

BioReg-Analysis

Participant-specific Laplacian EEG dynamic-factor analysis

Complete methods, equations, terminology, code traceability, and all accepted fits

Technical report v13

8 September 2026

Scientific status. This report documents the 15-participant, three-band, participant-specific *capacity* analysis made from surface-Laplacian EEG measurements. All 45 participant-band fits converged. The report deliberately does not call this participant-held-out validation: each model is fitted to the same participant trajectory that it describes. Participant p21 is excluded before fitting by a frozen, model-independent signal-quality rule.

Repository used: `/Users/veronique/BASE/Code/BioReg-Analysis`

Accepted result: `output/development-individual-laplacian-excluding-p21-k5-phase-closure-v13`

Contents

1	Purpose, scope, and the result being documented	1
1.1	Purpose	1
1.2	What is established, and what is not	1
1.3	Three analysis layers	2
2	Dataset, participants, phases, and cycles	3
2.1	Participants and retained electrode dimensions	3
2.2	What a phase means	3
2.3	What a cycle means	4
2.4	Why p21 is excluded	4
3	EEG measurement construction	5
3.1	Frequency bands and physical windows	5
3.2	Surface-Laplacian reference	5
3.3	Analytic band signal and normalized complex state	5
3.4	Window complex-state matrix	6
3.5	Absolute band activity	6
3.6	Entropy quantities	7
4	Geometry-preserving representation of the full matrix	8
4.1	Why raw matrix entries are not regressed	8
4.2	Regularized matrix logarithm	8
4.3	Hermitian vectorization	8
4.4	Baseline centering and observation vector	9
5	Participant-specific dynamic-factor model	10
5.1	Separate band models	10
5.2	Observation equation	10
5.3	Transition equation	10
5.4	Disturbance and recovery drive	11
5.5	Missing windows	12
6	Fitting method and objective function	13
6.1	What is fitted	13
6.2	Kalman prediction and update	13
6.3	The FCN	14
6.4	Generalized expectation-maximization	14
6.5	Parameter errors	14

7	From a predicted matrix to plotted entropy	15
7.1	Nonlinear uncertainty propagation	15
7.2	Phase-specific empirical residual distribution	15
7.3	Explicit entropy mean closure	15
8	Diagnostics and how to read the plots	17
8.1	Normalized residual	17
8.2	Chi-square and nominal degrees of freedom	17
8.3	Conditional versus free-running	17
8.4	Line and panel key	18
9	Accepted participant-specific result	19
9.1	Numerical summary	19
9.2	Interpretation	19
9.3	Scientific stopping and advancement conditions	20
10	Reproduction and code traceability	21
10.1	Exact input and output artifacts	21
10.2	Reproduction command	21
10.3	Code map	22
11	Glossary of technical terms	23
A	All conditional participant-band fit plots	26
B	Protocol constants used in this result	72

Chapter 1

Purpose, scope, and the result being documented

1.1 Purpose

The analysis asks a concrete model-capacity question: for one participant and one EEG frequency band, can a low-dimensional stochastic dynamical system describe the time evolution of (i) absolute band activity and (ii) the complete complex spatial EEG matrix? Separate models are fitted for theta, alpha, and beta. There is no cross-band coupling.

The data source is the *Neurophysiological Dataset of Stress Resilience During Human-Computer Interaction*, version 1.0.0, distributed by PhysioNet[1]. BioReg-EEG is the authoritative Java measurement engine. BioReg-Analysis consumes the resulting audited matrices and performs the exploratory model fit. Python does not recompute the scientific quantities in this workflow.

1.2 What is established, and what is not

The accepted result establishes **descriptive participant-specific capacity**: each participant-band model sees all unique accepted baseline, disturbance, and recovery observations from that participant, and the fitted one-step distribution describes those observations with a coordinate-space reduced chi-square close to one. Across the independently reported entropy quantities S_{vN} and S_{diag} , 33,758 of 35,418 fitted points (95.313%) lie within two fitted residual standard deviations. Every one of the 270 participant-by-band-by-phase-by-entropy groups has at least 90% coverage; the minimum is 91.124%, tied by p15 alpha disturbance S_{vN} and p15 theta disturbance S_{diag} .

This result does *not* yet establish any of the following:

- prediction of a participant who was not used for fitting;
- causal effects of the disturbance;
- a universal biological interpretation of the latent factors;
- equivalence to, or calibration of, the IRAM equations;
- independent test-retest reliability of Cycle 1 versus Cycle 2; or
- that the free-running trajectory is accurate simply because its high-dimensional aggregate chi-square is near one.

The exact participant-specific latent axes differ between fits. Raw entries of the latent matrices therefore must not be compared across participants. Comparable outputs are reconstructed EEG matrices, the derived entropies, residuals, and fit diagnostics.

1.3 Three analysis layers

1. **Measurement layer (BioReg-EEG).** Filter voltages, apply the surface Laplacian, construct complex analytic states, form the matrix ρ , and export activity, covariance, coherence, and entropy summaries.
2. **Model layer (BioReg-Analysis).** Convert ρ into geometry-preserving real coordinates, fit one continuous stochastic state-space model for each participant and band, and reconstruct predicted matrices.
3. **Report layer (BioReg-Analysis).** Propagate fitted coordinate uncertainty through nonlinear entropy functions, apply the explicit phase-specific descriptive closure, calculate residuals and chi-square, and draw the plots in this report.

Chapter 2

Dataset, participants, phases, and cycles

2.1 Participants and retained electrode dimensions

The present run contains 15 participants: p01, p02, p03, p04, p10, p11, p12, p13, p15, p16, p17, p18, p19, p20, and p22. Each participant keeps the electrodes that survived the frozen BioReg-EEG quality procedure. Consequently the matrix dimension d_i is the retained electrode count for participant i , not a fixed value of 16 and not the latent dimension k .

Table 2.1: Participant dimensions and fit diagnostics. Conditional and free-running columns give the minimum–maximum χ^2/NDF across theta, alpha, and beta.

