Estimating fMRI Timescale Maps

Gabriel Riegner [1*], Samuel Davenport [2], Bradley Voytek [1,3,4], Armin Schwartzman [1,2]

[1]Halicioğlu Data Science Institute, University of California San Diego

[2]Division of Biostatistics, University of California San Diego

[3]Department of Cognitive Science, University of California San Diego

[4]Neurosciences Graduate Program, University of California San Diego

[*]Correspondence: gariegner@ucsd.edu

Abstract

Brain activity unfolds over hierarchical timescales that reflect how brain regions integrate and process information, linking functional and structural organization. While timescale studies are prevalent, existing estimation methods rely on the restrictive assumption of exponentially decaying temporal autocorrelation and only provide point estimates without standard errors, limiting statistical inference. In this paper, we formalize and evaluate two methods for mapping timescales in resting-state fMRI: a time-domain fit of an autoregressive (AR1) model and an autocorrelation-domain fit of an exponential decay model. Rather than assuming exponential autocorrelation decay, we define timescales by projecting the fMRI time series onto these approximating models, requiring only stationarity and mixing conditions while incorporating robust standard errors to account for model misspecification. We introduce theoretical properties of timescale estimators and show parameter recovery in realistic simulations, as well as applications to fMRI from the Human Connectome Project. Comparatively, the time-domain method produces more accurate estimates under model misspecification, remains computationally efficient for high-dimensional fMRI data, and yields maps aligned with known functional brain organization. In this work, we show valid statistical inference on fMRI timescale maps, and provide Python implementations of all methods.

Keywords: time-domain linear model, autocorrelation-domain nonlinear model, uncertainty quantification, statistical inference, human connectome project, functional brain organization

1 Introduction

1.1 fMRI Timescale Maps

Neural processes span multiple timescales, from millisecond synaptic events to slower activity coordinating distributed brain networks (Buzsáki, 2004). Timescales index the integration of neural activity over time, defined by exponential decay in autocorrelation that emerges from stochastic models of neural activity (Murray et al., 2014, Supplementary Mathematical Note). When mapped across the cortex, these local temporal timescales reveal a large-scale spatial organization. Multimodal evidence links these timescale differences to intrinsic brain hierarchies, reflecting how distinct regions integrate and process information over time. Timescale maps of the brain align with functional hierarchy – sensory areas that process rapidly changing stimuli show shorter timescales than association areas involved in cognitive processes that unfold over longer durations (Raut et al., 2020; Gao et al., 2020; Hasson et al., 2008; Murray et al., 2014; Stephens et al., 2013). This hierarchy is also associated with anatomical organization including myelination levels and gene expression patterns, as shown in studies using human electrophysiology, MEG, and gene expression profiling (Gao et al., 2020; Shafiei et al., 2023).

Computational modeling by Li and Wang (2022) suggests that hierarchical timescales emerge from: (i) brain-wide gradients in synaptic excitation strength, (ii) electrophysiological differences between excitatory and inhibitory neurons, and (iii) balance between distant excitatory and local inhibitory inputs. In addition to these intrinsic mechanisms, there is growing evidence that neuronal timescales are dynamic and modulated by experimental manipulations or behavioral demands. For example, pharmacological agents like propofol and serotonergic drugs alter intrinsic timescales, affecting the temporal integration of information in the brain (Huang et al., 2018; Shinn et al., 2023). Timescale changes have also been observed during development, sleep deprivation, wakefulness, neuropsychiatric disorders (autism and schizophrenia), and naturalistic behaviors (Martin-Burgos et al., 2024; Meisel et al., 2017; Watanabe et al., 2019; Wengler et al., 2020; Manea et al., 2024). These findings demonstrate that timescales are broadly relevant to both structural and functional properties of the brain.

Seminal research on timescales has primarily used invasive electrophysiology in non-human animals (Murray et al., 2014; Cirillo et al., 2018; Nougaret et al., 2021; Manea et al., 2022; Spitmaan et al., 2020; Trepka et al., 2024). While these methods provide high temporal resolution for studying neural activity up to the single-neuron level, they are limited by sparse spatial sampling. Investigating the large-scale spatial organization of timescale maps requires non-invasive methods like resting-state functional MRI (rfMRI), which measures spontaneous fluctuations in the blood oxygen level-dependent (BOLD) signal. Unlike techniques with sparser spatial coverage such as EEG or MEG, rfMRI provides dense, whole-brain coverage, though its power is concentrated in the much slower (<0.1 Hz) frequency range (Raut et al., 2020; He, 2011). This is because the BOLD signal reflects hemodynamic changes, which act as a low-pass filter on neural activity, and because the signal-to-noise ratio is highest in this infraslow frequency band. Although an indirect measure, the BOLD signal reflects functionally significant fluctuations associated with underlying electrophysiological activity (Logothetis, 2008), making it a valuable tool for investigating high spatial-resolution cortical timescale maps. Studies have shown that rfMRI-derived timescale maps align spatially with those from other imaging modalities across human and animal models (Raut et al., 2020; Shafiei et al., 2020; Lurie et al., 2024). The present study will focus on rfMRI from the Human Connectome Project dataset (Van Essen et al., 2013).

1.2 Current methods

A key consideration is the representation of the time series – whether in its original time domain, or transformed into another domain for analysis. Timescales are generally estimated using three main methods corresponding to neural data represented in the (i) time domain, (ii) autocorrelation domain,

or (iii) frequency domain. The most common is the autocorrelation domain, where timescales are defined by fitting an exponential decay model to the sample autocorrelation function (ACF) (Rossi-Pool et al., 2021; Cirillo et al., 2018; Ito et al., 2020; Runyan et al., 2017; Zeraati et al., 2022; Nougaret et al., 2021; Wasmuht et al., 2018; Müller et al., 2020; Maisson et al., 2021; Li and Wang, 2022; Shafiei et al., 2020). Similar approaches compute timescales directly from the sample ACF as the sum of positive autocorrelations (Wengler et al., 2020; Manea et al., 2022; Watanabe et al., 2019), or by identifying where the sample ACF crosses a specified threshold (Wengler et al., 2020; Zilio et al., 2021). Alternatively, the time-domain method uses a first-order autoregressive (AR1) model to estimate timescales directly from time series (Kaneoke et al., 2012; Meisel et al., 2017; Huang et al., 2018; Lurie et al., 2024; Shinn et al., 2023; Shafiei et al., 2020; Spitmaan et al., 2020; Trepka et al., 2024), and has shown better test-retest reliability than autocorrelation-domain methods for rfMRI (Huang et al., 2018). Finally, frequency-domain methods operate on the premise that timescales are properties of the brain's aperiodic (non-rhythmic) activity. These methods therefore use the power-spectral density to estimate and remove oscillatory components to mitigate their influence on timescale estimates (Donoghue et al., 2020; Gao et al., 2020; Manea et al., 2024; Zeraati et al., 2022; Fallon et al., 2020). This paper focuses on time- and autocorrelation-domain methods, avoiding external criteria that attempt to separate periodic from aperiodic contributions to the BOLD signal. Because amplitude-modulated rhythmic activity can mimic aperiodic properties (He et al., 2010; He, 2011), distinguishing the true generative process is often circular. Instead, we introduce a timescale estimator that is robust regardless of whether the underlying process is strictly aperiodic, damped-oscillatory, or a mixture of both, and we validate this robustness using simulations across these regimes.

1.3 Problem Statement and Proposed Solution

A key challenge in applied timescale research is the lack of standardized model definitions, where diverse approaches have led to inconsistent findings across studies (Zeraati et al., 2022; Fallon et al., 2020; Shafiei et al., 2020). Many parameterization methods rely on restrictive assumptions, such as exponential autocorrelation decay, which may bias timescale estimates and (more often) their standard errors (Woolrich et al., 2001; Zeraati et al., 2022; Raut et al., 2019). Additionally, the distributional properties of these methods are often ignored, resulting in studies reporting only point estimates without quantifying uncertainty, which limits statistical inference and hypothesis testing.

To address these issues, this paper formalizes and evaluates two commonly applied timescale methods in rfMRI: the time-domain fit of an autoregressive (AR1) model and the autocorrelation-domain fit of an exponential decay model. The goal is to estimate accurate timescale maps that enable robust statistical testing and inference across brain regions. This work offers the following contributions: (i) The assumptions are generalized to include processes that are both stationary *and* mixing, not only those with exponential autocorrelation decay. (ii) Robust standard errors account for the inevitable mismatch between the data-generating process and fitted model, enabling valid inference despite model misspecification. (iii) Theoretical properties demonstrate that both time- and autocorrelation-domain estimators converge to different values due to their distinct definitions, and are consistent (i.e., converge in probability to the method-specific estimand) and asymptotically normal. (iv) Simulations confirm that both methods yield unbiased estimates across autoregressive and realistic settings, with standard errors that are robust to non-exponential autocorrelation decay. (v) Empirical analysis of rfMRI from the Human Connectome Project shows that both approaches yield similar t-ratio maps (timescales relative to their standard errors), revealing a hierarchical organization of timescales across the cortex. While this hierarchy aligns with prior point estimate maps, our approach improves interpretability by accounting for uncertainty, ensuring that observed patterns are not artifacts of sampling noise or model misspecification. (vi) Comparative Monte Carlo simulations and empirical data analyses show that the time-domain method performs as well as, and often better than, the autocorrelation-domain method, with improved computational efficiency for high-

dimensional fMRI data.

The proposed methods address important limitations in fMRI timescale research by providing rigorous statistical methods that move beyond point estimates to incorporate uncertainty quantification. This work establishes a methodological foundation for reliable inference, organized as follows. First, Section 2 formalizes the statistical models and their properties. Numbered equations indicate important expressions that are referenced throughout, while the others provide intermediate steps that can be skipped without loss of continuity. These models are validated against varying degrees of realistic misspecification in Section 3, followed by a demonstration of their empirical utility in Section 4 using the Human Connectome Project dataset. Finally, Section 5 addresses practical considerations for the interpretation of these estimates in the context of fMRI — detailing the influence of hemodynamics, measurement noise, and periodicity — to support future research investigating the functional and structural organization of timescales in the brain.

2 Methods

There are varying definitions and conventions in the study of time series data and stochastic processes more generally. To simplify our presentation, this work closely follows the Hansen (2022) *Econometrics* textbook as a primary reference. Throughout, we will frequently refer the reader to specific results from this textbook.

2.1 Assumptions

Let $\{X_t\}_{t \in \{\ldots, -2, -1, 0, 1, 2, \ldots\}}$ be a discrete-time stochastic process that is Weakly stationary and Strong Mixing (abbreviated as "stationary" and "mixing" throughout), and let $x = \{x_1, x_2, \ldots, x_T\}$ where x_t denotes an observed sample of X_t. For simplicity, assume $\mathbb{E}[X_t] = 0$.

2.1.1 Weakly stationary

Weak stationarity implies a constant mean and variance (independent of time index t), and an autocovariance function that only depends on time lag k, $\gamma_k = \text{cov}[X_t, X_{t-k}] = \mathbb{E}[X_t X_{t-k}]$. For analysis, we use a normalized measure of the autocovariances, the *autocorrelation function (ACF)*:

$$\rho_k = \text{corr}(X_t, X_{t-k}) = \gamma_k / \gamma_0, \tag{1}$$

where γ_k is the autocovariance at lag k and γ_0 is the variance. We emphasize that ρ_k represents *temporal* autocorrelation, not spatial autocorrelation often referenced in the fMRI literature. As a correlation, $|\rho_k| \leq 1$, but stationarity alone does not guarantee decay of ρ_k with increasing time lag; constant or periodic processes can maintain nonzero correlations indefinitely.

