Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Precision-weighted updating shows cross-task consistency as a bounded computational marker when precision gain is sufficiently large

[redacted] [redacted]

hzauluofunan@163.com

Huazhong Agriculture University https://orcid.org/0009-0008-6314-928X

Research Article

Keywords: computational psychiatry, precision weighting, hierarchical Gaussian filter, Balloon Analogue Risk Task, Columbia Card Task, hierarchical transfer cross-validation, Ornstein-Uhlenbeck process, identifiability

Posted Date: September 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10857398/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Precision-weighted updating shows cross-task consistency as a bounded computational marker when precision gain is sufficiently large

Running title: Precision-weighted updating when gain is large

Authors and affiliations

Ke Wan^{1,} (sole and corresponding author)

^{1} Huazhong Agricultural University, Wuhan, Hubei, China

Correspondence: Ke Wan, Huazhong Agricultural University, Wuhan, Hubei, China. E-mail: hzaoluofunan@163.com

Abstract

Background. Single-task fitted parameters are often treated as portable traits. We asked whether a precision-weighting coefficient γ in CPSP-8 generalizes across tasks under preregistered hierarchical transfer cross-validation (HTCV).

Methods. Archived NUTS posteriors came from OpenNeuro BART (ds000030; N=173) and CCTHot (ds004636; N=104); HierPrior (ds008083; N=43) supplied a behavioural γ probe. Ornstein–Uhlenbeck parameters were frozen. Consistency required a 95% HDI containing 0. Post-review checks (not used to change H1–H4) were ROPE, conjugate MAP+SE sensitivity, WAIC/LOO, age/sex IPW, and a non-null family-separability grid. Choice and RT RMSE were scored against dedicated models.

Results. Only γ transferred under preregistered NUTS (BART–CCT 0.3835, 95% HDI [−0.0084, 0.7826]; all three HDIs contain 0). This is not ROPE equivalence and fails on a conjugate grid (0/36 cells contain 0). WAIC favoured task-specific γ. Other decision parameters failed H1. γ recovered at 960 trials (median r≈0.66) but was family-inseparable from RW at γ=0. Non-null probes separated at γ≥1.0 with ≥960 trials (or γ≥1.5 at 480), not at γ≤0.5. θ stayed below r=0.3. CPSP-8 lost layer-wise RMSE to dedicated models.

Conclusions. γ is a bounded, pipeline-sensitive cross-task scalar, not a biomarker. Family-level precision-weighting identity holds in silico only when gain is large. Other decision parameters are mode switches; the OU layer remains frozen.

Keywords. computational psychiatry; precision weighting; hierarchical Gaussian filter; Balloon Analogue Risk Task; Columbia Card Task; hierarchical transfer cross-validation; Ornstein-Uhlenbeck process; identifiability

Introduction

Computational psychiatry aims to replace syndromal labels with generative parameters that explain how beliefs, values, and actions go awry [1,2,39]. The promise is mechanistic: if a learning rate, an inverse temperature, or a volatility coupling can be estimated from behaviour, then that quantity might serve as a transdiagnostic marker, a treatment target, or-more ambitiously-an endophenotype lying between genes and symptoms [3]. The risk is circular: a parameter that is only identifiable inside one experimental graph can be narrated as a person-level trait simply because the software returned a number. Distinguishing portable computational phenotypes from task-bound regression weights is therefore not a secondary psychometric chore. It is the condition under which computational psychiatry can claim to be measuring anything that outlives a particular session [4,5,40].

Wilson and Collins' s "Ten simple rules for the computational modeling of behavioral data" placed parameter recovery at the centre of that condition [4]. Their recipe is deliberately unromantic. Simulate data from known generating values; refit the same model; demand a tight, unbiased correspondence between true and recovered parameters before interpreting individual differences. Recovery failure is diagnostic of

underpowered designs, flat likelihoods, or code error; it is not a licence to proceed with a more colourful story. Complementary work has insisted that predictive fit is not falsification: a model can win a likelihood contest and still be unable to generate the behavioural effect it claims to explain [6]. Together these rules define a minimum standard: identifiability first, generative adequacy second, and only then claims about people rather than tasks.

That standard is routinely failed in practice. Pike and colleagues showed, in a large multi-task sample, that winning-model parameters can recover well and even show acceptable test-retest reliability while remaining essentially unrelated across tasks that nominally tap the same construct-except within a narrow class of inverse-temperature-like quantities-and only weakly related to symptoms [5,40]. An earlier meta-analysis of reinforcement-learning studies in mood and anxiety disorders likewise found that apparent patient-control differences are sensitive to model class and fitting pipeline [7]. The resulting tension is structural rather than anecdotal. Tasks are designed to identify a particular computation (sequential pumping under unknown explosion hazard; turning cards under known gain/loss probabilities; choosing colours and stakes). Parameters inherit that design. Task specificity versus cross-task generalization is therefore not a software bug to be patched by a larger optimizer. It is the empirical question of whether a computational quantity is an endophenotype or a context effect [4,5,8].

The classical psychiatric endophenotype was defined to solve a related problem one level up. Gottesman and Gould required a marker that is associated with illness, heritable, state-independent in the relevant sense, co-segregating in families, and present at higher rates in unaffected relatives than in the general population [3]. The construct was never a synonym for "biomarker that differs between cases and controls." It was a bridge between genotype and distal phenotype, intended to be closer to biology than a DSM checklist and closer to behaviour than a single molecular assay. Cognitive endophenotypes (working memory, P300, smooth-pursuit gain) occupy one layer of that bridge; neural endophenotypes (resting connectivity, receptor maps) occupy another [3,9]. Neither layer automatically licenses a fitted softmax temperature as an endophenotype. A computational endophenotype must earn the title by showing, at minimum, that the same generative quantity is recoverable, related to illness-relevant variation, and not wholly rewritten by the next experimental graph [1,2,5,40].

Huys, Maia and Frank framed computational psychiatry as the attempt to identify such quantities-prediction-error weights, precision terms, Pavlovian biases-inside explicit generative models rather than inside black-box classifiers [1,39]. Stephan and Mathys likewise argued that hierarchical Bayesian models of perception and learning provide a common currency for connecting neuronal message passing to clinical phenomena [2]. The emerging notion of a computational endophenotype therefore sits between Gottesman's familial marker and a diagnostic cut-off: it is a parameter of a process model that (i) can be recovered in individuals, (ii) is not reducible to a single task's regressors, and (iii) is a candidate for, rather than a demonstration of, heritability and co-segregation [1-3]. The present study addresses (i) and (ii) for one precision-weighting parameter and explicitly does not claim (iii).

The Hierarchical Gaussian Filter (HGF) supplies the most influential formalization of precision-weighted updating in this literature [10,11]. In the HGF, beliefs at each level are Gaussian; higher levels encode log-volatility; and the learning rate at a given level is the ratio of precisions, so that prediction errors are weighted by how confidently they should be believed. Mathys and colleagues derived one-step variational updates that make this hierarchy estimable from trial sequences [10,11]. The same arithmetic is the workhorse of predictive coding and active inference: ascending prediction errors are gain-modulated by expected precision, and descending predictions are contextualized by that gain [12-15,41]. Kanai and colleagues located this precision engineering in a distinction between driving (first-order, content) and modulatory (second-order, precision) connections, with the pulvinar as a candidate subcortical coordinator [15]. Clark's review of predictive processing made the same claim in cognitive terms: the brain is not a passive feature detector but a hierarchy that constantly reweights sensory evidence against priors [14].

Clinical applications followed immediately and exposed the identification problem. Powers, Mathys and Corlett showed that overweighting of perceptual priors under Pavlovian conditioning can generate conditioned hallucinations, tying a precision imbalance to a psychotic-like percept in a causal experimental design [16]. Iglesias and colleagues linked midbrain responses to precision-weighted prediction errors

during volatility learning [17]. Lawson, Mathys and Rees reported that autistic adults show reduced precision on volatility estimates [18]. Mikus, Lamm and Mathys later used a generalized HGF to phenotype schizotypal traits and schizophrenia, associating uncertainty about volatility with positive-symptom variation and rigidity about observational noise with negative symptoms [19]. Weber and colleagues generalized the HGF further (gHGF), allowing nonlinear value coupling alongside volatility coupling and thereby bringing the filter closer to generic predictive-coding networks [20]. These successes do not cancel a recurring practical difficulty: volatility and precision parameters are often weakly identified from short behavioural series, TAPAS-style toolboxes must be applied with care in clinical samples, and a two-level precision circuit fitted to binary choices is not automatically the Mathys hierarchy [11,16,21,22]. The challenge is not that precision weighting is theoretically empty. It is that the mapping from a clinical interview to an identifiable ω or γ is a design problem, not a naming problem [4,21].

