel) (—), along with point estimates (- - -) and 95% credible intervals (- · - ·) from a random sample of $N=25$ conditional MA(2) time series of length $n=300$.

Table A.2

Mean (*standard deviation*) coverage of 95% credible intervals for band-collapsed measures from 100 random samples of N conditional MA(2) time series of length n using hyperparamter G .

n	N	G	f_1^{HF}	f_2^{HF}	f_3^{HF}	$f_1^{LF/HF}$	$f_2^{LF/HF}$	$f_3^{LF/HF}$	$\rho_{12}^{2,HF}$	$\rho_{23}^{2,HF}$	$\rho_{13}^{2,HF}$
300	25	10^5	.943	.968	.958	.969	.965	.961	.968	.951	.963
			(.112)	(.073)	(.097)	(.084)	(.078)	(.093)	(.058)	(.083)	(.078)
		10^{10}	.943	.968	.958	.969	.965	.961	.968	.951	.963
			(.112)	(.073)	(.097)	(.084)	(.078)	(.093)	(.058)	(.083)	(.078)
		10^{15}	.943	.968	.958	.969	.965	.961	.968	.951	.963
			(.112)	(.073)	(.097)	(.084)	(.078)	(.093)	(.058)	(.083)	(.078)
500	25	10^5	.968	.949	.969	.965	.974	.967	.968	.950	.963
			(.070)	(.087)	(.070)	(.097)	(.062)	(.069)	(.068)	(.086)	(.083)
		10^{10}	.968	.949	.969	.965	.974	.967	.968	.950	.963
			(.070)	(.087)	(.070)	(.097)	(.062)	(.069)	(.068)	(.086)	(.083)
		10^{15}	.968	.949	.969	.965	.974	.967	.968	.950	.963
			(.070)	(.087)	(.070)	(.097)	(.062)	(.069)	(.068)	(.086)	(.083)
300	50	10^5	.948	.961	.960	.962	.945	.945	.962	.961	.960
			(.107)	(.077)	(.073)	(.092)	(.110)	(.106)	(.082)	(.073)	(.081)
		10^{10}	.948	.961	.960	.962	.945	.945	.962	.961	.960
			(.107)	(.077)	(.073)	(.092)	(.110)	(.106)	(.082)	(.073)	(.081)
		10^{15}	.948	.961	.960	.962	.945	.945	.962	.961	.960
			(.107)	(.077)	(.073)	(.092)	(.110)	(.106)	(.082)	(.073)	(.081)

n	N	G	f_1^{HF}	f_2^{HF}	f_3^{HF}	$f_1^{LF/HF}$	$f_2^{LF/HF}$	$f_3^{LF/HF}$	$\rho_{12}^{2,HF}$	$\rho_{23}^{2,HF}$	$\rho_{13}^{2,HF}$
			(.107)	(.077)	(.073)	(.092)	(.110)	(.106)	(.082)	(.073)	(.081)
500	50	10^5	.967	.956	.971	.954	.955	.963	.959	.953	.969
			(.070)	(.079)	(.064)	(.102)	(.093)	(.078)	(.080)	(.088)	(.059)
		10^5	.967	.956	.971	.954	.955	.963	.959	.953	.969
			(.070)	(.079)	(.064)	(.102)	(.093)	(.078)	(.080)	(.088)	(.059)

To compare the performance of the proposed procedure to existing approaches, we also computed two two-stage estimators of within-period band-collapsed measures. In the first stage, periodograms were calculated and summed within HF and LF bands for each of the $P = 3$ series for each of the N subjects to obtain raw subject-specific estimates. The raw subject-specific HF and LF/HF estimates were then smoothed across u . For the first estimator, smoothing was achieved by fitting a cubic smoothing spline with smoothing parameter selected through generalized cross-validation (GCV) (Gu, 2013). For the second estimator, smoothing was achieved through local linear regression with plug-in bandwidth (Loader, 1999). We defined the integrated square error (ISE) of an estimate $\hat{f}_p^{HF}$ of f_p^{HF} as

$$\int_0^1 \left[\hat{f}_p^{HF}(u) - f_p^{HF}(u) \right]^2 du.$$

The ISEs for $\hat{f}_p^{LF/HF}$ and $\hat{\rho}_{pq}^{2,HF}$ were similarly defined. The mean and standard deviation of the ISEs are presented in Table A.3. As expected, the ISE of each estimator improved with an increase in either n or N . The insensitivity of the proposed procedure to choice of hyperparameter that was observed through indistinguishable coverage probabilities was also observed in the ISE; the ISE under $G = 10^5$ and $G = 10^{10}$ were identical up to at least three significant digits. In each setting, the proposed estimator had smaller mean ISE compared to the two-stage procedures.

