Frequency-Specific Operant Learning in Neurofeedback Reveals Distinct Cortical Mechanisms: Evidence from Double-Blind ERSP and ERP Dissociations

Andrew Hill, PhD
Peak Brain Institute, Los Angeles, CA

Draft — April 2026

Abstract

Background. Neurofeedback reliably alters EEG activity, but the cortical mechanism by which reward-contingent feedback shapes oscillatory dynamics remains unresolved. In particular, no study has examined reward-locked event-related spectral perturbations (ERSP) under double-blind, active-placebo-controlled conditions.

Methods. Forty participants underwent five training sessions and a 3–5 week retention session of single-channel EEG biofeedback (C3 SMR 12–15 Hz, $n = 8$; C3 Beta 15–18 Hz, $n = 8$; C4 SMR 12–15 Hz, $n = 8$; active-placebo sham, $n = 16$), with concurrent 64-channel EEG recording. ERSP was computed from reward-locked epochs (approximately 600–700 trials per session) using Morlet wavelets (3–40 Hz) across four sessions.

Results. Active groups produced frequency-specific event-related desynchronization (ERD) in the rewarded band (pooled Active vs Sham $d = -1.23$, $p_{\text{adj}} = 0.001$; C3 Beta and C4 SMR FDR-significant, $|d| \geq 1.12$; C3 SMR trended, $d = -0.80$, $p_{\text{adj}} = 0.081$), absent in sham. A double dissociation emerged at C3: beta training produced the strongest ERD ($d = -2.38$), whereas C3 SMR training produced the largest P2 suppression ($d = -1.33$, $\text{BF}_{01} = 0.10$; smaller P2 at the trained electrode). Only SMR groups showed lasting plasticity, with increased eyes-closed alpha at follow-up (C3 SMR $d = 0.97$; C4 SMR $d = 0.78$) and significant across-session accumulation ($\beta = 1.44$, $p = 0.004$). ERD magnitude predicted long-term resting-state change ($r = 0.54$, $p = 0.009$) but not within-session shifts ($r = -0.09$, $p = 0.67$), dissociating transient from consolidating effects. An ICA-based sensitivity analysis confirmed convergence of all primary findings.

Conclusions. Neurofeedback engages frequency-specific, contingency-dependent cortical mechanisms. The ERD–P2 dissociation suggests that beta and SMR training recruit distinct circuits (C3 Beta: local cortical ERD with preserved P2; C3 SMR: putative thalamocortical relay ERD with suppressed P2 at the trained site) with different capacities for consolidation. These findings establish a multi-timescale model in which immediate reward-locked desynchronization drives durable plasticity only when supported by deeper circuit dynamics.

1. Introduction

Neurofeedback (the provision of real-time feedback about one's own brain activity) has demonstrated moderate-to-large effect sizes in some meta-analyses (Arns et al., 2009; Van Doren et al., 2019), though well-controlled trials with probably-blinded outcomes have produced smaller effects in ADHD (Cortese et al., 2016). Yet the field remains uncertain about *how* the brain integrates a reward signal yoked to its own neural oscillations and produces lasting changes in cortical physiology. Large sham-controlled trials have produced equivocal efficacy results (Schönenberg et al., 2017; Arnold et al., 2021), in part because the field lacks a mechanistic account to guide protocol design. Sitaram et al. (2017) identified this gap as a central priority; Thibault and Raz (2017) argued that many reported effects could be entirely attributable to placebo. While the field debates whether neurofeedback outperforms sham on behavioral outcomes, no study has examined what occurs in the brain at the moment of reward delivery during training itself.

The operant conditioning framework provides a testable hypothesis: if the brain learns to control its own oscillatory power through reward-contingent feedback, then measurable, frequency-specific cortical responses should emerge, time-locked to reward delivery. Sterman (1977, 2000) demonstrated that cats and humans could be operantly conditioned to increase sensorimotor rhythm (SMR, 12–15 Hz) power when rewarded for doing so. Fetz (1969, 2007) established that individual cortical neurons could be conditioned through contingent reward. Ros et al. (2010) have argued that neurofeedback engages the same neuroplasticity mechanisms underlying motor learning, and Sherlin et al. (2011; updated by Kerson, Sherlin, & Davelaar, 2025) extended this reasoning to oscillatory control. Davelaar (2018) formalized the first two stages of a multi-stage computational theory (implicit operant learning and structural thalamocortical change); a subsequent synthesis extends this with a third stage of interoceptive homeostasis (Kerson, Sherlin & Davelaar, 2025). Yet until recently, the analytical tools available have been poorly suited to visualizing frequency-specific learning signals with sufficient temporal and spectral resolution.

Event-Related Spectral Perturbation (ERSP) preserves both time and frequency simultaneously, revealing how oscillatory power evolves on millisecond timescales around significant events (Makeig, 1993; Pfurtscheller & Lopes da Silva, 1999; Grandchamp & Delorme, 2011). The distinction between phase-locked and non-phase-locked responses (Tallon-Baudry et al., 1996) is central to this framework. The ERSP framework describes event-related desynchronization (ERD) and synchronization (ERS), transient decreases or increases in oscillatory coherence, providing the resolution needed to ask: *Does the brain produce a frequency-specific, temporally-precise oscillatory response contingent on the learning signal, only when feedback is genuinely yoked to the target neural activity?*

Answering this question rigorously requires overcoming a substantial methodological hurdle: blinding. Unlike pharmacological trials, where inert placebos are standard, blinding neurofeedback participants to whether they are receiving real or sham feedback is notoriously difficult. Schönenberg et al. (2017) and Ros et al. (2020) have found that most "sham" conditions fail rigorous scrutiny: they either provide intermittent real feedback, lack temporal coherence, or employ crude pseudorandom sequences that participants can distinguish. Thibault et al. (2016) showed that

blinding failure could entirely account for reported effects. Arnold et al. (2013, 2021) and Sorger et al. (2019) have articulated standards for active-placebo control: the sham must provide real-time, artifact-merged EEG data with matched temporal resolution and amplitude scaling, with feedback contingency hidden, such that participants cannot reliably distinguish active from sham training. Only with such controls in place can one credibly claim that observed changes in brain activity are contingent on the *veridical* learning signal.

The present study addresses both gaps by examining the brain's *automatic*, reward-locked spectral response during operant conditioning, without instructing subjects to perform any specific mental task. Unlike prior work that examines ERPs or spectral changes *before and after* neurofeedback (Salansky et al., 2024, 2025; Egner & Gruzelier, 2001, 2004), that studies ERD during *voluntary* motor imagery aided by feedback (Muraoka et al., 2024), or that evaluates frequency specificity using only tonic power from a single training frequency (Dessy et al., 2020), we examine the reward-locked ERSP itself. Four groups (N = 40: C3 SMR n=8, C3 Beta n=8, C4 SMR n=8, Sham n=16) underwent double-blind, active-placebo-controlled neurofeedback with concurrent 64-channel EEG across five training sessions plus a retention visit (~3–5 weeks post-training). We hypothesized that active groups would produce frequency-specific ERD in the rewarded band, time-locked to reward delivery, absent in sham. These data were originally analyzed descriptively (Hill, 2012); we reanalyze them here using cluster-based permutation testing, Bayesian inference, mixed-effects models, and standardized effect sizes.

