Losses resulting from deliberate exploration trigger beta oscillations in frontal cortex

Boris V. Chernyshev*, Kristina I. Pultsina, Vera D. Tretyakova, Aleksandra S. Miasnikova, Andrey O. Prokofyev, Galina L. Kozunova, Tatiana A. Stroganova

Center for Neurocognitive Research (MEG Center), Moscow State University of Psychology and Education, Moscow, Russia

Corresponding author:

Boris V. Chernyshev,
Center for Neurocognitive Research (MEG Center),
Moscow State University of Psychology and Education,
Moscow, Russia;
b_chernysh@mail.ru

Abstract

We examined the neural signature of directed exploration by contrasting MEG beta (16-30 Hz) power changes between disadvantageous and advantageous choices in the two-choice probabilistic reward task. Both types of choices were made when our participants learned the probabilistic contingency between choices and their outcomes, i.e., acquired the inner model of choice value. Therefore, rare disadvantageous choices might serve exploratory, environment-probing purposes.

The study brought two main findings. Firstly, decision making leading to disadvantageous choices took more time and evidenced greater large-scale suppression of beta oscillations than its advantageous alternative. Additional neural resources required by disadvantageous decisions strongly suggest their deliberately explorative nature. Secondly, an outcome of disadvantageous and advantageous choices had qualitatively different impact on feedback-related beta oscillations. Only losses, but not gains, resulting from the disadvantageous choice were followed by late beta synchronization in frontal cortex.

Our results are consistent with the role of frontal beta oscillations in the stabilization of neural representations for selected behavioral rule when exploratory strategy conflicts with value-based behavior. Punishment for exploratory choice being congruent with its low value in the reward history is more likely to strengthen, through punishment-related beta oscillations, the representation of its competitor – the inner utility model.

Introduction

In professional and everyday life – for example, in search of a new job – people often face an exploration-exploitation dilemma. Searching for a new opportunity at the expense of the familiar ones can be costly, as tangible outcomes of exploration may only come in the distant future and no profits are warranted. On the other hand, sticking to the choices that proved valuable in the past discourages subjects from pursuing learning and development. A purposeful search for information at the expense of immediate reward is a characteristic feature of directed exploration, which became topic of considerable interest the last decade[1-7]. Although switching between exploitative and deliberate explorative strategies constitutes a core foundation of human adaptive behavior, its neural underpinnings remain largely unknown (but see [5,8,9]).

People actively interrogate their environment using 'question-and-answer' strategies[8]. Yet this is not an easy challenge: choosing a risky and potentially non-rewarding option involves a conflict, since participants have to inhibit the tendency to choose a rewarded safe alternative[10,11]. This conflict arises between at least two simultaneously active competing internal models, or 'task sets'[9,12] – one being a predominant response tendency (exploitation), and the other – its conscious alternative (exploration). Our recent pupillometric study lends support to this assumption[13]: we found that such explorative choices compared to exploitative ones are accompanied by larger pupil dilation and longer decision time. We speculated that this state of conflict supposedly entails an increase in the degree of processing required to make the deliberately explorative decisions. Given that such a selection was shown to carry greater load on information processing and/or attentional capacities[8], it should make a deliberately explorative decision more difficult than the decision to choose an advantageous option. To validate this explanation, one needs to compare cortical activation accompanying the decision process that leads to either exploitative or to explorative choices.

Importantly, any action outcome – even a negative one – has epistemic value since it reduces internal uncertainty[14]. Since deliberate explorative choices are in fact quests for information, they likely involve altered evaluation of the outcome received after one's choice. Indeed, for example, even modest variations in the context – such as changing the subject's beliefs about the probability of rewards and punishments – can alter cortical responses to feedback signals (e.g., Marco-Pallares, et al. [15]).

Here, we sought to evaluate the impact of directed exploration on cortical signatures of decision-making and feedback evaluation processes. To this end, we

investigated the dynamics of magnetoencephalographic β-band oscillations (16-30 Hz) while contrasting directed explorative choices and exploitative ones in the probabilistic two-alternative task. Well-documented properties of β oscillations make them an informative and accessible measure allowing us to pursue the above goals.

Based on numerous findings on β power suppression (event-related desynchronization, β-ERD) as a measure of cortical activation strength during decision-making[16-18], we predicted that decision-making leading to disadvantageous directed explorative choices would be accompanied by greater β-ERD than decision-making resulting in exploitative choice.

Concerning β-oscillations over the frontal cortex (frontal event-related synchronization, frontal β-ERS) following action outcome evaluation, the previous literature offers two accounts. The first account relies on the evidence specifically linking this oscillatory activity with reward, especially with an unexpected reward in human subjects[15,19]. This view is supported by the findings on greater frontal β-ERS for positive compared to negative feedback[20-22] and for larger compared to smaller reward magnitudes[15]. It was proposed that feedback-related frontal β-ERS reflects a signal involved in reward processing and underlying memory formation by signaling which events are better than expected (for review, see Marco-Pallares, et al. [23]). Based on the 'reward account' one can predict that in our study, a reward will be followed by frontal β-ERS compared to punishment regardless of the type of the choice and will be relatively greater for explorative compared with exploitative choices – simply because for explorative choices a positive reward is less expected.

Still, the dependence of frontal β-oscillations on some specific task conditions is difficult to explain using the straightforward 'reward account'. For example, there is evidence that in some tasks, a negative feedback signal (or omission of a positive one) contributes to the built-up of the inner state promoting the increase in frontal β-oscillations[24,25].

A complementary account, which accumulates numerous findings on frontal β-synchronization in both humans and animals, proposes that these oscillations play a role in re-activation and strengthening of the cognitive set, which is more likely to result in a favorable outcome and will be used in future actions[26-28]. The term ‘cognitive set' here is very similar to the term ‘internal model’, which, according to K. Friston’s predictive coding theory, accumulates experience-based knowledge and generates prior beliefs that are used to make inferences/predictions regarding incoming external events[29,30]. From a predictive

coding perspective, for the selection of a to-be-strengthened cognitive set/internal model, both positive and negative outcomes may be used as just pieces of evidence weighted against recent reward history that formed the inner utility model[31,32]. Depending on how reliable these different sources of information are, the selection process may lend the greatest weight to the inner model that ensures optimal guidance of future actions.

The 'predictive coding' account produces predictions different from those expected on the basis of the 'reward' account. In this case, appearance of frontal β-ERS will depend on how the two competing models, that were simultaneously active during decision-making regarding this choice, will be re-evaluated after receiving feedback. A reward for a rare, explorative, deliberately disadvantageous choice, which is discrepant with the previous reward history, hardly provides a piece of evidence reliable enough to ensure strengthening of the model governing this choice. Punishment for such a choice, on the other hand, being congruent with its low value in the reward history, is more likely to strengthen the representation of its competitor – the predominant utility model. This might result in the occurrence of frontal β-ERS after a punishment but not a reward for an explorative choice. In other words, directed explorative choices may be expected to differ from exploitative ones by a weak impact of a positive outcome, and a strong impact of a negative outcome.

Therefore, by evaluating the predictions of the two accounts regarding frontal β-ERS in deliberate explorative and exploitative choices, we hoped to get a better insight into the nature of brain processes resulting from the appraisal of feedback information during switching from exploitation to exploration strategy.

To sum up, we here hypothesize that disadvantageous choices in the two-alternative probabilistic reward task represent a deliberate switching from the exploitative to the directed explorative strategy that triggers a conflict with a prepotent response tendency. The modulation of brain activity in the β-range could exhibit specific features during both decision-making and feedback-related processes associated with such choices. First, we sought to test the prediction that a putatively more effortful arrival to the exploitative decision would involve greater cortical activation indexed by deeper β-ERD compared with exploitative decisions. Second, by contrasting post-feedback modulations of β-oscillations for the exploratory and exploitatory choices, we aimed to test the predictions made based on the two accounts of how prior knowledge about the probability of reward and punishment influences feedback processing in the brain.

Results

Behavioral data

In the current study, we were interested in directed exploration. We argue that disadvantageous choices can be ascertained as directed exploration only on condition that participants had acquired an appropriate utility model of the task (otherwise, they probably made most of their choices randomly). Thus, only trials satisfying the learning criteria were included in the analyses (for behavioral statistics, see Supplementary Table S1).

Factor Choice type was significant for response time (RT) (LMM, $F_{(3,3519)} = 17.09$, $p<0.001$) (Supplementary Table S2). RT was significantly longer for the LP condition compared with all the other conditions, and RTs for pre-LP and post-LP conditions were significantly longer compared with the HP condition (Tukey HSD, all p's<0.01) (Fig. 1B).

Figure 1

Data-Driven Approach to Choose Time Windows and Areas of Interest for the Analysis of Beta-Band Power Modulations

We expected that (1) brain processes related to decision-making will be different depending on the upcoming choice type (e.g., exploration vs. exploitation), and that (2) feedback sign will be processed differentially depending on the choice type that has been committed (contrasts of losses vs. gains will be different for exploratory and exploitatory trials).

In order to locate the expected effects in time and space, we combined a hypothesis-driven approach with a data-driven one. At the first step, we averaged response-locked β-band power data over consecutive 200-ms windows, in each sensor independently. We subjected these data to statistical analysis using linear models with mixed effects (LMM), with Choice type, Feedback and their interaction included in the model as fixed effects. Significance was FDR-corrected for 102 sensors x 17 time windows. We visualized significant sensors for each factor and their interaction on topographical maps represented in Fig. 2.

Figure 2

In order to test our hypotheses, we focused on the three time intervals with corresponding sensor topography. A profound lasting effect of the Choice type factor was present before and during the behavioral response (Fig. 2, upper row). Since we were

interested in the brain processes related to decision-making leading to initiation of the behavioral response, we restricted the further analysis to the pre-response part of the effect within -900 – -300 ms relative to the response onset, thus aiming to avoid a possible overlap with effects related to stimulus processing and movement execution.

