Individual expectations and adaptive prediction of musical pleasure

Amaç Erdem | amace@bu.edu | ORCID 0009-0000-1711-0036

Abstract Musical pleasure varies with both the sound and the listener. This protocol asks whether reported expectations explain that variation and whether earlier responses improve prediction for an individual. After rating 20 familiar classical works, participants complete one of two controlled tasks: an uncertainty–surprise task or a randomized test of expectation updating. All then hear 24 transfer excerpts selected by four policies. Matched models test the contributions of personal references, expectation reports and MI features beyond an equally informed acoustic learner. Three primary claims share familywise error control. A 60-registration pilot within a cap of 1,000 will establish feasibility, measurement properties, model recovery and a justified confirmatory sample. The protocol separates evidence that listeners learn from evidence that a model learns to predict them. No participant results are reported.

Prospective author manuscript. All figures are conceptual or analytic illustrations.

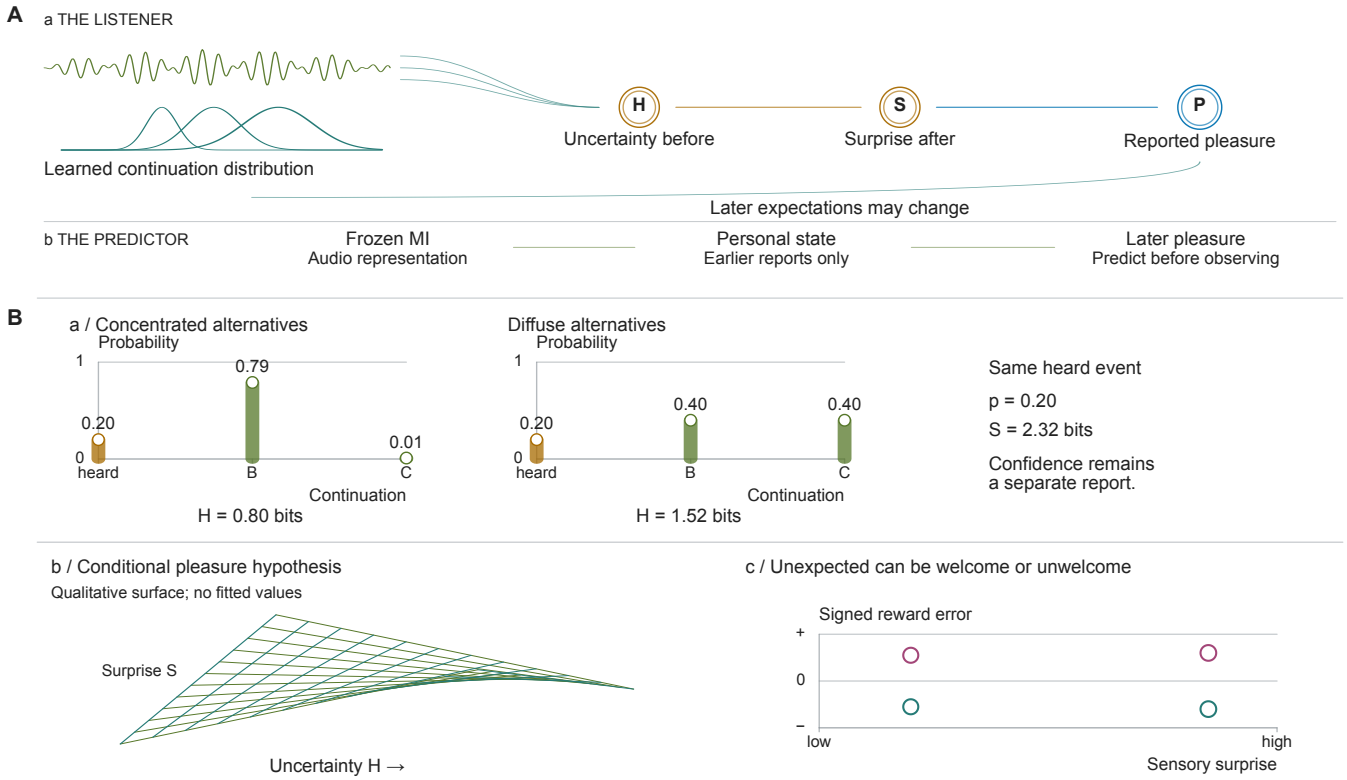


Fig. 1. Expectation and prediction **A,** Changes in a listener's expectations and updates to a predictor are separate learning processes. **B,** The example distributions give equal surprise (2.32 bits) at different entropies (0.80 and 1.52 bits). The pleasure surface is qualitative: it is neither a fit to Cheung et al.[2] nor the MI PUPF equation. The normalized axes do not show every feasible distribution. Sensory surprise can occur with either positive or negative value error.

Expectations, learning and musical pleasure

A listener brings a history of musical experience to each new passage. Predictive accounts propose that this history shapes expectations and, through them, pleasure[1]. Findings on uncertainty and surprise[2], intermediate predictability[3] and individual preferences[22] suggest that the same event need not be equally pleasurable for every listener. The present protocol asks whether reported expectations explain this variation and whether a model can use earlier observations to predict later pleasure.

Three questions require different evidence. Does uncertainty moderate the association between surprise and pleasure? Does a randomized change in exposure alter expectations? Do personal references and expectation reports improve prediction beyond an equally informed acoustic model? The study combines an association measured within a controlled task, a randomized behavioural comparison and an out-of-sample prediction test.

For listener i on trial t, $q_{it}(k)$ denotes the reported probability of each of three demonstrated continuations, and e_t is the continuation that occurs. Entropy and surprise are defined as

$$H_{it} = -\sum_{k=1}^{3} q_{it}(k)\log_2 q_{it}(k), \qquad S_{it} = -\log_2 q_{it}(e_t).$$

These measures describe expectations within the task; confidence is recorded separately. Expected enjoyment V_{it} and experienced pleasure Y_{it} use the same 0–100 scale. Their difference, $\delta_{it} = Y_{it} - V_{it}$, is a signed value error. The model's prediction residual, $r_{it} = Y_{it} - \widehat{Y}_{it}$, has a different meaning. Neither quantity is sensory surprise. A fall in entropy alone could reflect overconfidence; evidence of learning requires changes in later expectations, choices or accuracy.

