PREPROCESSING SENSITIVITY AS A SIGNAL IN BRAIN REPRESENTATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

We treat a brain representation's response to preprocessing as an observable feature in its own right. For fMRI, we measure how a joint time-delay spectral score changes when five selected aCompCor nuisance regressors are added to the same recording. The score measures excess concentration in two leading modes relative to channelwise phase-randomized references. In a post hoc analysis of 201 Alzheimer's Disease Neuroimaging Initiative participants, this paired response has an estimated adjusted slope of +0.472 percentage points per point of clinical severity (CDR-SB; HC3 standard error 0.247), while the original baseline score has no established association. The response coefficient remains similar when common temporal regressors change both pipelines' absolute scores, in a sensitivity analysis on the same cohort. We then characterize what this response can identify. A Gaussian construction shows that a related population statistic detects collective changes while every channel's complete univariate law remains fixed. Yet observation changes alone can produce a nonmonotone response at fixed latent dynamics. Observationally identical models further show that no decoder of the same inputs can resolve the latent source without additional assumptions. Five controlled simulation families connect these mechanisms to the finite-sample spectral measurement. Together, the findings make preprocessing response a concrete target for representation analysis, linking an estimated clinical association to explicit limits on latent interpretation.

1 INTRODUCTION

Brain dynamics are expressed jointly across regions and time. A useful representation should therefore capture organization that is absent from a collection of separate regional summaries. Yet the measured trajectories also depend on motion regression, nuisance components and temporal filtering. This creates two linked questions: which collective property does a representation detect, and how does its response to an observation model relate to clinical state?

Variability due to preprocessing and broader analysis choices is well established in neuroimaging (Botvinik-Nezer et al., 2020; Ciric et al., 2017; Lindquist et al., 2019). Treating this variability only as a choice of a preferred pipeline misses a measurable object: the *paired response* of a representation when a declared operator changes on the same recording. This response may have individual or clinical structure even when the original representation does not exhibit the same association. Its interpretation nevertheless requires care: a response can reflect altered latent organization, altered observation, or the information removed by the nuisance model.

We investigate this distinction using a joint delay representation of seven cortical networks. Our statistic J is the excess leading-two spectral energy fraction above a finite channelwise phase reference. The delay embedding and spectral fraction are established constructions; our contribution is the linked characterization and empirical evaluation of their *operator response*. In particular, we ask whether clinically associated variation in $\Delta J = J_{\mathrm{A5}} - J_{\mathrm{BASE}}$ persists under a common temporal nuisance control, and use controlled worlds to establish what such a result can mean.

The analysis produces three connected findings. First, the same statistic detects collective changes in a Gaussian model whose complete single-channel laws are unchanged. Second, nuisance projection and normalized concentration create an observation-only mechanism for a nonmonotone response, even at fixed latent dynamics. Third, a positive clinical association of the paired response survives a

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

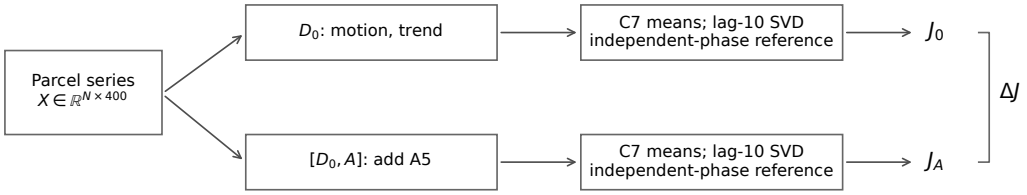

Figure 1: **The response is paired within each recording.** The same retained parcel series enter joint BASE and BASE+A5 nuisance regressions. Both branches then use the same seven-network aggregation, normalized delay score and channelwise phase-reference procedure; their difference ΔJ is one observable per person. The clinical analysis relates these paired responses to contemporaneous CDR-SB across 201 participants after fixed covariate adjustment. The diagram defines a measurement path, not a neural-source attribution.

common-DCT control on the same 201 participants. The absolute scores change under this control, while the estimated clinical contrast remains close to its original value. We establish this contrast directly, rather than infer it from the difference between two significance decisions.

These findings motivate evaluating joint representations through both their response to preprocessing and their clinical information. We first define the paired spectral observable, derive its collective sensitivity and observation dependence, and test its clinical association. We then extend the evaluation to future severity: window-FC dispersion improves matched two-year CDR-SB prediction in 48 MCI participants from the same programme. This complementary task uses a different representation and endpoint, allowing us to examine predictive value alongside the spectral-response association. Throughout, explicit operators, reference laws and matched comparisons make the evidence interpretable.

2 Related work

Time-delay coordinates and singular-spectrum decompositions provide established ways to represent trajectories from observed time series (Broomhead & King, 1986). HAVOK combines a Hankel representation with a fitted dynamical model (Brunton et al., 2017). Our joint delay representation uses this spectral vocabulary to summarize collective organization across cortical networks; it does not estimate a Koopman operator or an attractor dimension. Temporal Mapper offers a complementary systems-level approach: it constructs directed transition graphs from multivariate activity and temporal adjacency, with validation against a specified mechanistic model (Zhang et al., 2023). A concentration statistic and a transition graph therefore describe different properties of a trajectory. Neither their use nor a low-dimensional representation alone establishes deterministic chaos in BOLD.

Surrogate-data methods make the reference hypothesis part of the measurement (Schreiber & Schmitz, 2000). Independent channel phases preserve individual periodograms while changing cross-channel phase relations. Applying the same frequency-wise phase shift across channels instead preserves the cross-spectrum (Prichard & Theiler, 1994); these references answer different questions. Our spectral score targets organization relative to the former reference, including organization compatible with stationary linear Gaussian models. This distinction matters because apparently richer connectivity summaries can reproduce information already present in a static covariance model (Novelli & Razi, 2022). We consequently separate sensitivity to collective structure from claims of nonlinear dynamics or biological origin.

The dependence of fMRI results on processing is also established. Prior work benchmarks confound regression (Ciric et al., 2017), evaluates network-construction pipelines across several criteria (Luppi et al., 2024), and optimizes individual task-fMRI pipelines using prediction and reproducibility (Churchill et al., 2015). Spectral mismatch and sequential projections can reintroduce components removed earlier; joint regression supplies a known remedy (Hallquist et al., 2013; Lindquist

et al., 2019). These findings motivate explicit operators and the common temporal-basis control here. Furthermore, identifiable motion and physiological signatures can coexist with improved participant identification after denoising (Xifra-Porxas et al., 2021). Reproducibility and neural specificity thus remain distinct questions.

Representation identifiability requires explicit assumptions: Yao et al. (2025) distinguish latent-variable from causal-graph identification through invariance principles. Our preprocessing views do not establish their latent-invariance and observation-map assumptions. We use these distinctions to evaluate the clinical association of a specified preprocessing response. The contribution is the linked operator-response analysis and its controlled interpretation, built from established projection, spectral and surrogate methods.

3 A JOINT REPRESENTATION AND ITS OPERATOR RESPONSE

Let $X \in \mathbb{R}^{N\times 400}$ contain retained parcel time series on one temporal support. A fixed matrix $W \in \mathbb{R}^{400\times 7}$ averages parcels within the seven cortical networks. For a nuisance design D, let

$$R_D = I - U_D U_D^\top, \qquad Y_D = R_D X W, \tag{1}$$

where U_D is an orthonormal basis of $\mathrm{col}(D)$. We construct the joint projector on all 400 parcels and then aggregate. Since W is fixed, $R_D(XW) = (R_D X)W$; aggregation does not create an additional temporal operation.

Joint delayed geometry. Center each column of Y_D and divide by its root mean square, giving $\widetilde{Y}_D$. For $L = 10$, form the channel-major block-Hankel matrix

$$H_L(\widetilde{Y}_D)_{(c-1)L+\ell+1,t} = \widetilde{Y}_{D,t+\ell,c}, \quad c = 1,\dots,7, \quad \ell = 0,\dots,L-1, \tag{2}$$

with $t = 1,\dots,N-L+1$. There is no additional row centering. For singular values $s_1 \geq s_2 \geq \cdots$, define

$$\rho_2(Y) = \frac{s_1^2 + s_2^2}{\sum_k s_k^2}, \qquad J(Y) = \rho_2(Y) - \frac{1}{B}\sum_{b=1}^{B} \rho_2(Y^{(b)}), \qquad B = 199. \tag{3}$$

Each $Y^{(b)}$ randomizes the Fourier phases independently across channels while preserving their Fourier magnitudes, DC and real-valued Nyquist constraints. The same normalization and Hankel construction are applied to every reference. Thus J measures excess joint concentration relative to this specific phase-reference law. It is not an estimate of fractal dimension, entropy, nonlinear dynamics or anatomical whole-brain coverage. Finite nonperiodic Hankel edge weights are not preserved by the Fourier control.

Projection and finite Hankel edge weights can alter covariance and reference energy despite preserved Fourier magnitudes. We therefore distinguish the population construction from the finite estimator and evaluate its null offset explicitly (Appendices A.8 and D).

A paired change in observation. Write $D_0 = [\mathbf{1}, \mathrm{trend}, M24]$ for the baseline design and $A \in \mathbb{R}^{N\times 5}$ for the selected aCompCor columns. Both branches are joint regressions:

$$J_0 = J(R_{D_0} X W), \quad J_A = J(R_{[D_0,A]} X W), \quad \Delta J = J_A - J_0. \tag{4}$$

In the temporal sensitivity, the same delivered DCT columns C enter both designs, replacing D_0 by $D_0^C = [D_0, C]$. The cosine columns are restricted to the original retained indices, not regenerated for a shortened series. A, W, L, the reference ensemble size and the clinical design remain fixed.

For any common design D, the effective additional subspace is $\mathrm{col}(R_D A)$. With $P = R_D$ and $Q = R_{[D,A]}$, nested orthogonal projections imply

$$PQ = QP = Q, \qquad \|PX\|_F^2 = \|QX\|_F^2 + \|(P-Q)X\|_F^2. \tag{5}$$

This raw-energy identity does not order J after normalization, delay embedding and reference subtraction. In particular, ΔJ is not a literal amount of removed neural signal. Joint regression also avoids the reintroduction of removed components that can occur under naive sequential projections (Lindquist et al., 2019).

4 WHAT AN OPERATOR RESPONSE IDENTIFIES

Our question concerns collective organization and its dependence on observation. The following construction first establishes which collective changes the spectral observable detects, then shows how the same measurement responds to observation changes. Together these results specify what the clinical response measures and what additional information would identify its physical source.

Let T_L be the covariance of $L \geq 2$ consecutive observations of a stationary, zero-mean, unit-variance process, with eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$. Thus $\operatorname{tr} T_L = L$. For seven population-normalized channels, assume the separable delay covariance $C(r) \otimes T_L$, where $C(r) = (1-r)I_7 + r\mathbf{1}\mathbf{1}^\top$ and $r \in [0,1]$. Define its leading-two energy fraction minus that of independent channels with the same marginal covariance:

$$J_{\mathrm{pop}}(r) = \frac{\max\{5r\lambda_1,\ (\lambda_2 - \lambda_1) + 6r(\lambda_1 + \lambda_2)\}}{7L}. \tag{6}$$

This covariance functional is a tractable companion to the finite-record J, not its expectation or an exact phase-reference null.

The mechanism follows directly from the spatial spectrum. The common direction has eigenvalue $1 + 6r$, whereas each of its six orthogonal directions has eigenvalue $1 - r$. Tensoring with the temporal covariance gives leading modes $(1+6r)\lambda_1$ and $\max\{(1+6r)\lambda_2, (1-r)\lambda_1\}$, with total energy $7L$. Subtracting the independent-channel value $2\lambda_1/(7L)$ yields (6). Increasing shared organization transfers spectral weight toward the common direction without changing the marginal time-series law. The switch between the two candidates is an eigenvalue-ordering boundary, not evidence of a dynamical transition.

Proposition 1 (Collective sensitivity at fixed marginal laws)**.** *Let $Z_1, \ldots, Z_7, H$ be independent copies of the same stationary zero-mean, unit-variance Gaussian process. For $S_i^{(r)} = \sqrt{1-r}\, Z_i + \sqrt{r}\, H$, every individual channel has the same entire univariate finite-dimensional distribution for all r. Nevertheless, $J_{\mathrm{pop}}(r)$ in (6) is strictly increasing.*

Thus a joint observable can detect changes that every method based only on individual channel laws must miss. This example needs neither nonlinear dynamics nor chaos: it changes cross-channel organization within a Gaussian family. The proposition concerns the population construction; the controlled experiments test the actual finite-record measurement chain.

Proposition 2 (An observation-only response can reverse direction)**.** *Keep the latent target $Z = (Z_1, \ldots, Z_7)$ fixed and observe $X_i = Z_i + \sqrt{s}\, U$, with proxy $W = U + \tau V$, where U, V, Z_i are independent copies as above, $s \geq 0$, and $0 < \tau < \infty$. Population regression on W leaves common-component variance ts, where $t = \tau^2/(1+\tau^2)$. Consequently,*

$$\delta(s) = J_{\mathrm{pop}}\left(\frac{ts}{1+ts}\right) - J_{\mathrm{pop}}\left(\frac{s}{1+s}\right) < 0 \quad (s > 0), \qquad \delta(0) = \lim_{s\to\infty} \delta(s) = 0. \tag{7}$$

Hence δ is nonmonotone and has positive increments despite unchanged latent dynamics.

This is saturation of a normalized spectral fraction. When both branches' leading two modes share the common spatial direction, the contrast simplifies to $-K(1-t)s/[(1+s)(1+ts)]$, with $K = 6(\lambda_1 + \lambda_2)/(7L) > 0$; its slope becomes positive for $s > 1/\sqrt{t}$. Appendix A supplies the spectral ordering, branch condition, and exact finite-increment identity. The symmetric example does not isolate channel unit-RMS scaling from the subsequent energy-fraction normalization.

There is also an exact identification boundary for $s > 0$. Declare instead the latent target $S^\dagger = X/\sqrt{1+s}$, with observation gain $\sqrt{1+s}$, and retain the same proxy. Both models produce identical X, W for every realization, although the target's cross-channel correlation is zero in the first and $s/(1+s)$ in the second. Every statistic or learned decoder given only these identical observed inputs has identical outputs or output laws. Distinguishing the targets requires additional information or assumptions; calling a regressor *nuisance* does not supply them.

These constructions motivate testing the observable paired response directly. The following study fits its clinical coefficient on a common participant design, retaining covariance between the two branches.

5 STUDY DESIGN

Cohort selection before the new measurements. The ADNI source campaign contains 732 visits from 270 participants, selected by metadata to match three index-diagnosis groups in longitudinal depth. Its construction and remaining acquisition mismatch are detailed in Appendix B. For the present cross-sectional study, 615 visits passed a checked CDR-SB eligibility gate: a finite same-phase assessment within a nominal 90-day window, no contradictory nearest match or recorded assessment failure, and exact agreement with the six component scores. We selected the first eligible scan per person, giving 267 participants, then restricted those already selected scans to exactly TR= 3 s. The resulting 201-person panel comprises 64 CN, 70 MCI and 67 Dementia visits. We did not substitute a later TR= 3 scan for an earlier eligible scan. A native-lag all-TR analysis of 267 participants is retained separately.

The fixed TR makes the first-to-last lag span 27 seconds; record length still varies. The clinical assessment dates are nominal because assessment-date versus registry-fallback provenance is unresolved per row. The target is a contemporaneous retrospective association, not a forecast or pathology-confirmed disease classifier.

Observation provenance and quality. The parcel fields use the Schaefer 400-parcel, seven-network cortical atlas (Schaefer et al., 2018; Yeo et al., 2011). All 201 scans share the pinned fMRIPrep 25.2.5 image (Esteban et al., 2019; fMRIPrep developers, 2026). The aCompCor method derives nuisance components from anatomical noise regions (Behzadi et al., 2007). The actual runtime requests all aCompCor components; the downstream rule orders the stored combined-mask components by finite singular value and fixes five. All 1,005 selected components match this contract. This is a fixed-five selection, not a 50% explained-variance selection. The original source cosine columns occur in all 201 confound tables. Their absence from both original downstream joint designs motivated the shared-DCT sensitivity. Source and reportlet checks establish provenance; a new exhaustive visual tissue-mask review was not performed. Motion and signal-quality variables enter the clinical adjustment.

Clinical estimand and inferential history. Let c be the Clinical Dementia Rating Sum of Boxes (CDR-SB), a clinical severity measure (O'Bryant et al., 2008), and Z the unchanged participant-level nuisance design: age, sex, age-proxy indicator, record length, TR, recorded field-map status, study phase, site, scanner and three quality covariates (mean FD, fraction FD above 0.2 mm, median standardized DVARS). For $z = R_Z c$, the paired coefficient is

$$\widehat{\beta}_\Delta = \frac{z^\top \Delta J}{z^\top z} = \widehat{\beta}_A - \widehat{\beta}_0 = \widehat{\beta}_{\Delta\rho_2} - \widehat{\beta}_{\Delta\bar{\rho}_2^{\mathrm{ref}}}. \tag{8}$$

We report HC3 standard errors and one two-sided restricted-HC2 Rademacher wild-bootstrap test with full HC3 refits, using 19,999 draws for each separately frozen paired study. The original baseline analysis used a different residual-sign reference and remains a distinct result. The cohort had already contributed to research development: the paired analysis is a post-exposure discovery; the DCT comparison is a subsequent same-cohort sensitivity. Neither is an untouched confirmation. The old/new coefficient difference is descriptive, with no equivalence test or second difference test. Appendix E gives the exact algorithm, leverage treatment and known approximate-calibration limits.

The pairing also matters for uncertainty. With common full-model residuals u_A, u_0, leverage h_i, and coefficient weights $\omega = z/(z^\top z)$, the HC3 variance on nonsaturated rows is

$$\widehat{V}_\Delta = \sum_i \frac{\omega_i^2}{(1-h_i)^2}(u_{Ai} - u_{0i})^2. \tag{9}$$

The paired HC3 variance retains cross-branch residual covariance. Fourteen unit-leverage rows have verified zero clinical information and contribute zero; ranks and numerical guards are detailed in Appendix E.

Controlled mechanisms. Five Gaussian worlds vary one declared factor at a time: an inert null label, common observation amplitude, proxy quality, latent sharing, or target contamination of a proxy. We retain five levels per world, 16 independent realizations at each of two lengths ($N = 197, 788$), and all 800 scenario records. Every realization passes through 400-parcel joint BASE/A5 projection, C7 aggregation and the identical J estimator. The five synthetic proxies are exogenous; they do not simulate tissue PCA. Parameters were fixed before this panel and were not fitted to the human effect. The complete worlds and all curves appear in Appendix F.

6 RESULTS

6.1 THE CLINICAL ASSOCIATION LIES IN THE PAIRED OPERATOR RESPONSE

On the 201-person panel, mean original BASE concentration is 32.09% versus 15.67% in its phase reference, giving mean $J = 16.41$ percentage points. The excess is positive in every participant for both BASE and A5, but this is a descriptive comparison with each mean reference, not 201 individual significance tests. The original primary BASE–CDR-SB coefficient is -0.0443 points per CDR-SB point (partial $r = -0.0185$, reference $p = 0.8163$). The all-TR baseline sensitivity is likewise unestablished. We preserve this primary null.

