EXPERIMENTAL PROGRAMME · PROSPECTIVE PROTOCOL

Learning to predict musical pleasure

Amaç Erdem

Independent researcher · Correspondence: amace@bu.edu · ORCID 0009-0000-1711-0036

Can a model learn from the music a person selects and use that knowledge to predict how later music will feel? This protocol combines Musical Intelligence (MI)'s multidimensional representation with personal calibration, explicit expectation measurements and a scalable online experiment. Its central contribution is to test the predictive difference created by learning a listener: how much selected reference music, temporal context, mechanism-related features and subsequent experience each add to prediction of pleasure. A programme capped at 1,000 registrations, including a proposed 60-person pilot, begins with 20 familiar classical works selected by each participant. Controlled tasks then test uncertainty-dependent pleasure and exposure-dependent updating, followed by a common transfer block comparing personal MI, acoustic, random and mismatched-profile selection. A joint, multilevel analysis links expectation, uncertainty, surprise and pleasure while accounting for listeners, works and performances. Precomputed audio representations and frozen schedules permit reliable online collection; chronological replay evaluates learning using only earlier responses. Prespecified ablations distinguish the value of personal references, MI's additional dimensions and expectation-informed updating. Literature-based mechanisms are primary scientific targets alongside this predictive contribution. The outcome is a falsifiable research programme for explaining both common relationships and individual differences, rather than a claim that accurate recommendations alone validate a cognitive mechanism.

Design proposal, not collected results or a completed preregistration. Trial counts and launch criteria are specified below; ethics review, catalogue clearance and pilot qualification remain prerequisites.

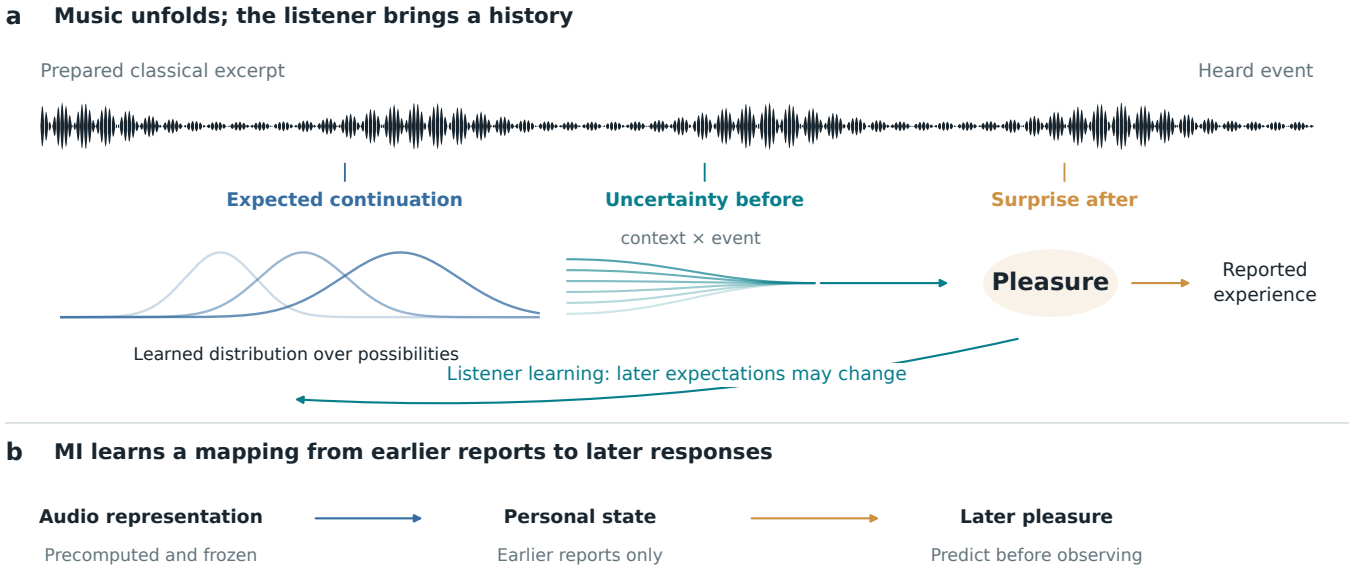


Fig. 1 | Pleasure is the outcome; expectation and learning supply testable explanations. a, A listener approaches an event with a distribution over possible continuations. Uncertainty precedes the event; surprise depends on the continuation that occurs. Pleasure is measured independently. **b,** A computational learner uses earlier observations to predict later responses. The upper loop concerns change in the listener; the lower loop concerns change in a predictor. Curves and waveforms are original schematic artwork, not recordings of participant responses.

The question is why this sound pleases this listener

A familiar melody can remain moving when its ending is known. An unfamiliar continuation can be delightful, indifferent or unwelcome. A useful account of musical pleasure must explain both facts without assigning a fixed reward to either predictability or surprise. Predictive accounts place the listener's history between the acoustic signal and its interpretation: experience changes what is expected, how uncertain that expectation is and what a deviation means[1]. The present programme turns that account into measurements whose failures would be informative.

Two aims have equal standing. The first tests mechanisms motivated by the literature, whether or not MI predicts them. The second quantifies what changes when MI learns from the listener's selected music and then from their responses: the improvement over a population model, an acoustic-only personal model and a model that never updates. These questions belong to one process spanning sound, personal history, expectation and reaction. Classical music supplies familiar reference material, unfamiliar works and multiple performances, allowing different layers of that process to be compared within a common online dataset.