Participant	d_i	fit points	missing	max iter.	conditional range	free-run range
p01	17	423	30	60	1.00311–1.00313	0.99919–1.00135
p02	28	438	5	17	1.00256–1.00259	1.00215–1.00249
p03	27	449	6	28	1.00250–1.00254	1.00107–1.00242
p04	29	237	216	45	1.00481–1.00543	1.00388–1.00461
p10	30	238	215	32	1.00461–1.00482	1.00445–1.00462
p11	32	369	71	22	1.00294–1.00299	1.00247–1.00270
p12	32	435	19	32	1.00256–1.00260	1.00135–1.00168
p13	32	409	46	18	1.00267–1.00273	1.00212–1.00259
p15	29	427	28	25	1.00259–1.00262	1.00156–1.00219
p16	32	260	9	27	1.00421–1.00440	1.00321–1.00415
p17	29	436	19	20	1.00265–1.00336	1.00162–1.00243
p18	31	444	10	24	1.00247–1.00252	1.00200–1.00239
p19	31	451	4	24	1.00247–1.00270	1.00162–1.00225
p20	29	440	14	31	1.00250–1.00252	1.00150–1.00226
p22	30	447	8	16	1.00252–1.00281	1.00193–1.00226

2.2 What a phase means

A **phase** is an experimentally labelled interval, not an EEG phase angle. The three labels are:
Baseline.

The pre-disturbance operating interval. It defines the participant-specific reference level and the initial latent operating point.

Disturbance.

The interval in which the experimental stressor or task disturbance is present.

Recovery.

The interval after removal of the disturbance. It tests the shape and timescale of return, persistence, or overshoot.

Within each phase, $s \in [0, 1]$ is the **phase position**: $s = 0$ means the beginning and $s = 1$ means the end. The current recordings contain phases of approximately 360 seconds each, but the code reads timing from each artifact instead of hard-coding that duration.

2.3 What a cycle means

A **cycle** is one ordered baseline–disturbance–recovery pass through the experimental protocol. It is not one period of a theta, alpha, or beta oscillation. It is also not a separate visit or an independent repeat measurement.

In this dataset the two stress cycles occur sequentially inside one recording. Cycle 2 follows Cycle 1 and can therefore contain habituation, sensitization, fatigue, or carryover. Most importantly, the source interval labelled Cycle-2 baseline repeats the same physical samples already labelled Cycle-1 recovery. Counting both would give the same EEG bytes two likelihood contributions.

The implemented **unique-cycle stitching** rule is therefore:

$$\mathcal{T}_i = [\text{Cycle 1 baseline, disturbance, recovery}] \parallel [\text{Cycle 2 disturbance, recovery}], \quad (2.1)$$

where $\parallel$ means chronological concatenation. Cycle-2 baseline is removed, and the latent state is not reset at the Cycle-2 boundary. No accepted EEG value is averaged or imputed during this operation. The implementation is `StateSpaceObservationBuilder.stitchUniqueCycles()` at line 349.

2.4 Why p21 is excluded

Participant p21 is excluded as a complete participant before any participant-band grouping, fitting, or residual calculation. A frozen model-independent V7 audit found sustained covariance-rank degeneracy in 40.255% of Cycle 2 and 96.316% of its recovery. Only seven recovery windows remained before inherited artifact screening, below the frozen minimum of 24. Cycle 1 is also excluded so that p21 cannot become a partial participant whose membership depends on which phase is being analyzed. The declaration and evidence are stored in

`protocol/participant-fit-quality-exclusions-v1.json`.

This is a quality exclusion, not a residual-based outlier removal.

Chapter 3

EEG measurement construction

3.1 Frequency bands and physical windows

The analysis uses three conventional bands independently:

Band	Frequency range	Model status
Theta	4–8 Hz	independent parameter set
Alpha	8–13 Hz	independent parameter set
Beta	13–30 Hz	independent parameter set

The current files have sampling rate $f_s \simeq 250$ Hz. A candidate window contains 1,000 samples (about 4 seconds), with a candidate step of 500 samples (about 2 seconds). The state-space fit prevents overlapping-window leakage by keeping every second candidate, so fitted grid points are about 4 seconds apart and do not share EEG samples. These numbers are metadata, not universal constants.

3.2 Surface-Laplacian reference

The **surface Laplacian** is a spatial high-pass reference. For electrode j with the four nearest retained electrodes $\mathcal{N}_4(j)$, the corrected voltage is

$$x_j^{\text{Lap}}(t) = x_j(t) - \frac{1}{4} \sum_{\ell \in \mathcal{N}_4(j)} x_\ell(t). \quad (3.1)$$

Here $x_j(t)$ is the as-recorded voltage at time sample t . The operation attenuates activity common to the head reference and emphasizes local curvature of the scalp field. It does not select a subset of electrodes for the fit; all retained electrodes remain. The nearest neighbors are fixed from montage coordinates. Current modeling fails closed if an as-recorded artifact is supplied. The method is implemented in `BioReg-EEG SurfaceLaplacian.java`.

3.3 Analytic band signal and normalized complex state

For each retained electrode j and band b , BioReg-EEG creates the complex analytic signal

$$z_j^{(b)}(t) = x_j^{(b)}(t) + i \mathcal{H}[x_j^{(b)}](t), \quad (3.2)$$

where $x_j^{(b)}$ is the band-limited real voltage, $i = \sqrt{-1}$, and $\mathcal{H}$ is the Hilbert transform, which supplies the quadrature (90-degree shifted) component. The implementation uses a frequency-domain analytic construction in `Fourier.analytic()`.

At one sample the electrode vector is normalized:

$$\psi_j^{(b)}(t) = \frac{z_j^{(b)}(t)}{\sqrt{\sum_{\ell=1}^{d_i} |z_\ell^{(b)}(t)|^2}}, \quad \sum_{j=1}^{d_i} |\psi_j^{(b)}(t)|^2 = 1. \quad (3.3)$$

The vector $\psi^{(b)}(t) \in \mathbb{C}^{d_i}$ represents relative amplitude and phase across electrodes at that instant. Its normalization deliberately removes overall band amplitude, which is retained separately below.

3.4 Window complex-state matrix

For accepted window w with N_w finite samples,

$$\rho_w^{(b)} = \frac{1}{N_w} \sum_{t \in w} |\psi^{(b)}(t)\rangle \langle \psi^{(b)}(t)|, \quad (3.4)$$

or, entry by entry,

$$\rho_{w,j\ell}^{(b)} = \frac{1}{N_w} \sum_{t \in w} \psi_j^{(b)}(t) \psi_\ell^{(b)*}(t). \quad (3.5)$$

The star means complex conjugation. The ket–bra notation $|\psi\rangle\langle\psi|$ is simply the outer product $\psi\psi^\dagger$. The dagger denotes conjugate transpose. By construction, ρ_w is:

- **Hermitian:** $\rho_w = \rho_w^\dagger$;
- **positive semidefinite:** $\mathbf{v}^\dagger \rho_w \mathbf{v} \geq 0$ for every $\mathbf{v}$;
- **trace one:** $\text{Tr } \rho_w = 1$; and
- generally mixed because it averages changing complex spatial states within the window.