The subsequently frozen *paired* coefficient is positive: $+0.4722$ points per CDR-SB point, HC3 SE 0.2470 and partial $r = 0.2332$ (Table 1). Its nominal wild-bootstrap $p = 0.0148$ comes from 295 exceedances among 19,999 draws with the plus-one correction. Directly testing ΔJ accounts for the covariance between branches. The observed-concentration contribution is $+0.6541$ and the phase-reference contribution $+0.1819$; their difference is the paired slope. Neither component receives an additional test.

Table 1: **The paired clinical response survives a shared temporal nuisance control.** All rows use the same 201 people and clinical nuisance span; slopes and SEs are in percentage points of J per CDR-SB point. †The initial BASE p uses the earlier residual-sign reference. Paired rows use the declared approximate wild bootstrap; their p values are nominal. The old/new paired difference is descriptive.

Clinical outcome	Slope	HC3 SE	Partial r	Reference p
BASE J (initial primary)	-0.0443	—	-0.0185	$0.8163^{\dagger}$
A5–BASE, original	+0.4722	0.2470	0.2332	0.0148
A5–BASE, common DCT	+0.4641	0.2210	0.2365	0.0076

6.2 SHARED DCT CHANGES SCORE LEVELS WHILE PRESERVING A POSITIVE CONTRAST

Adding the original cosine columns to both joint designs gives a paired slope of $+0.4641$ points per CDR-SB point (HC3 SE 0.2210; partial $r = 0.2365$). All 201 pairs remain eligible, all rank guards pass and A5 adds five directions in every scan. The nominal $p = 0.0076$ corresponds to 151 exceedances. The old/new point-estimate difference is -0.0081 points; closeness of estimates is not a test of equivalence or an independent replication.

This control changes the measurement: mean J becomes 15.286 for BASE and 12.229 for A5, compared with 16.415 and 13.794 originally. Figure 2 shows the participant-level adjusted relationship and the individual old/new responses. The new mean paired response is -3.057 points. Thus a positive clinical slope coexists with a negative average A5–BASE response. The new observed/reference slope decomposition is $0.7255 - 0.2614 = 0.4641$. Conditional phase Monte Carlo SE of the paired slope is 0.0038 points, much smaller than the participant-sampling HC3 SE.

The bootstrap is approximate. In the prior declared calibration, one heteroskedastic record-length world yielded 22/1,000 rejections at nominal 0.01 (59/1,000 at 0.05). The smaller DCT p therefore does not establish calibrated 1% confirmation or a superior pipeline. The empirical result is the positive paired coefficient and its persistence under this source-motivated control; the calibration limits qualify its inferential strength.

The decomposition $0.7255 - 0.2614 = 0.4641$ attributes the paired slope to excess concentration relative to each branch's own reference. Holding that reference fixed across operators would define a different response.

Independent numerical reconstruction. Independent OLS/HC3 and direct Hankel/SVD implementations reproduce the paired clinical arrays, all DCT bootstrap statistics and the declared controlled-world checks. Appendix **??** specifies the checked inputs and numerical discrepancies.

The clinical finding makes the representation response worth explaining; the controlled worlds now test what its collective sensitivity and observation dependence permit us to infer.

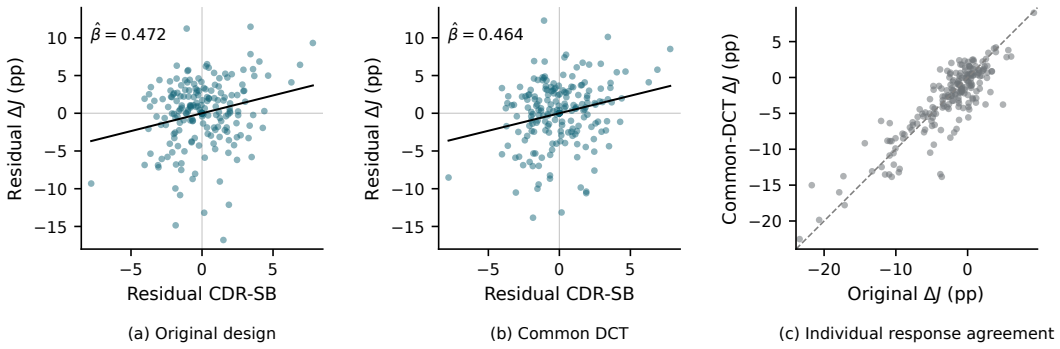

Figure 2: Adjusted clinical association and individual preprocessing responses. (a,b) Each point is one of the same 201 participants: CDR-SB and ΔJ are residualized against the fixed clinical nuisance design. Lines show the saved adjusted coefficients; uncertainty is reported in Table 1. Fourteen zero-information rows overlap at the origin; no jitter or participant exclusion is applied. (c) Each person's original and common-DCT ΔJ, with the equality diagonal. Similar cohort slopes need not imply identical individual responses. Units are percentage points (pp).

6.3 Collective sensitivity survives the complete finite pipeline

In the latent-sharing world, $S_i = \sqrt{1-r}Z_i + \sqrt{r}H$, all sources have the same stationary Gaussian AR(0.5) law. Each channel's full univariate path law is therefore unchanged as r varies, while its dependence with other channels changes. At $N = 197$, latent-oracle J rises from 0.489 to 24.482 percentage points as r increases from 0 to 0.8. With the observation component and proxy held fixed, measured A5 J rises from 3.041 to 21.711 points (Figure 3a). The direction is retained at $N = 788$. This is a positive control for collective sensitivity beyond changing regional spectra; it does not require nonlinearity.

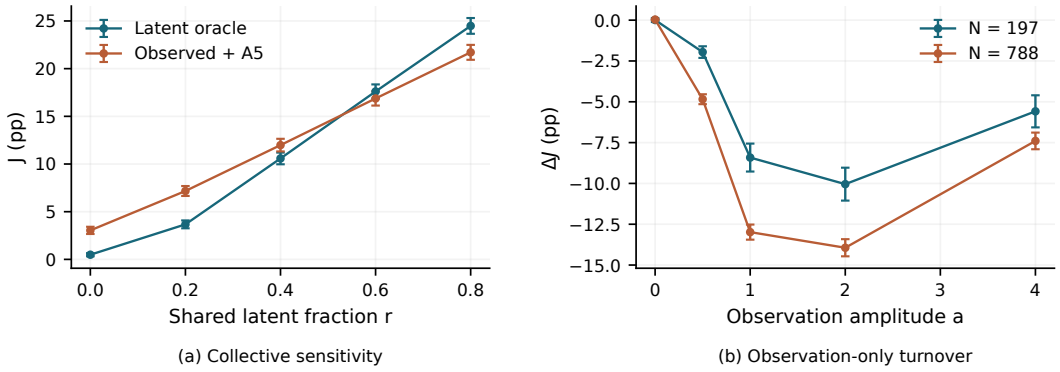

Figure 3: Controlled collective sensitivity and observation dependence. (a) Latent sharing changes joint organization at fixed univariate Gaussian path laws. (b) Increasing observation amplitude gives a nonmonotone response at fixed target. Panel (a) uses $N = 197$; (b) also shows $N = 788$. Target-retention and observational-alias readouts appear in Figure 5. All declared levels are retained. Bars show ± 1 simulation SE across 16 independent realizations, not clinical confidence intervals.

The observation-amplitude world holds $S = Z$ fixed and changes $X = Z + aU$, with proxy $U + 0.5V$. At $N = 197$, the contrast first decreases and then rises from -10.045 to -5.585 points as a increases from 2 to 4. The paired increment is $+4.460$ points (simulation SE 0.239); at $N = 788$ it is $+6.543$ (SE 0.151). The oracle measurement of the fixed latent target is unchanged across levels. The reversal is consistent with normalized spectral saturation; it is not an improvement in removal efficacy or a fit to the clinical slope.

Proxy contamination and observational aliasing explain why score recovery does not identify target preservation (Figure 5). Increasing proxy target loading from 0 to 4 lowers known target-energy retention from 0.9517 to 0.7231 while raising residual artifact-energy retention from 0.2024 to 0.8957. A rebound in J can therefore accompany poorer preservation and removal. The alias construction additionally gives different latent-oracle scores under identical observations and proxies. Appendix F retains both readouts and their interpretation.

The remaining worlds distinguish other mechanisms. Changing proxy quality changes J_A at exactly fixed X and J_0. The independent-channel null retains a small finite offset: at $N = 197$, mean BASE $J = 0.645$ and A5 $J = 0.581$ points, with mean contrast -0.065 (SE 0.117). At $N = 788$, the contrast is $+0.037$ (SE 0.024). Its inert levels are identical inputs, not extra independent replications. We retain these offsets rather than force finite J to zero or use a synthetic bias correction on human data.

6.4 Window-FC dispersion improves matched future-severity prediction

The spectral analysis establishes a clinical association of a joint representation's preprocessing response. We next ask whether joint variability can also contribute to future-severity prediction. For this task we use FC dispersion: the 21-coordinate vector of standard deviations of windowed correlations between seven networks, using 60-second windows at 30-second strides. In 48 index-MCI participants from the same ADNI programme, we evaluate whether this vector adds predictive information to clinical/technical covariates. This is a complementary assessment of another joint representation; the spectral ΔJ analysis concerns contemporaneous severity in 201 participants.

The target is the assessment nearest two years within 548–913 days, conditional on baseline CDR-SB measured within the preceding 90 days. Clinical/technical covariates and each augmented MRI model use identical participants, outcomes and five-fold participant splits, repeated for three seeds. Inner three-fold CV selects the PLS dimension and ridge penalty; transformations are fitted on training data only. Actual days to follow-up enter the baseline, so this is a retrospective conditional forecast. Appendix J gives the complete design and controls.

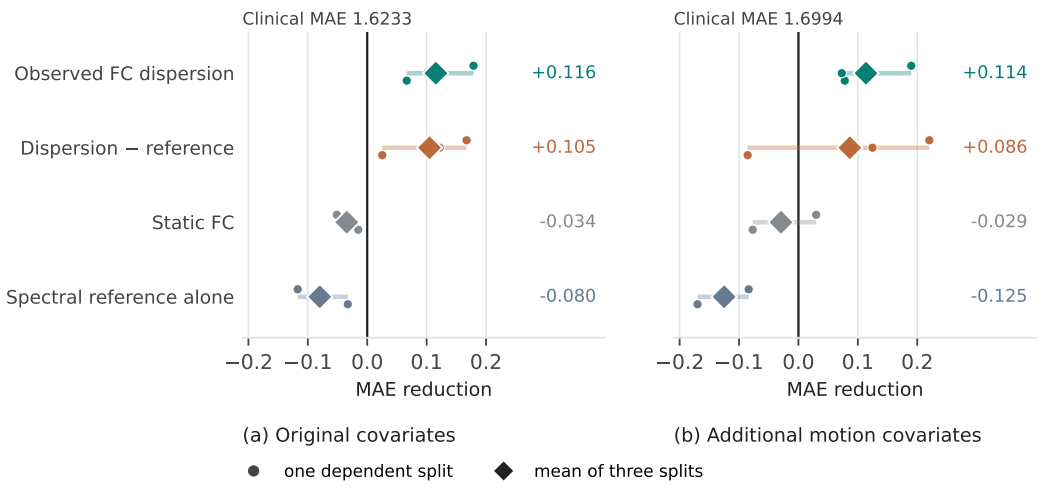

Figure 4: Matched two-year CDR-SB prediction errors. Positive values are paired reductions in two-year CDR-SB MAE relative to clinical/technical covariates alone. Small circles are the three dependent repeated nested-CV splits of the same 48 participants; diamonds are their means, and horizontal spans show their range, not confidence intervals. Both covariate tiers use the same participants and split identities. The mean cross-periodogram-preserving spectral reference does not recover the observed gain; the observed-minus-reference vector retains a mean advantage, but one extra-motion split is negative. Complete model comparisons appear in Table 7.

Adding observed dispersion lowers mean MAE from 1.6233 to 1.5077 CDR-SB points with the original covariates, and from 1.6994 to 1.5858 after additional motion adjustment (Figure 4). It improves on both the clinical baseline and matched static FC in all three splits under both covariate tiers. Static FC alone does not improve the mean error, and combining static FC with dispersion is worse on average than dispersion alone. Thus the positive result concerns a specific joint variability representation under a matched evaluation, rather than the mere addition of MRI features.

Each of 39 shared-phase reference records preserves the full-record cross-periodogram. Their mean dispersion vector does not recover the observed predictive gain. The observed-minus-reference vector retains a mean improvement of 0.1049/0.0865 points under original/additional-motion covariates, although one additional-motion split is negative. The comparison locates predictive information in the observed vector beyond its conditional spectral-reference mean. The gain is heterogeneous: 24 and 25 of 48 participants have lower individual error under original and additional-motion covariates, respectively, and most mean gain occurs among those without subsequent worsening. These dependent splits evaluate a previously explored candidate, so the result supports internal future-level prediction. Independent evaluation should test its generalization, especially among participants who deteriorate; attribution of the predictive information requires a separate source model.

7 Discussion

The central empirical result is a positive association between severity and the paired preprocessing response, with similar slopes under the original and common-DCT designs (Table 1). Its persistence despite changing absolute score levels supports evaluating the response as a feature in its own right.

The Gaussian construction and finite experiments establish collective sensitivity at fixed single-channel laws. The observation-only turnover and proxy controls explain how this sensitivity depends on the measurement process. Exact observational aliases identify the additional source information needed to interpret that response biologically. Together, these results connect the positive clinical association to a specified, reproducible observable.

The study design determines the next evaluation of this positive association. The analysis was developed in an already studied ADNI programme; the paired contrast is post-exposure and the common-DCT control is a later sensitivity on the same cohort. Clinical dates have unresolved source provenance at some rows, and the approximate wild-bootstrap reference has documented small-tail calibration limits. The five synthetic proxies do not reproduce tissue PCA or the full physiology of human scans. These considerations motivate two next tests: evaluate the frozen response in an independently specified cohort, and use additional source information to characterize what the selected nuisance components contain.

A reproducible response feature specifies its operator pair, representation and reference law. Independent-channel phases discount marginal spectra; a cross-spectrum-preserving reference asks a different question. These choices define what the clinical coefficient measures and should accompany any comparison across pipelines.

The positive severity slope coexists with a negative mean A5–BASE response. The observation-only turnover explains how these directions can coexist, motivating source measurements alongside clinical replication.

Replication should freeze participant selection, temporal support, nuisance-component selection, network map, lag length and reference law, and retain cross-branch covariance in the clinical fit. Source attribution additionally needs a justified proxy–target relation, such as external physiology or an exclusion restriction. These tests address reproducibility and neural interpretation, respectively.

The FC-dispersion result (Section 6.4) complements this account with lower future-level prediction error than matched clinical and static-FC baselines. The two empirical tasks connect representation evaluation to clinical information at different levels: an adjusted contemporaneous association for the spectral response, and an internally evaluated prognostic gain for windowed connectivity. Their common lesson is constructive: retain joint organization, specify the observation and reference models, and evaluate the resulting feature against the clinical comparator appropriate to its task. This yields reproducible questions about clinical association, predictive value and source attribution without requiring all three to be answered by one statistic.

REPRODUCIBILITY STATEMENT

Anonymous artifacts are available at https://anonymous.4open.science/r/preprocessing-sensitivity-review-D760/. They include aggregate results, portable measurement code, runnable controlled worlds, verification scripts and scoped proof sources. The appendices specify cohort selection, nuisance designs, references, inference and prediction. Appendix **??** distinguishes public execution from the restricted clinical pipeline and documents independent numerical checks. Frozen research records preserve protocols, source bindings and phase/bootstrap draws. Individual-level ADNI data and restricted cohort/fold ledgers are not redistributed.

ETHICS STATEMENT

This work analyzes existing ADNI data; it introduces no new acquisition or intervention. No additional institutional review board approval was required for this secondary analysis. The original studies and data-access procedures are described by the source initiatives (Jack et al., 2008; Weiner et al., 2017). We report cohort selection and the limits of generalization from longitudinal availability, acquisition heterogeneity and clinical assessment provenance. The Dementia category is not treated as independent confirmation of Alzheimer's pathology. The reported associations do not establish a validated diagnostic tool or a preferred treatment decision. We provide aggregate outputs and avoid distributing participant identifiers or restricted imaging data. Independent clinical replication and biological attribution remain separate scientific requirements.

AI USE STATEMENT

Generative AI assistants were used in literature search, formal verification, and text polishing.

REFERENCES

Yashar Behzadi, Khaled Restom, Joy Liau, and Thomas T. Liu. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, 37(1):90–101, 2007. doi: 10.1016/j.neuroimage.2007.04.042. URL https://doi.org/10.1016/j.neuroimage.2007.04.042.

Rotem Botvinik-Nezer, Felix Holzmeister, Colin F. Camerer, Anna Dreber, Juergen Huber, Magnus Johannesson, Michael Kirchler, Roni Iwanir, Jeanette A. Mumford, R. Alison Adcock, Paolo Avesani, Blazej M. Baczkowski, Aahana Bajracharya, Leah Bakst, Sheryl Ball, Marco Barilari, Nadège Bault, Derek Beaton, Julia Beitner, Roland G. Benoit, Ruud M. W. J. Berkers, Jamil P. Bhanji, Bharat B. Biswal, Sebastian Bobadilla-Suarez, Tiago Bortolini, Katherine L. Bottenhorn, Alexander Bowring, Senne Braem, Hayley R. Brooks, Emily G. Brudner, Cristian B. Calderon, Julia A. Camilleri, Jaime J. Castrellon, Luca Cecchetti, Edna C. Cieslik, Zachary J. Cole, Olivier Collignon, Robert W. Cox, William A. Cunningham, Stefan Czoschke, Kamalaker Dadi, Charles P. Davis, Alberto De Luca, Mauricio R. Delgado, Lysia Demetriou, Jeffrey B. Dennison, Xin Di, Erin W. Dickie, Ekaterina Dobryakova, Claire L. Donnat, Juergen Dukart, Niall W. Duncan, Joke Durnez, Amr Eed, Simon B. Eickhoff, Andrew Erhart, Laura Fontanesi, G. Matthew Fricke, Shiguang Fu, Adriana Galván, Remi Gau, Sarah Genon, Tristan Glatard, Enrico Glerean, Jelle J. Goeman, Sergej A. E. Golowin, Carlos González-García, Krzysztof J. Gorgolewski, Cheryl L. Grady, Mikella A. Green, João F. Guassi Moreira, Olivia Guest, Shabnam Hakimi, J. Paul Hamilton, Roeland Hancock, Giacomo Handjaras, Bronson B. Harry, Colin Hawco, Peer Herholz, Gabrielle Herman, Stephan Heunis, Felix Hoffstaedter, Jeremy Hogeveen, Susan Holmes, Chuan-Peng Hu, Scott A. Huettel, Matthew E. Hughes, Vittorio Iacovella, Alexandru D. Iordan, Peder M. Isager, Ayse I. Isik, Andrew Jahn, Matthew R. Johnson, Tom Johnstone, Michael J. E. Joseph, Anthony C. Juliano, Joseph W. Kable, Michalis Kassinopoulos, Cemal Koba, Xiang-Zhen Kong, Timothy R. Koscik, Nuri Erkut Kucukboyaci, Brice A. Kuhl, Sebastian Kupek, Angela R. Laird, Claus Lamm, Robert Langner, Nina Lauharatanahirun, Hongmi Lee, Sangil Lee, Alexander Leemans, Andrea Leo, Elise Lesage, Flora Li, Monica Y. C. Li, Phui Cheng Lim, Evan N. Lintz, Schuyler W. Liphardt, Annabel B.