Feedback-related effects were not uniform, thus we selected two separate effects: (1) an early feedback-related effect of Choice type around the feedback onset at 100 – 500 ms after feedback onset (Fig. 2, upper row); (2) a late feedback-related interaction of Choice type x Feedback at 500 – 900 ms after feedback onset (Fig. 2, bottom row). The latter effect apparently evidenced differential feedback sign processing depending upon the choice type.

For data visualization on topographical maps, event-related spectral perturbations and timecourses, see Supplementary Fig. S1-S4.

Decision-making effect

To further analyze the dependence of β power upon the choice type, we averaged β power across the cluster of sensors demonstrating the respective effect within the pre-response time window (Fig. 2). Strong and highly significant suppression of β power relative to baseline was observed during the decision-making period (Fig. 3A). A source-level analysis revealed the widespread β-ERD originating from the left frontal, left pericentral areas, as well as from the partial, occipital (both lateral and medial aspects), and medial temporal areas (Fig. 3E).

Figure 3

At the sensor level, the effect of choice type (Supplementary Table S3) was due to deeper β-ERD in LP trials compared with all the other trial types (HP, pre-LP and post-LP) (Tukey HSD, all p's<0.001) (Fig. 3B). At the source level, the significant difference between LP and HP conditions was localized to left pericentral, most of the partial areas and large parts of the occipital cortex (both on lateral and medial aspects), as well as posterior cingulum (Fig. 3F).

Thus, as we expected, decision-making regarding exploratory choices involved deeper suppression of β oscillations compared with exploitatory decisions.

In order to find out whether β suppression during decision-making is indeed related to greater cognitive load (which may explain delayed responding), we estimated regression between response time and β-band power. Indeed, we observed a highly significant negative

regression slope (LMM statistics, β_0 = 4.82, β_1 =-2.33, p< 0.001, R^2 = 0.41) (Fig. 3D). Negative slope indicates that the slower decision-making was, the more β power was suppressed during decision-making, thus suggesting greater cognitive load during slower decision-making.

Early effect of choice type on the feedback processing

Starting around the time of the visual feedback onset and sustaining up to 500 ms, the effect of Choice type on β power was observed in the occipito-parietal cluster of sensors (Fig. 2, Fig. 4A and Fig. 4D). For all choice types, early post-feedback β power in this cluster was suppressed relative to baseline (Fig. 4 A, B). Note that the concurrent rise of β power (β-ERS) in the anterior sensors (Fig. 4A) was not influenced by Choice type and started long before the feedback onset (Supplementary Fig. S1), thus hardly representing the early feedback-related response. The effect of Choice type (Supplementary Table S3) was mainly driven by stronger β-ERD on LP trials compared with all the other trial types (Tukey HSD, all p's <0.05) (Fig. 4B). Notably, the feedback sign (whether a participant gained or lost) did not influence β-ERD in any trial type (Supplementary Table S3). Source-level analysis of the LP vs. HP contrast revealed that the LP-related increase in the β-ERD mainly occupied cortical regions on the lateral and medial surface of the occipital and posterior part of parietal and occipital lobes (Fig. 4E).

Figure 4

Thus, as we predicted, exploratory choices boosted processing of the visual feedback signal, no matter whether their outcome was good or bad for a participant.

Late feedback-related effect

A significant Choice type x Feedback interaction distinguished the late feedback-related time interval (500 – 900 ms after feedback onset) (Fig. 2). This implies that, as we expected, β power modulations by positive and negative feedback depended upon choice type. Yet, surprisingly, this dependency manifested itself relatively late throughout the course of feedback-related brain response. For all trial types pooled together, we observed a significant β-ERS in the anterior areas, while significant β-ERD persisted over the posterior regions (Fig. 5A and D). Therefore, we proceeded with analyses in sensor subclusters manifesting significant ERD and ERS separately.

Figure 5

Posterior subcluster.

Similarly to the early post-feedback interval, a significant β-ERD relative to baseline was observed in the posterior sensors (Fig. 5A, Fig. 5D, Supplementary Table S3). For the posterior subcluster, LMM analysis confirmed the significance of Choice type, Feedback, and Choice type x Feedback interaction (Fig. 5A, Supplementary Table S3). The posterior β-ERD was deeper for LP trails compared to HP trials (Fig. 5B), i.e., it manifested the same, although less pronounced, LP-related bias as we observed in the early post-feedback interval. Choice type x Feedback interaction was mainly explained by significantly greater posterior β-ERD after positive compared with negative feedback on LP trials (Tukey HSD, $p < 0.001$), but not for the other choice types (Fig. 5B, C). On the contrary, on exploitatory HP trials sensitivity of posterior β-ERD to positive and negative feedback was opposite - with β-ERD being slightly yet significantly larger after losses compared with gains (Tukey HSD, $p = 0.03$) (Fig. 5C).

Thus, selectively increased posterior β-ERD indexed greater neural activation after losses for serial exploitatory HP choices and after wins in the exploratory LP choices. Given the higher probability of wins after HP choices and losses after LP choices, these findings suggested that the activation strength depended on the outcome expectancy, with greater activation accompanying late processing of the less probable outcome.

Anterior subcluster.

Sensors of the anterior subcluster exhibited significant β-ERS (Fig. 5A, Fig. 5D, Supplementary Table S3). Factors Choice type, Feedback and interaction Choice type x Feedback were all highly significant (Supplementary Table S3). The effect of the Choice type can be explained by the fact that anterior β-ERS was significantly greater for HP trials compared with pre-LP and LP trials (Fig. 5B, C). Since a choice of the advantageous stimulus has been rewarded more frequently than its risky alternative, this finding seemed to be in line with the role of the reward in anterior β-ERS induction for the specific, possibly, advantageous choices.

However, examination of Choice type x Feedback interaction by splitting the data by feedback sign (Fig. 5C) ruled out this explanation. Negative and positive feedback in the HP trials equally contributed to the anterior β-ERS, while in explorative LP trials β-ERS was driven exclusively by negative feedback (Fig. 5B). Note that in the LP trials, positive feedback

was even followed by β power suppression relatively to baseline – instead of its increase (Fig. 5B).

In addition, trials immediately preceding explorative LP choices (pre-LP), despite their superficial similarity with the advantageous HP choices, showed a significantly reduced ERS to feedback (Tukey HSD, $p < 0.001$), regardless of the feedback sign (Fig. 5C). Yet, the same advantageous choice in the post-LP trials, similarly to the LP trials themselves, induced greater anterior β-ERS to losses compared with gains (Tukey HSD, $p < 0.001$, Fig. 5C).

The high sensitivity of anterior β-ERS to punishment in the explorative choices (and to a lesser degree in post-explorative choices) strongly suggests that this specific form of oscillations was not driven by the sign of the feedback signal itself, but rather depended on its context, which, in our task, varies between highly expected (serial HP trials) and that involving conflict with the accumulated reward history (LP trials).

These findings on β power sensitivity to feedback sign on exploratory trials agree with the predictions made from the 'predictive coding' account of anterior β-ERS outlined in the Introduction.

However, none of the predictions can account for the finding that increased sensitivity of anterior β-ERS to punishment would sustain in the post-LP advantageous choices when a participant came back to the non-conflicting exploitative strategy. This finding implied that in addition to the accumulated history of trials and errors, there could be a short-term impact of the previous trial outcome on the anterior β-ERS. We explored this possibility in the following section.

The impact of the feedback presented on the previous trial

In order to evaluate the impact of the previous feedback (i.e., that was presented on the previous trial) on the late feedback-related effect evoked by losses and gains on the current trial, we ran LMM analysis with Previous feedback fixed factor additionally included into the model (Table 2, Fig. 6). Specifically, we were interested only in ERS effects (i.e. the anterior subcluster) and two choice types – HP and LP.

Figure 6

Triple interaction Choice type x Previous feedback x Feedback was significant ($F_{1, 3304} = 6.0$, $p = 0.015$). For exploitatory HP choices (Fig. 6, left panel), β-ERS induced by a current gain was slightly but significantly enhanced if the previous similarly advantageous choice was

expectedly rewarded rather than punished (Tukey HSD, $p = 0.02$), whereas β-ERS induced by a current loss remained unaffected by the previous outcome.

On the contrary, for the exploratory LP choices, the outcome of the previous advantageous choice significantly affected β-ERS in the case of a current loss (Tukey HSD, $p = 0.04$) but not a current gain (Tukey HSD, $p >> 0.05$) (Fig. 6, right panel). It should be stressed that gains for LP choices produced no β-ERS at all.

A strong β-ERS bound to anticipated monetary losses ensuing from objectively disadvantageous choices compared with gains, was highly significant if the previous advantageous choice was rightfully rewarded (Tukey HSD, $p < 0.001$), while after punishment it was at the margin of significance (Tukey HSD, $p = 0.04$).

In other words, for HP and LP choices similarly, anterior β-ERS was further enhanced if not only the outcome of the current choice but also the outcome of the previous one matched the prediction/expectation derived from the utility model acquired from the previous experience.

Discussion

Using a simple two-choice probabilistic reinforcement learning task, which was learnt by participants through trial and error, we investigated β-band power modulations related to rare disadvantageous choices conflicting with the acquired inner utility model.