This matrix is not the conventional voltage covariance. It is an amplitude-normalized complex spatial-state matrix. Construction and entropy reduction are in BioReg-EEG `StateSeries.windowState()` at line 197.

3.5 Absolute band activity

The conventional real-voltage covariance in window w is

$$C_{w,j\ell}^{(b)} = \frac{1}{N_w - 1} \sum_{t \in w} \left(x_j^{(b)}(t) - \bar{x}_{w,j}^{(b)} \right) \left(x_\ell^{(b)}(t) - \bar{x}_{w,\ell}^{(b)} \right), \quad (3.6)$$

where $\bar{x}_{w,j}^{(b)}$ is the within-window mean. The total absolute activity is its trace, and the fitted scalar is

$$a_w^{(b)} = \log \left(\max \left[\epsilon_a, \text{Tr } C_w^{(b)} \right] \right), \quad \epsilon_a = 10^{-12}. \quad (3.7)$$

The natural logarithm reduces positive skew. The plotted quantity is baseline-centered log activity, $a_w - \bar{a}_B$. The floor ϵ_a prevents log 0 and is a fixed numerical constant, not a physiological parameter.

3.6 Entropy quantities

Let λ_r be the eigenvalues of ρ_w . The **von Neumann entropy** is

$$S_{\text{vN}}(w) = -\text{Tr}(\rho_w \log \rho_w) = -\sum_r \lambda_r \log \lambda_r. \quad (3.8)$$

It measures spectral mixing of the complex-state matrix. A rank-one matrix has $S_{\text{vN}} = 0$; a more mixed spectrum has larger entropy. It is computed from the window matrix first. Averaging instantaneous entropy would be a different and generally incorrect operation because entropy is nonlinear.

The **diagonal entropy** in the electrode basis is

$$S_{\text{diag}}(w) = -\sum_{j=1}^{d_i} \rho_{w,jj} \log \rho_{w,jj}. \quad (3.9)$$

It measures how broadly the normalized activity is distributed over electrode labels. Unlike S_{vN} , it depends on the chosen electrode basis and on d_i .

The reported **entropy gap** is

$$G(w) = S_{\text{diag}}(w) - S_{\text{vN}}(w). \quad (3.10)$$

The gap is derived exactly from the two entropies. It is never fitted as a third independent datum. Calling it a “coherence gap” is shorthand within this project; it should not be confused with magnitude-squared spectral coherence. A change in G must be interpreted by inspecting S_{vN} and S_{diag} separately.

Chapter 4

Geometry-preserving representation of the full matrix

4.1 Why raw matrix entries are not regressed

Ordinary regression on the real and imaginary entries of ρ can predict a non-Hermitian matrix, a negative eigenvalue, or a trace different from one. The model instead maps each density matrix to an unconstrained real coordinate vector and maps predictions back through a matrix exponential.

4.2 Regularized matrix logarithm

First mix a fixed amount of the identity:

$$\rho_\epsilon = (1 - \epsilon)\rho + \epsilon \frac{I_d}{d}, \quad \epsilon = 10^{-5}. \quad (4.1)$$

I_d is the $d \times d$ identity matrix. This makes every eigenvalue strictly positive so that $\log \rho_\epsilon$ exists. It is a numerical regularizer, not fitted biology.

Remove the unidentifiable scalar trace component:

$$M = \log \rho_\epsilon - \frac{\text{Tr}(\log \rho_\epsilon)}{d} I_d. \quad (4.2)$$

M is Hermitian and traceless. The matrix logarithm is defined spectrally: if $\rho_\epsilon = V \text{diag}(\lambda_r) V^\dagger$, then $\log \rho_\epsilon = V \text{diag}(\log \lambda_r) V^\dagger$.

4.3 Hermitian vectorization

The map $\text{hvec}(M)$ contains $d - 1$ orthonormal diagonal coordinates and, for every $j < \ell$, the two coordinates

$$\sqrt{2} \text{Re } M_{j\ell}, \quad \sqrt{2} \text{Im } M_{j\ell}. \quad (4.3)$$

There are $d^2 - 1$ real coordinates. The factors $\sqrt{2}$ preserve the Frobenius inner product, meaning

$$\langle M_1, M_2 \rangle_F = \text{Tr}(M_1^\dagger M_2) = \text{hvec}(M_1)^T \text{hvec}(M_2). \quad (4.4)$$

The **Frobenius norm** is the square root of the sum of squared magnitudes of all matrix entries. “Frobenius mass” informally means the corresponding squared sum.

The inverse map constructs a traceless Hermitian $\widehat{M}$ and returns

$$\hat{\rho} = \frac{\exp(\widehat{M})}{\text{Tr}[\exp(\widehat{M})]}. \quad (4.5)$$

This guarantees a Hermitian, positive, trace-one prediction.

Implementation: `DensityMatrixGeometry.java`, lines 43 and 67.

4.4 Baseline centering and observation vector

For participant i , band b , and the initial Cycle-1 baseline set B_i ,

$$\bar{a}_{iB}^{(b)} = \frac{1}{|B_i|} \sum_{w \in B_i} a_{iw}^{(b)}, \quad (4.6)$$

$$\bar{\mathbf{m}}_{iB}^{(b)} = \frac{1}{|B_i|} \sum_{w \in B_i} \text{hvec}(M_{iw}^{(b)}). \quad (4.7)$$

The fitted observation vector is

$$\mathbf{y}_{it}^{(b)} = \begin{bmatrix} a_{it}^{(b)} - \bar{a}_{iB}^{(b)} \\ \text{hvec}(M_{it}^{(b)}) - \bar{\mathbf{m}}_{iB}^{(b)} \end{bmatrix} \in \mathbb{R}^{d_i^2}. \quad (4.8)$$

The first coordinate carries absolute activity; the remaining $d_i^2 - 1$ coordinates carry the complete normalized complex matrix. The baseline origin is learned only from Cycle-1 baseline. Its implementation is `StateSpaceObservationBuilder.java`, line 495.

Chapter 5

Participant-specific dynamic-factor model

5.1 Separate band models

For each participant i and each $b \in \{\theta, \alpha, \beta\}$, a distinct model is fitted. No parameter is shared across bands or participants in the present capacity analysis. The latent dimension is fixed at

$$k = 5. \tag{5.1}$$

Here k is the number of latent dynamic factors, not the number of electrodes. The observation dimension is $m_i = d_i^2$, ranging from 289 for p01 ($d = 17$) to 1,024 for participants with $d = 32$.