The choice of pleasure as the outcome also matters. Unexpectedness-related arousal[4], acoustically predicted valence and arousal[23], anticipation and reward value[5, 6], and auditory–motor memory[24] are related but distinct targets. Reward-learning[7] and naturalistic expectancy studies[8] motivate the present comparisons. The browser task measures reported experience and behaviour, so its findings will concern those levels of explanation. Familiarity, acoustic salience, nonlinear complexity preferences and individual rating habits remain plausible alternatives.

Participants, materials and procedure

Recruitment The proposed cap is 1,000 distinct, verified registrations meeting the initial eligibility criteria. The first 60 registrants will enter a feasibility and parameter-setting pilot and will be excluded from confirmatory analyses. This leaves at most 940 registrations before attrition. A registration counts once after verification, age eligibility and initial screening; consent is recorded separately. Withdrawals and incomplete baselines remain in the participant flow and do not reopen places. Reporting will cover invitations, verification, screening, consent, reference completion, readiness, task completion, transfer and analysable endpoints.

Participants must be at least 18, able to consent in the study language, and able to use supported headphones with a desktop or lap-

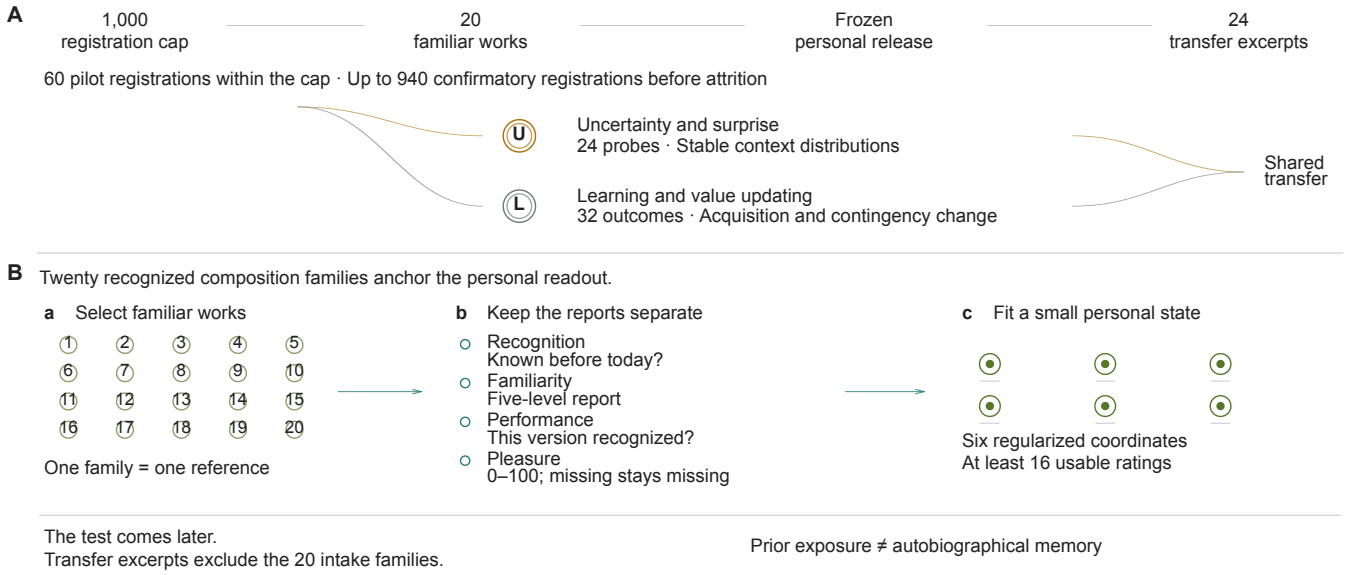


Fig. 2. Participants and personal references A, The proposed cap includes 60 pilot registrations, leaving at most 940 before readiness failures and attrition. Equal U/L assignment is followed by common transfer and four performance pairs. **B,** Twenty previously recognized families, with at least 16 usable pleasure ratings, initialize six regularized coordinates. Intake families are excluded from transfer. These selections describe familiarity and ratings; they do not measure autobiography or general musical ability.

top. Intake will record musical training, current practice, classical listening, language, age band, country, self-reported hearing difficulty and assistive hearing use, with missing or prefer-not-to-answer options where appropriate. Recruitment will span training and familiarity levels. The initial study will use one validated language; additional translations require cognitive interviewing and counterbalanced validation before data are pooled.

Catalogue Each stimulus will have separate identifiers for composition family, movement, arrangement, performance, recording master, excerpt and passage, together with duration, instrumentation and performer attribution. Movements cannot occupy separate intake positions, and re-encodings do not count as new performances. Candidate sources include individually cleared classical recordings and MAESTRO[18]. An audit will establish permitted use, territory, redistribution and attribution for each recording and derivative before the catalogue is frozen. Composition age and dataset access alone do not establish these permissions. The protocol does not propose downloading recordings from YouTube. Three performances per family are required only for the paired-performance subset.

Personal references During a separate intake visit, participants will use a searchable catalogue with 12-second previews to identify exactly 20 distinct families recognized before that day. Familiar but disliked works are eligible. For each reference, they will report prior recognition, five-level familiarity, recognition of the particular performance and preview pleasure on a 0–100 scale, with a separate unable-to-rate option. This procedure collects explicit familiarity and pleasure reports rather than inferring them from a streaming account. Repeated previews during onboarding cannot substitute for prior recognition.

Calibration requires at least 16 usable pleasure ratings. An incomplete intake can be saved, but missing rating decisions must be resolved before the session is prepared. Coefficients with little personal support will shrink towards the pilot-trained population prior. Search burden, completion rates and differences by listening background are pilot outcomes: requiring 20 familiar works may select a narrower population than a general listening study.

Experimental visit Following consent, browser/audio checks and intake, the system will prepare a fixed, versioned session. Participants will adjust playback to a comfortable level, complete headphone screening[19], and practise distinguishing likelihood, confidence, surprise and enjoyment. Persistent equipment or comprehension failures will be recorded before randomization. Screening assesses task suitability rather than calibrated acoustics. The visit then comprises one mechanistic arm, a break, a common transfer task and paired-performance comparisons. Combined intake and experiment time is provisionally 55–75 minutes; the pilot will provide the final estimate. Participants may pause between blocks or withdraw. Institutional responsibility, payment, recruitment arrangements, consent and contacts must be settled before invitations; compensation will not depend on ratings or agreement with the hypotheses.

