

Reviewed Preprint

v1 • September 26, 2025

Not revised

Reviewed Preprint

v2 • January 23, 2026

Revised by authors

✉ For correspondence:

paul.bogdan@duke.edu

Competing interests:

No competing interests declared

Funding:

See [page 24](#)

Reviewing editor:

Alex Fornito,
Monash University, Australia

© 2025, Bogdan et al. This article is

distributed under the terms of the

[Creative Commons Attribution](#)

[License](#), which permits unrestricted

use and redistribution provided that

the original author and source are

credited.

Spontaneous fluctuations in global connectivity reflect transitions between states of high and low prediction error

Paul C Bogdan¹✉, Shenyang Huang¹, Lifu Deng², Simon W Davis^{1,3}, Roberto Cabeza¹

¹Department of Psychology and Neuroscience, Duke University, Durham, United States • ²Cleveland Clinic,

Cleveland, United States • ³Department of Neurology, Duke University School of Medicine, Durham, United States

eLife Assessment

This study uses a **valuable** combination of functional magnetic resonance imaging and electroencephalography (EEG) to study brain activity related to prediction errors in relation to both sensorimotor and more complex cognitive functions. It provides **incomplete** evidence to suggest that prediction error minimisation drives brain activity across both types of processing and that elevated inter-regional functional coupling along a superior-inferior axis is associated with high prediction error, whereas coupling along a posterior-anterior axis is associated with low prediction error. The manuscript will be of interest to neuroscientists working on predictive coding and decision-making, but would benefit from more precise localisation of EEG sources and more rigorous statistical controls.

<https://doi.org/10.7554/eLife.107723.2.sa4>

Abstract

While numerous researchers claim that the minimization of prediction error (PE) is a general force underlying most brain functions, others argue instead that PE minimization drives low-level, sensory-related neuronal computations but not high-order, abstract cognitive operations. We investigated this issue using behavioral, fMRI, and EEG data. Studies 1A/1B examined semantic- and reward-processing PE using task-fMRI, yielding converging evidence of PE's global effects on large-scale connectivity: high-PE states broadly upregulated ventral-dorsal connectivity, and low-PE states upregulated posterior-anterior connectivity. Investigating whether these global patterns characterize cognition generally, Studies 2A/2B used resting-state fMRI and showed that individuals continuously fluctuate between ventral-dorsal (high-PE) and posterior-anterior (low-PE) dynamic connectivity states. Additionally, individual differences in PE task responses track differences in resting-state fluctuations, further endorsing that these fluctuations represent PE minimization at rest. Finally, Study 3 combined fMRI and EEG data, and the study found that the fMRI fluctuation amplitude correlates most strongly with EEG power at 3–6 Hz, consistent with the PE network fluctuations occurring at Delta/Theta oscillation speeds. This whole-brain layout and timeline together are consistent with high/low-PE fluctuations playing a role in integrative and general sub-second cognitive operations.

1. Introduction

The neuroscientific literature is replete with claims that one of the brain's primary functions is to model its environment, generate predictions, and refine its activity to minimize prediction error (PE). It is generally agreed that the brain uses PE mechanisms at neuronal or regional levels,^{1,2} and this idea has been useful in various low-level functional domains, including early vision¹ and

dopaminergic reward processing.³ Some theorists have further argued that PE propagates through perceptual pathways and can elicit downstream cognitive processes to minimize PE.^{4–7} For example, PE might be minimized via attentional re-orienting in parietal regions,^{8,9} memory encoding in the medial temporal lobes,^{10,11} and language learning or category-updating in temporal areas.^{12–14} More controversially, it has also been argued that abstract functions like cognitive processing of emotion,⁴ Theory of Mind,⁵ and others^{7,15} have likewise been described as PE minimization. Following this line of reasoning, even cognition without exogenous stimuli may reflect the minimization of PE. For instance, daydreaming about what food you will eat later can be seen as an effort to avoid future sub-optimal food choices. Coordinating cortical mechanisms for sensation, decision-making, and autobiographical memory, one can describe cognition as a series of mental operations seeking to minimize PE.

PE minimization may broadly coordinate brain functions of all sorts, including abstract cognitive functions. This includes the types of cognitive processes at play even in the absence of stimuli (e.g., while daydreaming). While it may seem counterintuitive to associate this type of cognition with PE – a concept often tied to external surprises – it has been proposed that the brain’s internal generative model is continuously active.^{16–18} Spontaneous thought, such as planning a future event or replaying a memory, is not a passive, low-PE process. Rather, it can be seen as a dynamic cycle of generating and resolving internal uncertainty. While daydreaming, you may be reminded of a past conversation, where you wish you had said something different. This situation contains uncertainty about what would have been the best thing to say. Wondering about what you wish you said can be viewed as resolving this uncertainty, in principle, forming a plan if the same situation ever arises again in the future. Each iteration of the simulated conversation repeatedly sparks and then resolves this type of uncertainty.

Ideas about PE operating at a global level and underlying abstract cognitive functions are difficult to evaluate experimentally, which is the root of the controversy about this view.^{19–21} In particular, it is challenging to isolate high-level and multifaceted cognition processes in a way that deters alternative explanations linked to other neural mechanisms or stimulus-specific functions. Yet, doing so would yield key evidence and provide direction for the variety of neurocognitive theories attempting to unify all levels of brain function under a PE lens.^{4–7,15} We pursued this goal using three studies (Figure 1 [↗](#)) that collectively targeted a specific question: Do the task-defined connectivity signatures of high vs. low PE also recur during rest, and if so, how does the brain transition between exhibiting high/low signatures?

First, we used task-fMRI to identify the functional connectivity signatures of high-PE and low-PE states. Few prior PE studies have targeted connectivity, contrasting the hundreds that have targeted BOLD.^{22,23} Yet, compared to measuring activity, tracking connectivity better taps into the core concepts of the integrative PE view that theorizes how signals communicate through the brain to propagate and/or minimize PE. For Study 1A, we conducted a task-fMRI study (N = 66) that elicited PE relative to pre-existing schemas and aimed to identify the connectivity patterns of low and high PE. For Study 1B, we analyzed a public task-fMRI study (N = 1000) that elicited reward PE using a gambling task. We expected to replicate the connectivity signatures of high and low PE seen in the schema task, showing that these networks emerge across starkly different PE tasks.

Second, we reasoned that if PE processing characterizes cognition more generally, then it should also occur in the absence of sensory stimuli. Studies 2A and 2B examined PE during resting-state fMRI (rs-fMRI) scans. Detecting PE processing during a resting-state fMRI scan required a template of PE processing.²⁴ Thus, we used the connectivity patterns of high and low PE identified in Studies 1A/1B, and searched the rs-fMRI data for these connectivity patterns. This strategy can link the Study 1 task manipulations and resting-state phenomena without *a priori* assumptions about individual brain regions’ exact roles.²⁵ One potential problem of this approach is known as “reverse inference”: Finding the same networks during a PE task and during rest does not mean that the networks’ engagement during rest necessarily reflects PE processing. To engage with this issue, we examined if the engagement of putative PE networks during rest and during a PE task




Figure 1. Task design and graphical summaries of analytic strategies.

(Study 1) Two task-fMRI studies were conducted to identify the connectivity patterns upregulated by high or low PE. **(A)** In each trial, participants saw a scene-object pair and were asked to report how likely it would be to find the object in the scene using a 4-point scale that was active while the object was on screen. **(B)** Separate participants completed a gambling task wherein each trial they made a binary choice (“?” screen) and were then told whether the choice caused them to earn money (upward arrow) or lose money (downward arrow). The “+” panel represents a 15 s fixation period between blocks. Both earning or losing money could elicit high or low PE, depending on whether it deviates from prior trials’ outcomes. **(Study 2)** **(A)** Using rs-fMRI data collected from Study 1A participants, analyses examined changes in connectivity over time to assess whether participants spontaneously fluctuate between the high and low PE signatures found in Study 1. **(B)** Further analyses examined whether slight differences in Study 1B participants’ task-fMRI responses to PE mirrored differences in their rs-fMRI connectivity states. **(Study 3)** With rs-EEG-fMRI data from another pool of participants, analyses modeled the amplitudes of fMRI connectivity fluctuations over time and correlated this with EEG oscillatory power at different frequencies to assess which frequency best tracked the fMRI patterns.

was significantly correlated across participants. Although this does not entirely address the reverse inference dilemma and can only produce correlative evidence, the present research nonetheless investigates these widely speculated upon PE ideas more directly than any prior work. Finally, in Study 3, we used rs-fMRI-EEG simultaneous recordings ($N = 21$) to investigate whether participants fluctuate between high/low PE states at speeds consistent with whole-brain integrative cognitive operations. Although no prior fMRI study has focused on PE under this lens, this perspective is motivated by EEG research. Specifically, in EEG studies on executive function, PE detection elicits a frontal negativity around 100–400 ms^{26,27} that is followed by a centroparietal positivity around 300–800, indexing PE minimization processes such as redirecting attention, changing behavior, and updating working memory.^{26,28,29} These results indicate that shifts from high-to low-PE states can occur within a few hundred milliseconds. This idea is consistent with evidence that PE modulates Delta (1–4 Hz) and Theta (4–8 Hz) power, which aligns with the timing of error-related negativity/positivity.³⁰ Hence, in Study 3, we expected to find that participants' high- and low-PE networks would fluctuate multiple times per second. Although such speed outpaces the temporal resolution of fMRI, correlating fluctuations in dynamic connectivity measured from fMRI data with EEG oscillations can provide an estimate of the fluctuations' speed. This interpretation of a correlation again runs up against issues related to reverse inference but would nonetheless serve as initial suggestive evidence that spontaneous transitions between network states occur rapidly.

Instead, sub-second fluctuations could reflect individual cognitive operations, such as working memory updating, attentional shifting, or other processes attributed to PE minimization. Moreover, we specifically expected that the fluctuations occur at Delta/Theta frequencies, which would also track the PE literature on oscillations and further validate that the resting-state patterns reflect PE minimization³⁰ – that is, despite concerns about reverse inference and interpreting rs-fMRI data, by demonstrating an additional link to Delta/Theta, this would constitute suggestive converging evidence for the PE interpretation.

In sum, we investigated the global and task-general/spontaneous role of PE in three steps. First, we used task-fMRI to identify large-scale brain networks associated with high and low PE. Second, we examined the recruitment of these networks during rs-fMRI, and although the problems related to reverse inference are impossible to overcome fully, we engage with this issue by linking rs-fMRI data directly to task-fMRI data of the same participants, which can provide suggestive evidence that the same neural mechanisms are at play in both. Finally, we examined rs-fMRI-EEG data to assess whether we find parallels consistent with the high/low-PE network fluctuations occurring at fast timescales suitable for the type of cognitive operations typically targeted by PE theories.

2. Results

2.1. Study 1 – Task-fMRI: Mapping the effects of prediction error on connectivity

2.1.1. Study 1A – Manipulating semantic PE

Participants completed a task wherein each trial they viewed a well-known scene, such as a farm. The scene was followed by either an object that was highly predictable given the scene (low PE; e.g. a tractor in a farm), an object that was less predictable but still possible (medium PE; e.g., an insect in a desert), or an object that was highly unpredictable (high PE; e.g. a poker table in a prison). Participants rated how likely it would be for each object to appear in its associated scene (1 = “*Very unlikely*”, 4 = “*Very likely*”), which confirmed that participants processed the stimuli as intended: Low-PE objects were rated as likely to appear in the preceding scene ($M = 3.67$, $SD = 0.20$), medium-PE objects as moderately likely ($M = 2.11$, $SD = 0.30$), and high-PE objects as unlikely ($M = 1.28$, $SD = 0.24$). Average response times for trials where participants gave the expected response were slightly slower for the high-PE condition ($M = 1.87$ s, $SD = 0.40$ s) compared to the

low-PE condition ($M = 1.79$ s, $SD = 0.41$ s; $t[65] = 2.54$, $p = .014$, $d = 0.27$). This response time gap is consistent with error-related slowing, although it is nominally small.³¹ Overall, these behavioral trends suggest the suitability of the data for understanding the neural correlates of PE.

In the fMRI data, we compared the three PE conditions' beta-series functional connectivity, aiming to identify network-level signatures of PE processing, from low to high. Network-level signatures could later be leveraged for detecting global trends in rs-fMRI data. After demonstrating that PE significantly influences brain-wide connectivity using Network-Based Statistic analysis (Supplemental Materials 2.1 [32](#)), we conducted a modularity analysis to study how specific groups of edges are all sensitive to high/low-PE information. For the modularity analysis, we first defined a connectome matrix of beta values, wherein each edge's value was the slope of a regression predicting that edge's strength from PE (coded as Low = -1, Medium = 0, High = +1; Figure 2A [33](#)). Modules representing positive and negative effects were then identified by applying a variant of Louvain's algorithm twice: once for a binary matrix of the 5% of edges showing the strongest Low > High effects and once for the analogous High > Low matrix (Figure 2B [34](#)).