CPSP (a three-layer threat-dissociation account used throughout this programme) takes that design problem as its starting point. Layer 1 is a continuous-time Ornstein-Uhlenbeck (OU) threat state: mean-reversion θ pulls a latent threat z_t toward a baseline z_0; diffusion σ_ou injects process noise; an input gain κ scales external threat drive. Layer 2 is a precision-gated decision map whose learning rate is modulated by γ. Layer 3 is a dissociation drift with gain κ_dissociation, decay λ, and onset. Layers 1 and 3 are theoretically primary in a threat-dissociation framework. Threat is treated as a mean-reverting latent, not a trial-wise dummy: after an insult, the state should relax at rate θ rather than jump and stick. Dissociation is treated as a delayed drift in precision, not as a personality score. That pairing is why OU parameters can be theoretically core while remaining unidentifiable in tasks that never show a threat-offset epoch or a slow precision collapse. In the Balloon Analogue Risk Task (BART) and the hot Columbia Card Task (CCT), those layers are currently behaviourally blind. Identification maps from stages v1-v5 and a 1920-trial virtual-subject experiment never recovered θ above Pearson r=0.3, so the M3 reduction froze the OU five-tuple at group means (θ=0.285, σ_ou=0.1425, κ=0.30, λ=0.35, onset=1.5). Freezing is an honesty constraint: it prevents narrating unidentified dynamics as patient-specific threat recovery. It is not a theoretical retraction of the OU layer.

The remaining empirical question is therefore confined to the decision layer. Which of {threat, baseline, α_0, γ, β, v_gain} generalizes across BART, CCT, and a behavioural probe collected in an EEG protocol, and which quantities are rewritten by the task graph? We answered that question with hierarchical transfer cross-validation (HTCV). Unlike leave-one-subject-out cross-validation, which asks whether a model predicts held-out people inside one likelihood, HTCV asks whether group-level posterior means are exchangeable across tasks. LOOCV can reward a richly parameterized model that overfits a single graph; HTCV fails that model if its hierarchical mean shifts when the graph changes. A 95% HDI of μ_BART-μ_CCT that contains zero is evidence that the parameter is not detectably task-rewritten at the hierarchical mean. The HDI is the shortest interval containing 95% of the posterior mass, not an equal-tailed interval, and is the natural Bayesian analogue of a consistency test when the scientific claim is "no systematic shift" rather than "a point null of exactly zero" [23,24,42]. Four preregistered hypotheses structured the analysis. H1: all of {threat, baseline, α_0, β, v_gain} are consistent across BART and CCT. H2: γ is consistent across BART, CCT, and the HierPrior behavioural probe. H3: the OU layer cannot be tested under the freeze. H4: CPSP-8 is not worse than dedicated models on both choice and RT under a fair RMSE rule. Criteria were locked before the HTCV run and were not altered after seeing results.

Methods

Participants

All analyses are secondary uses of de-identified OpenNeuro releases. No new human participants were recruited for this manuscript.

D1 used BART events from the UCLA Consortium for Neuropsychiatric Phenomics LA5c study (OpenNeuro ds000030) [25,26]. After preregistered quality control (completion ≥ 70%, balloons ≥ 8, RT <200 ms proportion ≤ 30%, exclusion of stereotyped single-action profiles) and restriction of the group contrast to diagnosis in {CONTROL, SCHZ}, the hierarchical sample comprised N=173 (CONTROL=124, SCHZ=49). Attention-deficit/hyperactivity and bipolar diagnoses were not entered into the schizophrenia-control

contrast. BART is a sequential risk task in which each pump increases potential reward and explosion hazard [27]. Only balloons with trial_type other than CONTROL (non-gain balloons) entered the likelihood.

D5 used the hot Columbia Card Task (CCTHot) from OpenNeuro ds004636 (N=104 after quality control; choices ≥ 20) [28,29]. CCTHot presents a deck with known gain amount, loss amount, and loss probability; the participant repeatedly chooses draw_card versus end_round [28]. Two sessions were concatenated within participant. The dataset has no diagnosis column; no patient-control contrast was tested on D5. We cite Rogers and colleagues' Cambridge Gamble Task as a related known-probability gambling procedure with colour choice and stake proportion [30], while emphasizing that D5 is CCTHot, not the Cambridge task: the archived event codes are draw_card / end_round, not colour-then-bet.

D2 used HierPrior (OpenNeuro ds008083; N=43). Continuous EEG was not on disk for this analysis. The γ series is a two-parameter behavioural probe extracted from choice events, not an event-related potential (ERP) component. Absolute γ on D2 is therefore a behavioural, not neural, quantity.

Model architecture

OU state layer (theoretical). Let z_t be a latent threat state. In continuous time,

$$dz_t = -\theta(z_t - z_0)\,dt + \kappa\, u_t\,dt + \sigma_{ou}\,dW_t,$$

where $\theta>0$ is the mean-reversion rate, z_0 is the steady-state mean in the absence of input, κ is threat-input gain, u_t is a task-derived drive, σ_{ou} is the diffusion coefficient, and W_t is a Wiener process. Euler-Maruyama discretization with step Δt yields

$$z_{t+1} = z_t - \theta(z_t - z_0)\,\Delta t + \kappa\, u_t\,\Delta t + \sigma_{ou}\sqrt{\Delta t}\,\varepsilon_t,\ \varepsilon_t \sim N(0,1).$$

Dissociation drift (theoretical). A precision-drift process with gain $\kappa_{dissociation}$, decay λ, and onset time (or threshold) onset describes a delayed, decaying shift in effective precision after a control-loss event. In the present task family these five quantities $\{\theta,\sigma_{ou},\kappa,\lambda,onset\}$ were frozen at M3 group means (see H3) and were not free parameters of the observation model.

Decision layer (archived Stage 4 observation model). CPSP-8 is a precision-gated delta-rule with a logistic choice probability and a Wald-inspired RT mean, not a three-alternative softmax. Let q_t in (0.02, 0.98) be the learned explosion/loss probability, initialized at $q_0=0.15$. On BART, the decision variable and pump probability are

$$dv_t=(1-q_t)-q_t\cdot threat-0.10\cdot baseline,\ P(pump_t)=\sigma(\beta\cdot dv_t),$$

where $\sigma(x) = 1 / (1 + e^{-x})$, threat scales the cost of the current q, baseline is a sensory/baseline preference, and $\beta>0$ is inverse temperature (choice determinism). Balloon-level likelihood uses n_ACCEPT log P(pump) on explosion balloons and n_ACCEPT log P(pump)+log(1-P(pump)) on cashed-out balloons [27]. Precision-gated updating is

$$\alpha_t = clip(\sigma(logit(\alpha_0) + \gamma\,(PE_t^2 - q_t(1-q_t))),\ 0.01,\ 0.90),$$

$$q_{t+1} = clip(q_t + \alpha_t \cdot PE_t,\ 0.02,\ 0.98),$$

with $PE_t=1_{exploded}-q_t$ and $logit(\alpha_0)=\log[\alpha_0/(1-\alpha_0)]$. Thus α_0 is the initial (tonic) learning rate and γ is the precision-weighting coefficient: it raises α when squared prediction errors exceed Bernoulli uncertainty $q(1-q)$ and lowers α otherwise. On CCTHot the same γ update is applied after each round, with

$$EV_t=v_{gain}\cdot(gain/20)\cdot p_{gain}\text{-}threat\cdot(|loss|/70)\cdot p_{loss}\text{-}baseline,\ P(draw_t)=\sigma(\beta\cdot EV_t).$$

The eight free decision parameters are therefore threat, baseline, α_0, γ, β, v_{gain}, plus RT scales a and t_0.