Transfer and performance comparisons The common task contains 24 previously unpresented 25-second excerpts from families outside the intake set. Six each will be selected by personal MI, an equally calibrated acoustic model, matched random sampling and MI using another person's profile. Participants may nevertheless know a work from outside the study, so familiarity will be reported after playback. Standardized fades and fixed loudness normalization will preserve internal dynamics. Clipping, noise, unmatched truncation and inconsistent channel layouts will be grounds for rejecting a stimulus. Loudness, roughness, tempo, onset density, tonal stability and instrumentation will be retained as controls.

Four additional pairs will compare two previously unheard performances beyond the intake reference. Passage, instrumentation and approximate position will be matched, and B/C order counterbalanced. Participants will rate both performances and indicate a preference, with an indifferent option. Arrangement changes fall outside this comparison and would require a separate study.

Controlled expectation and learning tasks

Both arms use cleared or newly rendered classical-motif continuations. Identical target endings are available across contexts. Intended probability distributions, seeded sequences, delivered exposures and counterbalanced outcome identities will be recorded separately. Later pleasure ratings will not determine the schedule or trigger additional trials.

Before an outcome, participants will allocate 100 probability points among three demonstrated audible alternatives, then report confidence and expected enjoyment. The context will be replayed after these reports are locked to restore a defined auditory lead-in; reporting delay and replay will be logged. Surprise and pleasure will be reported after the ending. The procedure requires no notation knowledge and uses the same response order in both arms.

Arm U: uncertainty and surprise Two visually labelled contexts have distributions (.20, .75, .05) and (.20, .40, .40), with outcome identities permuted across listeners. Twenty-four training outcomes (12 per context) precede 24 probes. Training and probe sequences do not overlap; context order is balanced, with six-second contexts and three-second endings. Arousal will be reported on eight prespecified probes balanced by context and outcome class.

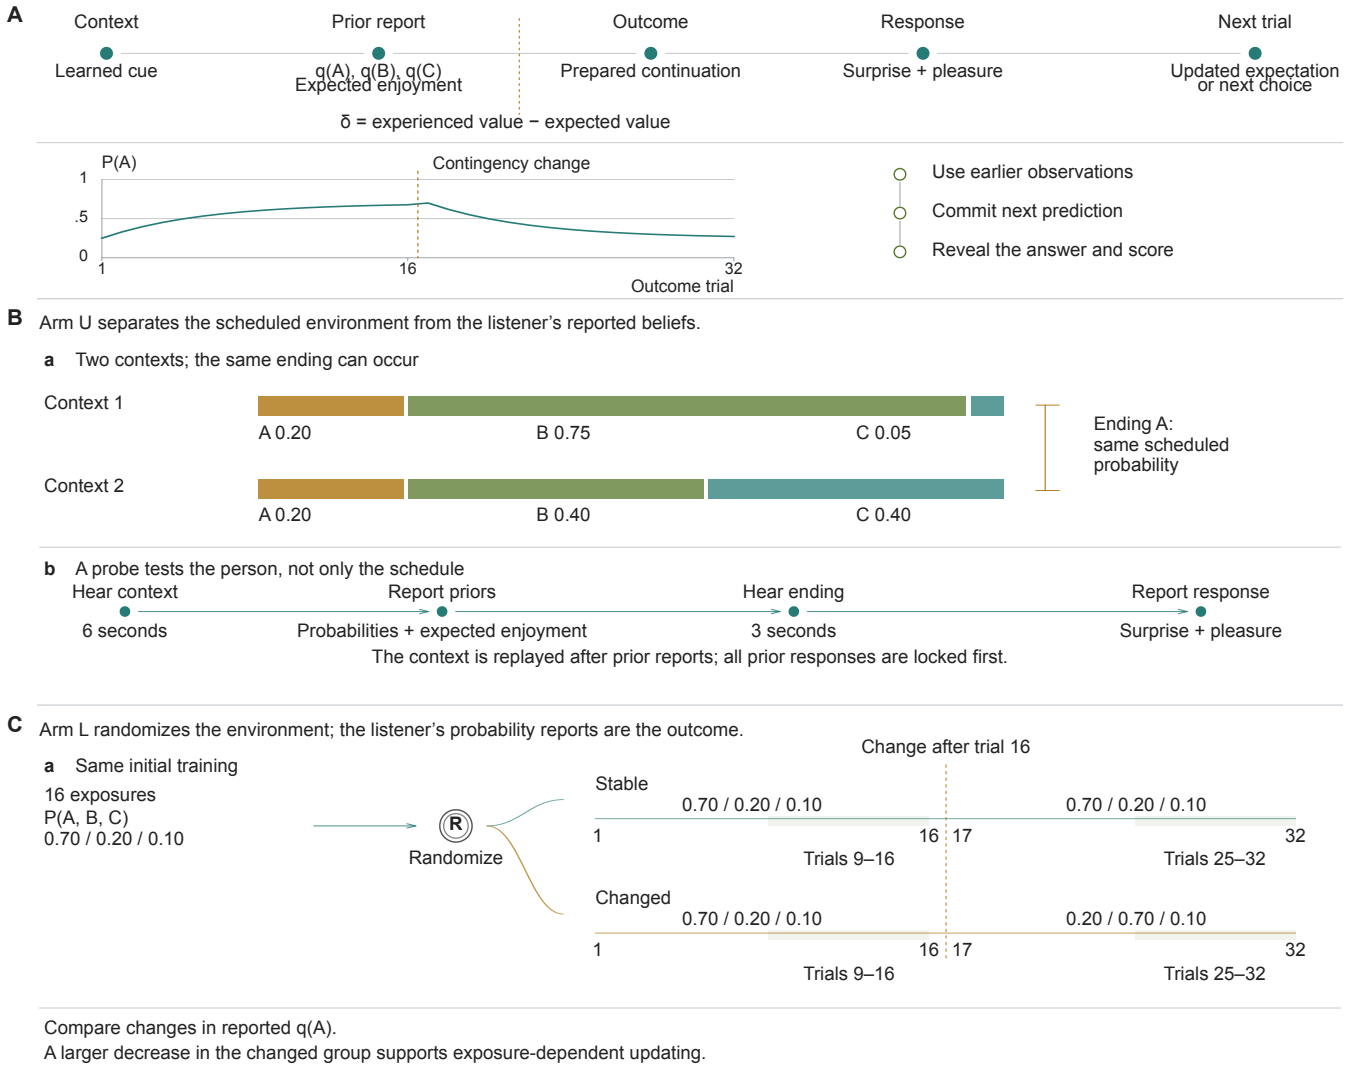


Fig. 3. Controlled tasks **A,** Prior reports are locked before the outcome, with experience reported afterwards. The adaptation curve is an analytic illustration. **B,** Arm U distinguishes scheduled probabilities from reported beliefs; the target ending has scheduled probability 0.20 in both contexts. **C,** Arm L compares stable and changed distributions after probe 16. Probabilities show the proposed schedules rather than observed learning. Offline replay retains the listening sequence used during collection.