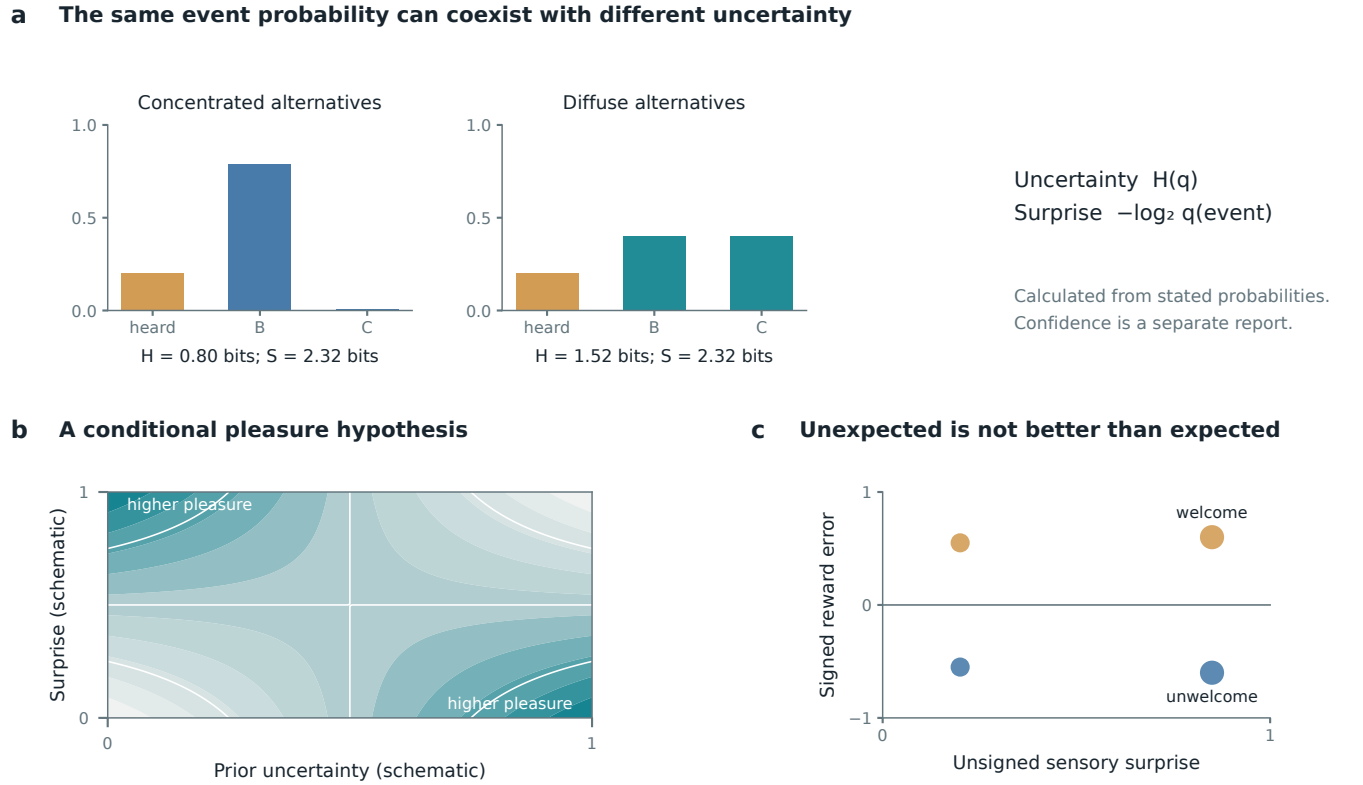


Fig. 2 | Uncertainty, surprise and value error are different quantities. a, Two exact three-outcome distributions assign the heard event probability 0.20, giving identical surprise of 2.32 bits but different entropy. **b,** A qualitative interaction motivated by Cheung et al.[2]: pleasurable surprise in a relatively certain context and pleasurable confirmation in a relatively uncertain context. The surface illustrates a hypothesis; it is neither a fit to their data nor MI's PUPF equation. The axes are normalized conceptual coordinates, not an exhaustive map of feasible probability distributions. **c,** High sensory surprise can accompany positive or negative reward error. Marker positions are illustrative.

What the literature does, and does not, predict

Cheung and colleagues reported that musical pleasure depended jointly on modelled uncertainty and surprise: high pleasure occurred for surprising chords under low uncertainty and unsurprising chords under high uncertainty[2]. This interaction motivates a conditional hypothesis, not a universal preference for violation. Gold and colleagues found preferences related to intermediate predictability and contextual information measures[3]. Their findings also require comparison with nonlinear main effects: an apparent interaction can otherwise inherit the shape of an omitted complexity preference. Whether these patterns transfer from the original materials to individual classical listeners remains an empirical question.

Unexpectedness is also associated with arousal[4]. Arousal, however, does not determine whether an event is liked. Evidence linking musical reward to anticipation, peak experience and the value of unfamiliar music further motivates separating expected enjoyment from experienced pleasure[5, 6]. A listener may strongly want to hear a resolution yet find the resolution less satisfying than expected. One slider cannot establish these distinctions.

Learning introduces a second time scale. A surprising outcome may change expectations about the next event. In a musical choice task, Gold and colleagues related reward prediction errors to subsequent behaviour and neural signals[7]. That result motivates testing value updating here; it does not establish a dopamine mechanism in a browser experiment. Naturalistic expectancy–liking results provide complementary context, under different measurement and task constraints[8].

An explicit language for expectation

For listener i on trial t, let $q_{it}(k)$ be the reported probability of continuation k among three audible alternatives. The entropy of this distribution and the surprise of the realized continuation e_t are

$$H_{it} = -\sum_{k=1}^{3} q_{it}(k) \log_2 q_{it}(k), \tag{1}$$

$$S_{it} = -\log_2 q_{it}(e_t). \tag{2}$$

These quantities describe the elicited three-option task, not the full musical expectations of a person. Subjective confidence is recorded separately. A confident choice alone is insufficient to identify the distribution or its entropy.

Expected enjoyment V_{it} and experienced pleasure Y_{it} use the same 0–100 anchors. A candidate signed reward error is $\delta_{it} = Y_{it} - V_{it}$. The model's prediction residual, $r_{it} = Y_{it} - \widehat{Y}_{it}$, is a third quantity. Neither residual is interchangeable with sensory surprise. Reduced entropy alone is also insufficient evidence of learning: it may indicate increasing overconfidence. Learning must appear in subsequent expectations, choices or predictive accuracy.

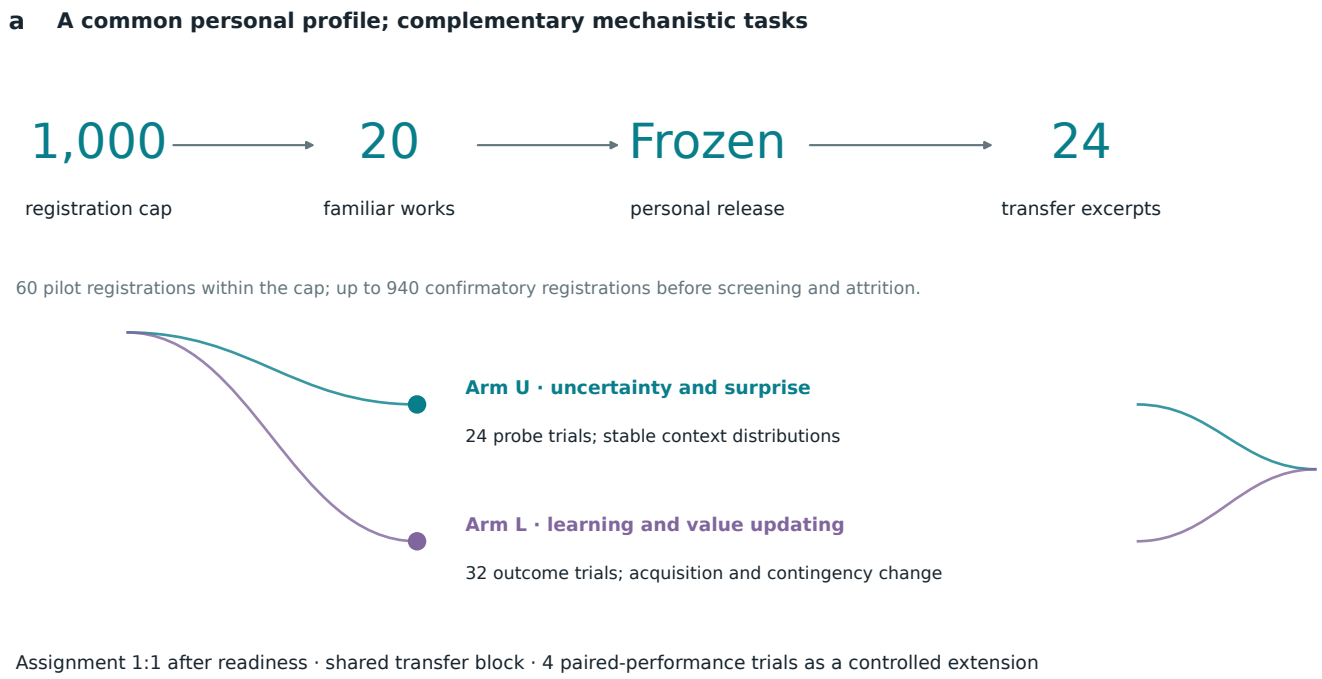


Fig. 3 | A personal reference profile connects controlled tasks to classical transfer. The cap is 1,000 verified eligible registrations, not 1,000 completed datasets. The proposed pilot uses the first 60 registrations; its participants are excluded from confirmatory estimates. Ready confirmatory participants are assigned equally to uncertainty (U) or learning (L) tasks, then complete the same 24-excerpt transfer block. Four paired-performance trials retain the earlier within-work design as a secondary control. Counts are proposed allocations, not observed recruitment or a power calculation.

