

Cognitive-Function Architecture Under Continuous Personality Measurement

Integrated evidence dossier: bell-curve retrofit, prevalence controls, exact stack-order tests, Big Five/HEXACO geometry, and BTIT/BTIT² qualification

Status: working manuscript / replication package

Prepared from the 2026 computational test series and cited external psychometric literature.

Central result: Continuous, center-weighted population personality scores do not by themselves rule out discrete/configurational cognitive-function architectures. Under a Big Five + HEXACO trait-geometry bridge, the canonical dominant → auxiliary → tertiary → inferior ordering ranked first among all 24 same-function permutations.

Executive Summary

This dossier consolidates a sequence of computational and literature-grounded tests aimed at a specific hypothesis: human personality may contain a relatively discrete/configurational cognitive-function architecture whose observable expression is continuous because of adaptation, environment, context, measurement noise, and other person-level modifiers. Four-letter MBTI codes are used only as shorthand or as bridges to published datasets; the computational primitives are the ordered Jungian function stacks themselves.

The first result is a compatibility result. A population composed of 16 ordered cognitive-function architectures crossed with 162 ordered tritype configurations can aggregate into smooth, center-weighted, non-bimodal personality-score distributions of the type reported by Bess and Harvey (2002). The result survives when the 16 architecture frequencies are frozen to published prevalence priors. This demonstrates that a smooth population curve is not logically sufficient to exclude discrete latent configurations.

The second result concerns specificity. Bell-curve shape alone does not recover the canonical function order: many alternative permutations fit the same marginal distributions as well or better. This negative result is important because it identifies the information missing from one-dimensional score curves.

The third result introduces multivariate external trait geometry. Using published Big Five and HEXACO correlations as empirical targets, exact function orders were compared against the other 23 permutations of the same four functions. Under both a fuller function-to-trait bridge and a more conservative bridge that removes the J/P-derived component, the canonical dominant → auxiliary → tertiary → inferior order ranked 1st of 24 on Big Five, 1st of 24 on HEXACO, and 1st jointly. The result remains directionally robust under bounded synthetic nurture/adaptation nuisance, although robustness is weaker in the conservative bridge.

The fourth result is methodological. BTIT/BTIT² repeated-measurement experiments demonstrate that an observer architecture can recover separate architecture-like and nurture/adaptation-like transition fields when those fields are experimentally present, and it follows the perturbed field when the control is swapped. This validates the testing machinery, not the claim that tritype literally equals nurture in real humans.

The combined evidence supports continued testing and eventual publication, but the strongest empirical claim still awaits a matched human dataset containing direct cognitive-function measures (or repeated function-level behavior), Big Five and/or HEXACO, tritype, and longitudinal/context measurements. The present dossier therefore distinguishes model demonstrations, bridge-based empirical support, and unresolved claims rather than treating them as equivalent.

Evidence Status at a Glance

Claim	Current status	Interpretation
Discrete/configurational architectures can yield continuous population curves	Demonstrated within constrained model	Strong compatibility result
Result survives frozen population prevalence	Demonstrated within constrained model	Strong compatibility result
Bell curve identifies canonical stack order	No	Marginal curves are non-identifying
Canonical exact stack order best fits Big Five + HEXACO bridge	Yes, in current bridge-based test	Positive evidence; not direct function measurement
Tritype literally equals nurture	No	Used only as adaptation/nurture proxy
BTIT ² can separate architecture-like vs nurture-like fields when controls exist	Yes, synthetically	Method validation, not human empirical proof
Canonical function order proven in humans	No	Requires direct matched longitudinal data

1. Research Question and Hypothesis

The program tests whether the apparent conflict between typological cognitive architecture and continuous psychometric traits is a false dichotomy. The working model is hierarchical rather than purely categorical or purely continuous.

Operational hypothesis:

- A person has a relatively stable cognitive-function architecture represented as an ordered stack (for example, Ne → Ti → Fe → Si).
- Tritype is treated as a bounded adaptation/nurture proxy for experimental purposes, not as an established biological/environmental decomposition.
- Context, life history, state, measurement noise, and architecture × adaptation interaction alter expressed behavior around that architecture.
- Population aggregation over many architectures and modifiers can therefore produce smooth continuous distributions even if the latent architecture is configurational.