2.1.2 Strong Mixing

Strong mixing (α-mixing) imposes stronger dependence constraints than stationarity while still allowing for a wide set of stochastic processes. By definition, a process is strong mixing if $\alpha(\ell) \to 0$ as $\ell \to \infty$, where $\alpha(\ell)$ measures the *maximum dependence* between events separated by ℓ time points. This controls the full joint dependence structure and implies distant events are uncorrelated. Strong mixing also implies ergodicity (Hansen, 2022, Chapter 14.12), which ensures consistent estimation by the ergodic theorem (Hansen, 2022, Theorem 14.9). Additionally, if a mixing process has $r > 2$ finite moments $\mathbb{E}|X_t|^r < \infty$ and its mixing coefficients satisfy $\sum_{\ell=0}^{\infty} \alpha(\ell)^{1-2/r} < \infty$, then its autocorrelations decay sufficiently fast for application of the central limit theorem for dependent data (Hansen, 2022,

Theorem 14.15). These conditions enable asymptotic theory for estimation and inference on timescale maps (see Estimator Properties).

As introduced by Murray et al. (2014), the timescale τ represents the lag where exponentially decaying autocorrelations reach $1/e \approx 0.37$, analogous to the time constants of many physical systems, and often referred to as e-folding time in other fields. While it provides an intuitive description of the memory or persistence of that process, assuming an exponential function imposes stricter constraints than strong mixing, which alone does not prescribe any specific type of decay (exponential, linear, damped periodic, etc.). Despite this, the timescale serves as a useful summary statistic for the duration of temporal dependence, enabling direct comparisons across brain regions even when the true decay structure is more complex. This highlights an important distinction between the data-generating process and the simplified parametric model used to describe the timescale at which such a process becomes decorrelated. In the present paper, we adopt broad assumptions, requiring only that the process is stationary and mixing, to account for cases where the ACF decay may not be strictly exponential.

2.1.3 Estimation and Inference under Misspecification

We approximate the dominant exponential decay in autocorrelations by a single timescale parameter τ, and formally evaluate two timescale methods that are commonly applied across neuroimaging modalities (fMRI, EEG, ECoG, MEG). The time-domain (TD) linear model estimated with linear least squares (Kaneoke et al., 2012; Meisel et al., 2017; Huang et al., 2018; Lurie et al., 2024; Shinn et al., 2023; Shafiei et al., 2020; Spitmaan et al., 2020; Trepka et al., 2024), and the autocorrelation-domain (AD) nonlinear model estimated with nonlinear least squares (Rossi-Pool et al., 2021; Cirillo et al., 2018; Ito et al., 2020; Runyan et al., 2017; Zeraati et al., 2022; Nougaret et al., 2021; Wasmuht et al., 2018; Müller et al., 2020; Maisson et al., 2021; Li and Wang, 2022; Shafiei et al., 2020). Both belong to the subclass of M-estimators defined by minimizing a sample squared error loss function. This approach seeks the best approximation in terms of squared error and does not require full specification of the data-generating distribution, relying on method of moments estimation.

Throughout, a superscript star (e.g. τ^*) denotes the pseudo-true parameter corresponding to the best fitting parameter in a possibly misspecified model. If correctly specified, τ^* is the true timescale; otherwise, it is an approximation. Acknowledging that the data-generating process and fitted model will likely be different in practice, we introduce standard errors that account for this mismatch, enabling valid inference despite misspecification of the autocorrelation decay. Under the above assumptions and relevant moment conditions, the M-estimators are consistent for their respective pseudo-true parameters and asymptotically normal (see Estimator Properties). Further, they are invariant under the reparameterization needed to define timescales, with asymptotic normality preserved by the delta method, ensuring valid inference for transformed parameters.

2.2 Timescale Definitions

2.2.1 Time-Domain Linear Model

A first-order autoregressive model (AR1) provides a linear approximation of the timescale. The AR1 model:

$$X_t = \phi_{\text{TD}} X_{t-1} + e_t, \tag{2}$$

defines the process as a linear regression between X_t and X_{t-1} with *iid* innovations e_t with variance $\sigma_{\text{TD}}^2 = \mathbb{E}[e_t^2]$. In the autocorrelation domain, it implies that the theoretical ACF decays exponentially at a rate determined by ϕ_{TD}, such that $\rho_k = \phi_{\text{TD}}^k$ (Hansen, 2022, Chapter 14.22). For a stationary process

with $|\phi_{\text{TD}}| < 1$, the exponential decay rate can be directly obtained from ϕ_{TD}, with a timescale τ_{TD} equal to the lag at which the AR1-projected ACF reaches $1/e \approx 0.37$, that is, $\rho_\tau = \phi_{\text{TD}}^\tau = 1/e$, resulting in $\tau_{\text{TD}} = g(\phi_{\text{TD}}) = -1/\log(|\phi_{\text{TD}}|)$. Here, the timescale τ_{TD} is expressed as a nonlinear function of ϕ_{TD}, denoted by $g(\phi_{\text{TD}})$. This defines τ_{TD} to be a real number even though the ACF only includes integer indices. The absolute value allows for $\phi_{\text{TD}} < 0$ and is introduced as a modeling assumption to ensure the timescale captures the monotonic decay rate of autocorrelations, consistent with the approach of Murray et al. (2014).

The parameter ϕ_{TD}^* is the value that minimizes the expected squared error function $S(\phi_{\text{TD}})$:

$$S(\phi_{\text{TD}}) = \mathbb{E}[(X_t - \phi_{\text{TD}}X_{t-1})^2], \quad \phi_{\text{TD}}^* = \underset{\phi_{\text{TD}}}{\text{argmin}}\; S(\phi_{\text{TD}}).$$

$S(\phi_{\text{TD}})$ is minimized by taking its derivative with respect to ϕ_{TD}, setting it to zero, and solving for ϕ_{TD}^*:

$$\frac{d}{d\phi_{\text{TD}}}S(\phi_{\text{TD}}) = -2\mathbb{E}[X_{t-1}(X_t - \phi_{\text{TD}}X_{t-1})] = 0.$$

Differentiating the quadratic function yields a linear equation in ϕ_{TD}, and solving this results in a closed-form expression for the optimal ϕ_{TD}^*. Therefore, ϕ_{TD}^* is defined by *linear projection* and the timescale parameter τ_{TD}^* by a change of variable:

$$\phi_{\text{TD}}^* = (\mathbb{E}[X_{t-1}^2])^{-1}(\mathbb{E}[X_t X_{t-1}]), \tag{3}$$

$$\tau_{\text{TD}}^* = g(\phi_{\text{TD}}^*) = -\frac{1}{\log(|\phi_{\text{TD}}^*|)}. \tag{4}$$

Thus, the timescale represents the e-folding time derived from the best-fitting AR1 model under squared error loss. Since X_t is stationary with finite variance, the parameters ϕ_{TD}^* and τ_{TD}^* defined by projection are unique; in fact, any approximating AR1 model is identifiable if $\mathbb{E}[X_{t-1}^2]$ is nonzero and finite (Hansen, 2022, Theorem 14.28). Importantly, the observed process X_t need not follow the AR1 model from equation (2) with *iid* innovations e_t. This allows for projection errors, i.e., residuals from linear projection onto the working model:

$$e_t^* = X_t - \phi_{\text{TD}}^* X_{t-1},$$

which may exhibit unequal variance and autocorrelation. This applies to any stationary and mixing process, even if the true data-generating process is not AR1, so the resulting fit is an AR1 projection.

2.2.2 Autocorrelation-Domain Nonlinear Model

Alternatively, timescales can be defined in the autocorrelation domain by an exponential decay function, as introduced by Murray et al. (2014). We formalize this by projecting the K-lagged sequence of autocorrelations $\rho_1, \ldots, \rho_K$ from (1) onto the exponential approximation $\rho_k = \phi_{\text{AD}}^k$ with decay parameter ϕ_{AD}, minimizing the sum of squared errors averaged over fixed K:

$$S_K(\phi_{\text{AD}}) = \frac{1}{K}\sum_{k=1}^{K}\left(\rho_k - \phi_{\text{AD}}^k\right)^2, \tag{5}$$

$$\phi_{\text{AD}}^* = \underset{\phi_{\text{AD}}}{\text{argmin}}\; S_K(\phi_{\text{AD}}). \tag{6}$$

Unlike the Time-Domain Linear Model, which effectively projects using only a single ($K = 1$) lag of the ACF, this definition captures exponential decay across a sequence of multiple (K) lags. Consequently, the relationship between the pseudo-true parameters depends on the process specification:

$\phi_{\text{AD}}^* = \phi_{\text{TD}}^*$ if correctly specified (AR1), whereas $\phi_{\text{AD}}^* \neq \phi_{\text{TD}}^*$ if misspecified. Crucially, because ρ_k and ϕ_{AD}^* are deterministic population parameters rather than random variables, the projection residuals $\rho_k - \phi_{\text{AD}}^{*k}$ are fixed quantities distinct from the stochastic errors in the time domain. Under correct specification, these residuals vanish to zero; under misspecification, they quantify the systematic discrepancy between the complex true autocorrelation structure (e.g., multi-exponential or damped periodic) and the projected single-parameter exponential decay.

$S_K(\phi_{\text{AD}})$ is minimized by taking its derivative with respect to ϕ_{AD}, setting it to zero, and solving for ϕ_{AD}:

$$\frac{d}{d\phi_{\text{AD}}} S_K(\phi_{\text{AD}}) = -\frac{2}{K} \sum_{k=1}^{K} k\phi_{\text{AD}}^{k-1}(\rho_k - \phi_{\text{AD}}^k) = 0.$$

However, this first-order condition sums nonlinear terms across K lags, resulting in a high-order equation with no general closed-form solution. As a result, the pseudo-true parameter ϕ_{AD}^* is defined implicitly as the solution to this nonlinear equation and computed using numerical optimization. As in the time-domain case, the corresponding timescale is defined as the e-folding time, expressed by the change of variable:

$$\tau_{\text{AD}}^* = g(\phi_{\text{AD}}^*) = -\frac{1}{\log(|\phi_{\text{AD}}^*|)}. \tag{7}$$

2.3 Timescale Estimation

2.3.1 Time-Domain Linear Estimator

Given observations $x_1, \ldots, x_T$ are demeaned, the linear least squares estimator of the Time-Domain Linear Model is obtained by replacing the expectations in equation (3) using method of moments. It has the following closed-form expression:

$$\hat{\phi}_{\text{TD}}^* = \left(\sum_{t=2}^{T+1} x_{t-1}^2 \right)^{-1} \left(\sum_{t=2}^{T} x_t x_{t-1} \right), \tag{8}$$

$$\hat{\tau}_{\text{TD}}^* = g(\hat{\phi}_{\text{TD}}^*) = -\frac{1}{\log(|\hat{\phi}_{\text{TD}}^*|)}, \tag{9}$$

where (8) and (9) are the sample versions of the population parameters from equations (3) and (4), respectively (Hansen, 2022, Chapter 14.3).