A design with two kinds of generalization

The intake asks which music the participant already knows. A curated catalogue presents searchable composition families with a short reference preview. The participant confirms prior recognition and rates the reference performance for exactly 20 distinct families. Recognition and liking are separate responses: an intensely disliked but familiar work still belongs in the profile. The task does not reconstruct general listening history, and it does not require a streaming account.

Twenty familiar works anchor person-specific calibration, but their availability cannot be assumed for every registrant. Catalogue search burden, the proportion reaching 20 and differences by musical background are pilot outcomes. People unable to complete this requirement remain visible in recruitment accounting. Replacing unfamiliarity with repeated onboarding exposure would change the construct and is not allowed.

The controlled tasks examine expectation under repeatable exposure. Arm U studies how the relationship between surprise and pleasure changes with prior uncertainty. Arm L studies whether expectations respond to a change in the environment rather than to repetition alone. Both arms use short continuations derived from classical motifs, using independently cleared or newly rendered materials, with identical target endings available in multiple contexts. They therefore provide stronger control over event probabilities than whole, familiar recordings can provide.

The transfer block then asks whether the account remains useful in more natural listening. Each participant hears 24 previously unpresented, 25-second excerpts from the eligible classical pool. Six are selected by personal MI, six by an equally calibrated acoustic model, six by matched random sampling and six by MI using another participant's reference profile. Selection labels are concealed. All excerpts come from the same eligibility rules; the comparison concerns policies, not a privileged MI catalogue versus an unrestricted random catalogue.

Four additional pairs compare performances of passages from four intake works. The reference performance heard during registration is not reused; passage, instrumentation and approximate segment position are matched. This comparison holds composition more nearly constant and tests whether performance-sensitive predictions survive that control. It is retained as a valuable extension rather than allowed to replace the broader question about pleasure and prediction.

What would count as a contribution from MI?

MI contributes explicit candidate representations, declared dependencies and testable mappings[9]. Its module names are not measurements of the corresponding human processes. The strongest result would combine three forms of evidence: reliable behavioural relationships, model quantities that track those relationships and improved prediction of later pleasure beyond equally informed alternatives. A successful recommendation policy alone meets only the last question's practical selection component. Conversely, a replicated uncertainty–surprise interaction with poor MI agreement would support the scientific relationship while identifying an implementation failure.

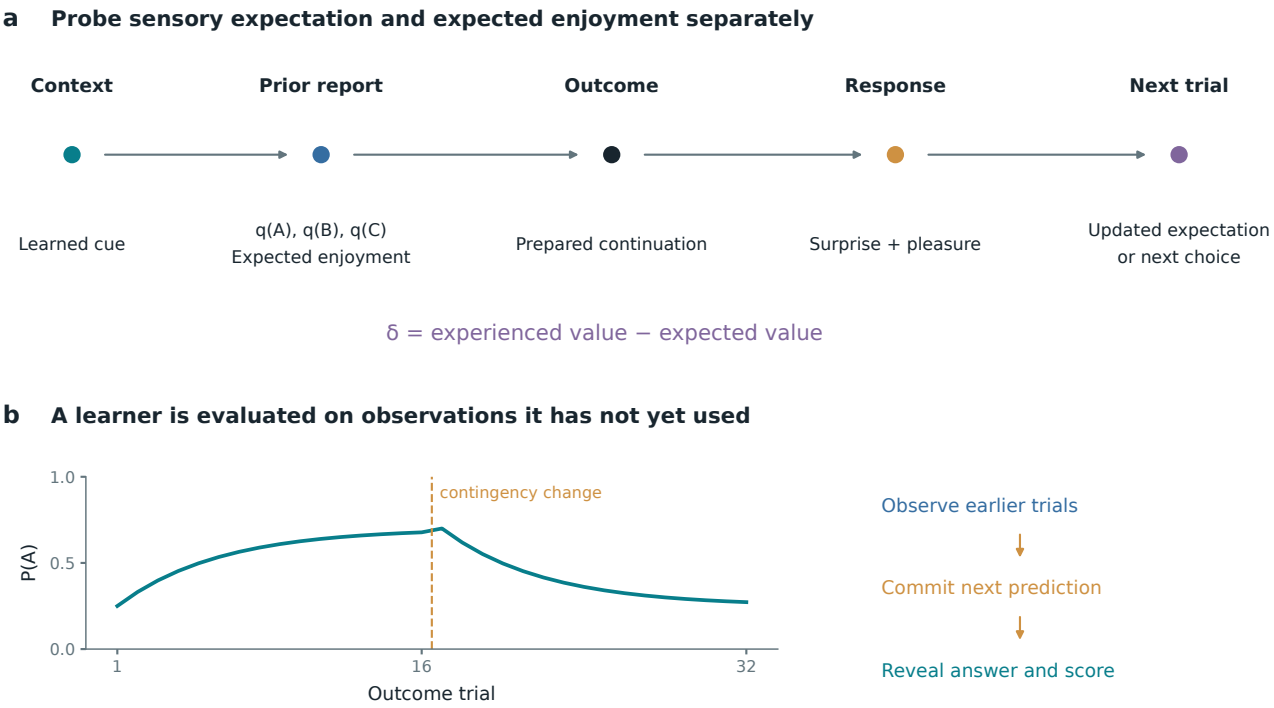


Fig. 4 | Learning must be measured across an information boundary. a, Sensory continuation probabilities and expected enjoyment are elicited before the outcome; surprise and pleasure follow it. Later reports test updating. **b,** The curve is an analytic illustration of adaptation after a contingency change, not an estimated learning curve. Offline replay commits each prediction before revealing that trial's answer. The audio schedule remains fixed throughout the online session; replay does not retroactively make the experiment adaptive.

Three confirmatory questions

The first confirmatory question concerns the uncertainty–surprise interaction. Arm U provides two learned context distributions over three continuations, with outcome identities permuted across participants. The primary model includes listener-centred entropy, surprise, their interaction, quadratic main effects, acoustic covariates and participant and stimulus effects. The predicted interaction is negative: the surprise slope is positive at low uncertainty and negative at high uncertainty. Slopes are evaluated at the pooled 25th and 75th percentiles of listener-centred pre-event entropy, defined without looking at pleasure outcomes. An interaction without these slopes supports conditionality, but does not reproduce the full crossed pattern.

The second question concerns listener learning. Within Arm L, random assignment determines whether the continuation distribution changes after outcome trial 16 or remains stable. The primary endpoint is the change in reported probability of the initially dominant continuation between trials 9–16 and 25–32. A larger decrease in the change group than the stable group supports exposure-dependent updating. This contrast is evaluated regardless of pleasure. Improvement in subsequent logarithmic prediction score is a complementary measure that can distinguish calibrated learning from a change in confidence.

