Supplementary Information

Soluble amyloid composition links proximal APP processing, fibrillar burden and the p-tau181 response

Analytical scope and reporting hierarchy

This supplement documents the estimation targets, uncertainty procedures, supporting analyses and inconclusive comparisons within the same study. The main text prioritizes peptide-only coordinate construction, biological orientation by independently measured readouts and formal information-retention analyses; model-compatibility, conditional, selection and additional sensitivity analyses are reported here in full. The analyses were developed iteratively and were not preregistered. A positive result in one model and a non-significant result in another are not by themselves evidence of a difference. Direct contrasts, confidence intervals and held-out performance are distinguished throughout.

Supplementary Methods

Peptide-only coordinates, scaling and dependence

The primary peptide decomposition used the 380-person parent cohort and was constructed exclusively from same-assay Aβ38, Aβ40 and Aβ42 measurements. Equal-SD scaling within the 375-person p-tau181 sample was used for surface coefficients and directions. Scaling changes coefficient units, not the fitted subspace or evidence for conditional association. Expanded p-tau models retain parent-coordinate scaling; their coefficients should not be compared numerically with the rescaled gradient without conversion. Because D42 is cross-fitted rather than an in-sample ordinary least-squares residual, exact algebraic orthogonality to T is not guaranteed. Empirical T-D42 dependence was therefore summarized with Pearson and Spearman correlations in the parent, p-tau181 and PET samples. Stability was evaluated by repeating the same repeated five-fold cross-fitting design across 101 alternative split sequences. These correlations are descriptive checks of coordinate geometry rather than tests that the population correlation equals zero. The complement of a learned score is geometrically perpendicular in the specified scaled metric, not statistically independent.

Specified-model compatibility checks

The five standardized log markers were residualized on age, sex, education, APOE dosage and diagnosis and re-standardized before maximum-likelihood one-factor fitting. Loadings were 0.644 (BACE), 0.514 (sAPPβ), 0.964 (Aβ38), 0.987 (Aβ40) and 0.559 (Aβ42). Off-diagonal residual-correlation RMSE was 0.025. Simulations reintroduced fitted covariate means, retained the observed PET availability pattern and rebuilt the complete coordinate pipeline. PET generation from the factor was calibrated to the observed biomarker-block incremental variance, with base R^2=0.2296 and full R^2=0.6194. Parameters were plug-in estimates. A separate conditional check held the measured design fixed and wild-bootstrapped outcomes from a reduced mean excluding D42. These tests do not address every nonlinear latent or time-dependent one-dimensional model.

p-tau181 response-surface estimand

The descriptive colour field in main Fig. 3 uses a two-dimensional Gaussian kernel smoother of covariate-standardized p-tau181, with leave-one-out-selected bandwidth 0.38 SD. Its dashed 27 pg/mL contour is distinct from the solid confidence boundary, which comes from the adjusted linear log-p-tau181 model. The orange marker is a descriptive location within the region satisfying the simultaneous-bound criterion, not an inferentially estimated onset location. Exponentiating a fitted mean on the log-p-tau scale yields a geometric mean at fixed mean covariates, not an arithmetic mean or an individual probability of positivity. The 150-by-150 grid spans coordinate 1st-99th percentiles and restricts evaluation to the convex hull with local Kish

effective N at least 20 (Gaussian bandwidth 0.38 SD). A maximum-standardized-excursion bootstrap supplies an approximate simultaneous band conditional on the coordinate design and support. The 27 pg/mL reference is a published Elecsys value, not an onset threshold discovered here. Area fractions describe grid area, not patient prevalence. Failure of the lower bound to exceed the reference does not establish a below-reference mean.

Information-retention, resampling and prediction

Rank-one versus rank-two least squares compared the 2-by-3 coordinate-related coefficient matrix after residualizing covariates, with equal weight for standardized log-BACE activity, log-sAPPβ and log-p-tau181. Covariate coefficients remained unrestricted. The restricted-null wild bootstrap used one sign per participant for the entire residual vector and a leverage correction; it did not assume diagonal outcome covariance. The corrected probability is not an exact finite-sample permutation probability. Primary rank inference conditions on the supplied coordinates; the separate 2,000-replicate reconstruction bootstrap accounts for coordinate estimation in the 380-person parent cohort. Main predictive transformations are training-only. Repeated-CV ranges or win counts are not independent-replicate inference.