2.3.2 Autocorrelation-Domain Nonlinear Estimator

The nonlinear least squares estimator of the Autocorrelation-Domain Nonlinear Model is fit to the ACF, so the time series needs to be first transformed into the autocorrelation domain. For a finite-length and centered time series, the population ACF from equation (1) is estimated by:

$$\hat{\rho}_k = \hat{\gamma}_0^{-1}\hat{\gamma}_k = \left(\sum_{t=1}^{T} x_t^2 \right)^{-1} \left(\sum_{t=k+1}^{T} x_t x_{t-k} \right), \tag{10}$$

where $\hat{\gamma}_k$ is the sample covariance at lag k and $\hat{\gamma}_0$ is the sample variance. Under the mixing assumption, the population ACF (1) approaches zero as lag k increases. However, sampling variability may

yield non-zero autocorrelations even when true values are zero. To mitigate this, the sample ACF estimator (10) imposes a bias towards zero by scaling the autocovariances ($\hat{\gamma}_k$, calculated using $T-k$ terms) by the total sample variance ($\hat{\gamma}_0$, calculated using all T timepoints). Marriott and Pope (1954) demonstrated that this estimator minimizes bias and variance compared to other standard approaches in finite samples.

The exponential decay parameter ϕ_{AD}^* that minimizes the cost function $S_K(\phi_{\text{AD}})$ in equation (5) is estimated by minimizing the sample analog $\widehat{S}_K(\phi_{\text{AD}})$:

$$\widehat{S}_K(\phi_{\text{AD}}) = \frac{1}{K} \sum_{k=1}^{K} (\hat{\rho}_k - \phi_{\text{AD}}^k)^2, \tag{11}$$

$$\hat{\phi}_{\text{AD}}^* = \underset{\phi_{\text{AD}}}{\text{argmin}} \ \widehat{S}_K(\phi_{\text{AD}}), \tag{12}$$

$$\hat{\tau}_{\text{AD}}^* = g(\hat{\phi}_{\text{AD}}^*) = -\frac{1}{\log(|\hat{\phi}_{\text{AD}}^*|)}. \tag{13}$$

Here, (12) and (13) are the sample versions of the population parameters from equations (6) and (7), respectively. In this paper we use the Levenberg-Marquardt algorithm to iteratively update the estimate of $\hat{\phi}_{\text{AD}}^*$ until convergence (i.e., when the step size goes below a 10^{-6} tolerance) (Levenberg, 1944).

2.4 Estimator Properties

In this section, we describe the large-sample properties of both the Time-Domain Linear Model and Autocorrelation-Domain Nonlinear Model, focusing on the consistency and asymptotic distribution of their respective estimators. Under general conditions — when the time-domain method is applied to a process that is not AR1, or the autocorrelation-domain method is applied to a decay process that is not exponential — we demonstrate that the asymptotic distribution is Gaussian, with a limiting variance that can be consistently estimated. Specific expressions for this variance and its estimator are provided after in sections Standard Error of the Estimators and Estimation of Standard Errors. This allows for the construction of hypothesis tests and confidence intervals across timescale maps of the brain.

2.4.1 Time-Domain Method

Following the description in Hansen (2022, Theorem 14.29), the ergodic theorem shows that mixing (which implies ergodicity) is sufficient for consistent estimation. Since X_t is stationary and ergodic, so too are $X_t X_{t-1}$ and X_{t-1}^2, and therefore:

$$\frac{1}{T} \sum_{t=2}^{T} x_t x_{t-1} \xrightarrow{p} \mathbb{E}[X_t X_{t-1}], \quad \frac{1}{T} \sum_{t=2}^{T} x_{t-1}^2 \xrightarrow{p} \mathbb{E}[X_{t-1}^2] \text{ as } T \to \infty.$$

Applying the continuous mapping theorem yields:

$$\hat{\phi}_{\text{TD}}^* = \left(\frac{1}{T} \sum_{t=2}^{T} x_{t-1}^2 \right)^{-1} \left(\frac{1}{T} \sum_{t=2}^{T} x_t x_{t-1} \right) \xrightarrow{p} \mathbb{E}[X_{t-1}^2]^{-1} \mathbb{E}[X_t X_{t-1}] = \phi_{\text{TD}}^* \text{ as } T \to \infty.$$

This shows that the coefficients of the Time-Domain Linear Model can be consistently estimated by least squares, for any stationary and mixing process with parameters defined by projection in equation (3).

Following Hansen (2022, Theorem 14.33), the asymptotic distribution under general dependence states that the behavior of the estimators of ϕ_{TD}^* and τ_{TD}^* can be approximated using a central limit theorem for correlated observations:

$$\sqrt{T}(\hat{\phi}_{\text{TD}}^* - \phi_{\text{TD}}^*) \xrightarrow{d} \mathcal{N}(0, q_{\text{TD}}^{-1}\, \omega_{\text{TD}}\, q_{\text{TD}}^{-1}) \text{ as } T \to \infty.$$

The estimator converges to a Gaussian distribution with a sandwich-form variance with components q_{TD} and ω_{TD} defined below in equation (14). And by the delta method we obtain the asymptotic distribution for the timescale given the transformation from equation (4):

$$\sqrt{T}(\hat{\tau}_{\text{TD}}^* - \tau_{\text{TD}}^*) \xrightarrow{d} \mathcal{N}\left(0, q_{\text{TD}}^{-1}\, \omega_{\text{TD}}\, q_{\text{TD}}^{-1} \left[\frac{d}{d\phi_{\text{TD}}}g(\phi_{\text{TD}})\bigg|_{\phi_{\text{TD}}=\phi_{\text{TD}}^*}\right]^2\right) \text{ as } T \to \infty.$$

2.4.2 Autocorrelation-Domain Method

To show consistent estimation, the sample autocorrelation $\hat{\rho}_k$ (10) is a ratio of sample moments in the time domain, and under the ergodic and continuous mapping theorems (Hansen, 2022, Theorem 14.9):

$$\hat{\rho}_k \xrightarrow{p} \rho_k \text{ as } T \to \infty \text{ for all } k \in \{1, \ldots, K\}.$$

Consequently, the objective function (11) is itself a sample average, such that for any fixed K and ϕ_{AD}:

$$\widehat{S}_K(\phi_{\text{AD}}) = \frac{1}{K}\sum_{k=1}^{K}(\hat{\rho}_k - \phi_{\text{AD}}^k)^2 \xrightarrow{p} \frac{1}{K}\sum_{k=1}^{K}(\rho_k - \phi_{\text{AD}}^k)^2 = S_K(\phi_{\text{AD}}) \text{ as } T \to \infty.$$

Further, if the minimizer ϕ_{AD}^* is unique, $S_K(\phi_{\text{AD}}) > S_K(\phi_{\text{AD}}^*)$ for all $\phi_{\text{AD}} \neq \phi_{\text{AD}}^*$ (Hansen, 2022, Theorem 22.1), then the sample minimizer from (12) converges to the population minimum:

$$\hat{\phi}_{\text{AD}}^* \xrightarrow{p} \phi_{\text{AD}}^* \text{ as } T \to \infty.$$

This shows that the parameters of the Autocorrelation-Domain Nonlinear Model can be consistently estimated by least squares. For this theoretical convergence to hold in practice, we assume that the numerical tolerance of the Levenberg-Marquardt algorithm approaches zero for large T.

With the additional assumption that the objective function (5) is Lipschitz-continuous for ϕ_{AD} near ϕ_{AD}^*, following Hansen (2022, Theorem 23.2), we can approximate the behavior of the estimators of ϕ_{AD}^* and τ_{AD}^* adapting the central limit theorem for nonlinear least squares. Since the estimator $\hat{\phi}_{\text{AD}}^*$ satisfies the first-order condition $\frac{d}{d\phi_{\text{AD}}}\widehat{S}_K(\hat{\phi}_{\text{AD}}^*) = 0$, a mean-value expansion of the first derivative (score function) around the population parameter ϕ_{AD}^* yields the linear approximation:

$$\sqrt{T}(\hat{\phi}_{\text{AD}}^* - \phi_{\text{AD}}^*) \approx -\left(\frac{d^2}{d\phi_{\text{AD}}^2}\widehat{S}_K(\phi_{\text{AD}}^*)\right)^{-1}\left(\sqrt{T}\frac{d}{d\phi_{\text{AD}}}\widehat{S}_K(\phi_{\text{AD}}^*)\right).$$

By the ergodic theorem, the sample second derivative (the first term) converges in probability to the deterministic curvature constant q_{AD}. In addition, the central limit theorem for mixing processes (Hansen, 2022, Theorem 14.33) implies that the scaled score (the second term) converges in distribution to a Gaussian with variance ω_{AD}. Combining these results yields the limiting distribution:

$$\sqrt{T}(\hat{\phi}_{\text{AD}}^* - \phi_{\text{AD}}^*) \xrightarrow{d} \mathcal{N}\left(0, q_{\text{AD}}^{-1}\, \omega_{\text{AD}}\, q_{\text{AD}}^{-1}\right) \text{ as } T \to \infty.$$

Thus, the autocorrelation-domain estimator exhibits Gaussian asymptotic behavior governed by a sandwich variance composed of q_{AD} and ω_{AD}, which are defined below in equation (16). Finally, the asymptotic distribution of the timescale $\hat{\tau}^*_{\text{AD}}$ follows from the delta method, accounting for the nonlinear transformation in equation (7):

$$\sqrt{T}(\hat{\tau}^*_{\text{AD}} - \tau^*_{\text{AD}}) \xrightarrow{d} \mathcal{N}\left(0, q_{\text{AD}}^{-1} \, \omega_{\text{AD}} \, q_{\text{AD}}^{-1} \left[\frac{d}{d\phi_{\text{AD}}} g(\phi_{\text{AD}}) \Big|_{\phi_{\text{AD}} = \phi^*_{\text{AD}}}\right]^2\right) \text{ as } T \to \infty.$$

2.5 Standard Error of the Estimators

2.5.1 Time-Domain Standard Error

Having established in Estimator Properties that the estimators are asymptotically normal, we now derive the specific variance components required for inference. When the data-generating process deviates from AR1 such that the errors are dependent and/or heteroskedastic, the usual (naive) standard errors will be biased. As a result, confidence intervals or hypothesis tests that rely on them are invalid. To correct for this, the Newey and West (1987) (NW) expression takes a sandwich form and explicitly accounts for misspecification by summing the covariance structure of the errors, ensuring that the resulting standard errors are asymptotically valid (Hansen, 2022, Theorem 14.32).