The study brought two main findings. First, decision making leading to exploratory choices took more time and evidenced greater large-scale suppression of β oscillations than its advantageous alternative. Crucially, the magnitude of a large-scale β suppression, when a subject formed his/her decision reliably, predicted decision costs (with response time taken as a measurable proxy of this internal variable) on a single trial basis. Given multiple evidence associating the strength of β suppression with greater cognitive and attentional efforts[16-18], this finding strongly suggests that such a choice requires additional resources to overcome the internal utility model favoring the advantageous alternative and strongly supports the hypothesis of its deliberately explorative nature[13,33].

Second, the most intriguing finding concerns β oscillations induced by feedback received after deliberately explorative choices. The main difference between directed explorative and exploitative choices was the impact of their negative versus positive outcome on β oscillations arising during post-evaluation processes (Fig. 5). When a choice was made within a sequence of similarly advantageous choices (HP choice), both monetary gains and losses produced a late synchronization of frontal β oscillations. These feedback-

related β oscillations were reduced to a baseline if a participant was going to switch from such an advantageous strategy toward the explorative one (pre-LP choice). Remarkably, only losses – but not gains – were accompanied by reliable frontal β synchronization after committing directed exploration (LP choice).

Therefore, in our task, there was no consistent relationship between the presence or absence of frontal β oscillations and the reward or punishment for a committed choice. Therefore, merely maintaining the 'reward-induced' state, as the 'reward' hypothesis implies[23], is not sufficient to generate enhanced β power. Although the context-dependency of punishment-related frontal β was previously mentioned by some authors (e.g., Yaple, et al. [25]), this conclusion was derived from the findings obtained in different research paradigms. Here, for the first time, we directly contrasted frontal β resulting from the outcomes of two choice types committed by the same participants within the same task on the trial-by-trial basis. The solely punishment-related β synchronization observed for exploratory choices cannot be explained or predicted by previously suggested contextual factors such as successful punishment avoidance[34], task difficulty[35], emotional salience of the feedback signal[24], reward expectancy[19,36-39]. It is possible that the bias of the literature toward a role of β oscillations in positive feedback processing derives in part from experimental biases, e.g., nonequivalence of the relevance of positive and negative outcomes for the future behavior (as in HajiHosseini and Holroyd [40]) and/or a complete lack of response-feedback contingency, i.e., a purely random structure of rewards and punishments for either choice[19,36].

Being contradictory to the 'pure reward' hypothesis, our results are instead consistent with the role of β oscillations in the brain state resulting from the matching processes between the outcome of a current choice and accumulated experience of gains and losses, i.e., the inner utility model.

Both positive and negative feedback signals used in our task could induce an increase in β power with a latency of five to nine hundred milliseconds. However, late β ERS clearly depended on the congruency between the feedback received and the prediction based on the previously acquired utility model. Specifically, when a participant's choice complies with the inner utility model (i.e., if the advantageous stimulus with the higher objective reward probability was repeatedly chosen), late frontal β increases regardless of the actual feedback that a participant has received. Given the probabilistic nature of the model, the precision of its prediction is not very narrow (see e.g., Yon and Frith [31]), and a single loss resulting from the objectively advantageous choice does not lead to a tangible

mismatch. By contrast, explorative choices committed against the predominant value-based response tendency trigger frontal β synchronization only after a monetary loss, which provides strong evidence for validity of the inner utility model. In this case, a reward rather than punishment for a preceding advantageous choice, being also congruent with an inner model's prediction, increases β-ERS induced by a current loss (Fig. 6). In other words, for both exploitative and explorative choices, maximal frontal β-ERS accompanies those feedback signals that reliably confirm the validity of the existing value-based model – even if the current choice was driven by its exploratory alternative. This type of β oscillations thus appears as a categorical information and not as a scalar measure of discrepancy between outcome and expectation.

The relatively long latency (500-900 ms) of the feedback-related frontal β-ERS places strong constraints on its potential functional significance. It is unlikely that the differential response of β-ERS to reward and punishment in our task is causally involved in the feedback valence evaluation, as it occurs substantially after the 'feedback processing' time window, which is usually between 200 and 500 ms relative to feedback onset[15,37]; but see also ref.[25], and ref.[41], see also ref.[24] for a later latency of differential effect of feedback valence on frontal β under specific task requirements).

We offer here that there is a more sophisticated account of feedback-related frontal β, which is focused on the post-evaluative memory-based replay of 'cognitive set'[26] or rule representation[28,42,43] with its possible role in the strengthening of the re-activated cortical representation. In line with this idea, the cortical topography of the late frontal anterior β-ERS (Fig. 5) reveals involvement of a widely distributed set of cortical areas that are mainly located on the dorsal and medial surface of the frontal cortex including anterior and dorsal cingulum, supplementary motor area, inferior frontal gyrus, anterior insula and ventromedial frontal cortex. All these structures constitute cortical nodes of the cingulate-basal ganglia-limbic circle, strongly implicated in the stabilization of neural representations for selected behavioral rule through reverberating activity in the β-range[28,44].

Our proposal here is closely related to this idea, yet places it in a context of exploration-exploitation dilemma. We suggest that during outcome evaluation of a deliberately explorative choice, a representation of the previously rewarded choice strategy (utility model) is maintained in the memory buffer alongside with that for the voluntary decision, which shifts current choice priorities in favor of information seeking, i. e. testing the alternative model, or 'task set'[9,12]. Such competing representations are thought to be re-activated whenever the current task induces conflict in processing preceding response

selection, but also in feedback evaluation[45]. The evaluation process lends the greatest weight to the external feedback that is most reliable, i.e., matches the previously accumulated experience (see e.g., Yon and Frith [31]), aka the inner utility model in our experimental settings. The further stabilization of the winning utility model is reflected in the punishment-related β oscillations for alternative explorative choice, and it would compromise future repetitions of the explorative behavioral program in the unchanging environment. A rare gain for such choices does not provide sufficient evidence to strengthen either model; therefore, it does not induce β synchronization. This leaves a window of opportunity for 'set shifting', i.e., for adjustment of value-based model in the future, which is vital in case environment changes. We predict that using the same task in the 'reversal learning' context will reveal selectively increased reward-related β oscillations for a number of explorative choices made after the reversal of previously acquired stimulus–response associations.

To conclude, our results suggest that emergence of late punishment-related β-ERS for deliberately explorative choices reflects a replay of a value-based model representation that possesses greater predictive capacity in an unchanging environment and will be strengthened for future use.

Material and methods

Participants

Sixty healthy participants (30 males, 30 females), aged 25.6±4.8 years (M ± SD) and having no neurological disorders or severe visual impairments were enrolled in the study. This study was conducted in accordance with the Declaration of Helsinki; the study protocol was approved by the Ethics Committee of Moscow State University of Psychology and Education. All participants signed written informed consent before the experiment.

Procedure

A modified probabilistic learning task[46,47] was rendered as a computer game consisting of six blocks. Each pair of visual stimuli comprised two images of the same Hiragana hieroglyph (1.54 × 1.44°) rotated at two different angles (Fig. 1A). Participants were instructed to make two-alternative choices between two stimuli presented on the screen simultaneously (a new pair was used in each block). Within each block, one of the stimuli was associated with greater probability of better outcomes (monetary gains on 60-70% of

trials). The other stimulus was associated with greater probability of worse outcomes (monetary gains on 40-30% of trials, with the rest of the trials leading to losses).

Participants' behavioral responses were recorded with the use of a handheld MEG-compatible fiber optic button response pad (CurrentDesigns, Philadelphia, PA, USA). Visual feedback was presented for 500 ms on each trial, with a delay of 1000 ms after the behavioral response. The intertrial interval varied from 700 to 1400 ms in a quasi-random order (flat distribution). Each of six experimental blocks comprised 40 trials with the overall duration of approximately 5 minutes per block. The experiment was carried out using the Presentation 14.4 software (Neurobehavioral systems, Inc., Albany, CA, USA).

MEG data acquisition and preprocessing

MEG was recorded in a magnetically shielded room (AK3b, Vacuumschmelze GmbH, Hanau, Germany). We used the dc-SQUID Neuromag VectorView system (Elekta-Neuromag, Helsinki, Finland). Sampling rate was 1000 Hz, and the passband was 0.03–330 Hz.

We applied the temporal signal space separation method (tSSS)[48] to the raw data using MaxFilter (Elekta Neuromag software). We removed cardiac artifacts and artifacts related to eye movements from continuous data using the ICA method implemented in MNE-Python software. We extracted response-locked epochs from -1750 to 2750 ms relative to the button press.

Trial selection

In the current study, we primarily addressed directed exploration, i.e. voluntary goal-directed commission of disadvantageous choices purposefully aimed at gaining information[4,6]. We assumed that directed exploration can be ascertained only if a subject has acquired a utility model of the task. Thus, within each block, we extracted only those trials that met the following learning criteria: (1) these trials went after four consecutive choices of the advantageous stimulus and until the end of the block; (2) the percentage of advantageous choices within the selected trials was above 65%.

A recent pupillometric study[13] revealed that advantageous choices that immediately preceded and immediately followed exploratory choices significantly differed from the advantageous choices committed within the periods of continuous exploitation. Thus, we used the following four levels of Choice type factor (Fig. 1; see also Supplementary Table S1): (1) the 'high-payoff' choice (HP) – the advantageous choice that was preceded and followed by advantageous choices (such sequences of uninterrupted advantageous choices were

considered as a stable 'exploitative' preference for the advantageous stimulus); (2) the advantageous choice preceding the 'low-payoff' choice (pre-LP); (3) the 'low-payoff' choice (LP) – the disadvantageous 'explorative' choice that was preceded and followed by advantageous choices; (4) the advantageous choice following the 'low-payoff' choice (post-LP).

In addition, we classified trials depending upon the feedback received on the current trial (Feedback factor: 'loss' and 'gain') and on the previous trial (Previous feedback factor: 'punishment' and 'reward').