Zero probability reports will be retained in the source data. For logarithmic scoring only, probabilities will be floored at .01 and renormalized; .005 and .02 floors will provide sensitivity analyses. The pilot will assess probability-allocation comprehension, entropy–surprise coverage and floor/ceiling behaviour before waveform joins and timing are fixed. Although the target has scheduled probability .20 in both contexts, reported surprise may differ. Subjective entropy and surprise are therefore measured variables, not independently randomized factors.

The primary model will predict 0–100 pleasure from listener-centred and scaled H, S, HS, H^2, S^2, context, order, reference familiarity and acoustic controls. Participant and continuation-family intercepts will account for clustering, with participant H/S slopes where identifiable. If the fit is singular, random-effect correlations will be removed first, followed by the least-supported slope; both intercepts will remain.

The directional prediction is a negative interaction coefficient, accompanied by a positive surprise slope at the pooled 25th percentile and a negative slope at the 75th percentile of centred pre-event entropy. These percentiles will be defined without outcome data. Joint uncertainty will be reported because an interaction alone does not establish both slope signs. A bounded-outcome model will assess scale sensitivity. Since H and S come from the same probability report, their errors may be correlated. Pilot and recovery analyses will examine reliability, occupied support, scoring floors and nonlinear-main-effect alternatives. Quantiles, lags and polynomial order will remain fixed during confirmation.

Arm L: exposure-dependent updating Sixteen training exposures from (.70, .20, .10) precede 32 probes. After probe 16, half the participants, assigned at random, will encounter (.20, .70, .10); the remainder will continue with the original distribution. Outcome identities will be independently counterbalanced. For the initially dominant outcome A, $\Delta\bar{q}_i(A)$ compares probes 25–32 with probes 9–16. The primary assignment effect is $\tau_L = E[\Delta\bar{q}_i(A)|Z_i = 1] - E[\Delta\bar{q}_i(A)|Z_i = 0]$.

A condition-by-phase model will include participant intercepts, baseline probability, counterbalanced identity and realized-exposure imbalance. Post-assignment pleasure will not be treated as a pretreatment covariate. A larger decrease in reported probability under changed exposure would support updating, whether or not pleasure changes. Analyses will report completion by condition, baseline predictors of missingness and all-available-data sensitivities. Next-outcome log score, adaptation rate and signed value learning are secondary. Short sequences warrant hierarchical uncertainty estimates rather than individual learning-rate rankings.

Personal prediction and stimulus selection

Preparation and replay Each versioned audio representation will be computed once. Global mappings, dimension reductions, module selection, learner priors, update rules and tuning grids will be frozen before confirmatory recruitment. Personal preparation uses the 20-reference intake and declared background to save schedules, fixed predictions, policy assignments, seeds and content hashes. A prepared session requires no live MI, GPU or LLM inference.

Personal updating will be evaluated by chronological offline replay. At each transfer trial, the learner can use intake, the completed mechanistic task and earlier transfer responses. Its state and predic-