---

2. Methods

2.1 Participants and Design

Forty healthy adults with no prior neurofeedback experience were recruited from the University of California, Los Angeles undergraduate community. Participants were randomly assigned to one of four groups by an unaffiliated researcher: C3 SMR (n = 8; reward 12–15 Hz at C3), C3 Beta (n = 8; reward 15–18 Hz at C3), C4 SMR (n = 8; reward 12–15 Hz at C4), or active-placebo sham (n = 16). All 40 participants were included in analyses. Data were collected as part of the first author's doctoral research (Hill, 2012); the present ERSP-focused reanalysis was designed in 2025–2026 using statistical methods not available at the time of original data collection, with the analysis plan documented prior to re-accessing raw BDF files. Because the first author's 2012 dissertation analyzed these data descriptively, the present reanalysis is not a pre-registered replication. The pre-documented analysis plan mitigates but does not eliminate prior-exposure bias; a prospective replication is the natural next step. The study was approved by the UCLA Institutional Review Board; all participants provided written informed consent.

2.2 Session Schedule

Participants completed five training sessions over approximately one week (sessions 1–5, alternating days), plus a follow-up session approximately 3–5 weeks post-training (Session 6). BioSemi 64-

channel EEG was recorded concurrently with biofeedback on sessions 1, 3, 5, and 6. Sessions 2 and 4 used the biofeedback system only. ERSP analyses use all four BioSemi sessions.

2.3 EEG Acquisition

A BioSemi ActiveTwo system recorded 64-channel EEG (active Ag/AgCl electrodes, 10-20 layout, 512 Hz) with bilateral ear references and six ExG channels. Concurrent single-channel biofeedback was acquired via a ProComp+ Infiniti (Thought Technology Ltd.) at the training site (C3-A1 or C4-A2), feeding into EEGer software. The ProComp+ is optically isolated from the BioSemi system. EEGer transmitted TTL triggers to the BioSemi status channel at each reward event onset.

2.4 Neurofeedback Protocol

Training used a three-band protocol: reward (SMR 12–15 Hz or Beta 15–18 Hz), inhibit theta (4–7 Hz), and inhibit high beta (22–40 Hz). Reward required simultaneous threshold achievement in all three bands for 500 ms continuously. Reward events consisted of a simultaneous 125 ms auditory beep and picture-grid reveal (both triggered at reward onset; the concurrent visual event occurs within the P2 analysis window). Auto-thresholding recalibrated every 30 seconds, targeting ~70% reward rate. Each 30-minute session yielded approximately 600–700 reward events. Reward rates did not differ across groups (mean = 618 ± 45 events/session; $F(3,36) = 0.29$, $p = 0.83$). Trial counts entering the reward-locked ERP average (§2.6) therefore did not differ by group, eliminating power-asymmetry confounds in the P2 contrast.

2.5 Active-Placebo Sham

The sham was developed in collaboration with EEGer Inc. and constitutes an active placebo that eliminates contingency while preserving realistic non-specific effects. Fifteen pre-recorded 3-minute clean EEG segments were assembled in random order to produce a 30-minute sham signal, scaled in real time to match the participant's broadband EEG amplitude (the fifteen sham segments were amplitude-matched to each other and to the live raw signal before blending). Blinks, muscle artifacts, and electrode noise from the participant's live EEG were merged into the sham display, making it visually indistinguishable from veridical feedback. All reward thresholds and events were derived from the sham signal. Configuration was password-protected and hidden; group assignment was performed by an unaffiliated researcher. This implementation passes three critical tests: (a) the signal looks alive (artifact merging), (b) the reward rate is plausible (auto-thresholding on EEG-like source), and (c) the operator cannot detect condition (hidden configuration). (a) and (b) are infrastructural design properties; (c) is a design-enforced blinding feature rather than a measured blinding-integrity check, since trainer-guess data were not systematically recorded in the original protocol. To our knowledge, this remains the most rigorous neurofeedback sham in the published literature (see Supplementary S5.3 for sham taxonomy).

2.6 Preprocessing and ERSP Computation

Raw BDF files were processed in MNE-Python (Gramfort et al., 2013). Data were high-pass filtered at 0.16 Hz, notch filtered at 60 Hz, and re-referenced to the common average. Reward-locked epochs (−500 to +1500 ms) were baseline-corrected (−100 to 0 ms) and artifact-rejected using an automated statistical procedure (see Supplementary S1 for full details), retaining 94–98% of epochs.

ERSP was computed using Morlet wavelets (3–40 Hz) with single-trial normalization (Grandchamp & Delorme, 2011). The ERD magnitude metric (mean dB in the subject's reward band within the 200–800 ms post-reward window at C3) served as the primary dependent variable. ERP analysis used the same epochs bandpass-filtered 0.5–30 Hz. Resting-state power spectral density was estimated by Welch's method.

2.7 Statistical Analysis

Between-group ERD differences were tested with a linear mixed-effects model (ERD ~ Group × Session, random intercept by subject) and planned contrasts (independent t-tests with Cohen's d, Bayes factors, and FDR correction). Cluster-based permutation tests (Maris & Oostenveld, 2007; 1000 permutations) controlled family-wise error across the time-frequency plane. ERP components (P50, N1, P2) were tested with 4 Group × 4 Session mixed ANOVAs. Resting-state change scores (Session 6 minus Session 1) were compared between groups with FDR-corrected t-tests; a growth curve model tested pre-session baseline trajectories. Full statistical specifications are in Supplementary S2.

3. Results

3.1 Frequency-Specific, Contingency-Dependent ERD

The frequency crossover (Figure 2) provides the most direct demonstration of specificity: C3 SMR and C3 Beta, trained at the same electrode with reward bands differing by only 3 Hz, produced ERD peaking at different frequencies. Sham showed a flat profile. Any non-specific explanation would predict identical spectral responses; the crossover is difficult to reconcile with non-specific accounts. Per-subject peak ERD frequencies confirmed the separation quantitatively: C3 SMR peaked at 13.8 ± 2.8 Hz, C3 Beta at 16.1 ± 1.8 Hz ($d = -0.89$; Supplementary S3.6).

Grand-average ERSP heatmaps at C3 (Figure 1; C4 in Supplementary Figure S2) confirmed this pattern across all sessions: active groups produced frequency-specific ERD in the reward band (C3 Beta at 15–18 Hz, C3 SMR and C4 SMR at 12–15 Hz) while sham showed no reward-band ERD despite identical auditory stimulation. C3 Beta and C4 SMR reached FDR-corrected significance; C3 SMR trended in the same direction (see contrasts below). All groups, including sham, produced theta ERS (4–7 Hz, 100–500 ms), likely reflecting a mixture of early auditory-cortical processing and mid-frontal feedback/reward theta (Cavanagh & Frank, 2014); the latter is expected even in sham because sham participants still receive the tone as (placebo) feedback.

Planned contrasts confirmed a large pooled Active vs Sham effect (d = −1.23, p_adj = 0.001, BF01 = 0.01). The pooled contrast tests the contingency-dependence hypothesis; group-specific contrasts below separately test frequency and site specificity. C3 Beta showed the strongest individual contrast (d = −2.38 vs sham); C4 SMR was significant (d = −1.12); C3 SMR trended (d = −0.80, p_adj = 0.081). An LME model (ERD ~ Group × Session, random intercept by subject; 157 of 160 possible observations, 3 missing due to excluded sessions) revealed a large Group effect (η_p^2 = 0.38) but no Session or Interaction effects, confirming the ERD is immediate and stable, present from Session 1 (one-way F(3,36) = 4.72, p = 0.007, η_p^2 = 0.28; Supplementary §S3.10) and maintained through the 3–5 week retention visit. Paired retention tests (Session 5 vs 6) showed no significant change for any group (Supplementary Table S1).