Several participants were excluded from the analyses because of the lack of advantageous choices satisfying the learning criterion (2 participants) and due to violations of the behavioral procedure during recording (2 participants). Additionally, in order to keep the design more balanced, we excluded participants that had no trials in at least one of the eight conditions Choice type x Feedback (16 participants). Thus, statistical analyses were done in 40 participants.

We evaluated the dependence of response time (RT) upon choice type using the following LMM model, which included Choice type factor (four levels: HP, pre-LP, LP, post-LP), Previous feedback factor, i.e. outcome of the previous trial (two levels: punishment and reward) and their interactions as the fixed effects, and Subject as a random factor:

$$\text{Response time} \sim \text{Choice type} * \text{Previous feedback} + (1 \mid \text{Subject}), \quad (1)$$

where Response time is log-transformed RT.

Time-frequency analysis of the MEG data at the sensor level

We restricted our analyses of the β frequency range to 16-30 Hz[26,49]. We performed all the analyses using MNE Python open-source software[50] and custom scripts in Python. At sensor level, all time-frequency analyses were performed at single-trial level. Discrete prolate spheroidal sequences (DPSS) multitaper analysis was applied to estimate spectra; this method is efficient for small numbers of trials[51]. For baseline correction, we extracted baseline values from stimulus-locked epochs in the HP condition, -350 to -50 ms relative to fixation cross onset.

In order to reduce the dimensionality of the data, we combined sensors by adding time-frequency data within pairs of orthogonal planar gradientometers. Thus, at sensor level, we report power estimates for 102 such combined gradientometers.

Statistical analysis

We applied linear mixed models (LMM) using the 'lmer' function[52] available in the lme4 package for R. We checked the data for normality of the distribution of residuals with the normal quantile plot. To check the data for homoscedasticity we drew the residual plot.

We used linear mixed effects models (LMM) at a single-trial level rather than repeated measures ANOVA at the grand-average level for several reasons. The LMM method is robust to imbalanced designs – thus, missing data can be handled without listwise deletion of cases; moreover, different numbers of trials per condition are less problematic than in traditional ANOVA applied at the grand-average level (e.g., Kliegl, et al. [53]). LMM is suitable for large numbers of repeated measurements per participant, thus making LMM especially appropriate to analyze data from individual trials: such a statistical approach allows accounting for intertrial variability[54,55].

We used a data-driven approach towards selection of time windows and locations. Thus, the analysis was performed in two steps, as described below.

At the first step, we aimed to identify effects:

(1) We averaged power of β oscillations over consecutive 200-ms windows.

(2) We applied a linear model with mixed effects (LMM) to single-trial data at each of 17 time points within each of 102 sensors independently. The model included Choice type factor (four levels: HP, pre-LP, LP, post-LP), Feedback factor (two levels: loss and gain) and their interactions as fixed effects, and Subject as a random factor:

$$\text{Beta power} \sim \text{Choice type} * \text{Feedback} + (1 \mid \text{Subject}) \qquad (2)$$

(3) We applied correction for multiple comparisons using the false discovery rate (FDR) method[56] to 1734 observations (17 time points x 102 sensors).

(4) Next, guided by our hypotheses and by visual inspection of the topographical plots with significant sensors indicated, we selected time windows, during which spatial distribution of significant sensors comprised distinct yet relatively stationary patterns.

(5) For further analyses, we considered only those sensors that were persistently significant throughout all timeframes within the selected time windows.

(6) In order to make sure that the effects to be analyzed represented either ERD or ERS, we retained for further analyses only those sensors that were significant against baseline for all conditions pooled (t-test, $p < 0.05$, uncorrected). If within the given time window some sensors manifested ERD while other displayed ERS, we assigned two separate clusters of sensors correspondingly.

At the second step, we performed an in-depth data analysis. We collapsed MEG data across all timeframes within a corresponding time window and across sensors within the corresponding cluster (see above), thus obtaining one value of β-band power per each trial.

For the analysis of effects preceding feedback onset, we applied to single-trial data the following LMM with the Choice type factor as the fixed effect, and Subject as a random factor:

$$\text{Beta power} \sim \text{Choice type} + (1 \mid \text{Subject}) \quad (3)$$

For effects following feedback onset, we additionally included Feedback factor:

$$\text{Beta power} \sim \text{Choice type} * \text{Feedback} + (1 \mid \text{Subject}) \quad (4)$$

At this stage, we did not focus on the main effects and their interaction, which expectedly mirrored the findings obtained at the previous stage (see Supplementary Table S3). Instead, we aimed to find out which statistical contrasts created the effects found at the previous stage. Thus, we complemented the basic LMM analysis with post-hoc pairwise comparisons using Tukey honestly significant difference (HSD) test (emmeans package version 1.6.0)[57].

In order to ascertain whether the power of β oscillations preceding the behavioral response reflects greater cognitive load during decision-making, we evaluated regressions at a single-trial level. We used the following LMM model:

$$\text{Beta power} \sim \text{Response time} * \text{Choice type} * \text{Previous feedback} + (1 \mid \text{Subject}) \quad (5)$$

where Beta power is the power of β oscillations averaged over the time window within corresponding significant sensors, and Response time is log-transformed RT.

Finally, we aimed to evaluate how the immediate history of punishments and rewards in the previous trial affected the brain response to the feedback on the current trial. We addressed the influence of the previous feedback because it affects behavioral and pupillometric measures on the next trial[13]. An additional factor Previous feedback corresponding to the outcome of the previous trial (two levels: punishment and reward) was included into the LMM model:

$$\text{Beta power} \sim \text{Feedback} * \text{Previous feedback} * \text{Choice type} + (1 \mid \text{Subject}) \quad (6)$$

Source level analysis

Participants underwent MRI scanning with a 1.5T Philips Intera system (voxel size 1 x 1 x 1 mm, T1-weighted images). We reconstructed single-layer (inner skull) boundary-element models of cortical gray matter using individual structural MRIs, with a watershed

segmentation algorithm[58] implemented in FreeSurfer 4.3 software (Martinos Center for Biomedical Imaging, Charlestown, MA, United States).

We used the 'mne.setup_source_space' function implemented in the MNE-Python open-source software[50] to create the source space. In order to compute the source space for the β-band power (16-30 Hz), we used 'mne.minimum_norm.source_band_induced_power' function (implemented in the MNE-Python open-source software), with Morlet wavelets (8 cycles per wavelet).

We used a standardized low-resolution brain electromagnetic tomography (sLORETA) localization method[59]. To illustrate spatial cortical localization of β-ERD and β-ERS, we rendered baseline-corrected power of β oscillations for all conditions collapsed. We used the averaged data to illustrate spatial cortical localization of the pairwise contrasts. We assessed significance of the difference in the averaged β-band power, for each vertex independently (pairwise t-test, $p<0.05$, uncorrected).

A more detailed description of the methods is provided in the Supplementary information.

Acknowledgements

This study was supported by Russian Science Foundation (project # 20-18-00252).

We are grateful to Ksenia E. Sayfulina for developing scripts for statistical analysis, and to Ksenia E. Sayfulina and Anastasia Yu. Nikolaeva for invaluable help with data collection.

Author contributions statement

G. L. K. and T.A.S. contributed to conception and design of the study. G. L. K. and A.O.P. conducted the experiments. B.V.C., K.I.P., V.D.T., and A.S.M. analyzed the results and prepared illustrations. B.V.C. and T.A.S. wrote the first draft of the manuscript. B.V.C., K.I.P., V.D.T., A.S.M., A.O.P., G.L.K., T.A.S. wrote sections of the manuscript. All authors reviewed the manuscript.

Additional Information

The authors declare no competing financial or non-financial interests.

Data availability

The datasets generated and/or analyzed during the current study are available from https://figshare.com/projects/Losses_resulting_from_deliberate_exploration_trigger_beta_oscillations_in_frontal_cortex/149792.

Code availability

The code (custom scripts) generated during the current study are available from the corresponding author on reasonable request.

References

1 Payzan-LeNestour, E. & Bossaerts, P. Do not bet on the unknown versus try to find out more: estimation uncertainty and "unexpected uncertainty" both modulate exploration. *Front. Neurosci.* **6** (2012).

2 Schulz, E. & Gershman, S. J. The algorithmic architecture of exploration in the human brain. *Curr. Opin. Neurobiol.* **55**, 7-14 (2019).

3 Schwartenbeck, P. *et al.* Computational mechanisms of curiosity and goal-directed exploration. *Elife* **8**, e41703 (2019).

4 Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A. & Cohen, J. D. Humans Use Directed and Random Exploration to Solve the Explore-Exploit Dilemma. *J. Exp. Psychol. Gen.* **143**, 2074-2081 (2014).

5 Zajkowski, W. K., Kossut, M. & Wilson, R. C. A causal role for right frontopolar cortex in directed, but not random, exploration. *Elife* **6**, e27430 (2017).

6 Wilson, R. C., Bonawitz, E., Costa, V. D. & Ebitz, R. B. Balancing exploration and exploitation with information and randomization. *Curr. Opin. Behav. Sci.* **38**, 49-56 (2021).

7 Dubois, M. *et al.* Human complex exploration strategies are enriched by noradrenaline-modulated heuristics. *Elife* **10**, ARTN e59907 (2021).

8 Gottlieb, J. & Oudeyer, P.-Y. Towards a neuroscience of active sampling and curiosity. *Nat. Rev. Neurosci.* **19**, 758-770 (2018).

9 Koechlin, E. in *The Cognitive Neurosciences, 6th ed.(Cambridge, MA: MIT Press, 2020)* (eds D. Poeppel, G. R. Mangun, & M. S. Gazzaniga) Ch. 38, 451-468 (MIT Press, 2020).