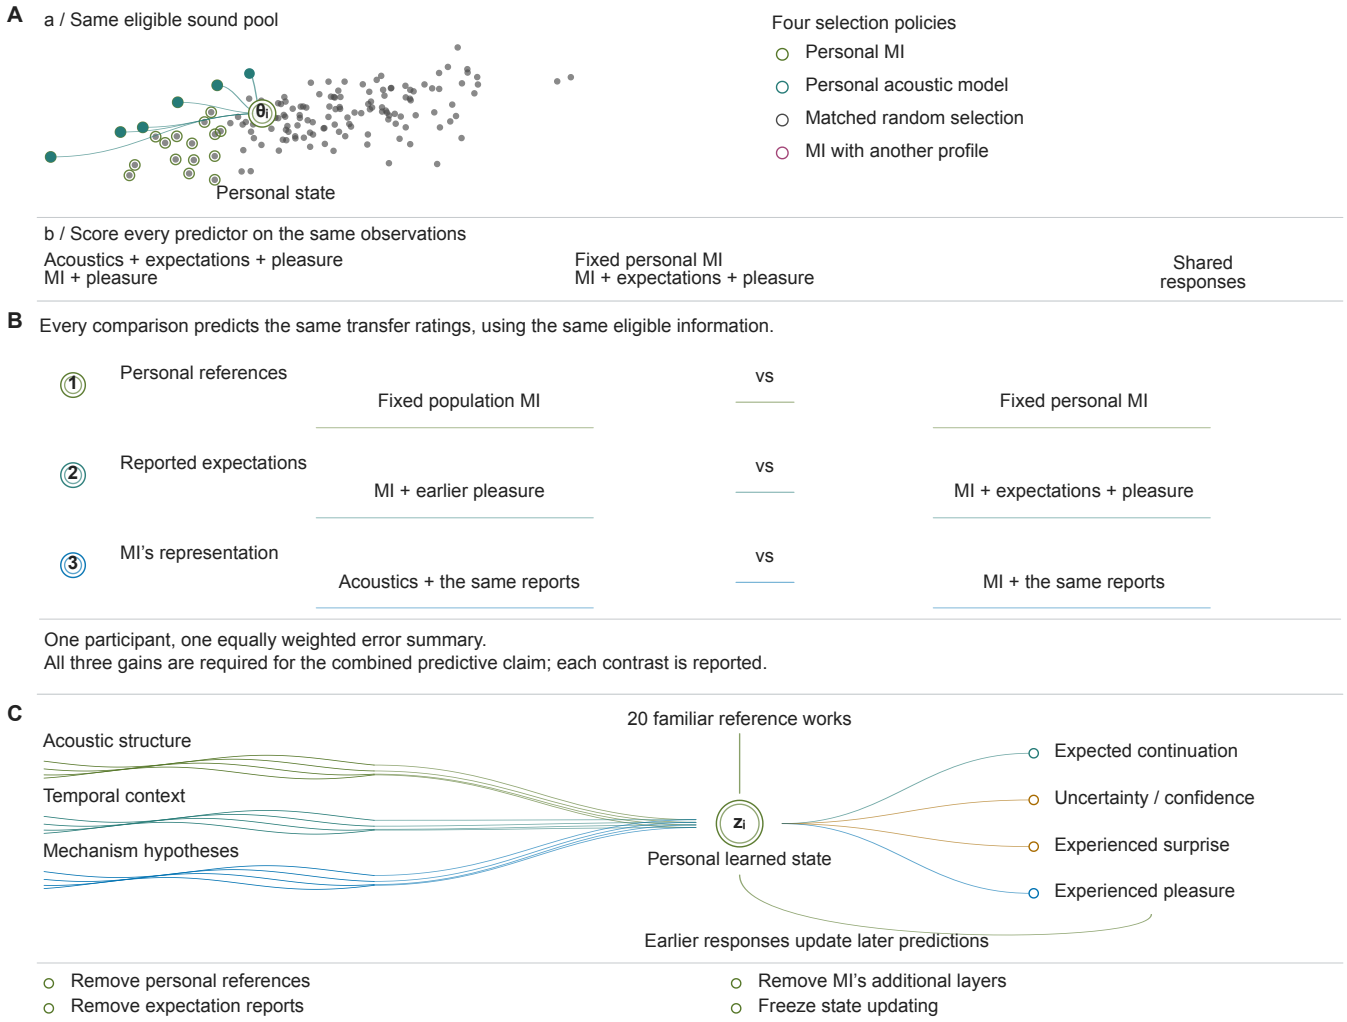


Fig. 4. Personal prediction **A,** Four selection policies draw from the same eligible pool. Display coordinates are synthetic, and all predictors are evaluated on the same observations. **B,** The composite predictive claim requires improvement in all three matched comparisons; each estimate is also reported. **C,** References and earlier responses inform personal state. Ablations remove history, representation layers, expectation reports or updating. The connections represent computations, not neural pathways.

tion are saved before the current response is revealed. This permits a test of learning about the listener without changing what that listener heard. Predictions for known recordings may use the full prepared audio; an account of time-local anticipation would require a stricter sound-time boundary. Predictions committed before collection and predictions generated in replay will remain distinguishable.

Model specification Let x_j^A, x_j^T, x_j^M denote acoustic, temporal and mechanism features, B_i intake, z_{it} personal state and o_{it} observed outcomes. The model has the form $z_{i0} = G(B_i; \theta)$, $z_{i,t+1} = U(z_{it}, o_{it}; \theta)$ and $\widehat{o}_{it} = F(x_j^A, x_j^T, x_j^M, z_{it}; \theta)$. These expressions define the proposed information flow. The functions, priors and global parameters θ will be instantiated and frozen at pilot exit; uncollected outcomes will be masked.

The primary linear readout uses six standardized coordinates: two pilot-catalogue principal components from each feature group. The acoustic comparator uses six acoustic components. A population ridge model supplies coefficient priors, and personal fitting minimizes reference squared error with shrinkage. The penalty will be selected from $\{.1, 1, 10, 100\}$ by participant- and composition-grouped pilot validation. Transforms, priors, coefficients and software hashes will be retained. Nonlinear and richer state-space learners are secondary and will receive matched tuning budgets.

The pleasure-only learner updates cumulatively after each eligible earlier rating. An expectation-informed version adds a global pilot-fitted adjustment based on earlier mean entropy and surprise, expected enjoyment, confidence calibration and changes in continuation accuracy. Mechanistic-task summaries initialize transfer. Later surprise and pleasure reports update only quantities that are measured again; they do not supply unreported continuation probabilities or additional personal coefficients. Acoustic comparators receive the same reports, background, timing and fitting resources. Replay and parameter-recovery checks will precede confirmation. Shuffled personal states and permuted expectation reports will test whether any improvement depends on the correct listener's information.

Primary predictive comparison On identical eligible trials $\mathcal{T}_i$, define $L_i(m) = |\mathcal{T}_i|^{-1} \sum_{t \in \mathcal{T}_i} |Y_{it} - \widehat{Y}_{it}^{(m)}|$ and $\Delta_{a,b} = E_i[L_i(a) - L_i(b)]$. Positive differences favour the enriched model, with equal weight for each listener. The composite claim requires three improvements: fixed personal MI over fixed population MI; expectation-plus-pleasure MI over pleasure-only MI; and expectation-informed MI over an equally updated acoustic learner. The intersection–union p value is the maximum of the three one-sided p values. Each estimate and interval will also be reported. Holm correction at $\alpha = .05$ covers this composite and the U and L tests. All predictions will use the same frozen clipping rule at 0–100, with personal-mean/intercept and repetition/order controls retained.

Selection-policy comparison The eligible pool will be randomly partitioned into disjoint candidate sets before ranking, with membership rotated across participants. Concealed policy labels will be interleaved in balanced six-trial blocks, yielding six excerpts per policy. Random selection will match instrumentation, duration and broad familiarity rather than MI score. MI and acoustic policies will share the profile and ranking budget. Swapped profiles will be matched on training and catalogue coverage without disclosing identities. A session will not be released if a quota cannot be filled.

The comparison estimates selection-policy value within the released pool, rather than a module's causal contribution with audio held fixed. Generalization to another catalogue distribution requires known selection probabilities and adequate support. Policy value, predictive accuracy and mechanism evidence will be reported separately.