The scalar ERD contrast between C3 SMR and C4 SMR was null (d = 0.34, BF01 = 3.13; Supplementary Table S1), indicating equivalent magnitude at the training electrode. However, a cluster permutation test on the Active–Sham time-frequency difference confirmed a significant cluster at C3 (p = 0.002; Supplementary Figure S4) that did not survive correction at C4 (min p = 0.13; Supplementary Figure S5). Site specificity thus appears in the topographic distribution rather than the scalar magnitude, consistent with a deeper-generator prediction in which thalamocortical ERD would be more spatially diffuse (see Discussion). ERD stability across sessions is shown in Supplementary Figure S6. The effect was consistent across individuals, with minimal overlap between active and sham distributions (Supplementary Figure S7). Detailed contrasts, LME coefficients, and retention statistics are in Supplementary S3.

3.2 ERD–P2 Double Dissociation

The P2 ERP component at C3 showed a large group effect (F(3,33) = 5.02, p = 0.006, η_p^2 = 0.31; Figure 3), stable across sessions (Group × Session interaction n.s.), paralleling the ERD's temporal profile. The non-significant interaction is consistent with a between-group site/frequency signature of the operant mechanism rather than a training-induced morphological change, though at these sample sizes we cannot definitively distinguish stable-from-Session-1 from slowly-converging-across-sessions (Session 1 one-way F(3,36) = 0.94, p = 0.43; Supplementary §S3.9). But the groups driving the two effects were different.

The ERD was strongest in C3 Beta (d = −2.38 vs sham). The P2 was most affected by C3 SMR, which produced a substantially *smaller* mean P2 amplitude at the trained electrode than sham (d = −1.33, BF01 = 0.10, strong Bayesian evidence for a real difference), while C3 Beta (d = −0.27) and C4 SMR (d = 0.19) did not differ from sham. Site-specific suppression was striking: C3 SMR vs C4 SMR d = −1.44 (p_adj = 0.026). The frequency-specificity contrast trended in the *opposite* direction from the ERD (C3 SMR vs C3 Beta d = −1.07, p_adj = 0.079 for the P2; cf. the ERD frequency-specificity contrast p_adj = 0.039, opposite direction). Raw group means at C3: C3 SMR +0.30 µV, C3 Beta +1.43 µV, C4 SMR +1.79 µV, Sham +1.64 µV (n = 37 with complete session data; df = 3, 33). Three participants had missing session-level ERP data (108 S3, 110 S5, 131 S3).

These opposing patterns form a double dissociation: **Beta → stronger ERD, preserved P2; C3**

SMR → weaker ERD, suppressed P2 at the trained site. Both are contingency-dependent (absent in sham), but they fractionate spectral power modulation (ERD) from the phasic cortical response to the reward stimulus (P2). Per-subject ERD and P2 amplitudes were uncorrelated ($r = -0.06$, $p = 0.79$, $n = 24$ active; Supplementary S3.7), confirming these are independent mechanisms rather than different views of a single process. Two active subjects had missing Session 6 resting-state data, reducing n to 22 for the resting-state prediction analysis reported in §3.3. P50 was null across all sites; N1 showed a session main effect with no group effect, likely reflecting habituation to the reward tone over repeated sessions. Topographic maps confirming lateralized ERD are in Supplementary Figure S3.

3.3 Three-Timescale Plasticity

Three timescales of plasticity were dissociable, each with a different specificity profile.

Immediate (trial-locked). The ERD was stable across sessions (LME interaction n.s.) and retained at 3–5 weeks. Per-subject ERD magnitude during training predicted eyes-closed resting-state power change at follow-up ($r = 0.54$, $R^2 = 0.29$, $p = 0.009$, $n = 22$ active subjects; Figure 4).

Within-session (minutes). A single training session transiently increased EC alpha power in all groups ($d = 0.5$–1.1), including sham at borderline ($d = 0.53$, $p = 0.053$). ERD magnitude did not predict within-session shift size ($r = -0.09$, $p = 0.67$; Supplementary S4.2), dissociating the operant signal from the transient rebound (Supplementary S4.1).

Across-session (days to weeks). Pre-session EC alpha trajectories diverged by group (Supplementary Figure S10). An LME growth curve confirmed SMR-specific accumulation: C3 SMR × Session $\beta = 1.44$ ($p = 0.004$), C4 SMR × Session $\beta = 1.24$ ($p = 0.012$); C3 Beta and sham showed flat or declining slopes (raw values in Supplementary S4.3, LME coefficients in S4.4). Eyes-closed alpha at follow-up was significantly higher in both SMR groups versus sham (C3 SMR $d = 0.97$, $p_adj = 0.012$; C4 SMR $d = 0.78$, $p_adj = 0.028$). C3 Beta showed no change ($d = -0.002$). Eyes-open was null for all groups (Supplementary S4.5–S4.6).

3.4 Summary of Mechanistic Profiles

The full battery reveals distinct mechanistic profiles (Table 1).

Measure	C3 SMR	C3 Beta	C4 SMR	Sham
ERD vs sham (d)[a]	−0.80	**−2.38**	−1.12	—
P2 vs sham (d)[a]	**−1.33**	−0.27	0.19	—
ERD–P2 correlation	$r = -0.06$, $p = 0.79$ ($n = 24$ active)			
Within-session EC alpha (d)	0.93	1.10	0.94	0.53
EC alpha growth (LME β)	**1.44**	−0.49	**1.24**	0.39
EC alpha S6-S1 vs sham (d)	**0.97**	−0.002	**0.78**	—

[a] Negative d indicates Active < Sham; for P2, this corresponds to a smaller evoked peak in the active group.

The within-session row is the key: all groups show comparable transient alpha increases (d = 0.53–1.10), yet only the SMR groups consolidate those shifts into cumulative growth. No single measure captures the full mechanistic picture; the implications for circuit-level interpretation are developed in Discussion §4.3.

3.5 Sensitivity Analysis

To assess whether results depended on the choice of preprocessing pipeline, we repeated the primary ERD analyses using an ICA-based pipeline (extended Infomax ICA with ICLabel auto-rejection of non-brain components; see Supplementary S1). Trial retention was comparable across pipelines (minimal median = 583, ICA median = 572 per session). Despite similar trial counts, the minimal pipeline produced a somewhat larger Active vs Sham effect (d = −1.23, p = 0.0004) than ICA (d = −1.03, p = 0.0025), and the LME Group effect was larger ($\eta_p^2 = 0.38$ vs 0.26; Supplementary Figure S11). Delorme (2023) benchmarked phase-locked ERP detectability and showed that trial rejection generally reduces detection power rather than improving it; we extend this argument to induced ERD. The minimal pipeline's statistical artifact rejection preserves more of the reward-locked ERD signal, whereas ICA component removal reduces detection power. All primary findings (frequency specificity, contingency dependence, and the ERD–P2 dissociation) converge across both pipelines.