10 Cogliati Dezza, I., Yu, A. J., Cleeremans, A. & Alexander, W. Learning the value of information and reward over time when solving exploration-exploitation problems. *Sci. Rep.* **7**, 16919 (2017).

11 Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876-879 (2006).

12 Domenech, P., Rheims, S. & Koechlin, E. Neural mechanisms resolving exploitation-exploration dilemmas in the medial prefrontal cortex. *Science* **369**, eabb0184 (2020).

13 Kozunova, G. L. *et al.* Pupil dilation and response slowing distinguish deliberate explorative choices in the probabilistic learning task. *Cogn. Affect. Behav. Neurosci.* **22**, 1108-1129 (2022).

14 Parr, T. & Friston, K. J. Uncertainty, epistemics and active inference. *J. R. Soc. Interface* **14**, 20170376 (2017).

15 Marco-Pallares, J. *et al.* Human oscillatory activity associated to reward processing in a gambling task. *Neuropsychologia* **46**, 241-248 (2008).

16 Pavlova, A. A. *et al.* Effortful verb retrieval from semantic memory drives beta suppression in mesial frontal regions involved in action initiation. *Hum. Brain. Mapp.* (2019).

17 Scharinger, C., Soutschek, A., Schubert, T. & Gerjets, P. Comparison of the working memory load in n-back and working memory span tasks by means of EEG frequency band power and P300 amplitude. *Front. Hum. Neurosci.* **11**, 6 (2017).

18 Tafuro, A., Ambrosini, E., Puccioni, O. & Vallesi, A. Brain oscillations in cognitive control: A cross-sectional study with a spatial stroop task. *Neuropsychologia* **133**, 107190 (2019).

19 Mas-Herrero, E., Ripolles, P., HajiHosseini, A., Rodriguez-Fornells, A. & Marco-Pallares, J. Beta oscillations and reward processing: Coupling oscillatory activity and hemodynamic responses. *Neuroimage* **119**, 13-19 (2015).

20 Novikov, N. A. *et al.* Slow and Fast Responses: Two Mechanisms of Trial Outcome Processing Revealed by EEG Oscillations. *Front. Hum. Neurosci.* **11**, Artn 218 (2017).

21 Cohen, M. X., Wilmes, K. A. & van de Vijver, I. Cortical electrophysiological network dynamics of feedback learning. *Trends Cogn. Sci.* **15**, 558-566 (2011).

22 Weiss, S. & Mueller, H. M. "Too many betas do not spoil the broth": the role of beta brain oscillations in language processing. *Front. Psychol.* **3**, Artn 201 (2012).

23 Marco-Pallares, J., Munte, T. F. & Rodriguez-Fornells, A. The role of high-frequency oscillatory activity in reward processing and learning. *Neurosci. Biobehav. Rev.* **49**, 1-7 (2015).

24 Leicht, G. *et al.* Relationship between Oscillatory Neuronal Activity during Reward Processing and Trait Impulsivity and Sensation Seeking. *PLoS One* **8**, e83414 (2013).

25 Yaple, Z. *et al.* Power of Feedback-Induced Beta Oscillations Reflect Omission of Rewards: Evidence From an EEG Gambling Study. *Front. Neurosci.* **12**, Artn 776 (2018).

26 Engel, A. K. & Fries, P. Beta-band oscillations--signalling the status quo? *Curr. Opin. Neurobiol.* **20**, 156-165 (2010).

27 Miller, E. K., Lundqvist, M. & Bastos, A. M. Working Memory 2.0. *Neuron* **100**, 463-475 (2018).

28 Brincat, S. L. & Miller, E. K. Prefrontal Cortex Networks Shift from External to Internal Modes during Learning. *J. Neurosci.* **36**, 9739-9754 (2016).

29 Friston, K. J. & Kiebel, S. Predictive coding under the free-energy principle. *Philos. T. Roy. Soc. B.* **364**, 1211-1221 (2009).

30 Summerfield, C. & de Lange, F. P. Expectation in perceptual decision making: neural and computational mechanisms. *Nat. Rev. Neurosci.* **15**, 745-756 (2014).

31 Yon, D. & Frith, C. D. Precision and the Bayesian brain. *Curr. Biol.* **31**, R1026-R1032 (2021).

32 Kennerley, S. W., Walton, M. E., Behrens, T. E. J., Buckley, M. J. & Rushworth, M. F. S. Optimal decision making and the anterior cingulate cortex. *Nat. Neurosci.* **9**, 940-947 (2006).

33 Ellerby, Z. W. & Tunney, R. J. The Effects of Heuristics and Apophenia on Probabilistic Choice. *Adv. Cogn. Psychol.* **13**, 280-295 (2017).

34 Hamel, R. *et al.* Added value of money on motor performance feedback: Increased left central beta-band power for rewards and fronto-central theta-band power for punishments. *NeuroImage* **179**, 63-78 (2018).

35 Billeke, P. *et al.* Human Anterior Insula Encodes Performance Feedback and Relays Prediction Error to the Medial Prefrontal Cortex. *Cerebral Cortex* **30**, 4011-4025 (2020).

36 HajiHosseini, A., Rodriguez-Fornells, A. & Marco-Pallares, J. The role of beta-gamma oscillations in unexpected rewards processing. *Neuroimage* **60**, 1678-1685 (2012).

37 Cohen, M. X., Elger, C. E. & Ranganath, C. Reward expectation modulates feedback-related negativity and EEG spectra. *Neuroimage* **35**, 968-978 (2007).

38 Donamayor, N., Schoenfeld, M. A. & Munte, T. F. Magneto- and electroencephalographic manifestations of reward anticipation and delivery. *Neuroimage* **62**, 17-29 (2012).

39 Ramakrishnan, A. *et al.* Cortical neurons multiplex reward-related signals along with sensory and motor information. *PNAS* **114**, E4841-E4850 (2017).

40 HajiHosseini, A. & Holroyd, C. B. Reward feedback stimuli elicit high-beta EEG oscillations in human dorsolateral prefrontal cortex. *Sci. Rep.* **5**, Artn 13021 (2015).

41 Alicart, H., Cucurell, D. & Marco-Pallarés, J. Gossip information increases reward-related oscillatory activity. *NeuroImage* **210**, 116520 (2020).

42 Brincat, S. L. & Miller, E. K. Frequency-specific hippocampal-prefrontal interactions during associative learning. *Nat. Neurosci.* **18**, 576-581 (2015).

43 Buschman, T. J., Denovellis, E. L., Diogo, C., Bullock, D. & Miller, E. K. Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron* **76**, 838-846 (2012).

44 Leventhal, D. K. *et al.* Basal Ganglia Beta Oscillations Accompany Cue Utilization. *Neuron* **73**, 523-536 (2012).

45 Botvinick, M. M., Cohen, J. D. & Carter, C. S. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn. Sci.* **8**, 539-546 (2004).

46 Frank, M. J., Seeberger, L. C. & O'Reilly, R. C. By Carrot or by Stick: Cognitive Reinforcement Learning in Parkinsonism. *Science* **306**, 1940-1943 (2004).

47 Kozunova, G., Voronin, N., Venidiktov, V. & Stroganova, T. Reinforcement learning: a role of immediate feedback and internal model. *Zh. Vyssh. Nerv. Deiat. Im. I. P. Pavlova* **68**, 602-613 (2018).

48 Taulu, S., Simola, J. & Kajola, M. Applications of the signal space separation method. *IEEE Transactions on Signal Processing* **53**, 3359-3372 (2005).

49 Kilavik, B. E., Zaepffel, M., Brovelli, A., MacKay, W. A. & Riehle, A. The ups and downs of beta oscillations in sensorimotor cortex. *Exp. Neurol.* **245**, 15-26 (2013).

50 Gramfort, A. *et al.* MEG and EEG data analysis with MNE-Python. *Front. Neurosci.*, 267 (2013).

51 Thomson, D. J. Spectrum estimation and harmonic analysis. *Proc. IEEE* **70**, 1055-1096 (1982).

52 Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* **67**, 1-48 (2015).

53 Kliegl, R., Wei, P., Dambacher, M., Yan, M. & Zhou, X. Experimental Effects and Individual Differences in Linear Mixed Models: Estimating the Relationship between Spatial, Object, and Attraction Effects in Visual Attention. *Front. Psychol.* **1** (2011).

54 Tibon, R. & Levy, D. A. Striking a balance: analyzing unbalanced event-related potential data. *Front. Psychol.* **6** (2015).

55 Vossen, H., Van Breukelen, G., Hermens, H., Van Os, J. & Lousberg, R. More potential in statistical analyses of event-related potentials: A mixed regression approach. *Int. J. Methods Psychiatr. Res.* **20**, e56-e68 (2011).

56 Benjamini, Y. & Yekutieli, D. The Control of the False Discovery Rate in Multiple Testing under Dependency. *Ann. Stat.* **29**, 1165-1188 (2001).

57 Lenth, R. *et al. emmeans: estimated marginal means, aka least-squares means. R package version 1.4.5*, <https://cran.microsoft.com/snapshot/2020-10-12/web/packages/emmeans/emmeans.pdf> (2020).

58 Ségonne, F. *et al.* A hybrid approach to the skull stripping problem in MRI. *Neuroimage* **22**, 1060-1075 (2004).

59 Pascual-Marqui, R. D. Standardized low-resolution brain electromagnetic tomography (sLORETA): technical details. *Methods Find. Exp. Clin. Pharmacol.* **24 Suppl D**, 5-12 (2002).