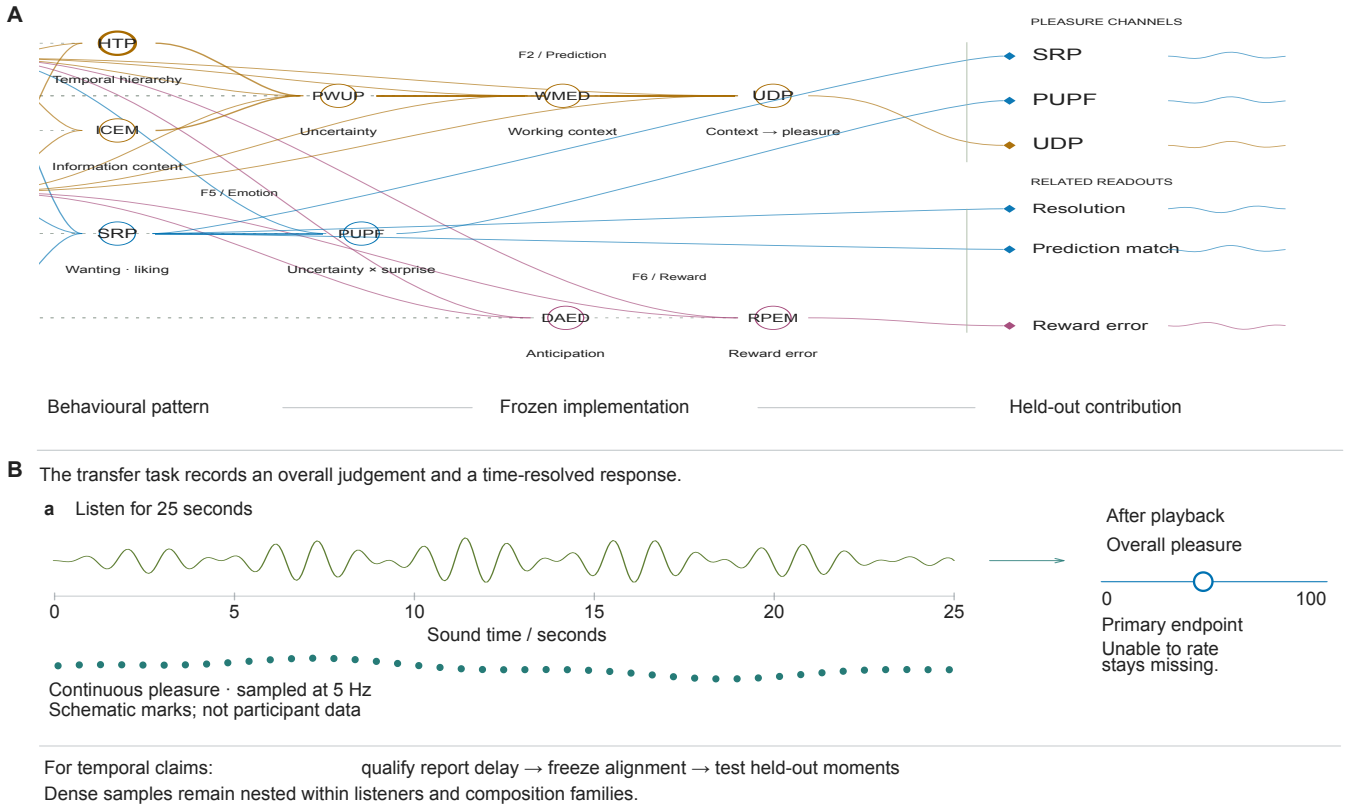


Fig. 5. Mechanisms and measurements A, The nine-module topology retains the original function colours, six dependencies within the selected set, incoming acoustic and temporal inputs, and outgoing readouts. No additional upstream mechanisms are assigned to RPEM or DAED. **B,** A 5-Hz trace is recorded during each 25-second excerpt; the immediate overall rating is primary. Traces and slider positions are schematic. Samples remain nested within listeners and composition families.

Mechanism comparisons and secondary outcomes

Nine computations from the executable MI architecture[9] will be evaluated with their definitions fixed. Their names identify implementations; construct validity depends on the measurements and comparisons below. These analyses are secondary or exploratory.

HTP and ICEM Hierarchical Temporal Prediction is motivated by auditory integration and learned sequence prediction[10, 11]. Its stability, pitch, timbre and onset inputs nevertheless differ from a listener's reported probability distribution. Information Content Emotion Model measures acoustic variation rather than $-\log_2 q(e)$. Arm U will test whether these outputs correspond to elicited expectation and surprise beyond acoustic change. The sampled arousal reports permit a behavioural comparison, without providing cardiac or skin-conductance evidence.

PWUP and UDP Precision-Weighted Uncertainty Processing scales raw error as $E(1 - .5g)$ using gate g. Uncertainty-dependent neural deviance responses[12] motivate testing this gate, but do not establish its empirical direction. Uncertainty-Driven Pleasure contains an upstream uncertainty/confirmation branch and a separate acoustic pleasure index; each increment will be tested separately. Accounts of pleasure in uncertain music[13] provide a rationale for these tests rather than validation of either computation.

SRP and PUPF Striatal Reward Prediction supplies audio-derived wanting, liking and tension proxies. Evidence on musical tension[14] motivates separating tension, relief and liking in interpretation. Pleasure-Uncertainty Prediction Function combines acoustic proxies and pleasantness; its frozen equation is not a fitted reproduction of an empirical uncertainty–surprise surface. Disagreement with behaviour will be reported as a limitation of that readout.

DAED, RPEM and WMED Dopamine Anticipation-Experience Dissociation and Reward Prediction Error in Music are acoustic computations. Centred RPEM differs from reported $Y - V$ and does not itself update a personal learner. Working Memory-Entrainment Dissociation draws on evidence separating timing and memory[15, 16]. The core task measures neither working-memory capacity nor neural entrainment, so these outputs will be compared only with the observed behavioural targets.

Joint models and ablations A secondary explanatory model will combine probability reports, unexpectedness and pleasure with participant and composition effects, and performance nested within composition. Probability allocations will use a logistic-normal model and affect ratings bounded models. Expected enjoyment, surprise and signed value error will retain separate definitions; H/S derived from one report are not independent observations. Predictors include acoustic controls, temporal/mechanism features, response dependence and prespecified moderation by training and familiarity.

Equal-capacity, equal-budget ablations will compare acoustic, acoustic-plus-temporal and full representations. Further comparisons will remove personal calibration, freeze updating, or omit probability and expected-value reports while retaining pleasure. Correlated outputs may provide overlapping or incremental information without permitting unique causal attribution. A reference-dose analysis will use 0, 5, 10 and 20 families under ten preregistered subset orderings, evaluated on the same excerpts, with missing ratings and effective reference counts reported.

Value learning and additional assays Secondary analyses will compare $V_{t+1} = (1 - \alpha)V_t + \alpha Y_t$ with fixed-value, repetition, unsigned-surprise and experienced-reward alternatives. The target is next-trial enjoyment or choice under identical prior information. Directly correlating $Y_t - V_t$ with $V_{t+1} - V_t$ would introduce algebraic coupling; learning rates will instead be regularized, with hyperparameters chosen in the pilot.

Separately consented extensions could compare resolved, delayed and withheld endings, measuring pre-ending wanting/enjoyment and post-ending relief/liking. Other extensions would separate anticipation and experience response blocks, or cross rhythmic regularity with two-/four-item pattern memory while measuring tapping error and recognition separately. Until these assays are conducted, interpretations of SRP, DAED and WMED remain limited to the core measurements.