Given that X_t is stationary and mixing, so too are the errors from equation (2) since these properties are preserved by finite transformations (Hansen, 2022, Theorem 14.2 and Theorem 14.12). Consequently, the autocovariances of the errors vanish as the time lag increases (see Assumptions). Further, because the timescale estimator $\hat{\tau}^*_{\text{TD}}$ is given by the nonlinear function $g(\hat{\phi}^*_{\text{TD}})$, its theoretical standard error can be approximated by the delta method, which uses the derivative with respect to ϕ_{TD} evaluated at the pseudo-true parameter value ϕ^*_{TD}:

$$\text{se}_{\text{NW}}(\hat{\phi}^*_{\text{TD}}) = \sqrt{q_{\text{TD}}^{-1} \, \omega_{\text{TD}} \, q_{\text{TD}}^{-1}}, \quad \text{se}_{\text{NW}}(\hat{\tau}^*_{\text{TD}}) \approx \text{se}_{\text{NW}}(\hat{\phi}^*_{\text{TD}}) \cdot \frac{d}{d\phi_{\text{TD}}} g(\phi_{\text{TD}}) \Big|_{\phi_{\text{TD}} = \phi^*_{\text{TD}}},$$

where

$$q_{\text{TD}} = \mathbb{E}[X_{t-1}^2] \quad \text{and} \quad \omega_{\text{TD}} = \sum_{\ell=-\infty}^{\infty} \mathbb{E}[(X_{t-1} \cdot e_t^*)(X_{t-1-\ell} \cdot e_{t-\ell}^*)]. \tag{14}$$

These moments, q_{TD} and ω_{TD}, are stationary and therefore do not depend on time index t. The covariance terms in ω_{TD} capture deviations in the error structure from the standard *iid* case. For the special case of correct specification, when X_t is a true AR1 process, the standard error of $\hat{\phi}^*_{\text{TD}}$ reduces to the usual formula:

$$\text{se}_{\text{Naive}}(\hat{\phi}^*_{\text{TD}}) = \sqrt{\sigma_{\text{TD}}^2 q_{\text{TD}}^{-1}}, \quad \text{se}_{\text{Naive}}(\hat{\tau}^*_{\text{TD}}) \approx \text{se}_{\text{Naive}}(\hat{\phi}^*_{\text{TD}}) \cdot \frac{d}{d\phi_{\text{TD}}} g(\phi_{\text{TD}}) \Big|_{\phi_{\text{TD}} = \hat{\phi}^*_{\text{TD}}}, \tag{15}$$

where σ_{TD}^2 is the time-domain error variance.

2.5.2 Autocorrelation-Domain Standard Error

Equation (5) models autocorrelation decay using a parametric exponential function. The standard errors proposed below account for potential misspecification of this decay (e.g., if the true process is not strictly exponential) and are asymptotically valid for any stationary and mixing process.

Following Hansen (2022, Chapter 22.8 and Chapter 23.5), the theoretical standard error of the estimator $\hat{\phi}^*_{\text{AD}}$ takes a Newey and West (1987) (NW) sandwich form that reflects both the curvature of the squared loss function and the time-domain covariance of the scores. The standard error for the timescale $\hat{\tau}^*_{\text{AD}}$ is then approximated via the delta method:

$$\text{se}_{\text{NW}}(\hat{\phi}^*_{\text{AD}}) = \sqrt{q^{-1}_{\text{AD}} \, \omega_{\text{AD}} \, q^{-1}_{\text{AD}}}, \quad \text{se}_{\text{NW}}(\hat{\tau}^*_{\text{AD}}) \approx \text{se}_{\text{NW}}(\hat{\phi}^*_{\text{AD}}) \cdot \frac{d}{d\phi_{\text{AD}}} g(\phi_{\text{AD}})\Big|_{\phi_{\text{AD}}=\phi^*_{\text{AD}}}.$$

The curvature term q_{AD} is the sum of the squared first derivatives of the regression function ϕ^k_{AD} with respect to ϕ_{AD}, evaluated at the minimizer ϕ^*_{AD}. The variance term ω_{AD} is the long-run variance of the normalized score process u_t (see Appendix for the derivation). This process is a weighted sum of the time-domain score contributions across all K lags, where the weights are given by the regression function derivative:

$$q_{\text{AD}} = \sum_{k=1}^{K}(k\phi^{*k-1}_{\text{AD}})^2, \quad \omega_{\text{AD}} = \sum_{\ell=-\infty}^{\infty} \mathbb{E}[u_t u_{t-\ell}], \quad u_t = \frac{1}{\gamma_0}\sum_{k=1}^{K}(k\phi^{*k-1}_{\text{AD}})(X_t X_{t-k} - \phi^{*k}_{\text{AD}} X_t^2). \tag{16}$$

Defining ω_{AD} in terms of the score process u_t maps the joint dependence structure of the autocorrelation sequence $\rho_1, \ldots, \rho_K$ onto a scalar process. Of note, u_t is defined in the time domain since $\rho_k = \gamma_k/\gamma_0$ is a function of moments of X_t. This linearized representation ensures that the long-run variance ω_{AD} captures the relevant dependencies, including both the correlations within and between lags. For a stationary process with $|\phi^*_{\text{AD}}| < 1$, the weights $k\phi^{*k-1}_{\text{AD}}$ in the score process decrease rapidly with lag k, meaning the asymptotic variance is largely determined by the variability of the lower-order sample autocorrelations. Under the special case of correct AR1 specification, higher-order lags provide no additional information as they are deterministic functions of the first lag ρ_1. In this case, the autocorrelation-domain standard error is asymptotically equivalent to the time-domain standard error in (15), replacing ϕ^*_{TD} with ϕ^*_{AD}.

2.6 Estimation of Standard Errors

2.6.1 Time-Domain Standard Error Estimator

The sample standard error estimator takes the form:

$$\widehat{\text{se}}_{\text{NW}}(\hat{\phi}^*_{\text{TD}}) = \sqrt{\hat{q}^{-1}_{\text{TD}} \, \hat{\omega}_{\text{TD}} \, \hat{q}^{-1}_{\text{TD}}}, \quad \widehat{\text{se}}_{\text{NW}}(\hat{\tau}^*_{\text{TD}}) \approx \widehat{\text{se}}_{\text{NW}}(\hat{\phi}^*_{\text{TD}}) \cdot \frac{d}{d\phi_{\text{TD}}} g(\phi_{\text{TD}})\Big|_{\phi_{\text{TD}}=\hat{\phi}^*_{\text{TD}}},$$

where

$$\hat{q}_{\text{TD}} = \frac{1}{T}\sum_{t=2}^{T} x^2_{t-1} \quad \text{and} \quad \hat{\omega}_{\text{TD}} = \sum_{\ell=-M}^{M}\left(1 - \frac{|\ell|}{M+1}\right)\frac{1}{T}\sum_{1 \le t-\ell \le T}(x_{t-1} \cdot \hat{e}^*_t)(x_{t-1-\ell} \cdot \hat{e}^*_{t-\ell}).$$

This estimator calculates a weighted sum of the regression scores $x_{t-1} \cdot \hat{e}^*_t$, where $\hat{e}^*_t = x_t - \hat{\phi}^*_{\text{TD}} \cdot x_{t-1}$. The true ω_{TD} is approximated by $\hat{\omega}_{\text{TD}}$ by taking a finite sum of the regression score covariances up to lag M, where M is the lag-truncation (or bandwidth). The weights used in the sum decrease linearly with lag ℓ, following a Bartlett kernel (Newey and West, 1987). This kernel not only ensures the standard errors remain non-negative but also regularizes $\hat{\omega}_{\text{TD}}$ to change smoothly with M (Hansen, 2022, Chapter 14.35).

For comparison we also include the naive estimator which simplifies under correct specification:

$$\widehat{\text{se}}_{\text{Naive}}(\hat{\phi}^*_{\text{TD}}) = \sqrt{\hat{\sigma}^2_{\text{TD}}\hat{q}^{-1}_{\text{TD}}}, \quad \widehat{\text{se}}_{\text{Naive}}(\hat{\tau}^*_{\text{TD}}) \approx \widehat{\text{se}}_{\text{Naive}}(\hat{\phi}^*_{\text{TD}}) \cdot \frac{d}{d\phi_{\text{TD}}} g(\phi_{\text{TD}})\Big|_{\phi_{\text{TD}}=\hat{\phi}^*_{\text{TD}}}, \tag{17}$$

where $\hat{\sigma}^2_{\text{TD}} = 1/T \sum_{t=2}^{T} \hat{e}^{*2}_t$ is an estimate of the time-domain error variance.

2.6.2 Autocorrelation-Domain Standard Error Estimator

The sample standard error estimator takes the form:

$$\widehat{\text{se}}_{\text{NW}}(\hat{\phi}^*_{\text{AD}}) = \sqrt{\hat{q}_{\text{AD}}^{-1}\,\hat{\omega}_{\text{AD}}\,\hat{q}_{\text{AD}}^{-1}}, \quad \widehat{\text{se}}_{\text{NW}}(\hat{\tau}^*_{\text{AD}}) \approx \widehat{\text{se}}_{\text{NW}}(\hat{\phi}^*_{\text{AD}}) \cdot \frac{d}{d\phi_{\text{AD}}} g(\phi_{\text{AD}})\Big|_{\phi_{\text{AD}}=\hat{\phi}^*_{\text{AD}}},$$

where the curvature $\hat{q}_{\text{AD}}$ is estimated by the sum of the squared first derivatives evaluated at $\hat{\phi}^*_{\text{AD}}$, and the long-run variance ω_{AD} is obtained by applying a Newey–West estimator to the sample normalized score process $\hat{u}_t$:

$$\hat{q}_{\text{AD}} = \sum_{k=1}^{K}(k\hat{\phi}^{*\,k-1}_{\text{AD}})^2, \quad \hat{\omega}_{\text{AD}} = \sum_{\ell=-M}^{M}\left(1 - \frac{|\ell|}{M+1}\right)\frac{1}{T}\sum_{1\leq t-\ell\leq T}\hat{u}_t\hat{u}_{t-\ell}, \quad \hat{u}_t = \frac{1}{\hat{\gamma}_0}\sum_{k=1}^{K}(k\hat{\phi}^{*\,k-1}_{\text{AD}})(x_t x_{t-k} - \hat{\phi}^{*\,k}_{\text{AD}}x_t^2).$$

The estimator $\hat{\omega}_{\text{AD}}$ sums sample autocovariances of $\hat{u}_t$ up to lag M using a Bartlett kernel. This explicitly accounts for serial dependence in x_t that is not captured by the fitted autocorrelation model, ensuring robust inference under misspecification (e.g., non-exponential decay). In the case of correct specification, the equation simplifies to (17), replacing $\hat{\phi}^*_{\text{TD}}$ with $\hat{\phi}^*_{\text{AD}}$.

3 Simulations

3.1 Simulation Settings

Monte Carlo simulations with $B = 10,000$ replications were used to evaluate time- and autocorrelation-domain methods. Time series realizations $x = \{x_1, x_2, ..., x_T\}$ with $T = 4800$ were based on three data-generating processes, each characterized by a different autocorrelation structure: AR1, AR2, and autocorrelations derived from rfMRI data. All autocorrelation structures shared the same time-domain projection parameter (ϕ^*_{TD}) for comparable timescales. That is, there is always a ϕ^*_{TD} value that represents the time-domain AR1 model projection, even if the time series was generated by a more complex process. To define a feasible parameter range for simulation, we referred to the Human Connectome Project (HCP) dataset, where $\hat{\phi}^*_{\text{TD}}$ estimates ranged from +0.1 to +0.8. Accordingly, autocorrelation strength was varied using five positive ϕ^*_{TD} values (0.1 to 0.8) with corresponding timescales τ^*_{TD} from equation (4). This design resulted in a total of 15 simulation settings (three data-generating models $\times$ five autocorrelation strengths). For each setting, estimator performance was assessed by relative root mean squared error (rRMSE):

$$\text{rRMSE}(\hat{\tau}^*) = \frac{\sqrt{\frac{1}{B}\sum_{b=1}^{B}(\hat{\tau}^*_b - \tau^*)^2}}{\tau^*}, \quad \text{rRMSE}(\widehat{\text{se}}(\hat{\tau}^*)) = \frac{\sqrt{\frac{1}{B}\sum_{b=1}^{B}\left(\widehat{\text{se}}(\hat{\tau}^*_b) - \text{sd}(\hat{\tau}^*)\right)^2}}{\text{sd}(\hat{\tau}^*)},$$

where $\text{sd}(\hat{\tau}^*) = \sqrt{\frac{1}{B}\sum_{b=1}^{B}(\hat{\tau}^*_b - \bar{\tau}^*)^2}$ is the empirical standard deviation of the timescale estimator across B simulation replications, and $\bar{\tau}^* = \frac{1}{B}\sum_{b=1}^{B}\hat{\tau}^*_b$. The rRMSE captures both bias and variance of timescale estimates and their standard errors as a proportion of the true value, facilitating comparison across parameter ranges and simulation settings. Cutoffs of <10% rRMSE for timescales and <20% for standard errors are shown to indicate "good" estimation, following applied time series conventions (Jadon et al., 2024).