RT layer. The theoretical Wald (inverse-Gaussian) density for a boundary $a>0$, drift v, and non-decision time t_0 is

$$f(RT=t|\ a,v,t_0)=(a)/(\sqrt{2\pi(t-t_0)^3})\exp(-((a-v(t-t_0))^2)/(2(t-t_0))),\ t>t_0.$$

The archived likelihood does not evaluate this density. It uses a log-normal approximation around a Wald-like mean,

log(median RT) ~ N(log(t_0+a/(v_gain|dv|+0.05)), 0.35^2),

with explosion-row RTs of zero excluded. We report this approximation because it is the model that generated the numbers; we do not retrofit a full Wald PDF after seeing results.

Datasets

OpenNeuro accession numbers and sample sizes are as above. D1 BART maps ACCEPT → risky, CASHOUT → safe, EXPLODE → a forced outcome (not a chosen exploratory act) [25-27]. D5 CCTHot maps draw_card → risky and end_round → safe; clicked_on_loss_card==1 is a forced loss outcome still counted as a draw. There is no independent pass event code; Stage 4 mapped non-draw actions to safe (high-risk mapping; Supplementary Materials). D2 γ is behavioural only.

HTCV procedure

HTCV was preregistered at 2026-08-24T08:24:30+08:00 (HTCV_prereg_v2.md) and run 08:44:01-08:45:07 the same day, after Stage 4 scope (07:31:00) and before this manuscript' s figures. D1 group means used sample-size pooling of CONTROL and SCHZ hierarchical means. For each decision parameter, 8,000 independent posterior draws formed Δ=μ_BART,pool-μ_CCT (and the two γ contrasts with D2). Consistency required that the shortest 95% interval of Δ contain 0. H1 required this for all of {threat, baseline, α_0, β, v_gain}. H2 required it for three γ contrasts (BART-CCT, BART-EEG, CCT-EEG). H3 was declared untestable under the OU freeze. H4 required PPC coverage and a fair RMSE rule on both tasks. Criteria were not altered after seeing results. HTCV is not leave-one-out predictive cross-validation: it tests exchangeability of hierarchical means across likelihoods, which is the relevant claim for a portable computational marker [4,5,23].

Post-review diagnostics (locked after H1–H4)

ROPE used ε ∈ {0.10, 0.15, 0.20} on the archived NUTS Δγ samples; practical equivalence required the 95% HDI to lie entirely inside ±ε. The sensitivity grid recomputed conjugate hierarchical means from archived subject MAP+SE for all eight CPSP-8 parameters, with prior SD {archived, ×0.5, ×2}, SE floors {0.02, 0.05, 0.10}, and pooling {sample-size weighted, equal, CONTROL-only, SCHZ-only} (3 × 3 × 4 = 36 cells per parameter). WAIC and a raw importance LOO-IC compared a shared-μ_γ model with a task-specific-μ_γ model, treating MAP+SE as observations (not trial-level likelihoods). Age and sex were read from ds000030, ds004636, and ds008083 participants tables, matched to fitted IDs. Propensity of BART versus CCT membership was a logistic model of age and sex; CCT subjects were IPW-weighted toward the BART covariate distribution; a linear meta-regression predicted MAP γ from task + age + sex. The non-null family-separability grid used the D1c precision-gated simulator: true γ ∈ {0.25, 0.5, 1.0, 1.5, 2.0}, trial counts {480, 960, 1920}, seeds {42, 7, 123}, 10 cohorts × 100 virtual subjects (450 cells, 45,000 simulated subjects). Nested comparison was CPSP-8 with γ free versus γ pinned at 0; BF = exp[0.5(BIC_nested − BIC_free)].

Virtual-subject recovery

Read-only archives were: D1d (480 vs 960 trials, three seeds) for γ recovery; D1e nested point-null γ=0 versus RW for family separability; M3 nine-parameter recovery with OU frozen; and θ recovery from v1-v5, M1-960, and M1-1920. Subject-level true-versus-posterior scatter arrays were not archived; Figure 4 therefore plots Pearson r, not fabricated clouds.

Model comparison

Competitors shared the same inclusions, the same MAP-then-summary pipeline, and independent test summaries [6,8]. (i) RW: two parameters (α,β), q← q+α(outcome-q), logistic choice, no RT likelihood [31]. (ii) Logistic regression (D5 only): P(draw)=σ(b_0+b_1gain+b_2|loss|+b_3 p_loss). (iii) DDM: (v,a,t_0) with P(pump)=σ(v) and the same RT mean form without learning [32,33]. (iv) HGF-2: a simplified two-level precision circuit with (ω,ϑ,κ) and choice temperature frozen at 4.0. Required wording: this competitor is an IsolatedHGFChoice-style binary precision update written for these tasks; it is not the Mathys hierarchical Gaussian filter and is not a TAPAS inversion [10,11,21]. RMSE was computed separately on the choice summary (adjusted pumps on BART; cards drawn on CCT) and on median RT; layers were never blended into

a single "better model" score for the H4 claim.

Statistical inference

Participant-level maps used L-BFGS-B on the balloon/round sequential negative log-likelihood. Hierarchical inference used NUTS [34] in a measurement-error model y_i ~ N(μ_g(i), sqrt(σ^2+se_i^2)) with four chains × 2000 draws, warmup/tune=1000, target_accept=0.90, cores=1, progressbar=False, seed 42. Preregistered diagnostics were R-hat<1.01 and effective sample size >400; if a chain failed those thresholds the posterior was still reported without altering priors, mappings, or decision rules. Group-mean priors (natural scale) were μ_threat ~ N(0,1); μ_baseline ~ N(1.0,0.5) truncated >0; μ_α_0 ~ N(0.2,0.1) truncated to (0.01,0.9); μ_γ ~ N(1.5,0.8) truncated >0; μ_β ~ N(4.0,2.0) truncated >0; μ_v_gain ~ N(1.5,0.5) truncated >0; μ_a ~ N(1.2,0.3) truncated >0; μ_t_0 ~ N(0.3,0.1) truncated >0; σ_k ~ HalfNormal(0.5). Laplace standard errors from the MAP Hessian were floored at 0.05; Hessian failure set se=0.30. The 95% HDI is the shortest interval containing 95% of the sampled posterior, not an equal-tailed credible interval [23,24,42]. D2 used the same NUTS settings on the two-parameter probe. This manuscript reads archived posteriors for the preregistered H1–H4 tests; post-review grids used those archives plus the virtual-subject simulator and did not refit Stage-4 CPSP-8 NUTS.

Results

Numbers below are copied from openneuro_validation/derivatives_htcv/htcv_result.json and the recovery JSON files cited in each subsection. No value was rounded away from those archives except where a table uses three-decimal display of a four-decimal HDI already reported in the Abstract.

H1. Task-specific decision parameters (Figure 1; Table 1)

H1 required every non-γ decision parameter to be consistent across BART and CCT. All five tests failed. Figure 1 shows the HDI forest; Table 1 lists pooled BART means, CCT means, differences, and shortest 95% HDIs.

Threat

BART μ=-0.828, CCT μ=0.049, Δ=-0.878, 95% HDI [-1.153, -0.588], excludes 0. Threat cost of the current q is expressed on BART (unknown explosion hazard) and near zero on CCTHot (explicit p_loss). Verdict: inconsistent.

Baseline

BART μ=0.938, CCT μ=0.032, Δ=0.906, 95% HDI [0.767, 1.047], excludes 0. The additive preference term is absorbed by BART's pump-versus-cash structure and is essentially unused when CCT supplies gain, loss, and probability as separate channels. Verdict: inconsistent.

Initial learning rate α_0

BART μ=0.187, CCT μ=0.046, Δ=0.141, 95% HDI [0.114, 0.171], excludes 0. Tonic learning is higher when explosion probability must be inferred than when card statistics are displayed. Verdict: inconsistent.

Inverse temperature β

BART μ=1.769, CCT μ=0.001, Δ=1.769, 95% HDI [1.584, 1.949], excludes 0. Choice determinism does not transfer. This is the same parameter that, within D1 only, is the sole SCHZ≠CONTROL contrast (hierarchical difference 95% HDI [-0.8083, -0.0670], excludes 0). A within-task group difference is not a cross-task endophenotype, and β is not "higher certainty" in patients. Verdict: inconsistent across tasks.

Value gain v_gain

BART μ=1.900, CCT μ=0.019, Δ=1.882, 95% HDI [1.757, 2.005], excludes 0. Drift scaling of the RT mean and the CCT expected-value term do not occupy the same numerical range. Verdict: inconsistent.

H1 is therefore false: the decision layer, except γ, is rewritten by the task graph.