Figures and figure legends

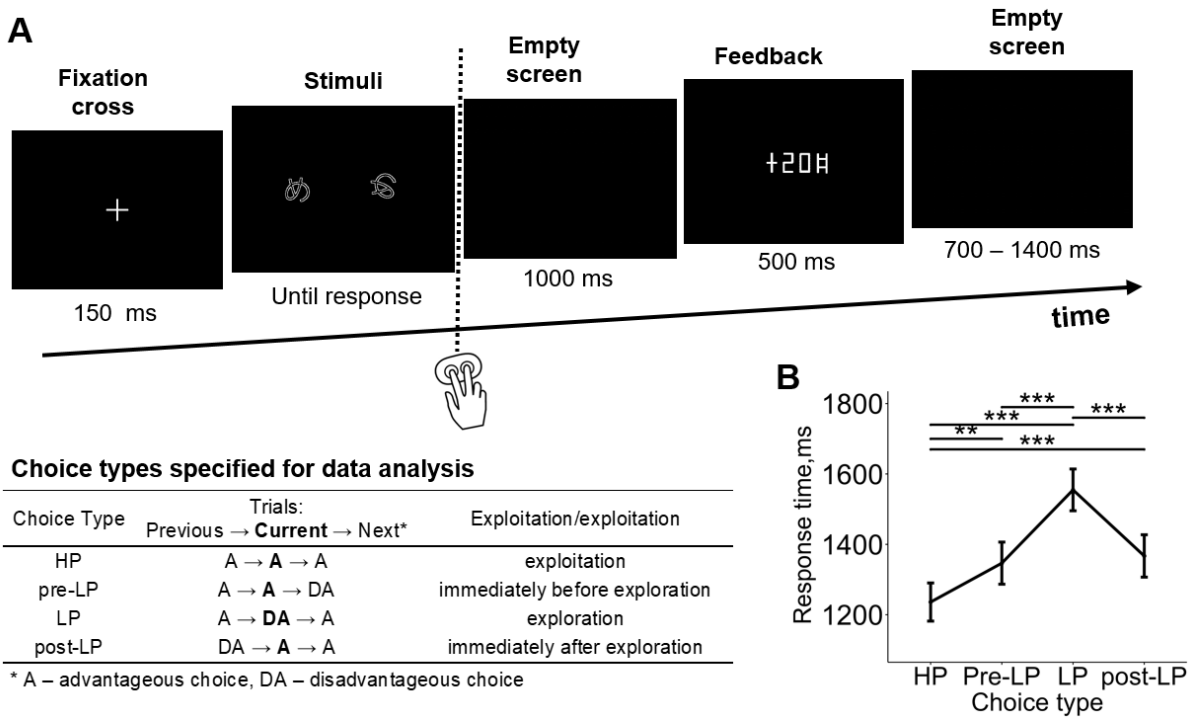

Choice types specified for data analysis

Choice Type	Trials: Previous → **Current** → Next*	Exploitation/exploitation
HP	A → **A** → A	exploitation
pre-LP	A → **A** → DA	immediately before exploration
LP	A → **DA** → A	exploration
post-LP	DA → **A** → A	immediately after exploration

* A – advantageous choice, DA – disadvantageous choice

Fig. 1. Probabilistic two-alternative gambling task: experimental design and behavioral statistics.

(A) Experimental procedure. After a briefly presented fixation cross, a pair of visual stimuli was displayed on the screen. Participants learned by trial-and-error to select (by pressing one of the two buttons) the more advantageous of the two options, i.e., the one bringing monetary gains more often than the disadvantageous option, using the probabilistic feedback. After each choice, the number of points earned was displayed with a 1000-ms delay. Duration of each event is indicated under respective screens. The trials were organized into six blocks with different stimuli and reward structures between blocks. In each block, we took into analysis only 'after learning' periods, during which participants exhibited a stable behavioral preference for the advantageous stimulus. The table explains the further division of advantageous and disadvantageous choices into four choice types used in further analyses: high-payoff (HP) – advantageous/exploitatory choices committed in a row of similar advantageous choices, low-payoff (LP) – rare disadvantageous/exploratory choices, pre-LP and post LP – advantageous choices that were committed immediately before and immediately after an exploratory LP choice. See text for details.

(B) Response time as a function of choice type. Points and error bars on graphs represent M ± SEM across single trials in all subjects.

** – $p<0.01$, *** – $p<0.001$ (Tukey HSD).

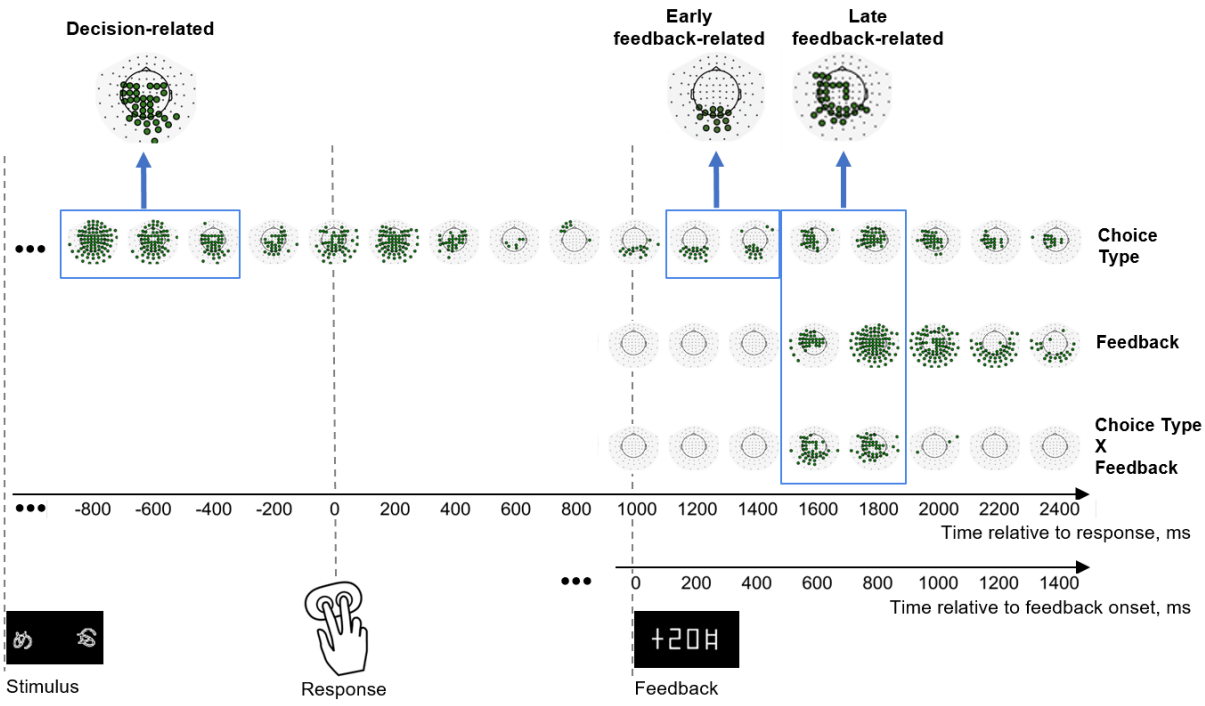

Fig. 2. Effects of Choice type, Feedback, and their interaction on β power change in the consecutive 200 msec time intervals throughout an entire trial. Timelines at the bottom relate to response onset (zero point on the upper timeline) and to feedback onset (zero point on the lower timeline). The onsets of each event (stimulus presentation, choice selection, feedback) are marked by respective pictograms. Effect names are indicated at the right of the figure. The sequence of topographical maps corresponding to each effect demonstrates significant MEG sensors derived from LMM statistics (filled circles; $p<0.05$, FDR-corrected for 17 time windows x 102 sensors). Blue frames mark time windows used in further analyses of the effects (decision-related, early feedback-related, and late feedback-related), with inserts at the top of each rectangle representing clusters of sensors taken for analysis. See Methods for details.