During transfer, continuous pleasure will be sampled at 5 Hz. Overall pleasure (0–100, with unable-to-rate distinct), unexpectedness and recognition will follow each excerpt. The overall rating, rather than the trace maximum, is the primary endpoint. Candidate pleasurable moments will be specified from prepared signals

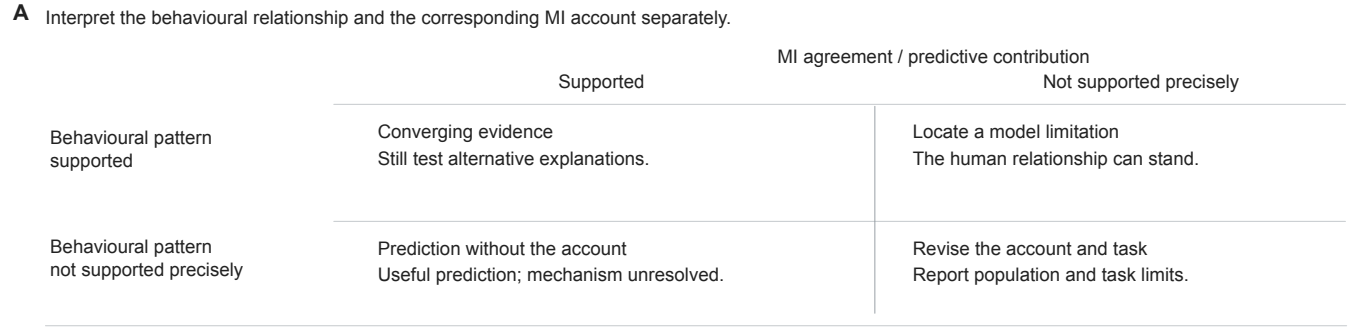


Fig. 6. Possible interpretations Behavioural support and the corresponding MI prediction are evaluated separately. Agreement still requires comparison with alternatives. Behavioural support with MI failure identifies a model limitation; prediction without the proposed behavioural pattern leaves the explanation unresolved. Precise incompatibility can narrow the operational account, whereas wide intervals remain inconclusive. The four outcomes are hypothetical.

and compared with matched moments within and across excerpts. Reporting lag and smoothing will be fixed after pilot qualification, not optimized on test traces. Audio, browser-monotonic and server-receipt clocks will be recorded separately, with latency and variability assessed on representative hardware[20].

Analysis and interpretation

Uncertainty and multiplicity Paired predictive differences will use a crossed participant-by-composition bootstrap with 5,000 replicates, retaining family assignments across occurrences. A predeclared cluster wild-bootstrap sensitivity will address sparse work coverage; simulations will assess error control for the resampling procedure. Mean squared error, calibration slope and signed bias are secondary outcomes. Policy pleasure will be analysed with participant/work effects, order and prespecified familiarity controls. False-discovery-rate correction will apply within secondary mechanism, policy and temporal families. Exploratory findings will require independent confirmation; neither dense traces nor repeated works increase the number of independent listeners.

Recovery and power Simulations will represent bounded ratings, participant/composition dependence, incomplete baselines, actual allocation, clustering, attrition and differential dropout. Null scenarios include acoustic-only pleasure, scale differences, nonlinear main effects without an interaction, and calibrated expectations without a hedonic contribution. Positive controls will span relevant effects and nuisance variation. Each analysis must recover its intended target without falsely recovering another[25]. Inputs, seeds, Monte Carlo uncertainty and failed scenarios will be retained. Individual coefficients will not be interpreted as psychological traits without recovery and independent reliability evidence.

Pilot estimates will inform intake and comprehension feasibility, H/S coverage, variance, dependence, missingness, balance and timing. The target is at least 90% power at the most stringent Holm threshold, with Type I error assessed under null and misspecified scenarios. Initial meaningful-effect thresholds are a standardized interaction magnitude of .10, a changed-versus-stable probability difference of .10, and a one-point reduction in absolute prediction error on the 0–100 scale for each predictive contrast. These are planning choices to be justified or amended using reliability and substantive interpretation before confirmation. They are not literature effect estimates. Simulations will examine 300, 500, 700 and 900 analysable participants. If attainable precision is inadequate, scope or recruitment will be revised and documented in the pilot-exit amendment. The registration cap does not establish achieved power, and nonsignificance will not be interpreted as equivalence.

Missing observations and exclusions Consent, duplicate identity, technical delivery and task-quality rules will be specified before analysis. Low or constant pleasure and disagreement with MI will not be exclusion criteria. Primary prediction requires at least 18/24 valid transfer ratings; arm summaries require at least 75% valid probes and observations in both phases or contexts. All-valid-observation sensitivities and intake predictors of completion will be reported.

Invalid trials include failed audio start, stalls, hidden pages during the target period and outcomes inconsistent with the released schedule. Late upload is acceptable if local timing and verified receipt remain intact. Replayed audio is a new exposure, excluded from primary first-exposure analyses but retained for sensitivity analyses; it cannot replace the original response. Primary outcomes will not be imputed or forward-filled. Any baseline-covariate imputation will use pilot-frozen rules and missingness indicators.

Interpretation Behavioural learning and computational personalization are distinct aims. An uncertainty–surprise pattern without an MI increment would limit the representation or readout. Improved prediction without that pattern would support predictive utility while leaving the proposed explanation unresolved. Precise incompatibility with both would challenge this operational account; wide intervals would remain inconclusive. Failure to transfer from three-option tasks to natural music is itself a substantive result.

Familiarity, acoustic salience, reporting reactivity and response style remain alternatives despite the planned controls. The 20-family intake requirement constrains the population, and broader styles or cultures require new catalogue coverage and task validation. Behavioural ratings cannot identify neural mediators, neurotransmitters, autobiography or general musical ability. A failure in this task would not, by itself, refute predictive-processing theory.

Ethics and availability Consent will describe the tasks, burden, stored information, separation of contacts from research records, withdrawal and limits after irreversible aggregation. Pseudonymous histories remain personal data. Approved code, permitted catalogue records, synthetic fixtures and sufficiently de-identified summaries will be released; contact details, consent records and restricted audio will remain controlled. Access conditions will be stated explicitly, consistent with FAIR principles[21]. Reporting will include participant flow, primary tests, policy comparisons, null results and disagreements with the implementation.