Nonlinear single-index and held-out PET comparisons

The nonlinear sensitivity optimized one common linear score while allowing separate polynomial responses, with shared maximum degree from one to three and regularization chosen by inner cross-validation. It is not an exhaustive comparison with nonlinear latent curves. For the PET complement analysis, the 284-person non-PET training set and 91-person PET set were disjoint for score learning. Parent-derived coordinate preprocessing was inherited; this is not an external-cohort validation or a fully held-out preprocessing pipeline. Two-stratum bootstrap repeats both score learning and PET regression. PET selection weighting was not rerun for this extension.

Conditional analyses and biological interpretation

Conditioning on a response positively associated with both coordinates can generate an inverse association between them. Residualization against T or T+D42 imposes orthogonality by construction and therefore helps explain model behaviour, but cannot independently prove a molecular mechanism. No model in this study measures peptide flux, a fibrillar conversion rate, oligomer concentration or a causal inhibitory effect of BACE.

PET selection and influence audits

Inverse-probability weighting was restricted to 288 CN/MCI participants, with 91 subsequently observed by PET. Truncated stabilized weights ranged from 0.492 to 2.309; effective sample size was 77.4 and the maximum absolute standardized mean difference fell from 0.308 to 0.055. These diagnostics concern measured selection factors and do not eliminate unmeasured attrition or restore baseline-dementia support. Participant-level influence on the primary D42-minus-T PET contrast was assessed by omitting each of the 91 PET participants in turn, refitting the full HC3 model and recalculating the contrast. The influence audit was then repeated across 101 alternative D42 cross-fitting sequences. Sign stability and the range of leave-one-out contrasts were treated as descriptive robustness summaries; Cook's distance and leverage were inspected in the reference reconstruction but were not used for participant exclusion.

Supplementary Table S1. Complete proximal association estimates

Outcome	Proximal marker	Coefficient	95% CI	P	Block R^2
T	BACE activity	0.516	0.433 to 0.598	1.67×10^{-34}	0.468
T	sAPPβ	0.313	0.220 to 0.406	3.76×10^{-11}	0.468
D42	BACE activity	0.049	-0.040 to 0.137	0.281	0.006
D42	sAPPβ	0.043	-0.038 to 0.124	0.301	0.006

Incremental R^2 is a block contribution, not partial R^2 and not the variance of an identified biological process.

Supplementary Table S2. Specified-model compatibility checks

Statistic	Data-generating condition	Estimate / median	2.5th-97.5th percentiles	Directional empirical tail P
Upstream separation	Observed	0.463	-	-
Upstream separation	Fitted one-factor model	0.470	0.395 to 0.542	0.583
PET D42 coefficient	Observed	36.920	-	-
PET D42 coefficient	Fitted one-factor model	-0.394	-8.775 to 7.911	0.00020
PET T coefficient	Observed	5.712	-	-
PET T coefficient	Fitted one-factor model	29.945	20.075 to 39.749	-
PET D42-minus-T contrast	Observed	31.208	-	-
PET D42-minus-T contrast	Fitted one-factor model	-30.434	-43.215 to -17.138	0.00020
PET D42 coefficient	Conditional PET model check	0.020	-15.821 to 15.624	0.00005
PET D42-minus-T contrast	Conditional PET model check	-3.817	-24.712 to 17.216	0.00035

The proximal separation is reproduced by the fitted benchmark; this is a guardrail on interpretation, not evidence against the measured PET contrast. Directional probabilities use the +1 correction.

Supplementary Table S3. Relative Aβ38:Aβ40 composition

Analysis	Estimate	95% CI	P value	N
OOF R^2, T only	0.254	-	-	380
OOF R^2, T + G38:40	0.268	-	-	380
Δ OOF R^2 from G38:40	0.013	-	-	380
G38:40 → z(log Aβ42)	-0.130	-0.215 to -0.044	0.003	380
Aβ42 → Aβ38, conditional	-0.032	-0.076 to 0.012	0.154	380
Expanded PET: D42	37.53	25.51 to 49.55	9.49×10^{-10}	91
Expanded PET: T	4.58	-6.26 to 15.43	0.407	91
Expanded PET: G38:40	-4.65	-14.75 to 5.45	0.367	91
Expanded PET: D42-minus-T contrast	32.94	14.04 to 51.85	0.00064	91

G38:40 summarizes relative shorter-peptide composition. Failure to account for D42 does not localize D42 to a particular molecular step.