Losecaat Vermeer, Bradley C. Love, Michael L. Mack, Norberto Malpica, Theo Marins, Camille Maumet, Kelsey McDonald, Joseph T. McGuire, Helena Melero, Adriana S. Méndez Leal, Benjamin Meyer, Kristin N. Meyer, Glad Mihai, Georgios D. Mitsis, Jorge Moll, Dylan M. Nielson, Gustav Nilsonne, Michael P. Notter, Emanuele Olivetti, Adrian I. Onicas, Paolo Papale, Kaustubh R. Patil, Jonathan E. Peelle, Alexandre Pérez, Doris Pischedda, Jean-Baptiste Poline, Yanina Prystauka, Shruti Ray, Patricia A. Reuter-Lorenz, Richard C. Reynolds, Emiliano Ricciardi, Jenny R. Rieck, Anais M. Rodriguez-Thompson, Anthony Romyn, Taylor Salo, Gregory R. Samanez-Larkin, Emilio Sanz-Morales, Margaret L. Schlichting, Douglas H. Schultz, Qiang Shen, Margaret A. Sheridan, Jennifer A. Silvers, Kenny Skagerlund, Alec Smith, David V. Smith, Peter Sokol-Hessner, Simon R. Steinkamp, Sarah M. Tashjian, Bertrand Thirion, John N. Thorp, Gustav Tinghög, Loreen Tisdall, Steven H. Tompson, Claudio Toro-Serey, Juan Jesus Torre Tresols, Leonardo Tozzi, Vuong Truong, Luca Turella, Anna E. van 't Veer, Tom Verguts, Jean M. Vettel, Sagana Vijayarajah, Khoi Vo, Matthew B. Wall, Wouter D. Weeda, Susanne Weis, David J. White, David Wisniewski, Alba Xifra-Porxas, Emily A. Yearling, Sangsuk Yoon, Rui Yuan, Kenneth S. L. Yuen, Lei Zhang, Xu Zhang, Joshua E. Zosky, Thomas E. Nichols, Russell A. Poldrack, and Tom Schonberg. Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582(7810):84–88, 2020. doi: 10.1038/s41586-020-2314-9. URL https://doi.org/10.1038/s41586-020-2314-9.

D. S. Broomhead and Gregory P. King. Extracting qualitative dynamics from experimental data. *Physica D: Nonlinear Phenomena*, 20(2–3):217–236, 1986. doi: 10.1016/0167-2789(86)90031-X. URL https://doi.org/10.1016/0167-2789(86)90031-X.

Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, Eurika Kaiser, and J. Nathan Kutz. Chaos as an intermittently forced linear system. *Nature Communications*, 8(1):19, 2017. doi: 10.1038/s41467-017-00030-8. URL https://doi.org/10.1038/s41467-017-00030-8.

Nathan W. Churchill, Robyn Spring, Babak Afshin-Pour, Fan Dong, and Stephen C. Strother. An Automated, Adaptive Framework for Optimizing Preprocessing Pipelines in Task-Based Functional MRI. *PLOS ONE*, 10(7):e0131520, 2015. doi: 10.1371/journal.pone.0131520. URL https://doi.org/10.1371/journal.pone.0131520. Correction: https://doi.org/10.1371/journal.pone.0145594.

Rastko Ciric, Daniel H. Wolf, Jonathan D. Power, David R. Roalf, Graham L. Baum, Kosha Ruparel, Russell T. Shinohara, Mark A. Elliott, Simon B. Eickhoff, Christos Davatzikos, Ruben C. Gur, Raquel E. Gur, Danielle S. Bassett, and Theodore D. Satterthwaite. Benchmarking of participant-level confound regression strategies for the control of motion artifact in studies of functional connectivity. *NeuroImage*, 154:174–187, 2017. doi: 10.1016/j.neuroimage.2017.03.020. URL https://doi.org/10.1016/j.neuroimage.2017.03.020.

Russell Davidson and Emmanuel Flachaire. The wild bootstrap, tamed at last. *Journal of Econometrics*, 146(1):162–169, 2008. doi: 10.1016/j.jeconom.2008.08.003. URL https://doi.org/10.1016/j.jeconom.2008.08.003.

Oscar Esteban, Christopher J. Markiewicz, Ross W. Blair, Craig A. Moodie, A. Ilkay Isik, Asier Erramuzpe, James D. Kent, Mathias Goncalves, Elizabeth DuPre, Madeleine Snyder, Hiroyuki Oya, Satrajit S. Ghosh, Jessey Wright, Joke Durnez, Russell A. Poldrack, and Krzysztof J. Gorgolewski. fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature Methods*, 16(1):111–116, 2019. doi: 10.1038/s41592-018-0235-4. URL https://doi.org/10.1038/s41592-018-0235-4.

fMRIPrep developers. *fMRIPrep 25.2.5 documentation: Outputs*. NiPreps, 2026. URL https://fmriprep.org/en/25.2.5/outputs.html. Version-specific documentation; accessed September 25, 2026. Year denotes access year.

Michael N. Hallquist, Kai Hwang, and Beatriz Luna. The nuisance of nuisance regression: Spectral misspecification in a common approach to resting-state fMRI preprocessing reintroduces noise and obscures functional connectivity. *NeuroImage*, 82:208–225, 2013. doi: 10.1016/j.neuroimage.2013.05.116. URL https://doi.org/10.1016/j.neuroimage.2013.05.116.

Clifford R. Jack, Jr., Matt A. Bernstein, Nick C. Fox, Paul Thompson, Gene Alexander, Danielle Harvey, Bret Borowski, Paula J. Britson, Jennifer L. Whitwell, Chadwick Ward, Anders M. Dale, Joel P. Felmlee, Jeffrey L. Gunter, Derek L.G. Hill, Ron Killiany, Norbert Schuff, Sabrina Fox-Bosetti, Chen Lin, Colin Studholme, Charles S. DeCarli, Gunnar Krueger, Heidi A. Ward, Gregory J. Metzger, Katherine T. Scott, Richard Mallozzi, Daniel Blezek, Joshua Levy, Josef P. Debbins, Adam S. Fleisher, Marilyn Albert, Robert Green, George Bartzokis, Gary Glover, John Mugler, and Michael W. Weiner. The Alzheimer's Disease Neuroimaging Initiative (ADNI): MRI methods. *Journal of Magnetic Resonance Imaging*, 27(4):685–691, 2008. doi: 10.1002/jmri.21049. URL https://doi.org/10.1002/jmri.21049.

Martin A. Lindquist, Stephan Geuter, Tor D. Wager, and Brian S. Caffo. Modular preprocessing pipelines can reintroduce artifacts into fMRI data. *Human Brain Mapping*, 40(8):2358–2376, 2019. doi: 10.1002/hbm.24528. URL https://doi.org/10.1002/hbm.24528.

Andrea I. Luppi, Helena M. Gellersen, Zhen-Qi Liu, Alexander R. D. Peattie, Anne E. Manktelow, Ram Adapa, Adrian M. Owen, Lorina Naci, David K. Menon, Stavros I. Dimitriadis, and Emmanuel A. Stamatakis. Systematic evaluation of fMRI data-processing pipelines for consistent functional connectomics. *Nature Communications*, 15(1):4745, 2024. doi: 10.1038/s41467-024-48781-5. URL https://doi.org/10.1038/s41467-024-48781-5.

James G MacKinnon and Halbert White. Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3):305–325, 1985. doi: 10.1016/0304-4076(85)90158-7. URL https://doi.org/10.1016/0304-4076(85)90158-7.

Leonardo Novelli and Adeel Razi. A mathematical perspective on edge-centric brain functional connectivity. *Nature Communications*, 13(1):2693, 2022. doi: 10.1038/s41467-022-29775-7. URL https://doi.org/10.1038/s41467-022-29775-7.

Sid E. O'Bryant, Stephen C. Waring, C. Munro Cullum, James Hall, Laura Lacritz, Paul J. Massman, Philip J. Lupo, Joan S. Reisch, Rachelle Doody, and Texas Alzheimer's Research Consortium. Staging dementia using Clinical Dementia Rating Scale Sum of Boxes scores: a Texas Alzheimer's research consortium study. *Archives of Neurology*, 65(8):1091–1095, 2008. doi: 10.1001/archneur.65.8.1091. URL https://doi.org/10.1001/archneur.65.8.1091.

Dean Prichard and James Theiler. Generating surrogate data for time series with several simultaneously measured variables. *Physical Review Letters*, 73(7):951–954, 1994. doi: 10.1103/PhysRevLett.73.951. URL https://doi.org/10.1103/PhysRevLett.73.951.

Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and B T Thomas Yeo. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28(9):3095–3114, 2018. doi: 10.1093/cercor/bhx179. URL https://doi.org/10.1093/cercor/bhx179.

Thomas Schreiber and Andreas Schmitz. Surrogate time series. *Physica D: Nonlinear Phenomena*, 142(3–4):346–382, 2000. doi: 10.1016/S0167-2789(00)00043-9. URL https://doi.org/10.1016/S0167-2789(00)00043-9.

The mathlib Community. The Lean Mathematical Library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2020, pp. 367–381. Association for Computing Machinery, 2020. doi: 10.1145/3372885.3373824. URL https://doi.org/10.1145/3372885.3373824.

Michael W. Weiner, Dallas P. Veitch, Paul S. Aisen, Laurel A. Beckett, Nigel J. Cairns, Robert C. Green, Danielle Harvey, Clifford R. Jack, Jr., William Jagust, John C. Morris, Ronald C. Petersen, Jennifer Salazar, Andrew J. Saykin, Leslie M. Shaw, Arthur W. Toga, John Q. Trojanowski, and Alzheimer's Disease Neuroimaging Initiative. The Alzheimer's Disease Neuroimaging Initiative 3: Continued innovation for clinical trial improvement. *Alzheimer's & Dementia*, 13(5):561–571, 2017. doi: 10.1016/j.jalz.2016.10.006. URL https://doi.org/10.1016/j.jalz.2016.10.006.

Alba Xifra-Porxas, Michalis Kassinopoulos, and Georgios D Mitsis. Physiological and motion signatures in static and time-varying functional connectivity and their subject identifiability. *eLife*, 10:e62324, 2021. doi: 10.7554/eLife.62324. URL https://doi.org/10.7554/eLife.62324.

Dingling Yao, Dario Rancati, Riccardo Cadei, Marco Fumero, and Francesco Locatello. Unifying Causal Representation Learning with the Invariance Principle. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/85381f4549b5ddf1d48e2e287d7d3d15-Abstract-Conference.html.

B. T. Thomas Yeo, Fenna M. Krienen, Jorge Sepulcre, Mert R. Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L. Roffman, Jordan W. Smoller, Lilla Zöllei, Jonathan R. Polimeni, Bruce Fischl, Hesheng Liu, and Randy L. Buckner. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165, 2011. doi: 10.1152/jn.00338.2011. URL https://doi.org/10.1152/jn.00338.2011.

Mengsen Zhang, Samir Chowdhury, and Manish Saggar. Temporal Mapper: Transition networks in simulated and real neural dynamics. *Network Neuroscience*, 7(2):431–460, 2023. doi: 10.1162/netn_a_00301. URL https://doi.org/10.1162/netn_a_00301.

A DEFINITIONS, PROOFS, AND IDENTIFICATION BOUNDARIES

The appendices follow the evidence supporting the main results: definitions and proofs, cohort and measurement details, inference, and the complete controlled worlds. We then document the related representations explored in the same programme, including the full forecast design, matched controls, spatial context and endpoint-specific outcomes. The final appendix maps these analyses to the supplied reproducibility artifacts.

A.1 FINITE-RECORD OBSERVABLE

Let $F \in \mathbb{R}^{N\times p}$ be a retained parcel record. Here $p = 400$; a fixed matrix $W_{\text{net}} \in \mathbb{R}^{p\times 7}$ averages parcels within each nonempty network. Given a nuisance design D including an intercept, let Π_D be its column-space orthogonal projector and set

$$R_D = I_N - \Pi_D, \qquad Y_D = R_D F W_{\text{net}}.$$

The definition depends on the column space, not on uniqueness of nuisance coefficients. The implementation also specifies a numerical rank cutoff. For $Y \in \mathbb{R}^{N\times 7}$, define

$$\bar{Y}_j = \frac{1}{N}\sum_{t=0}^{N-1} Y_{tj}, \quad s_j^2 = \frac{1}{N}\sum_{t=0}^{N-1}(Y_{tj} - \bar{Y}_j)^2, \quad \widetilde{Y}_{tj} = (Y_{tj} - \bar{Y}_j)/s_j.$$

All s_j must be positive. For $L \geq 2$, $N > L$, the vertical joint delay matrix is

$$H_L(Y)_{jL+\ell,t} = \widetilde{Y}_{t+\ell,j}, \quad j = 0,\dots,6,\ \ell = 0,\dots,L-1,\ t = 0,\dots,N-L.$$

Its rows are not centered again. With ordered singular values σ_k,

$$\widehat{\rho}_2(Y) = \frac{\sigma_1(H_L(Y))^2 + \sigma_2(H_L(Y))^2}{\|H_L(Y)\|_F^2}.$$

The implementation uses $L = 10$, hence lags $0,\dots,9$. Let $\mathcal{S}_b$ independently randomize each channel's non-DC Fourier phases while preserving its Fourier magnitudes; the real Nyquist coefficient is also retained when N is even. Centering and unit-RMS normalization are applied inside $\widehat{\rho}_2$ to each resulting record. The score is

$$J_B(Y) = \widehat{\rho}_2(Y) - \frac{1}{B}\sum_{b=1}^{B}\widehat{\rho}_2(\mathcal{S}_b Y), \qquad B = 199. \tag{10}$$

This is excess leading-two joint-delay energy relative to a marginal-spectrum reference. Independent-channel phases do not preserve the cross-spectrum. The score does not specifically isolate higher-order, nonlinear, or causal interactions, and is not a fractal-dimension or attractor-dimension estimator.

A.2 POPULATION SPECTRAL CALCULATION

Let $\gamma(h)$ be a stationary autocovariance with $\gamma(0) = 1$. Then $T_L = (\gamma(\ell - k))_{\ell,k=0}^{L-1}$ is positive semidefinite, has trace L, and has eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$. If a population field satisfies

$$\mathrm{Cov}(Y_i(t), Y_j(t+h)) = C(r)_{ij}\gamma(h), \qquad C(r) = (1-r)I_7 + r\mathbf{1}\mathbf{1}^\top, \quad 0 \leq r \leq 1,$$

its channel-major delay-vector covariance is $C(r) \otimes T_L$.

Lemma 1 (Leading two population eigenvalues)**.** *Set $A = 1 + 6r$ and $B_0 = 1 - r$. The leading two eigenvalues of $C(r) \otimes T_L$, counting multiplicity, are*

$$A\lambda_1, \qquad \max\{A\lambda_2, B_0\lambda_1\}.$$

The leading-two fraction of $I_7 \otimes T_L$ is $2\lambda_1/(7L)$.

Proof. The vector $\mathbf{1}/\sqrt{7}$ is a spatial eigenvector with eigenvalue A; its six-dimensional orthogonal complement has eigenvalue B_0. Tensor products with temporal eigenvectors yield $A\lambda_j$ and six copies of $B_0\lambda_j$. Since $A \geq B_0 \geq 0$, the largest is $A\lambda_1$. After removing one occurrence, every remaining eigenvalue is bounded by either $A\lambda_2$ or $B_0\lambda_1$, both of which occur. This covers ties and zero eigenvalues. The trace is $7L$. For independent channels, each λ_j occurs seven times, so the two largest are both λ_1. □

Subtracting the independent-channel fraction gives (6). Its two affine candidates have slopes $5\lambda_1 > 0$ and $6(\lambda_1 + \lambda_2) > 0$, because $\operatorname{tr} T_L > 0$. Their maximum is strictly increasing and equals zero at $r = 0$. Both leading eigenvalues can be taken from the common spatial direction exactly when

$$(1 + 6r)\lambda_2 \geq (1 - r)\lambda_1 \quad \Longleftrightarrow \quad r \geq r_* := \frac{\lambda_1 - \lambda_2}{\lambda_1 + 6\lambda_2}. \tag{11}$$

This is an eigenvalue-ordering boundary, not a dynamical phase transition. If $\lambda_2 = 0$, then $r_* = 1$; a common-branch formula must not be used at finite $r < 1$. The maximum expression remains valid.

A.3 Fixed marginal laws with changing collective organization

Proof of Proposition 1. Independence and the common autocovariance give

$$\operatorname{Cov}(S_i^{(r)}(t), S_j^{(r)}(t + h)) = \begin{cases} \gamma(h), & i = j, \\ r\gamma(h), & i \neq j. \end{cases}$$

Each channel is a zero-mean Gaussian process with autocovariance γ, irrespective of r. Its mean and covariance determine each finite-dimensional distribution. The joint delay covariance is $C(r) \otimes T_L$, and the preceding calculation proves strict increase of J_{pop}. □

Any observable depending only on the individual channel laws is constant in this construction. This does not assert that finite-record univariate estimators return identical values on different realizations.

A.4 Observation-only turnover

Proof of Proposition 2. Population regression of $X_i = Z_i + \sqrt{s}\, U$ on $W = U + \tau V$ gives

$$b = \frac{\operatorname{Cov}(X_i, W)}{\operatorname{Var}(W)} = \frac{\sqrt{s}}{1 + \tau^2}, \qquad X_i - bW = Z_i + \sqrt{s}\left(U - \frac{W}{1 + \tau^2}\right).$$

The parenthesized process has autocovariance $t\gamma(h)$, where $t = \tau^2/(1 + \tau^2) \in (0, 1)$, and is independent of every Z_i. Population normalization therefore gives correlations

$$r_B(s) = \frac{s}{1 + s}, \qquad r_A(s) = \frac{ts}{1 + ts},$$

with unchanged marginal autocovariance γ. For $s > 0$, $r_A(s) < r_B(s)$, so strict increase of J_{pop} yields $\delta(s) < 0$. At zero both correlations vanish; at infinity both tend to one. Continuity gives $\delta(0) = \lim_{s\to\infty} \delta(s) = 0$. Thus δ cannot be nondecreasing, as its positive-s values are negative; it cannot be nonincreasing, as its limit is zero. For any $s_1 > 0$, that limit supplies an $s_2 > s_1$ with $\delta(s_2) > \delta(s_1)$. □

The proof uses population regression coefficients, not the finite simultaneous regressions in the experiments. If both $r_A(s)$ and $r_B(s)$ satisfy (11) on the interval of interest, put $K = 6(\lambda_1 + \lambda_2)/(7L) > 0$. Then

$$\delta(s) = -\frac{K(1 - t)s}{(1 + s)(1 + ts)}, \tag{12}$$

$$\delta'(s) = \frac{K(1 - t)(ts^2 - 1)}{(1 + s)^2(1 + ts)^2}, \tag{13}$$

$$\delta(s_2) - \delta(s_1) = \frac{K(1 - t)(s_2 - s_1)(ts_1 s_2 - 1)}{(1 + s_1)(1 + ts_1)(1 + s_2)(1 + ts_2)}. \tag{14}$$

These identities follow by a common denominator and differentiation. In particular, $0 < s_1 < s_2$ and $ts_1 s_2 > 1$ give a positive increment while both contrasts remain negative. If the spectral branch changes, use (6), not this single-branch expression.