Table A.3

Mean (*standard deviation*) of the integrated square error (ISE) of band-collapsed measures from 100 random samples of N independent conditional MA(2) time series of length n . Estimates were obtained using the proposed procedure with tuning parameter $G = 10^5$ (Bayes: 10^5) and $G = 10^{10}$ (Bayes: 10^{10}) and two-stage estimators using smoothing splines (2-Stage: Spline) and local linear regression (2-Stage: LOESS). Values are reported $\times 10^3$ for HF measures and $\times 10^5$ for LF/HF measures.

n	N	Estimator	f_1^{HF}	f_2^{HF}	f_3^{HF}	$f_1^{LF/HF}$	$f_2^{LF/HF}$	$f_3^{LF/HF}$
300	25	Bayes: 10^5	6.76	6.75	6.57	2.98	3.28	4.31
			(4.96)	(6.03)	(6.30)	(2.23)	(2.45)	(3.34)

<i>n</i>	<i>N</i>	Estimator	f_1^{HF}	f_2^{HF}	f_3^{HF}	$f_1^{LF/HF}$	$f_2^{LF/HF}$	$f_3^{LF/HF}$
500	25	Bayes: 10^{10}	6.76	6.75	6.57	2.98	3.28	4.31
			(4.96)	(6.03)	(6.30)	(2.23)	(2.45)	(3.34)
		2-Stage: Spline	12.11	13.14	12.82	8.75	10.06	8.92
			(12.70)	(18.56)	(12.20)	(8.29)	(8.29)	(7.00)
		2-Stage: LOESS	10.31	11.25	11.68	10.01	10.86	10.18
			(6.97)	(14.06)	(7.89)	(7.36)	(7.36)	(7.97)
	50	Bayes: 10^5	3.75	4.57	3.75	1.86	1.79	2.62
			(2.94)	(3.54)	(2.82)	(1.68)	(1.29)	(1.89)
		Bayes: 10^{10}	3.75	4.57	3.75	1.86	1.79	2.62
			(2.94)	(3.54)	(2.82)	(1.68)	(1.29)	(1.89)
		2-Stage: Spline	7.87	8.05	7.15	5.26	5.04	4.60
			(7.76)	(8.63)	(7.27)	(4.35)	(4.80)	(3.44)
300	50	2-Stage: LOESS	7.93	7.69	7.21	5.94	5.34	5.42
			(6.32)	(5.29)	(5.61)	(3.84)	(3.50)	(2.87)
		Bayes: 10^5	3.33	3.61	3.59	1.51	1.83	2.51
			(2.69)	(2.84)	(2.97)	(1.37)	(1.39)	(1.81)
		Bayes: 10^{10}	3.33	3.61	3.59	1.51	1.83	2.51
			(2.69)	(2.84)	(2.97)	(1.37)	(1.39)	(1.81)
	50	2-Stage: Spline	6.64	7.09	6.84	4.88	5.39	5.57
			(7.70)	(10.34)	(6.63)	(4.15)	(4.41)	(4.35)
		2-Stage: LOESS	5.50	5.60	5.91	5.31	5.62	5.79
			(4.00)	(6.35)	(4.21)	(3.64)	(3.48)	(3.85)
	500	Bayes: 10^5	1.74	2.12	2.06	1.00	1.13	1.44
			(1.25)	(1.50)	(1.36)	(0.83)	(1.13)	(0.92)
		Bayes: 10^{10}	1.74	2.12	2.06	1.00	1.13	1.44
			(1.25)	(1.50)	(1.36)	(0.83)	(1.13)	(0.92)
500	50	2-Stage: Spline	4.52	4.87	3.84	2.87	3.19	2.84
			(5.00)	(5.31)	(4.31)	(2.56)	(3.07)	(2.32)
		2-Stage: LOESS	3.85	3.89	3.47	2.81	3.04	2.79
			(3.13)	(2.89)	(2.41)	(1.92)	(2.18)	(1.57)

B Details of the Sampling Scheme

In this appendix we provide more details about the sampling scheme outlined in Section 4.4, assuming $P=3$. As in Section 4.4, q'_{jm} is the row of Q corresponding to u_j and ω_m . The DFT of the p th series from the j th subject at Fourier frequency ω_m is denoted by Y_{pjm} . The $k\ell$ element of $\Theta(\omega_m, u_j)$ defined in Section 4.2 is expressed as

$$\theta_{k\ell jm} = q'_{jm} \eta_{r k \ell} + i q'_{jm} \eta_{i k \ell}, \quad k > \ell = 1, \dots, P-1, \quad (\text{B.1})$$

where the i in the second term on the right-hand side of (B.1) is the unit imaginary number.