Supplementary Table S4. p-tau181 association models

Model / term	Coefficient	95% CI / added R^2	P
T + D42, equal scale: T	0.518	0.454 to 0.581	3.69×10^{-57}
T + D42, equal scale: D42	0.472	0.394 to 0.550	1.35×10^{-32}
Expanded proximal model: T	0.383	Parent-coordinate scale	1.1×10^{-14}
Expanded proximal model: D42	0.456	Parent-coordinate scale	2.8×10^{-34}
Expanded proximal model: BACE	0.197	Standardized log marker	4.6×10^{-6}
Expanded proximal model: sAPPβ	0.094	Standardized log marker	0.017
T + D42 beyond covariates	-	Added R^2 = 0.425	-
BACE + sAPPβ beyond T + D42	-	Added R^2 = 0.030	-

All rows refer to N=375 with clinical/genetic adjustment. Expanded-coordinate coefficients retain their parent-cohort scaling; only the first two rows define the equal-SD response gradient. Expanded-model coefficient intervals are not imputed from P values.

Supplementary Table S5. Regression directions and paired angular contrasts

Quantity	Point angle	95% bootstrap interval
BACE	4.6	-4.1 to 13.7
sAPPβ	6.9	-5.4 to 20.7
p-tau181	42.4	36.5 to 48.2
Later PET	88.7	75.4 to 98.1
p-tau181 minus BACE	37.8	28.4 to 47.1
p-tau181 minus sAPPβ	35.5	22.0 to 48.1
PET minus p-tau181, same 91	45.5	30.9 to 58.2

Degrees measured from +T toward +D42. An angle near zero indicates alignment with T, not lack of a biomarker association. Paired contrasts use matched resampling within the corresponding sample.

Supplementary Table S6. Surface model and reference-level sensitivity

Quantity	Result
Linear model R^2 / AIC / BIC	0.6145 / 723.77 / 759.11
Quadratic model R^2 / AIC / BIC	0.6186 / 725.70 / 772.83
Joint quadratic HC3 test	P=0.355
Quadratic-minus-linear CV RMSE	+0.00276
27 pg/mL: supported area above lower bound	44.5%
28.69 pg/mL: supported area above lower bound	38.3%
34.61 pg/mL: supported area above lower bound	20.2%
Relative balance B at fixed M	Coefficient -0.032; P=0.384

M=(T+D42)/sqrt(2), B=(D42-T)/sqrt(2) are signed linear reparameterizations, not a norm or kinetic quantities. Stricter thresholds necessarily produce nested regions; nesting is not independent validation. No nonlinear-switch claim follows from these analyses.

Supplementary Table S7. Ratio benchmark and linear-rank comparison

Quantity	Result
Ratio-depletion p-tau model R^2	0.414
Ratio depletion + T model R^2	0.613
T coefficient given ratio depletion	0.467 (95% CI 0.402-0.532)
Rank-two mean in-sample R^2 gain	0.0236
Parent-reconstruction bootstrap 95% interval	0.0140 to 0.0355
Restricted-null bootstrap / exceedances	20,000 / 0
+1-corrected tail probability	0.000050
Joint OOF R^2, rank 1 / rank 2	0.372 / 0.394
p-tau181 OOF R^2, rank 1 / rank 2	0.556 / 0.597

The modest mean prediction gain should not be interpreted as a large improvement in global disease prediction. Rank-two evidence concerns a common linear projection of measured coordinates.

Supplementary Table S8. Rank and nonlinear sensitivities

Analysis	Result
BACE + sAPPβ only	Wild-bootstrap P=0.766
BACE + p-tau181	Wild-bootstrap P=0.00020
sAPPβ + p-tau181	Wild-bootstrap P=0.00020
Diagnosis omitted	Wild-bootstrap P=0.00020
CN stratum	Wild-bootstrap P=0.006
MCI stratum	Wild-bootstrap P=0.00020
Dementia stratum	Wild-bootstrap P=0.312
Tuned nonlinear index: joint OOF R^2	0.376
Two-coordinate linear model: joint OOF R^2	0.394

Stratum-specific significance differences do not test effect heterogeneity. The nonlinear comparison uses the same 30 outer repetitions, not the 100 repetitions used for the main rank-prediction summary.

Supplementary Table S9. PET information in the complementary direction

Quantity	Result
Score-training N / PET-evaluation N	284 / 91 (disjoint for score learning)
Complement coefficient, Centiloids/SD	30.8
HC3 95% interval / P	17.5-44.1 / 5.67×10^{-6}
Two-stratum bootstrap interval	20.3-40.7
R^2 without / with complement	0.351 / 0.620
Mean OOF RMSE without / with complement	45.1 / 35.2 Centiloids

The complementary coordinate is perpendicular in the training-standardized metric. Parent-derived preprocessing and single-cohort selection limit claims of external validation.