5.2 Observation equation

$$\mathbf{y}_t = W\mathbf{z}_t + \boldsymbol{\nu}_t, \quad \boldsymbol{\nu}_t \sim \mathcal{N}(\mathbf{0}, R). \tag{5.2}$$

Definitions:

$\mathbf{y}_t \in \mathbb{R}^m$

Measured baseline-centered log activity plus full matrix coordinates at time point t .

$\mathbf{z}_t \in \mathbb{R}^k$

Unobserved, or latent, dynamic-factor state.

$W \in \mathbb{R}^{m \times k}$

Loading matrix that maps latent factors to EEG observation coordinates.

$\boldsymbol{\nu}_t$ Observation noise: matrix structure not represented by the k factors plus measurement variability.

R Observation-noise covariance. The implemented R is diagonal because estimating a dense $m \times m$ covariance is not supported by the available time points.

$\mathcal{N}(\mathbf{0}, R)$

Multivariate normal distribution with zero mean and covariance R .

The first loading column is reserved for scaled log activity. The remaining $k - 1 = 4$ columns are the leading principal directions of the matrix-coordinate block. Principal component analysis (PCA) is an orthogonal decomposition that orders directions by observed variance[6]. Each direction receives a deterministic sign so equivalent sign flips do not change artifact labeling. The fixed loading construction is in `StateSpaceFitter.fixedLoading()` near line 685.

5.3 Transition equation

The participant-specific state evolves as

$$\mathbf{z}_t = \mathbf{m}_i + A(\mathbf{z}_{t-1} - \mathbf{m}_i) + D\mathbf{u}_t + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, Q). \quad (5.3)$$

Definitions:

$\mathbf{m}_i \in \mathbb{R}^k$

Participant's stationary latent baseline state, estimated from Cycle-1 baseline projected scores.

$A \in \mathbb{R}^{k \times k}$

Transition matrix. It determines persistence, cross-factor influence, and return toward $\mathbf{m}_i$.

$D \in \mathbb{R}^{k \times 6}$

Phase-drive coefficient matrix.

$\mathbf{u}_t \in \mathbb{R}^6$

Known basis vector determined by experimental phase and within-phase position.

$\boldsymbol{\eta}_t$ Process noise, meaning stochastic state evolution not explained by A and $D\mathbf{u}_t$.

$Q \in \mathbb{R}^{k \times k}$

Process-noise covariance.

When $\mathbf{u}_t = \mathbf{0}$, the deterministic fixed point is $\mathbf{m}_i$. Stability is enforced by

$$\rho(A) \leq 1 - \delta = 0.999, \quad (5.4)$$

where $\rho(A)$ is the spectral radius (largest absolute eigenvalue) and $\delta = 0.001$ is the protocol stability margin. A radius below one makes old deviations decay in the absence of drive.

5.4 Disturbance and recovery drive

For baseline, $\mathbf{u}(s) = \mathbf{0}$. For disturbance,

$$\mathbf{u}_D(s) = \begin{bmatrix} (1-s)^3 \\ 3s(1-s)^2 \\ 3s^2(1-s) \\ s^3 \\ 0 \\ 0 \end{bmatrix}. \quad (5.5)$$

The first four entries are the cubic Bernstein basis[7]. Their fitted linear combination can represent a smooth cubic response without fixing its sign, peak, or monotonicity in advance.

For recovery,

$$\mathbf{u}_R(s) = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \exp(-s/\lambda_f) \\ \exp(-s/\lambda_s) \end{bmatrix}, \quad \lambda_f = 0.15, \quad \lambda_s = 0.60. \quad (5.6)$$

λ_f and λ_s are fractions of the recovery duration. If a recovery phase lasts T_R seconds, the physical time constants are $\tau_f = 0.15T_R$ and $\tau_s = 0.60T_R$. The exponentials supply fast and slow decaying components; the entries of D determine their direction and magnitude. The implementation is `RecoveryDrive.java`.

5.5 Missing windows

Rejected artifact windows remain explicit grid points with no observation. The transition equation propagates across them, but they contribute no measurement update and no fit score. A run longer than eight consecutive missing model steps is rejected by the protocol. This preserves elapsed physical time and prevents the false lag-one pairing that would result from deleting gaps.

Chapter 6

Fitting method and objective function

6.1 What is fitted

For each participant-band, the current capacity routine jointly estimates

$$\Theta = \{A, D, Q, R, W, \mathbf{m}_i\}. \quad (6.1)$$

W begins in a deterministic EEG-defined activity-plus-PCA basis and remains fixed during the generalized expectation-maximization iterations. A , D , Q , and diagonal R are updated. The participant baseline is treated as a fixed effect. A between-participant covariance Ω cannot be identified from one participant, so it is held at the numerical covariance floor and is explicitly reported as not estimated.

6.2 Kalman prediction and update

The Kalman filter[2] recursively computes the Gaussian one-step prediction. Given filtered state mean $\hat{\mathbf{z}}_{t-1|t-1}$ and covariance $P_{t-1|t-1}$,

$$\hat{\mathbf{z}}_{t|t-1} = \mathbf{m}_i + A(\hat{\mathbf{z}}_{t-1|t-1} - \mathbf{m}_i) + D\mathbf{u}_t, \quad (6.2)$$

$$P_{t|t-1} = AP_{t-1|t-1}A^T + Q, \quad (6.3)$$

$$\hat{\mathbf{y}}_{t|t-1} = W\hat{\mathbf{z}}_{t|t-1}, \quad (6.4)$$

$$S_t = WP_{t|t-1}W^T + R. \quad (6.5)$$

The innovation, or new prediction error, is

$$\mathbf{e}_t = \mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}. \quad (6.6)$$

When an observation is available, the posterior update is

$$K_t = P_{t|t-1}W^TS_t^{-1}, \quad (6.7)$$

$$\hat{\mathbf{z}}_{t|t} = \hat{\mathbf{z}}_{t|t-1} + K_t\mathbf{e}_t, \quad (6.8)$$

$$P_{t|t} = (I - K_tW)P_{t|t-1}, \quad (6.9)$$

where K_t is the Kalman gain. The implementation uses an equivalent information-form calculation and the matrix inversion lemma, so only $k \times k$ matrices are inverted even when $m = d^2$ is large.

For parameter estimation, the forward filter is followed by the Rauch–Tung–Striebel backward smoother[3]. A smoother uses observations before and after a time point to estimate its latent state; it is appropriate inside maximum-likelihood fitting but must not be confused with a real-time forecast.