A.5 EXACT OBSERVATIONAL ALIASES

Fix all processes and auxiliary observed inputs on the same probability space. Model O declares target $S^O = Z$, observation component $\sqrt{s}\,U\mathbf{1}^\top$, and proxy $W = U + \tau V$. Model N declares

$$S^N = \frac{Z + \sqrt{s}\,U\mathbf{1}^\top}{\sqrt{1+s}}, \qquad G^N = \sqrt{1+s}, \qquad X^N = G^N S^N = X^O,$$

with the same proxy and all other observed inputs. Both targets have unit-variance channels and marginal autocovariance γ, but their cross-channel correlations for $s > 0$ are respectively zero and $s/(1+s)$.

The observation equality is pointwise. Any deterministic pipeline $f(X, W, \text{auxiliary inputs})$ returns the same result in both models. A randomized pipeline has the same output distribution; coupling its seed gives pointwise equality. This includes a learned rule given the same observed training and test information. It cannot return both distinct target values correctly on identical input. If s varies with an observed covariate, the construction holds conditionally on that covariate: one model keeps the target law fixed while the other changes its collective law.

This is a nonidentifiability construction, not a complete model of human physiology. It permits the proxy to contain part of the target in Model N. An independently justified exclusion restriction, external measurement, or intervention could rule out one assignment. A tissue label alone does not supply such a restriction. The finite experiment uses shared observed bytes for aliases; numerical division followed by multiplication need not itself be bitwise exact.

A.6 NESTED PROJECTIONS WITH A COMMON TEMPORAL BASIS

Let $D_0 = [\mathbf{1}, \mathrm{trend}, M24, C]$ and $D_1 = [D_0, A5]$. Set $P = I - \Pi_{D_0}$ and $Q = I - \Pi_{D_1}$. Inclusion of nuisance spaces gives reverse inclusion of residual spaces. Consequently

$$P^\top = P, \quad Q^\top = Q, \quad P^2 = P, \quad Q^2 = Q, \quad PQ = QP = Q.$$

The effective added projector $E = P - Q$ is self-adjoint and

$$E^2 = P - PQ - QP + Q = P - Q = E, \qquad QE = EQ = 0.$$

It projects onto the part of the former residual space removed by the additional nuisance span, with $\operatorname{rank} E = \operatorname{rank} D_1 - \operatorname{rank} D_0$. For every finite parcel matrix F,

$$PF = QF + EF, \qquad \langle QF, EF\rangle_F = 0, \qquad \|PF\|_F^2 = \|QF\|_F^2 + \|EF\|_F^2. \tag{15}$$

Expand the squared Frobenius norm and use $Q^\top E = 0$ to obtain the identity. Fixed aggregation commutes with projection by associativity: $(PF)W_{\mathrm{net}} = P(FW_{\mathrm{net}})$.

This identity orders total unnormalized energy. It supplies no ordering for a statistic that additionally divides by channel scales, reorders eigenvalues, normalizes their sum, and subtracts an operator-dependent reference. Separately applying noncommuting residual projectors is not generally projection onto the combined nuisance span. The experiment uses one joint design per branch and restricts stored cosine columns to the original retained indices, without regenerating them on a different window.

A.7 ONE WEIGHTED PAIRED CLINICAL CONTRAST

Let the rows index the same n people for both operators. Fix severity $g \in \mathbb{R}^n$ and a common clinical nuisance projector Π_Z. Put $R = I - \Pi_Z$, $c = Rg$, $d = c^\top c > 0$, and $a = c/d$. The OLS coefficient of g in a fit of y on g, Z is $\widehat{\beta}(y) = a^\top y$: the normal equation against c is $c^\top(y - Zb - \widehat{\beta}g) = 0$, with $c^\top Z = 0$ and $c^\top g = d$. Rank deficiency among nuisance columns does not matter if their space is fixed and $d > 0$.

For $J_o = \rho_o - \bar{q}_o$, linearity gives

$$\widehat{\beta}(J_A - J_B) = \widehat{\beta}(J_A) - \widehat{\beta}(J_B) = \widehat{\beta}(\rho_A - \rho_B) - \widehat{\beta}(\bar{q}_A - \bar{q}_B). \tag{16}$$

This uses fixed weights and a common outcome scale. It does not hold in general for differences of partial correlations, whose denominators depend on the outcome. A fixed positive-definite WLS

weighting matrix V gives the analogous formula with the V-orthogonal severity residual and $a = Vc/(c^\top Vc)$. The reported analysis uses OLS.

Let M be the common full-model residual projector, $u_o = MJ_o$, and $h_i = 1 - M_{ii}$ the fitted-model leverage. For active rows $h_i < 1$, put $v_i = [a_i/(1-h_i)]^2$. The HC3 variance estimator satisfies

$$\widehat{V}_\Delta = \sum_i v_i (u_{Ai} - u_{Bi})^2 = \widehat{V}_A + \widehat{V}_B - 2\sum_i v_i u_{Ai} u_{Bi}.$$

Directly fitting the paired difference retains cross-operator covariance. An implementation may assign a saturated row zero contribution only after verifying zero clinical information and corresponding zero residual support; unsupported saturation is rejected. This algebraic identity is not a calibration theorem for HC3 or the wild bootstrap.

A.8 Finite-reference uncertainty and the population boundary

Condition on the observed records and fixed clinical design. Write $e_o = \bar{q}_o - \mathbb{E}_{\mathrm{MC}}[\bar{q}_o]$; reference simulation induces coefficient error $a^\top(e_B - e_A)$. If $\Sigma_{oo'} = \mathrm{Cov}_{\mathrm{MC}}(e_o, e_{o'})$, then

$$V_{\mathrm{MC},\Delta} = a^\top(\Sigma_{AA} + \Sigma_{BB} - \Sigma_{AB} - \Sigma_{BA})a.$$

These matrices include all cross-person terms. Only with independent simulation streams between people and operators does this reduce to

$$\widehat{V}_{\mathrm{MC},\Delta} = \sum_i a_i^2 (s_{Ai}^2 + s_{Bi}^2)/B,$$

where each s^2 is the sample variance across independent reference draws, with denominator $B - 1$. For deliberately shared old/new phases within an operator, use the sample variance of paired reference-score differences divided by B, retaining their covariance. Distinct integer seeds describe implementation, not a proof of probabilistic independence. These conditional Monte Carlo standard errors are separate from uncertainty across people; they are not population confidence intervals and are not automatically added to HC3.

Finally, (6) does not assert $\mathbb{E}[J_B] = J_{\mathrm{pop}}$. The ordered-eigenvalue ratio is nonlinear; overlapping delay vectors have finite boundaries; and each independent-phase reference does not have covariance exactly $I_7 \otimes T_L$. Even before sample scaling, conditioning on a finite common residual projector can replace temporal covariance T_N by $PT_NP^\top$, which need not be stationary Toeplitz. Phase randomization need not be an exact symmetry of that conditional law. The controlled finite-record experiments, including their nonzero null offsets, supply distinct evidence. No finite-null mean, asymptotic convergence, or clinical type-I calibration is proved by the population calculation.

A.9 Scope of the optional mechanized supplement

`NormalizationAlias.lean` checks attenuation, its conditional secant sign, pointwise observation alias, and consequences for a supplied decoder. `NestedResidual.lean` checks effective projection and energy identities under supplied self-adjointness, idempotence, and nesting assumptions. `PairedFWL.lean` checks the coefficient formula conditional on normal equations, paired/observed-reference linearity, and finite covariance identities retaining cross terms.

These files do not formalize the Gaussian process construction, tensor eigenvalue ordering, numerical SVD rank selection, finite-J_B behavior, bootstrap calibration, or biological interpretation. The mathematical arguments appear above; empirical and numerical assertions require their separate evidence. Formal finite covariance uses an ensemble-size denominator, not the sample-variance denominator one less. Machine checking establishes the stated conditional algebra, not its empirical applicability.

B Cohort construction and clinical support

The source catalogue contains 4,347 BOLD series spanning 4,257 visits from 1,764 ADNI participants. A metadata gate required a dated diagnostic classification within a nominal 90-day window,

an accepted diagnostic-concordance tier, exact source linkage, an anatomical-reference candidate, positive TR and at least 90 candidate volumes. The last condition is an input-support requirement, not a universal guarantee of sufficient data for every estimator. The gate retained 3,810 visits from 1,698 participants; 907 participants had at least two eligible visits.

The longitudinal construction retained all 90 eligible Dementia histories, with index same-study T1 support and age/recorded-sex support, and selected 90 CN and 90 MCI histories with exactly matched visit-depth distributions. Within each depth stratum, selection minimized acquisition and demographic mismatch without using brain-derived outcomes. Each index group contains 49 participants with two visits, 23 with three, 13 with four and five with five, totaling 244 visits per group. All eligible visits of the selected people were retained. Selection was diagnostic- and metadata-aware; it was not diagnosis-blind or random sampling from ADNI.

Table 2: Selection stages and endpoint-specific populations. The 201-person panel is selected from already chosen first eligible visits, not by searching later visits for TR$= 3$ s.

Stage	Participants	Observations
Source catalogue	1,764	4,257 visits
Metadata-eligible source	1,698	3,810 visits
At least two eligible visits	907	3,019 visits
Depth-matched longitudinal campaign	270	732 visits
Checked CDR-SB visits	—	615 visits
First eligible visit per person	267	267 visits
Exactly TR$= 3$ s among these visits	201	201 visits

Exact longitudinal-depth matching did not fully match acquisition. Relative to the Dementia anchors, CN and MCI matches shared site in 33/90 and 41/90 pairs, scanner model in 69/90 and 69/90, phase in 89/90 and 89/90, and field-map status in 89/90 and 88/90. Median absolute age differences were 3.96 and 3.45 years. These discrepancies remain part of the observation process and motivate the declared adjustment; covariate inclusion does not prove that all selection or acquisition effects have been removed.

The later cross-sectional CDR-SB gate required a finite nearest same-phase assessment, a nominal interval of at most 90 days, no inconsistent closest match, no known assessment QC failure or explicit not-done flag, and exact agreement between CDR-SB and the sum of six component scores. Three of the 270 people had no eligible visit: one lacked same-phase CDR, one had only assessments beyond the nominal window, and one failed the component-sum check. No additional person among the 267 was excluded for incomplete declared covariates. Of the selected visits, 251 were the first campaign visit and 16 a later first eligible visit. Ordering was by scan date and then the fixed record identifier, not by the score.

The 267-person panel has 87 CN, 91 MCI and 89 Dementia visits. Its exact-TR subset has 64, 70 and 67, respectively; these contemporaneous labels need not equal the index strata. The primary subset spans ADNI2 (78), ADNI3 (120) and ADNI4 (3). An age-from-birth-year proxy indicator is present for 3/201 and 43/267 people. Clinical-date reconstruction does not distinguish assessment dates from registry fallbacks for every record. No temporal availability claim for clinical prediction follows from these nominal dates.

C Observation model and numerical support

The seven network sizes in the fixed atlas are 52, 91, 46, 26, 47, 77 and 61 parcels, in the order Control, Default, Dorsal Attention, Limbic, Salience/Ventral Attention, Somatomotor and Visual. Network aggregation is an arithmetic mean over the fixed membership. The representation concerns these cortical averages and their temporal dependence, rather than a complete neural state.

M24 contains the six motion paths, backward derivatives, squares and squared derivatives in the retained source order. Both temporal designs contain an intercept and a linear trend. The small number of permitted undefined leading motion-derivative values is handled by the original source rule; arbitrary missing confound values are not imputed. The retained indicator must be finite and exactly

binary, with a contiguous suffix and common support across the parcel field and every regressor. Historical Boolean and numeric 0/1 encodings are equivalent after this check.

All selected aCompCor columns are labeled retained, anatomical CompCor, combined mask in the source JSON. The same pinned fMRIPrep 25.2.5 image produced all 201 scans. Runtime all-component output is followed by the fixed-five downstream selector; generic 50% boilerplate does not describe that runtime choice. The actual native-field component-generation path removes non-steady-state rows, applies the supplied temporal cosine/mean model, scales voxel time series and computes a temporal SVD. The downstream fixed five are conditional inputs in this analysis, not newly estimated components in the DCT sensitivity.

The common-DCT models use all existing cosine columns in source TSV order. There are 3 columns for 3 people, 5 for 75, 8 for 122, and 9 for one. The selected time indices are identical to the original H1 analysis. Rebuilding a DCT basis at the shortened length can change its span and is not the operation evaluated here. All 201 paired designs add five effective A5 directions after the common span. Minimum algebraic residual dimensions are 68 for BASE+DCT and 63 for A5+DCT; these are not effective independent sample sizes.

For temporal projection, the implementation centers candidate columns, screens numerically empty columns, normalizes retained columns and constructs an SVD basis with the explicit source cutoff. Ranks must agree at tolerance multipliers 0.1, 1, 10 and 100. The clinical design has its own fixed column-space convention. We compare subspaces and projected fields rather than basis signs. The DCT execution requires positive residual dimension and relative symmetry, idempotence, nesting, aggregation and raw-energy errors at most 10^{-10}. The maximum observed errors were 1.12×10^{-15} for projector identities, 9.00×10^{-16} for aggregation and 3.51×10^{-14} for raw-energy decomposition. There is no gate on the sign of J, the clinical coefficient or its p value.

Input and component provenance were checked for all 201 scans; all 603 ROI, CompCor-variance and segmentation reportlets matched their source records. Tissue maps were present, but this campaign did not perform a new exhaustive visual or voxelwise review of all masks. Source consistency is not a claim that a tissue component contains no neural signal. Independent physiological information sufficient to resolve that interpretation was not established by the scoped source inventory.

The retained masks remove the leading non-steady-state prefix rather than censoring arbitrary internal time points. Zero through five leading volumes are removed in 77, 79, 35, 7, 2 and 1 scans, respectively. Motion is adjusted by the stated regressors and continuous participant-level summaries; this analysis does not include a new low-motion-only or internally censored sensitivity. Exhaustive selection-shift or observation-process sensitivity analyses for this endpoint are also not established by the source audit.

D PHASE REFERENCES AND FINITE-RECORD INTERPRETATION

For each normalized $N \times 7$ field, we take the real FFT along time. Each non-constrained frequency/channel phase is drawn independently from the declared uniform phase law; DC is unchanged and, for even N, the real Nyquist coefficient is unchanged. The inverse transform uses the original N. Each reference then undergoes the same centering, RMS scaling and lag-10 embedding as the observation. The leading-two fraction is computed from squared singular values, with a Gram-eigenvalue implementation independently checked against direct SVD.

Preserving the channel periodograms preserves the chosen Fourier amplitudes, not all finite non-periodic delay statistics. A sample near the center of the record enters more Hankel columns than an endpoint. Across the original 468 observation/operator cases and their references, total Hankel-energy ratios ranged approximately from 0.915 to 1.207; channelwise ratios ranged from 0.743 to 1.378. These changes coexist with Fourier-magnitude errors near machine precision. No alternate boundary convention was introduced into the reported frozen analysis.

The 199-reference standard error estimates Monte Carlo uncertainty of the phase mean conditional on the observed field and reference law. It is distinct from uncertainty across participants. The original human BASE and A5 streams are disjoint by person and operator. For the DCT sensitivity, each new branch reuses its corresponding old branch's phase seeds, giving common random numbers for old/new comparisons; BASE and A5 remain separate. Simulation worlds instead deliberately

share phase streams across operators and levels within each independent bank. The variance of a paired reference difference is used when streams are shared; independent-variance addition is used for the declared independent branches.

The measurement can detect linear cross-channel organization. Independent channel phases do not preserve the full cross-spectrum. Consequently, positive J is not evidence for nonlinearity beyond a matched full cross-spectrum. Harmonic and finite-record qualifications further bound any interpretation as attractor dimension. A Gaussian latent oracle can itself have a finite positive average J; the synthetic null offset was not subtracted from the human scores.

E CLINICAL COEFFICIENT, UNCERTAINTY AND CALIBRATION

The raw participant-level design has 74 columns and numerical rank 63 on the fixed 201 rows. Adding residualized CDR-SB gives full rank 64 and residual dimension 137. Fourteen rows have unit nuisance leverage and algebraically zero residual clinical information. They remain in the cohort; their zero contributions are verified rather than handled by dividing by a near-zero leverage denominator. The clinical coefficient is well defined because $z^\top z > 0$.

Let Q_Z be an orthonormal basis of the clinical nuisance span, $z = (I - Q_Z Q_Z^\top)c$, $d = z^\top z$, and $V = [Q_Z, z/\sqrt{d}]$. For an outcome $y = \Delta J$, the estimate is $\widehat{\beta} = z^\top y/d$. Write $h_i = \|V_{i,:}\|^2$, $h_i^0 = \|(Q_Z)_{i,:}\|^2$, full residual $e = (I - VV^\top)y$ and reduced residual $e^0 = (I - Q_Z Q_Z^\top)y$. On nonsaturated rows the HC3 variance is

$$\widehat{\mathrm{Var}}_{\mathrm{HC3}}(\widehat{\beta}) = \sum_i \left(\frac{z_i}{d(1-h_i)} \right)^2 e_i^2. \tag{17}$$

The restricted wild bootstrap generates independent Rademacher signs $\xi_i^{(b)}$ and pseudo-outcomes

$$y_i^{*(b)} = \widehat{y}_i^0 + \xi_i^{(b)} \frac{e_i^0}{\sqrt{1-h_i^0}}. \tag{18}$$

Every draw is fully refitted and studentized by its own HC3 SE. The fitted nuisance mean may be omitted algebraically because the full fit annihilates it; the independent reconstruction checks both forms. With $t = \widehat{\beta}/\widehat{\mathrm{SE}}_{\mathrm{HC3}}$, the two-sided reported reference value is

$$p = \frac{1 + \#\{b : |t^{*(b)}| \geq |t|\}}{1 + 19{,}999}, \tag{19}$$

with the fixed floating-point comparison tolerance in the implementation. Both paired studies use seed 20260918801. The DCT reference has sample SD 0.81893; its p is not a value inferred from a Student-t table or a normal interval applied to HC3 SE. HC3 and wild-bootstrap references draw on established methods (MacKinnon & White, 1985; Davidson & Flachaire, 2008); those references do not establish exact calibration for this application.

The original baseline clinical study used residual-sign randomization and partial-r^2 rather than this studentized wild statistic. Its p values are retained under that method and are not silently recomputed to optimize the narrative. The later paired contrast tests a different estimand. All data had previously participated in research development; study-local freezes prevent subsequent changes within a declared analysis but do not erase that history or establish project-wide multiple-testing control.

The prior finite simulation qualification included a heteroskedastic record-length world with 59/1,000 rejections at nominal 0.05 and 22/1,000 at nominal 0.01. These are a material limit on small-tail interpretation. The controlled observation/latent panel in Section 6 is not a calibration of the clinical null: an observation component associated with clinical state can induce a real nonzero observable β_Δ even when a latent-neural target does not change. A clinical-null calibration must specify a zero coefficient for the actual tested observable.

For paired slopes, HC3 covariance is retained:

$$\widehat{\mathrm{Var}}(\widehat{\beta}_A - \widehat{\beta}_0) = \widehat{\mathrm{Var}}(\widehat{\beta}_A) + \widehat{\mathrm{Var}}(\widehat{\beta}_0) - 2\widehat{\mathrm{Cov}}(\widehat{\beta}_A, \widehat{\beta}_0). \tag{20}$$

Conditional phase Monte Carlo uncertainty is propagated separately through the weights z_i/d. In the DCT comparison the paired slope MCSE is 0.003812 percentage points per CDR-SB point,

Table 3: All declared controlled worlds. The first proxy is supplemented by four independent AR(0.5) proxies. The observation is $X = S + O$ before synthetic motion. Scalars a, q, r, k are model parameters, not clinical severity.