Table 1. Hierarchical means and HTCV differences μ_BART,pool-μ_CCT. Source: htcv_result.json tests. except γ from gamma.bart_cct; means from d1_posterior_extract.csv / d5_posterior_extract.csv.

Parameter	BART μ	CCT μ	Δ	95% HDI	Contains 0	Verdict
threat	-0.828	0.049	-0.878	[-1.153, -0.588]	No	Inconsistent
baseline	0.938	0.032	0.906	[0.767, 1.047]	No	Inconsistent
α_0	0.187	0.046	0.141	[0.114, 0.171]	No	Inconsistent
β	1.769	0.001	1.769	[1.584, 1.949]	No	Inconsistent
v_gain	1.900	0.019	1.882	[1.757, 2.005]	No	Inconsistent
γ	1.482	1.101	0.384	[-0.008, 0.783]	Yes	Consistent

H2. Cross-task and cross-modal consistency of γ (Figure 2; Table 2)

All three preregistered γ contrasts had HDIs containing 0 (Figure 2; Table 2). D2 hierarchical mean μ_γ=1.423, 95% HDI [1.286, 1.554]. Absolute γ is not interpreted as "CCT participants have stronger precision gating." H2 is true inside this family under the preregistered NUTS HDI-contains-zero rule, subject to the degrade sentence and to the post-review ROPE, conjugate-grid, and WAIC results below.

Table 2. γ contrasts. Source: htcv_result.json gamma.. "EEG" denotes the HierPrior behavioural probe (D2), not an ERP parameter.

Contrast	Δ	95% HDI	Contains 0
BART vs CCT	0.384	[-0.008, 0.783]	Yes
BART vs EEG	0.059	[-0.203, 0.298]	Yes
CCT vs EEG	-0.324	[-0.690, 0.041]	Yes

Post-review: ROPE, pipeline grid, WAIC/LOO, demographics, non-null family separability

Archived NUTS BART–CCT Δγ: mean 0.379, 95% HDI [−0.011, 0.775], contains 0. Posterior mass $|\Delta\gamma|<0.10$ / 0.15 / 0.20 = 0.079 / 0.124 / 0.182. HDI+ROPE decision at all three ε: undecided (HDI overlaps ROPE and is not contained).

Conjugate MAP+SE grid specified by the review (prior archived/half/double; SE floor {0.02, 0.05, 0.10}; pooling including SCHZ-only): 36 cells per parameter. For γ, 0/36 cells had HDI containing 0 (frac_match_archived = 0). Non-γ decision parameters that already excluded 0 under NUTS remained excluding 0 on 36/36 or 35/36 cells (threat, baseline, α_0, β, v_gain match 1.00; a match 0.972). t_0 contained 0 on 36/36, matching the archived NUTS verdict.

WAIC/LOO on the measurement-error layer (N=277 MAP+SE rows, BART+CCT): shared-γ WAIC = 1565.3, task-specific WAIC = 1418.7, ΔWAIC (shared − specific) = 146.6; ΔLOO-IC = 146.7. The shared-γ model is not preferred.

Demographics of fitted IDs: BART age 33.01 (SD 9.10), 41.0% female, MAP γ 1.491 (N=173); CCT age 23.95 (SD 5.67), 65.0% female, MAP γ 1.757 (N=103); EEG probe age 24.95 (SD 4.43), 64.3% female, MAP γ 1.485 (N=42). Logistic propensity of BART membership: age coefficient 0.162, female coefficient −0.713. IPW-weighted MAP γ: BART 1.491, CCT 1.823, Δ = −0.333; bootstrap 95% HDI [−0.789, 0.192], contains 0. Meta-regression of MAP γ on task + age + sex: is_BART −0.281, age 0.00286, female 0.049.

Table 2b. Non-null family separability. Generator: precision-gated delta-rule. Nested BF of γ-free vs γ=0. 10 cohorts × 100 subjects × 3 seeds (30 cells per row). Source: results/pw_gamma_positivity_result.json.

γ_true	trials	median BF	P(BF>3)	P(BF>10)
0.25	480	0.023	0.012	0.002
0.25	960	0.018	0.018	0.009
0.25	1920	0.021	0.038	0.020
0.50	480	0.071	0.071	0.033

0.50	960	0.127	0.160	0.101
0.50	1920	0.923	0.389	0.301
1.00	480	2.24	0.466	0.351
1.00	960	93.3	0.741	0.652
1.00	1920	3.31×10^5	0.929	0.887
1.50	480	340	0.810	0.729
1.50	960	4.69×10^6	0.960	0.931
1.50	1920	4.85×10^8	0.995	0.993
2.00	480	2.04×10^5	0.942	0.902
2.00	960	4.85×10^8	0.993	0.985
2.00	1920	4.85×10^8	1.000	1.000

D1e (true γ=0) remains BF ≈ 1.05–1.09. P(BF>3) first exceeds 0.5 at γ=1.0 with 960 trials. At 480 trials the 0.5 threshold is first crossed at γ=1.5.

H3. OU layer untestable under freeze

Because {θ,σ_ou,κ,λ,onset} were frozen at group means, no posterior contrast on these quantities exists. H3 is therefore not a negative empirical finding about threat recovery in patients; it is a declaration of unfalsifiability under the honesty constraint. Virtual-subject θ recovery never exceeded r=0.3 (see below), which is why the freeze was imposed before HTCV, not after seeing H1/H2.

H4. Fair model comparison (Figure 3; Table 3)

PPC coverage was 2/3 summary checks on BART and 3/3 on CCT. Fair RMSE, however, favoured dedicated models on every layer (Figure 3; Table 3). HGF-2 choice RMSE was 2.8699 on D1 and 0.5215 on D5-better than CPSP-8 on choice, still not a Mathys HGF, and without an RT layer. H4 overall is false (h4.d1_not_better=0; h4.d5_not_worse=0). Required wording: CPSP-8 jointly predicts choice and RT, but neither layer is better than a dedicated model.

Table 3. Layer-wise RMSE. Source: htcv_result.json models.d1 / models.d5.

Task	Layer	CPSP-8	Dedicated winner	HGF-2
D1 BART	choice RMSE	3.6253	RW 1.6158	2.8699
D1 BART	RT RMSE	0.0712	DDM 0.0615	-
D5 CCT	choice RMSE	1.1783	logistic 0.0082	0.5215
D5 CCT	RT RMSE	0.3591	DDM 0.3167	-

CCT pass has no independent event code; non-draw mapped to safe remains a high-risk mapping and is a limitation on D5 choice RMSE, including the near-zero logistic error.

Virtual-subject recovery (Figure 4)

D1d 960-trial γ recovery: Pearson r=0.612 / 0.660 / 0.701 (seeds in R960.corr.gamma, R960_7, R960_123); median ≈ 0.66. The 480-trial condition is weaker (archived alongside 960; Figure 4 reports the 960 standard). M3 (OU frozen, nine-parameter observation model) γ r=0.688 / 0.697 / 0.712. D1e: under true γ=0, nested Bayes-factor medians were ≈ 1.05-1.09 [43]; e1_ok=false, e2_ok=false (RW family inseparable). θ recovery: v1 profile range 0.0000, Pearson r=-0.0168; v2 r=0.242; v3 r=0.202; M1-960 r=0.293/0.260/0.299; M1-1920 r=0.260. A companion discrete-choice Fisher analysis found a non-positive-definite observed Hessian under softmax-plus-Gaussian-RT, so archived Cramér-Rao numbers from that Hessian are not a legal bound; behavioural blindness of θ rests on the r<0.3 recoveries, not on an inverted information matrix.

Discussion

Precision weighting as a cross-task cognitive endophenotype

In this family, only γ transferred. That is a bounded positive result, closer in spirit to treating HGF tonic volatility ω as a transdiagnostic candidate [11,16,19] than to a diagnostic cut-off. Gottesman and Gould's endophenotype requires heritability, state-independence, and familial co-segregation [3]. We did not measure those. What we measured is the computational prerequisite: a parameter that is not wholly rewritten when the task graph changes from unknown explosion hazard to known card statistics to a hierarchical-prior probe [4,5,40]. Huys, Maia and Frank argued that computational psychiatry earns its keep when parameters index processes rather than tasks [1,39]. γ meets that weaker standard here; it does not yet meet the classical endophenotype standard, and it must not be sold as a schizophrenia assay.