In the *AR1 setting*, the data-generating process matches the fitted model, with time series simulated from the AR1 model:

$$x_t = \phi^*_{\text{TD}}x_{t-1} + e_t, \quad e_t \overset{\text{iid}}{\sim} \mathcal{N}(0,1). \tag{18}$$

The *AR2 setting* misspecifies the order of the autoregression by introducing a mismatch between the AR2 data-generating process and AR1 fitting, with five pairs of AR2 coefficients selected so that the ϕ_{TD}^* projection matched the above setting (see Figure 2 Panel A). The following model was used:

$$x_t = \phi_1 x_{t-1} + \phi_2 x_{t-2} + e_t, \quad e_t \overset{\text{iid}}{\sim} \mathcal{N}(0, 1). \tag{19}$$

The *HCP setting* did not follow an autoregressive process, and instead used empirical ACFs from five brain regions of subject #100610 in the HCP dataset (see Dataset Description). These regions were selected to match the ϕ_{TD}^* projections above. To simulate time series with the same autocorrelation structure as the empirical data, we sampled from a multivariate normal distribution $\mathcal{N}(0, \hat{\Sigma})$, where $\hat{\Sigma} \in \mathbb{R}^{T \times T}$ is the covariance matrix constructed from the sample ACFs. Under stationarity, $\hat{\Sigma}$ is Toeplitz, so its k^{th} off-diagonal equals the sample ACF at lag k: $\hat{\Sigma}_{i,j} = \hat{\rho}_{|i-j|}$ for indices $i, j = \{1, \ldots, T\}$. A Cholesky decomposition of the covariance matrix $\hat{\Sigma} = LL^\top$, where L is a lower triangular matrix, was multiplied with Gaussian white noise:

$$x_t = \sum_{j=1}^{t} L_{t,j} e_j, \quad e_j \overset{\text{iid}}{\sim} \mathcal{N}(0, 1). \tag{20}$$

3.2 Simulation Results

3.2.1 Results for Autoregressive Simulations

Figure 1: *AR1 time-domain simulations.*

(**A**) Simulation Setting: solid lines show the simulated ACFs of the process defined in equation (18); dashed lines show the AR1-projected ACFs which are overlapping as both follow AR1. Horizontal line marks the timescale where the AR1-projected ACF reaches $1/e \approx 0.37$. (**B**) Timescale Estimates: vertical lines show true timescales τ^*; histograms show estimates $\hat{\tau}^*$ across $B = 10,000$ replications; points show rRMSE versus a 10% error line. (**Row 1**): Time-Domain Linear Estimator. (**Row 2**): Autocorrelation-Domain Nonlinear Estimator. (**C**) Naive and (**D**) Newey-West Standard Errors: vertical lines show true standard errors se($\hat{\tau}^*$), i.e., standard deviations of the sampling distributions from panel B; histograms show standard error estimates $\widehat{se}(\hat{\tau}^*)$; points show rRMSE versus a 20% error line. (**Row 1**): Time-Domain Standard Error Estimator. (**Row 2**): Autocorrelation-Domain Standard Error Estimator.

AR1 simulations (**Figure 1**) serve as a baseline for correct specification, where the data-generating process defined in (18) aligns with the fitted models (**Panel A**). **Panel B** demonstrates that while estimation variance increases with the true timescale, rRMSE remains below the conventional 10% threshold. Regarding uncertainty quantification (**Panels C-D**), the time-domain estimator (**Row 1**) shows minimal difference between Naive and Newey-West approaches because of correct specification. However, the autocorrelation-domain estimator (**Row 2**) exhibits a downward bias that is mitigated by the Newey-West correction. This bias likely occurs because fitting multiple noisy lags in the autocorrelation domain introduces additional sampling variability that the naive formula does not capture. Consistent with this, the autocorrelation-domain method produces point estimates and standard errors with higher overall variance than the time-domain method.

Figure 2: *AR2 time-domain simulations.*

(**A**) Simulation Setting (**Left**): triangle shows the AR2 stationary region in the (ϕ_1, ϕ_2) plane, with the periodic/aperiodic boundary at $\phi_2 = -\phi_1^2/4$. Dashed lines within the triangle illustrate the preimage under the AR2-to-AR1 projection mapping; specifically, all (ϕ_1, ϕ_2) pairs on a given dashed line correspond to AR2 processes that project to the same AR1. Points show five selected AR2 (ϕ_1, ϕ_2) pairs, with their corresponding AR1 projections shown in the colorbar. (**Right**): solid lines plot the simulated ACFs from the process defined in equation (19); dashed lines show the ACFs based on AR1-projected models, which are not overlapping because the simulated AR2 differs from its AR1 fit. (**B**) Timescale Estimates: vertical lines show true timescales τ^*; histograms show estimates $\hat{\tau}^*$ across $B = 10,000$ replications; points show rRMSE versus a 10% error line. (**Row 1**): Time-Domain Linear Estimator. (**Row 2**): Autocorrelation-Domain Nonlinear Estimator. (**C**) Naive and (**D**) Newey-West Standard Errors: vertical lines show true standard errors $\mathrm{se}(\hat{\tau}^*)$, i.e., standard deviations of the sampling distributions from panel B; histograms show standard error estimates $\widehat{\mathrm{se}}(\hat{\tau}^*)$; points show rRMSE versus a 20% error line. (**Row 1**): Time-Domain Standard Error Estimator. (**Row 2**): Autocorrelation-Domain Standard Error Estimator.

AR2 simulations (**Figure 2**) assess estimator performance under model misspecification, using a higher-order process defined in (19) that deviates from the fitted AR1 model (**Panel A**). **Panel B** shows that the two estimators converge to different timescales due to their distinct definitions. In terms of standard errors, **Panel C** shows a consistent underestimation by the naive estimator across both methods because of misspecification. **Panel D** demonstrates that this bias is largely mitigated by the Newey-West estimator.

3.2.2 Results for Realistic rfMRI Simulations

Figure 3: *Realistic rfMRI time-domain simulations.*

(**A**) Simulation Setting: solid lines show estimated ACFs from five selected brain regions of HCP subject #100610 used to generate the process define in equation (20); dashed lines show the AR1-projected ACFs, which are not overlapping because the empirically estimated ACF differs from its AR1 fit. (**B**) Timescale Estimates: vertical lines show true timescales τ^*; histograms show estimates $\hat{\tau}^*$ across $B = 10{,}000$ replications; points show rRMSE versus a 10% error line. (**Row 1**): Time-Domain Linear Estimator. (**Row 2**): Autocorrelation-Domain Nonlinear Estimator. (**C**) Naive and (**D**) Newey-West Standard Errors: vertical lines show true standard errors se$(\hat{\tau}^*)$, i.e., standard deviations of the sampling distributions from panel B; histograms show standard error estimates $\widehat{se}(\hat{\tau}^*)$; points show rRMSE versus a 20% error line. (**Row 1**): Time-Domain Standard Error Estimator. (**Row 2**): Autocorrelation-Domain Standard Error Estimator.

Realistic rfMRI simulations (**Figure 3**) evaluate estimator performance using the data-generating process defined in (20) with realistic autocorrelation derived from five distinct brain regions. Similar to the AR2 setting, this analysis examines the impact of model misspecification, as the empirical ACFs differ from the projected AR1 fits (**Panel A**). **Panel B** confirms that time- and autocorrelation-domain estimators capture distinct timescale values due to their respective definitions. **Panels C-D** further demonstrate that naive standard errors are systematically underestimated, a bias that is substantially mitigated by the Newey-West correction.

4 Data Analysis

4.1 Dataset Description

Resting fMRI (rfMRI) scans were provided by the Human Connectome Project (HCP), WU-Minn Consortium (led by principal investigators David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers supporting the NIH Blueprint for Neuroscience Research, and by the McDonnell Center for Systems Neuroscience at Washington University (Van Essen et al., 2013).

Informed consent was obtained from all participants. Two subsets of the dataset were used: one for methods development and defining realistic simulation parameters (see Simulations), and the other for estimating high-resolution timescale maps of the cortex.

The *methods development subset* included 10 subjects (#100004 - #101410) scanned with a 3T gradient-echo EPI sequence (TR=720ms, slice thickness=2mm). Each subject completed four 15-minute runs (4800 timepoints total), preprocessed with standard steps including motion regression and artifact removal with ICA-FIX (see Glasser et al. (2013); Salimi-Khorshidi et al. (2014) for details). To further mitigate systemic low-frequency physiological oscillations driven by respiration and CO_2 variation, we applied RAPIDtide (Frederick, 2025). This step accounts for the regionally varying lags of these global signals, caused by differences in blood arrival times, which can otherwise mimic or inflate regionally specific autocorrelation. The four runs were concatenated into a single continuous time series, under the assumption of weak stationarity (constant mean and variance) across denoised runs. The resulting dataset dimensions were {10 subjects, 4800 timepoints, 300 regions}. The *timescale mapping subset* included 180 subjects scanned with a 7T gradient-echo EPI sequence (TR=1000ms, slice thickness=1.6mm) over four 16-minute runs (3600 timepoints total), using the same preprocessing and concatenation steps. Functional data were analyzed on the cortical surface downsampled to 2mm spatial resolution, yielding a dataset with the dimensions {180 subjects, 3600 timepoints, 64984 vertices}. The time- and autocorrelation-domain methods were fit to each vertex independently, a mass-univariate analysis approach that resulted in subject-level spatial maps of temporal timescale estimates and their standard errors.

4.1.1 Group-level Analysis

Group-level maps combined individual timescales and standard errors, accounting for within- and between-subject variability. For the group-level analysis, we used the *timescale mapping subset* of $N = 180$ subjects $T = 3600$ concatenated timepoints, distinct from the *methods development subset* of $N = 10$ subjects described above. While remaining within the mass-univariate framework, for simplicity, we express the group timescale at a single cortical vertex:

$$\hat{\tau}_n \text{ for } n \in \{1, 2, ..., N\},$$

$$\bar{\tau} = \frac{1}{N} \sum_{n=1}^{N} \hat{\tau}_n.$$

The group-level standard error for the timescale is given by the law of total variance:

$$\widehat{\mathsf{se}}(\bar{\tau}) = \sqrt{\frac{1}{N} \sum_{n=1}^{N} \widehat{\mathsf{se}}(\hat{\tau}_n)^2 + \frac{1}{N} \sum_{n=1}^{N} (\hat{\tau}_n - \bar{\tau})^2}.$$

Here, the first term under the square root is the within-individual variance and the second term is the between-individual variance.