World	Latent target S	Observation O; first proxy	Grid
Null	Z	0; V	$0, .25, .5, .75, 1$
O	Z	aU; $U + .5V$	$a = 0, .5, 1, 2, 4$
Q	Z	U; $\sqrt{q}U + \sqrt{1-q}V$	$q = 0, .25, .5, .75, 1$
L	$\sqrt{1-r}Z + \sqrt{r}H$	U; $U + .5V$	$r = 0, .2, .4, .6, .8$
C	$\sqrt{.75}Z + .5H$	U; $U + kH + .5V$	$k = 0, .5, 1, 2, 4$

compared with HC3 SE 0.220965. The maximum covariance bound for the phase MCSE is 0.005326 points. The descriptive old/new slope change has common-phase MCSE 0.003786 points; this is neither its participant-sampling SE nor an equivalence test.

F CONTROLLED WORLDS AND COMPLETE RESULT CONTEXT

All non-motion sources are independent stationary Gaussian AR(0.5) processes with unit variance, independent unit-normal initial state and innovation SD $\sqrt{1 - 0.5^2}$. There is no hemodynamic response, fractional process or chaotic system in this generator. Z has seven independent columns; U, V, H are scalar paths broadcast with loading one to each network. The same underlying paths are shared across declared levels within a bank to define paired counterfactuals. The 16 banks at each length are independent; the two lengths are not nested prefixes.

Synthetic motion comprises six stationary unit-variance AR(0.3) paths, their backward differences with first difference zero, squares and squared differences. Both operators include intercept, trend and these M24 columns. A5 adds the five declared proxies by a joint regression. At parcel level, each network shares its declared field; additional 0.3-scale parcel AR variation has its sample network mean removed at every time point. Consequently it aggregates to zero before and after any common linear temporal projection. This tests the full 400-to-seven implementation, not a realistic model of within-network anatomy.

The latent oracle applies the same normalization/Hankel/phase estimator directly to the known target without nuisance regression. It has privileged information unavailable in the human data. In the observation world, the alternative assignment $S^\dagger = Z + aU$ uses the identical observed bytes and proxy, but its oracle differs because the designated target differs. The normalized-target construction in Appendix A.5 expresses the same ambiguity without reducing it to a common amplitude change. **Proxy content changes the meaning of score recovery.** The contamination world makes the ambiguity quantitative. Its known target is $S = \sqrt{0.75}Z + 0.5H$, the observation adds U, and the first proxy is $U + kH + 0.5V$. Thus increasing k changes what the regressor removes while leaving the target and observed field fixed. At $N = 197$, increasing k from 0 to 4 lowers known target-energy retention from 0.9517 to 0.7231, while residual artifact-energy retention rises from 0.2024 to 0.8957. Figure 5a displays both energy readouts. The rebound of J therefore coexists with loss of target energy and poorer artifact removal. These component-specific readouts are available because the generator supplies the latent paths; they cannot be reconstructed from a higher observed score alone. Cross-energy terms are retained, so target and artifact energies are not assumed to add without covariance. This experiment directly rules out maximizing J as a sufficient criterion for choosing a target-preserving pipeline.

Figure 5b makes the identification boundary visible in the same finite measurement. The two oracle curves use different designated targets, but the observed field and proxy are identical at each amplitude. An oracle has privileged access to its target; an observed-data decoder does not. The divergent curves therefore illustrate missing identifying information, rather than failure of one chosen estimator. The normalized target construction in Appendix A.5 retains this ambiguity while fixing marginal target variance.

Changing proxy quality in Q leaves the observed field and BASE fixed. At $q = 1$ the known common observation component is removed to numerical precision. World C complements this

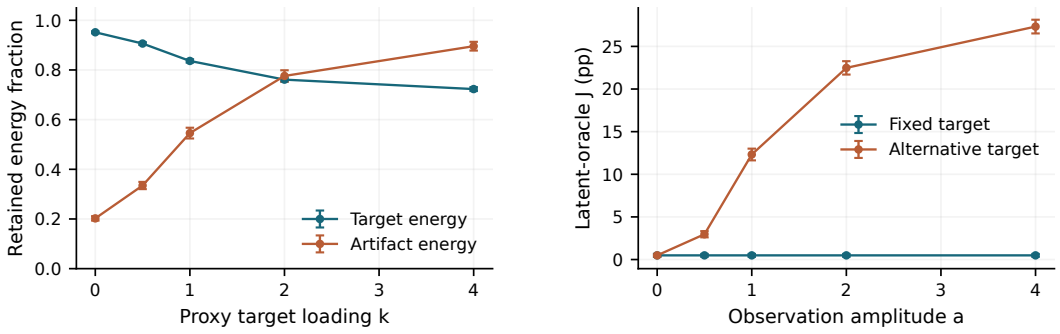

Figure 5: Proxy contamination and observational ambiguity. (a) Increasing proxy target loading decreases retained target energy and increases residual artifact energy; each fraction uses its own component-energy denominator. (b) Identical observed fields and proxies admit different designated latent targets and oracle scores. Both panels use $N = 197$ and retain all declared levels; bars show ± 1 simulation SE across 16 independent realizations. These privileged readouts explain why a higher observed score alone does not establish target preservation or latent identification.

control by changing the proxy's target overlap, with the component-specific retention results shown in Figure 5a.

Within each bank, reference phases are coupled across levels, operators and worlds. Simulation SEs of adjacent changes use their paired replicate differences. They include generator variability and finite-reference variability; conditional reference MCSE is not added a second time. No p values, participant confidence intervals or fitted clinical effects are attached to these worlds. The null's inert levels are identical by design. The population companion uses $C_{\text{normalized}} \otimes T_{10}$ with $(T_{10})_{ij} = 0.5^{|i-j|}$ and an independent-channel covariance reference. Its nonlinear spectral fraction is not the expectation of the finite projected and normalized estimator.

All 800 scenario summaries and paired contrasts were independently reconstructed. A fixed first-replicate subset covers 160 score cases at both lengths, each with all 199 references computed using literal Hankel construction and direct SVD; the maximum reference-score discrepancy is 4.72×10^{-16}. This is independent numerical reconstruction of fixed generated records, not 800 independent biological observations.

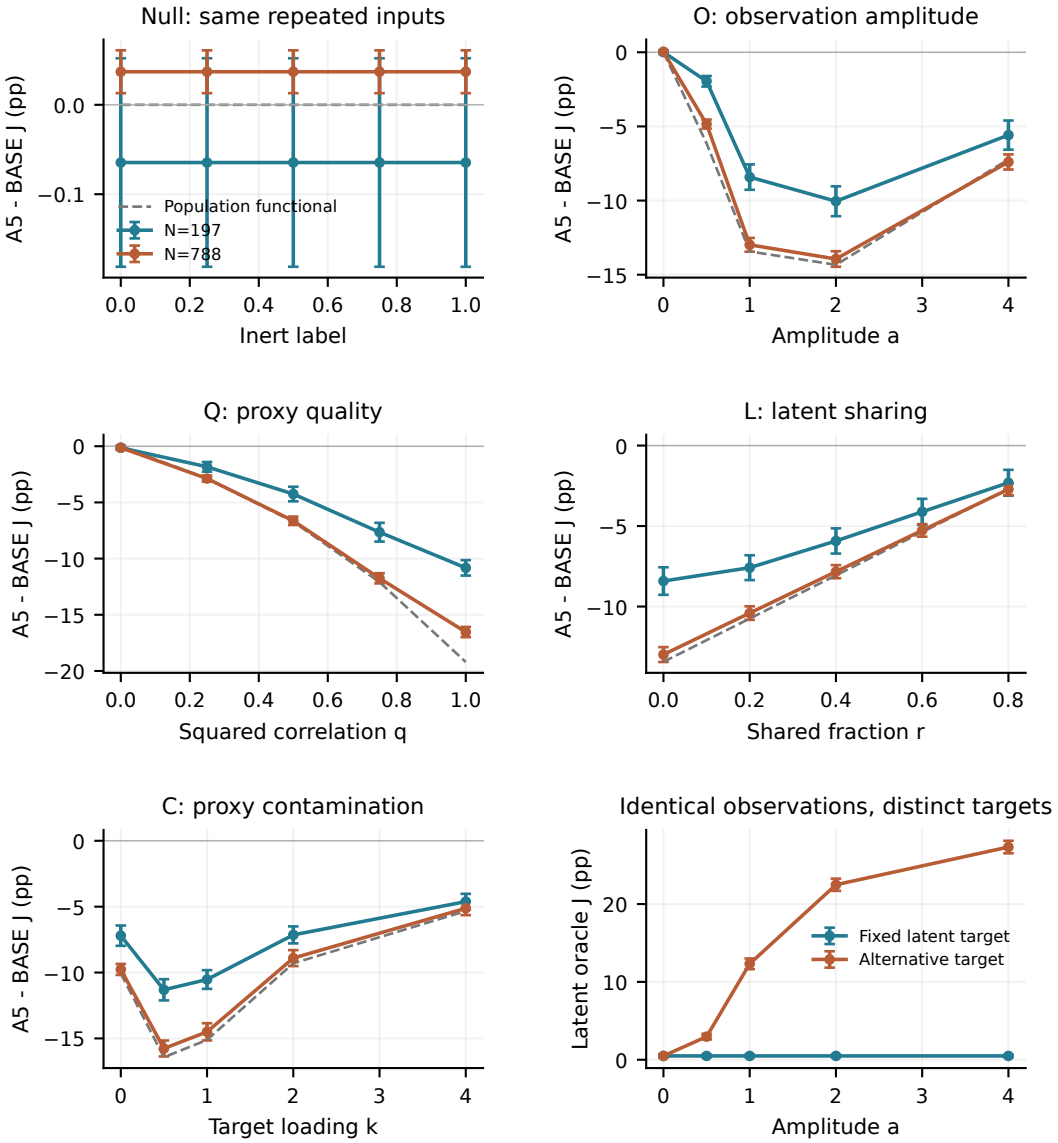

Figure 6: **Complete controlled response curves and observational alias.** All five world grids at $N = 197$ and $N = 788$ are retained; 16 independent realizations per length. Bars are ± 1 simulation SE. Dashed lines show the population covariance functional, not $\mathbb{E}[J]$ for the finite estimator. The bottom-right panel compares two privileged target readouts under exactly identical observations, $N = 197$. Repeated null labels do not add independent samples. These models use exogenous proxies and do not replay tissue PCA.

Table 4: Complementary temporal experiments use a different observable. None establishes repeatability, recoverability or clinical replication of the main spectral contrast.

Question	Preserved result and scope
Recurrence beyond cross-spectrum	Mean J_{rec} is 1.5203 pp for BASE and 1.1733 pp for A5, positive in 165/201 and 158/201 people. Clinical partial r is 0.11479 ($p = 0.30805$) and -0.03342 ($p = 0.73445$); clinical associations are unestablished.
Static monotone marginal transforms	Per-channel normal-score control retains recurrence excess: BASE 1.3000 pp and A5 1.0907 pp. This changes geometry and spectrum and does not eliminate all motion or physiological explanations. No new clinical test.
Intervisit recurrence agreement	Among 78 matched pairs, 54 are conditionally informative. Rank-recurrence BASE $r = 0.001604$, reference $p = 0.49995$; A5 $r = 0.047908$, $p = 0.37955$. Intervals span 77–1,006 days, not short-interval test–retest.
Synthetic recoverability	At the primary $N = 197$, noise level one, the tested recurrence readout does not recover the simulated between-system variation. An observable product-memory readout and the privileged latent oracle recover structure under the same declared experiment. Joint information can remain accessible when one scalar fails.

G Related temporal questions and their boundaries

The research program also evaluated temporal recurrence beyond a full cross-spectrum reference. We denote that separate statistic J_{rec}; it is not the spectral J used in the main paper. These results are retained to locate the scope of the current claims.

These controls preserve positive collective/temporal evidence and negative endpoint-specific results. They do not license a claim that every scalar summary of a dynamic system should detect its memory, or that a failure of one readout eliminates all individual dynamics. Conversely, the clinical spectral association in the present paper does not retrospectively turn the recurrence clinical nulls into positives. The harmonic and full-cross-spectrum distinctions in Appendix D remain relevant when interpreting the word "dynamics."

H Analysis sequence and status

Table 5: Successive studies and their inferential roles. A protocol frozen before a local computation is not a claim that the dataset was untouched by the research program.

Study	Role and history
Initial spectral evaluation	First eligible panel fixed from metadata. BASE clinical association primary; A5 and all-TR analyses declared sensitivities. Existing cohort already used in development.
Direct paired contrast	A5–BASE clinical coefficient tested after the original branch results were known. One new frozen paired endpoint; post-exposure discovery.
Controlled mechanisms	Same finite spectral statistic, five fixed worlds and all levels. Constructed after human contrast; no parameters fitted to human effect size or labels.
Common-DCT sensitivity	Motivated by component-generation/source audit after the preceding results. Same people, fixed A5 and clinical covariates; original cosine basis added to both branches. One new test; old/new difference descriptive.

Other clinical cohorts and earlier preprocessing definitions are not pooled into the current 201-person estimates. Independent replication requires new eligibility and exposure verification, rather than relabeling another split or visit of already studied people.

I Exploratory local directional organization

The preceding sections explain the spectral response. This extension examines where joint variation occurs across parcels, decomposing local directional variability into strength and orientation. It

provides spatial context for the programme before the following appendix develops the window-FC representation used in the forecast.

Population and exposure. This later analysis uses the existing 732-scan ADNI programme without new preprocessing. The common first-eligible-MRI clinical panel contains 269 people (89 CN, 90 MCI, 90 dementia); one CN index scan is too short for two local windows at 32 seconds and is not replaced by a later visit. The exact-TR=3-second panel contains 206 people (68 CN, 71 MCI, 67 dementia). These panels overlap with, but differ from, the 201-person concentration analysis. The 16-second scale is secondary. The normalized readout was specified after observing the weighted results and is explicitly post hoc.

Spatially resolved measurement. M24 jointly regresses an intercept, linear trend and 24 motion regressors; M24+A5 additionally includes the first five supplied aCompCor columns. For both operators, Morlet-6 coefficients at periods 16 and 32 seconds define parcel power shares $p_i(t) = |z_i(t)|^2 / \sum_j |z_j(t)|^2$. Samples outside twice the cone-of-influence boundary are excluded. Let $x(t)$ be the share vector centered over the valid interval. In complete 60-second windows with 30-second stride and 12-second lag,

$$A_w = \frac{1}{2m_w} \sum_t [x(t)x(t+\ell)^\top - x(t+\ell)x(t)^\top].$$

Durations are rounded by $\lfloor \text{seconds}/\text{TR} + 1/2 \rfloor$. All 400 parcels are retained. This is a lagged association of power shares, not causal transfer or complex-phase topology.

Exact decomposition. With $r_w = \|A_w\|_F$, $\overline{r} = \operatorname{mean} r_w$ and window mean $\overline{A}$,

$$V = \operatorname{mean} \|A_w - \overline{A}\|_F^2 = S + O, \quad S = \operatorname{Var}(r_w), \quad O = \overline{r}^{\,2} - \|\overline{A}\|_F^2.$$

For nonzero norms, $\|A_w - A_v\|_F^2 = (r_w - r_v)^2 + r_w r_v \|U_w - U_v\|_F^2$, where $U_w = A_w / r_w$. A zero norm has zero angular weight and no assigned direction. For isometric packed edges $a_{we} = \sqrt{2}(A_w)_{ij}$, $i < j$, split each nonnegative contribution

$$O_e = \frac{\overline{r}}{n} \sum_{w: r_w > 0} \left(\frac{a_{we}}{\sqrt{r_w}} - \sqrt{r_w} \frac{\overline{a}_e}{\overline{r}} \right)^2$$

equally between its incident parcels. The 400 regional contributions sum to O. They measure changes of the full edge pattern, including redistribution between edges, not only sign reversals.

References and clinical designs. Nineteen common-phase references on each cleaned 400-channel BOLD record preserve its complete cross-periodogram; the entire nonlinear wavelet/share calculation is repeated. This does not preserve spectra after the share transformation. Clinical excess is observed value minus the mean reference value. All three diagnosis contrasts use common panels and unadjusted, acquisition-adjusted and acquisition/motion-adjusted (AQ) designs. AQ includes age, sex, acquisition, site/scanner/phase and motion information. There are 18 saturated design rows in the native panel and 14 in the TR panel. Standard errors are descriptive; the 19 references do not provide calibrated significance.

Weighted and normalized findings. At 16 seconds, O accounts for about 68% of the native adjusted dementia–CN excess-V contrast under either operator, and 84%/78% for dementia–MCI. Weighted signs remain positive in the TR panel. However, $O(cA) = c^2 O(A)$ for $c > 0$. The post-hoc readout

$$N = \frac{O}{\overline{r}^{\,2}} = 1 - \left\| \frac{\sum_w r_w U_w}{\sum_w r_w} \right\|_F^2$$

removes common matrix scale while retaining relative window weights. Each reference is normalized before averaging. Zero denominators would be unsupported; none occurs on the selected panels.

The normalized dementia–CN excess at 32 seconds is positive under both operators and panels (native: $+0.004145$ and $+0.008371$), with substantial uncertainty. Normalized dementia–MCI at 16 seconds does not retain one sign in all sensitivities. The primary weighted 32-second result is not robust across processing. Thus the descriptive pattern is scale- and normalization-dependent, not evidence for universally greater or smaller neural chaos.

Diagnostic increment. The regional DFA/Buchstaber comparator with AQ is refitted on the same 269 people as every augmentation. Five repeated nested 5-outer/4-inner person splits use training-only transformations and bandwidths, with a comparator-only option in every kernel mixture. The

Table 6: Descriptive AQ dementia–CN excess contrasts at the secondary 16-second scale. Parentheses show one descriptive SE, not confidence limits. The normalized readout is post hoc.

Panel	Operator	O (10^{-8})	N (10^{-3})
Native	M24	$+5.215$ (4.000)	-7.042 (3.262)
Native	M24+A5	$+4.075$ (3.450)	-5.993 (3.153)
TR=3	M24	$+5.613$ (4.378)	-4.589 (2.783)
TR=3	M24+A5	$+2.334$ (3.865)	-3.020 (2.540)

additional blocks are excess strength and all 400 weighted orientation contributions at both periods. Normalized post-hoc features are not retrospectively substituted.

Contrast	Operator	Regional comparator	With strength/orientation
CN–dementia	M24	.599750	.597903
CN–dementia	M24+A5	.550112	.550112
CN–MCI	M24	.430262	.434782
CN–MCI	M24+A5	.470012	.455506
MCI–dementia	M24	.553531	.529901
MCI–dementia	M24+A5	.522494	.512642

These are mean pooled out-of-fold AUCs. The primary increment is not supported; within-fold AUC does not rescue it. Repetitions are dependent and are not clinical replications. This result does not erase the descriptive decomposition or validate a diagnostic marker.

Verification and remaining scope. All 2,926 supported operator/period conditions and 55,594 references are retained; previous local metrics reproduce exactly. Independent checks reconstruct every observed decomposition, regional sums and reported predictive outputs. Wolfram checks finite identities; Lean checks the finite pair identity under explicit squared-norm premises and its normalized-vector form, with nonnegativity conditional on a Cauchy bound. Neither validates clinical inference. Since O and N are unchanged by reordering the window matrices, temporal transition order and complex-phase topology remain distinct questions. Full aggregate coefficients and predictive summaries accompany the source package.