Figures

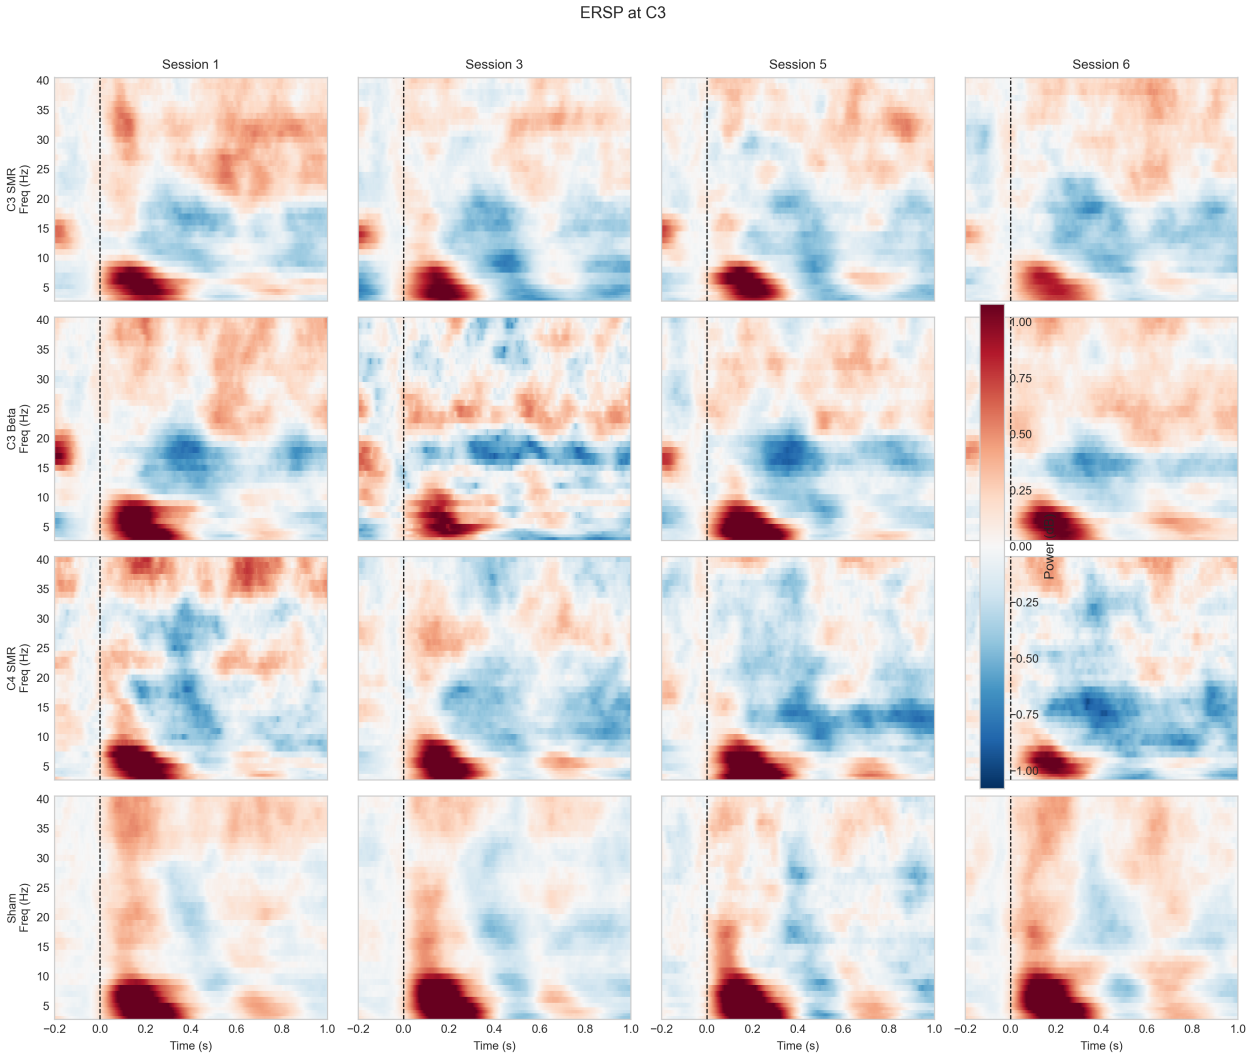

Figure 1: **Figure 1.** Grand-average ERSP heatmaps at C3 (4 groups × 4 sessions). Active groups show frequency-specific ERD in the reward band; sham shows theta ERS only.

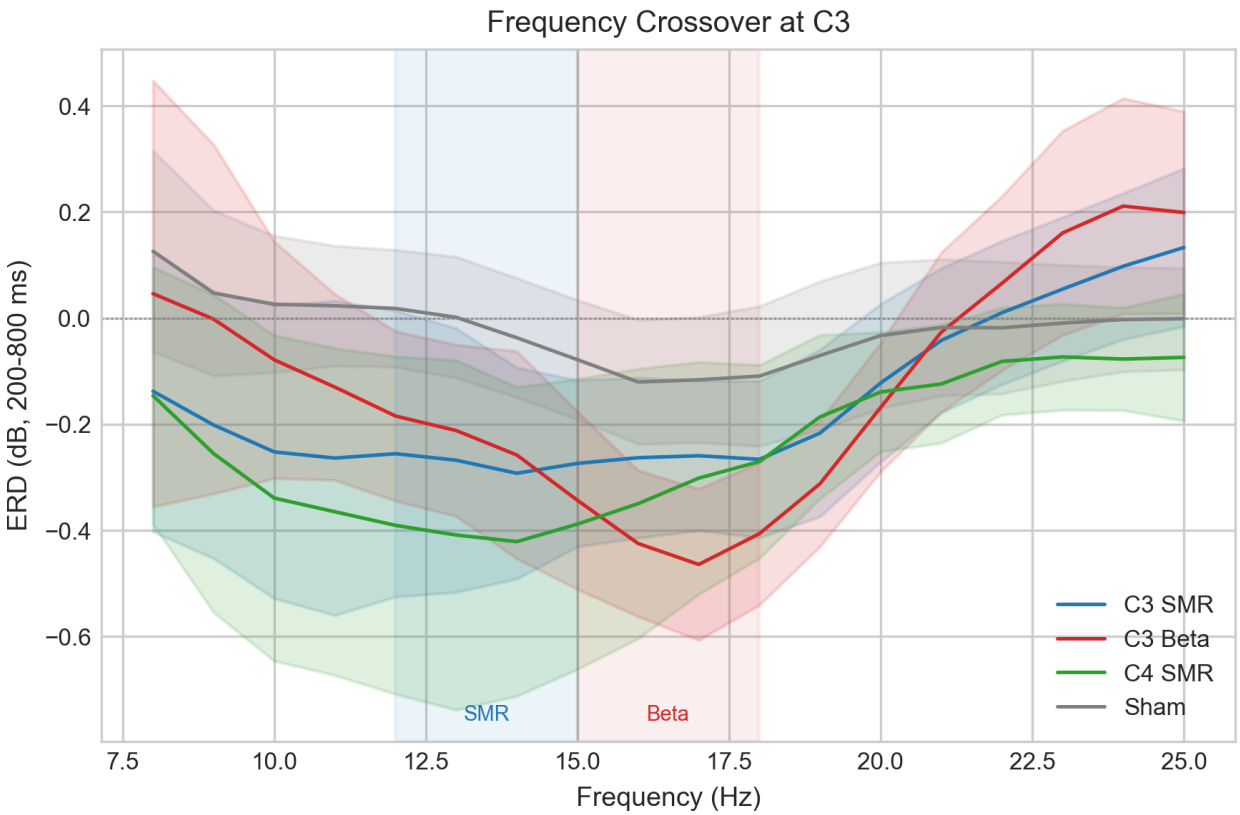

Figure 2: **Figure 2.** Frequency crossover at C3. C3 SMR ERD peaks at 12–15 Hz; C3 Beta peaks at 15–18 Hz; sham is flat.