The diagonal elements of $\Psi^{-1}(\omega, u)$ are expressed as $\psi_{kkjm}^{-1} = \exp(\mathbf{q}_{jm}' \boldsymbol{\eta}_{dkk})$, $k = 1, \dots, P$. To aid presentation and simplify notation, the superscript for iteration number is suppressed and all derived distributions are conditional on the current values of all other parameters.

Drawing the Basis Function Coefficient Vectors

The conditional posterior distribution of $\boldsymbol{\eta}_{ck\ell}$ $c = r, i$, $k > \ell = 1, \dots, P-1$, is multivariate normal, $N(\boldsymbol{\mu}_{ck\ell}, \Sigma_{ck\ell})$. In what follows we provide expressions for $\boldsymbol{\mu}_{ck\ell}$ and $\Sigma_{ck\ell}$

Mean vectors and covariance matrices for $\boldsymbol{\eta}_{i21}$ and $\boldsymbol{\eta}_{r21}$

$$\begin{aligned}\sum_{c21}^{-1} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{11jm}^{-1} |Y_{2jm}|^2 \mathbf{q}_{jm} \mathbf{q}_{jm}' + D_{c21}^{-1}, \quad c=r, i \\ \sum_{r21}^{-1} \boldsymbol{\mu}_{r21} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{11jm}^{-1} \Re\{Y_{1jm} Y_{2jm}^* - \theta_{31jm}^* Y_{2jm} Y_{3jm}^*\} \mathbf{q}_{jm} \\ \sum_{i21}^{-1} \boldsymbol{\mu}_{i21} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{11jm}^{-1} \Im\{Y_{1jm} Y_{2jm}^* + \theta_{31jm}^* Y_{2jm} Y_{3jm}^*\} \mathbf{q}_{jm}.\end{aligned}$$

Note that ψ_{11jm}^{-1} and θ_{31jm} are evaluated at their current values.

Mean vectors and covariance matrices for $\boldsymbol{\eta}_{i31}$ and $\boldsymbol{\eta}_{r31}$

$$\begin{aligned}\sum_{c31}^{-1} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{11jm}^{-1} |Y_{3jm}|^2 \mathbf{q}_{jm} \mathbf{q}_{jm}' + D_{c31}^{-1}, \quad c=r, i \\ \sum_{r31}^{-1} \boldsymbol{\mu}_{r31} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{11jm}^{-1} \Re\{Y_{1jm} Y_{3jm}^* - \theta_{21jm}^* Y_{2jm}^* Y_{3jm}\} \mathbf{q}_{jm} \\ \sum_{i31}^{-1} \boldsymbol{\mu}_{i31} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{11jm}^{-1} \Im\{Y_{1jm} Y_{3jm}^* + \theta_{21jm}^* Y_{2jm}^* Y_{3jm}\} \mathbf{q}_{jm}.\end{aligned}$$

Mean vectors and covariance matrices for $\boldsymbol{\eta}_{i32}$ and $\boldsymbol{\eta}_{r32}$

$$\begin{aligned}\sum_{c32}^{-1} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{22jm}^{-1} |Y_{3jm}|^2 \mathbf{q}_{jm} \mathbf{q}_{jm}' + D_{c32}^{-1}, \quad c=r, i \\ \sum_{r32}^{-1} \boldsymbol{\mu}_{r32} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{22jm}^{-1} \Re\{Y_{2jm}^* Y_{3jm}\} \mathbf{q}_{jm} \\ \sum_{i32}^{-1} \boldsymbol{\mu}_{i32} &= 2 \sum_{j=1}^N \sum_{m=1}^M \psi_{22jm}^{-1} \Im\{Y_{2jm} Y_{3jm}^*\} \mathbf{q}_{jm}.\end{aligned}$$

The basis function coefficient vectors $\boldsymbol{\eta}_{dkk}$, $k = 1, \dots, P$, are drawn from $p(\boldsymbol{\eta}_{dkk} / Q, \mathbf{v}_k, D_{dkk})$, given in Section 4.3. The entries v_{kjm} of $\mathbf{v}_k$ for $k = 1, 2, 3$ are as follows.

$$\begin{aligned} v_{1jm} &= |Y_{1jm}|^2 + |\theta_{21jm} Y_{2jm}|^2 + |\theta_{31jm} Y_{3jm}|^2 \\ &\quad - 2\Re\{\theta_{21jm} Y_{1jm}^* Y_{2jm} + \theta_{31jm} Y_{1jm}^* Y_{3jm} - \theta_{21jm}^* \theta_{31jm} Y_{2jm}^* Y_{3jm}\}. \\ v_{2jm} &= |Y_{2jm}|^2 + |\theta_{32jm} Y_{3jm}|^2 - 2\Re\{\theta_{32jm} Y_{2jm}^* Y_{3jm}\}. \\ v_{3jm} &= |Y_{3jm}|^2. \end{aligned}$$