For visualization, brain-wide t-statistic maps tested whether timescales exceeded 0.5 seconds under the null hypothesis $H_0 : \tau \le 0.5$, by comparing the empirical group mean $\bar{\tau}$ to the null value 0.5:

$$t(\bar{\tau}) = \frac{\bar{\tau} - 0.5}{\widehat{\mathsf{se}}(\bar{\tau})}.$$

Similarly, relative standard error (RSE) maps visualize the spatial precision of timescale estimates to

show the reciprocal of the unadjusted (zero-valued null) t-ratio:

$$\text{rse}(\bar{\tau}) = \frac{\widehat{\text{se}}(\bar{\tau})}{\bar{\tau}}.$$

4.2 Data Analysis Results

Figure 4: *Human Connectome Project subject-level timescale maps.*

(**A-D**) Cortical surface maps from HCP subject #100610. Displays show lateral-left, lateral-right, medial-left, and medial-right views, plotted using the *StatBrainz* package (Davenport, 2025). The upper bounds on the colorbars are set for each panel at the 99th percentile of cortical map values. For each panel, the top row shows results of the time-domain method, and the bottom row the autocorrelation-domain method. (**A**) Timescale estimates: maps display the timescales (in seconds) estimated at each vertex. (**B**) Newey-West standard errors: shows the spatial distribution of standard errors, where smaller values indicate greater estimation precision. (**C**) T-statistics: unthresholded and uncorrected t-ratios testing where timescales exceed 0.5 seconds. (**D**) Relative Standard Errors (RSEs): relative reliability of estimates, where low RSE (near zero) indicates high precision with small uncertainty.

4.2.1 Results for rfMRI Timescale Maps

Subject-level maps (**Figure 4**) were generated by mass-univariate fitting of time- and autocorrelation-domain estimators to cortical surface data at the individual vertex level (64984 vertices). **Panel A** presents the spatial distribution of timescale estimates, which show that autocorrelation-domain estimates tend to yield larger timescales than the time-domain estimator, consistent with simulation results. **Panel B** shows the corresponding maps of standard errors, which are aligned with the timescale maps, consistent with the simulation finding that larger timescales are associated with greater sampling variability. **Panel C** depicts the relative ratio of timescale estimates to their standard errors (i.e., t-statistics), testing whether the timescales significantly exceed a half second ($H_0 : \tau \leq 0.5$). Despite larger standard errors for higher timescales, these regions still exhibit higher t-statistics. **Panel D** shows low RSEs across much of the brain indicating high estimation reliability.

Figure 5: *Human Connectome Project group-level timescale maps.*

(**A-D**) Cortical surface maps from $N = 180$ HCP subjects. Displays show lateral-left, lateral-right, medial-left, and medial-right views, plotted using the *StatBrainz* package (Davenport, 2025). The upper bounds on the colorbars are set for each panel at the 99th percentile of cortical map values. For each panel, the top row shows results of the time-domain method, and the bottom row the autocorrelation-domain method. (**A**) Timescale estimates: maps display the group-level timescales (in seconds) at each vertex, averaged over subjects. (**B**) Newey-West standard errors: group-level spatial distribution of estimates, accounting for within- and between-subject variability. Smaller values indicate greater precision. (**C**) T-statistics: unthresholded and uncorrected t-ratios testing where group-level timescales exceed 0.5 seconds. (**D**) Relative standard errors (RSEs): relative reliability of estimates, where low RSE (near zero) indicates high precision with small uncertainty across subjects. (**E**) Timescale organization by brain network: maps display brain networks from the Yeo 7 Network Atlas, ordered by the network-averaged t-statistics (from panel C). This ordering is the same for time- and autocorrelation-domain methods, and highlights the hierarchical organization of timescales, progressing from sensory networks (e.g., somatomotor and limbic in blues) to association networks (e.g., frontoparietal and default mode in reds).

Group-level timescale maps (**Figure 5**) were generated by combining individual estimates at the vertex level to account for both within- and between-individual variability, providing an aggregate view of timescale distributions across subjects. The smooth appearance of these timescale maps reflects the high spatial resolution of the vertex-level analysis and the averaging across subjects. **Panel A** shows that the average timescale maps for both time- and autocorrelation-domain estimators are smoother than the individual maps, displaying a well-organized spatial pattern across the cortex; notably autocorrelation-domain estimates are generally larger than time-domain estimates. **Panel B** presents the standard error maps, which combine variances from within-subject Newey-West estimates and between-subject timescale estimates. As expected from the simulation results, the standard errors are larger in the autocorrelation domain than in the time domain. **Panel C** depicts t-statistics testing whether timescales exceed half a second, while **Panel D** plots relative standard errors (RSEs), providing complementary views of estimation reliability. Interpreting these maps together

clarifies how signal strength interacts with uncertainty: although standard errors theoretically increase with the true timescale (as shown in sections 2.5 and 3), these empirical results demonstrate that the estimated timescale magnitude grows faster than its associated error. Consequently, regions with longer timescales actually exhibit higher signal-to-noise ratios (larger t-statistics). In contrast, regions with shorter timescales, such as the limbic network (comprising orbitalfrontal cortex and anterior temporal cortex), exhibit greater relative uncertainty and lower signal-to-noise ratios (Yeo et al., 2011; Du et al., 2024).

Panel E highlights the spatial organization of timescales into networks by mapping the t-statistic at each vertex to one of seven networks from the Yeo et al. (2011) atlas. The ordering {limbic, somatomotor, ventral attention, visual, dorsal attention, default, frontoparietal} aligns with the sensory-to-association axis of brain organization and is consistent with an extensive literature on the hierarchical organization of timescales (Murray et al., 2014; Stephens et al., 2013; Raut et al., 2020; Gao et al., 2020; Hasson et al., 2008). Sensory networks (limbic and somatomotor) have short timescales, followed by attentional networks (ventral and dorsal attention), while higher-order association networks (default mode and frontoparietal) show the longest timescales. However, interpreting these network differences requires caution, particularly because regions in the limbic network are especially susceptible to measurement noise which can conflate timescale estimation (see Effects of Measurement Noise).

This empirical analysis highlights methodological considerations for estimating fMRI timescale maps. (i) Both time- and autocorrelation-domain methods produce similar maps, but diverge at extremes – the autocorrelation-domain estimator yields larger timescales. The corresponding standard errors in the autocorrelation domain are also larger, so the resulting t-ratio maps remain similar between the two methods, highlighting why point estimates can be misleading without considering standard errors. Likewise, estimates are on average more precise for the time-domain estimator, consistent with simulation results. (ii) In this mass-univariate analysis, the computational cost of the time-domain estimator is substantially lower than autocorrelation-domain estimator because of its simple analytical solution. (iii) The t-statistic maps organized by brain network exhibit a clear hierarchical timescale organization, reflecting how networks integrate and process information over time. And this pattern is consistent regardless of the estimation method. (iv) As a sensitivity check, we repeated the analyses without added RAPIDtide denoising step and found that it had only a small influence on local vertex-level timescale estimates, and it did not alter the global network-level organization (e.g., the relative ordering of networks by timescale). Taken together, these findings suggest that the time-domain method may be preferable for large-scale neuroimaging studies due to its computational efficiency and higher precision, while producing maps consistent with previously reported timescale hierarchies, and that are robust to two different denoising strategies.

5   Practical Considerations for fMRI Timescale Estimation

In this section, we provide recommendations for practitioners seeking to estimate fMRI timescale maps. We advocate for the Time-Domain Linear Estimator over the Autocorrelation-Domain Nonlinear Estimator, as it is more computationally efficient while maintaining robust performance in Simulations. However, a critical distinction must be drawn between the *statistical validity* of the estimator and the *interpretability* of the resulting timescale maps, which inevitably reflect the composite nature of the observed fMRI BOLD signals. The mixed metabolic and neuronal origins of the hemodynamic signal complicate mechanistic interpretations (Raut et al., 2020; He, 2011). Consequently, interpreting fMRI timescales as proxies for *intrinsic neural timescales* requires disentangling latent neural dynamics from confounding sources of variability. The following subsections detail how the hemodynamic response function, measurement noise, and periodicity specifically impact these interpretations.

5.1 Effect of the Hemodynamic Response Function

The fMRI BOLD signal is a composite measure reflecting underlying neural activity filtered through a complex vascular system. While neural populations exhibit the characteristic exponential decay in autocorrelation that defines the intrinsic neural timescale (Murray et al., 2014, Supplementary Mathematical Note), the hemodynamic response function (HRF) acts as a temporal low-pass filter on this latent neural activity. Forward modeling indicates that while this filtering inflates absolute timescale values, it generally preserves the relative spatial hierarchy, recovering the sensory-to-association axis observed empirically (Watanabe et al., 2019).

However, reliance on a canonical HRF may oversimplify the relationship between neural and vascular dynamics. Recent evidence demonstrates that HRF shape varies across both brain regions and individuals (Bailes et al., 2023; Rangaprakash et al., 2018). This variability introduces a potential confound where observed timescale differences may stem from vascular origins, such as regional heterogeneity in blood flow latency, rather than differences in neural processing. A method proposed by Wu et al. (2021) for HRF estimation and deconvolution offers a promising route to disentangle these effects from the resting-state BOLD signal. Nevertheless, a fundamental identifiability problem remains in that both a wider HRF and a longer neural timescale can produce nearly identical observed BOLD signals. Resolving this ambiguity requires moving beyond observational studies, and future work should leverage experimental manipulations and/or concurrent EEG-fMRI to tease apart vascular and neural timescales.

5.2 Effects of Measurement Noise

Figure 6: *AR1 with added measurement noise simulations.*

(**A**) Simulation Setting: solid lines show the simulated ACFs of the AR1 + measurement noise process defined in equation (21) with a fixed timescale and increasing measurement noise variance σ_η^2; the dashed black line shows the AR1-projected ACF with no noise variance ($\sigma_\eta^2 = 0$). Increasing noise (red-to-blue) increases the rate of autocorrelation decay even when the latent timescale is fixed. (**B**) Timescale and Standard Error Estimates: vertical lines show the true latent neural timescale ($\tau_{TD} = 4.48$ for all settings); histograms show estimates across $B = 10,000$ replications; points show rRMSE. (**Row 1**) Time-Domain Linear Estimator, where rRMSE is plotted versus a 10% error line. (**Row 2**) Time-Domain Standard Error Estimator (Newey-West), where rRMSE is plotted versus a 20% error line.

While the HRF tends to lengthen timescales, measurement noise typically biases estimates in the opposite direction. In the fMRI setting, the observed time series is rarely generated by a clean autoregressive process; rather, it is a sum of distinct stochastic processes capturing neural activity, hemody-

namic filtering, and measurement noise. In such settings, regions with identical latent neural timescales can yield divergent observed fMRI timescales depending on their local signal-to-noise ratio. This raises an interpretational concern where the fitted timescale may be statistically well-estimated yet fail to correspond to the intended parameter of neural integration if measurement noise dominates.