The theoretical reason γ is a plausible core is the same reason precision is central in predictive coding. Belief updating is not a fixed α; it is a gain on prediction error that should rise when the world is more surprising than the current uncertainty budget and fall when it is not [10-15,41]. Powers, Mathys and Corlett's conditioned-hallucination experiment remains the clearest causal demonstration that overweighting priors (a precision imbalance) can produce a psychotic-like percept [16]. Our γ is not their prior precision, and IsolatedHGFChoice is not their HGF. The analogy is architectural: a scalar that gates how much a delta-rule may learn from a squared error. In Friston's hierarchical inference, that scalar is expected precision; in the HGF it is a function of ω and ϑ; in CPSP-8 it is γ [10-13]. Cross-task stability of γ is therefore evidence that this gain, unlike threat loadings and inverse temperatures, is not merely a description of one graph.

Task specificity as cognitive mode switching, not software failure

Threat, baseline, α_0, β, and v_gain failed H1 by large, decisive HDIs. We read this as theory, not implementation failure. BART is a learning problem: explosion probability is hidden, pumps accumulate, and cash-out is the only safe act [27]. CCTHot is an information-use problem: gain, loss, and probability are on the screen, and the affective "hot" procedure is designed to engage sequential, arousing risk rather than a one-shot gamble [28]. Context effects in decision theory are the rule, not the exception [8,35]. A parameter that absorbs BART's hidden hazard (threat, α_0) has little work to do when p_loss is displayed; a β that determinizes a learned dv is not the same object as a β on an explicit expected value. Pike and colleagues' finding that only inverse-temperature-like quantities sometimes correlate across reward tasks, and even then weakly, is the broader empirical backdrop [5]. H1 false is therefore informative: CPSP's decision layer contains both a portable gain (γ) and a set of mode switches. Calling the latter "unstable" would mistake a changed computation for a failed fit.

OU: theoretically central, empirically blind

Behavioural blindness is not theoretical vacancy. An OU threat state is the natural continuous-time model for a quantity that should relax after an insult and wander under process noise; dissociation drift is the natural model for a delayed precision failure. Those claims do not become false because BART and CCT do not identify θ. They become untestable, which is why they were frozen. The freeze tuple (θ=0.285, σ_ou=0.1425, κ=0.30, λ=0.35, onset=1.5) is a group-level plug-in. It is forbidden to narrate it as "patients' θ." A discrete-choice Fisher analysis was run as an auxiliary check; the observed Hessian was not positive definite, so it cannot be inverted into a Cramér-Rao bound. The behavioural evidence remains the recovery correlations below 0.3. The predicted remedies are external anchors-neural time series, ecological momentary assessment (EMA) of threat and dissociation, or tasks with explicit threat-offset epochs-not another 1920 mixed BART trials of the same likelihood [2,21].

Joint models, dedicated models, and the absence of a free lunch

CPSP-8's explanatory advantage is architectural: one generative story for threat, precision-gated learning, and RT means. Its predictive disadvantage is the no-free-lunch fact that a specialist likelihood wins on the scalar it was built to predict [36]. RW has two parameters and no RT; logistic regression on CCT is almost a restatement of the displayed features; DDM is an evidence-accumulation specialist [32,33]. Palminteri, Wyart and Koechlin warned that winning a comparison is not the same as generating the phenomenon of interest [6]. We therefore refuse the sentence "CPSP is better." The long-term value of joint modelling in computational psychiatry is constraint: a parameter that must serve choice and RT and a second task is less free to overfit one session [1,2,6]. That value is not a leaderboard medal.

Clinical translation, with the brakes on

γ is a candidate cross-task cognitive marker in schizophrenia-spectrum research only in the weak sense that it transferred in a sample that includes SCHZ on BART and a mixed community sample on CCT. We did not show that γ differs by diagnosis (Stage 4′ s γ group contrast was not the surviving clinical effect), nor that it predicts conversion, nor that it is heritable [3,19]. β is the only SCHZ≠CONTROL parameter in the D1 hierarchical model (95% HDI [-0.8083, -0.0670]). Inverse temperature differences can arise from exploration, lapsing, utility scaling, or mapping artefacts; they must not be rewritten as "patients decide more deterministically." CCT high-risk mapping (non-draw → safe) further limits any claim that D5 choice error is a clinical phenotype. Translation, if it comes, will need neural or EMA anchors, larger and more representative samples, and a design that can separate γ from RW at the family level [44].

A further clinical caution concerns what is not in the HTCV table. CCTHot has no diagnosis column, so γ′ s cross-task stability is not a replication of a schizophrenia effect; it is a replication of a computational scalar across graphs. BART′ s SCHZ contrast lives on β, a parameter that failed H1. Prior BART studies in schizophrenia have reported behavioural risk aversion, not a portable computational marker [37,38]. Using a task-specific inverse temperature as if it were a portable cognitive endophenotype would invert the logic of this paper. Using a portable γ as if it were a diagnostic likelihood ratio would ignore D1e. The clinically responsible sentence is therefore conjunctive: γ is the only decision parameter that transfers, and it is still family-inseparable from RW, and β is the only parameter that differs by diagnosis inside BART, and that difference is not a certainty narrative.

The degrade sentence, in full

Precision-gated delta-rule: individual γ is recoverable; family-inseparable from RW at the family level when the generator is the RW limit (γ=0); mechanism absence is not detectable (one-sided: presence can be justified, absence cannot be justified). Non-null probes show family separability when true $\gamma \geq 1.0$ at ≥960 trials (or $\gamma \geq 1.5$ at 480 trials); they do not show it at $\gamma \leq 0.5$.

The methodological meaning is strict. D1d and M3 justify talking about person-level γ inside this observation model. D1e Bayes factors near 1 deny talking about γ as a mechanism distinct from RW when the truth is the RW-generate limit [43]. The non-null grid supplies the missing positive test: family separability is empirical when precision gain is large enough, and remains architectural when it is not. A non-significant nested test cannot prove that precision gating is absent; it can only fail to prove that it is present. One-sided detectability is an honesty constraint on rhetoric, parallel to the OU freeze: we do not convert a flat likelihood into a negative ontology.

Cross-task HTCV contrasts are between independent cohorts and may conflate task structure with sample composition. We did not harmonize units across decision variables; raw differences may partially reflect scale rather than cognitive differences. Within-subject replication across BART and CCTHot was not tested; cross-cohort exchangeability does not guarantee individual-level portability. HTCV used MAP-then-summary handoff; Laplace approximation error remains unquantified.

Limitations and future directions

Several limitations are design features rather than afterthoughts. (ROPE) Formal equivalence of γ across tasks is not established: the BART–CCT 95% HDI overlaps ROPE(±0.20) but is not contained in it. (Pipeline) The preregistered NUTS HDI contains 0; the reviewer conjugate MAP+SE grid (prior half/double, SE floor {0.02, 0.05, 0.10}, pooling including SCHZ-only) yielded 0/36 γ cells containing 0. Cross-task HDI-contains-zero is therefore pipeline-sensitive. (Shared vs specific) WAIC/LOO on MAP+SE preferred task-specific μ_γ (ΔWAIC = 146.6). (Cohorts) Cross-task HTCV contrasts are between independent cohorts and may conflate task structure with sample composition. Age/sex IPW bootstrap HDI for MAP Δγ contained 0, but the samples differ (BART older, fewer women). (Scale) We did not harmonize units across decision variables; raw differences may partially reflect scale rather than cognitive differences. (Within-subject) Within-subject replication across BART and CCTHot was not tested. (Handoff) HTCV used MAP-then-summary handoff; Laplace approximation error remains unquantified. (Family) D1e probed separability only in the RW-generate limit (γ=0). The non-null grid is a BIC approximation, not a nested

NUTS BF. (i) D2 γ is a behavioural probe in an EEG protocol, not an ERP parameter; "cross-modal" here means cross-likelihood, not brain-behaviour. (ii) CCTHot has no independent pass code; mapping non-draw to safe is high risk and inflates the apparent simplicity of logistic RMSE. (iii) No subject-level recovery scatter was archived; Figure 4 cannot be a cloud plot. (iv) D1 is a United States mixed clinical BART sample from a phenomics consortium [25]; it is not a first-episode, medication-naive, or multi-site schizophrenia cohort. (v) IsolatedHGFChoice must never be cited as a Mathys HGF result. (vi) Discrete Fisher information under the choice-RT likelihood is not a usable Cramér-Rao bound. The OU freeze remains an honesty constraint: {θ, σ_ou, κ, λ, onset} stay at group means because BART/CCT do not identify them.