Model shorthand: observed trait profile = cognitive-function architecture + adaptation/nurture proxy + interaction + context/residual.

1.1 Nomenclature and the Unit Being Tested

The four-letter MBTI type is not the primitive variable in the computational tests. It is used only as a shorthand identifier and, where necessary, as a bridge to published prevalence or correlation data. The tested architecture is the ordered function stack itself. The canonical order used in the experiments follows the alternating-attitude convention shown in Appendix A. The Myers & Briggs Foundation describes type dynamics as an ordered sequence of dominant, auxiliary, tertiary, and inferior processes, while also noting debate about tertiary attitude/orientation [2].

2. The Continuous-Curve Challenge

Bess and Harvey (2002) examined approximately 12,000 respondents scored with item response theory. Earlier reports of bimodality were shown to be artifacts of using too few quadrature points; with 50 quadrature points, the four MBTI score distributions became strongly center-weighted rather than bimodal [1]. This is an important challenge to any naive claim that psychological types must appear as separated peaks in observed score distributions.

The research question here is narrower: does a center-weighted distribution logically or statistically exclude a population generated from discrete/configurational cognitive architectures? The answer from the retrofit experiments

is no. However, compatibility is not the same as identification; a model that can reproduce a curve is not thereby proven to be the true generating mechanism.

3. Study Series

3.1 Bell-Curve Retrofit: Cognitive Stacks × Ordered Tritype

The first retrofit crossed the 16 ordered cognitive-function stacks with 162 ordered tritypes, yielding 2,592 latent configurations. Because the accessible Bess-Harvey paper reports the qualitative center-weighted result but not the original respondent-level scores or numeric histogram bins, the target was a standardized normal-shaped surrogate. This makes the experiment a compatibility test, not a literal re-fit of the original ~12,000 participants.

Regime	Within-state SD / architecture	Tritype / architecture	Max JS	All 4 unimodal?
Strict identifiable	0.50	0.50	0.01297	No
Moderate overlap	1.00	0.68	0.00109	Yes
Loose overlap	1.50	0.30	0.00022	Yes




Figure 1. Curve mismatch falls sharply once within-configuration expression is allowed to overlap substantially. The vertical scale is logarithmic.

Interpretation: discrete/configurational latent states can generate a smooth aggregate distribution when individual expression overlaps across configurations. The strongest form of the result is therefore not “people form isolated type islands,” but “latent architecture can be configurational while phenotype is continuous.”

3.2 Frozen Published Prevalence

The second retrofit removed a major degree of freedom: the 16 architecture frequencies were no longer estimated by the optimizer. Instead, they were frozen to a national normative prevalence prior ($n=3,009$) and, separately, to a larger 2025 aggregate sensitivity prior (~44,278 classifications). Under the moderate-overlap ceiling, both priors still yielded four unimodal center-weighted marginals.

Prevalence prior	Tritype / architecture	Within-state SD / architecture	Stack+tritype max JS	Stack-only max JS	Mismatch reduction
National normative	0.868	1.000	0.001232	0.002538	51.5%
2025 aggregate sensitivity	1.000	1.000	0.000842	0.002743	69.3%




Figure 2. The continuous-curve compatibility result survives when architecture population weights are frozen rather than fitted.

This strengthens the compatibility result: the optimizer was not merely inventing convenient type frequencies. At the same time, the successful fits sit at or near the maximum moderate-overlap bound, which reinforces the interpretation of substantial continuous expression around any latent configuration.

3.3 What the Bell Curve Cannot Tell Us: Exact Function Order

To test order rather than simple function membership, the exact four-function set for each architecture was held fixed while all 24 literal permutations competed. For the ENTP shorthand, for example, Ne → Ti → Fe → Si competed against Ti → Ne → Fe → Si, Ne → Fe → Ti → Si, and the other 21 permutations of the same four functions.

Curve-fit prior	Canonical order rank
National normative prevalence	20 / 24
2025 aggregate sensitivity	18 / 24

The canonical order did not win the marginal-curve contest. This is not evidence that the order is false; it shows that one-dimensional population marginals are non-identifying with respect to internal stack order. This negative result motivated the move to multivariate Big Five and HEXACO geometry.