The third question concerns the predictive difference made by learning a person. It requires three complementary improvements on shared transfer observations: fixed personal MI over fixed population MI, expectation-informed updating MI over pleasure-only updating MI, and expectation-informed updating MI over an equally updated acoustic learner receiving the same reports. The first isolates learning from the 20 selected references; the second tests the added value of expectation; the third tests MI's additional representation. One-sided superiority tests for the three prespecified mean-absolute-error differences are combined by an intersection–union rule; the composite p value is their maximum. Each difference and its interval are also reported independently, so a partial success remains interpretable. Together with the uncertainty and listener-learning tests, this constitutes one Holm-corrected family of three primary claims at $\alpha = 0.05$.

Value learning is a distinguishable extension

Arm L also measures expected enjoyment before and experienced pleasure after each outcome. An error-driven value learner,

$$V_{t+1} = (1 - \alpha)V_t + \alpha Y_t, \tag{3}$$

is compared with fixed value, repetition, unsigned surprise and experienced-reward alternatives. Merely finding a correlation between δ_t and $V_{t+1} - V_t$ is inadequate because both differences contain V_t. The decisive comparison scores next-trial expected enjoyment or choice out of sample, with each competitor given the same preceding information. Individual learning rates are regularized, and global hyperparameters are selected using pilot data only. This is a secondary test of a reward-learning account, not a direct test of the current RPEM equation or neural reward circuitry.

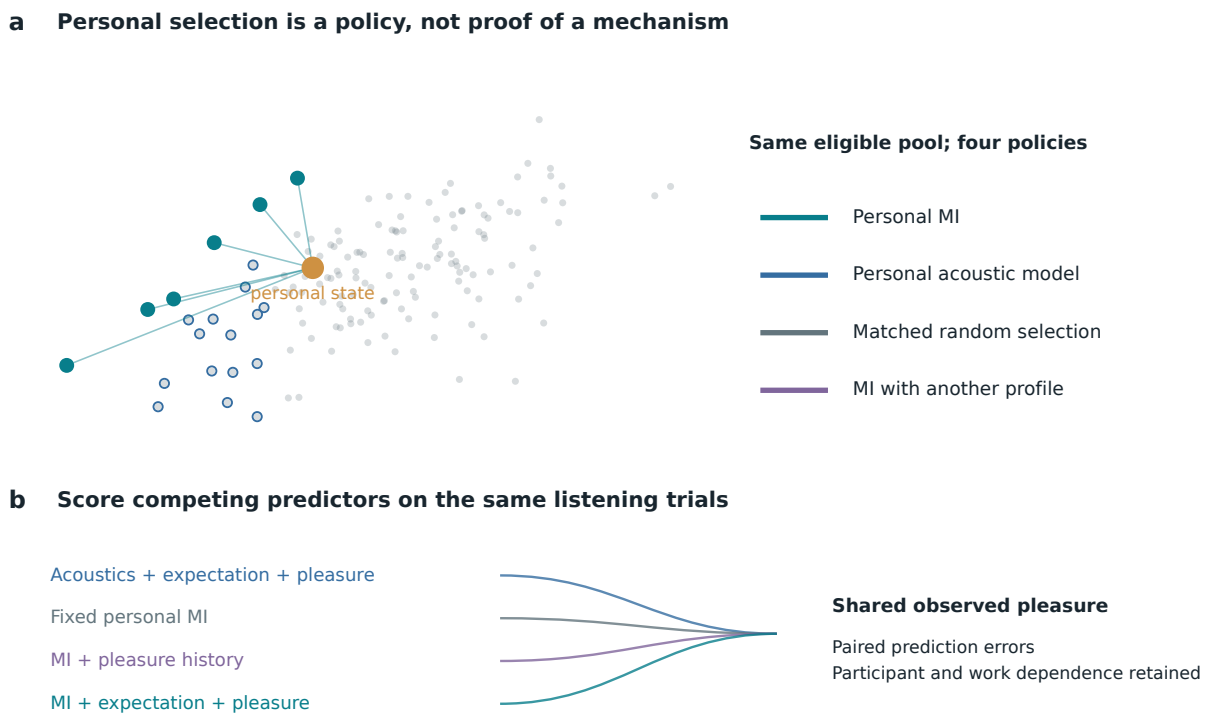


Fig. 5 | Selection and prediction require different controls. a, A personal profile informs selection from a shared pool. Positions and connecting curves are synthetic display coordinates; they do not represent a fitted embedding or observed pleasure. **b,** Competing predictors are scored on identical listening observations, regardless of which policy selected the excerpt. Policy value and prediction accuracy are analysed separately. The expectation-informed acoustic comparison receives the same reports as the expectation-informed MI model.

Predicting a person without seeing their future

The representation is computed once for each versioned audio excerpt. Participant-specific preparation uses the 20 reference ratings and declared background variables to initialize a regularized personal mapping. Global mappings and dimensionality reductions are fixed before confirmatory recruitment. The release includes the stimulus schedule, fixed predictions, random seed and all relevant content hashes. During listening, no MI, GPU or LLM inference is required.

An updating predictor is evaluated later through chronological replay. For each transfer trial, it receives only the participant's intake, their preceding mechanistic task and responses from earlier transfer trials. Its state is saved, the prediction is committed, and only then is the current response revealed. The representation may describe the entire prepared audio excerpt, including its later acoustics, because the computational task is prediction of a listener's response to a known stimulus. Such access does not establish that the listener anticipated that future audio. Claims about human anticipation instead use measurements made before the target event.

The expectation-informed state combines the participant's continuation reports, calibration and updating parameters with past pleasure. Fixed global functions map this state and the current audio features to predicted pleasure. The dimensionality, parameter priors, update equations and regularization grid must be executable and frozen at the pilot exit. At most six personal coefficients are estimated from the 20 reference ratings; unsupported coefficients shrink to the group model. Module output selection cannot be optimized on confirmatory outcomes. A shuffled-person state and an expectation-report permutation supply secondary specificity checks.

Does the selected music reach the listener?

The policy comparison estimates mean pleasure under each of the four prepared policies. Policy slots are interleaved in concealed, balanced six-trial blocks. Candidate sets are disjoint across policies for an individual, drawn by a recorded random partition before ranking; partition membership rotates across participants. This prevents duplicate listening and gives each policy comparable opportunity. The random policy is frequency-matched to the eligible pool on instrumentation, duration and broad familiarity strata, but is not matched on MI score. MI and acoustic policies use the same profile and ranking budget. Mismatched profiles are paired within broad training and catalogue-coverage strata without sharing identity data.

Because the selected excerpts differ, this comparison estimates the performance of a selection policy, not the causal effect of a module value at fixed audio. A secondary within-work contrast and identical-audio context probes address different confounds. Policy superiority, mechanism support and predictive superiority are reported independently, including when their conclusions diverge.

Temporal precision adds another requirement. Candidate pleasurable moments are declared from prepared signals before response inspection. Their local pleasure changes are compared with matched control moments within and across excerpts. Report delay and temporal smoothing are calibrated in the pilot, and lag choices are then frozen. Moment-level results remain secondary until the timing assay qualifies; excerpt-level ratings are the primary pleasure endpoint.