The vectors $\boldsymbol{\eta}_{dkk}$, $k = 1, \dots, P$, are generated independently via a Metropolis-Hastings step with a multivariate t proposal distribution, $t(\hat{\boldsymbol{\eta}}_{dkk}, \hat{\Sigma}_{dkk})$, where

$$\hat{\boldsymbol{\eta}}_{dkk} = \arg \max_{\boldsymbol{\eta}_{dkk}} \log p(\boldsymbol{\eta}_{dkk} | Q, \mathbf{v}_k, D_{dkk})$$

and

$$\hat{\Sigma}_{dkk} = \left[-\frac{\partial^2}{\partial \boldsymbol{\eta}_{dkk} \partial \boldsymbol{\eta}_{dkk}'} \log p(\boldsymbol{\eta}_{dkk} | Q, \mathbf{v}_k, D_{dkk}) \right]_{\boldsymbol{\eta}_{dkk} = \hat{\boldsymbol{\eta}}_{dkk}}^{-1}.$$

The gradient and Hessian of $\log p(\boldsymbol{\eta}_{dkk} / Q, \mathbf{v}_k, D_{dkk})$ are given by

$$\sum_{j=1}^N \sum_{m=1}^M \left[1 - v_{kjm} \exp(\mathbf{q}_{jm}' \boldsymbol{\eta}_{dkk}) \right] \mathbf{q}_{jm} - D_{dkk}^{-1} \boldsymbol{\eta}_{dkk}$$

and

$$-\sum_{j=1}^N \sum_{m=1}^M v_{kjm} \exp(\mathbf{q}_{jm}' \boldsymbol{\eta}_{dkk}) \mathbf{q}_{jm} \mathbf{q}_{jm}' - D_{dkk}^{-1},$$

respectively.

Drawing the Smoothing Parameters

Details are given below for the smoothing parameters associated with the real part of $\theta_{k\ell}(\omega, u)$. The details for the rest of the smoothing parameters are similar. The smoothing parameters $\tau_{\beta_{rk\ell}}^2$, $\tau_{\gamma_{rk\ell}}^2$ and $\tau_{\delta_{rk\ell}}^2$ are drawn independently for $k > \ell = 1, \dots, P-1$, as follows.

$$\begin{aligned}\tau_{\beta rkl}^2 &\stackrel{\text{ind}}{\sim} IG((n_b + \nu)/2, \mathbf{b}'_{rkl} \mathbf{b}_{rkl}/2 + \nu/g_{\beta rkl}) \\ \tau_{\gamma rkl}^2 &\stackrel{\text{ind}}{\sim} IG((n_c + \nu)/2, \mathbf{c}'_{rkl} \mathbf{c}_{rkl}/2 + \nu/g_{\gamma rkl}) \\ \tau_{\delta rkl}^2 &\stackrel{\text{ind}}{\sim} IG((n_d + \nu)/2, \mathbf{d}'_{rkl} \mathbf{d}_{rkl}/2 + \nu/g_{\delta rkl}),\end{aligned}$$

where

$$\begin{aligned}g_{\beta rkl} &\stackrel{\text{ind}}{\sim} IG((\nu + 1)/2, \nu/\tau_{\beta rkl}^2 + 1/G^2) \\ g_{\gamma rkl} &\stackrel{\text{ind}}{\sim} IG((\nu + 1)/2, \nu/\tau_{\gamma rkl}^2 + 1/G^2) \\ g_{\delta rkl} &\stackrel{\text{ind}}{\sim} IG((\nu + 1)/2, \nu/\tau_{\delta rkl}^2 + 1/G^2).\end{aligned}$$

Figure 1.
Detrended HRV time series during the first three periods of NREM from two subjects.
Subject 1 reported a TIB of 357.67 minutes and subject 2 reported a TIB of 521.00 minutes.

Figure 2.

Estimated conditional log-spectra for each period of NREM (top panel) and estimated logit of conditional coherence between each period (bottom panel).

Figure 3.

Estimated conditional HF (top panel), $\hat{f}_p^{HF}$, and LF/HF (bottom panel), $\hat{f}_p^{LF/HF}$, as functions of TIB with 95% pointwise credible intervals for each period of NREM.

Figure 4.

Estimated logit of integrated HF coherence (top panel), $\text{logit}(\hat{\rho}_{pq}^{2, HF})$, between each NREM period as functions of TIB and their first derivatives (bottom panel), $\hat{D}_{pq}^{HF}$, with pointwise 95% credible intervals.