J EXPLORATORY SPATIAL CONNECTIVITY AND FUTURE-SEVERITY EXTENSION

This appendix develops the FC-dispersion forecast introduced in Section 6.4. We define its 21 connection-specific features and spectral reference, specify the matched 48-person evaluation, then examine prediction controls, temporal sampling, clinical heterogeneity and spatial patterns. These analyses explain the positive forecast comparison and identify which aspects warrant further evaluation.

J.1 SPATIAL REPRESENTATION AND CONDITIONAL CONTROLS

Spatially resolved dispersion. For each window, calculate the seven-network Pearson correlation matrix. The primary width and stride are $w = \max(8, \mathrm{round}(60/\mathrm{TR}))$ and $h = \max(1, \mathrm{round}(30/\mathrm{TR}))$ samples. For a selected set I of $K > 1$ windows, the representation has 21 coordinates:

$$v_I(e) = \left[\frac{1}{K-1} \sum_{t \in I} (r_{te} - \bar{r}_{I,e})^2 \right]^{1/2}, \quad e = (a, b),\ a < b. \tag{21}$$

It retains the identity of the varying connection but not the order of variations. It is not a measure of directed neural flow. The matched static representation contains the 21 full-record correlations.

Conditional spectral reference. For each prepared residual record, multiply its discrete Fourier coefficient vector at frequency f by one common random phase $e^{i\phi_f}$. Keep DC unchanged and the even-length Nyquist coefficient real; invert to a real time series. The identity

$$(e^{i\phi_f} \widehat{Y}_f)(e^{i\phi_f} \widehat{Y}_f)^* = \widehat{Y}_f \widehat{Y}_f^* \tag{22}$$

preserves the cross-periodogram and full-record covariance. We use 39 deterministic realizations per participant, applying the original windows to each. Compare the observed vector v, reference mean m, residual $v - m$, and their concatenation $[m, v - m]$. Disjoint 19- and 20-realization means assess Monte Carlo sensitivity. This conditional reference does not preserve every local artifact, marginal distribution or nonperiodic boundary condition; it is not a fully specified neural null. Shared phases differ from the independent-channel phases used for the main spectral observable.

Temporal support. The earlier eight-window comparison changes both span and raw-data reuse. The stricter follow-up uses four nonoverlapping windows of the same width w. Compact starts are $a + jw$, $j = 0, \ldots, 3$, with $a = 0$, $\lfloor (N - 4w)/2 \rfloor$ or $N - 4w$. Distributed starts are $\lfloor j(N - w)/3 \rfloor$. All 48 records support $N \geq 4w$, so every selection uses exactly $4w$ distinct frames. Span and spacing can change, but window number, width and unique-frame count cannot. Full-record nuisance fitting remains common: these are retrospective representation tests, not independently preprocessed short acquisitions.

J.2 FINITE WINDOW GEOMETRY

For one edge let $Z = (Z_1, \ldots, Z_K)$ denote window estimates with finite second moments, mean vector μ and covariance Σ. No independence or stationarity is assumed.

Proposition 3 (Finite window geometry). *For $K > 1$, $S^2 = (K-1)^{-1} \sum_i (Z_i - \bar{Z})^2$ satisfies*

$$\mathbb{E}S^2 = \frac{\operatorname{tr}\Sigma - K^{-1}\mathbf{1}^\top \Sigma \mathbf{1} + \sum_i (\mu_i - \bar{\mu})^2}{K - 1}. \tag{23}$$

If all means agree and all marginal variances equal σ^2, then

$$\mathbb{E}S^2 = \sigma^2 - \frac{\sum_{i \neq j} \Sigma_{ij}}{K(K-1)}. \tag{24}$$

Proof. Use $\sum_i (Z_i - \bar{Z})^2 = \sum_i Z_i^2 - K\bar{Z}^2$, take expectations, and substitute $\mathbb{E}Z_i^2 = \Sigma_{ii} + \mu_i^2$ and $\mathbb{E}\bar{Z}^2 = K^{-2}\mathbf{1}^\top \Sigma \mathbf{1} + \bar{\mu}^2$. The equal-marginal case cancels the mean term and separates the diagonal covariance sum. □

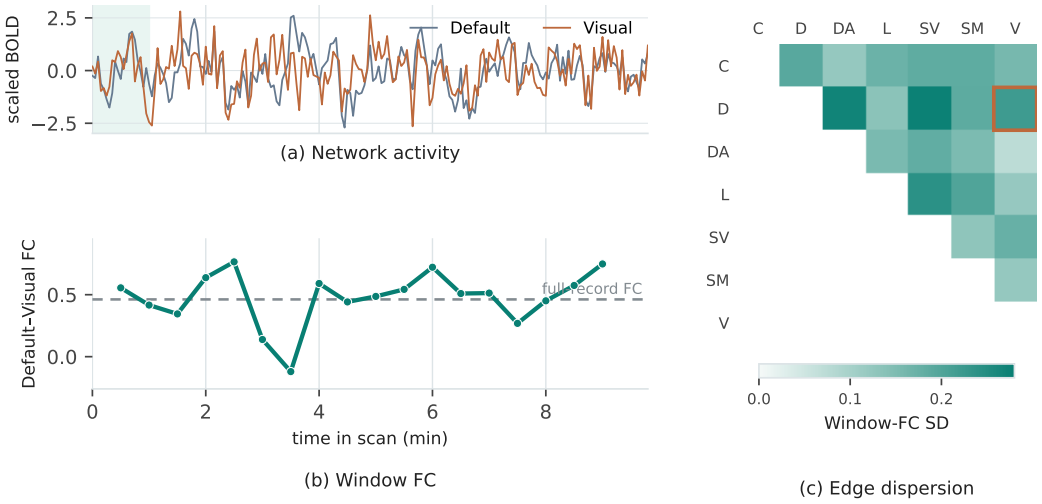

Figure 7: Construction of the window-FC dispersion vector. An anonymized scan was chosen by median retained-frame count and source order, without using its clinical outcome. The example has 197 frames, TR 3 s, and 18 overlapping 60 s windows advanced by 30 s; the first window is shaded. The Default–Visual pair illustrates local FC estimates (teal), its full-record FC (dashed grey), and one coordinate of the 21-edge dispersion map. The full-record value need not equal the mean window FC. Signals are scaled only for display. This single scan explains the measurement, not the clinical effect. Network abbreviations: C, Control; D, Default; DA, Dorsal Attention; L, Limbic; SV, Salience/Ventral Attention; SM, Somatomotor; V, Visual.

The identity is elementary but exposes a relevant ambiguity: changing window spacing or overlap changes covariance among estimates, even at the same window count. Changing local means adds another term. It does not predict the sign of a change in clinical MAE, nor does it directly identify $\mathbb{E}S$. We use it to design controls, not to interpret every dispersion difference as biology. Known-covariance numerical worlds, Wolfram checks and a scoped Lean formalization are provided in Appendix L; Lean verifies the algebra conditional on the stated moment equalities, not the entire probability derivation.

J.3 Clinical evaluation design

Source cohort and endpoint-specific panels. The ADNI programme comprises 732 visits from 270 participants, selected by metadata for three index-diagnosis groups and longitudinal depth. It is not a representative prevalence sample. The index-diagnosis panel has 90 CN, 90 MCI and 90 dementia participants; later visits are not substituted for unavailable index fields. We use the Schaefer 400-parcel/seven-network cortical atlas (Schaefer et al., 2018; Yeo et al., 2011). The main spectral-response analysis uses a distinct 201-person, exact-TR3 panel. This extension uses a prognostic subset of 48 index-MCI participants with eligible two-year CDR-SB. The panels overlap within one research programme. Exact selection, acquisition provenance and clinical-date limitations appear in Appendices B and M.

Repeated clinical maps. We retain a fixed inventory of 137 fields in six families: power/spatial shares, DFA, window-FC dispersion, spectral-geometric projections, joint-delay spectra and ordinal entropy. For each coordinate we estimate diagnosis contrasts with the declared demographic/acquisition/quality adjustment (AQ). All three contrasts—dementia–CN, MCI–CN and dementia–MCI—are retained. The full native panel has 270 index and 462 repeat visits; some fields have 267 participants. Within each index/repeat analysis, each person has total weight one. We compare the coefficient vectors using

$$C(\widehat{\beta}_0, \widehat{\beta}_R) = \frac{\widehat{\beta}_0^\top \widehat{\beta}_R}{\|\widehat{\beta}_0\| \|\widehat{\beta}_R\|}. \tag{25}$$

This describes direction, not effect size, significance or diagnostic accuracy. The 3,288 declared combinations retain current/index-fixed diagnoses, adjustment and exact-TR3 sensitivities; they are dependent comparisons.

Two-year prediction with fixed participants. Baseline CDR-SB is the latest eligible preceding assessment within 90 days of MRI; the target is the assessment nearest 730 days within $[548, 913]$. We predict the future level conditional on the baseline level, not dementia conversion. The clinical/technical baseline additionally includes age, sex, TR, retained sample count, mean FD, scanner, study phase and actual days to assessment. Thus the evaluation conditions on a retrospectively observed horizon. Additional-motion sensitivity adds the fraction of frames with FD above 0.2 mm and median standardized DVARS to *both* baseline and MRI models.

Five-fold participant CV is repeated for three fixed seeds. Inside each training fold, three-fold CV selects 2, 5 or 10 PLS components and ridge regularization from $\{0.1, 1, 10, 100\}$. Centering, scaling, categorical encoding and supervised projection use training data only. All representations and baselines use the same participants, outcomes and splits. The FC candidate was identified during earlier analyses; subsequent controls are post-selection sensitivity analyses, not an untouched test set.

Matched static comparator. On these same 48 participants we compare clinical covariates alone, covariates with 21 full-record Pearson FC edges, covariates with 21 window-dispersion edges, and covariates with their 42-coordinate concatenation. Each MRI representation uses the same train-only PLS/ridge selection. The declared primary contrast is combined minus static-FC MAE; dispersion versus static FC is a separate comparison. Both covariate sets are retained, and all four historical clinical/dispersion branches reproduce their saved predictions exactly. This is a post-selection comparison of specified learners, not a conditional-independence test.

Matched controls and temporal support. Ninety-nine joint row permutations of MRI repeat inner selection in all 24 family/projection/complexity models from the prior search. We report both matched-candidate and best-of-24 surrogate frequencies; the row shuffle breaks MRI–covariate dependence and is not a conditional-null p-value. A separate 38 within-person order-surrogate analysis preserves the FC-state multiset while altering order.

For temporal coverage we refit the same FC model on all windows, the first/middle/last eight consecutive windows, or eight distributed indices $\lfloor j(W-1)/7 \rfloor$, $j = 0, \ldots, 7$. This holds the window count fixed across shortened representations, not their dependence or span. One participant has only eight windows, so all selections coincide. A complementary diagnostic locks each full-record outer model and changes only the held-out measurement. It does not establish that a model can be trained on shortened acquisitions.

Clinical and spatial sensitivity. We rank paired participant-level OOF losses without dropping anyone from the primary analysis. Deleting a loss contribution is distinguished from refitting. Subgroup readouts cover baseline severity, subsequent worsening, study phase and scanner; future-outcome groups are descriptive and cannot be used at prediction time. A predeclared tighter baseline rule (at most 30 preceding days) is assessed both with fixed predictions and full nested retraining on the eligible subset. Original folds are restricted, not reselected.

For explicit signed maps, we additionally compare 21 static and 21 dispersion edges for all three index-diagnosis pairs. Each participant contributes one index vector and one mean over later visits, with equal person weight. These supplementary unadjusted standardized differences are distinct from the adjusted 137-field analysis. No coordinate selection is based on the prediction outcome.

J.4 Spatial variability improves future clinical severity prediction

Figure 4 makes the clinical comparison explicit: zero means no improvement over clinical covariates, and every observed-dispersion split lies on the improving side in both tiers. Dispersion also improves on matched static FC in all three splits (Table 7). Adding dispersion to static FC improves average MAE by 0.0639 with original covariates (two of three splits) and 0.1182 with extra motion (all three), but the combined model is worse on average than dispersion alone. Dispersion alone provides the strongest mean performance among these matched representations.

Historical matched MRI row-shuffle controls give 1/99 reference draws reaching the candidate gain and 5/99 for the best of the 24 searched configurations. These diagnostic frequencies are not conditional p-values and do not account for every historical research choice. Leave-

Table 7: Two-year CDR-SB mean MAE on the same 48 participants and three repeated nested-CV splits. Every row includes clinical/technical covariates; lower is better. The reference uses 39 shared-phase realizations.

Representation	Original	Extra motion
Clinical only	1.6233	1.6994
Static FC	1.6577	1.7289
Observed dispersion	**1.5077**	**1.5858**
Static + dispersion	1.5938	1.6107
Mean spectral reference	1.7032	1.8246
Observed − reference	1.5184	1.6129
Reference + residual	1.6806	1.7312

one-edge-out losses identify sensitivity to Limbic–Visual, Limbic–Salience/Ventral Attention and Salience/Ventral Attention–Somatomotor coordinates. This localizes candidate contributions for subsequent study; all 21 edges remain in the predictor.

J.5 Spectral controls help localize the predictive information

The mean reference predictor is worse than the clinical baseline in all three splits in both tiers. The residual $v - m$ retains a mean improvement; the original-covariate improvement occurs in all three splits, while the extra-motion improvement occurs in two. Combining the mean and residual does not improve on observed dispersion. A linear invertible re-expression of features need not give the same predictions after regularized supervised compression.

The observed-minus-reference comparison provides a concrete lead: clinically useful variation remains after subtracting the average dispersion expected under this conditional spectral reference. This identifies a representation to investigate, while the broader mechanistic interpretation remains open. Finite reference precision matters: disjoint 19/20-draw means give MAEs 1.8025/1.7144 (original) and 1.8469/1.7884 (extra motion). The feature-mean Monte Carlo SE has RMS 0.00612, and predictor convergence is not established. No best realization was selected.

J.6 Temporal sampling determines which representation is retained

The earlier eight-window results are 1.6358/1.6452/1.6520 for first/middle/last selections and 1.5629 for distributed windows, compared with full-record dispersion at 1.5077. Equal window count does not equal unique-frame count or dependence. Under the stricter four-window design, which holds unique frames fixed, original-covariate MAEs are 1.6650/1.6491/1.6250 for compact selections and 1.7318 for distributed selection. Extra-motion values are 1.7200/1.7040/1.6993 and 1.7565. Together, the two experiments distinguish a useful eight-window sampling result from a universal benefit of greater span: distributing four nonoverlapping windows does not improve performance. The covariance identity explains why number, dependence and local means belong in the design of such comparisons. The four-window experiment is sparser, so the pair of experiments guides further acquisition studies rather than isolating one cause of the eight-window result.

J.7 Clinical sensitivity identifies robustness and the next target

Restricting baseline assessment to 30 days before MRI retains 38 participants. Full nested retraining gives clinical/dispersion MAE 1.7031/1.5947 with original covariates and 1.7604/1.6268 with extra motion. Improvements occur in two and three splits, respectively. The mean improvement therefore persists after tightening the clinical timing criterion and retraining.

Mean gain remains positive when any one participant's fixed OOF loss contribution is removed (minimum 0.0775/0.0766), but only 24/25 of 48 people have lower error with dispersion (Figure 8). Removing the five largest positive contributions reverses the mean gain; this is an influence description without model refitting, not a proposed exclusion. Among 29 people whose CDR-SB worsens, gains are only 0.0181/-0.0127; among 19 without worsening they are 0.2644/0.3064. The supported endpoint is future level, not established detection of deterioration. ADNI2 and ADNI3 also

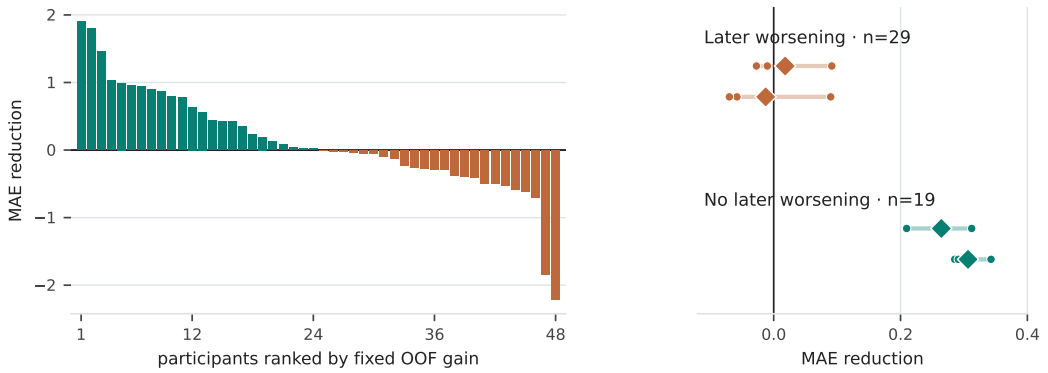

(a) Individual OOF error differences

(b) Future-outcome groups

Figure 8: Participant-level error differences and future-outcome strata. Left: fixed out-of-fold clinical-minus-dispersion absolute-error contributions for all 48 participants, ranked without excluding anyone; 24 have lower error with dispersion under original covariates. Right: subgroup means (diamonds) and the three dependent repeated splits (circles) for participants whose CDR-SB later worsened (29) or did not (19), under original covariates (upper mark) and extra motion adjustment (lower mark). This grouping uses the future outcome and is descriptive, not available to a prospective predictor. Contributions are not refitted after deletion and do not measure treatment or individual clinical benefit.

differ: original-covariate gains are 0.2729 and -0.0149 (extra motion: 0.2216 and 0.0150). These descriptive differences identify two priorities for the next evaluation: performance among people who deteriorate and consistency across acquisition strata.

J.8 Spatial clinical differences recur across visits

The forecast controls assess prediction on 48 participants. To describe the same feature coordinates in a broader clinical context, we now examine group maps in the 270-person index panel and their alignment across later visits. These maps describe spatial organization rather than attribute the fitted forecast. In the adjusted 137-field inventory, all 56 native-panel power/spatial-share dementia–CN map cosines are positive, with median 0.719; all 36 spectral-geometric variants are also positive. MCI-related contrasts depend more on representation and TR restriction. These direction comparisons characterize recurring group-level spatial structure across dependent field variants.

The explicit 21-edge maps retain 270 index observations and all 462 repeats, averaging repeats within person. Figure 9 exposes the coordinates of the dementia–CN comparison: both static FC and dispersion have positive and negative edge differences, and neither is a uniform shift across all network pairs. The dispersion pattern is distinct from static FC. These are unadjusted group descriptions, not a map of predictive feature contributions.

Index–repeat direction cosines for static FC are 0.905, 0.633 and 0.955 for dementia–CN, MCI–CN and dementia–MCI. Dispersion maps give 0.551, 0.283 and 0.569. Figure 10 shows all 21 edge coordinates for each contrast as index/repeat pairs. The weaker MCI–CN alignment is visible rather than hidden inside a dense heatmap. Directional recurrence does not imply stable effect magnitude or individual diagnosis; these descriptive maps complement the adjusted 137-field inventory.