Future work should (a) anchor γ and, separately, θ to neural or EMA series rather than to more of the same balloons; (b) redesign tasks with threat-offset epochs if the OU layer is to be identified; (c) replace BIC-BF with nested NUTS on the same non-null grid; (d) collect within-subject BART+CCT; (e) test Gottesman criteria (relatives, heritability) only after (a)–(d); (f) extend HTCV transdiagnostically rather than treating SCHZ as the sole clinical pole [1,5,19,40].

Conclusions

Cross-task generalization is not all-or-none. Under preregistered NUTS HTCV, the CPSP-8 update coefficient γ is the only decision scalar whose hierarchical-mean difference HDIs contain 0; ROPE equivalence is not established, and a conjugate MAP+SE grid specified after review does not reproduce HDI-contains-zero. Family-level identity as precision weighting rather than a learning-rate-like scalar is supported in silico when $\gamma \geq 1.0$ at ≥960 trials (or $\gamma \geq 1.5$ at 480 trials) and is not supported at $\gamma=0$ (D1e) or $\gamma \leq 0.5$. Threat, baseline, α_0, β, and v_gain are mode switches; the OU layer awaits an observation model that can see it. Joint prediction of choice and RT is not victory over RW or DDM. The honest claims are smaller than a diagnostic biomarker and larger than a single-task fit.

Data and code availability

Behavioural events were obtained from OpenNeuro under CC0 and are cited with accession numbers: UCLA Consortium for Neuropsychiatric Phenomics LA5c BART, ds000030 (https://openneuro.org/datasets/ds000030; doi:10.18112/openneuro.ds000030.v1.0.0) [25,26]; hot Columbia Card Task, ds004636 (https://openneuro.org/datasets/ds004636; doi:10.18112/openneuro.ds004636.v1.0.4) [28,29]; HierPrior behavioural events, ds008083 (https://openneuro.org/datasets/ds008083; doi:10.18112/openneuro.ds008083.v1.0.1). Large source data remain in those community databases. Preregistered HTCV and recovery numbers were read from archived posterior files listed in the Supplementary Materials. Post-review grids, download scripts, and analysis code are in manuscripts/gamma_endophenotype/code/ (download_openneuro.py, analysis_pipeline.py) with numerical products in manuscripts/gamma_endophenotype/results/. Core model files (cpsp_framework_core.py, cpsp_2027_core.py) were not modified.

Ethics statement

This study did not recruit new human participants and did not collect new identifiable data. All analyses used de-identified public releases. For ds000030, participants gave written informed consent under procedures approved by the University of California Los Angeles Institutional Review Board, as described in the dataset record and by Poldrack and colleagues [25]. For ds004636, original data collection was approved by the Stanford Institutional Review Board (protocol 39322) [29]. For ds008083, only de-identified behavioural events released on OpenNeuro were used; no EEG time series were analysed. Secondary analysis of these public de-identified records does not constitute new human-subjects research.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The original data collections were funded as described in the OpenNeuro records and source publications [25,29]; those funders had no role in the present secondary analysis.

Competing interests

The author declares that there are no competing interests.

Author contributions

K.W. conceived the study, performed the analyses, and wrote the manuscript.

Acknowledgements

I thank the OpenNeuro curators and the original investigators of ds000030, ds004636, and ds008083, and acknowledge those resources in accordance with their data-use terms.

Supplementary materials

Supplementary tables, the download script, the analysis pipeline, archived JSON/NPZ paths, and post-review numerical products accompany this manuscript in manuscripts/gamma_endophenotype/.

References

[1] Huys QJM, Maia TV, Frank MJ. Computational psychiatry as a bridge from neuroscience to clinical applications. Nat Neurosci. 2016;19(3):404-413. doi:10.1038/nn.4238

[2] Stephan KE, Mathys C. Computational approaches to psychiatry. Curr Opin Neurobiol. 2014;25:85-92. doi:10.1016/j.conb.2013.12.007

[3] Gottesman II, Gould TD. The endophenotype concept in psychiatry: etymology and strategic intentions. Am J Psychiatry. 2003;160(4):636-645. doi:10.1176/appi.ajp.160.4.636

[4] Wilson RC, Collins AGE. Ten simple rules for the computational modeling of behavioral data. eLife. 2019;8:e49547. doi:10.7554/eLife.49547

[5] Pike AC, Board R, Travers E, Croal M, Lam YC, Valton V, Robinson OJ, Roiser JP. The recoverability, reliability, and generalizability of reward processing parameters and relation to mental health symptoms. Psychol Med. 2026;56:e238. doi:10.1017/S0033291726105340

[6] Palminteri S, Wyart V, Koechlin E. The importance of falsification in computational cognitive modeling. Trends Cogn Sci. 2017;21(6):425-433. doi:10.1016/j.tics.2017.03.011

[7] Pike AC, Robinson OJ. Reinforcement learning in patients with mood and anxiety disorders vs control individuals: a systematic review and meta-analysis. JAMA Psychiatry. 2022;79(4):313-322. doi:10.1001/jamapsychiatry.2022.0051

[8] Adams RA, Huys QJM, Roiser JP. Computational psychiatry: towards a mathematically informed understanding of mental illness. J Neurol Neurosurg Psychiatry. 2016;87(1):53-63. doi:10.1136/jnnp-2015-310737

[9] Braff DL, Freedman R, Schork NJ, Gottesman II. Deconstructing schizophrenia: an overview of the use of endophenotypes in order to understand a complex disorder. Schizophr Bull. 2007;33(1):21-32. doi:10.1093/schbul/sbl049

[10] Mathys C, Daunizeau J, Friston KJ, Stephan KE. A Bayesian foundation for individual learning under uncertainty. Front Hum Neurosci. 2011;5:39. doi:10.3389/fnhum.2011.00039

[11] Mathys CD, Lomakina EI, Daunizeau J, Iglesias S, Brodersen KH, Friston KJ, Stephan KE. Uncertainty in perception and the Hierarchical Gaussian Filter. Front Hum Neurosci. 2014;8:825. doi:10.3389/fnhum.2014.00825

[12] Friston K. A theory of cortical responses. Philos Trans R Soc Lond B Biol Sci. 2005;360(1456):815-836. doi:10.1098/rstb.2005.1622

[13] Friston K. The free-energy principle: a unified brain theory? Nat Rev Neurosci. 2010;11(2):127-138. doi:10.1038/nrn2787

[14] Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behav Brain Sci. 2013;36(3):181-204. doi:10.1017/S0140525X12000477

[15] Kanai R, Komura Y, Shipp S, Friston K. Cerebral hierarchies: predictive processing, precision and the pulvinar. Philos Trans R Soc Lond B Biol Sci. 2015;370(1668):20140169. doi:10.1098/rstb.2014.0169

[16] Powers AR, Mathys C, Corlett PR. Pavlovian conditioning-induced hallucinations result from overweighting of perceptual priors. Science. 2017;357(6351):596-600. doi:10.1126/science.aan3458

[17] Iglesias S, Mathys C, Brodersen KH, Kasper L, Piccirelli M, den Ouden HEM, Stephan KE. Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. Neuron. 2013;80(2):519-530. doi:10.1016/j.neuron.2013.09.009

[18] Lawson RP, Mathys C, Rees G. Adults with autism overestimate the volatility of the sensory environment. Nat Neurosci. 2017;20(9):1293-1299. doi:10.1038/nn.4615

[19] Mikus N, Lamm C, Mathys C. Computational phenotyping of aberrant belief updating in individuals with schizotypal traits and schizophrenia. Biol Psychiatry. 2025;97(2):188-197. doi:10.1016/j.biopsych.2024.08.021

[20] Weber LA, Waade PT, Legrand N, Moller AH, Stephan KE, Mathys C. The generalized Hierarchical Gaussian Filter. eLife. 2025;14:RP110174. doi:10.7554/eLife.110174

[21] Frässle S, Aponte EA, Raman S, Harrison SJ, Heinzle J, Iglesias S, Kasper L, Lomakina EI, Mathys C, Müller-Schrader M, Pereira I, Reiter A, Schöbi D, Yao Y, Diaconescu AO, Brodersen KH, Stephan KE. TAPAS: an open-source software package for translational neuromodeling and computational psychiatry. Front Psychiatry. 2021;12:680811. doi:10.3389/fpsyt.2021.680811

[22] Katthagen T, Fromm S, Wieland L, Schlagenhauf F. Models of dynamic belief updating in psychosis: a review across different computational approaches. Front Psychiatry. 2022;13:814111. doi:10.3389/fpsyt.2022.814111

[23] Kruschke JK. Bayesian estimation supersedes the t test. J Exp Psychol Gen. 2013;142(2):573-603. doi:10.1037/a0029146

[24] Kruschke JK. Doing Bayesian data analysis: a tutorial with R, JAGS, and Stan. 2nd ed. London: Academic Press; 2015. ISBN: 978-0-12-405888-0.