Roughly, the high/low PE modules appear as connections among quadrants, consisting of the occipital cortex (OC), inferior parietal lobule (IPL), anterior temporal lobe (ATL), and dorsolateral prefrontal cortex (LPFC, i.e., the inferior frontal gyrus [IFG] & middle frontal gyrus [MFG]) (Figure 2C [35](#)-2E [36](#)). That is, high PE enhanced ventral-dorsal connectivity (OC-IPL & ATL-LPFC) whereas low PE enhanced posterior-anterior connectivity (OC-ATL & IPL-LPFC) (Figure 2F [37](#)). This four-quadrant description omits some ROIs included in the modules (e.g., the lateral temporal component of the IPL-LPFC module; Figure 2C [38](#)).

Some ROIs appear in Figure 2C [39](#) but are excluded from the four targeted quadrants (Figures 2C & 2D [40](#)) – e.g., posterior inferior temporal lobe and fusiform ROIs are excluded from the OC-IPL module, and middle temporal gyrus ROIs are excluded from the IPL-LPFC modules. These exclusions, in favor of a four-quadrant interpretation, are motivated by existing knowledge of prominent structural pathways among these quadrants. This interpretation is also supported by classifier-based analyses showing connectivity within each quadrant is significantly influenced by PE (Supplemental Materials 2.2 [41](#)), along with analyses of single-region activity showing that these areas also respond to PE independently (Supplemental Materials 3 [42](#)). Hence, we proceeded with further analyses of these quadrants' connections, which summarize PE's global brain effects.

This four-quadrant setup also imparts analytical benefits. First, this simplified structure may better generalize across PE tasks, and Study 1B would aim to replicate these results with a different design. Second, the four quadrants mean that each ROI contributes to both the posterior-anterior and ventral-dorsal modules, which would benefit later analyses and rules out confounds such as PE eliciting increased/decreased connectivity between an ROI and the rest of the brain. An additional, less key benefit is that this setup allows more easily evaluating whether the same phenomena arise using a different atlas (Supplemental Materials 7 [43](#)).

2.1.2. Study 1B – Manipulating reward PE

To confirm that the ventral-dorsal and posterior-anterior networks are indeed general signatures of high and low PE, respectively, we investigated their occurrence during a very different PE task: the gambling task from the Human Connectome Project ($N = 1000$).^{33,34} In each trial of this task, participants made a binary choice, then won or lost money as pre-determined. Our analyses modeled PE as the absolute difference between a trial's monetary outcome and participants' expectations given recent prior trials (i.e., relative to an exponential moving average). This strategy assigns a continuous PE measure to each trial, which was then split at the median to define high/low PE conditions. As a validation check, we confirmed that participants' response times were slower after loss trials with high-PE feedback ($M = 422$ ms) compared to after loss trials with low-PE feedback ($M = 418$ ms) ($t[999] = 7.46$, $p < .0001$). This is consistent with expected post-error/surprise slowing,³¹ although given how slight the response-time difference was (4 ms), it was not expected to reflect any meaningful confound.




Figure 2. Modularity behind the enhancing and weakening effects of prediction error.

(A) Each neocortical connection was submitted to a regression predicting connectivity strength based on PE. For clearer visualization, the shown 46 x 46 matrix simplifies the original 210 x 210 matrix by averaging ROIs per anatomical labels. Supplemental Figure S1 [\[link\]](#) shows the exact ROIs used on a glass brain, illustrating their precise anatomy. (B) Modularity was assessed in terms of the 5% of edges showing the most positive coefficients and 5% showing the most negative coefficients. (C) Two high-PE modules and two low-PE modules are shown. (D & E) The four modules can be described anatomically as connections among four anatomical quadrants. (F) Analyzing these anatomically defined connections shows a significant Connectivity-Direction x PE interaction (β [standardized] = .20, $p < .001$). Boxes indicate quantiles.

Analyses evaluated the consistency between the connectivity effects of PE in the present gambling task and those seen earlier in the semantic PE task. Here, functional connectivity was computed for high/low PE conditions. Then, a paired t-test between the two conditions was performed for each edge, which yielded a connectome-wide high/low contrast. This contrast matrix paralleled that of Study 1A, as evidenced by element-wise correlations of the matrices' bottom triangles (Figures 2A & 3A, $r = .27$, $p < .0001$). Critically, the present data also reproduced Study 1A's Connectivity Direction $\times$ PE interaction ($F[1, 999] = 68.41$, $p < .0001$; Figure 3B). This prominent correlation and reproduction point to strong task-general qualities of the identified global PE phenomena.

In sum, Study 1A identified distinct global connectivity patterns linked to PE: High PE upregulates a ventral-dorsal network whereas low PE upregulates a posterior-anterior network. Study 1B reproduced these signatures using a very different PE task. This regularity points to PE processing as prompting general connectome-wide mechanisms. Yet, it remains to be tested whether PE processing applies to high-level cognition more broadly, including in the absence of stimuli specifically designed to elicit PE. Thus, Study 2 investigated whether the ventral-dorsal and posterior-anterior networks are engaged when stimuli are absent, as would be expected if PE minimization is a fundamental function of the human brain.

2.2. Study 2 – Rs-fMRI: Identifying connectivity fluctuations tied to prediction error

2.2.1. Study 2A – Dynamic connectivity and global states

We assessed whether the above correlates of high/low PE emerge as dynamic connectivity states in rs-fMRI data from Study 1A participants. In other words, we essentially asked whether resting-state participants are sometimes in low PE states and sometimes in high PE states, which would be consistent with spontaneous PE processing in the absence of stimuli. To formally test this, we first measured *time-varying connectivity* on a TR-by-TR basis.³⁵ For example, OC-IPL connectivity strength in the first TR is the product of OC and IPL standardized activations in said TR ($OC_{BOLD} \times IPL_{BOLD}$). Time-varying connectivity was computed for the edges among adjacent portions of the earlier four quadrants (Figure 4A). In turn, edge-edge correlations³⁵ were calculated to form the Figure 4B matrix; this matrix shows, for example, that strong L OC-IPL connectivity coincides with weak L OC-ATL connectivity. Submitting this matrix to Louvain's algorithm, in turn, yielded the two modules in Figure 4C: Here, module 1 indicates that, when IPL-LPFC connectivity is strong, then OC-ATL tends to also be strong, and these two sets of edges make up a global state of strong posterior-anterior connectivity. Likewise, module 2 shows that OC-IPL and ATL-LPFC connectivity make up a ventral-dorsal state. These emerging states overlap strikingly with the previous task effects of PE, suggesting that rs-fMRI scans exhibit fluctuations that resemble the signatures of low- and high-PE states.

2.2.2. Study 2B – Linking task and resting-state PE

Study 2B had two goals. The first goal was to replicate the results of Study 2A in a different group of rs-fMRI participants ($N = 1000$). Fulfilling this goal, Study 2B yielded the same posterior-anterior and ventral-dorsal connectivity patterns during rs-fMRI (Figure 5A).

The second goal of Study 2B was to engage with the aforementioned reverse inference issue: Finding the same networks during a PE task and during rest does not imply that the networks' engagement during rest necessarily reflects PE processing. By establishing these networks using two very different PE tasks (Studies 1A/1B), this deters non-PE interpretations of the networks specific to just one task. To further address this issue, we considered the evidence that the anatomical organization of functional networks varies between individuals.³⁶ Hence, if the recruitment of ventral-dorsal/posterior-anterior networks during rest is indeed related to PE, then a participant's unique connectivity response to PE in the reward-processing task should match their specific patterns of resting-state fluctuation. Our focus is not on the overall strength of the task contrast and overall fluctuation magnitude at rest; the theoretical justification for these being linked seems questionable, as PE minimization is presumed to characterize every cognitive

Figure 3. Reproducing the posterior-anterior/ventral-dorsal PE effects using a gambling task.

(A) This matrix represents high/low PE paired t-tests for each connection. Like Figure 2, a 46 x 46 regional matrix is shown rather than a 210 x 210 ROI matrix for clarity. (B) Low PE enhances posterior-anterior connectivity while high PE enhances ventral-dorsal connectivity but note the High > Low PE ventral-dorsal effect is driven by ATL-LPFC connectivity, and the OC-IPL effect is minuscule. For the bars, the connectivity data was normalized by subtracting means across conditions. Boxes indicate quantiles.





Figure 4. Dynamic connectivity modules.

(A) Time-varying connectivity was computed between adjacent portions of the four bilateral quadrants from earlier. Given that each quadrant is composed of many ROIs, this involved computing time-varying connectivity for individual ROI-ROI pairs and averaging them to produce overall estimates. (B) Based on the time-varying connectivity of each edge, an edge-edge correlation matrix was produced. (C) Computing modularity of the edge-edge correlation matrix yielded the two illustrated modules.






Figure 5. Replication of the connectivity modules and subject-specific analyses.

Performing the Study 2A modularity analysis using the Study 1B participants yields a similar pair of modules as in [Figure 4](#) but with slight changes to the interhemispheric connections.

process of every single person's brain. Yet, there should still be slight anatomical differences in where exactly these networks are positioned, and these differences should be consistent across the task-fMRI and rs-fMRI data. We tested whether this was indeed the case.

To evaluate anatomical correspondence across task-fMRI and rs-fMRI, we analyzed these effects for subsets of the quadrants. From each quadrant, we randomly selected a pair of bilateral ROIs (one L ROI and its symmetric R ROI). Then, using just these eight ROIs, we calculated the strength of each participant's task-fMRI effect (Connectivity Direction $\times$ PE interaction) along with the magnitude of their rs-fMRI posterior-anterior/ventral-dorsal fluctuation (mean absolute difference of time-varying connectivity, $E[|PA - VD|]$; see Study 3 and Methods). These two measures were computed for all 3,432 possible sets of eight ROIs. For each participant, the within-subject correlation was computed with respect to the 3,432 ROI sets with the expectation that these would be positive on average. A positive correlation would indicate that ROIs associated with stronger PE effects in the task are linked to stronger fluctuation at rest. This was indeed found and demonstrated to be significantly above zero by submitting the 1000 participants' correlations to a one-sample t-test (mean $r = .021$, $t[999] = 3.31$, $p < .001$). Note that this nominally low correlation likely stems from the short task-fMRI scan length (only six minutes), which in turn causes the split-half reliability of a participant's task contrast to be just .08 (the reliability of the rs-fMRI measure was much higher, see Supplemental Materials 4 [4](#)). Hence, a nominally low correlation is to be expected given that reliability sets an upper limit on the possible correlation. This rs-fMRI effect is specific to the PE contrast in the task-fMRI data; if the same analysis is performed with respect to a win/loss task-fMRI contrast, this produces an insignificant task-fMRI $\times$ rs-fMRI association (mean $r = .006$, $t[999] = 1.18$, $p = .24$). Hence, a participant's unique pattern of connectivity elicited by PE matches their fluctuations in resting-state scans.

To be clear, this does not entirely dissuade concerns about reverse inference, which would require a type of causal manipulation that is difficult (if not impossible) to perform in a resting state scan. Nonetheless, these results provide further evidence consistent with our interpretation that the resting brain spontaneously fluctuates between high/low PE network states.

Studies 2A and 2B show that, even in the absence of stimuli, participants fluctuate between global connectivity states corresponding to high and low PE. Overall, these results endorse claims that abstract cognition is rooted in PE minimization. However, claims of PE unifying psychological concepts generally refer to mental operations that can occur fairly quickly – e.g., emotional responses, logical inferences, memory retrieval, attentional shifting, or working memory updating. This may seem to create a discrepancy with the results thus far, which are all based on fMRI with its slow temporal resolution. Yet, slow measurement does not imply that underlying neural processes are slow. Study 3 investigated this further.

2.3. Study 3 – Rs-fMRI-EEG: Measuring fluctuation frequency

To investigate the temporal properties of our putative global PE networks shifting between high/low PE states, we conducted analyses using fMRI-EEG simultaneous recordings from participants at rest. Our overarching strategy for detecting the fluctuation rate involved first modeling the amplitude of the fMRI connectivity fluctuations on a TR-by-TR basis. Then, we correlated this amplitude with EEG oscillatory power at different frequencies to assess which best aligned with the fMRI patterns (Figures 6A & 6B [4](#)). We validated this strategy with simulations (Supplemental Materials 5 [4](#)). In brief, we specifically simulated a neural oscillator that varies in amplitude over time (Figures 6C-E [4](#)) and reflects opposing posterior-anterior (PA) and ventral-dorsal (VD) connectivity (Figure 6F [4](#)). fMRI can measure time-varying connectivity but only at a slow scale (i.e., modeled as the average connectivity across a TR; Figure 6G [4](#)). Nonetheless, the frequency of the oscillator producing these fMRI measurements can be discerned by linking the magnitude of fMRI fluctuations ($|PA - VD|$ in Figure 6H [4](#)) to EEG oscillatory power at different frequencies (Figures 6I-K [4](#)). If VD and PA strengths fluctuate at 3 Hz, then $|VD - PA|$ in fMRI will most strongly correlate with 3 Hz power in EEG. Notice how the dots in Figure 6H [4](#) follow the dots in Figure 6J [4](#) (3 Hz) better than the dots in Figure 6I [4](#) (0.5 Hz) or Figure 6K [4](#) (10 Hz); this visual comparison is intended for illustrative purposes only, and quantitative analyses are

provided in Supplemental Materials 5 [↗](#). This same logic can be applied to experimental data to determine the frequency at which individuals fluctuate between posterior-anterior and ventral-dorsal dynamic connectivity states.