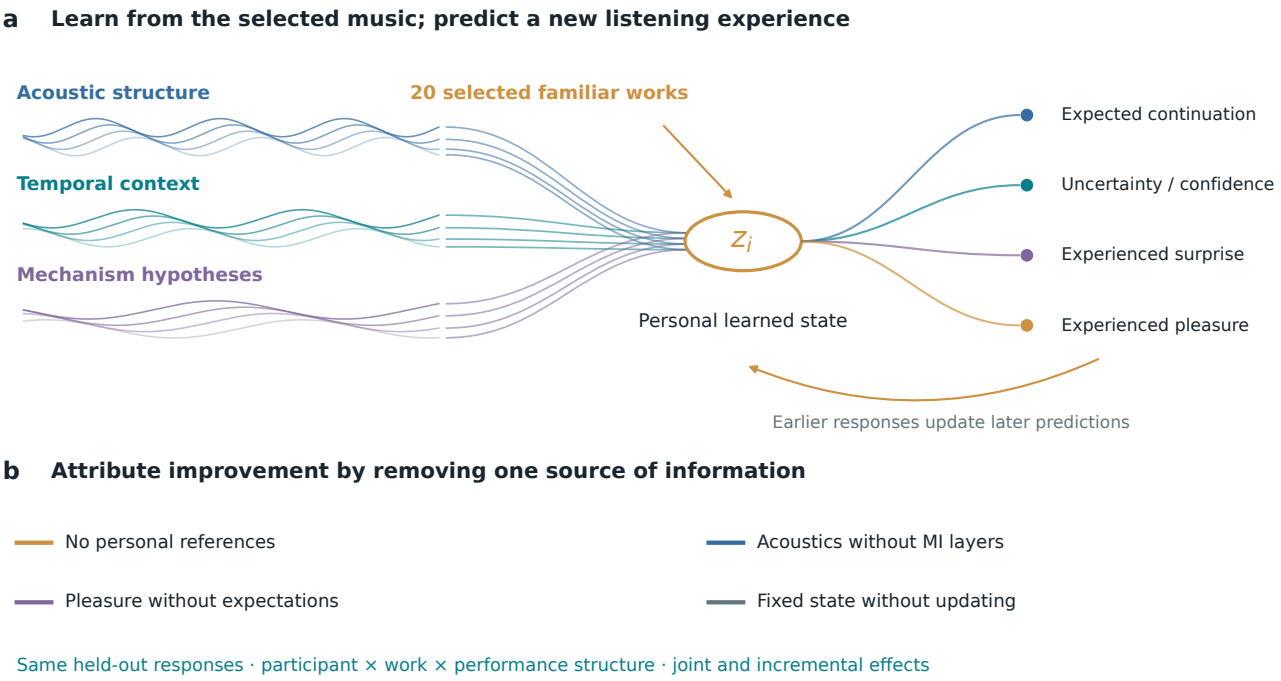


Fig. 6 | The central test is the gain from a learned personal model across multiple layers. a, Selected reference music initializes a personal state. Acoustic structure, temporal context and mechanism-related representations jointly inform predictions of expectation and reaction. Earlier responses can subsequently update the state. Lines illustrate model structure, not measured neural pathways or fitted trajectories. **b,** Information ablations quantify what is lost when personal references, MI's additional layers, expectation reports or state updating are removed. All variants are scored on the same held-out observations.

A holistic model with identifiable contributions

The unit of explanation is the listening process, rather than an isolated correlation between one module and one rating. A musical event has acoustic properties, a temporal context and a relationship to the listener's previous experience. Those layers jointly constrain expected continuations, uncertainty, surprise and value. MI makes this joint account implementable by supplying a common representation on which personal learning and competing hypotheses can operate. The online method supplies repeated observations across people and music; the dataset connects them without losing their distinct sources of variation.

Let x_j^A, x_j^T and x_j^M denote acoustic, temporal and mechanism-related feature groups for excerpt j. Let B_i contain the listener's selected reference works and reports, and let z_{it} be the personal state before trial t. The proposed model factorization is

$$z_{i0} = G(B_i; \theta), \tag{4}$$

$$z_{i,t+1} = U(z_{it}, o_{it}; \theta), \tag{5}$$

$$\widehat{o}_{it} = F(x_j^A, x_j^T, x_j^M, z_{it}; \theta), \tag{6}$$

where o_{it} contains the outcomes actually collected in that task. Global parameters θ are fixed before confirmatory observation. The equations specify the information structure; the pilot-exit implementation must instantiate the constrained functions and parameter priors. They do not describe an already trained MI listener model. Missing outcome types are masked, not imputed as if every task measured every construct.

The personal reference set is scientifically consequential. It provides a small, participant-chosen sample of the relationship between familiarity, structure and preference. Fixed population MI and fixed personal MI differ only in whether that calibration is available. A secondary dose analysis initializes otherwise identical models with 0, 5, 10 and 20 reference families, using ten preregistered random subset orderings. All variants predict the same new excerpts. This yields a learning curve for the information gained from selected music, without confusing greater intake size with an easier test set. Missing reference ratings and effective calibration counts are reported explicitly.

Test layers together, then remove them

The explanatory analysis jointly models probability reports, unexpectedness and pleasure, with participant and composition effects and performance nested within composition. A logistic-normal observation model respects the probability simplex; bounded rating models describe reported affect. Entropy and surprise are derived from the same distribution, not treated as independent observations. Temporal and mechanism features enter alongside acoustic controls. Training and familiarity moderate prespecified paths. Expected enjoyment, surprise and signed value error retain distinct roles, while response dependence is estimated within a common model.

This joint model is a secondary explanatory analysis, while the primary contrasts remain fixed and interpretable. Layer ablations compare acoustic features alone, acoustic plus temporal features, and the full representation under the same capacity and tuning budget. Personalization ablations remove reference calibration; learning ablations freeze state updates; expectation ablations retain pleasure history but remove probability and expected-value reports. A shuffled personal state tests whether the improvement depends on whose history was learned. These comparisons identify incremental and overlapping information; correlated module outputs do not permit unique causal attribution merely because their names differ.

Scale makes this strategy useful across heterogeneous listeners and works. Dense repeated measures support estimates of within-person change; many people support analysis of individual differences; diverse compositions support transfer beyond one stimulus. Their effective sample sizes remain different. The intended contribution is coherent evidence across these levels, with shared measurements and explicit countermodels, rather than a collection of disconnected positive associations.