J.9 Clinical change probes distinguish information from utility

A subsequent paired-visit analysis asks a different question from future prognosis: whether two observed scans help describe signed or absolute CDR-SB change. On 264 eligible same-TR intervals from 190 previously studied people, a scalar ordinal-entropy distance weakly lowers held-participant error for absolute change (-0.0106 CDR-SB points; descriptive person-bootstrap interval $[-0.0325, 0.0118]$) relative to clinical/technical and FC-distance controls. Fixed signed 14-group

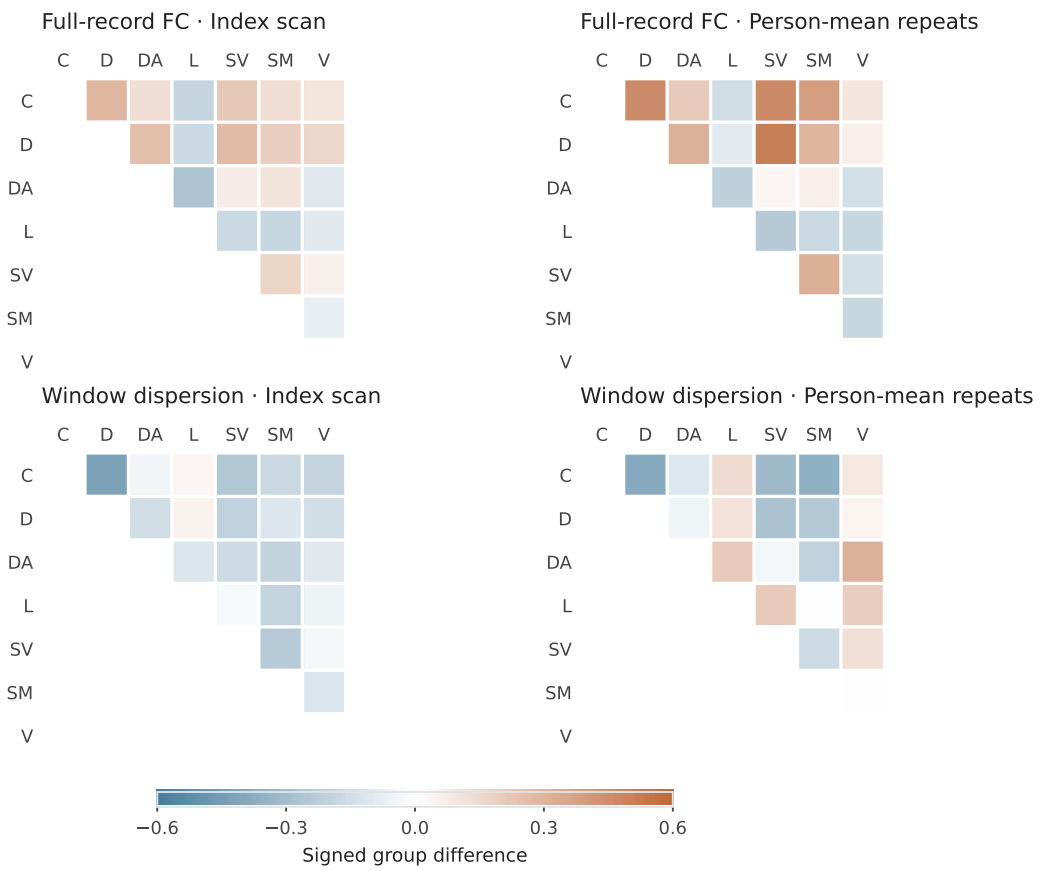

Figure 9: Signed dementia–CN network-pair differences. Upper-triangular matrices show signed unadjusted standardized differences for all 21 pairs of seven networks at index and person-mean repeat visits, separately for static FC and windowed dispersion. All 90 people in each index-diagnosis group are retained. Blue denotes lower values in dementia than CN; orange denotes higher values. This is spatial context, not an adjusted estimate, anatomical tract map, or predictive attribution. Abbreviations follow Figure 7. Full coordinate maps for all three contrasts appear in Appendix K.

and within-group C400 spatial representations instead worsen signed-change error by $+0.0736$ and $+0.0900$ points relative to controls that also include full C400 total variation. The within-group norm nevertheless identifies a person's true endpoint over a matched other-person endpoint in 96.1% of 409 comparisons. Thus scan distinctiveness alone did not yield clinical gain under these fitted representations; these same-cohort, different-target tests neither replicate nor negate the separate 48-person forecast (Appendix M.9).

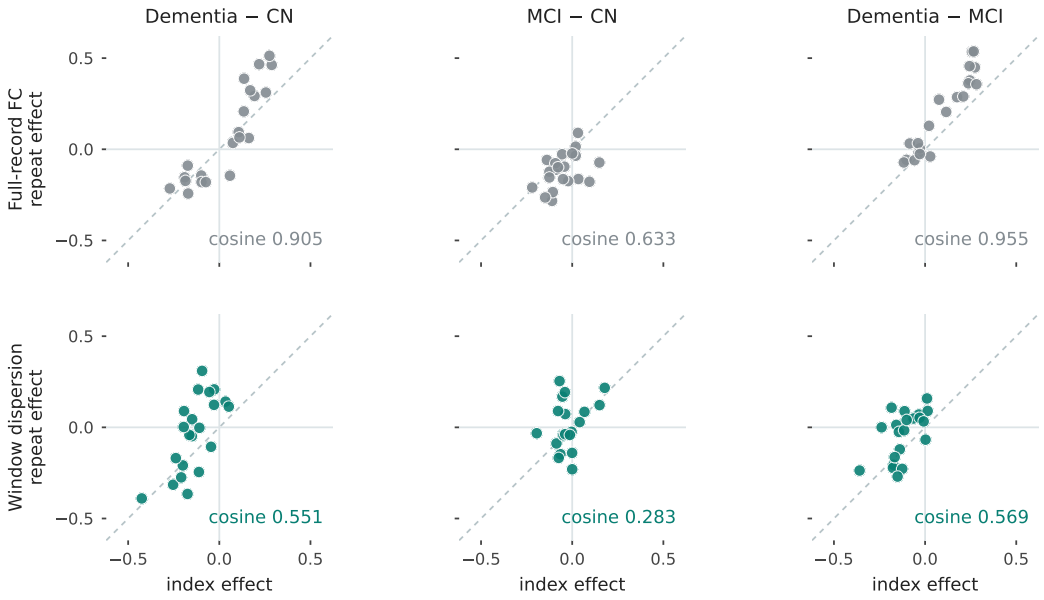

Figure 10: Index–repeat alignment of spatial group differences. Each dot is one network pair's signed unadjusted standardized group difference at the index visit (horizontal) and person-mean repeat visits (vertical). Both axes have the same scale; the dashed diagonal marks equal index/repeat effects. Cosines describe alignment of the complete 21-coordinate vectors, not correlations or inferential probabilities. Diagnosis labels are fixed from index visits and all three contrasts are shown; repeated visits are not independent cohorts.

K COMPLETE SPATIAL AND PREDICTION DISPLAYS

Appendix J separates the later dispersion measurement, clinical comparison, network map and index–repeat alignment so that each figure answers one question. The coordinate heatmap and absolute-error display below provide all declared contrasts and model rows alongside the selected comparisons in the main text.

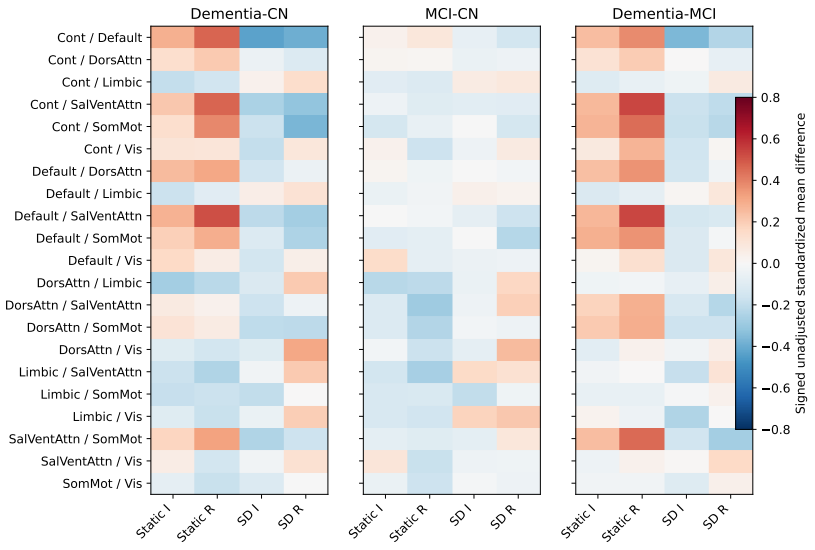

Figure 11: **All spatial contrasts and coordinates.** Signed unadjusted standardized differences for all 21 network pairs, all three index-diagnosis contrasts, and index (I) and person-mean repeat (R) views of full-record static FC and window dispersion. Each contrast has 90 people per index-diagnosis group. This display supplements the network matrices and recurrence plots in Appendix J.

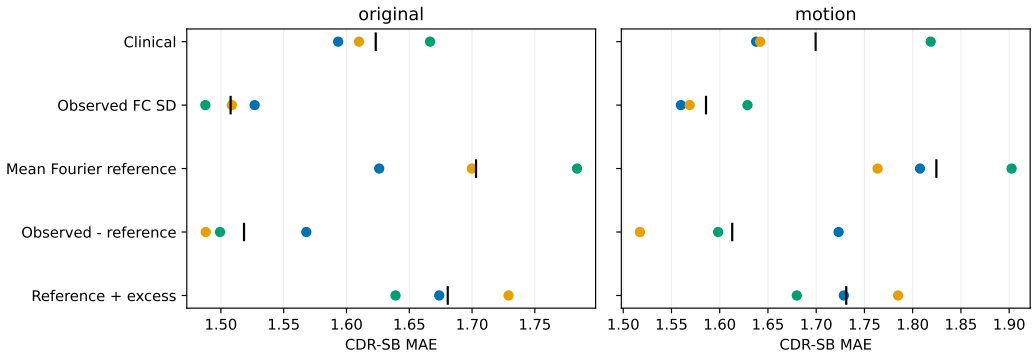

Figure 12: **Absolute future-CDR errors for the dispersion spectral-reference comparison.** Colored points are the three dependent repeated nested-CV splits of the same 48 participants; black marks are their means, not uncertainty intervals. Both covariate tiers and every declared spectral-model row appear here. Figure 4 presents selected matched contrasts as paired gains over the clinical baseline.

L COMPLETION CONTROLS AND CLINICAL HETEROGENEITY

All controls use the existing 48-person panel and frozen folds. Shared-frequency phases preserve the cross-periodogram; all 48 prepared records reproduce original features with maximum error 2.18×10^{-13} and preserve static FC to 4.00×10^{-15}. The 360 outer predictions were reconstructed from saved projections and independently solved ridge systems (maximum difference 3.26×10^{-13}). The 60 tighter-date outer models were checked by refitting PLS and independently solving ridge (difference 2.94×10^{-14}). These checks do not independently reselect all inner hyperparameters.

For $K > 1$ finite-second-moment window estimates, Eq. 23 follows directly from second moments. Supplied `Coverage.lean` verifies the algebra given explicit moment premises; `Coverage.wl` checks the corresponding symbolic identities. Six Gaussian known-covariance worlds check expected variance at two spacings and three correlation levels. These are measurement fixtures, not models of neural disease. Lean version 4.26.0 and pinned Mathlib commit `2df2f0150c275ad53cb3c90f7c98ec15a56a1a67` were used, with no `sorry` or extra axioms.

The signed spatial maps are shown in Figure 9; all coordinate values are also supplied in `supplement/completion/SPATIAL_MAPS.json`.

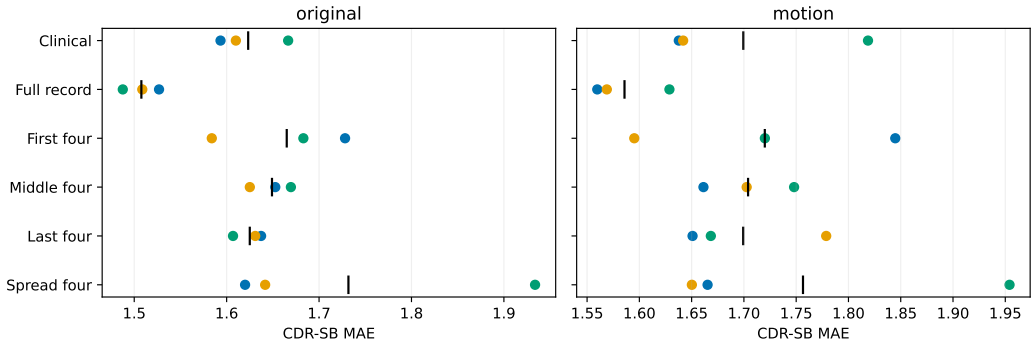

Figure 13: Four nonoverlapping windows, identical unique-frame budget within each person. Distributed sampling does not improve the mean clinical result in this test.

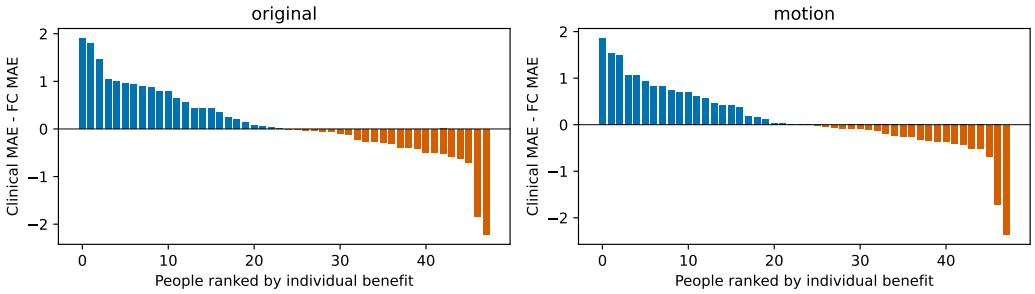

Figure 14: Paired individual benefit in fixed OOF losses, ranked for description. Ranking is not used to select participants or retrain models.

Interpretive limits and full outcomes. Aggregate files under `supplement/completion` retain every whole-panel model branch, signed spatial coordinate, known world and audit. Clinical/acquisition-stratum tables and leave-one-out extrema are withheld from the public artifact to avoid disclosing individual derived results. Observed-minus-reference features need not isolate physiology, and performance variation between disjoint reference averages prevents a convergence claim. The equal-frame experiment does not make the residual preprocessing local to each sampled

Table 8: All new representation results: MAE per dependent repeated split and their mean.

Branch	Split 1	Split 2	Split 3	Mean
original__clinical	1.5934	1.6100	1.6666	1.6233
original__full	1.5269	1.5087	1.4876	1.5077
original__static	1.6085	1.6468	1.7178	1.6577
original__reference	1.6261	1.6998	1.7836	1.7032
original__ref_half1	1.7450	1.7105	1.9521	1.8025
original__ref_half2	1.6521	1.6935	1.7975	1.7144
original__first4	1.7281	1.5839	1.6830	1.6650
original__middle4	1.6526	1.6251	1.6695	1.6491
original__last4	1.6373	1.6309	1.6069	1.6250
original__spread4	1.6199	1.6417	1.9337	1.7318
original__excess	1.5680	1.4879	1.4993	1.5184
original__ref_excess	1.6739	1.7290	1.6390	1.6806
motion__clinical	1.6377	1.6419	1.8187	1.6994
motion__full	1.5597	1.5689	1.6288	1.5858
motion__static	1.7147	1.6830	1.7889	1.7289
motion__reference	1.8076	1.7636	1.9025	1.8246
motion__ref_half1	1.7804	1.7615	1.9988	1.8469
motion__ref_half2	1.7592	1.7547	1.8513	1.7884
motion__first4	1.8448	1.5951	1.7201	1.7200
motion__middle4	1.6613	1.7026	1.7480	1.7040
motion__last4	1.6509	1.7787	1.6684	1.6993
motion__spread4	1.6653	1.6502	1.9541	1.7565
motion__excess	1.7232	1.5173	1.5984	1.6129
motion__ref_excess	1.7287	1.7849	1.6799	1.7312

interval. No untouched people remain within the released cohort; no new external validation cohort was established.

M Expanded clinical evaluation and all declared controls

M.1 Panel selection and analysis history

The 732-scan source programme contains 270 participants with two to five visits each. The diagnostic index panel uses the earliest eligible MRI with compatible DXSUM diagnosis within 90 days; a later visit is not substituted after feature inspection. It has 90 CN, 90 MCI and 90 dementia participants. The repeated-field study includes 462 later visits (135 CN, 149 MCI, 178 dementia by current visit diagnosis). A participant can change group; repeat counts are visits, not disjoint people. Some fields have 267 index participants and 457 repeats. The exact-TR3 common panel has 187 index participants and 330 repeats. Those changes mix acquisition restriction and cohort composition.

The fixed 137-field inventory crosses two domains (native and exact TR3), two diagnosis definitions (current and index-fixed), two adjustment levels, and three contrasts, giving 3,288 rows. Every row is supplied in `supplement/clinical_inventory.json`, including negative and discordant results. Native coordinates are not automatically harmonized across TR. Coordinate-wise regression includes the contrast, an indicator for the third group, and the frozen nuisance design. AQ includes age at the first root MRI, sex, an age-proxy indicator, elapsed years from that MRI, TR, retained sample count, recorded field-map status, phase, site-prefix and scanner indicators, mean FD, fraction FD above 0.2 mm, and median standardized DVARS. The categorical level dictionary is fixed from the declared cohort; these are descriptive adjusted maps, not cross-validated predictions. Each person has total weight one within each analysis; clinical-map cosines are descriptive. Their directional stability cannot establish individual discrimination.

The prognostic panel is a different subset of index-MCI participants with eligible preceding baseline and future CDR-SB assessments. Full local clinical tables, including assessments between MRI visits, were used. The current two-year CDR analysis has 48 participants, not 732 observations. Baseline is past-only within 90 days; target selection minimizes distance to 730 days within [548,913]. Actual follow-up days enter the design, so this is conditional retrospective prediction rather than an

operational model with a guaranteed future assessment date. MMSE and FAQ analyses belong to prior exploration: they are not silently pooled into the present target. The two-year FAQ panel had only nine eligible people and was unsupported under its declared minimum-30 rule.

The candidate followed earlier PCA/PLS and horizon exploration. The fixed same-participant tests here constrain that candidate but do not reset the exposure history. None of the repetitions, later visits of the same people, operator sensitivities or new report versions is described as an independent untouched cohort. The supplementary analyses preserve the different 201-person spectral panel and its clinical-date caveats; the newer full-table endpoint selection does not retrospectively resolve the old panel's date provenance.

M.2 Exact learning and measurement definitions

All MRI blocks are standardized within training data, with a training SD at or below 10^{-12} replaced by one for model scaling, and divided by the square root of block dimension. PLSRegression uses `scale=False`, up to ten components, `max_iter=1000`, `tol=1e-8`. Its scores are concatenated with the clinical design, and the combined design is centered/scaled again within the training fold for ridge. The target is centered and an unpenalized intercept restored. Ridge solves $(X^\top X+\alpha I)b = X^\top(y-\bar{y})$. Inner mean absolute errors are pooled over validation participants; ties are resolved by the ordered tuple (loss, component count, alpha). Categories are enumerated from the training fold; unseen categories have all-zero indicators. The clinical-only model independently selects alpha in the same inner folds.