Our analyses used a public dataset of fMRI-EEG simultaneous recordings from participants at rest ($N = 21$)³⁷. We first measured time-varying connectivity for the ventral-dorsal and posterior-anterior edges, as in Study 2B. We calculated the absolute difference between the two at each TR, $|PA - VD|$ to generate connectivity-amplitude time series. Turning to the EEG data, we measured oscillatory power at each frequency from 0.5 to 50 Hz (bin size = 0.5 Hz). To match the power time series with fMRI's temporal resolution, the power data were averaged by TR (e.g., averaging 0-2099 ms to align with the first 2.1 s TR volume) and convolved with the hemodynamic response function.³⁸

Spearman correlations were measured between the connectivity-amplitude time series and each frequency's power time series. Group statistics and effect sizes were then established by submitting each participant's correlation to a one-sample t-test relative to zero. The analysis was first done using power averaged across frontal electrodes, as these are the typical focus of PE research on oscillations.^{30,39–41} Correlations revealed that connectivity-fluctuation amplitude was most associated with frontal Delta/Theta and specifically 3-6 Hz power (Figure 7A [↗](#), Table 1 [↗](#)).

Similar albeit weaker trends emerged when correlating fluctuation amplitude with power averaged across a larger pool of electrodes, including central and parietal ones (Figure 7B [↗](#)). Confirmatory analyses demonstrated that these oscillatory links specifically reflect associations with the posterior-anterior/ventral-dorsal absolute difference and not the signed difference or the overall level of connectivity (Supplemental Materials 6 [↗](#)). These patterns are most consistent with a characteristic timescale near 3–6 Hz for the amplitude of the putative high/low-PE fluctuations. This is notably consistent with established links between PE and Delta/Theta^{30,39} and is further consistent with an interpretation in which these fluctuations relate to PE-related processing during rest.

3. Discussion

The present research conducted task-fMRI, rs-fMRI, and rs-fMRI-EEG studies to clarify whether PE elicits global connectivity effects and whether the signatures of PE processing arise spontaneously during rest. This investigation carries implications for how PE minimization may characterize abstract task-general cognitive processes. *First*, the two task-fMRI studies (Studies 1A/1B) yielded converging evidence that high-PE stimuli upregulate connectivity between ventral and dorsal regions (OC-IPL and ATL-LPFC), whereas low-PE stimuli upregulate posterior-anterior connectivity (OC-ATL and IPL-LPFC). *Second*, the two rs-fMRI studies (Studies 2A/2B) indicated that the brain, even while at rest, shifts between global states of strong posterior-anterior connectivity (low-PE signature) and strong ventral-dorsal connectivity (high-PE signature). Importantly, the recruitment of ventral-dorsal and posterior-anterior networks during rest correlated with the recruitment of these networks during a PE-processing task across participants. *Finally*, the rs-fMRI-EEG study (Study 3) showed that these connectivity fluctuations correlate primarily with Delta/Theta oscillatory (EEG) power. Although there are different ways to interpret this correlation, it is consistent with high/low PE states generally fluctuating at 3-6 Hz during rest. Below, we discuss these three studies' findings.

The connectivity patterns put forth are, for the most part, not spatially novel and instead overlap heavily with prior functional and anatomical findings. Detecting that high PE upregulates ventral-dorsal (OC-IPL & ATL-LPFC) connectivity matches the handful of prior studies on how PE influences interactions between regions. On the posterior side of the ventral-dorsal network, finding that high PE enhanced OC-IPL connectivity is consistent with PE prompting attentional shifts by the parietal cortex.^{8,9} For this interpretation and those upcoming, note that these mechanisms will likely only apply to subsets of the full ROIs; the OC and IPL are functionally heterogeneous structures, and our analyses focused more on broad-scale patterns rather than drawing neuroanatomically narrow conclusions. On the anterior side of the ventral-dorsal




Figure 6. Correlating connectivity fluctuations and oscillatory power.

(A) If posterior-anterior (PA) and ventral-dorsal (VD) time-varying connectivity measurements at each TR reflect samples from an oscillator, then their absolute difference ($|PA - VD|$) will correlate with the oscillator's true amplitude (A) even if the oscillator's frequency outpaces fMRI's temporal resolution. (B) Correlations between $|PA-VD|$ and amplitude/power at different frequencies allow assessing the frequency of the oscillator generating the PA/VD fluctuations; amplitude (A) and power (A^2) are equivalent for the analyses, which used Spearman correlations. The fMRI-EEG analytic strategy is motivated by simulations: (C) One sine wave is used as a simulation input. (D) An autoregressive time series is also used as an input. (E) The product of the sine wave and autoregressive time series defined the simulated oscillator, which represents how fluctuations between states may vary in intensity over time. (F) Ventral-dorsal (VD) connectivity is defined as the sine-amplitude product, and posterior-anterior (PA) connectivity is defined as the inverse of this product. (G) Representing brain data recorded using fMRI, the VD and PA dots here are the average over the TR windows. No hemodynamic response function convolution was applied when making this figure for the sake of clarity, but convolution was performed for the actual simulation results in Supplemental Materials 5 [5](#). (H) Using the PA and VD dots, the $|PA - VD|$ differences were computed. (I-K) Green lines represent the time-frequency decomposition of the sine-amplitude product. The dots represent the power averaged for each TR window. The first TR window for 0.5 Hz power was removed as it was susceptible to edge artifacts.

Figure 7. Correlations between fluctuation amplitude and oscillatory power.

For all participants, Spearman correlations were measured between connectivity-fluctuation amplitude and power at every 0.5 Hz interval from 0-50 Hz. Visualized here, each dot represents the Cohen's *d* effect size from a one-sample t-test assessing whether the mean correlation surpassed zero. This was done using either **(A)** power averaged across the shown frontal electrodes or **(B)** power averaged across the shown frontal/central/ parietal pool of electrodes.




Table 1. Mean correlations between connectivity fluctuation and frontal oscillatory power.

The table lists the Spearman correlations between fluctuation amplitude and the frontal power, averaged across an established frequency band (e.g., Delta is the average of 1 Hz power, 1.5 Hz power, ... 3.5 Hz power). The effects for Delta, Theta, Beta, and Gamma remain significant if significance testing is instead performed using permutation-testing and with family-wise error rate correction ($p_{\text{corrected}} < .05$). The significance stars represent uncorrected significance levels from a standard t-test. *, $p < .05$; **, $p < .01$; ***, $p < .0001$.

Frequency ranges		Mean correlation	Cohen's <i>d</i>
Delta	1-3.5 Hz	.042 [.027, .056]	1.24 ***
Theta	4-7.5 Hz	.041 [.025, .056]	1.13 ***
Alpha	8-12.5 Hz	.029 [.012, .047]	0.71 **
Beta	13-29.5 Hz	.029 [-.002, .059]	0.41 *
Gamma	30-50 Hz	.037 [.010, .064]	0.59 *

network, the association between high PE and ATL-LPFC connectivity aligns with findings on PE signals flowing between ATL and PFC.^{42,43} Given ATL's role in encoding semantics and abstract information,^{13,44–46} ATL-LPFC interactions could reflect conflicts between incoming stimuli and existing mental models detected in ATL and propagated to PFC,⁴³ and/or signals from the PFC to update categories in ATL.⁴³ Altogether, ventral-dorsal connectivity could underlie interactions between temporal lobe representations and frontoparietal control processes, with OC-IPL connectivity reflecting more perceptual interactions and ATL-LPFC connectivity representing more semantic interactions.

In contrast with high PE, low PE was tied to greater connectivity between posterior and anterior brain regions. On the ventral side of the posterior-anterior network, low PE enhanced OC-ATL connectivity along the ventral stream. This finding is consistent with perception studies showing how the ventral stream encodes stimulus predictions (schemas) that guide processing.⁴⁷ Strong OC-ATL connectivity under low PE could reflect information automatically flowing along the ventral pathway, in the absence of the frontoparietal control processes demanded by high PE. On the dorsal side of the posterior-anterior network, low PE boosted IPL-LPFC connectivity, which points to the relevancy of the frontoparietal network, although this network boasts many functions and its exact role for low PE processing is unclear.⁴⁸ Overall, these posterior-anterior patterns recapitulate well-known dorsal frontoparietal and ventral stream pathways that have been central to the focus of attention and visual processing, respectively.

These functional pathways, which form the backbone of our PE topology, also recapitulate respective structural dorsal (superior longitudinal fasciculus; SLF) and ventral (inferior longitudinal fasciculus, ILF, uncinate fasciculus) pathways. Supporting these inferences, the SLF has been associated with cognitive control and attention,^{49,50} and the ILF has been associated with the semantic processing of objects,⁵¹ faces,⁵² and words.⁵³ It is therefore unsurprising that low PE is associated with such well-worn, structurally connected pathways, and future studies may help to shed light on how major fiber pathways support a global topology of prediction error. Furthermore, that high PE is associated with a connectivity pattern that runs contralaterally to these well-worn pathways is also consistent with the idea that novel connectivity patterns are necessary to form coherent associations between schematically incongruent stimuli. Nonetheless, the consistencies between functional and structural topologies serve to build an inferential bridge linking the cognitive processes at play in PE tasks to resting state.

Our rs-fMRI investigation examined whether resting dynamics resemble the task-defined connectivity signatures of high vs. low PE, independent of the type of stimulus encountered. The resting-state analyses indeed found that, even at rest, participants' brains fluctuated between strong ventral-dorsal connectivity and strong posterior-anterior connectivity, consistent with shifts between states of high and low PE. This conclusion is based on correlative/observational evidence and so may be controversial as it relies on reverse inference; it is impossible to establish that the task networks only reflect PE and that the rs-fMRI patterns are not due to some other variable manipulated by the task. Nonetheless, several key concerns are addressed. First, we established these signatures across two distinct PE designs (Studies 1A/1B), constraining the possible interpretations of the networks. Second, we linked the recruitment of ventral-dorsal and posterior-anterior networks during rest to the engagement of these networks during a PE task in the same participants (Study 2B), creating a sharp parallel between the task effects and resting-state functioning. Although further work is warranted to confirm the function of ventral-dorsal and posterior-anterior during rest, the present research remains, we believe, the most thorough investigation of the type of abstract PE processes that lie at the core of contemporary theories.

A key finding supporting our interpretation is the significant link between individual differences in task-evoked PE responses and resting-state fluctuations. One might initially view the effect size of this correspondence ($r = .021$) as modest. However, this interpretation must be contextualized by the considerable measurement noise inherent in short task-fMRI scans; the split-half reliability of the task contrast was only .08. This low reliability imposes a severe statistical ceiling on any possible task-rest correlation. Therefore, detecting a highly significant ($p < .001$) relationship despite this constraint provides robust evidence for a genuine link. Furthermore, our analytical

approach, which iteratively examined thousands of ROI subsets rather than one aggregated network, was intentionally granular. The goal was not simply to correlate two global measures, but to test for a person-specific anatomical correspondence – that is, whether the specific parts of an individual's network most sensitive to PE during the task are the same parts that fluctuate most strongly at rest. An aggregate analysis would obscure this critical spatial specificity. Taken together, this granular analysis provides compelling evidence for an anatomically consistent fingerprint of PE processing that bridges task-evoked activity and spontaneous resting-state dynamics, strengthening our central claim.