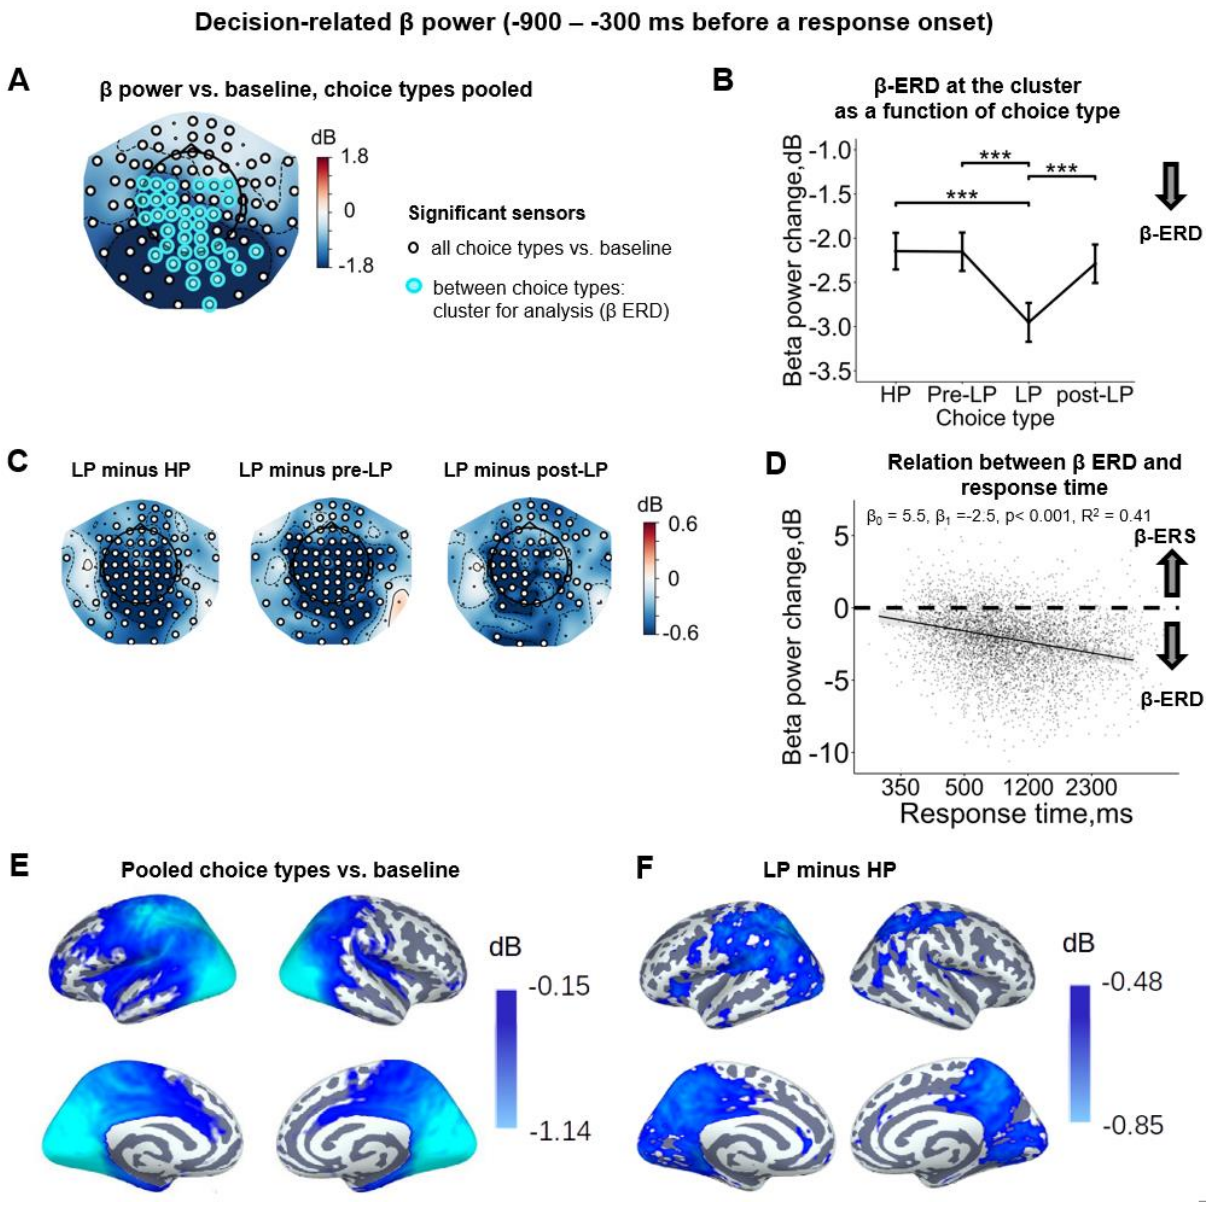

Fig. 3. Enhanced decision-related β power suppression (β-ERD) distinguishes exploratory LP choice from exploitatory choice types.

(A) Topographic map representing β-band power change relative to baseline averaged over the time interval of -900 – -300 ms before a response onset, all choice types pooled. Small black open circles mark sensors, in which β power is significantly different from baseline (t-test, p<0.05, uncorrected); sensors included in the cluster differentiating choice types according to LMM analysis (cf. Fig 3) are additionally marked by colored rings. Here and hereafter, blue and red colors denote β-ERD and β-ERS respectively.

(B) Decision-related β-band power change averaged over the cluster of significant sensors (colored rings in panel A), as a function of Choice Type. The β power change for each choice type show averages across sensors per significant cluster marked by the colored rings in panel A. Points and error bars on graphs represent M ± SEM in single trials in all subjects. *** – $p<0.001$ (LMM, Tukey HSD test).

(C) Topographic maps of the pairwise difference in β power change between the LP and other choice types at the sensor level. Significant sensors (LMM, Tukey HSD test, $p<0.05$, FDR-corrected) are indicated by small black open circles.

(D) Linear regression of response time to decision-related β-band power change averaged across sensors of significant cluster (colored rings in panel A), all choice types pooled. The line represents linear regression; shaded areas depict 95% confidence interval. Dots on scatterplots correspond to β power change in individual trials in all participants.

(E and F) Spatial distribution of cortical sources underlying decision-related β-ERD: (E) relative to baseline, all conditions pooled; and (F) increased β-ERD for LP as compared to HP choices. Only significant vertices are shown (t-test, $p < 0.05$, uncorrected).

For event-related spectral perturbations and timecourses of β power of the LP vs. HP contrast, see Supplementary Fig. S2.

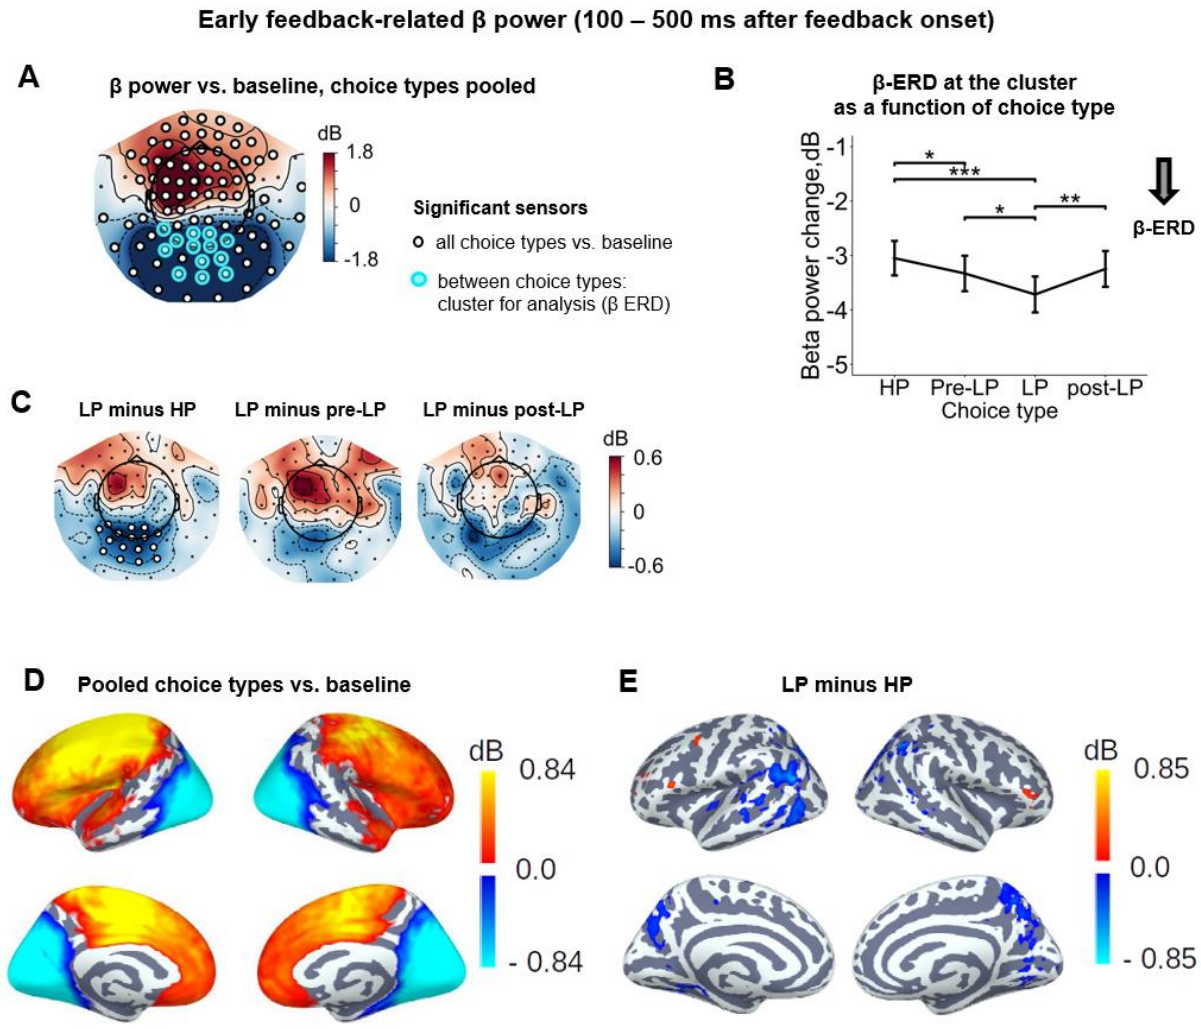

Fig. 4. Enhanced early feedback-related β-ERD at posterior cortex distinguishes exploratory LP choice from exploitatory choice types.

(A) Topographic map of feedback-related β-band power change relative to baseline averaged across the time interval 100 – 500 ms after feedback onset, all choice types pooled. Blue and red colors denote β-ERD and β-ERS respectively. Small black open circles mark sensors, in which feedback-related β power is significantly different from baseline (t-test, $p<0.05$, uncorrected); sensors included in the cluster differentiating feedback-related β power between choice types according to LMM analysis (cf. Fig. 3) are additionally marked by colored rings. Note that β-ERD in the posterior cluster significantly varies between choice types, while β-ERS in the anterior sensors does not differentiate between choices.

(B) Early feedback-related β-band power change averaged over the cluster of significant sensors (colored rings in panel A), as a function of Choice Type. Points and error bars on graphs represent M ± SEM across single trials in all subjects.
* – $p<0.05$, ** – $p<0.01$, *** – $p<0.001$ (LMM, Tukey HSD test).