6.3 The FCN

The **FCN** is the scalar function minimized by the fit. Here it is the summed marginal Gaussian negative log likelihood over FIT observations:

$$\text{FCN}(\Theta) = \frac{1}{2} \sum_{t \in \mathcal{F}} \left[m \log(2\pi) + \log |S_t| + \mathbf{e}_t^T S_t^{-1} \mathbf{e}_t \right]. \quad (6.10)$$

$\mathcal{F}$ is the set of accepted unique FIT points, $|S_t|$ is the determinant, and $\mathbf{e}_t^T S_t^{-1} \mathbf{e}_t$ is the squared Mahalanobis innovation. The determinant term penalizes inflating uncertainty merely to shrink normalized residuals. A negative final FCN is possible for continuous densities whose fitted density exceeds one in the chosen units; FCN differences, not its sign, are meaningful.

The FCN implementation is `StateSpaceFitter.marginalFcn()` at line 1202 and `KalmanFilter.innovationNll()` near line 257.

6.4 Generalized expectation-maximization

Expectation-maximization (EM) alternates between estimating latent-state expectations and updating model parameters[4]. This code uses generalized EM[5]: an update need not be the exact maximum, but it must not increase the exact FCN. If a proposed update worsens the FCN, the step is repeatedly halved, up to 24 times. Ridge 10^{-6} stabilizes the regression normal equations.

The nominal protocol ceiling is 750 iterations. The participant report raises only the numerical ceiling to 5,000 while leaving the scientific tolerances unchanged:

$$\frac{\text{FCN}_{r-1} - \text{FCN}_r}{\max(1, |\text{FCN}_{r-1}|)} \leq 10^{-6}, \quad (6.11)$$

$$\max |\Theta_r - \Theta_{r-1}| \leq 10^{-5}. \quad (6.12)$$

A flat likelihood ridge may stop after eight consecutive negligible FCN changes. That is numerical convergence, not proof that every raw parameter entry is identifiable. In the accepted run all 45 fits converged, no stability boundary was active, and no numerical guardrail activation was reported.

6.5 Parameter errors

The files `fitted_parameters_with_errors.csv` report conditional standard errors:

- A and D : conditional regression-information errors using smoothed state moments;
- W : descriptive fixed-basis regression errors;
- diagonal R : the Gaussian variance approximation;
- Q : Gaussian covariance-element approximation; and
- $\mathbf{m}_i$: Cycle-1 baseline projected-score standard error.

These errors condition on the fitted latent representation and on temporally overlapping assumptions inside the estimator. They are identifiability diagnostics, not population confidence intervals. A participant bootstrap would be needed for population-level uncertainty. The code is `ParameterUncertainty.estimate()` at line 70.

Chapter 7

From a predicted matrix to plotted entropy

7.1 Nonlinear uncertainty propagation

The state model is Gaussian in log-matrix coordinates, but entropy is nonlinear:

$$\mathbb{E}[h(\mathbf{Y})] \neq h(\mathbb{E}[\mathbf{Y}]) \quad (7.1)$$

for entropy function h . The plot therefore distinguishes:

Central fitted.

Entropy of the matrix reconstructed from the central predicted coordinate vector.

Fitted expectation.

Mean entropy after the fitted coordinate residual distribution is propagated through matrix reconstruction and entropy.

Measurement sigma.

Standard deviation induced by fitted observation noise around the measured coordinates.

Predictive sigma.

Standard deviation of the fitted predictive entropy distribution.

The deterministic propagation uses fixed antithetic Gaussian draws: each random displacement $\mathbf{v}$ is paired with $-\mathbf{v}$, reducing odd-order Monte Carlo error[8].

Implementation: `EntropyUncertaintyPropagation.java`.

7.2 Phase-specific empirical residual distribution

The clinical closure audit showed that a single stationary coordinate-residual distribution left phase-dependent entropy offsets. The accepted descriptive report therefore forms the coordinate residuals separately for baseline, disturbance, and recovery:

$$\mathbf{r}_t^{(\phi)} = \mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}, \quad t \in \phi. \quad (7.2)$$

For a predicted coordinate center $\mathbf{c}_t$, every residual vector from the same phase is added to $\mathbf{c}_t$ and transformed back to a physical matrix. This preserves the empirical cross-coordinate pattern without estimating an enormous dense residual covariance.

7.3 Explicit entropy mean closure

After nonlinear propagation, two phase-specific observation intercepts are fitted:

$$\delta_{\text{vN},\phi} = \frac{1}{N_\phi} \sum_{t \in \phi} [S_{\text{vN},t}^{\text{obs}} - S_{\text{vN},t}^{\text{pred}}], \quad (7.3)$$

$$\delta_{\text{diag},\phi} = \frac{1}{N_\phi} \sum_{t \in \phi} [S_{\text{diag},t}^{\text{obs}} - S_{\text{diag},t}^{\text{pred}}]. \quad (7.4)$$

S_{vN} and S_{diag} are fitted separately because they represent different quantities. The gap intercept is derived:

$$\delta_{G,\phi} = \delta_{\text{diag},\phi} - \delta_{\text{vN},\phi}. \quad (7.5)$$

Residual scales are the Bessel-corrected sample standard deviations after subtracting these means. Because all FIT points from the same participant are used, phase-mean residuals are zero by construction to floating-point precision. This correction solves an observation-closure mismatch for a capacity plot; it is not an independent validation success. Held-out use would require estimating these intercepts on training participants and freezing them before opening the held-out participant. The implementation is `PhaseMeanEntropyCalibration.fit()` at line 69.

Chapter 8

Diagnostics and how to read the plots

8.1 Normalized residual

For plotted quantity q ,

$$z_t^{(q)} = \frac{q_t^{\text{obs}} - q_t^{\text{fit}}}{\sigma_{t,\text{pred}}^{(q)}}. \quad (8.1)$$

This z is a standardized residual, not the latent state $\mathbf{z}_t$. The horizontal zero line is exact agreement. Values between -2 and $+2$ fall inside the plotted two-sigma predictive band. Under a correct independent Gaussian model, about 95.45% are expected inside; temporal dependence and estimated residual scales mean that number is a guide rather than an exact binomial test.

8.2 Chi-square and nominal degrees of freedom

The coordinate-space predictive chi-square is

$$\chi^2 = \sum_{t \in \mathcal{F}} \mathbf{e}_t^T \mathbf{S}_t^{-1} \mathbf{e}_t. \quad (8.2)$$

The displayed nominal parameter count is

$$N_{\text{par}} = k^2 + 6k + \frac{k(k+1)}{2} + m + \frac{k(k+1)}{2} + k, \quad (8.3)$$

corresponding to A, D, Q , diagonal R , stored Ω , and one participant baseline state. W is excluded because it is a deterministic fixed EEG basis, and latent states are integrated by the filter rather than counted as free regression coefficients. Thus

$$\text{NDF}_{\text{nom}} = N_{\text{scalar observations}} - N_{\text{par}}, \quad \chi_{\text{red}}^2 = \frac{\chi^2}{\text{NDF}_{\text{nom}}}. \quad (8.4)$$

“Nominal” is essential: the matrix coordinates and time points are not perfectly independent, so this is an audit convention rather than an exact textbook regression degree of freedom. A value near one is necessary for calibrated descriptive fit but is not sufficient for scientific validation.