While our claims about resting-state cognition will often be speculative, it is worth considering what minimization of PE could look like in a resting-state scan, or more succinctly during the cognitions that we expect to be active during rest.^{54,55} One possibility is that PE could compel episodic replay or episodic simulation, consistent with the notion that perceptual decoupling and episodic memory retrieval are necessary conditions for the self-generation of a train of thought unrelated to the external environment.^{56,57} For instance, in their stream of thought, a participant's working memory may disengage from their immediate scanner environment, and retrieve information about a recent discussion with a colleague. In turn, the participant may wonder how this discussion resolved (PE detection; high PE), prompting them to retrieve a similar experience to resolve this uncertainty (PE response; low PE). This description casts PE as a spark for episodic replay.¹⁵ In turn, the participant may consider that their comment was rude (PE detection) and predict how their colleague will treat them in the future (PE response). To be clear, these types of mental processes are not exclusive to resting-state scans and could likewise occur when participants experience an event live. Yet, by showing resting dynamics that resemble task-defined PE signatures, we suggest that the currently identified networks are not necessarily limited to lower-level sensory violations. In turn, the present findings support theories about how many aspects of executive function and high-level cognition are rooted in PE minimization.^{4–7,58}

Finally, our third approach for investigating the global PE hypothesis was to examine if high and low PE networks fluctuate sufficiently fast for cognitive operations. As noted before, executive function EEG studies suggest that PE detection^{26,27} and PE minimization^{26,28,29} are separated by a few hundred milliseconds. This result matches evidence that PE modulates Delta (1–4 Hz) and Theta (4–8 Hz) power.³⁰ Thus, we expected high and low PE networks would fluctuate multiple times per second. Indeed, the fMRI-EEG correlations in Study 3 suggest that the ventral-dorsal and posterior-anterior networks fluctuate at a speed around 3–6 Hz. We specifically interpret the fMRI-EEG correlation as reflecting fluctuation speed because we correlated EEG oscillatory power with the fluctuation amplitude computed from fMRI data. Simply correlating EEG power with the average connectivity or the signed difference between posterior-anterior and ventral-dorsal connectivity yields null results (Supplemental Materials 6 [↗](#)), suggesting that this is a very particular association, and viewing it as capturing fluctuation amplitude provides a parsimonious explanation. Yet, this correlation may be interpreted in other ways. For example, resting-state Theta is also a signature of drowsiness,⁵⁹ which may correlate with PE processing, but perhaps should be understood as some other mechanism. Additionally, Theta is widely seen as a sign of dorsal anterior cingulate cortex activity,³⁹ and it is unclear how to reconcile this with our claims about network fluctuations. Nonetheless, as we show with simulations (Supplemental Materials 5 [↗](#)), a correlation between slow fMRI network fluctuations and fast EEG Delta/Theta oscillations is also consistent with a common global neural process oscillating rapidly and eliciting both measures.

Some prior studies have correlated EEG spectral signatures with fMRI functional connectivity in resting-state data,⁶⁰ including by correlating fMRI measures with power, which was likewise averaged over intervals and convolved with a hemodynamic response function.³⁸ However, our work was unique because we specifically correlated power with the absolute difference between connectivity measures after establishing that those measures correspond to dynamic connectivity states. In turn, our interpretation shifts. Whereas these prior papers claimed, for instance, that an

inverse correlation between functional connectivity and alpha indicates ties to alertness,³⁸ we put forth that the correlations with oscillations reveal the frequency of the high/low PE-state fluctuations – an interpretation supported by our simulations.

Interpreting connectome-wide phenomena as changing multiple times per second may seem unlikely, given that much network research is based on fMRI and its slow temporal resolution. However, prior studies incorporating data from other modalities have likewise argued that, despite usually being measured slowly, network states can fluctuate at sub-second frequencies.^{61,62} Theoretically, the identified association between PE topology fluctuations and power in Delta/Theta frequencies agrees with the premise that these fluctuations reflect ongoing human thought. Such thoughts can be conceptualized as repeated manipulations of working memory, which can occur up to five times per second in humans⁶³ and is indexed by Delta/Theta oscillations.⁶⁴ Further research on other variables remains necessary to draw firm conclusions about specific cognitive mechanisms, although these timelines are consistent with the types of functions discussed by psychological PE theories – e.g., working memory states prompting memory recall or affective labeling to minimize PE could conceivably occur in ~200 ms.^{4,15}

To conclude, PE elicits brain-wide effects on connectivity, upregulating ventral-dorsal connectivity while down-regulating posterior-anterior connectivity. These patterns resemble global connectivity signatures seen in resting-state participants, and correlations between fMRI and EEG data yield associations, consistent with participants fluctuating between high-PE (ventral-dorsal) and low-PE (posterior-anterior) states at 3-6 Hz. Although definitively testing these ideas is challenging, given that rs-fMRI is defined by the absence of any causal manipulations, our results provide evidence consistent with PE minimization playing a role beyond stimulus processing. These findings hence empirical foundation for the deluge of theories on the role of PE in abstract cognition. PE minimization may be the *telos* of the brain, while attentional reorienting, episodic retrieval, working memory updating, and other psychological processes may be how PE minimization is accomplished. In this way, the present results link neuroscientific literature while also joining this work with more cognitively focused theories of the mind.

4. Materials and Methods

4.1. Study 1A – Task-fMRI

4.1.1. Participants

The Study 1A task-fMRI analyses were performed using data from a larger study on semantic processing, memory, and aging. Participants were recruited from the local community and screened for no history of brain injury or mild cognitive impairment. Sixty-six healthy participants completed the PE task, which included a younger adult subset ($N = 35$, $M_{age} = 22.8$ [SD = 3.3]; sex: 64% female, 36% male) and an older adult subset ($N = 31$, $M_{age} = 71.5$ [SD = 4.5], 64% female, 36% male). The two groups were not compared here. Of the 66 participants, two encountered scanner malfunctions that caused one run of data loss from each, but these partial datasets were still included in the analyses. There were no exclusions. All participants provided informed consent, the manuscript includes no information that could be used to identify participants, and the research was approved by the Duke University Health System Institutional Review Board.

4.1.2. Task design

Participants viewed pictures of common scenes and objects. Each trial began with a picture of a scene (e.g., a farm or desert) shown for 3 seconds. Then, a blank screen appeared for an average of 3 seconds (jittered), followed by a picture of an object shown for 4 seconds. While the object was displayed, participants used a 4-point scale to indicate how likely it would be to find the object in the scene (1 = “*Very unlikely*”, 4 = “*Very likely*”). Across three runs, participants completed 114 trials: 38 high PE, 38 medium PE, and 38 low PE. For the high-PE condition, the object-scene pairs were incongruent – e.g., a *poker table* (object) likely would not be found in a *prison* (scene) because prisoners’ leisure is restricted. For the medium PE condition, the object and scene were meant to be conceptually orthogonal – e.g., a *bug* could be found in a *desert* but not necessarily. For the low

PE condition, the pairs were congruent – e.g., a *tractor* operates on a *farm*. All participants saw the same 114 objects and 114 scenes but counterbalanced under different configurations. For instance, the *tractor* was low PE for one-third of participants, medium PE for one-third, and high PE for one-third. The Results response times based on the 77.2% of trials where participants gave the expected response: 1 likelihood for high PE and 4 likelihood for low PE. Of note, the medium PE condition yielded slower responses than both conditions ($M = 2.33$ s, $SD = 0.43$) (comparisons of high/low to medium show a large difference, $ps < .0001$, $ds > 1.96$). These response time patterns suggest that the medium condition was more difficult, but low vs. high differences in the brain data would not be attributable to difficulty. An analysis that treats PE as a gradient from low-medium-high would be well positioned to establish that effects genuinely reflect PE rather than task difficulty, which would instead be reflected by strong medium PE connectivity. As part of the aforementioned wider memory study, on a previous day, participants completed a task where they viewed the objects without proceeding scenes. Additionally, on the day following the PE task, participants completed a surprise memory retrieval test. However, neither of these tasks' data were analyzed here.

4.1.3. MRI acquisition

MRI data were collected using a General Electric 3T MR750 scanner and an 8-channel head coil. Anatomical images were acquired using a T1-weighted echo-planar sequence (96 slices at $0.9 \times 0.9 \times 1.9$ mm³). Functional images were acquired using an echo-planar imaging sequence (repetition time = 2000 ms, echo time = 30 ms, field of view = 19.2 cm, 36 oblique slices with voxel dimensions of $3 \times 3 \times 3$ mm). Stimuli were projected onto a mirror at the back of the scanner bore, and responses were recorded using a four-button fiber-optic response box (Current Designs, Philadelphia, PA, USA). Functional resting-state images were collected from the participants using the same parameters (210 volumes, 7 minutes).

The below descriptions of anatomical and functional preprocessing were automatically generated by fMRIPrep with the express intention that users should copy and paste this text into their manuscripts unchanged.

4.1.4. Anatomical data preprocessing

A total of two T1-weighted (T1w) images were found within the input BIDS dataset. All of them were corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection,⁶⁵ distributed with ANTs 2.3.3.⁶⁶ The T1w-reference was then skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as a target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM), and gray-matter (GM) was performed on the brain-extracted T1w using fast (FSL 6.0.5.1).⁶⁷ An anatomical T1w-reference map was computed after registration of two T1w images (after INU-correction) using mri_robust_template (FreeSurfer 7.3.2).⁶⁸

Brain surfaces were reconstructed using recon-all (FreeSurfer 7.3.2),⁶⁹ and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter of Mindboggle.⁷⁰ Volume-based spatial normalization to one standard space (MN152NLin2009cAsym) was performed through nonlinear registration with antsRegistration (ANTs 2.3.3), using brain-extracted versions of both T1w reference and the T1w template. The following template was selected for spatial normalization and accessed with TemplateFlow (23.0.0):⁷¹ ICBM 152 Nonlinear Asymmetrical template version 2009c.⁷²

4.1.5. Functional data preprocessing

For each of the seven BOLD runs found per subject (across all tasks and sessions), the following preprocessing was performed. [The mention of “seven” BOLD runs here reflects other tasks collected as part of the aforementioned memory study.] First, a reference volume and its skull-stripped version were generated using a custom methodology of fMRIPrep. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding

rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt (FSL 6.0.5.1, Jenkinson et al. 2002). BOLD runs were slice-time corrected to 0.972s (0.5 of slice acquisition range 0s-1.94s) using 3dTshift from AFNI.⁷³

The BOLD time series (including slice-timing correction when applied) were resampled onto their original, native space by applying the transforms to correct for head motion. These resampled BOLD time-series will be referred to as preprocessed BOLD in original space, or just preprocessed BOLD. The BOLD reference was then co-registered to the T1w reference using bbregister (FreeSurfer) which implements boundary-based registration.⁷⁴ Co-registration was configured with six degrees of freedom.

Several confounding time series were calculated based on the preprocessed BOLD: framewise displacement (FD), DVARS, and three region-wise global signals. FD was computed using two formulations following Power et al. (absolute sum of relative motions)⁷⁵ and Jenkinson (relative root mean square displacement between affines).⁷⁶ FD and DVARS are calculated for each functional run, both using their implementations in Nipype (following the definitions by Power et al.⁷⁵). The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Principal components are estimated after high-pass filtering the preprocessed BOLD time series (using a discrete cosine filter with 128s cut-off) for the two CompCor variants: temporal (tCompCor) and anatomical (aCompCor). tCompCor components are then calculated from the top 2% variable voxels within the brain mask. For aCompCor, three probabilistic masks (CSF, WM, and combined CSF+WM) are generated in anatomical space. The implementation differs from that of Behzadi et al. in that instead of eroding the masks by two pixels on BOLD space, a mask of pixels that likely contain a volume fraction of GM is subtracted from the aCompCor masks. This mask is obtained by dilating a GM mask extracted from the FreeSurfer's aseg segmentation, and it ensures components are not extracted from voxels containing a minimal fraction of GM. Finally, these masks are resampled into BOLD space and binarized by thresholding at 0.99 (as in the original implementation). Components are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the k components with the largest singular values are retained, such that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration.

The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each.⁷⁷ Additional nuisance time series are calculated by means of principal components analysis of the signal found within a thin band (crown) of voxels around the edge of the brain, as proposed by Patriat, Reynolds, and Birn.⁷⁸ The BOLD time series were resampled into standard space, generating a preprocessed BOLD run in MNI152Nlin2009cAsym space. First, a reference volume and its skull-stripped version were generated using a custom methodology of fMRIPrep. All resamplings can be performed with a single interpolation step by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using antsApplyTransforms (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels.⁷⁹ Non-gridded (surface) resamplings were performed using mri_vol2surf (FreeSurfer). Many internal operations of fMRIPrep use Nilearn 0.9.1,⁸⁰ mostly within the functional processing workflow. For more details of the pipeline, see the section corresponding to workflows in fMRIPrep's documentation.

4.1.6. Functional connectivity

Calculating ROI-ROI functional connectivity required first measuring each voxel's BOLD response in each trial. This was done using first-level general linear models with the Least Squares Separate approach by Mumford et al.,⁸¹ which involves fitting a separate regression for each trial. The regression included a boxcar signal spanning the trial's object presentation period, and this regressor's coefficient represents the trial's BOLD response. The regression also included a boxcar

signal covering the presentation time of all other stimuli (i.e., every other object and every scene). Additionally, the linear models included six translation/rotation regressors, three other covariates for head motion (FD, DVARS, and RMSD), and covariates for mean global, white-matter, and cerebrospinal signals.

The voxelwise measures of activity (beta coefficients) were organized and averaged in terms of the 210 neocortical ROIs of the Brainnetome Atlas (i.e., the cortex sans the hippocampus); the exact ROIs employed for the main analyses involving four anatomical quadrants are shown in [Supplemental Figure S1](#). ROI-ROI functional connectivity was calculated independently for the high/medium/low PE conditions as the Pearson correlation between two vectors describing trial-wise activation estimates for each pair of ROIs (beta-series correlations).⁸² Notably, the Brainnetome Atlas provides anatomical labels for each ROI. However, it does not provide an “ATL” label, which was defined here as the eight most anterior temporal gyral ROIs, as our group has previously done in semantic processing research.⁴⁵ The data on ROI-ROI connectivity was used for the modularity analysis (described next) and for the classification-based analyses, which evaluated connectivity effects while assessing statistical significance with multiple-comparison correction (methods and results provided in [Supplemental Materials 2](#)). The analyses were also conducted and reproduced using the Schaefer atlas ([Supplemental Materials 7](#)).