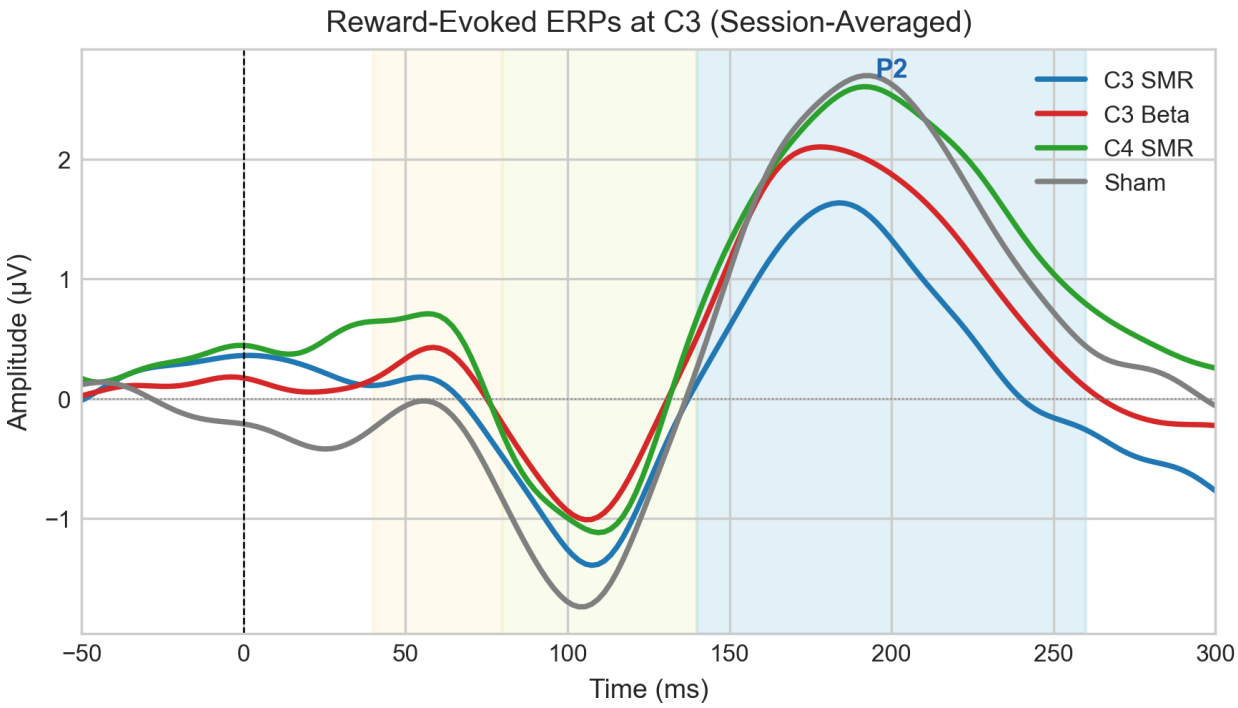

Figure 3: **Figure 3.** Session-averaged reward-evoked ERPs at C3. The P2 (shaded) is suppressed in C3 SMR relative to the other groups, including sham, the opposite of the ERD pattern (which is strongest in C3 Beta). Full waveforms at C3, C4, Pz in Supplementary Figure S9.

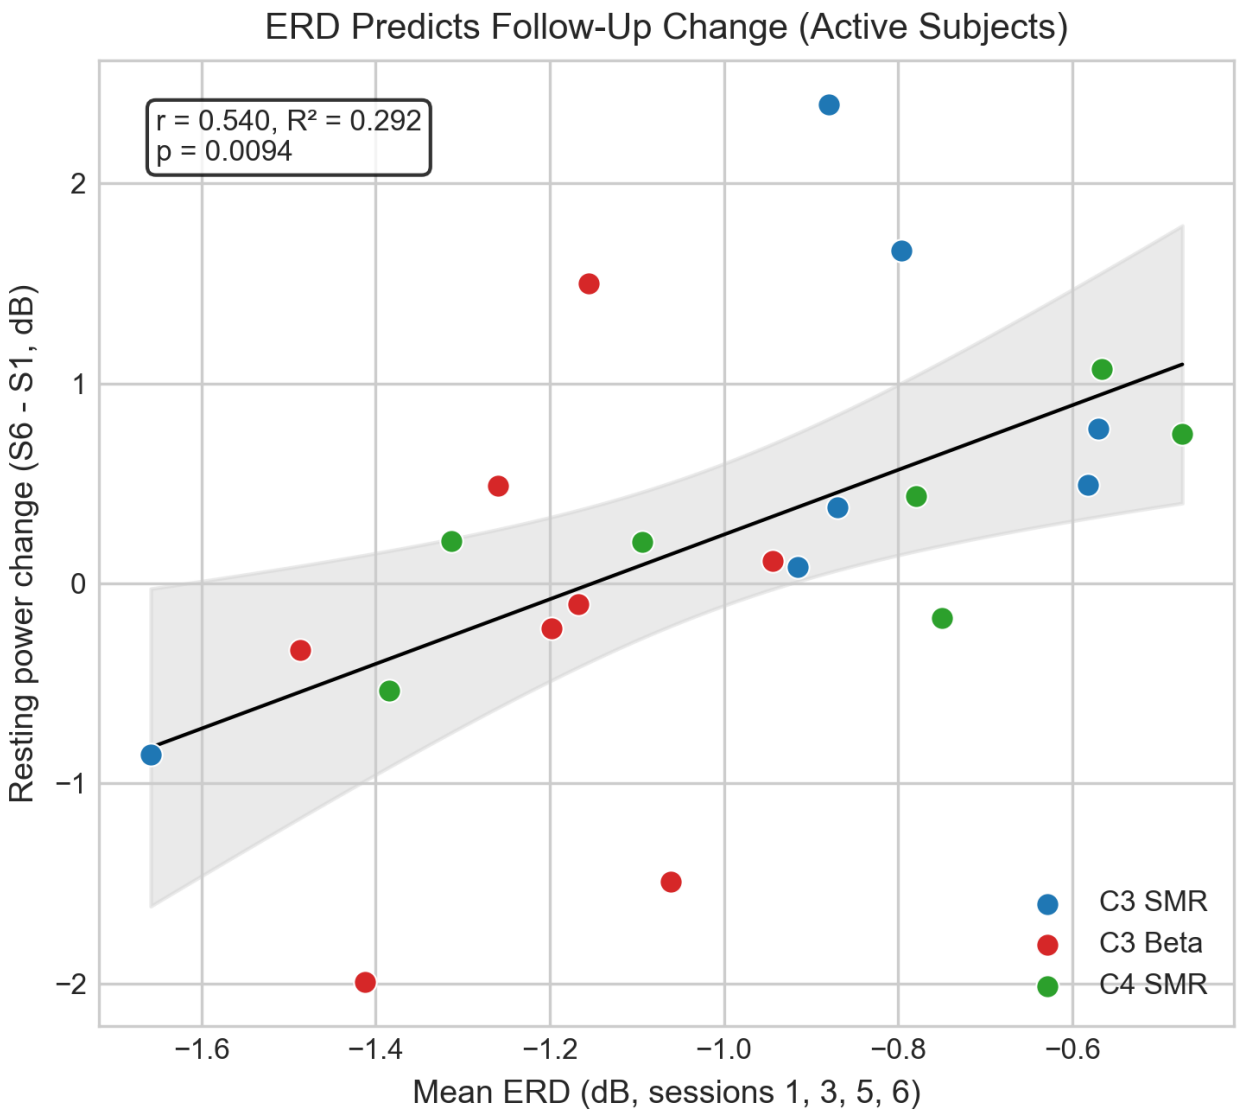

Figure 4: **Figure 4.** Per-subject ERD during training vs eyes-closed resting-state change at follow-up ($r = 0.54$, $p = 0.009$, $n = 22$).