8.3 Conditional versus free-running

Conditional one-step fit.

At each accepted time, predict using earlier information, score the observation, and then let that

observation update the state used at the next time. This is the distribution optimized by the FCN and is what the 45 figures show.

Baseline-calibrated free run.

Initialize from baseline and propagate without assimilating displayed trajectory observations. Its initial covariance is the stationary solution $P = APA^T + Q$. This is a harder diagnostic of autonomous line shape.

The aggregate free-running χ^2/NDF is also close to one in this high-dimensional fit, but most scalar coordinates are small and observation noise dominates the sum. It must therefore be inspected together with the time-series plots and cannot be treated as held-out evidence.

8.4 Line and panel key

In every appendix figure:

- teal points are accepted EEG observations;
- vertical teal bars are measurement standard deviations;
- the solid coral line is the fitted expected value;
- the dashed purple line is the central reconstructed-matrix value when distinct;
- shaded background identifies baseline, disturbance, and recovery;
- each lower panel shows the normalized residual from [equation \(8.1\)](#); and
- gaps in points correspond to explicit missing/artifact-rejected windows, not interpolation.

Chapter 9

Accepted participant-specific result

9.1 Numerical summary

Quantity	Accepted result
Participants	15; p21 excluded before fitting
Independent fits	$45 = 15 \text{ participants} \times 3 \text{ bands}$
Convergence	45/45; no recorded fit failures
Reference	surface Laplacian for every input
Electrode dimension	participant-specific, $17 \leq d_i \leq 32$
Latent dimension	$k = 5$
Continuous trajectory	Cycle 1 plus unique Cycle 2, no latent reset
Conditional χ^2/NDF	1.002465 to 1.005427
Free-running diagnostic χ^2/NDF	0.999194 to 1.004616
Independent entropy two-sigma coverage	95.313% overall
Minimum group coverage	91.124%; p15 alpha disturbance S_{vN} and p15 theta disturbance S_{diag}
Groups below 90% coverage	0 of 270
Phase-mean residual	at most 7.8×10^{-13} in magnitude, expected because phase intercepts use those FIT points

The conditional coordinate-space agreement is uniform across participants and bands. Participants p04 and p10 have many more explicit missing positions, which reduces their fit-point count and slightly raises the reduced chi-square, but the fits still converge and remain close to one. Isolated entropy residuals outside two sigma remain visible; a scientifically plausible Gaussian residual model should not force every point inside the band.

9.2 Interpretation

The result supports the limited conclusion that a five-factor, phase-driven, participant-specific stochastic model can provide a calibrated one-step description of the retained Laplacian EEG coordinate series after explicit phase-specific entropy observation closure. It does not show that one parameter set generalizes to another participant. It also does not show that the phase-specific entropy offsets can be predicted from baseline.

The all-participant plots are included so agreement and remaining structure can be assessed directly

rather than inferred only from scalar scores. Particular attention should be paid to sustained runs of residuals with one sign, recovery-specific widening, and differences between S_{vN} and S_{diag} .

9.3 Scientific stopping and advancement conditions

The present capacity stage is complete. Before calling the model predictive, the following sequence is required:

1. Freeze the current measurement definitions, quality gate, cycle stitching, and entropy closure rule.
2. Establish simulation recovery across declared dimensions, noise regimes, and missing-window patterns on an independently generated synthetic validation set.
3. In participant-held-out validation, fit all population quantities and phase entropy offsets without the held-out participant.
4. Use only that participant's Cycle-1 baseline for permitted individual calibration.
5. Predict disturbance and recovery without refitting and compare against persistence, baseline-hold, phase-mean, and no-phase-input null models.
6. Require preregistered improvement and coverage thresholds. A descriptive one-step fit alone cannot open this gate.

Failure to beat the nulls on held-out participants would mean that the participant-specific capacity is real but does not generalize with the current predictors. It would not justify adding cross-band coupling until the simpler residual failure is identified and replicated.

Chapter 10

Reproduction and code traceability

10.1 Exact input and output artifacts

All paths below are relative to the BioReg-Analysis root. Hashes are SHA-256 digests of the exact declared files.

```
Measurement input:
output/development-complex-state-laplacian-frozen16-v1/complex-state

Measurement manifest:
output/development-complex-state-laplacian-frozen16-v1/complex_state_manifest.csv

Protocol snapshot:
output/eeg-dynamic-factor-simulation-v4/protocol_snapshot.json
SHA-256: 8eabcf59a7ea6fa0aa2d25199eb828bfc3af1f004db44ba8afafd58cff5223bc

Quality gate:
protocol/participant-fit-quality-exclusions-v1.json
SHA-256: 2f9f41f2faecab0bceca54721a19502040ad2ea99641cd1a45ffa15108405c4e

Accepted result:
output/development-individual-laplacian-excluding-p21-k5-phase-closure-v13

Run summary:
output/development-individual-laplacian-excluding-p21-k5-phase-closure-v13/
individual_participant_fit_summary.json
```

10.2 Reproduction command

Build the authoritative measurement library and analysis JAR, then run the exact participant list. The quality gate is still supplied and independently enforces the p21 exclusion.

```
cd /Users/veronique/BASE/Code/BioReg-EEG
mvn clean install

cd /Users/veronique/BASE/Code/BioReg-Analysis
mvn clean test package

java -cp target/BioReg-Analysis-0.1.0-SNAPSHOT-all.jar \
  org.cas.bioreg.analysis.statespace.report.IndividualParticipantFitReport \
  output/eeg-dynamic-factor-simulation-v4/protocol_snapshot.json \
  output/development-complex-state-laplacian-frozen16-v1/complex-state \
  protocol/participant-fit-quality-exclusions-v1.json \
```

```
output/development-individual-laplacian-excluding-p21-k5-phase-closure-v13 \
theta,alpha,beta \
5 \
p01,p02,p03,p04,p10,p11,p12,p13,p15,p16,p17,p18,p19,p20,p22 \
5000
```

This command overwrites files of the same name in the selected output directory. Use a new output directory if preserving the accepted artifact byte-for-byte is required.