4.1.7. Modularity

Our goal in testing modularity and our approach contrasts typical modularity studies, which identify task-independent patterns by examining the connectivity averaged across conditions and participants. Here, modularity was tested using a beta matrix representing the effect of PE on each connection. For each edge, a regression was fit predicting connectivity strength, while treating PE as a continuous predictor (low = -1, medium = 0, high = +1). The regression was analogous to a paired t-test but with three points rather than two. That is, each participant contributed one sample per PE condition, and the regression included a dummy variable for each participant (66 participants = 65 subject variables). This was separately done for each edge in the connectome to construct the aforementioned beta matrix. Evaluating modularity required preparing binary matrices. Specifically, the beta matrix was filtered to only the 5% of edges with the most positive coefficients (high > low PE) or the 5% of edges with the most negative coefficients (low > high PE). The high-filtered and low-filtered binary matrices were then submitted to the Leiden algorithm – an optimized form of Louvain’s algorithm, which addresses the risk of disconnected communities and has become increasingly recommended.⁸³

4.2. Study 1B – Task-fMRI

4.2.1. Participants and task design

Analyses were conducted on 1,000 participant’s gambling task data from the S1200 HCP dataset ($M_{age} = 28.7$, 57% female, 43% male); this dataset is described by Barch et al.³⁴ The task is based on the design by Delgado et al.⁸⁴ with some changes to the timing procedures.^{34,85} In each trial, participants were shown a question mark and asked to press a button to make a choice (“2” or “3”) within 1.5 s. After making a choice, participants immediately saw a fixation cross for the rest of the 1.5 s period. Then, the fixation continued for another 1 s. Afterward, participants saw an upward green arrow in 43.75% of trials (signifying a \$1.00 reward) or a downward red arrow in 43.75% of trials (signifying a \$0.50 loss). In 12.5% of trials, participants encountered a screen indicating a neutral outcome (no gain or loss). Participants were told that their choice determined the outcome, but the choice actually had no effect, and each trial’s outcome was predetermined. The monetary difference between rewards/losses was meant to account for loss aversion causing losses to be perceived as worse than gains. Participants completed two runs of the task, each with 32 trials separated into four blocks.

4.2.2. Task design and modeling PE

To model PE, analyses defined a participant’s prediction in each trial as an exponential moving average, $EMA[t] = \alpha \times outcome + (1 - \alpha) \times EMA[t - 1]$. Outcomes were defined as +1, losses as -1, and neutral outcomes as 0. The learning rate was defined at 0.3, although testing 0.2 or 0.5 showed that

this does not change whether the Connectivity Direction $\times$ PE interaction effect is significant (still $p < .0001$). The prediction at the start of each run was defined as zero. PE was computed as the absolute difference between a trial's outcome and the prediction based on previous trials.

This continuous measure of PE was converted into binary high-PE and low-PE conditions to support the functional connectivity analyses. Binarization involved dividing trials based on the median PE. This median split was done separately for gain and loss trials to ensure they were in equal quantities across the high/low PE conditions. The neutral trials were omitted from the connectivity analysis as they were few in number. The first trial of the run was also omitted as there were no prior trials on which to base a meaningful prediction. Although a median split will lead to roughly equal numbers of trials in the high/low PE conditions, to ensure that these two had exactly the same number of trials (rather than possibly one condition having one or two more), trials at the end of the task were omitted for balancing as needed. To validate these conditions and the binary split, participants' reaction times were compared post-high-PE-loss and post-low-PE-loss trials using a paired t-test, which was expected to show post-error slowing.³¹

4.2.3. Imaging protocol and preprocessing

Details on the imaging protocols are provided by Van Essen et al.³³ Of note, the functional voxels were small (2 mm isotropic) and the temporal resolution was very fast (0.72 s TR time). The downloaded dataset had already been "minimally preprocessed" by the HCP, which included the removal of spatial artifacts/distortions, segmentation, within-subject registration, and normalization to a common brain template.⁸⁶ Although the imaging and preprocessing protocols differ in Study 1B from Study 1A, this is acceptable given that the goal for Study 1B was not to be a precise replication. Additionally, adhering to the already-performed preprocessing for a public dataset benefits reproducibility.

4.2.4. Functional connectivity

Analysis of functional connectivity paralleled that of Study 1A but with two changes. First, single-trial beta estimates were computed for each trial using a Least Squares All rather than a Least Squares Separate approach.⁸¹ The Study 1B dataset is much larger than the Study 1A dataset, and the Least Squares All approach is much less computationally intensive; preliminary analyses suggested that the produced beta values are largely similar. Second, each single-trial-beta regression model included the motion parameters provided by the HCP (six translation/rotation regressors and their derivatives) along with a covariate for mean global signal; note that the Study 1A preprocessing did not include motion derivatives, but this was done here because the HCP provided them, and we sought consistency with their standard procedures. Functional connectivity estimates were calculated separately for high-PE and low-PE conditions, for each pair among the 210 ROIs by using beta-series connectivity, as above.⁸² The contrast matrix was correlated with that of Study 1A (Figure 2A [↗](#)). The Results section reports the correlation taken between the bottom triangle of a 46x46 matrix, wherein ROIs were averaged based on their Brainnetome label (e.g., averaging all L MFG ROIs to a single region); this is the matrix shown in Figures 2A & 3A [↗](#). This averaging yields a more stable measure, given that individual signal within ROIs can vary wildly due to noise. As reported, the correlation was fairly strong ($r = .27$, $p < .0001$). If the correlation is instead taken with respect to the 210x210 matrix, there remains a clear correlation albeit a weaker one ($r = .14$, $p < .0001$).

4.3. Study 2A – Resting-state-fMRI

4.3.1. Data and preprocessing

Rs-fMRI data were collected from the 66 participants from Study 1A with the same acquisition protocols, and also as above, with the six motion covariates regressed out. However, whereas the single-trial models for the task-fMRI analysis included a global-signal covariate, this was not done for the rs-fMRI analysis because global signal regression would strongly bias ROI-ROI and edge-edge correlations toward negative values.^{87,88} This would be an issue for the rs-fMRI analyses, as they focused on pure correlations rather than contrasts between conditions. Yet, without global signal regression, connectivity is strongly biased toward positive correlations.⁸⁹ Hence, our

resting-state analyses analyzed only a subset of the brain (OC, ATL, IPL, LPFC) and removed effects linked to connectivity with regions outside these subsets. This was done by data cleaning with temporal CompCor (see Section 4.1.5 [\[35\]](#)) and with other procedures detailed in the next subsection.

4.3.2. Edge-edge correlations

Our analysis of the rs-fMRI data focused on how connectivity strength fluctuates over time and how groups of edges fluctuate in sync or out of sync. Of the emerging techniques for addressing dynamic connectivity,⁹⁰ we chose to employ *time-varying connectivity* and *edge-edge correlations*.^{35,91,92} Whereas traditional connectivity studies have focused on the correlation between activation time series extracted from pairs of ROIs, this direction models connectivity itself as a time series and quantifies the relations between edges' time series. Time-varying connectivity is modeled by decomposing the traditional Pearson correlation.^{35,91,92} That is, a Pearson correlation entails z-standardizing two vectors and taking the average of their element-wise products. For time-varying connectivity, the element-wise products are calculated but not averaged, which means that the two input time series have created a new time series of the same length. For example, in the first volume, if OC activity is high and IPL connectivity is high, the OC-IPL connectivity in the first volume is said to be strong. After time-varying connectivity has been calculated for each edge in the connectome, the values were averaged in terms of the major anatomical quadrants shown in Figure 3A [\[35\]](#). For example, the OC-IPL entry represents the average across every possible pair of OC and IPL ROIs.

In turn, Pearson correlations can be computed between pairs of the edge sets (e.g., correlating OC-IPL $\times$ OC-ATL). This edge-edge correlation represents how the two edges shift over time in sync (positive correlation) or out of sync (negative correlation). As mentioned in the prior subsection, these analyses required accounting for global signal. Preliminary tests showed that if global signal effects are unaccounted, this would promote a very positive correlation between edges, particularly for pairs of edges linked to the same ROI (i.e., global signal particularly enhances correlations between ROI₁-ROI₂ $\times$ ROI₁-ROI₃). Yet, if traditional global signal regression was done during preprocessing, this would induce a very negative correlation. Hence, to achieve an unbiased estimate, rather than measuring edge-edge connectivity as plain Pearson correlations, they were measured as partial correlations. These partial correlations essentially measured edge-edge fluctuations relative to each edge's baseline correlation with the rest of the connectome. For this, analyses gathered the 144 neocortical ROIs outside the four quadrants (non-OC, -IPL, -ATL, or -LPFC) and calculated the mean time-varying connectivity between each octant with these 144 ROIs (e.g., the mean time-varying connectivity between the L OC and the rest of the brain) (a octant is the unilateral version of a quadrant; see Figure 4 [\[35\]](#)). These eight measures of time-varying connectivity, one with each octant, were then partialled out of each edge-edge correlation. This partial correlation strategy is analogous to existing methods for unbiased global signal regression designed for classic ROI-ROI connectivity, which regress out signal from regions not analyzed. This approach appeared to be successful, given that in the correlation matrix shown in Figure 4 [\[35\]](#), the mean correlation is $r = .02$, suggesting the absence of pressure toward positive or negative correlations.

Edge-edge correlations were measured between each of the twelve connections shown in the Figure 4 [\[35\]](#) cuboid to produce an edge-edge correlation matrix, on which modularity was tested. In this context, modules constitute sets of edges that are all strong in the same volumes (positive edge-edge correlations within modules). Modularity also is based on fluctuation between modules (negative correlations between modules). For the modularity analysis, the variant of Louvain's algorithm was used as above but now, rather than just filtering the matrix to the 5% of strongest edges, the matrix was filtered to the 50% of strongest edges. This much more liberal approach was necessary as the edge-edge correlation matrix was much smaller (i.e., for a 12 $\times$ 12 matrix, using just 5% of edges would mean that most rows do not have a single preserved entry).

4.4. Study 2B – Resting-state-fMRI

4.4.1. Data and preprocessing

Rs-fMRI data were used from the 1,000 HCP participants of Study 1B, which were collected using the same acquisition protocols and had already been submitted to the HCP minimal preprocessing pipeline. Preprocessing was meant to be similar to that of Study 2A, including regressing out the six motion parameters and using temporal CompCor but not global signal regression. Afterward, the rs-fMRI data were submitted to the edge-edge modularity analysis described above for Study 2A, attempting to reproduce the original modules

4.4.2. Analysis

Study 2B also incorporated a new analysis, considering subsets of the four anatomical quadrants (visualized in [Figure 5](#)). Specifically, the analysis was predicated on repeatedly selecting a pair of bilateral ROIs (one left and its symmetric right) at random from the four quadrants. This selection yields eight ROIs, and the strength of the Study 1B task-fMRI effect and the strength of the rs-fMRI fluctuation were computed for each participant. The task-fMRI effect constituted the Connectivity Direction $\times$ PE interaction, $(PA_{\text{High PE}} - PA_{\text{Low PE}}) - (VD_{\text{High PE}} - VD_{\text{Low PE}})$; these values refer to a single participant's posterior-anterior (PA) and ventral-dorsal (VD) effects for a given condition. The rs-fMRI effect was computed using time time-varying (see above, [Section 4.3.2](#)), and the effect was defined as the mean of the TR-by-TR absolute difference between posterior-anterior and ventral-dorsal connectivity, $E[|z(PA) - z(VD)|]$; posterior-anterior and ventral-dorsal connectivity were each z-scored to ensure that the measure does not merely reflect overall variability in participants' BOLD data. This absolute difference measure would later be used and discussed more prominently in Study 3, and the simulations (detailed in full in [Supplemental Materials 5](#)) demonstrate and provide intuition for why this quantity captures the degree of fluctuation between two states in resting-state fMRI data.