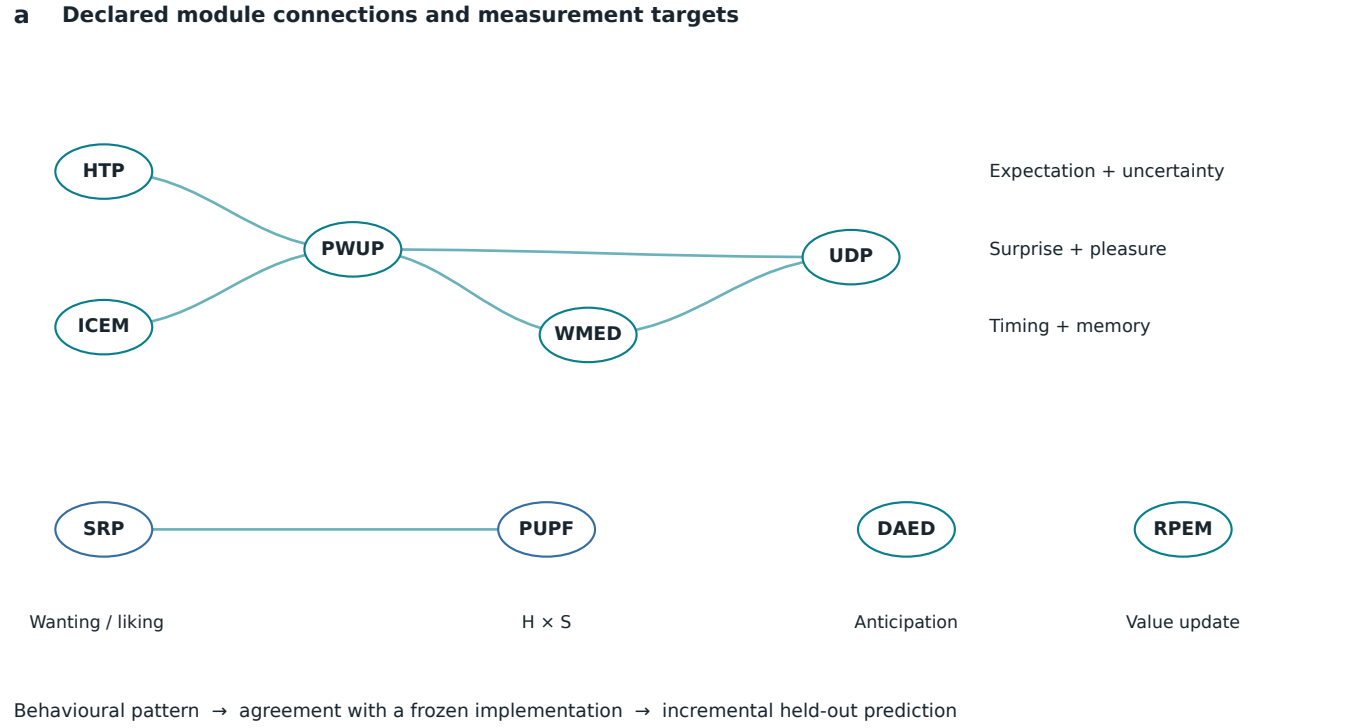


Fig. 7 | The mechanism audit distinguishes declared computation from empirical explanation. Six connections are declared among the nine audited modules. These curves reproduce the registry's dependency structure, not a neural circuit. RPEM and DAED are not given invented upstream edges. Labels below the modules identify candidate behavioural measurements. The evidence sequence at the bottom separates observing a human pattern, reproducing it with a fixed implementation and adding held-out predictive information.

A mechanism can fail at more than one level

HTP and ICEM. Hierarchical Temporal Prediction is motivated by work on auditory integration and learned sequence prediction[10, 11]. Those are related but distinct empirical claims. Current HTP outputs mix acoustic stability, pitch, timbre and onset cues; they do not constitute a listener's probability distribution. Information Content Emotion Model outputs likewise derive from acoustic variation rather than $-\log_2 q(e)$. Arm U tests whether these scores track elicited expectation and surprise beyond acoustic change. Arousal is collected on a counterbalanced subset of probes, separately from pleasure; no cardiac or skin-conductance validation is claimed.

PWUP and UDP. Precision-Weighted Uncertainty Processing attenuates its raw error according to $E(1 - 0.5g)$, where g is its estimated gate. Its name does not determine the expected empirical direction. Neural deviance responses have been found to depend on musical uncertainty[12], but an online pleasure effect and this particular gate require independent tests. Uncertainty-Driven Pleasure contains an upstream uncertainty/confirmation branch and a separate acoustic pleasure index. Their incremental contributions are evaluated separately. A theory of pleasure from uncertain music[13] supplies motivation, not validation of either equation.

SRP and PUPF. Striatal Reward Prediction supplies audio-derived wanting, liking and tension-related proxies. Tension and its neural correlates have been studied directly[14]; tension, relief and liking must still be distinguished in this task. Pleasure-Uncertainty Prediction Function includes interactions among acoustic proxies and pleasantness. It is not a fitted copy of the empirical uncertainty–surprise surface. The preregistered comparison must retain the frozen implementation even if it disagrees with the behavioural pattern.

DAED, RPEM and WMED. Dopamine Anticipation-Experience Dissociation and Reward Prediction Error in Music carry mechanistic names, but their current outputs are acoustic computations. RPEM's centred channel is not the participant's reported $Y - V$, and it does not itself update a personal learner. Working Memory-Entrainment Dissociation concerns separable timing and memory accounts motivated by studies using tapping and neural measures[15, 16]. This protocol's core blocks do not measure working-memory capacity or neural entrainment. Additional timing, memory, wanting and anticipation assays are specified as extensions below. Without them, those channels can be evaluated only as predictors of the measured outcomes.

The audit therefore covers all nine modules while restricting confirmatory claims to the three questions above. A wide interval around a null effect is inconclusive. A precise incompatible pattern challenges the operationalized hypothesis in this population and task. It does not by itself refute predictive processing as a general theory of music.

Methods

Design status and registration

This is a prospective, pilot-dependent protocol and analysis specification. It contains no participant results, institutional approval or completed public preregistration. A Stage 1-style structure makes hypotheses and planned analyses inspectable before data collection[17]; it does not imply acceptance as a Registered Report. The architecture manuscript and O-Lab systems companion remain separate records. The former describes executable representations; the latter specifies how this study would be run and reproduced.

The target is 1,000 distinct, verified, minimally eligible registrations. The first 60 form a feasibility and parameter-setting pilot; the remaining cap is 940, before screening, readiness failures or attrition. This allocation is a new planning choice, not a claim that 940 observations provide adequate power. A registration is counted once after verification, age eligibility and agreement to initial screening; study consent is a separately versioned event. Withdrawals and incomplete baselines do not disappear from the flow or reopen recruitment slots. Counts are reported for invitations, verification, screening, consent, completion of 20 references, release readiness, each task, transfer completion and analysable endpoints.

Adults aged at least 18 who can consent in the study language and use supported headphones and a desktop or laptop may enter. Recruitment should span musical training and classical familiarity, with strata recorded rather than inferred from demographic stereotypes. Training years, current practice, typical classical listening, language, age band, country, self-reported hearing difficulty and assistive hearing use are collected with explicit missing/prefer-not-to-answer categories where appropriate. Hearing reports characterize the sample; technical failure is assessed separately. No diagnosis is inferred. The initial study uses one validated language version; later translations require cognitive interviewing and counterbalanced validation before pooling.

Catalogue identity, rights and recognition

The catalogue distinguishes composition family, movement, arrangement, performance, recording master and excerpt. Different encodings of one master are not independent performances. Movements of the same family cannot fill multiple positions in the 20-work intake. Reference and transfer selections record exact passage, duration, instrumentation, performer attribution and prior exposure within the study.

Candidate sources include per-recording rights-cleared classical collections and the MAESTRO research dataset[18]. Dataset access or an old composition's public-domain status does not settle the rights of a particular recording, its derivative excerpt or web redistribution. Every admitted asset requires a retained rights record covering the proposed use, territory and attribution. A licence-restricted research recording is not automatically suitable for unrestricted public playback. The confirmatory catalogue is frozen only after this audit. No YouTube downloading is part of the proposed workflow.