[25] Poldrack RA, Congdon E, Triplett W, Gorgolewski KJ, Karlsgodt KH, Mumford JA, Sabb FW, Freimer NB, London ED, Cannon TD, Bilder RM. A phenome-wide examination of neural and cognitive function. Sci Data. 2016;3:160110. doi:10.1038/sdata.2016.110

[26] Markiewicz CJ, Gorgolewski KJ, Feingold F, Blair R, Halchenko YO, Miller E, Hardcastle N, Wexler J, Esteban O, Goncalves M, Jwa A, Poldrack R. The OpenNeuro resource for sharing of neuroscience data. eLife. 2021;10:e71774. doi:10.7554/eLife.71774

[27] Lejuez CW, Read JP, Kahler CW, Richards JB, Ramsey SE, Stuart GL, Strong DR, Brown RA. Evaluation of a behavioral measure of risk taking: the Balloon Analogue Risk Task (BART). J Exp Psychol Appl. 2002;8(2):75-84. doi:10.1037/1076-898X.8.2.75

[28] Figner B, Mackinlay RJ, Wilkening F, Weber EU. Affective and deliberative processes in risky choice: age differences in risk taking in the Columbia Card Task. J Exp Psychol Learn Mem Cogn. 2009;35(3):709-730. doi:10.1037/a0014983

[29] Eisenberg IW, Bissett PG, Enkavi AZ, Li J, MacKinnon DP, Marsch LA, Poldrack RA. Uncovering the structure of self-regulation through data-driven ontology discovery. Nat Commun. 2019;10:2319. doi:10.1038/s41467-019-10301-1

[30] Rogers RD, Everitt BJ, Baldacchino A, Blackshaw AJ, Swainson R, Wynne K, Baker NB, Hunter J, Carthy T, Booker E, London M, Deakin JFW, Sahakian BJ, Robbins TW. Dissociable deficits in the decision-making cognition of chronic amphetamine abusers, opiate abusers, patients with focal damage to prefrontal cortex,

and tryptophan-depleted normal volunteers: evidence for monoaminergic mechanisms. Neuropsychopharmacology. 1999;20(4):322-339. doi:10.1016/S0893-133X(98)00091-8

[31] Rescorla RA, Wagner AR. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Black AH, Prokasy WF, editors. Classical conditioning II: current research and theory. New York: Appleton-Century-Crofts; 1972. p. 64-99.

[32] Ratcliff R, McKoon G. The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. 2008;20(4):873-922. doi:10.1162/neco.2008.12-06-420

[33] Ratcliff R, Smith PL, Brown SD, McKoon G. Diffusion decision model: current issues and history. Trends Cogn Sci. 2016;20(4):260-281. doi:10.1016/j.tics.2016.01.007

[34] Hoffman MD, Gelman A. The No-U-Turn Sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. J Mach Learn Res. 2014;15:1593-1623. doi:10.48550/arXiv.1111.4246

[35] Tversky A, Kahneman D. The framing of decisions and the psychology of choice. Science. 1981;211(4481):453-458. doi:10.1126/science.7455683

[36] Wolpert DH, Macready WG. No free lunch theorems for optimization. IEEE Trans Evol Comput. 1997;1(1):67-82. doi:10.1109/4235.585893

[37] Reddy LF, Lee J, Davis MC, Altshuler L, Glahn DC, Miklowitz DJ, Green MF. Impulsivity and risk taking in bipolar disorder and schizophrenia. Neuropsychopharmacology. 2014;39(2):456-463. doi:10.1038/npp.2013.218

[38] Cheng GLF, Tang JCY, Li FWS, Lau EYY, Lee TMC. Schizophrenia and risk-taking: impaired reward but preserved punishment processing. Schizophr Res. 2012;136(1-3):122-127. doi:10.1016/j.schres.2012.01.002

[39] Montague PR, Dolan RJ, Friston KJ, Dayan P. Computational psychiatry. Trends Cogn Sci. 2012;16(1):72-80. doi:10.1016/j.tics.2011.11.018

[40] Karvelis P, Paulus MP, Diaconescu AO. Individual differences in computational psychiatry: a review of current challenges. Neurosci Biobehav Rev. 2023;148:105137. doi:10.1016/j.neubiorev.2023.105137

[41] Rao RPN, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci. 1999;2(1):79-87. doi:10.1038/4580

[42] Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. Bayesian data analysis. 3rd ed. Boca Raton: CRC Press; 2013. doi:10.1201/9780429258480

[43] Wagenmakers EJ, Marsman M, Jamil T, Ly A, Verhagen J, Love J, Selker R, Gronau QF, Smira M, Epskamp S, Matzke D, Rouder JN, Morey RD. Bayesian inference for psychology. Part I: theoretical advantages and practical ramifications. Psychon Bull Rev. 2018;25(1):35-57. doi:10.3758/s13423-017-1343-3

[44] Stephan KE, Iglesias S, Heinzle J, Diaconescu AO. Translational perspectives for computational neuroimaging. Neuron. 2015;87(4):716-732. doi:10.1016/j.neuron.2015.07.008

Figure captions

Figure 1. Hierarchical transfer of CPSP-8 decision parameters between BART and CCT. Forest plot of posterior differences μ_BART,pool-μ_CCT with shortest 95% HDIs. Intervals excluding zero indicate task-specific expression (threat, baseline, α_0, β, v_gain). The γ interval contains zero. Data: OpenNeuro ds000030 (N=173) and ds004636 (N=104); source htcv_result.json.

Figure 2. Precision-weighting γ across BART, CCT, and a HierPrior behavioural probe. Three HTCV contrasts; all 95% HDIs contain 0. The probe (D2, N=43, ds008083) is behavioural, not ERP. Source: htcv_result.json gamma..

Figure 3. Layer-wise fair RMSE. Choice and RT errors for CPSP-8 versus RW, logistic regression, DDM, and simplified HGF-2. HGF-2 is not the Mathys HGF. CPSP-8 does not win either layer. Source: htcv_result.json models..

Figure 4. Virtual-subject recovery of γ versus blindness of θ. Pearson r for γ at 960 trials (three seeds) and M3, and θ from v1 through M1-1920. Subject-level scatter arrays were not archived; panels are r summaries, not fabricated clouds.