The task-fMRI and rs-fMRI measures were computed for each participant and for each selection of eight ROIs (3,432 in total). They were then submitted to correlations and group-level analyses. Specifically, within-ROI-set (subject-by-subject) correlation was computed for each set of ROIs, and the correlation coefficients were likewise submitted to a one-sample t-test assessing whether the mean correlation significantly surpassed zero. In addition, a multilevel regression was performed which accounted for subjects' overall task-resting association with respect to the full quadrants – i.e., to show that the results indeed reflect within-subject heterogeneity in the anatomy of their PE networks rather than just their intensity. All 3,432 million data points (3,432 sets $\times$ 1,000 participants) were submitted to the following multilevel regression: $\text{rest} \sim 1 + \text{task} + (1 + \text{task} | \text{ROI set}) + (1 | \text{subject})$. The use of random slopes for the ROI-set-level grouping makes this analogous to the t-test done on the correlations for each ROI set. The results report the significance of the *task* fixed effect, computed using the *lmerTest* R package.⁹³

4.4.3. Links between network fluctuations and behavior

We considered whether the extent of PE-related network expression states during resting-state is behaviorally relevant. We specifically investigated whether individual differences in the overall magnitude of resting-state fluctuations could predict individual difference measures, provided with the HCP dataset. This yielded a significant association with age, whereby older participants tended to display weaker fluctuations. However, associations with cognitive measures were limited. A full description of these analyses is provided in [Supplemental Materials 8](#).

4.5. Study 3 – Rs-fMRI-EEG

4.5.1. Participants and preprocessing

The fMRI-EEG analyses used the public dataset provided by Telesford and colleagues³⁷, which consists of fMRI and EEG data simultaneously recorded from 22 participants ($M_{\text{age}} = 36.8$ [22–51], 50% male, although one participant's data was unusable, see below). Participants generally completed two 10-minute resting-state scans. Participants also completed two 10-minute scans where they viewed the *Inscapes* video – an animation featuring shapes that slowly move and

transform, which is meant to mirror resting-state conditions while discouraging movement and drowsiness.⁹⁴ Data for one resting-state scan and one Inscapes scan were missing from five participants. Also, in our analyses, we encountered difficulties temporally aligning ten scans' fMRI data and EEG data. Hence, the analyses were done using 68 scans, spread across 21 participants. The dataset also includes scans where participants viewed movies with higher cognitive demands (e.g., Hollywood films), but these were not used. For ease of reproducibility, the analyses used the preprocessed version of the dataset provided by Telesford et al.³⁷; see their report for a description of the acquisition and preprocessing procedures. Of note, the voxel size was 3.469 x 3.469 x 3.330 mm and the TR length was 2100 ms. The preprocessing by Telesford et al. notably included gradient artifact removal, which introduced discontinuities to the EEG data. For the fMRI data, we further preprocessed the data using temporal CompCor, as otherwise, as was done for the resting-state Studies 2A and 2B and as without CompCor, the edges all strongly correlated.

4.5.2. Time-frequency analysis

The time-frequency analyses were done using electrode data recorded at 5000 Hz, which we downsampled to 250 Hz. The analyses first targeted five midline frontal electrodes (F3, F1, Fz, F2, F4; BioSemi64 layout), given that this frontal row is typically the focus of executive-function PE research on oscillations.^{30,40,41} Second, given the large size of the posterior-anterior/ventral-dorsal fMRI pattern, analyses were also done with the frontal and the frontocentral (FC3, FC1, FCz, FC2, FC4), central (C3, C1, Cz, C2, C4), centroparietal (CP3, CP1, CPz, CP2, CP4), and parietal (P3, P1, Pz, P2, P4) electrodes. As part of their preprocessing, Telesford et al.³⁷ omitted some electrodes' data (different across participants, $M = 3.2$ electrodes omitted). For the time-frequency analysis, complex Morlet seven-cycle wavelets were defined for each frequency band 0.5-50 Hz in 100 linearly spaced 0.5 Hz steps. Each electrode's time series was convolved using the wavelets, and the resulting spectrograms were averaged across electrodes for each participant. The discontinuities in the data generated during preprocessing, as mentioned above, create border artifacts in the time-frequency data, and the next section describes how these were dealt with.

4.5.3. Fluctuation-power correlation

Using the fMRI data, time-varying connectivity was calculated for each edge among the quadrant pairs illustrated earlier (e.g., edges between each of the OC ROIs with each of the ATL ROIs). Connectivity was averaged by quadrant pair as above, yielding overall time-varying measurements for OC-ATL, IPL-LPFC, OC-IPL, and ATL-LPFC connectivity. Partialling out edge-edge correlations with the rest of the brain, as done for the rs-fMRI studies, was not necessary here, as our focus was not on comparing positive vs. negative correlations while computing modularity. At each TR, the absolute difference ($|PA - VD|$) was calculated between posterior-anterior connectivity (OC-ATL + IPL-LPFC) and ventral-dorsal connectivity (OC-IPL + ATL-LPFC). This time series of absolute differences represents how the intensity of the posterior-anterior vs. ventral-dorsal fluctuations shifts over time. That is, fluctuation amplitude was considered equivalently high in TRs when (i) posterior-anterior connectivity was strong while ventral-dorsal was weak or (ii) ventral-dorsal connectivity was strong while posterior-anterior was weak. This measure is discussed in [Supplemental Materials 5](#) [↗](#), and simulations are provided showing how this measure captures the amplitude of an underlying neural oscillator, even if the oscillator's frequency outpaces fMRI's temporal resolution. This absolute difference was correlated with oscillatory power at different frequency bands. The simulations in [Supplemental Materials 5](#) [↗](#) also show why measuring such correlations shed light on the frequency underlying the fMRI connectivity fluctuations.

In the EEG data, the measures of oscillatory power were organized by TR and averaged. For instance, at the fourth volume, the oscillatory power of a given frequency was averaged across 6300-8399 ms. That is, the spectrograms' temporal resolutions were reduced from 250 Hz to 1/2.1 Hz. To account for EEG discontinuities introduced during preprocessing, and the potential for discontinuities to create border artifacts, TR intervals overlapping with discontinuous segments were identified along with their neighbors (e.g., if TR 4 overlaps with a discontinuous 500 ms segment, then TRs 3-5 are marked). Oscillatory power in these TRs was interpolated linearly (e.g.,

for TRs 3-5, interpolated using TRs 2 and 6). Finally, before oscillatory power was correlated with the connectivity-fluctuation differences, the interpolated spectrograms were convolved with respect to time using a hemodynamic response function. This accounts for the lag in the fMRI data between neural activation and the later BOLD response. After organization, the analysis simply consisted of measuring the Spearman correlation between the connectivity-fluctuation difference and oscillatory power at each frequency bin (e.g., correlate 1 Hz, 1.5 Hz, etc.). For each participant, this correlation was measured for each of their four scans, then Fisher z-transformed, and averaged. Each participant's mean correlation was submitted to a one-sample t-test assessing whether the correlation significantly surpassed zero.

To confirm the significance of our fMRI-EEG correlations with a non-parametric approach, we performed a group-level permutation-test. For each of 1,000 permutations, we phase-randomized the fMRI fluctuation amplitude time series. Specifically, we randomized the Fourier phases of the |PA-VD| series (within run), while retaining the original amplitude spectrum; inverse transforms yielded real surrogates with the same power spectrum. This procedure breaks the true temporal relationship between the fMRI and EEG data while preserving its structure. We then re-computed the mean Spearman correlation for each frequency band using this phase-randomized data. We evaluated significance using a family-wise error correction approach that accounts for us analyzing five oscillatory power bands. We thus create a null distribution composed of the *maximum* correlation value observed across all frequency bands from each permutation. Our observed correlations were then tested for significance against this distribution of maximums.

Data availability

The datasets used for Studies 1B and 2B are publicly available from the Human Connectome Project website (<https://db.humanconnectome.org>). The dataset for Study 3 is also publicly available (https://fcon_1000.projects.nitrc.org/indi/retro/nat_view.html). The data collected for Studies 1A and 2A will be shared upon request via email to the corresponding author, Paul C. Bogdan, or principal investigators Roberto Cabeza and Simon W. Davis. Note that the results supporting the main conclusions of Studies 1A and 2A were each replicated using the freely available Study 1B and 2B datasets.

Code availability

The Python and R code used for every study has been made publicly available (<https://github.com/paulcbogdan/FlucPE>). This repository also contains the result images directly generated via the analysis code.

Acknowledgements

This work was supported by the National Institutes of Health, R01-AG066901, R01-AG075417.

Additional files

[Supplemental Materials](#)

Additional information

Funding

Funder	Grant reference number	Author
HHS National Institutes of Health (NIH)	R01-AG066901	Roberto Cabeza Simon W Davis
HHS National Institutes of Health (NIH)	R01-AG075417	Roberto Cabeza Simon W Davis

Author ORCID iDs

Paul C Bogdan: <https://orcid.org/0000-0002-4362-6084>

Shenyang Huang: <https://orcid.org/0000-0002-5494-925X>

Simon W Davis: <https://orcid.org/0000-0002-5943-0756>

References

1. **Rao R. P.**, Ballard D. H (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* **2**:79-87
2. **Walsh K. S.**, McGovern D. P., Clark A., O'Connell R. G (2020) Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Ann N Y Acad Sci* **1464**:242-268
3. **Niv Y.**, Schoenbaum G (2008) Dialogues on prediction errors. *Trends Cogn Sci* **12**:265-272
4. **Barrett L. F.** (2017) The theory of constructed emotion: an active inference account of interoception and categorization (vol 12, pg 17, 2017). *Soc Cogn Affect Neurosci* **12**:1833-1833
5. **Ondobaka S.**, Kilner J., Friston K. J (2017) The role of interoceptive inference in theory of mind. *Brain Cogn* **112**:64-68
6. **Theriault J. E.**, Young L., Barrett L. F (2020) The sense of should: A biologically-based framework for modeling social pressure. *Phys Life Rev*
7. **Lühe Von Der**, et al. (2016) Interpersonal predictive coding, not action perception, is impaired in autism. *Philos Trans R Soc B Biol Sci* **371**
8. **Bowling J. T.**, Friston K. J., Hopfinger J. B (2020) Top-down versus bottom-up attention differentially modulate frontal-parietal connectivity. *Hum Brain Mapp* **41**:928-942
9. **Den Ouden H. E.**, Kok P., De Lange F. P. (2012) How prediction errors shape perception, attention, and motivation. *Front Psychol* **3**
10. **Bein O.**, Duncan K., Davachi L (2020) Mnemonic prediction errors bias hippocampal states. *Nat Commun* **11**:3451
11. **Sinclair A. H.**, Manalili G. M., Brunec I. K., Adcock R. A., Barense M. D (2021) Prediction errors disrupt hippocampal representations and update episodic memories. *Proc Natl Acad Sci* **118**:e2117625118
12. **Heilbron M.**, Armeni K., Schoffelen J.-M., Hagoort P., De Lange F. P (2022) A hierarchy of linguistic predictions during natural language comprehension. *Proc Natl Acad Sci* **119**:e2201968119
13. **Matsumoto A.**, Soshi T., Fujimaki N., Ihara A. S (2021) Distinctive responses in anterior temporal lobe and ventrolateral prefrontal cortex during categorization of semantic information. *Sci Rep* **11**
14. **Davis T.**, Love B. C., Preston A. R (2012) Learning the exception to the rule: Model-based fMRI reveals specialized representations for surprising category members. *Cereb Cortex* **22**:260-273
15. **Barron H. C.**, Aukstulewicz R., Friston K. J (2020) Prediction and memory: A predictive coding account. *Prog Neurobiol* **192**
16. **Friston K. J.** (2010) The free-energy principle: a unified brain theory?. *Nat Rev Neurosci* **11**:127-138
17. **Feldman H.**, Friston K. J. (2010) Attention, uncertainty, and free-energy. *Front Hum Neurosci* **4**:215
18. **Friston K. J.**, et al. (2015) Active inference and epistemic value. *Cogn Neurosci* **6**:187-214
19. **Brown R. L.**, Brusse C., Huebner B., Pain R (2020) Unification at the cost of realism and precision. *Behav Brain Sci* **43**
20. **Dołęga K.**, Schlicht T., Dennett D. C (2020) Explaining or redefining mindreading?. *Behav Brain Sci* **43**
21. **Zefferman M. R.**, Smaldino P. E (2020) Integrating models of cognition and culture will require a bit more math. *Behav Brain Sci* **43**
22. **Corlett P. R.**, Mollick J. A., Kober H (2022) Meta-analysis of human prediction error for incentives, perception, cognition, and action. *Neuropsychopharmacology* **47**:1339-1349