Outer five-fold seeds are 732211, 732212 and 732213. Inner three-fold seeds are the outer seed plus fold index plus 100. Participants are never split across train/test within an analysis. Repeated splits share people and are not independent experiments. The frozen row order is a restricted participant ordering; reproducing folds requires authorized access to its ledger, not only the published seeds. Our anonymous supplement supplies definitions and aggregate results, not restricted records.

For each stored FC matrix the upper triangle follows network order Control, Default, Dorsal Attention, Limbic, Salience/Ventral Attention, Somatomotor, Visual. Correlations are computed on windowed network-average residuals after joint nuisance projection. Full record lengths yield 8, 12 or 18 windows. Native rounding gives 20-sample windows/10-sample strides at TR3 and 99/49 at TR0.607. Eight consecutive windows span 270 or 268.294 seconds. Distributed eight windows retain the full temporal span and alter overlap; they do not simulate a shorter acquisition.

M.3 All temporal-coverage comparisons

Table 9: All 20 coverage configurations. Original/motion denote the two covariate sets. Fixedfull-to rows replace the held-out measurement in a full-record model without refitting. All other rows are nested refits. Three seed columns are repeated CV, not independent cohorts.

Configuration	Seed 1	Seed 2	Seed 3	Mean MAE
original__clinical	1.5934	1.6100	1.6666	1.6233
original__fixedfull_to_first8	1.6767	1.6230	1.6234	1.6410
original__fixedfull_to_middle8	1.7551	1.7500	1.6421	1.7157
original__fixedfull_to_last8	1.5121	1.5512	1.5421	1.5352
original__fixedfull_to_spread8	1.5628	1.4848	1.5295	1.5257
original__full	1.5269	1.5087	1.4876	1.5077
original__first8	1.5364	1.5605	1.8105	1.6358
original__middle8	1.6098	1.6558	1.6701	1.6452
original__last8	1.5986	1.6406	1.7168	1.6520
original__spread8	1.4786	1.5810	1.6291	1.5629
motion__clinical	1.6377	1.6419	1.8187	1.6994
motion__fixedfull_to_first8	1.7134	1.6702	1.7109	1.6981
motion__fixedfull_to_middle8	1.7966	1.7625	1.7186	1.7592
motion__fixedfull_to_last8	1.5378	1.5975	1.6533	1.5962
motion__fixedfull_to_spread8	1.5812	1.5301	1.6471	1.5861
motion__full	1.5597	1.5689	1.6288	1.5858
motion__first8	1.5744	1.6179	1.9415	1.7112
motion__middle8	1.6325	1.6847	1.8258	1.7143
motion__last8	1.6268	1.6563	1.6994	1.6608
motion__spread8	1.5013	1.6547	1.6475	1.6012

M.4 Matched candidate controls and preprocessing sensitivity

The 99 MRI row-shuffle seeds are 733001–733099; one permutation is shared across all blocks/families. All 24 combinations of six families, PCA/PLS, and single/multiple components are retrained with inner selection, preserving clinical outcomes and covariates. The unpermuted pass reproduces all 24 previous metric rows. The FC candidate improves MAE by 0.11557; the median matched-shuffle difference from baseline is +0.16328, and the median best-of-24 difference is -0.01854. Counts at least as good as the candidate are 1/99 and 5/99 respectively. These are not conditional p-values: the shuffle also breaks MRI–covariate dependence, and the search envelope excludes other historically considered horizons/outcomes. All 2,400 aggregate rows are supplied, rather than only the favorable comparison.

On the same 48 participants, A5 FC gives mean MAE 1.5540 and the paired A5–M24 dispersion response gives 1.5615; both improve against the matching original clinical baseline in all three splits. M24's 1.5077 remains the primary candidate. This operator comparison is exploratory and does not certify a neurally optimal preprocessing scheme.

M.5 Matched static FC and dispersion comparison

The same 48 participants, outcomes, M24 operator and participant splits are used for eight configurations: clinical, clinical plus static FC, clinical plus dispersion, and clinical plus their concatenation, each under original and extra-motion covariates. Static FC consists of all 21 full-record Pearson edges in the same seven-network order. The combined representation is one 42-coordinate MRI block, standardised within training data and divided by $\sqrt{42}$ before the existing PLS stage. All models retain the earlier component/regularisation grid; no edge selection is introduced.

The declared primary combined-minus-static MAE differences across the three original splits are $+0.0364, -0.0491, -0.1791$ (mean -0.0639); with extra motion they are $-0.0457, -0.0920, -0.2169$ (mean -0.1182). Dispersion alone beats static FC in all three splits under both covariate sets, whereas concatenation does not beat dispersion alone on average. All eight configurations and per-split metrics are supplied in `supplement/static_comparison_results.json`.

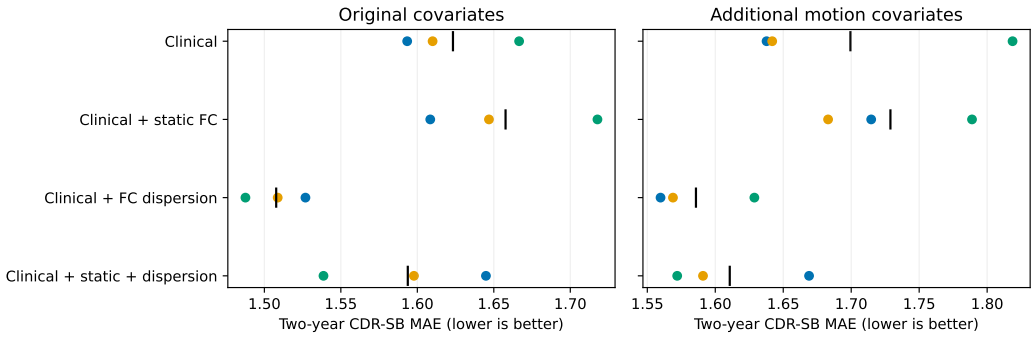

Figure 15: All matched static/dispersion configurations. Coloured points are the three repeated splits of the same participants; black marks are their means, not uncertainty intervals.

Before human fitting, a rank-one synthetic fixture exposed numerical failure of the inherited ten-component PLS extraction. An explicit applicability guard was added before any new outcome was inspected: the centred training MRI block must have at least ten singular values exceeding 10^{-10} times the largest. Otherwise the run stops as unsupported without participant deletion or component retuning. All real training folds passed. This qualification amendment leaves the four historical clinical/dispersion predictions unchanged exactly; no human rank failure was encountered.

Static FC and window dispersion were independently reconstructed from prepared parcel time series, the common M24 residual projection and fixed network membership. Maximum differences from the archived features were 7.14×10^{-14} and 2.18×10^{-13}. All 120 outer predictions were

reconstructed from saved projections and original features; each ridge stage was additionally solved through an augmented least-squares system. This is not an independent PLS implementation. The comparison remains retrospective and post-selection; improved performance in this fitted representation does not prove neural nonstationarity or superiority over every static-FC learner.

M.6 All edge ablations

Weights below are raw-coordinate regression effects multiplied by the training SD, averaged across 15 outer fits. Signs summarize fits, not confidence intervals. Each deletion repeats all inner choices. Correlated edges can substitute, and ablation changes the representation; neither coefficient nor deletion effect is causal. In particular, removing Default–Limbic or Dorsal Attention–Limbic improves MAE in all three splits, whereas some consistently signed large coefficients have little deletion effect.

M.7 Order features: definition, controls and negative result

For a W-window edge trajectory, let Z_{te} be centered and divided by its population SD. The order analysis rejects constant/nonfinite coordinates rather than interpreting them dynamically. At lag L, persistence and antisymmetry are

$$P_e = 1 - \frac{W-1}{2W(W-L)} \sum_{t=1}^{W-L} (Z_{t+L,e} - Z_{te})^2, \quad A_{ef} = \frac{1}{W-L} \sum_{t=1}^{W-L} (Z_{te} Z_{t+L,f} - Z_{tf} Z_{t+L,e}).$$

The primary $L = 3$ corresponds to 90 or 89.229 seconds. The original $L = 2$ non-overlap input gate failed for four TR0.607 records (one shared sample); this was corrected before human model fitting, retaining smaller lags as sensitivities. P is not the conventional ACF estimator; it is reversal-even and has zero expectation under a uniform permutation of fixed window indices. A is reversal-odd and concerns ordered pairs of FC edges, not anatomical neural flow.

All three additions (21 persistence coordinates, 210 antisymmetry coordinates, or both) are fitted jointly with the dispersion block. This changes the shared projection, so lack of gain need not mean absence of information in raw dynamics. Mean MAE is 1.7000, 1.5982 and 1.7043 respectively. The combined model at smaller lags gives 1.6156 and 1.8333, also worse than full dispersion. Nineteen full-window permutations (734001–734019) and 19 permutations of adjacent two-window blocks (735001–735019) preserve each person's joint-state multiset. Matched counts no worse than the observed P/A/combined results are 13/7/10 of 19 for full shuffling and 11/4/10 for blocks. Best-of-three counts are 18/10/18 and 19/9/19. These surrogate frequencies are diagnostic; they do not establish physiological exchangeability. All 131 aggregate configurations are supplied.

The identity $\sum_{i,j}(z_i - z_j)^2 = 2W \sum_i z_i^2 - 2(\sum_i z_i)^2$ underlies the finite-count normalization. The supplied additional Lean file verifies this finite-set identity, a scalar normalization and reversed-pair algebra. Wolfram verifies corresponding symbolic cases. The probability argument over permutations and clinical claims are not fully formalized.

M.8 Other negative clinical results and scope

The earlier three-class CN/MCI/dementia experiment uses 270 index participants and retains all nine distinct prediction outputs for joint-delay/DFA/technical combinations. Balanced accuracy spans approximately 0.28–0.36; no useful general diagnostic classifier is established by that experiment. Its aggregate confusion matrices and metrics are supplied. This does not negate group-map structure or independently test the later FC prognosis model.

An earlier phase-attribution change (last minus first four minutes, primary 32-second scale, M24, adjusted dementia–CN) had index coefficient -0.032612 (descriptive SE 0.011169). On 171 strictly later visits from 129 overlapping people the coefficient is $+0.000744$ (SE 0.012974). That particular directional contrast did not repeat. Its narrower selection and different observable are not treated as a failure of all phase topology or all joint dynamics. The inherited parcel-orientation appendix concerns a separate earlier descriptive extension and is not a confirmation of this later phase result.

Table 10: All 21 edge deletions. Positive final column means worse prediction after removal. Full per-seed values are supplied as aggregate JSON.

Edge	Mean SD effect	Signs $+/-$	MAE deletion $-$ full
Cont / Default	-0.041	6/9	-0.006
Cont / DorsAttn	-0.004	9/6	+0.030
Cont / Limbic	+0.164	14/1	-0.001
Cont / SalVentAttn	+0.017	7/8	+0.002
Cont / SomMot	+0.225	15/0	+0.073
Cont / Vis	-0.120	0/15	+0.027
Default / DorsAttn	-0.099	2/13	-0.005
Default / Limbic	-0.062	7/8	-0.016
Default / SalVentAttn	-0.236	0/15	+0.008
Default / SomMot	+0.063	11/4	+0.009
Default / Vis	+0.181	13/2	-0.010
DorsAttn / Limbic	+0.061	12/3	-0.024
DorsAttn / SalVentAttn	-0.135	3/12	+0.088
DorsAttn / SomMot	-0.074	5/10	+0.036
DorsAttn / Vis	-0.172	1/14	+0.031
Limbic / SalVentAttn	-0.347	0/15	+0.103
Limbic / SomMot	+0.035	11/4	-0.009
Limbic / Vis	+0.391	15/0	+0.143
SalVentAttn / SomMot	+0.212	15/0	+0.094
SalVentAttn / Vis	-0.066	3/12	+0.010
SomMot / Vis	+0.156	15/0	+0.031

M.9 POST-ASSEMBLY PAIRED-VISIT REPRESENTATION CHALLENGE

These experiments were completed after the original prognostic candidate and static-FC comparison. They reuse the ADNI732 programme and therefore do not constitute a held-out cohort. Their target is *retrospective change between two observed visits*, not prediction from an index MRI. An eligible pair requires two independently recorded CDR-SB values under the checked 90-day rule, positive clinical elapsed time and the same fMRI TR. There are 264 pairs from 190 people (79 CN, 87 MCI and 98 dementia pairs by initial stage), with 121 increases, 103 zero changes and 40 decreases in CDR-SB. All folds are person-disjoint but the three shuffled five-fold seeds reuse these people. Person-equal weighting, train-only scaling and fixed ridge $\alpha = 10$ are common to the paired tests; bootstrap intervals below are descriptive because the cohort and method families have already been examined.

The first comparison adds one of seven declared scalar distances to baseline clinical/time/technical covariates and three FC distances, using each method's own supported panel. The primary outcome is $|\Delta\text{CDR-SB}|$; negative ΔMAE is better. These are limited representatives of broader method families, not full phase topology or a universal fractal/entropy search:

Added scalar	Pairs / people	ΔMAE	Descriptive person-bootstrap95
C400 spatial-share TV	264 / 190	$+0.0093$	$[+0.0042, +0.0147]$
DFA log-fluctuation shape	223 / 161	$+0.0091$	$[+0.0031, +0.0155]$
Normalized delay energy	264 / 190	$+0.0097$	$[+0.0019, +0.0181]$
Phase H_1 scalar	256 / 186	$+0.0021$	$[-0.0056, +0.0111]$
C400 ordinal-entropy distance	264 / 190	-0.0106	$[-0.0325, +0.0118]$
Full DFA log-fluctuation level	223 / 161	$+0.0010$	$[-0.0048, +0.0075]$
Absolute phase holonomy	256 / 186	$+0.0075$	$[+0.0037, +0.0116]$

Ordinal entropy is the only favourable overall scalar increment and improves all three repeated splits, yet its interval includes zero. Its absolute-change MAE is 0.9320, near the clinical/technical-only 0.9340 and below the FC-control 0.9426. A stage breakdown performed *after* the main result suggests MCI/dementia rather than CN may drive that pattern; it is a hypothesis, not a confirmed interaction.

The signed spatial challenge holds a different, fixed comparator: baseline plus the same three FC distances *and full C400 spatial-share TV*. A 14-group hemisphere-by-network signed mass vector and the signed residual after subtracting its group means were defined without CDR-SB selection. The former discards a median 58.5% of full TV within groups; the latter retains a median 91.3% of squared spatial-change energy on 416 same-TR pairs. Among 409 matched true/other-person endpoint comparisons, the residual's L^1 norm places the true endpoint closer 96.1% of the time, versus 96.3% for full TV and 76.0% for the 14-group TV. The residual norm still tracks absolute motion change (Spearman $\rho = 0.287$), and its near-equality to full TV does not show independent scalar information. Exact finite rescaling examples keep full/window Pearson FC and full TV fixed while changing the signed field; these establish information distinction, not clinical relevance.

For signed ΔCDR-SB on the same 264 pairs/190 people, adding the 13 independent 14-group coordinates increases person-held-out MAE by $+0.0736$ points (descriptive interval $[+0.0333, +0.1204]$). Adding all 400 within-group signed coordinates, globally trace-scaled *within each training fold* to total weighted variance 13 rather than standardized parcel by parcel, increases it by $+0.0900$ $[+0.0300, +0.1432]$; combining group and residual directions yields $+0.1537$ $[+0.0839, +0.2228]$. All three seeds worsen under each tested construction. For the absolute-change sensitivity, the within-group addition is $+0.0973$ $[+0.0429, +0.1495]$. The changed-scanner subgroup of 23 people is descriptively favourable for the within-group model, whereas the 177 same-scanner people worsen; this is not an interaction claim or a rescue of the overall result. The source-bound reports and independent audits check 462 pair constructions, all 15 person-disjoint folds, 1,584 out-of-fold records and a separate weighted ridge solution. Their aggregate machine-readable results are supplied as `supplement/recent_representation_checks.json`; restricted individual records are not redistributed.

These paired tests sharpen the distinction between a mathematically richer observable, individual scan separability and clinical generalization. They do not revise the future-severity FC result on 48 MCI participants, prove neural origins of the spatial field, or reject other phase, Hurst/DFA, fractal,

entropy and dynamical-system constructions. In particular the ordinal stage signal still requires a separately defined, explicitly exploratory test rather than selection on these out-of-fold errors.

Other separately defined probes broaden the search without producing a common method ranking. A complete ordinal-pattern transition law changes under time reversal while full FC, 23 window-FC dispersion summaries and scalar ordinal entropy remain fixed (maximum L^1 change 0.2616); its 252-coordinate addition worsens held-person three-class diagnosis log-loss relative to scalar entropy by $+0.0450$ nats on 270 people (descriptive interval $[+0.0222, +0.0692]$). A 35-coordinate finite-scale Haar construction likewise distinguishes equal full and window FC, but its declared surrogate-residual addition worsens the separate diagnosis log-loss by $+0.0862$ nats on 228 supported people $[+0.0420, +0.1316]$; it is not a full wavelet-leader or multifractal test. Finally, a 12-coordinate M04 graph summary adds $+0.0727$ MAE $[-0.0165, +0.1585]$ to its own matched FC/window control on the previously studied 48-person future-CDR panel. M04 is a deterministic function of full C400 FC and cannot distinguish identical full FC. These targets, supports and baselines differ; each result limits its stated parameterization, not the broader method family. The source-locked aggregates are also in `supplement/recent_representation_checks.json`.

Another outcome-blind spatial decomposition uses seven network masses and a start-composition-weighted signed C400 residual, rather than the 14-group orthogonal residual fitted above. Its scalar norm places the true endpoint closer than a matched other-person endpoint in 97.1% of 409 pairs (full TV: 96.3%), but correlates with absolute motion change ($\rho = 0.301$). The signed field is mathematically distinct from full TV; the similar scalar-norm result alone does not establish an incremental scalar signal. This separate measurement has not had a clinical fit or signed-direction repeatability gate and reuses the same people and pseudo-pairs, so it is not independent replication of the 14-group result.

Other prospectively specified probes retain both favourable leads and unfavourable fits. Adding the A5/32-second phase-holonomy magnitude to FC/window/technical controls lowers three-class diagnosis log-loss by 0.00963 nats on 267 supported first-visit people, with a descriptive interval $[-0.01981, +0.00022]$ for the added-minus-control difference; diagnostic gain is unconfirmed and the holonomy itself lacks a matched common-spectrum control. On a separate 48-person future-CDR fit, Schaefer-block mean/SD features are mathematically distinguishable from equal C7 full/window FC but add $+0.0325$ MAE $[-0.0546, +0.1262]$ to their own M24 FC/window control. A W96 imaginary-coherency summary also distinguishes equal-FC constructions but adds $+0.1319$ MAE $[+0.0567, +0.2135]$ in its M24 future-CDR fit; its A5 counterpart is uncertain. These absolute MAEs must not be compared directly with Table 7, which has a different fitted baseline.

Time-reversal witnesses for seven-network switching and quadratic drift likewise preserve specified full/window FC while changing their respective dynamic operators. The tested 28-coordinate switching residual adds $+0.0237$ MAE $[+0.0022, +0.0442]$ to its own FC/window control on 65 year-one MCI people. The quadratic-drift residual gives $+0.0047$ $[-0.0261, +0.0339]$ on that panel. In a different within-scan task, own-network history improves held-block prediction over a constant mean in 720/732 visits, while joint history beats own history in 434/732; its predefined diagnosis contrasts are unresolved. None of these results rejects continuous dynamics or memory as mathematical families. The aggregate supplement includes source-file and manifest hashes for these releases without pooling their non-equivalent clinical outcomes.