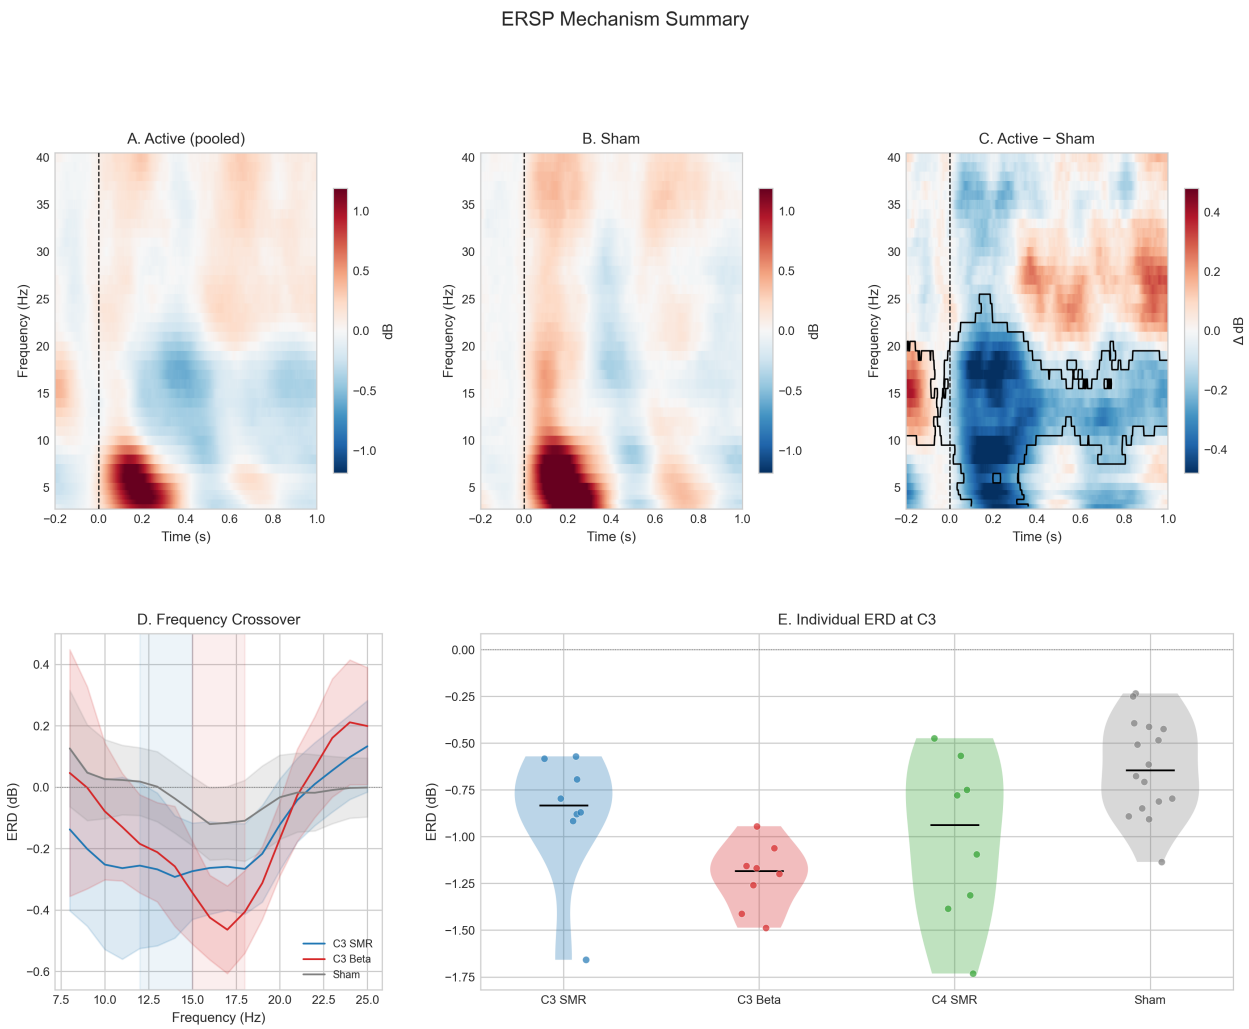

Figure 5: **Figure 5.** Multi-panel composite summary of the reward-locked ERSP mechanism.

4. Discussion

No single finding in this study would be conclusive at these sample sizes. The evidentiary weight derives from convergence: four independent patterns (frequency-specific ERD, contingency dependence, the ERD–P2 double dissociation, and SMR-specific resting-state growth) align within the same 40 subjects under double-blind conditions. Each eliminates a different class of alternative explanation; together, they constrain interpretation to a degree that no single contrast could achieve alone.

4.1 The Sensory-Operant Dissociation and the Sham

The brain's reward-locked response contains two distinct signals: a theta-band ERS reflecting a mixture of auditory-cortical processing and mid-frontal feedback/reward theta (Cavanagh & Frank, 2014), present in all groups including sham, and a frequency-specific ERD in the operantly conditioned band that appears only when reward delivery is genuinely contingent on oscillatory power.

The active-placebo sham is critical to this dissociation. The sham produced non-specific effects (theta ERS, behavioral practice) but failed to produce reward-band ERD, with Bayesian evidence favoring the null (BF01 > 3 for all sham reward-band contrasts). This clarifies a current debate about whether sham NF is "partially active" (Supplementary S5.3): it is (the sham produces real sensory-evoked responses), but the contingency-dependent operant mechanism is absent.

Why should reviewers trust this dissociation? The sham signal is pre-recorded EEG with artifact merging from the participant's live signal, scaled to match real-time amplitude. The reward criterion operates on this sham signal, not on the participant's actual brain activity. There is no pathway by which the participant's oscillatory state can influence reward delivery, so reward-band ERD *cannot* emerge through contingency. Theta ERS, by contrast, is a sensory response to the auditory beep itself; it requires only that the participant hear the tone, which all groups do. The theta-present, ERD-absent pattern in sham is exactly what the design predicts. This is the specific signature the Kerson, Sherlin, and Davelaar (2025) framework predicts for non-contingent reinforcement schedules (sensory processing preserved, operant signature absent), and empirically constrains the debate about what "active placebo" sham actually contains.

4.2 Frequency Specificity and the Double Dissociation

The frequency crossover provides the most direct test of specificity. As reported in §3.1, two groups trained at the same electrode with reward bands 3 Hz apart produced ERD peaking at different frequencies, a result no non-specific account readily accommodates. This answers Dessy et al. (2020), whose single-reward-frequency SMR design at Cz could not contrast competing trained frequencies within the same paradigm (Supplementary S5.1).

The ERD–P2 dissociation extends the frequency crossover. Both measures are contingency-dependent (absent in sham) but fractionate by group in opposite directions: C3 Beta shows the strongest ERD without changing the P2, while C3 SMR shows moderate ERD with a markedly

suppressed P2 at the trained electrode. The ERD indexes spectral power change (phase-locked and induced combined); the P2 indexes the phasic evoked response to the reward tone. That C3 SMR training reduces the P2 specifically at the trained site suggests a mechanism that damps the phasic cortical response to the operantly-trained stimulus, a gating profile distinct from the cortical desynchronization engaged by Beta training.

4.3 Circuit Depth Determines the Outcome Profile

The double dissociation, combined with the resting-state data, aligns with a model in which the reward frequency indexes circuit depth, and circuit depth shapes which measures respond and whether effects persist.

SMR at 12–15 Hz is classically tied to thalamocortical relay circuitry implicated in sleep-spindle generation and in sensorimotor gating (Sterman, 1977, 2000; Steriade, McCormick & Sejnowski, 1993; De Gennaro & Ferrara, 2003). If SMR training engages this relay-based thalamocortical mode, a plausible downstream consequence at the trained site is *attenuation* of the phasic evoked response to the operantly-relevant stimulus. The C3 SMR result matches this prediction: moderate ERD at the trained electrode combined with a markedly reduced P2, at that electrode but not at C4. Repeated engagement of this thalamocortical loop may reshape tonic resting-state dynamics. The gating analogy has a logical limit: the reward beep is the operantly-relevant stimulus, not an irrelevant one, so the relay-state interpretation must either confine the gating to the brief reward-delivery window (testable by ITC; §4.5) or be understood as a correlate of consolidated training rather than a sleep-spindle-style gating mechanism. An alternative reading that also fits the data is consolidation-driven attentional automatization: as subjects internalize the SMR-to-reward mapping, the reward beep becomes less phasically salient and the evoked P2 diminishes, a mechanism that would co-predict the SMR-specific across-session alpha growth. ITC (phase versus power) is the adjudicating measurement; the present data do not distinguish the two.