23. Fouragnan E., Retzler C., Philiastides M. G (2018) Separate neural representations of prediction error valence and surprise: Evidence from an fMRI meta-analysis. *Hum Brain Mapp* **39**:2887-2906
24. Gonzalez-Castillo J., et al. (2015) Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns. *Proc Natl Acad Sci* **112**:8762-8767
25. Smith S. M., et al. (2009) Correspondence of the brain's functional architecture during activation and rest. *Proc Natl Acad Sci* **106**:13040-13045
26. Chase H. W., Swainson R., Durham L., Benham L., Cools R (2011) Feedback-related negativity codes prediction error but not behavioral adjustment during probabilistic reversal learning. *J Cogn Neurosci* **23**:936-946
27. Holroyd C. B., Coles M. G (2002) The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev* **109**
28. Polich J (2007) Updating P300: an integrative theory of P3a and P3b. *Clin Neurophysiol* **118**:2128-2148
29. Shalgi S., Barkan I., Deouell L. Y (2009) On the positive side of error processing: error-awareness positivity revisited. *Eur J Neurosci* **29**:1522-1532
30. Cavanagh J. F (2015) Cortical delta activity reflects reward prediction error and related behavioral adjustments, but at different times. *NeuroImage* **110**:205-216
31. Notebaert W., et al. (2009) Post-error slowing: an orienting account. *Cognition* **111**:275-279
32. Zalesky A., Fornito A., Bullmore E. T (2010) Network-based statistic: identifying differences in brain networks. *Neuroimage* **53**:1197-1207
33. Van Essen D. C., et al. (2013) The WU-Minn human connectome project: an overview. *Neuroimage* **80**:62-79
34. Barch D. M., et al. (2013) Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage* **80**:169-189
35. Faskowitz J., Esfahlani F. Z., Jo Y., Sporns O., Betzel R. F (2020) Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nat Neurosci* **23**:1644-1654
36. Finn E. S., et al. (2015) Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nat Neurosci* **18**:1664-1671
37. Telesford Q. K., et al. (2023) An open-access dataset of naturalistic viewing using simultaneous EEG-fMRI. *Sci Data* **10**
38. Chang C., Liu Z., Chen M. C., Liu X., Duyn J. H (2013) EEG correlates of time-varying BOLD functional connectivity. *Neuroimage* **72**:227-236
39. Cavanagh J. F., Frank M. J (2014) Frontal theta as a mechanism for cognitive control. *Trends Cogn Sci* **18**:414-421
40. Hoy C. W., Steiner S. C., Knight R. T (2021) Single-trial modeling separates multiple overlapping prediction errors during reward processing in human EEG. *Commun. Biol* **4**
41. Neo P. S.-H., Shadli S. M., McNaughton N., Sellbom M (2024) Midfrontal theta reactivity to conflict and error are linked to externalizing and internalizing respectively. *Personal Neurosci* **7**:e8
42. Aukstulewicz R., et al. (2017) The cumulative effects of predictability on synaptic gain in the auditory processing stream. *J Neurosci* **37**:6751-6760
43. Chao Z. C., Takaura K., Wang L., Fujii N., Dehaene S (2018) Large-scale cortical networks for hierarchical prediction and prediction error in the primate brain. *Neuron* **100**:1252-1266.e3
44. Davis S. W., et al. (2021) Visual and semantic representations predict subsequent memory in perceptual and conceptual memory tests. *Cereb Cortex* **31**:974-992
45. Deng L., et al. (2021) Age-related dedifferentiation and hyperdifferentiation of perceptual and mnemonic representations. *Neurobiol Aging* **106**:55-67

46. Binder J. R., et al. (2011) Mapping anterior temporal lobe language areas with fMRI: a multicenter normative study. *Neuroimage* **54**:1465-1475
47. Shao Z., Beck D. M (2024) Is attention necessary for the representational advantage of good exemplars over bad exemplars?. *Eur J Neurosci* **59**:2353-2372
48. Scolari M., Seidl-Rathkopf K. N., Kastner S (2015) Functions of the human frontoparietal attention network: Evidence from neuroimaging. *Curr Opin Behav Sci* **1**:32-39
49. Bennett I. J., Motes M. A., Rao N. K., Rypma B (2012) White matter tract integrity predicts visual search performance in young and older adults. *Neurobiol Aging* **33**:433-e21
50. Lissek S., Schlaffke L., Tegenthoff M (2024) Microstructural properties of attention-related white matter tracts are associated with the renewal effect of extinction. *Behav Brain Res* **471**
51. Papagno C., et al. (2023) Deficits in naming pictures of objects are associated with glioma infiltration of the inferior longitudinal fasciculus: A study with diffusion MRI tractography, volumetric MRI, and neuropsychology. *Hum Brain Mapp* **44**:4011-4027
52. Burkhardt E., et al. (2023) The central role of the left inferior longitudinal fasciculus in the face-name retrieval network. *Hum Brain Mapp* **44**:3254-3270
53. Shin J., et al. (2019) Inferior longitudinal fasciculus' role in visual processing and language comprehension: a combined MEG-DTI study. *Front Neurosci* **13**
54. Gonzalez-Castillo J., Kam J. W., Hoy C. W., Bandettini P. A (2021) How to interpret resting-state fMRI: ask your participants. *J Neurosci* **41**:1130-1141
55. Pezzulo G., Zorzi M., Corbetta M (2021) The secret life of predictive brains: what's spontaneous activity for?. *Trends Cogn Sci* **25**:730-743
56. Liu Y., Dolan R. J., Kurth-Nelson Z., Behrens T. E (2019) Human replay spontaneously reorganizes experience. *Cell* **178**:640-652
57. Schacter D. L., Addis D. R., Buckner R. L (2008) Episodic simulation of future events: Concepts, data, and applications. *Ann N Y Acad Sci* **1124**:39-60
58. De Filippo R., Schmitz D (2024) Synthetic surprise as the foundation of the psychedelic experience. *Neurosci Biobehav Rev* **105538**
59. Strijkstra A. M., Beersma D. G., Drayer B., Halbesma N., Daan S (2003) Subjective sleepiness correlates negatively with global alpha (8–12 Hz) and positively with central frontal theta (4–8 Hz) frequencies in the human resting awake electroencephalogram. *Neurosci Lett* **340**:17-20
60. Allen E., Damaraju E., Eichele T., Wu L., Calhoun V. D (2018) EEG signatures of dynamic functional network connectivity states. *Brain Topogr* **31**:101-116
61. Hunyadi B., Woolrich M. W., Quinn A. J., Vidaurre D., De Vos M (2019) A dynamic system of brain networks revealed by fast transient EEG fluctuations and their fMRI correlates. *Neuroimage* **185**:72-82
62. Baker A. P., et al. (2014) Fast transient networks in spontaneous human brain activity. *eLife* **3**:e01867 <https://doi.org/10.7554/eLife.01867>
63. Oberauer K., Souza A. S (2020) How fast can people refresh and rehearse information in working memory?. *Mem Cognit* **48**:1442-1459
64. de Vries I. E., Slagter H. A., Olivers C. N (2020) Oscillatory control over representational states in working memory. *Trends Cogn Sci* **24**:150-162
65. Tustison N. J., et al. (2010) N4ITK: improved N3 bias correction. *IEEE Trans Med Imaging* **29**:1310-1320
66. Avants B. B., Epstein C. L., Grossman M., Gee J. C (2008) Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med Image Anal* **12**:26-41
67. Zhang Y., Brady M., Smith S (2001) Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging* **20**:45-57

68. Reuter M., Rosas H. D., Fischl B (2010) Highly accurate inverse consistent registration: a robust approach. *Neuroimage* **53**:1181-1196
69. Dale A. M., Fischl B., Sereno M. I (1999) Cortical surface-based analysis: I. Segmentation and surface reconstruction. *Neuroimage* **9**:179-194
70. Klein A., et al. (2017) Mindboggling morphometry of human brains. *PLoS Comput Biol* **13**:e1005350
71. Ciric R., et al. (2022) TemplateFlow: FAIR-sharing of multi-scale, multi-species brain models. *Nat Methods* **19**:1568-1571
72. Fonov V. S., Evans A. C., McKinstry R. C., Almlí C. R., Collins D (2009) Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage* **47**
73. Cox R. W., Hyde J. S (1997) Software tools for analysis and visualization of fMRI data. *NMR Biomed Int J Devoted Dev Appl Magn Reson Vivo* **10**:171-178
74. Greve D. N., Fischl B (2009) Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* **48**:63-72
75. Power J. D., et al. (2014) Methods to detect, characterize, and remove motion artifact in resting state fMRI. *Neuroimage* **84**:320-341
76. Jenkinson M., Bannister P., Brady M., Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**:825-841
77. Satterthwaite T. D., et al. (2013) An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage* **64**:240-256
78. Patriat R., Reynolds R. C., Birn R. M (2017) An improved model of motion-related signal changes in fMRI. *Neuroimage* **144**:74-82
79. Lanczos C (1964) Evaluation of noisy data. *J Soc Ind Appl Math Ser B Numer Anal* **1**:76-85
80. Abraham A., et al. (2014) Machine learning for neuroimaging with scikit-learn. *Front. Neuroinformatics* **8**
81. Mumford J. A., Turner B. O., Ashby F. G., Poldrack R. A (2012) Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *Neuroimage* **59**:2636-2643
82. Rissman J., Gazzaley A., D'Esposito M (2004) Measuring functional connectivity during distinct stages of a cognitive task. *Neuroimage* **23**:752-763
83. Traag V. A., Waltman L., Van Eck N. J (2019) From Louvain to Leiden: guaranteeing well-connected communities. *Sci Rep* **9**:5233
84. Delgado M. R., Nystrom L. E., Fissell C., Noll D., Fiez J. A (2000) Tracking the hemodynamic responses to reward and punishment in the striatum. *J Neurophysiol* **84**:3072-3077
85. Wu-Minn H (2017) 1200 subjects data release reference manual. URL <https://www.humanconnectome.org> **565**:2
86. Glasser M. F., et al. (2013) The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* **80**:105-124
87. Murphy K., Birn R. M., Handwerker D. A., Jones T. B., Bandettini P. A (2009) The impact of global signal regression on resting state correlations: are anti-correlated networks introduced?. *Neuroimage* **44**:893-905
88. Murphy K., Fox M. D (2017) Towards a consensus regarding global signal regression for resting state functional connectivity MRI. *Neuroimage* **154**:169-173
89. Birn R. M., et al. (2013) The effect of scan length on the reliability of resting-state fMRI connectivity estimates. *Neuroimage* **83**:550-558
90. Cohen J. R (2018) The behavioral and cognitive relevance of time-varying, dynamic changes in functional connectivity. *NeuroImage* **180**:515-525

91. Betzel R. F., Faskowitz J., Sporns O (2023) Living on the edge: network neuroscience beyond nodes. *Trends Cogn Sci*
92. Lurie D. J., et al. (2020) Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Netw Neurosci* **4**:30-69
93. Kuznetsova A., Brockhoff P. B., Christensen R. H. B (2017) lmerTest package: tests in linear mixed effects models. *J Stat Softw* **82**
94. Vanderwal T., Kelly C., Eilbott J., Mayes L. C., Castellanos F. X (2015) Inscapes: A movie paradigm to improve compliance in functional magnetic resonance imaging. *Neuroimage* **122**:222-232

Peer reviews

Reviewer #1 (Public review):

Summary:

This study investigates whether prediction error extends beyond lower-order sensory-motor processes to include higher-order cognitive functions. Evidence is drawn from both task-based and resting-state fMRI, with addition of resting-state EEG-fMRI to examine power spectral correlates. The results partially support the existence of dissociable connectivity patterns: stronger ventral-dorsal connectivity is associated with high prediction error, while posterior-anterior connectivity is linked to low prediction error. Furthermore, spontaneous switching between these connectivity patterns was observed at rest and correlated with subtle intersubject behavioral variability.

Strengths:

Studying prediction error from the lens of network connectivity provides new insights into predictive coding frameworks. The combination of various independent datasets to tackle the question adds strength, including two well-powered fMRI task datasets, resting-state fMRI interpreted in relation to behavioral measures, as well as EEG-fMRI.

Minor Weakness:

The lack of spatial specificity of sensor-level EEG somewhat limits the inferences that can be obtained in terms of how the fMRI network processes and the EEG power fluctuations relate to each other.

While the language no longer suggests a strong overlap of the source of the two signals, several scenarios remain open (e.g., the higher-order fMRI networks being the source of the EEG oscillations, or the networks controlling the EEG oscillations expressed in lower-order cortices, or a third process driving both the observations in fMRI networks and EEG oscillations...) and somewhat weaken interpretability of this section.

Comments on revisions:

My prior recommendations have been mostly addressed.

Questions remaining about the NBS results:

The authors write about the NBS cluster: "Visual examination of the cluster roughly points to the same four posterior-anterior and ventral-dorsal modules identified formally in main-text". I think it might be good to add quantification, not just visual inspection. The size of the significant NBS cluster should be reported. What proportion of the edges that passed uncorrected threshold and entered NBS were part of the NBS cluster? Put simply, I don't think any edges beyond those passing NBS-based correction should be interpreted or used downstream in the manuscript.

Also, NBS is not typically used by collapsing over effects in two effect directions, but the authors use NBS on the absolute value of Z. I understand the logic of the general manuscript focusing on strength rather than direction, but here I am wondering about the methodological validity. I believe that the editor who is an expert on the methodology may be able to comment on the validity of this approach (as opposed to running two separate NBS analyses for the two directions of effect).

<https://doi.org/10.7554/eLife.107723.2.sa3>

Reviewer #2 (Public review):

Summary:

This paper investigates putative networks associated with prediction errors in task-based and resting state fMRI. It attempts to test the idea that prediction errors minimisation includes abstract cognitive functions, referred to as global prediction error hypothesis, by establishing a parallel between networks found in task-based fMRI where prediction errors are elicited in a controlled manner and those networks that emerge during "resting state".