3.4 Direct Rank-1 / Rank-2 Checks and Longitudinal Evidence

A direct but imperfect historical check comes from three studies summarized in the 1985 MBTI manual: the dominant preference score exceeded the auxiliary score in 55% of $n=101$, 53% of $n=62$, and 44% of $n=160$. Using nearest-integer counts from the rounded percentages gives approximately $159/323 = 49.2\%$, with a two-sided binomial $p \approx .824$ against 50%. Preference-score magnitude is not a direct measure of function use or development, so this does not cleanly falsify the hierarchy; it simply fails to supply rank-1 > rank-2 evidence.

Salter, Evans, and Forney (1997) followed 99 master’s students over 20 months and reported that dominant Thinking and dominant Sensing did not retest as well as dominant Intuition and dominant Feeling; they suggested environmental pressures as a major influence [5]. This finding is consistent with the need to control adaptation/context rather than assuming function-role expression is invariant.

3.5 Big Five + HEXACO Exact-Order Test

The strongest order-specific result in the present sequence uses two independent continuous trait systems as multivariate targets. A Big Five meta-analysis summarized seven samples totaling 2,243 participants [3]. A HEXACO-60 concurrent-validity study included 1,016 participants and published MBTI–HEXACO correlations [4]. Neither source measured the eight Jungian functions directly, so the experiment required an explicit function-to-trait bridge.

Two bridges were used. The full bridge combined function attitude (extraverted/introverted), process family (N/S/T/F), and the empirical external judging/perceiving phenotype for extraverted functions. The minimal bridge removed the J/P-derived contribution to reduce circularity. In both bridges, nominal role strengths were fixed at dominant=1.00, auxiliary=0.75, tertiary=0.45, inferior=0.25, and all 24 exact same-function orders competed.

Bridge	Trait target	Canonical exact-order rank
Full	HEXACO	1 / 24
Full	Big Five	1 / 24
Full	Joint	1 / 24
Minimal	HEXACO	1 / 24
Minimal	Big Five	1 / 24
Minimal	Joint	1 / 24

This is the first test in the sequence to produce positive order-specific evidence: under the stated bridge, the canonical dominant → auxiliary → tertiary → inferior ordering was the best-fitting of all 24 literal permutations in both Big Five and HEXACO geometry.




Figure 3. Canonical exact-order win rate under structured synthetic adaptation/nurture nuisance. The conservative minimal bridge remains less decisive than the full bridge.

3.6 Nurture / Adaptation Sensitivity

Because neither Big Five nor HEXACO source measured tritype, the study cannot perform true participant-level nurture residualization. Instead, v5 injects structured latent adaptation/tritype-like offsets and perturbs role strengths, rerunning all 24 exact orders 1,000 times per nuisance level. This is a sensitivity analysis, not empirical proof that tritype is nurture.

Nuisance / architecture	Full: exact win	Minimal: exact win	Full: top-3	Minimal: top-3
0.00	96.4%	61.4%	100.0%	100.0%
0.50	91.1%	59.3%	100.0%	100.0%
1.00	71.7%	46.8%	98.3%	91.6%

The full bridge is highly robust through moderate nuisance. The minimal bridge still places the canonical order among the strongest candidates but exact-win robustness degrades as unmeasured adaptation becomes as large as the architecture signal. This is precisely why direct participant-level control is needed before publication as a strong empirical claim.

3.7 BTIT / BTIT² Repeated-Measurement Qualification

A separate synthetic program tested whether an upper observer could distinguish a persistent architecture-like field from an adaptive/nurture-like field while watching lower-BTIT state transitions. The full controlled topology contained 16 architectures × 162 ordered tritypes = 2,592 states. Candidate neighborhoods were cardinality-balanced so tritype did not receive more opportunities simply because it has more configurations.