Study status and disclosures This protocol is being prepared for Stage 1 review; it is not a completed preregistration or an accepted Registered Report[17, 26]. Recruitment requires an executable analysis, justified power calculation, institutional responsibility, consent, compensation and approvals. The companion measurement and inference papers describe supporting methods and do not provide corroborating results. All figures are conceptual or analytic illustrations. Editable sources accompany this edition; participant records, restricted audio and deployment credentials are excluded.

Amaç Erdem is MI's originator and the proposed sole author. Affiliation, funding, competing interests and final contributor statements remain to be confirmed before submission. Generative-AI assistance was used for editorial work, literature retrieval, figure code and document production. Scientific content and the disclosure submitted to a journal require the author's approval.

References

[1] Vuust, P., Heggli, O. A., Friston, K. J. & Kringelbach, M. L. Music in the brain. *Nature Reviews Neuroscience* **23**, 287–305 (2022). doi:10.1038/s41583-022-00578-5.

[2] Cheung, V. K. M. et al. Uncertainty and surprise jointly predict musical pleasure and amygdala, hippocampus, and auditory cortex activity. *Current Biology* **29**, 4084–4092.e4 (2019). doi:10.1016/j.cub.2019.09.067.

[3] Gold, B. P. et al. Predictability and uncertainty in the pleasure of music: a reward for learning? *Journal of Neuroscience* **39**, 9397–9409 (2019). doi:10.1523/JNEUROSCI.0428-19.2019.

[4] Egermann, H., Pearce, M. T., Wiggins, G. A. & McAdams, S. Probabilistic models of expectation violation predict psychophysiological emotional responses to live concert music. *Cognitive, Affective, & Behavioral Neuroscience* **13**, 533–553 (2013). doi:10.3758/s13415-013-0161-y.

[5] Salimpoor, V. N. et al. Anatomically distinct dopamine release during anticipation and experience of peak emotion to music. *Nature Neuroscience* **14**, 257–262 (2011). doi:10.1038/nn.2726.

[6] Salimpoor, V. N. et al. Interactions between the nucleus accumbens and auditory cortices predict music reward value. *Science* **340**, 216–219 (2013). doi:10.1126/science.1231059.

[7] Gold, B. P. et al. Musical reward prediction errors engage the nucleus accumbens and motivate learning. *Proceedings of the National Academy of Sciences USA* **116**, 3310–3315 (2019). doi:10.1073/pnas.1809855116.

[8] Gold, B. P. et al. Auditory and reward structures reflect the pleasure of musical expectancies during naturalistic listening. *Frontiers in Neuroscience* **17**, 1209398 (2023). doi:10.3389/fnins.2023.1209398.

[9] Erdem, A. Musical Intelligence: A Glass-Box Architecture for Computational Models of Music Listening. Author manuscript, 8 October 2026. Dated architecture record.

[10] Norman-Haignere, S. V. et al. Multiscale temporal integration organizes hierarchical computation in human auditory cortex. *Nature Human Behaviour* **6**, 455–469 (2022). doi:10.1038/s41562-021-01261-y.

[11] Bonetti, L. et al. Spatiotemporal brain hierarchies of auditory memory recognition and predictive coding. *Nature Communications* **15**, 4313 (2024). doi:10.1038/s41467-024-48302-4.

[12] Quiroga-Martinez, D. R. et al. Reduced prediction error responses in high- as compared to low-uncertainty musical contexts. *Cortex* **120**, 181–200 (2019). doi:10.1016/j.cortex.2019.06.010.

[13] Mencke, I., Omigie, D., Wald-Fuhrmann, M. & Brattico, E. Atonal music: can uncertainty lead to pleasure? *Frontiers in Neuroscience* **12**, 979 (2019). doi:10.3389/fnins.2018.00979.

[14] Lehne, M., Rohrmeier, M. & Koelsch, S. Tension-related activity in the orbitofrontal cortex and amygdala: an fMRI study with music. *Social Cognitive and Affective Neuroscience* **9**, 1515–1523 (2014; online 2013). doi:10.1093/scan/nst141.

[15] Noboa, M. de L., Kertész, C. & Honbolygó, F. Neural entrainment to the beat and working memory predict sensorimotor synchronization skills. *Scientific Reports* **15**, 10466 (2025). doi:10.1038/s41598-025-93948-9.

[16] Yuan, Y. et al. Decoding auditory working memory load from EEG alpha oscillations. *Psychophysiology* **62**, e70210 (2025). doi:10.1111/psyp.70210.

[17] Nature Human Behaviour. Registered Reports: submission guidelines. Publisher guidance, accessed 10 October 2026.

[18] Hawthorne, C. et al. Enabling factorized piano music modeling and generation with the MAESTRO dataset. *ICLR* (2019). Dataset, versions and licence.

[19] Woods, K. J. P., Siegel, M. H., Traer, J. & McDermott, J. H. Headphone screening to facilitate web-based auditory experiments. *Attention, Perception, & Psychophysics* **79**, 2064–2072 (2017). doi:10.3758/s13414-017-1361-2.

[20] Bridges, D., Pitiot, A., MacAskill, M. R. & Peirce, J. W. The timing mega-study: comparing a range of experiment generators, both lab-based and online. *PeerJ* **8**, e9414 (2020). doi:10.7717/peerj.9414.

[21] Wilkinson, M. D. et al. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* **3**, 160018 (2016). doi:10.1038/sdata.2016.18.

[22] Mas-Herrero, E. & Marco-Pallarés, J. Predictive processes shape individual musical preferences. *Proceedings of the National Academy of Sciences USA* **122**, e2500494122 (2025). doi:10.1073/pnas.2500494122.

[23] Gu, Y., Shao, C., Li, J. & Fan, Y. Interpretable deep learning reveals distinct spectral and temporal drivers of perceived musical emotion. *Scientific Reports* **16**, 4205 (2026). doi:10.1038/s41598-025-34238-2.

[24] Albury, A. W., Bianco, R., Johnson, A. P. & Penhune, V. B. Stimulus predictability and liking enhance auditory–motor encoding and memory for melodies. *Annals of the New York Academy of Sciences* **1560**, e70310 (2026). doi:10.1111/nyas.70310.

[25] Wilson, R. C. & Collins, A. G. E. Ten simple rules for the computational modeling of behavioral data. *eLife* **8**, e49547 (2019). doi:10.7554/eLife.49547.

[26] Music Perception. Registered Reports: guidelines for authors. Official author guidelines (accessed 10 October 2026).