Beta at 15–18 Hz is supported by substantial local-cortical mechanisms, including transient beta events arising from pyramidal-cell dendritic integration (Sherman et al., 2016; Shin et al., 2017), with thalamocortical and basal-ganglia contributions also documented (Bonaiuto et al., 2021). A local-cortical ERD without the relay-gating component is consistent with the C3 Beta profile: strongest ERD, preserved P2. But the absence of relay architecture means the local-cortical loop does not propagate into tonic resting-state reorganization in the present data.

The ERD alone is therefore an incomplete index of neurofeedback mechanism. A design measuring only the ERD would rank Beta highest; measuring only the P2 would rank C3 SMR highest in magnitude of change (in the suppression direction). Both measures are needed to differentiate the circuits engaged by each protocol.

We emphasize that the present scalp-level data motivate but do not directly test this circuit-level account. The predictions outlined below provide concrete falsification criteria: source localization should show thalamic involvement for SMR-band ERD and purely cortical generators for beta-band ERD; ITC in the reward band should be higher for SMR than Beta groups (Supplementary S5.2).

4.4 A Consolidation Account

The data reveal a plasticity arc across three timescales. The ERD is the immediate mechanism, present at full strength from Session 1, stable across sessions, maintained at the 3–5 week retention visit. Each training session also produces a transient EC alpha increase (d = 0.5–1.1) in all groups including sham, representing a non-specific post-training rebound. ERD magnitude does not predict within-session shift size ($r = -0.09$, $p = 0.67$). This null result argues against the possibility that stronger ERD produces larger transient shifts that then consolidate. Instead, the ERD and the transient rebound appear to be independent processes, with specificity arising at the consolidation step. The LME growth curve shows that only SMR groups' baselines ratchet upward across sessions (C3 SMR × Session $\beta = 1.44$, $p = 0.004$; C4 SMR $\beta = 1.24$, $p = 0.012$), while C3 Beta, despite within-session alpha shifts as large as any active group (d = 1.10), shows a *declining* baseline. Each session perturbs the system; the data suggest that only the deeper circuits engaged by SMR training consolidate the perturbation.

The arc: **immediate ERD** (trial-locked, stable) → **transient within-session rebound** (universal, non-specific) → **cumulative baseline drift** (SMR-specific, persisting to 3-5 week follow-up, d = 0.97 vs sham). We propose the following mapping onto Davelaar's stages, recognizing that our temporal grain (single-trial ERD, across-session growth, multi-week consolidation) differs from Davelaar's original division: Stage 1 (implicit basal-ganglia-mediated operant learning; Davelaar, 2018) maps onto the immediate ERD; Stage 2 (circuit-level consolidation of repeated perturbation; Davelaar, 2018) maps onto the across-session accumulation seen only in SMR groups; Stage 3 (interoceptive homeostasis; Kerson, Sherlin & Davelaar, 2025) addresses the transition from externally reinforced to self-sustained regulation that the present data cannot test but that the framework predicts. The clinical implication is that the key variable is not session count but *consolidability*, the capacity of each session's transient shift to persist offline. While this sequence is supported by the present dissociations, additional longitudinal work is needed to confirm the causal structure. If confirmed, the model implies that protocol selection should be guided not by session count alone but by the consolidability of the circuit engaged.

4.5 Limitations, Future Directions, and Broader Implications

Group sizes were modest (8–16). However, the intensive protocol means each participant contributes approximately 3,000 reward-locked trials across training sessions, providing substantial within-subject power. Effect sizes should be interpreted as upper bounds likely to attenuate in larger samples; the critical finding is the convergence of four independent patterns within the same participants, not the magnitude of any single contrast. The data were collected in 2010–2011 and originally analyzed descriptively (Hill, 2012); the present reanalysis applies a pre-specified modern framework to existing data. Fixed frequency bands (12–15 Hz, 15–18 Hz) were used rather than IAF-anchored bands; if anything, IAF-anchored protocols should produce stronger effects, making these results a conservative test. Five training sessions address mechanism, not clinical dose-response.

Four extensions follow: ITC analysis to test the phase-locking prediction; source localization to

test thalamic involvement; single-trial ERSP for responder subtypes; and a prospective replication with pre-registered effect sizes.

In response to pre-submission feedback, we tested for a classical Sterman-style post-reinforcement parietal alpha synchronization (Pz/P3/P4, 8–12 Hz, 800–1500 ms post-reward; Supplementary §S3.8) and did not detect one: pooled Active vs Sham d = +0.16, BF01 = 5.20 (moderate evidence for H0), with all three per-group contrasts against sham likewise null (BF01 = 3.30–4.66) and a whole-plane cluster permutation at Pz (3–20 Hz × 0–1500 ms) yielding no significant post-reward cluster. The operant signature in this paradigm therefore appears to be carried by the reward-band ERD itself rather than by a distinct post-reinforcement synchronization burst; whether PRS is genuinely absent in reward-contingent human NF or whether our paradigm masks it (the simultaneous picture-grid reveal dominates the parietal alpha response in this window; Supplementary §S3.8) and Sterman's original cat/human operant schedules had no such concurrent visual event, is an open question we cannot resolve here.

The reward consisted of a simultaneous auditory beep and picture-grid reveal, so the P2 group effects reflect combined auditory and visual evoked responses rather than a pure auditory P2. This methodological limit, shared with any paradigm delivering multimodal reward at a single time point, means the P2 group differences could partly reflect differential visual-attentional processing; future designs with temporally separated auditory and visual feedback would isolate the modality-specific contribution.

This work provides mechanistic evidence the field needs to move beyond the current efficacy stalemate. The reward-locked ERSP framework provides a principled basis for protocol design: if the mechanism is contingency-dependent, then optimizing contingency fidelity, not session count or gamification, should be the primary target.

References

Arnold, L.E., Lofthouse, N., Hersch, S., et al. (2013). EEG neurofeedback for ADHD: Double-blind sham-controlled randomized pilot feasibility trial. *Journal of Attention Disorders*, 17(5), 410–419.

Arnold, L.E., Arns, M., Barterian, J., et al. (2021). Double-blind placebo-controlled randomized clinical trial of neurofeedback for attention-deficit/hyperactivity disorder with 13-month follow-up. *Journal of the American Academy of Child & Adolescent Psychiatry*, 60(7), 841–855.

Arns, M., de Ridder, S., Strehl, U., Breteler, M., & Coenen, A. (2009). Efficacy of neurofeedback treatment in ADHD: the effects on inattention, impulsivity and hyperactivity: a meta-analysis. *Clinical EEG and Neuroscience*, 40(3), 180–189.

Bonaiuto, J.J., Little, S., Neymotin, S.A., et al. (2021). Laminar dynamics of high amplitude beta bursts in human motor cortex. *NeuroImage*, 242, 118479. https://doi.org/10.1016/j.neuroimage.2021.118479

Cavanagh, J.F., & Frank, M.J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, 18(8), 414–421. https://doi.org/10.1016/j.tics.2014.04.012

Cortese, S., Ferrin, M., Brandeis, D., et al. (2016). Neurofeedback for attention-deficit/hyperactivity disorder: meta-analysis of clinical and neuropsychological outcomes from randomized controlled trials. *Journal of the American Academy of Child & Adolescent Psychiatry*, 55(6), 444–455.

Davelaar, E.J. (2018). Mechanisms of neurofeedback: a computation-theoretic approach. *Neuroscience*, 378, 175–188.