10.3 Code map

Line numbers below refer to the source version used to produce this report. The two roots are:

```
/Users/veronique/BASE/Code/BioReg-EEG
/Users/veronique/BASE/Code/BioReg-Analysis
```

BioReg-EEG measurement code

- `SurfaceLaplacian.java`, class at line 26: four-nearest-neighbor surface-Laplacian operator, equation (3.1).
- `Fourier.java`, `analytic()`: frequency-domain band limiting and analytic quadrature signal.
- `StateSeries.java`, `windowState()` at line 197: normalized complex state, ρ , S_{vN} , S_{diag} , and derived gap.
- Same file, `windowVoltageCovariance()` at line 271: conventional voltage covariance C .
- Same file, `windowMagnitudeSquaredCoherence()` at line 308: time-local analytic coherence matrix; this is exported but not fitted as an independent target here.
- `ComplexStateExport.java`: JSON metadata plus binary ρ , covariance, and coherence matrices with provenance.

BioReg-Analysis model and reporting code

- `StateSpaceObservationBuilder.java`: line 54 requires Laplacian input.
- Same file, line 349: removes duplicate Cycle-2 baseline without a latent reset.
- Same file, line 450: selects the non-overlapping grid.
- Same file, line 495: learns the Cycle-1 origin.
- `DensityMatrixGeometry.java`: lines 43 and 67 implement the forward and inverse physical matrix map.
- `RecoveryDrive.java`: `input()` implements the four disturbance Bernstein components and two recovery exponentials.
- `StateSpaceFitter.java`: line 122 starts the participant full-trajectory fit; line 1202 evaluates equation (6.10).
- `KalmanFilter.java`: `run()` at line 75 performs conditional filtering, missing-window propagation, and smoothing; `runFreeRunning()` at line 95 performs the baseline-calibrated autonomous diagnostic.
- `ParameterUncertainty.java`, `estimate()` at line 70: conditional parameter standard errors.
- `EntropyUncertaintyPropagation.java`: matrix-consistent nonlinear entropy means and standard deviations.
- `PhaseMeanEntropyCalibration.java`, `fit()` at line 69: separate S_{vN} and S_{diag} phase intercepts and residual scales; gap remains derived.
- `SimulationFitPlot.java`: line 280 builds reconstructed observations and uncertainties; line 721 calculates innovation chi-square and nominal NDF.
- `IndividualParticipantFitReport.java`, `fitOne()` at line 228: orchestrates fitting, parameter tables, both prediction modes, and the run summary.

Chapter 11

Glossary of technical terms

Accepted window

A window that passes the frozen signal-quality rules and has finite measurements.

Analytic signal

A complex signal whose real part is the measured band-passed waveform and whose imaginary part is its Hilbert-transform quadrature.

Artifact

Non-neural or numerically unusable signal contamination, such as large transients, flat channels, or rank collapse.

Band

A frequency interval. This report uses theta 4–8 Hz, alpha 8–13 Hz, and beta 13–30 Hz.

Baseline

Pre-disturbance phase used to define the participant’s operating point.

Capacity analysis

A fit to the same participant observations it describes; it asks whether the model class can represent the trajectory, not whether it generalizes.

Channel / electrode

One recorded scalp voltage location. In this report the model matrix dimension equals the participant’s retained electrode count.

Chi-square

Sum of squared prediction errors weighted by the full predictive covariance, equation (8.2).

Conditional one-step prediction

Prediction of the next observation after earlier observations have updated the state.

Convergence

Satisfaction of numerical stopping rules for FCN change and/or a stable likelihood plateau. It is not identical to scientific validity.

Covariance

Matrix of joint variation. C is voltage covariance; Q is process covariance; R is observation covariance; P is latent-state uncertainty.

Cycle

One sequential experimental protocol pass. Here Cycle 2 follows Cycle 1 in the same recording and is not independent test–retest data.

Density matrix

The project’s name for the Hermitian, positive, trace-one complex state matrix ρ . It is a mathematical representation of EEG analytic states, not a claim that scalp EEG is a microscopic quantum state.

Disturbance

Experimental interval in which the stressor is present.

Drive

Known phase basis $\mathbf{u}_t$ multiplied by fitted coefficient matrix D .

Eigenvalue

Scalar λ satisfying $M\mathbf{v} = \lambda\mathbf{v}$. Eigenvalues of ρ determine S_{vN} .

Entropy

A nonlinear summary of a probability spectrum or diagonal distribution, measured here in nats because natural logarithms are used.

FCN

The scalar objective minimized by the fit, here the marginal Gaussian negative log likelihood.

Fixed effect

A participant-specific fitted quantity without an estimated population distribution in the current one-participant fit.

Free-running prediction

Propagation after baseline without assimilating the displayed later observations.

Frobenius norm

Matrix size $\|M\|_F = \sqrt{\sum_{j\ell} |M_{j\ell}|^2}$.

Gap

$G = S_{\text{diag}} - S_{\text{vN}}$, a derived diagnostic, not a third independent fit target.

Generalized EM

An EM variant in which each parameter update need only improve or preserve the actual objective.

Hermitian

Complex analogue of symmetric: $M = M^\dagger$.

Held out

Not used to estimate the parameters being tested. Nothing in the participant-specific capacity plots is participant-held-out.

Hilbert transform

Linear transform used to construct a quadrature signal and instantaneous analytic phase.

Identifiability

Ability of data and model to distinguish parameter values. Predictive trajectories can be identifiable even when rotated latent entries are not.

Innovation

One-step observation residual $\mathbf{e}_t$ before the current measurement update.

Kalman filter

Exact recursive Gaussian state estimator for a linear state-space model.

Latent

Unobserved. $\mathbf{z}_t$ is inferred from EEG through the state-space equations.

Loading

Matrix W mapping latent coordinates to EEG observation coordinates.

Mahalanobis distance

Squared residual weighted by inverse covariance, $\mathbf{e}^T S^{-1} \mathbf{e}$.

Matrix logarithm/exponential

Spectral functions used to move between constrained positive matrices and unconstrained Hermitian coordinates.

Missing point

A physical grid position retained in time but without an EEG observation or likelihood contribution.

Nat

Entropy unit produced by the natural logarithm. One bit equals $\log 2$ nats.

NDF

Nominal degrees of freedom, observations minus the declared nominal parameter count.

Observation noise

Variation in $\mathbf{y}_t$ not represented by the latent factors, modeled by R .

Operating point

Participant-specific baseline center $\mathbf{m}_i$ toward which unforced stable dynamics return.

Phase

Baseline, disturbance, or recovery experimental interval; distinct from an oscillatory phase angle.

Positive semidefinite

Matrix whose quadratic form is never negative; required for covariance and density matrices.