Qualification metric	Result
Architecture amplitude recovery, no interaction	$r = .833$
Nurture amplitude recovery, no interaction	$r = .820$
Dominant-lane recovery, no interaction	78.6%
Architecture amplitude recovery, strong interaction	$r = .743$
Nurture amplitude recovery, strong interaction	$r = .718$
Dominant-lane recovery, strong interaction	72.9%
Nurture-sensitive control: tritype-lane share	1.000 (SD .001)
Architecture-sensitive swapped control: architecture-lane share	.988 (SD .016)

The strongest control is the swap: when context sensitivity is deliberately moved from the tritype lane to the architecture lane, the inferred context-reactive direction follows the injected perturbation. Therefore the observer is not simply hard-coded to call “tritype = nurture.” However, blind clustering without control runs was weak (56.8% accuracy; adjusted Rand index .019). Real identification therefore requires repeated measurements and meaningful control conditions.

4. Integrated Interpretation




Figure 4. Qualitative evidence ladder. This figure summarizes the relative role of each test; values are interpretive and are not inferential statistics.

Taken together, the study sequence supports a hierarchical model more strongly than a simple “types versus traits” dichotomy. The mathematical retrofit establishes compatibility: discrete/configurational latent architectures can

aggregate into continuous trait distributions. Prevalence controls show that this result does not depend on freely choosing architecture frequencies. The exact-order bell-curve controls then show that marginal distributions are insufficient to recover internal stack order. Finally, multivariate Big Five/HEXACO geometry supplies information that the marginal curves lack and favors the canonical ordering under two bridge specifications.

The strongest defensible current statement is:

A discrete/configurational cognitive-function architecture is compatible with continuous personality measurement, and the canonical dominant → auxiliary → tertiary → inferior ordering is the best-fitting same-function permutation in the present Big Five + HEXACO bridge analysis. The result is promising but not yet a direct empirical confirmation of Jungian function order in humans.

5. What the Program Has Established — and What It Has Not

5.1 Established within the stated computational models

- Smooth, center-weighted personality marginals can be produced by populations built from discrete/configurational cognitive-function architectures.
- The compatibility survives frozen published architecture prevalence rather than requiring freely fitted type frequencies.
- A tritype/adaptation layer improves curve fit under the specified geometry and overlap bounds.
- Marginal curve shape alone cannot identify canonical function order.
- An exact-order test can recover a true role order in synthetic data when role-specific temporal signatures are actually present.
- BTIT²-style observation can recover controlled architecture-like and adaptation-like transition fields and passes swapped negative controls.

5.2 Empirical/bridge-based support

- Published Big Five geometry and published HEXACO geometry independently place the canonical exact order first among the 24 same-function permutations under the stated function-to-trait bridge.
- The order result survives removal of the strongest J/P-derived bridge component at baseline, but robustness to large unmeasured adaptation is materially weaker in that conservative specification.
- Historical dominant-function longitudinal evidence is mixed and suggests environmental/contextual influence is relevant rather than negligible.

5.3 Not yet established

- That tritype is literally the human nurture variable.
- That the eight cognitive functions have been directly measured as latent causal mechanisms.
- That canonical stack order would remain 1/24 in a participant-level dataset measuring functions directly, rather than reconstructing them through MBTI preference correlations.
- That the present results uniquely outperform all continuous latent-trait or alternative configurational models under out-of-sample likelihood criteria.
- That any of these results constitute biological proof of cognitive functions.

6. Recommended Decisive Replication

The next experiment should be designed to make the cognitive-function hypothesis fail if it is wrong. The dataset should be participant-level and preregistered, with the canonical order and all 23 literal same-function alternatives specified before analysis.

1. Recruit a sufficiently large sample with broad demographic coverage and repeat a meaningful subset longitudinally.
2. Collect direct cognitive-function measures or repeated behavior tasks designed to distinguish Ne/Ni, Se/Si, Te/Ti, and Fe/Fi. Do not infer functions solely from the four-letter code.
3. Collect Big Five and HEXACO in the same participants.
4. Collect ordered tritype plus external nurture/context covariates (age, education, culture, socioeconomic context, major life events, current stress, etc.). Tritype should be tested as a proxy, not assumed to exhaust nurture.
5. Use baseline, ordinary task, deliberate support, moderate stress, acute stress, and recovery conditions where feasible. Stack order should predict asymmetric temporal signatures if dominant/auxiliary/tertiary/inferior roles are real.
6. For each four-function architecture, compete all 24 exact orders under identical parameter budgets. No candidate receives additional flexibility.
7. Use participant-level cross-validation: fit on some people/context and predict held-out people/context. A canonical result must generalize, not merely fit in-sample.
8. Compare the canonical stack model not only against permutations, but also against function-set-without-order, generic continuous latent-factor models, and flexible mixture controls.
9. Predefine success criteria: canonical order should beat all 23 exact-order controls out of sample in both Big Five and HEXACO geometry, survive nurture/context residualization, and replicate across independent samples.