De Gennaro, L., & Ferrara, M. (2003). Sleep spindles: an overview. *Sleep Medicine Reviews*, 7(5), 423–440. https://doi.org/10.1053/smrv.2002.0252

Delorme, A. (2023). EEG is better left alone. *Scientific Reports*, 13, 2372.

Dessy, E., Mairesse, O., van Puyvelde, M., Cortoos, A., Neyt, X., & Pattyn, N. (2020). Train your brain? Can we really selectively train specific EEG frequencies with neurofeedback training. *Frontiers in Human Neuroscience*, 14, 22.

Egner, T., & Gruzelier, J.H. (2001). Learned self-regulation of EEG frequency components affects attention and event-related brain potentials in humans. *NeuroReport*, 12(18), 4155–4159.

Egner, T., & Gruzelier, J.H. (2004). EEG biofeedback of low beta band components: frequency-specific effects on variables of attention and event-related brain potentials. *Clinical Neurophysiology*, 115(1), 131–139.

Fetz, E.E. (1969). Operant conditioning of cortical unit activity. *Science*, 163(3870), 955–958.

Fetz, E.E. (2007). Volitional control of neural activity: implications for brain–computer interfaces. *Journal of Physiology*, 579(3), 571–579.

Gramfort, A., Luessi, M., Larson, E., et al. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7, 267.

Grandchamp, R., & Delorme, A. (2011). Single-trial normalization for event-related spectral decomposition reduces sensitivity to noisy trials. *Frontiers in Psychology*, 2, 236.

Hill, A.R. (2012). *Measuring & Modulating Hemispheric Attention.* Doctoral dissertation, University of California, Los Angeles. https://escholarship.org/uc/item/55c35139

Kerson, C., Sherlin, L.H., & Davelaar, E.J. (2025). Neurofeedback, biofeedback, and basic learning theory: revisiting the 2011 conceptual framework. *Applied Psychophysiology and Biofeedback.* https://doi.org/10.1007/s10484-025-09756-4

Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Research Reviews*, 29(2–3), 169–195. https://doi.org/10.1016/S0165-0173(98)00056-3

Makeig, S. (1993). Auditory event-related dynamics of the EEG spectrum and effects of exposure to tones. *Electroencephalography and Clinical Neurophysiology*, 86(4), 283–293.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190.

Muraoka, T., Iwama, S., & Ushiba, J. (2024). Neurofeedback-induced desynchronization of sensorimotor rhythm elicits pre-movement downregulation of intracortical inhibition. *Imaging Neuroscience.*

Pfurtscheller, G., & Lopes da Silva, F.H. (1999). Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clinical Neurophysiology*, 110(11), 1842–1857.

Ros, T., Munneke, M.A., Ruge, D., Gruzelier, J.H., & Rothwell, J.C. (2010). Endogenous control of waking brain rhythms induces neuroplasticity in humans. *European Journal of Neuroscience*, 31(4), 770–778.

Ros, T., Enriquez-Geppert, S., Zotev, V., et al. (2020). Consensus on the reporting and experimental design of clinical and cognitive-behavioural neurofeedback studies (CRED-nf checklist). *Brain*, 143(6), 1674–1685.

Salansky, N., et al. (2024). Sensori-motor neurofeedback improves inhibitory control and induces neural changes: a placebo-controlled, double-blind, event-related potentials study. *International Journal of Clinical and Health Psychology.*

Salansky, N., et al. (2025). Enhancing attentional processing through sensorimotor neurofeedback training: evidence from a placebo-controlled, double-blind, event-related potentials study. *NeuroImage.*

Schönenberg, M., Wiedemann, E., Schneidt, A., et al. (2017). Neurofeedback, sham neurofeedback, and cognitive-behavioural group therapy in adults with attention-deficit hyperactivity disorder: a triple-blind, randomised, controlled trial. *The Lancet Psychiatry*, 4(9), 673–684.

Sherlin, L.H., Arns, M., Lubar, J., et al. (2011). Neurofeedback and basic learning theory: implications for research and practice. *Journal of Neurotherapy*, 15(4), 292–304.

Sherman, M.A., Lee, S., Law, R., et al. (2016). Neural mechanisms of transient neocortical beta rhythms: converging evidence from humans, computational modeling, monkeys, and mice. *Proceedings of the National Academy of Sciences*, 113(33), E4885–E4894. https://doi.org/10.1073/pnas.1604135113

Shin, H., Law, R., Bhatt, S., Moore, C.I., Sherman, M.A., & Jones, S.R. (2017). The rate of transient beta frequency events predicts behavior across tasks and species. *eLife*, 6, e29086. https://doi.org/10.7554/eLife.29086

Sitaram, R., Ros, T., Stoeckel, L., et al. (2017). Closed-loop brain training: the science of neurofeedback. *Nature Reviews Neuroscience*, 18(2), 86–100.

Sorger, B., Scharnowski, F., Linden, D.E.J., Hampson, M., & Young, K.D. (2019). Control freaks: Towards optimal selection of control conditions for fMRI neurofeedback studies. *NeuroImage*, 186, 256–265.

Steriade, M., McCormick, D.A., & Sejnowski, T.J. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science*, 262(5134), 679–685. https://doi.org/10.1126/science.8235588

Sterman, M.B. (1977). Sensorimotor EEG operant conditioning: experimental and clinical effects. *Pavlovian Journal of Biological Science*, 12(2), 63–92.

Sterman, M.B. (2000). Basic concepts and clinical findings in the treatment of seizure disorders with EEG operant conditioning. *Clinical Electroencephalography*, 31(1), 45–55.

Tallon-Baudry, C., Bertrand, O., Delpuech, C., & Pernier, J. (1996). Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human. *Journal of Neuroscience*, 16(13), 4240–4249.

Thibault, R.T., Lifshitz, M., & Raz, A. (2016). The self-regulating brain and neurofeedback: Experimental science and clinical promise. *Cortex*, 74, 247–261. https://doi.org/10.1016/j.cortex.2015.10.024

Thibault, R.T., & Raz, A. (2017). The psychology of neurofeedback: Clinical intervention even if applied placebo. *American Psychologist*, 72(7), 679–688. https://doi.org/10.1037/amp0000118

Van Doren, J., Arns, M., Heinrich, H., Vollebregt, M.A., Strehl, U., & Loo, S.K. (2019). Sustained effects of neurofeedback in ADHD: a systematic review and meta-analysis. *European Child & Adolescent Psychiatry*, 28(3), 293–305.

Wyrwicka, W., & Sterman, M.B. (1968). Instrumental conditioning of sensorimotor cortex EEG spindles in the waking cat. *Physiology & Behavior*, 3(5), 703–707. https://doi.org/10.1016/0031-9384(68)90139-X

Acknowledgments

Data were collected at the Department of Psychology, University of California, Los Angeles, under IRB approval. The author thanks the original dissertation committee and the participants who contributed their time.

Conflict of Interest

Peak Brain Institute, with which the author is affiliated, provides clinical neurofeedback services. The data analyzed in this study were collected during the author's doctoral research at UCLA and are unrelated to Peak Brain Institute's clinical operations. This reanalysis was conducted independently without external funding.

Data and Code Availability

Analysis code is available at https://github.com/DocSalamandyr/Hill_2026_ERSP-neurofeedback-analysis. Derived data (ERSP matrices, ERP averages, resting-state PSD, and statistical outputs) are deposited at https://doi.org/10.5281/zenodo.19555777. The pre-specified analysis plan is posted to OSF. Raw EEG recordings (BioSemi BDF) are available from the corresponding author subject to a data use agreement, as the original informed consent did not include provisions for unrestricted public sharing.