(C) Topographic maps of the pairwise difference in early feedback-related β power change between the LP and other choice types at the sensor level. Significant sensors (LMM, Tukey HSD test, $p<0.05$, FDR-corrected) are indicated by small black open circles

(D and E) Spatial distribution of cortical sources underlying feedback-related β power change: (E) relative to baseline, all choice types pooled; and (D) increased β-ERD for LP as compared to HP choices. Only significant vertices are shown (t-test, $p < 0.05$, uncorrected).

For event-related spectral perturbations and timecourses of β power of the LP vs. HP contrast, see Supplementary Fig. S3.

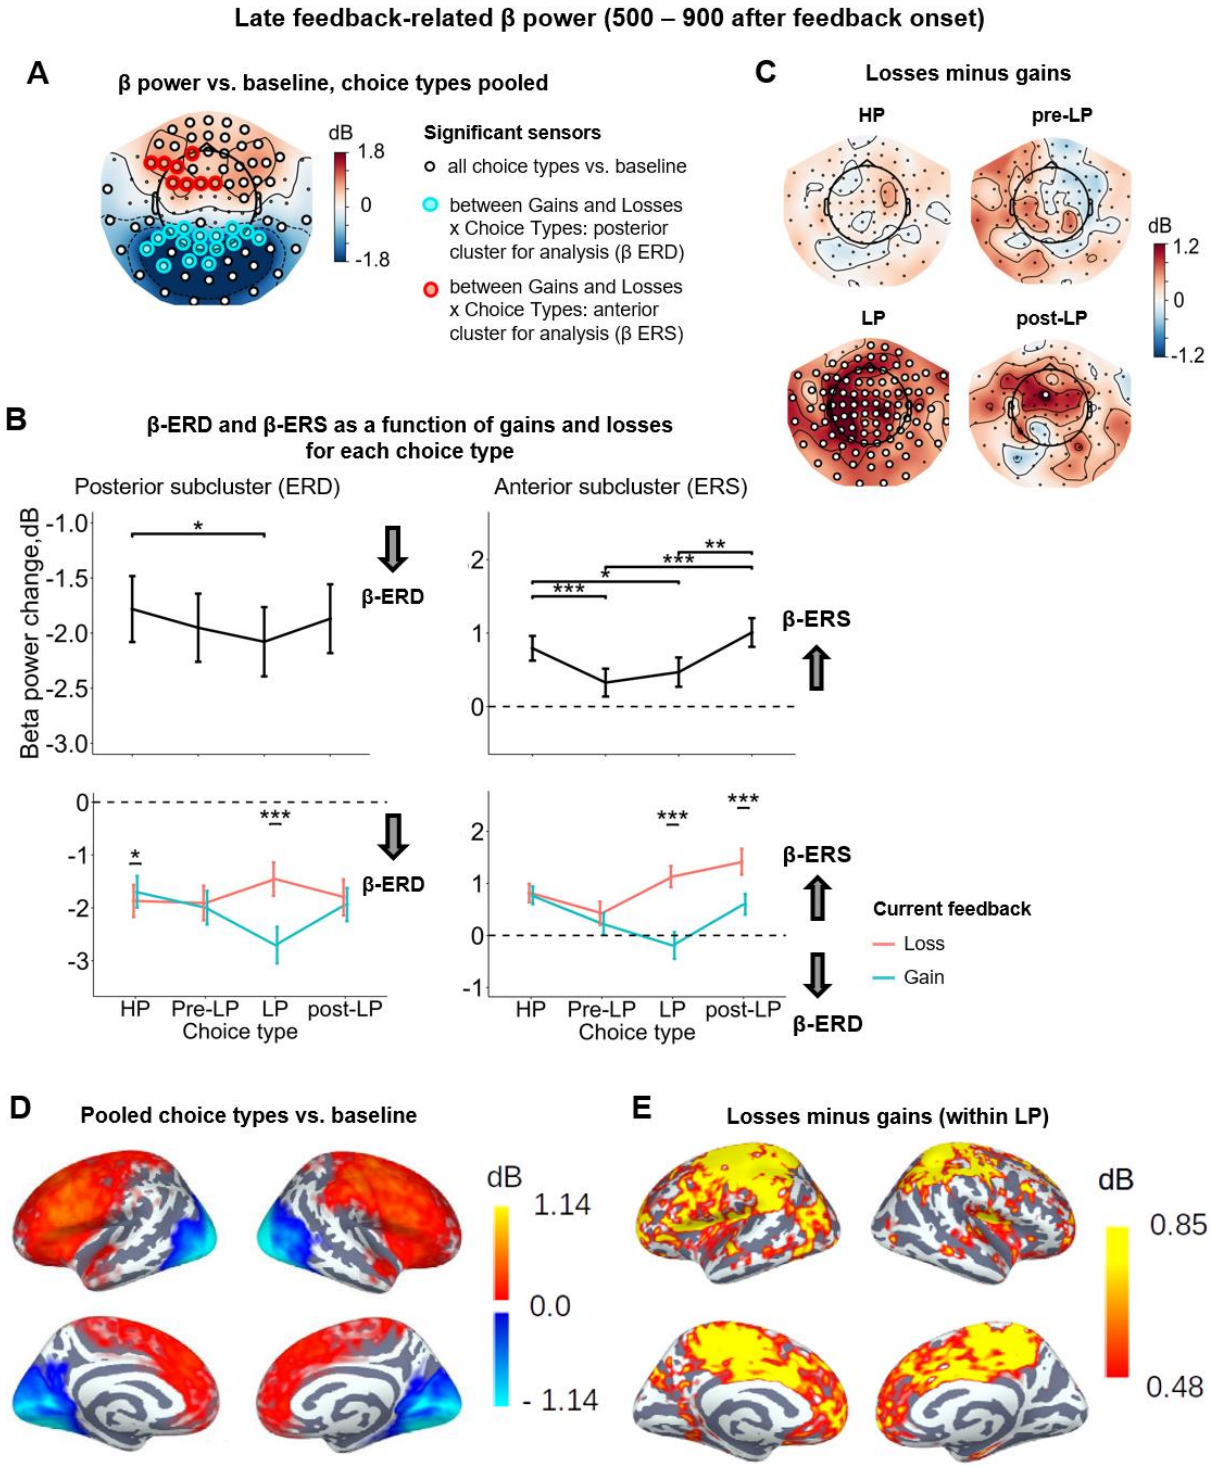

Fig. 5. Heightened sensitivity to losses versus gains in late feedback-related anterior β-ERS and posterior β-ERD distinguishes exploratory LP choices from exploitatory choice types.

(A) Topographic map of late feedback-related β-band power change relative to baseline averaged over the time interval 500 - 900 ms after feedback onset, all choice types pooled. Small black open circles mark sensors, in which late feedback-related β power is significantly different from baseline (t-test, $p<0.05$, uncorrected); sensors included in the cluster

differentiating feedback-related β power following gains and losses between choice types according to LMM analysis (cf. Fig. 3) are additionally marked by colored rings.

(B) Late feedback-related β-band power change averaged over the clusters of sensors (colored rings in panel A): anterior cluster displaying β-ERS and posterior cluster displaying β-ERD, as a function of Choice Type (upper row) and its interaction with gain or loss (lower row). Points and error bars on graphs represent M ± SEM across single trials in all subjects. * – $p<0.05$, *** – $p<0.001$ (LMM, Tukey HSD test).

(C) Topographic maps representing contrasts between losses and gains for each choice type separately. Significant sensors (LMM, Tukey HSD test, $p<0.05$, FDR-corrected) are indicated by small black open circles.

(D and E) Cortical distribution of sources demonstrating β power change: (D) relative to baseline, all conditions pooled, and (E) the contrast between losses and gains within the LP choices. Only significant vertices are shown (t-test, $p < 0.05$, uncorrected).

For event-related spectral perturbations and timecourses of β power of the losses vs. wins contrast under LP condition, see Supplementary Fig. S4.

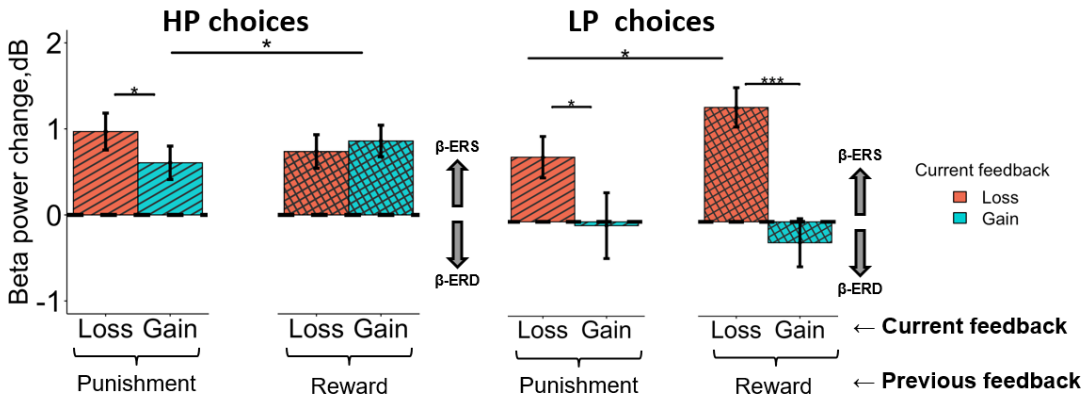

Fig. 6. Different impact of the previous choice outcome on anterior β-ERS induced by losses and gains in the LP and HP choices.

Late β band power change (500 – 900 ms after feedback onset) averaged across the anterior cluster (see Fig.5A) are shown for LP and HP choices as a function of two factors: a current feedback sign (loss or gain, red and blue rectangles respectively) and a previous feedback sign (punishment or reward indicated by horizontal brackets). Boxes and error bars on graphs represent M ± SEM across single trials in all subjects. Asterisks denote Tukey HSD p-values: *p<0.05, ***p<0.001.