7. Publication Framing

A publishable manuscript should lead with the compatibility problem rather than with a claim that cognitive functions have been proven. The strongest contribution is methodological: continuous observed traits do not logically exclude configurational latent architecture, and a sequence of increasingly restrictive tests can separate mere curve compatibility from order-specific evidence.

7.1 Suggested paper structure

Manuscript	Primary question	Natural framing
Paper 1: Compatibility / methods	Can discrete cognitive-function configurations generate continuous observed personality distributions under fixed prevalence and bounded adaptation?	Computational methods / psychometrics
Paper 2: Order-specific trait geometry	Does canonical exact function order outperform the other 23 same-function permutations against Big Five and HEXACO geometry?	Personality / measurement
Paper 3: Direct preregistered replication	Do direct function-level longitudinal measures recover the same order after nurture/context control?	Confirmatory empirical study

7.2 Claim language recommended for external review

- Use “compatible with,” “supports,” “best-fitting under the stated bridge,” and “requires direct replication.”
- Avoid “proved cognitive functions are real” until there is direct participant-level function measurement and independent replication.
- Report negative results prominently: bell curves do not identify order; direct dominant-vs-auxiliary score evidence is approximately chance-level; uncontrolled clustering does not cleanly recover nature/nurture lanes.
- Release the computational specification and exact alternative orders so third parties can rerun or attack the model.

8. Limitations

- The Bess-Harvey raw respondent data and numerical 50-point histograms were not available in the accessible source used here; the retrofit therefore targets a standardized center-weighted surrogate rather than the original microdata.
- The Big Five and HEXACO sources report correlations with MBTI preference dimensions, not direct measures of Ne/Ni/Se/Si/Te/Ti/Fe/Fi. The function-level order result depends on an explicit bridge from published trait directions to function vectors.
- Tritype is used as a structured adaptation/nurture proxy because it is central to the working personality model; it is not established as a literal decomposition of environmental causation.
- The canonical stack implementation fixes a specific function-attitude ordering. The Myers & Briggs Foundation notes continuing debate over tertiary attitude, so results pertain to the operational stack convention tested here [2].
- The 2025 aggregate prevalence prior is used only as a sensitivity prior and should not be presented as a representative population estimate without its full sampling source and methodology.
- BTIT/BTIT² components in these psychometric experiments are bounded adapters inspired by the route-selection/observer architecture. Their use does not by itself create independent empirical evidence; the conventional statistical lane remains necessary.
- Current order evidence is bridge-based and mostly in-sample to published correlation geometry. Stronger claims require participant-level out-of-sample prediction and independent replication.

9. External References

[1] Bess, T. L., & Harvey, R. J. (2002). Bimodal Score Distributions and the Myers-Briggs Type Indicator: Fact or Artifact? *Journal of Personality Assessment*, 78(1), 176–186. DOI: 10.1207/S15327752JPA7801_11.
<https://pubmed.ncbi.nlm.nih.gov/11936208/>

[2] Myers & Briggs Foundation. The Processes of Type Dynamics: Dominant, Auxiliary, Tertiary, and Inferior Processes. <https://www.myersbriggs.org/unique-features-of-myers-briggs/type-dynamics-processes/>

- [3] The Myers-Briggs Company. MBTI Facts: Big Five relationships; summarizes Arneson (2016) meta-analysis of 7 samples, combined N=2,243. <https://www.themyersbriggs.com/en-US/Support/MBTI-Facts>
- [4] Tarigan, M., & Fadillah, F. (2021). Uji Validitas Eksternal Tes Kepribadian HEXACO-60. *Intuisi: Jurnal Psikologi Ilmiah*, 13(1). N=1,016. DOI: 10.15294/intuisi.v13i1.28408. <https://journal.unnes.ac.id/nju/INTUISI/article/view/28408/0>
- [5] Salter, D. W., Evans, N. J., & Forney, D. S. (1997). Test-Retest of the Myers-Briggs Type Indicator: An Examination of Dominant Functioning. *Educational and Psychological Measurement*, 57(4), 590–597. DOI: 10.1177/0013164497057004005. <https://journals.sagepub.com/doi/10.1177/0013164497057004005>