Figures

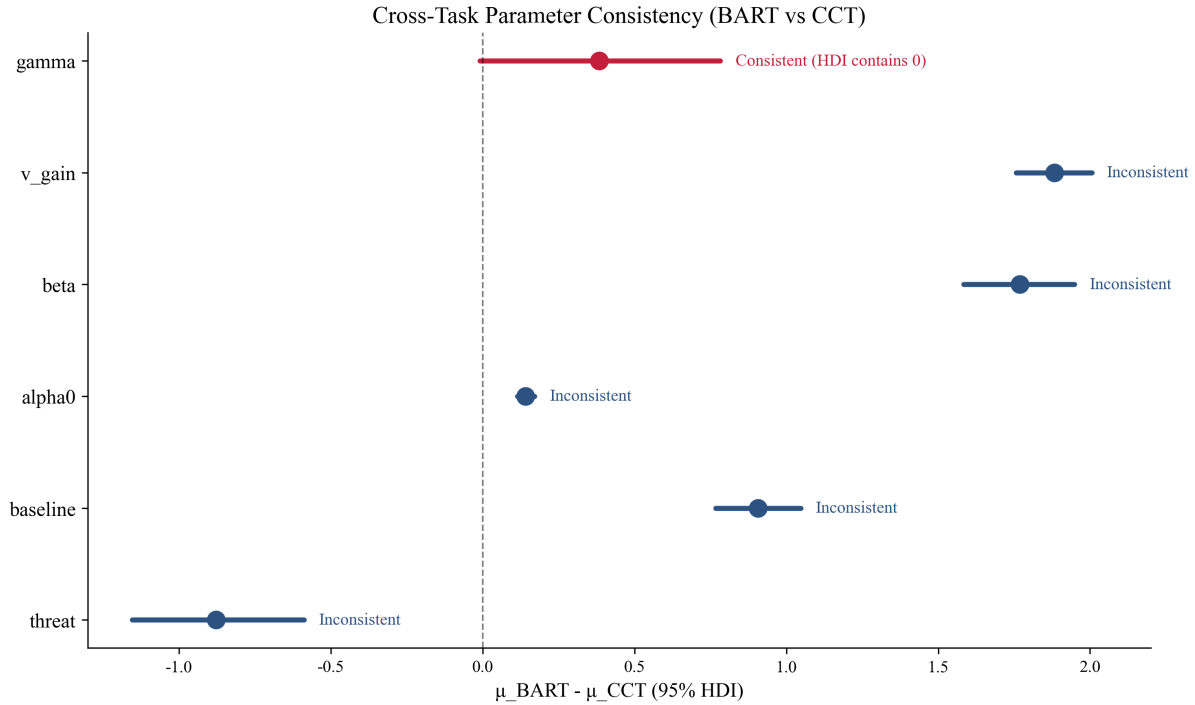


Figure 1. Hierarchical transfer of CPSP-8 decision parameters between BART and CCT.

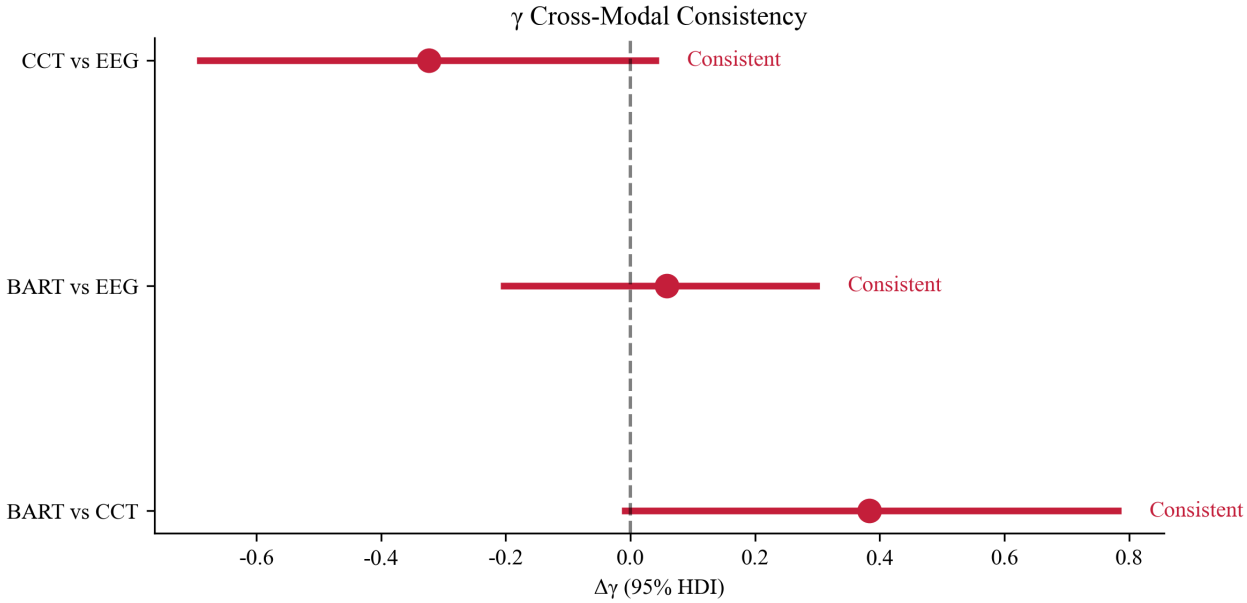


Figure 2. Precision-weighting gamma across BART, CCT, and a HierPrior behavioural probe.

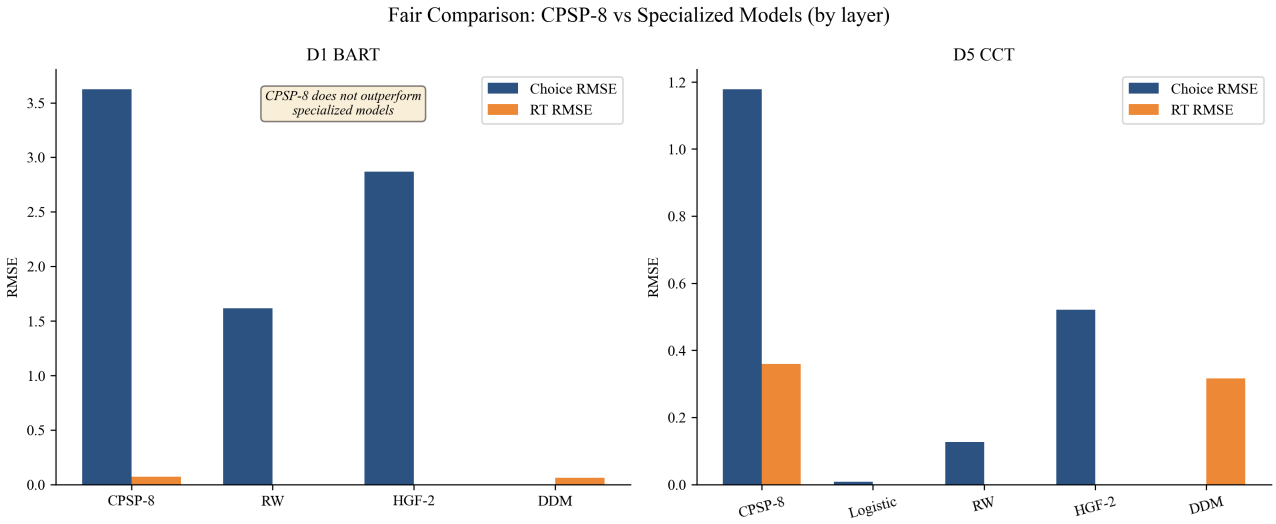


Figure 3. Layer-wise fair RMSE.

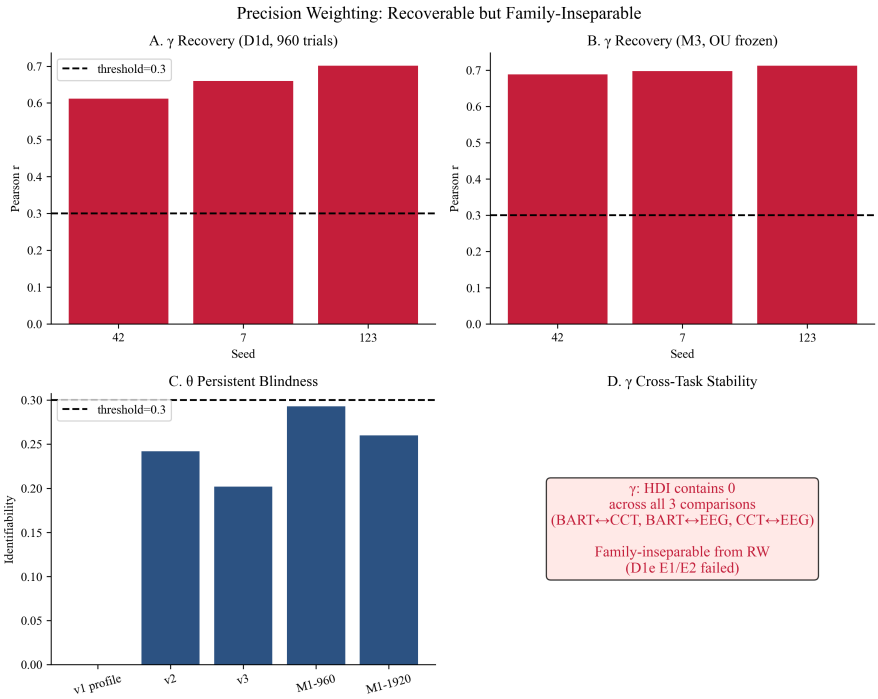


Figure 4. Virtual-subject recovery of gamma versus blindness of theta.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- 02supplementarycombined.pdf