To illustrate this limitation, consider the simulation setting where the latent neural process is generated as an AR1 with a fixed timescale $\tau_{\text{TD}}^* = 4.48$ ($\phi_{\text{TD}}^* = 0.8$):

$$z_t = \phi_{\text{TD}}^* z_{t-1} + e_t, \quad e_t \overset{\text{iid}}{\sim} \mathcal{N}(0,1).$$

Observations are formed by adding independent measurement noise of increasing magnitude to vary observation quality while holding the latent timescale constant:

$$x_t = z_t + \epsilon_t, \quad \epsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\epsilon^2), \quad \sigma_\epsilon^2 \in \{0, 0.5, 1.0, 1.5, 2.0\}, \tag{21}$$

Under this construction, increasing σ_ϵ^2 accelerates the rate of autocorrelation decay (**Figure 6 Panel A**). This occurs because noise inflates the lag-0 variance without altering cross-lag covariances, shrinking the autocorrelation ratio in equation (1). Consequently, the time-domain estimator in **Panel B row 1** diverges from the true latent neural timescale, reflecting instead the *pseudo-true timescale* of the misspecified process (see Estimation and Inference under Misspecification). Importantly, while the Newey-West standard errors in **Panel B row 2** ensure that inference on this pseudo-true timescale remains statistically valid, its interpretation should reflect the mixture of neural integration and confounding measurement noise.

This dependency emphasizes that the fitted timescale is effectively a confounded measure of both neural dynamics and measurement noise and the resulting timescale should be interpreted as a pseudo-true timescale under model misspecification. Therefore in practice, effective denoising is a prerequisite for separating neural timescales from artifacts such as head motion and physiological noise. Regarding motion, the proposed Time-Domain Linear Estimator offers a practical advantage; while standard strategies often require censoring followed by interpolation (Goldberg et al., 2024), this estimator relies solely on adjacent timepoint pairs (x_t, x_{t+1}). This allows motion-contaminated frames to be excluded without interpolation. Finally, even with careful denoising, interpretational caution is warranted in regions susceptible to signal loss, such as the orbital frontal cortex and the anterior inferior medial temporal lobe (Yeo et al., 2011; Du et al., 2024).

5.3 Effects of Periodicity

In addition to the HRF and measurement noise, we consider the influence of periodic processes, since neural timescales are generally considered to originate from the brain's aperiodic (non-rhythmic) activity. While monotonic autocorrelation decay is consistent with an AR1 process, latent neural activity may exhibit rhythmic dynamics (He et al., 2010; He, 2011). Furthermore, non-neural sources, such as respiration and CO_2 variation, can introduce systemic low-frequency oscillations (Tong et al., 2019; Korponay et al., 2024). These dynamics manifest as periodic fluctuations in the ACF, creating a case of model misspecification. To validate the robustness of the time-domain estimator in this setting, we simulated AR2 processes that display oscillations, holding the projected timescale constant ($\tau_{\text{TD}}^* = 4.48$) while increasing the oscillatory strength via the second-order coefficient $\sqrt{-\phi_2} \in \{0.5, 0.642, 0.758, 0.859, 0.949\}$.

As shown in **Figure 7 Panel A**, we selected parameters along a preimage of constant timescale but increasing periodicity. Despite the mismatch between the data generating process (AR2) and the estimator (AR1), the histograms in **Panel B row 1** are accurately centered on the pseudo-true timescale. This demonstrates that the estimator recovers the dominant decay rate even amidst substantial oscillations. Notably, increasing the strength of oscillations decreases the variance of the sampling distributions; as the process becomes more rhythmic, the time-domain error variance shrinks. Consistent

Figure 7: *AR2 with damped oscillation simulations.*

(**A**) Simulation Setting (**Left**): triangle shows the AR2 stationary region in the (ϕ_1, ϕ_2) plane, with the periodic/aperiodic boundary at $\phi_2 = -\phi_1^2/4$. Points show five selected AR2 (ϕ_1, ϕ_2) combinations that project to the same AR1 timescale, with increasing periodicity magnitude by adjusting $\sqrt{-\phi_2}$. (**Right**): solid lines show the simulated ACFs of the AR2 process defined in equation (19); the dashed black line shows the same AR1-projected ACF for all processes. (**B**) Timescale and Standard Error Estimates: vertical lines show the true timescale ($\tau_{TD} = 4.48$ for all settings); histograms show estimates across $B = 10,000$ replications; points show rRMSE. (**Row 1**) Time-Domain Linear Estimator, where rRMSE is plotted versus a 10% error line. (**Row 2**) Time-Domain Standard Error Estimator (Newey-West), where rRMSE is plotted versus a 20% error line.

with this, the Newey-West standard errors (**Panel B row 2**) accurately track this reduction in sampling variability.

Taken together, these results suggest that timescale inference can be reliably applied to oscillatory processes without explicitly disambiguating aperiodic and periodic components, a separation often required by frequency-domain methods (Donoghue et al., 2020; Gao et al., 2020). However, distinguishing neural from non-neural confounding remains critical for interpretation. We therefore recommend denoising methods such as RAPIDtide to remove confounding time-lagged oscillations, ensuring the estimated timescale more closely reflects neural rather than physiological processes (Frederick, 2025). That said, the maps reported in Data Analysis were robust to the exclusion of RAPIDtide denoising, suggesting the hierarchical organization of fMRI timescales is generalizable beyond this specific denoising strategy.

6   Conclusions

This study introduces statistical methods for mapping fMRI timescales, characterizing differences in local temporal autocorrelation across brain regions. We detail the large-sample properties of time- and autocorrelation-domain methods, showing that while both converge, they converge to different values due to their distinct parameter definitions. We also demonstrate that both estimators yield consistent standard errors under broad conditions, enabling reliable inference and hypothesis testing. This addresses a major limitation in fMRI timescale studies, which have typically reported point estimates without uncertainty measures.

Simulation results highlight differences in finite sample bias and variance between the methods. While all estimators are largely unbiased, the time-domain method performs as well as, or better than, the autocorrelation-domain method in terms of relative root mean square error (rRMSE). Standard

error estimates exhibit pronounced differences between naive and Newey-West corrected methods, particularly under misspecified settings such as AR2 processes and realistic rfMRI data. Overall, the time-domain method (2) offers greater computational efficiency and estimation stability, whereas the autocorrelation-domain method (5) accommodates longer-range autocorrelations at the cost of reduced accuracy. The application of Newey-West corrected standard errors enhances inference reliability by reducing uncertainty bias across both methods (Newey and West, 1987), which is particularly vital when the autocorrelation decay is not strictly exponential.

Applied to HCP rfMRI data, both methods yield timescale and t-statistic maps consistent with known functional hierarchies (Van Essen et al., 2013). These results align with prior work demonstrating larger timescales in associative versus sensory cortices (Raut et al., 2020; Shafiei et al., 2020; Lurie et al., 2024; Mitra et al., 2014; Kaneoke et al., 2012; Wengler et al., 2020; Shinn et al., 2023; Manea et al., 2022; Ito et al., 2020; Müller et al., 2020). The spatial patterns reinforce the hypothesis that hierarchical organization governs temporal processing across cortical regions. However, mechanistic interpretations must carefully distinguish intrinsic neural dynamics from the confounding effects of hemodynamics and measurement noise. Linking macroscopic fMRI timescales to their underlying physiology remains an important objective for future research.

In conclusion, we present robust rfMRI estimators for timescales and standard errors, enabling rigorous statistical comparisons across regions, conditions, and subjects. These methods advance the accuracy and interpretability of neural timescale maps, moving beyond point estimates to incorporate uncertainty for inference and testing. This work lays the methodological foundation for future research into the role of timescales in brain structure and function.

Appendix: Autocorrelation-Domain Score Process

This appendix derives the score process u_t from equation (16). We omit scaling factors (e.g., $1/K$), as they cancel out in the sandwich variance ratio $q_{\text{AD}}^{-1}\omega_{\text{AD}}q_{\text{AD}}^{-1}$.

The pseudo-true parameter ϕ_{AD}^* minimizes the population objective function $S_K(\phi_{\text{AD}})$ (5), satisfying the first-order condition:

$$\sum_{k=1}^{K} \psi_k(\phi_{\text{AD}}^*)(\rho_k - \phi_{\text{AD}}^{*k}) = 0, \quad \text{where } \psi_k(\phi_{\text{AD}}^*) = \frac{d}{d\phi_{\text{AD}}}(\phi_{\text{AD}}^k)\Big|_{\phi_{\text{AD}}=\phi_{\text{AD}}^*} = k\phi_{\text{AD}}^{*k-1}. \tag{22}$$

Using $\rho_k = \mathbb{E}[X_t X_{t-k}]/\gamma_0$ (1), we rewrite the residual term as an expectation:

$$\rho_k - \phi_{\text{AD}}^{*k} = \mathbb{E}\left[\frac{X_t X_{t-k} - \phi_{\text{AD}}^{*k} X_t^2}{\gamma_0}\right]. \tag{23}$$

Substituting (23) into (22) defines the score process u_t:

$$\mathbb{E}\left[\frac{1}{\gamma_0}\sum_{k=1}^{K} \psi_k(\phi_{\text{AD}}^*)(X_t X_{t-k} - \phi_{\text{AD}}^{*k} X_t^2)\right] = \mathbb{E}[u_t] = 0.$$

This matches the definition in equation (16). The asymptotic variance is then determined by the long-run variance of this score, $\omega_{\text{AD}} = \sum_{\ell=-\infty}^{\infty} \mathbb{E}[u_t u_{t-\ell}]$.

The sample score process $\hat{u}_t$ used for variance estimation is obtained by replacing population expectations with sample moments and the parameter ϕ_{AD}^* with the estimator $\hat{\phi}_{\text{AD}}^*$:

$$\hat{u}_t = \frac{1}{\hat{\gamma}_0}\sum_{k=1}^{K} k(\hat{\phi}_{\text{AD}}^*)^{k-1}(x_t x_{t-k} - (\hat{\phi}_{\text{AD}}^*)^k x_t^2).$$

Code and Data Availability

All simulation results and fMRI timescale maps, inclusive of the code by which they were derived, can be accessed on github.com/griegner/fmri-timescales. The code is under the open source MIT license, allowing access and reuse with attribution. The Human Connectome Project young adult dataset (ages 22-35; 2018 release) used in this study is publicly accessible under a data usage agreement, which describes specific terms for data use and sharing.

Author Contributions

Gabriel Riegner: Conceptualization, Methodology, Software, Validation, Formal Analysis, Writing - Original Draft **Samuel Davenport**: Conceptualization, Methodology, Formal Analysis, Writing - Review & Editing, Supervision **Bradley Voytek**: Conceptualization, Writing - Review & Editing, Supervision **Armin Schwartzman**: Conceptualization, Methodology, Formal Analysis, Writing - Review & Editing, Supervision, Project Administration, Funding Acquisition.

Disclosure Statement

The authors declare no conflicts of interest.

References

Bailes, S. M., Gomez, D. E., Setzer, B., and Lewis, L. D. (2023). Resting-state fMRI signals contain spectral signatures of local hemodynamic response timing. *eLife*, 12:e86453.

Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience*, 7(5):446–451.

Cirillo, R., Fascianelli, V., Ferrucci, L., and Genovesio, A. (2018). Neural Intrinsic Timescales in the Macaque Dorsal Premotor Cortex Predict the Strength of Spatial Response Coding. *iScience*, 10:203–210.