Process noise

Unmodeled random state evolution $\boldsymbol{\eta}_t$ with covariance Q .

Recovery

Post-disturbance experimental interval.

Residual

Observed minus fitted value. A normalized residual divides this difference by predictive sigma.

Ridge

Small diagonal regularization added to regression normal equations to avoid unstable inversion.

Sigma

Standard deviation. A two-sigma band spans fitted expectation plus or minus two standard deviations.

Smoother

State estimator that uses both earlier and later observations; here the Rauch–Tung–Striebel algorithm inside fitting.

Spectral radius

Largest absolute eigenvalue of A ; below one is the discrete-time stability condition used here.

Surface Laplacian

Spatial re-reference subtracting the mean of nearby electrodes, equation (3.1).

Trace

Sum of diagonal entries. $\text{Tr } \rho = 1$; $\text{Tr } C$ is total voltage variance.

Validation

Evaluation on information not used for the parameters being tested, with frozen criteria and appropriate null comparisons.

Appendix A

All conditional participant-band fit plots

The following 45 figures are the complete accepted conditional one-step set. They are ordered by participant and, within participant, theta, alpha, and beta. The exact numeric values under every plot are stored in the adjacent CSV file in the accepted result directory. No p21 plot is present because the quality exclusion is applied before fitting.




Figure A.1: Participant p01, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.



Figure A.2: Participant p01, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.3: Participant p01, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.4: Participant p02, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.5: Participant p02, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.6: Participant p02, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.7: Participant p03, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.8: Participant p03, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.9: Participant p03, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.10: Participant p04, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.11: Participant p04, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.12: Participant p04, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.13: Participant p10, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.14: Participant p10, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.15: Participant p10, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.16: Participant p11, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.17: Participant p11, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.18: Participant p11, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.19: Participant p12, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.20: Participant p12, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.21: Participant p12, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.22: Participant p13, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.23: Participant p13, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.24: Participant p13, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.25: Participant p15, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.26: Participant p15, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.27: Participant p15, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.28: Participant p16, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.29: Participant p16, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.30: Participant p16, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.31: Participant p17, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.32: Participant p17, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.33: Participant p17, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.34: Participant p18, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.35: Participant p18, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.36: Participant p18, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.37: Participant p19, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.38: Participant p19, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.39: Participant p19, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.40: Participant p20, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.



Figure A.41: Participant p20, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.42: Participant p20, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.43: Participant p22, theta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.44: Participant p22, alpha band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.




Figure A.45: Participant p22, beta band: conditional one-step participant-specific fit. Every accepted unique EEG window is fitted. The four main panels show baseline-centered log absolute band activity, S_{vN} , S_{diag} , and the derived gap $G = S_{diag} - S_{vN}$; each has a normalized-residual panel directly below it. Background color marks baseline, disturbance, and recovery. This is descriptive capacity evidence, not held-out validation.

Appendix B

Protocol constants used in this result

Protocol field	Value	Meaning
eigenvalueMixingEpsilon	10^{-5}	Identity mixing before matrix logarithm
amplitudeFloor	10^{-12}	Positive floor before log activity
covarianceFloor	10^{-6}	Numerical covariance eigenvalue/variance floor
stabilityMargin	0.001	Requires spectral radius at most 0.999
ridge	10^{-6}	Regression stabilization
maximumIterations	750 source; 5000 run ceiling	Source scientific protocol plus numerical participant-fit extension
relativeFcnTolerance	10^{-6}	Relative objective convergence scale
parameterTolerance	10^{-5}	Maximum dynamic-parameter convergence scale
maximumMissingSteps	8	Longest accepted consecutive missing grid run
fastTimeFraction	0.15	Fast recovery exponential time constant divided by phase length
slowTimeFraction	0.60	Slow recovery exponential time constant divided by phase length
Latent dimension	5	Number of participant-band dynamic factors
Surface-Laplacian neighbors	4	Nearest retained electrodes subtracted in equation (3.1)
EEG window	1000 samples, about 4 s	Matrix and activity measurement interval
Candidate overlap	50%	Exported candidates step by about 2 s
Model-grid stride	2 candidates	Non-overlapping fitted steps, about 4 s

Bibliography

- [1] S. Roy and J. Nuamah, “Neurophysiological Dataset of Stress Resilience During Human–Computer Interaction,” version 1.0.0, PhysioNet, 2026. DOI: [10.13026/x3vc-p627](https://doi.org/10.13026/x3vc-p627).
- [2] R. E. Kalman, “A New Approach to Linear Filtering and Prediction Problems,” *Journal of Basic Engineering*, 82(1), 35–45, 1960. DOI: [10.1115/1.3662552](https://doi.org/10.1115/1.3662552).
- [3] H. E. Rauch, F. Tung, and C. T. Striebel, “Maximum Likelihood Estimates of Linear Dynamic Systems,” *AIAA Journal*, 3(8), 1445–1450, 1965. DOI: [10.2514/3.3166](https://doi.org/10.2514/3.3166).
- [4] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum Likelihood from Incomplete Data via the EM Algorithm,” *Journal of the Royal Statistical Society B*, 39(1), 1–38, 1977. JSTOR [2984875](https://www.jstor.org/stable/2984875).
- [5] C. F. J. Wu, “On the Convergence Properties of the EM Algorithm,” *Annals of Statistics*, 11(1), 95–103, 1983. DOI: [10.1214/aos/1176346060](https://doi.org/10.1214/aos/1176346060).
- [6] H. Hotelling, “Analysis of a Complex of Statistical Variables into Principal Components,” *Journal of Educational Psychology*, 24, 417–441 and 498–520, 1933. DOI: [10.1037/h0071325](https://doi.org/10.1037/h0071325).
- [7] R. T. Farouki, “The Bernstein Polynomial Basis: A Centennial Retrospective,” *Computer Aided Geometric Design*, 29(6), 379–419, 2012. DOI: [10.1016/j.cagd.2012.03.001](https://doi.org/10.1016/j.cagd.2012.03.001).
- [8] J. M. Hammersley and D. C. Handscomb, *Monte Carlo Methods*, Methuen, 1964, section 5.4.
- [9] N. J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, 2008. DOI: [10.1137/1.9780898717778](https://doi.org/10.1137/1.9780898717778).
- [10] F. Perrin, J. Pernier, O. Bertrand, and J. F. Echallier, “Spherical Splines for Scalp Potential and Current Density Mapping,” *Electroencephalography and Clinical Neurophysiology*, 72(2), 184–187, 1989. DOI: [10.1016/0013-4694\(89\)90180-6](https://doi.org/10.1016/0013-4694(89)90180-6).