An intake preview lasts 12 seconds. Participants answer whether they knew the work before today, how familiar it feels on a five-level scale, whether they recognize this particular performance and how pleasurable the preview is on a 0–100 scale. Exactly 20 distinct families with explicit prior recognition are retained. Familiarity with a composition does not imply familiarity with the selected performance. An "unable to rate" response remains missing; it is not converted to zero. At least 16 usable pleasure ratings are required to fit a personal calibration; otherwise the participant can finish intake but requires additional rating resolution before readiness. Model coefficients unsupported by the intake shrink toward the pilot-trained population prior.

Session order and task comprehension

After consent and browser/audio checks, participants complete intake in a separate visit. Preparation then produces the immutable release. The experimental visit starts with comfortable-volume adjustment, a headphone screening task[19], and comprehension trials demonstrating the difference between likelihood, confidence, surprise and enjoyment. The headphone task is a screening aid, not proof of calibrated acoustic delivery. Persistent equipment or comprehension failure is recorded before randomization.

The proposed experimental visit contains one mechanistic arm, a break, the common transfer block and the paired-performance extension. Intake and experiment are expected to require approximately 55–75 minutes combined, depending on catalogue search; this is a planning range that must be replaced by measured pilot completion times. Participants may pause between blocks or withdraw. Payment, recruitment channels, consent wording, responsible institution and contact arrangements must be finalized in the ethics-approved operational plan before invitations. No contingent bonus depends on reporting pleasure or reproducing an expected pattern.

Arm U: separating context and outcome

Two visually labelled contexts have underlying continuation distributions $(0.20, 0.75, 0.05)$ and $(0.20, 0.40, 0.40)$. Continuation identities are permuted across participants. A shared outcome can therefore have equal probability across contexts while the alternatives' dispersion differs. Other outcomes provide variation in surprise. This is not a fully orthogonal manipulation of subjective entropy and surprise: both derive from the same reported distribution, and their empirical support must be checked.

Participants first hear 24 training outcomes, 12 per context, from seeded schedules sampled from the declared distributions. Training and 24 subsequent probe trials use non-overlapping sequence realizations, balanced context order and identical available target endings. The seeds are sampled without filtering for eventual pleasure responses. A six-second context precedes a three-second continuation. Before each probe outcome, participants allocate 100 probability points among three previously demonstrated audible continuations using constrained controls. They also report confidence and expected enjoyment. No musical notation knowledge is required. After the continuation, they report surprise and pleasure. Arousal is additionally sampled on eight prespecified probes, balanced across context and outcome class.

The context is replayed once after prior reports to preserve a defined auditory lead-in; the delay and replay are logged. The same rule applies in both arms. Prior responses are locked before playback. A zero reported probability is retained as a meaningful response; for logarithmic scoring only, probabilities are floored at 0.01 and renormalized. Sensitivity analyses use 0.005 and 0.02. Pilot checks determine whether participants understand the simplex and whether the two contexts produce separable reported

entropy without floor/ceiling saturation. The precise timing and waveform joins are frozen after this qualification.

Arm L: changing an environment

After 16 initial training exposures with probabilities $(0.70, 0.20, 0.10)$, participants receive 32 outcome probes using the same prior/outcome reporting sequence. Half are assigned to a change condition: after trial 16 the probabilities become $(0.20, 0.70, 0.10)$. For the other half they remain stable. Outcome identities and change assignment are independently counterbalanced. Samples are drawn from seeded schedules; the intended distribution, realized sequence and each delivered exposure are all retained. No trial is inserted online to force an effect.

The primary probability-update contrast compares mean $q(A)$ over trials 9–16 with trials 25–32, where A is initially dominant. It is estimated with the randomized condition-by-phase interaction, adjusting for baseline probability and realized exposure imbalance. Secondary analyses compare next-outcome log score, adaptation rate and signed value-learning models. Learning is not inferred from lower uncertainty, increased pleasure or repetition alone. Because these short sequences may yield weak parameter recovery for some listeners, individual learning-rate estimates carry hierarchical uncertainty and are not used as diagnostic traits.

Transfer and within-work procedures

The common transfer pool excludes the participant's 20 intake families. This tests transfer across works; it does not ensure that every other work is unfamiliar outside the study. Recognition of each excerpt is reported after listening and enters the prespecified familiarity analysis. Excerpts are 25 seconds, with short standardized fades and a fixed loudness-normalization policy preserving internal dynamics. Clipping, excessive noise, unmatched truncation and inconsistent channels are rejected by catalogue QA. Loudness, roughness, tempo, onset density, tonal stability and instrumentation descriptors are retained as control variables.

Each of the four policies fills six slots from its randomly assigned eligible candidate subset. Eligibility and matching use only intake information and catalogue metadata. Failure to fill any subset makes the release unready; it does not silently relax criteria. Slot order is randomized within six-trial blocks subject to balanced policy counts across the complete session. A continuous pleasure control is sampled at 5 Hz during playback. Immediately after each excerpt, overall pleasure is rated from 0 (no pleasure) to 100 (extreme pleasure), with a distinct unable-to-rate option. A short rating of unexpectedness and recognition follows. Overall pleasure, not the maximum slider value, is the primary target.

The secondary within-work task chooses four intake families with two eligible, previously unheard performance excerpts beyond the intake reference. Each pair uses the same passage and instrumentation; B/C order is counterbalanced. Both receive pleasure ratings, followed by a forced preference with an indifferent option. Only this subcatalogue requires three eligible performances per family. Different arrangements are not pooled with performance changes; an arrangement study would require its own manipulation and analysis.

Candidate predictors and information access

The acoustic baseline uses the fixed descriptors above, summary dynamics and reference ratings. The MI candidate set is the versioned set of audited outputs from HTP, ICEM, PWUP, WMED, UDP, SRP, PUPF, DAED and RPEM. Both receive the same timing information, participant background and allowable reports. Comparisons include fixed acoustic and fixed MI mappings, pleasure-only updating MI, expectation-plus-pleasure MI and expectation-plus-pleasure acoustic mappings. A regularized linear readout is the primary function class; a nonlinear model with matched tuning budget is a secondary robustness analysis.

The primary personal readout uses six standardized coordinates. In the full model, two principal components come from each of the acoustic, temporal and mechanism feature groups; the acoustic comparison uses six acoustic components. Component transforms are learned on the pilot catalogue and frozen. A population ridge model provides the coefficient prior. Personal coefficients minimize squared error on available reference ratings plus a penalty for deviation from that prior, with strength selected from $\{0.1, 1, 10, 100\}$ by pilot validation grouped by participant and composition. This equal-dimensional design constrains calibration from 20 references. The pilot-exit record lists all transformations, priors, coefficients and software hashes.

The pleasure-only learner updates this penalized readout cumulatively after each eligible earlier rating. The expectation-informed learner additionally uses a global, pilot-fitted adjustment based on earlier mean entropy, mean surprise, expected enjoyment, confidence calibration and the change in continuation accuracy. These summaries use only completed mechanistic probes and are initialized before transfer; subsequent reported surprise and pleasure update only the features actually remeasured. They do not invent continuation probabilities for whole excerpts. Global adjustment coefficients are fixed; they do not add personal fit parameters. The acoustic learner receives precisely the same adjustment inputs and pilot fitting budget. A nonlinear or richer state-space learner belongs to the secondary analysis. The final executable specification must pass replay and parameter-recovery checks before confirmatory release.