Davenport, S. (2025). StatBrainz: A MATLAB package for analysing Brain Imaging data using statistics.

Donoghue, T., Haller, M., Peterson, E. J., Varma, P., Sebastian, P., Gao, R., Noto, T., Lara, A. H., Wallis, J. D., Knight, R. T., Shestyuk, A., and Voytek, B. (2020). Parameterizing neural power spectra into periodic and aperiodic components. *Nature Neuroscience*, 23(12):1655–1665.

Du, J., DiNicola, L. M., Angeli, P. A., Saadon-Grosman, N., Sun, W., Kaiser, S., Ladopoulou, J., Xue, A., Yeo, B. T. T., Eldaief, M. C., and Buckner, R. L. (2024). Organization of the human cerebral cortex estimated within individuals: networks, global topography, and function. *Journal of Neurophysiology*, 131(6):1014–1082.

Fallon, J., Ward, P. G. D., Parkes, L., Oldham, S., Arnatkevičiūtė, A., Fornito, A., and Fulcher, B. D. (2020). Timescales of spontaneous fMRI fluctuations relate to structural connectivity in the brain. *Network Neuroscience*, 4(3):788–806.

Frederick, B. (2025). Rapidtide: a suite of Python programs used to model, characterize, visualize, and remove time varying, physiological blood signals from fMRI and fNIRS datasets.

Gao, R., van den Brink, R. L., Pfeffer, T., and Voytek, B. (2020). Neuronal timescales are functionally dynamic and shaped by cortical microarchitecture. *eLife*, 9:e61277.

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Van Essen, D. C., and Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124.

Goldberg, A., Rosario, I., Power, J., Horga, G., and Wengler, K. (2024). Strategies for motion- and respiration-robust estimation of fMRI intrinsic neural timescales. *Imaging Neuroscience*, 2:imag–2–00326.

Hansen, B. E. (2022). *Econometrics*. Princeton University Press, Princeton.

Hasson, U., Yang, E., Vallines, I., Heeger, D. J., and Rubin, N. (2008). A Hierarchy of Temporal Receptive Windows in Human Cortex. *The Journal of Neuroscience*, 28(10):2539–2550.

He, B. J. (2011). Scale-Free Properties of the Functional Magnetic Resonance Imaging Signal during Rest and Task. *The Journal of Neuroscience*, 31(39):13786–13795.

He, B. J., Zempel, J. M., Snyder, A. Z., and Raichle, M. E. (2010). The Temporal Structures and Functional Significance of Scale-free Brain Activity. *Neuron*, 66(3):353–369.

Huang, Z., Liu, X., Mashour, G. A., and Hudetz, A. G. (2018). Timescales of Intrinsic BOLD Signal Dynamics and Functional Connectivity in Pharmacologic and Neuropathologic States of Unconsciousness. *The Journal of Neuroscience*, 38(9):2304–2317.

Ito, T., Hearne, L. J., and Cole, M. W. (2020). A cortical hierarchy of localized and distributed processes revealed via dissociation of task activations, connectivity changes, and intrinsic timescales. *NeuroImage*, 221:117141.

Jadon, A., Patil, A., and Jadon, S. (2024). A Comprehensive Survey of Regression Based Loss Functions for Time Series Forecasting. *Data Management, Analytics and Innovation*, 2.

Kaneoke, Y., Donishi, T., Iwatani, J., Ukai, S., Shinosaki, K., and Terada, M. (2012). Variance and Auto-correlation of the Spontaneous Slow Brain Activity. *PLoS ONE*, 7(5):e38131.

Korponay, C., Janes, A. C., and Frederick, B. B. (2024). Brain-wide functional connectivity artifactually inflates throughout functional magnetic resonance imaging scans. *Nature Human Behaviour*, 8(8):1568–1580.

Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168.

Li, S. and Wang, X.-J. (2022). Hierarchical timescales in the neocortex: Mathematical mechanism and biological insights. *Proceedings of the National Academy of Sciences*, 119(6):e2110274119.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197):869–878.

Lurie, D. J., Pappas, I., and D'Esposito, M. (2024). Cortical timescales and the modular organization of structural and functional brain networks. *Human Brain Mapping*, 45(2):e26587.

Maisson, D. J.-N., Cash-Padgett, T. V., Wang, M. Z., Hayden, B. Y., Heilbronner, S. R., and Zimmermann, J. (2021). Choice-relevant information transformation along a ventrodorsal axis in the medial prefrontal cortex. *Nature Communications*, 12(1):4830.

Manea, A. M., Zilverstand, A., Ugurbil, K., Heilbronner, S. R., and Zimmermann, J. (2022). Intrinsic timescales as an organizational principle of neural processing across the whole rhesus macaque brain. *eLife*, 11:e75540.

Manea, A. M. G., Maisson, D. J.-N., Voloh, B., Zilverstand, A., Hayden, B., and Zimmermann, J. (2024). Neural timescales reflect behavioral demands in freely moving rhesus macaques. *Nature Communications*, 15(1):2151.

Marriott and Pope (1954). Bias in the estimation of autocorrelations. *Biometrika*, 41(3/4):390–402.

Martin-Burgos, B., McPherson, T. S., Hammonds, R., Gao, R., Muotri, A. R., and Voytek, B. (2024). Development of neuronal timescales in human cortical organoids and rat hippocampus dissociated cultures. *Journal of Neurophysiology*, 132(3):757–764.

Meisel, C., Bailey, K., Achermann, P., and Plenz, D. (2017). Decline of long-range temporal correlations in the human brain during sustained wakefulness. *Scientific Reports*, 7(1):11825.

Mitra, A., Snyder, A. Z., Hacker, C. D., and Raichle, M. E. (2014). Lag structure in resting-state fMRI. *Journal of Neurophysiology*, 111(11):2374–2391.

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., and Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, 17(12):1661–1663.

Müller, E. J., Munn, B., Hearne, L. J., Smith, J. B., Fulcher, B., Arnatkevičiūtė, A., Lurie, D. J., Cocchi, L., and Shine, J. M. (2020). Core and matrix thalamic sub-populations relate to spatio-temporal cortical connectivity gradients. *NeuroImage*, 222:117224.

Newey, W. and West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelationconsistent covariance matrix. *Econometrica*, 55(3):703–08.

Nougaret, S., Fascianelli, V., Ravel, S., and Genovesio, A. (2021). Intrinsic timescales across the basal ganglia. *Scientific Reports*, 11(1):21395.

Rangaprakash, D., Wu, G., Marinazzo, D., Hu, X., and Deshpande, G. (2018). Hemodynamic response function (HRF) variability confounds resting-state fMRI functional connectivity. *Magnetic Resonance in Medicine*, 80(4):1697–1713.

Raut, R. V., Mitra, A., Snyder, A. Z., and Raichle, M. E. (2019). On time delay estimation and sampling error in resting-state fMRI. *NeuroImage*, 194:211–227.

Raut, R. V., Snyder, A. Z., and Raichle, M. E. (2020). Hierarchical dynamics as a macroscopic organizing principle of the human brain. *Proceedings of the National Academy of Sciences*, 117(34):20890–20897.

Rossi-Pool, R., Zainos, A., Alvarez, M., Parra, S., Zizumbo, J., and Romo, R. (2021). Invariant timescale hierarchy across the cortical somatosensory network. *Proceedings of the National Academy of Sciences*, 118(3):e2021843118.

Runyan, C. A., Piasini, E., Panzeri, S., and Harvey, C. D. (2017). Distinct timescales of population coding across cortex. *Nature*, 548(7665):92–96.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., and Smith, S. M. (2014). Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage*, 90:449–468.

Shafiei, G., Fulcher, B. D., Voytek, B., Satterthwaite, T. D., Baillet, S., and Misic, B. (2023). Neurophysiological signatures of cortical micro-architecture. *Nature Communications*, 14(1):6000.

Shafiei, G., Markello, R. D., Vos De Wael, R., Bernhardt, B. C., Fulcher, B. D., and Misic, B. (2020). Topographic gradients of intrinsic dynamics across neocortex. *eLife*, 9:e62116.

Shinn, M., Hu, A., Turner, L., Noble, S., Preller, K. H., Ji, J. L., Moujaes, F., Achard, S., Scheinost, D., Constable, R. T., Krystal, J. H., Vollenweider, F. X., Lee, D., Anticevic, A., Bullmore, E. T., and Murray, J. D. (2023). Functional brain networks reflect spatial and temporal autocorrelation. *Nature Neuroscience*, pages 1–12. Publisher: Nature Publishing Group.

Spitmaan, M., Seo, H., Lee, D., and Soltani, A. (2020). Multiple timescales of neural dynamics and integration of task-relevant signals across cortex. *Proceedings of the National Academy of Sciences*, 117(36):22522–22531.

Stephens, G. J., Honey, C. J., and Hasson, U. (2013). A place for time: the spatiotemporal structure of neural dynamics during natural audition. *Journal of Neurophysiology*, 110(9):2019–2026.

Tong, Y., Hocke, L. M., and Frederick, B. B. (2019). Low Frequency Systemic Hemodynamic "Noise" in Resting State BOLD fMRI: Characteristics, Causes, Implications, Mitigation Strategies, and Applications. *Frontiers in Neuroscience*, 13:787.

Trepka, E., Spitmaan, M., Qi, X.-L., Constantinidis, C., and Soltani, A. (2024). Training-Dependent Gradients of Timescales of Neural Dynamics in the Primate Prefrontal Cortex and Their Contributions to Working Memory. *The Journal of Neuroscience*, 44(2):e2442212023.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., and Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80:62–79.

Wasmuht, D. F., Spaak, E., Buschman, T. J., Miller, E. K., and Stokes, M. G. (2018). Intrinsic neuronal dynamics predict distinct functional roles during working memory. *Nature Communications*, 9(1):3499.

Watanabe, T., Rees, G., and Masuda, N. (2019). Atypical intrinsic neural timescale in autism. *eLife*, 8:e42256.

Wengler, K., Goldberg, A. T., Chahine, G., and Horga, G. (2020). Distinct hierarchical alterations of intrinsic neural timescales account for different manifestations of psychosis. *eLife*, 9:e56151.

Woolrich, M. W., Ripley, B. D., Brady, M., and Smith, S. M. (2001). Temporal Autocorrelation in Univariate Linear Modeling of FMRI Data. *NeuroImage*, 14(6):1370–1386.

Wu, G.-R., Colenbier, N., Van Den Bossche, S., Clauw, K., Johri, A., Tandon, M., and Marinazzo, D. (2021). rsHRF: A toolbox for resting-state HRF estimation and deconvolution. *NeuroImage*, 244:118591.

Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fischl, B., Liu, H., and Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165.

Zeraati, R., Engel, T. A., and Levina, A. (2022). A flexible Bayesian framework for unbiased estimation of timescales. *Nature Computational Science*, 2(3):193–204.

Zilio, F., Gomez-Pilar, J., Cao, S., Zhang, J., Zang, D., Qi, Z., Tan, J., Hiromi, T., Wu, X., Fogel, S., Huang, Z., Hohmann, M. R., Fomina, T., Synofzik, M., Grosse-Wentrup, M., Owen, A. M., and Northoff, G. (2021). Are intrinsic neural timescales related to sensory processing? Evidence from abnormal behavioral states. *NeuroImage*, 226:117579.