Supplementary Table S10. Conditional p-tau181 inversion

Conditional model / term	Estimate	95% CI or P
BACE-D42 before p-tau181 adjustment	+0.049	-0.040 to 0.137; P=0.281
BACE-D42 after p-tau181 adjustment	-0.167	-0.258 to -0.077
sAPPβ-D42 after p-tau181 adjustment	-0.070	-0.151 to 0.012
T-residualized BACE, p-tau-adjusted	-0.032	P=0.381
T-residualized sAPPβ, p-tau-adjusted	-0.003	P=0.937
Residual p-tau181 after T/D42 removal	-	P=0.702

These conditional checks explain coefficient behaviour. Orthogonalization is not independent biological evidence, and no protective or inhibitory BACE effect is inferred.

Supplementary Table S11. Full PET benchmark and sensitivity estimates

Analysis	Estimate	95% CI	P value	N
Primary D42	36.92	24.25 to 49.59	1.12×10^{-8}	91
Primary T	5.71	-5.27 to 16.69	0.308	91
D42-minus-T contrast	31.21	11.12 to 51.30	0.0023	91
Aβ42/Aβ40 ratio-depletion benchmark	36.76	23.76 to 49.76	2.97×10^{-8}	91
D38 negative control	0.26	-11.35 to 11.87	0.965	91
D42, joint p-tau181 model	25.90	11.71 to 40.09	3.47×10^{-4}	91
D42, IPW	36.00	21.14 to 50.87	2.07×10^{-6}	91
D42 × time, repeated PET	1.08	-1.78 to 3.95	0.458	59

The repeated-PET analysis involved 59 participants and 123 post-landmark observations. Its uncertainty precludes a supported accumulation-rate claim. Technical BACE/sAPPβ replicate-choice sensitivity preserved upstream separation across 3,000 selections: median 0.455, empirical interval 0.443-0.466; that interval concerns replicate choice, not participant sampling.

Supplementary Table S12. Coordinate-dependence and PET-influence audits

A. Empirical T-D42 dependence

Sample	N	Reference split-sequence correlations	Range across 101 alternative split sequences
Parent cohort	380	Pearson r=-0.0002; Spearman ρ=0.0142	r: -0.00117 to 0.00118; ρ: 0.01311 to 0.01531
p-tau181 sample	375	Pearson r=-0.0146; Spearman ρ=0.0044	r: -0.01547 to -0.01351; ρ: 0.00363 to 0.00550
Later-PET subset	91	Pearson r=-0.0162; Spearman ρ=-0.0104	r: -0.01758 to -0.01393; ρ: -0.01336 to -0.00519

Cross-fitted D42 is not constrained to be exactly orthogonal to T. The audit therefore reports empirical dependence directly. Across the 101 alternative split sequences, the out-of-fold Aβ42 prediction R^2 ranged from 0.2660 to 0.2687. Correlation P values are omitted because this is a descriptive geometry audit rather than a null-hypothesis test of zero dependence.

B. Participant-level influence on the later-PET D42-minus-T contrast

Quantity	Reference reconstruction	Across 101 alternative D42 split sequences
Positive leave-one-out contrasts	91/91	91/91 in every sequence
Leave-one-out contrast range, Centiloids/SD	29.68 to 39.45	29.55 to 39.63 overall
Lowest HC3 95% CI lower bound	8.78 Centiloids/SD	Not recalculated for each alternative sequence
Most influential omission (Cook's D)	Contrast 39.45; Cook's D=0.258	Influential-point ranking not used for the split-sequence audit

No participant exclusion was used in the primary PET analysis. In the reference reconstruction, all 91 leave-one-out HC3 confidence intervals for the contrast remained above zero; the omission with the largest Cook's distance strengthened rather than attenuated the contrast. The 101-sequence audit repeated coordinate construction and leave-one-out refitting to assess whether the influence conclusion depended on a particular cross-fitting partition.

Interpretive limits

A two-coordinate linear representation can retain more information even if the underlying observations arise from a nonlinear scalar trajectory. For example, one scalar s can generate (s, s-squared), whose linear map has rank two. The present study characterizes measured conditional information, not the intrinsic dimensionality of a molecular system. Similarly, the supported p-tau response region does not prove that people outside it are p-tau-negative or that an individual trajectory avoids tau pathology. Timing and mechanism require direct longitudinal and molecular measurements.