For each transfer trial, the expected-state variables summarize earlier elicited q, confidence, prediction calibration, adaptation and value reports. Current and later transfer answers are inaccessible until prediction commitment. The core comparison uses equal-capacity mappings so that an advantage is not merely the addition of more adjustable parameters. A fixed model's predictions are committed in the participant release; an updating model's chronological predictions are committed in the offline replay ledger. Both are clearly identified in all exports.

Primary analysis and dependence

Arm U fits a linear mixed model to 0–100 pleasure with centred H, S, HS, H^2, S^2, context, order, reference familiarity and acoustic controls. Participant and continuation-family intercepts are included; participant slopes for surprise and entropy are retained when identifiable. The preregistered singular-fit fallback first removes random-effect correlations, then the least-supported random slope, retaining participant and stimulus intercepts. The primary interaction test is directional; confidence intervals and both conditional slopes accompany it. A bounded-outcome model provides sensitivity analysis without replacing the primary endpoint.

Arm L fits the randomized condition-by-phase contrast to probability reports, with participant intercepts and counterbalanced

outcome identity. The effect of changed exposure is interpretable at the assignment level. Associations between an individual's entropy, surprise and pleasure are not automatically causal, because reported beliefs were not independently randomized.

For prediction, absolute errors are first averaged within participant across eligible shared transfer trials. Differences between competing predictors are paired on exactly the same observations. The primary estimand gives equal weight to each participant, not to the most complete or most highly sampled person. Confidence intervals and one-sided tests use a crossed participant-by-composition bootstrap with 5,000 replicates; composition-family assignments are retained across all occurrences. A predeclared cluster wild-bootstrap sensitivity analysis addresses sparse work coverage. Power simulations must assess the selected resampling procedure's error control before launch. Mean squared error, calibration slope and signed bias are secondary reports.

For policy value, a mixed model estimates mean pleasure differences among randomized policy slots with participant and work effects, order and prespecified familiarity controls. The estimand is the average outcome of these complete policies in the released pool. It is not a population-wide advantage over all possible classical music. Secondary comparisons are reported with false-discovery-rate control within mechanism, policy and temporal families; exploratory analyses are labelled and are not promoted to confirmatory status by a small p value.

Power, precision and pilot exit

The registration cap does not establish effective sample size. Pilot estimates must characterize baseline completion, task comprehension, entropy/surprise coverage, outcome variance, participant/work dependence, missingness and assignment balance. Simulation then recreates the actual allocation, stimulus overlap, attrition and proposed models. The target is at least 90% power for the smallest scientifically meaningful confirmatory effects under the worst Holm threshold, while verifying Type I error under null and misspecified scenarios.

Initial planning thresholds are an interaction of magnitude 0.10 in a model with standardized outcome and centred standardized entropy/surprise, a 0.10 change-group versus stable-group difference in reported continuation probability, and a one-point reduction in absolute pleasure-prediction error on the 0–100 scale for each predictive contrast. These are proposed meaningful-effect targets, not literature effect-size estimates. Pilot reliability and substantive interpretation must justify or amend them before confirmation, without selecting thresholds for favourable significance. Simulations include 300, 500, 700 and 900 analysable participants, plausible work clustering and differential dropout. If the achievable sample cannot distinguish the intended effects, the scope or recruitment design is revised before collection. The pilot-exit amendment, executable simulations and all changes remain beside this version; no achieved power is claimed here.

Missingness, interruptions and exclusion

All exclusions are based on consent, identity duplication, technical delivery or prespecified task-quality rules, not on the direction of pleasure or agreement with MI. No participant is removed merely for low or constant pleasure. The primary prediction set requires at least 18 of 24 valid transfer ratings; mechanistic arm summaries require at least 75% valid probes and non-empty reporting in both analysis phases or contexts. Analyses using all available valid observations are reported as sensitivity checks, with completion probabilities described by intake characteristics.

A trial is technically invalid if audio does not start, playback stalls, the page becomes hidden during the target segment, or the prescribed outcome differs from the release. Late upload alone is not failure if local timing and a verified receipt survive. A replayed excerpt constitutes another exposure. Its second response does not replace the first unnoticed; repeat exposures are excluded from the primary first-exposure analysis and retained for sensitivity analysis. Missing reports are never forward-filled. Missing baseline covariates may use pilot-frozen imputation with missingness indicators; outcomes are not imputed for the primary comparison.

Timing and optional mechanistic extensions

Browser measurement requires validation rather than confidence in timestamps[20]. The audio clock, browser monotonic clock and server receipt time are recorded separately. The pilot measures latency and variability on representative hardware. Pleasure traces are analysed at predeclared windows after an empirically qualified response lag. No lag or peak is optimized on the confirmatory trace being scored.

Additional assays remain secondary, separately consented extensions. A tension/relief task compares resolved, delayed and withheld endings using independently cleared matched audio, with wanting and expected enjoyment before the ending and relief and liking after it. An anticipation assay uses separate response blocks for pre-event wanting and experienced pleasure, testing two processes against a single delayed response. A timing/memory assay crosses rhythmic regularity with two- versus four-item pattern memory, measuring tapping error and recognition accuracy separately. These designs can interrogate SRP, DAED and WMED interpretations; until implemented and qualified, their absence limits claims rather than being filled by module labels.

Ethics, availability and reporting

Consent must explain the audio task, time commitment, stored information, contact separation, withdrawal procedure and limits after irreversible aggregation. The study collects behavioural self-reports, not neural or neurochemical measurements. Pseudonymous records remain potentially personal data. Public releases should contain approved analysis code, synthetic fixtures, catalogue metadata permitted by licence and sufficiently de-identified summaries; identifiable contacts, consent records and restricted audio remain controlled. FAIR documentation concerns findability and reuse under clear conditions, not unrestricted access[21].

The final report will publish participant flow, all three primary tests, policy comparisons and implementation disagreements, including null findings. Figures in the present manuscript are analytic or conceptual illustrations generated from editable source. They must not be reused as empirical result panels. Existing architecture outputs, prospective protocol choices and future observations are distinct provenance classes.

Discussion

The most informative outcome is not simply that an algorithm chooses pleasant music. The programme asks whether expectations explain when pleasure changes, whether exposure changes those expectations, and whether a transparent personal model

can use that change to anticipate another response. A behavioural replication with an unsuccessful MI mapping would locate a computational deficiency. Accurate pleasure prediction without the expected mechanism pattern would favour a useful predictor with an incomplete explanatory account. Failure of both, when measured precisely, would motivate revising the representation and the task rather than renaming a score.

The design also has clear limits. A three-option probability report samples only a small part of musical expectation; reference selection favours people who can recognize a classical catalogue; eliciting priors may change listening; and laboratory-style continuations do not capture the full experience of a symphony. The transfer block and within-work extension address different parts of that gap. Their value lies in preserving the same question across levels of control: what did this listener expect, what occurred, what changed in pleasure, and what did the listener learn?

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.

Author and artifact information

Conceptualization, protocol development, software-artifact design and writing: Amaç Erdem. This edition is an author-controlled draft. Institutional sponsorship, funding and the operational ethics record are to be declared before recruitment. No competing commercial deployment or participant outcome is asserted by this document.

The accompanying source package contains the editable manuscript, original PDF/SVG figure assets, figure-generation code, design configuration, source audit and visual-quality record. Diagrams and numerical examples are explicitly illustrative. Earlier protocol editions and the architecture manuscript are preserved as dated records. The editorial format takes scientific-figure legibility and continuous argument as design constraints; it does not reproduce a journal's branding or imply endorsement.