10. Internal Computational Study Series

Run	Study	Purpose / result
v1	Cognitive-Stack × Ordered-Tritype Curve Retrofit	Bell-curve compatibility; free architecture mixture weights; BTIT lane
v2	Fixed Published Prevalence Retrofit	Frozen prevalence priors; tritype sensitivity; order non-identification
v2-RM	BTIT ² Repeated-Measurement Personality Topology	Synthetic architecture/adaptation lane recovery and swapped controls
v4	Exact Cognitive-Function Order Test	All 24 literal orders; bell-curve non-identification; rank-1/rank-2 checks
v5	Big Five + HEXACO Exact Function Order	Canonical order 1/24 under full and minimal bridges; nurture sensitivity

Working conclusion: the hypothesis has moved from “compatible with the old curve” to “positively supported by multivariate trait geometry under explicit assumptions.” The remaining task is direct, participant-level, independently replicated measurement of the functions themselves.

Appendix A. Canonical Exact Cognitive-Function Stacks Used

Shorthand only	Operational exact stack
ENTP	Ne → Ti → Fe → Si
INTP	Ti → Ne → Si → Fe
ENTJ	Te → Ni → Se → Fi
INTJ	Ni → Te → Fi → Se
ENFP	Ne → Fi → Te → Si
INFP	Fi → Ne → Si → Te
ENFJ	Fe → Ni → Se → Ti
INFJ	Ni → Fe → Ti → Se
ESTP	Se → Ti → Fe → Ni
ISTP	Ti → Se → Ni → Fe
ESTJ	Te → Si → Ne → Fi
ISTJ	Si → Te → Fi → Ne
ESFP	Se → Fi → Te → Ni
ISFP	Fi → Se → Ni → Te
ESFJ	Fe → Si → Ne → Ti
ISFJ	Si → Fe → Ti → Ne

Note: the shorthand labels are not explanatory variables in the stack-order tests. They identify which published prevalence/correlation information is associated with a given exact function set.

Appendix B. Key Numeric Results

Test	Result
v1 moderate stack+tritype	Max JS 0.00109; all four unimodal
v1 moderate stack-only	Max JS 0.00333; all four unimodal
v2 national normative, stack+tritype	Max JS 0.001232; all four unimodal
v2 national normative, stack-only	Max JS 0.002538
v2 canonical order rank on marginals	20/24 normative; 18/24 aggregate
v4 dominant>auxiliary historical check	159/323 ≈ 49.2%; $p \approx .824$
BTIT ² mixed recovery, gamma=0	$r = .833$ architecture; $r = .820$ nurture; 78.6% dominant lane
BTIT ² swapped architecture-sensitive control	98.8% architecture-lane share
v5 full bridge exact order	1/24 Big Five; 1/24 HEXACO; 1/24 joint
v5 minimal bridge exact order	1/24 Big Five; 1/24 HEXACO; 1/24 joint
v5 full bridge, nuisance=0.50	91.1% exact win; 100% top-3
v5 minimal bridge, nuisance=0.50	59.3% exact win; 100% top-3
v5 full bridge, nuisance=1.00	71.7% exact win; 98.3% top-3
v5 minimal bridge, nuisance=1.00	46.8% exact win; 91.6% top-3

Appendix C. Replication Checklist

- Publish exact data transformations and all 24 same-function order candidates.
- Provide code and fixed random seeds for sensitivity analyses.
- Distinguish empirical data from synthetic nuisance/control data in every table.
- Pre-register direct function-level measures and function-order predictions.
- Collect participant-level tritype and external context variables rather than treating nurture as latent only.
- Use held-out participant/context prediction as the primary confirmation criterion.
- Replicate in a second independent sample before making a strong ontological claim about cognitive functions.