Strengths:

Clearly a lot of work and data went into this paper, including 2 task-based fMRI experiments and the resting state data for the same participants, as well as a third EEG-fMRI dataset. Overall well written with a couple of exceptions on clarity as per below and the methodology appears overall sound, with a couple of exceptions listed below that require further justification. It does a good job of acknowledging its own weakness.

Weaknesses:

The paper does a good job of acknowledging its greatest weakness, the fact that it relies heavily on reverse inference, but cannot quite resolve it. As the authors put, "finding the same networks during a prediction error task and during rest does not mean that the networks engagement during rest reflect prediction error processing". Again, the authors acknowledge the speculative nature of their claims in the discussion, but given that this is the key claim and essence of the paper, it is hard to see how the evidence is compelling to support that claim.

Given how uncontrolled cognition is during "resting-state" experiments, the parallel made with prediction errors elicited during a task designed to that effect is a little difficult to make. How often are people really surprised when their brains are "at rest", likely replaying a previously experienced event or planning future actions under their control? It seems to be more likely a very low prediction error scenario, if at all surprising.

The quantitative comparison between networks under task and rest was done on a small subset of the ROIs rather than on the full network - why? Noting how small the correlation between task and rest is ($r=0.021$) and that's only for part of the networks, the evidence is a little tenuous. Running the analysis for the full networks could strengthen the argument.

Looking at the results in Figure 2C, the four-quadrant description of the networks labelled for low and high PE appears a little simplistic. The authors state that this four-quadrant description omits some ROIs as motivated by prior knowledge. This would benefit from a more comprehensive justification. Which ROIs are excluded and what is the evidence for exclusion?

The EEG-fMRI analysis claiming 3-6Hz fluctuations for PE is hard to reconcile with the fact that fMRI captures activity that is a lot slower while some PEs are as fast as 150 ms. The

discussion acknowledges this but doesn't seem to resolve it - would benefit from a more comprehensive argument.

Comments on revisions:

The authors have done a good job of addressing the issues raised during the review process. There is one issue remaining that still required attention. In R2.4, when referring to "existing knowledge of prominent structural pathways among these quadrants" please cite the relevant literature.

<https://doi.org/10.7554/eLife.107723.2.sa2>

Reviewer #3 (Public review):

Summary:

Bogdan et al. present an intriguing investigation into the spontaneous dynamics of prediction error (PE)-related brain states. Using two independent fMRI tasks designed to elicit prediction and prediction error in separate participant samples, alongside both fMRI and EEG data, the authors identify convergent brain network patterns associated with high versus low PE. Notably, they further show that similar patterns can be detected during resting-state fMRI, suggesting that PE-related neural states may recur outside of explicit task demands.

Strengths:

The authors use a well-integrated analytic framework that combines multiple prediction tasks and brain imaging modalities. The inclusion of several datasets probing PE under different contexts strengthens the claim of generalizability across tasks and samples. The open sharing of code and data is commendable and will be valuable for future work seeking to build on this framework.

Weaknesses:

A central challenge of the manuscript lies in interpreting the functional significance of PE-related brain network states during rest. Demonstrating that a task-defined cognitive state recurs spontaneously is intriguing, but without clear links to behavior, individual traits, or experiential content during rest, it remains difficult to interpret what such spontaneous brain states tell us about the mind and brain. For example, it is unclear whether these states support future inference or learning, reflect offline predictive processing, or instead suggest state reinstatement due to a more general form of neural plasticity and circuit dynamics in the brain. Demonstrating any one of these downstream relationships would be valuable since it has the potential to inform our understanding of cognitive function or more general principles of neural organization.

I appreciate the authors' position that establishing the existence of such states is a necessary first step, and that future work may clarify their behavioral relevance. However, the current form makes it challenging to assess the conceptual advance of the present work in isolation.

Relatedly, in my previous review I raised questions about both across- and within-individual variability-for example, whether individuals who exhibit stronger or more distinct PE-related fluctuations at rest also show superior performance on prediction-related tasks (across-individual), or whether momentary increases in PE-network expression during tasks relate to faster or more accurate prediction (within-individual). The authors thoughtfully addressed this suggestion by conducting an individual-differences analysis correlating each participant's fluctuation amplitude with approximately 200 behavioral and trait measures from the HCP dataset.

The reported findings—a negative association with age and card-sorting performance, alongside a positive association with age-adjusted picture sequence memory—are interesting but difficult to interpret within a coherent functional framework. As presented, these results do not clearly support the idea that spontaneous PE-state fluctuations are related to enhancement in prediction, inference, or broader cognitive function. Instead, they raise the possibility that fluctuation amplitude may reflect more general factors (e.g., age) rather than a functionally meaningful PE-related process.

Overall, while the methodological contribution is strong, the manuscript would benefit from a clearer articulation of what functional conclusions can or cannot be drawn from the presence of spontaneous PE-related states, as well as a more cautious framing of their potential cognitive significance.

Further comments:

I appreciate that the authors took my earlier suggestions seriously and incorporated additional analyses examining behavioral relevance and permutation tests in the revision.

<https://doi.org/10.7554/eLife.107723.2.sa1>

Author response:

The following is the authors' response to the original reviews.

Public Reviews:

Reviewer #1 (Public review):

The Reviewer structured their review such that their first two recommendations specifically concerned the two major weaknesses they viewed in the initial submission. For clarity and concision, we have copied their recommendations to be placed immediately following their corresponding points on weaknesses.

Strengths:

Studying prediction error from the lens of network connectivity provides new insights into predictive coding frameworks. The combination of various independent datasets to tackle the question adds strength, including two well-powered fMRI task datasets, resting-state fMRI interpreted in relation to behavioral measures, as well as EEG-fMRI.

Weaknesses:

Major:

(R1.1) Lack of multiple comparisons correction for edge-wise contrast:

The analysis of connectivity differences across three levels of prediction error was conducted separately for approximately 22,000 edges (derived from 210 regions), yet no correction for multiple comparisons appears to have been applied. Then, modularity was applied to the top 5% of these edges. I do not believe that this approach is viable without correction. It does not help that a completely separate approach using SVMs was FDR-corrected for 210 regions.

[Later recommendation] Regarding the first major point: To address the issue of multiple comparisons in the edge-wise connectivity analysis, I recommend using the Network-Based Statistic (NBS; Zalesky et al., 2010). NBS is well-suited for identifying clusters

(analogous to modules) of edges that show statistically significant differences across the three prediction error levels, while appropriately correcting for multiple comparisons.

Thank you for bringing this up. We acknowledge that our modularity analysis does not evaluate statistical significance. Originally, the modularity analysis was meant to provide a connectome-wide summary of the connectivity effects, whereas the classification-based analysis was meant to address the need for statistical significance testing. However, as the reviewer points out, it would be better if significance were tested in a manner more analogous to the reported modules. As they suggest, we updated the Supplemental Materials (SM) to include the results of Network-Based Statistic analysis (SM p. 1-2):

“(2.1) Network-Based Statistic

Here, we evaluate whether PE significantly impacts connectivity at the network level using the Network-Based Statistic (NBS) approach.[1] NBS relied on the same regression data generated for the main-text analysis, whereby a regression is performed examining the effect of PE (Low = -1, Medium = 0, High = +1) on connectivity for each edge. This was done across the connectome, and for each edge, a z-score was computed. For NBS, we thresholded edges to $|Z| > 3.0$, which yielded one large network cluster, shown in Figure S3. The size of the cluster – i.e., number of edges – was significant ($p < .05$) per a permutation-test using 1,000 random shuffles of the condition data for each participant, as is standard.[1] These results demonstrate that the networklevel effects of PE on connectivity are significant. The main-text modularity analysis converts this large cluster into four modules, which are more interpretable and open the door to further analyses”.

We updated the Results to mention these findings before describing the modularity analysis (p. 8-9):

“After demonstrating that PE significantly influences brain-wide connectivity using Network-Based Statistic analysis (Supplemental Materials 2.1), we conducted a modularity analysis to study how specific groups of edges are all sensitive to high/low-PE information.”

(R1.2) Lack of spatial information in EEG:

The EEG data were not source-localized, and no connectivity analysis was performed. Instead, power fluctuations were averaged across a predefined set of electrodes based on a single prior study (reference 27), as well as across a broader set of electrodes. While the study correlates these EEG power fluctuations with fMRI network connectivity over time, such temporal correlations do not establish that the EEG oscillations originate from the corresponding network regions. For instance, the observed fronto-central theta power increases could plausibly originate from the dorsal anterior cingulate cortex (dACC), as consistently reported in the literature, rather than from a distributed network. The spatially agnostic nature of the EEG-fMRI correlation approach used here does not support interpretations tied to specific dorsal-ventral or anterior-posterior networks. Nonetheless, such interpretations are made throughout the manuscript, which overextends the conclusions that can be drawn from the data.

[Later recommendation] Regarding the second major point: I suggest either adopting a source-localized EEG approach to assess electrophysiological connectivity or revising all related sections to avoid implying spatial specificity or direct correspondence with fMRI-derived networks. The current approach, which relies on electrode-level power fluctuations, does not support claims about the spatial origin of EEG signals or their alignment with specific connectivity networks.

We thank the reviewer for this important point, which allows us to clarify the specific and distinct contributions of each imaging modality in our study. Our primary goal for Study 3 was to leverage the high temporal resolution of EEG to identify the characteristic frequency

at which the fMRI-defined global connectivity states fluctuate. The study was not designed to infer the spatial origin of these EEG signals, a task for which fMRI is better suited and which we addressed in Studies 1 and 2.

As the reviewer points out, fronto-central theta is generally associated with the dACC. We agree with this point entirely. We suspect that there is some process linking dACC activation to the identified network fluctuations – some type of relationship that does not manifest in our dynamic functional connectivity analyses – although this is only a hypothesis and one that is beyond the present scope.

We updated the Discussion to mention these points and acknowledge the ambiguity regarding the correlation between network fluctuation amplitude (fMRI) and Delta/Theta power (EEG) (p. 24):

“We specifically interpret the fMRI-EEG correlation as reflecting fluctuation speed because we correlated EEG oscillatory power with the fluctuation amplitude computed from fMRI data. Simply correlating EEG power with the average connectivity or the signed difference between posterior-anterior and ventral-dorsal connectivity yields null results (Supplemental Materials 6), suggesting that this is a very particular association, and viewing it as capturing fluctuation amplitude provides a parsimonious explanation. Yet, this correlation may be interpreted in other ways. For example, resting-state Theta is also a signature of drowsiness, [2] which may correlate with PE processing, but perhaps should be understood as some other mechanism. Additionally, Theta is widely seen as a sign of dorsal anterior cingulate cortex activity,³ and it is unclear how to reconcile this with our claims about network fluctuations. Nonetheless, as we show with simulations (Supplemental Materials 5), a correlation between slow fMRI network fluctuations and fast EEG Delta/Theta oscillations is also consistent with a common global neural process oscillating rapidly and eliciting both measures.”

Regarding source-localization, several papers have described known limitations of this strategy for drawing precise anatomical inferences,[4–6] and this seems unnecessary given that our fMRI analyses already provide more robust anatomical precision. We intentionally used EEG in our study for what it measures most robustly: millisecond-level temporal dynamics.

(R1.2a) Examples of problematic language include:

Line 134: "detection of network oscillations at fast speeds" - the current EEG approach does not measure networks.

This is an important issue. We acknowledge that our EEG approach does not directly measure fMRI-defined networks. Our claim is inferential, designed to estimate the temporal dynamics of the large-scale fMRI patterns we identified. The correlation between our fMRI-derived fluctuation amplitude ($|PA - VD|$) and 3-6 Hz EEG power provides suggestive evidence that the transitions between these network states occur at this frequency, rather than being a direct measurement of network oscillations.

To support the validity of this inference, we performed two key analyses (now in Supplemental Materials). First, a simulation study provides a proof-of-concept, confirming our method can recover the frequency of a fast underlying oscillator from slow fMRI and fast EEG data. Second, a specificity analysis shows the EEG correlation is unique to our measure of fluctuation amplitude and not to simpler measures like overall connectivity strength. These analyses demonstrate that our interpretation is more plausible than alternative explanations.

Overall, we have revised the manuscript to be more conservative in the language employed, such as presenting alternative explanations to the interpretations put forth based on correlative/observational evidence (e.g., our modifications above described in our response

to comment R1.2). In addition, we have made changes throughout the report to state the issues related to reverse inference more explicitly and to better communicate that the evidence is suggestive – please see our numerous changes described in our response to comment R3.1. For the statement that the reviewer specifically mentioned here, we revised it to be more cautious (p. 7):

“Although such speed outpaces the temporal resolution of fMRI, correlating fluctuations in dynamic connectivity measured from fMRI data with EEG oscillations can provide an estimate of the fluctuations’ speed. This interpretation of a correlation again runs up against issues related to reverse inference but would nonetheless serve as initial suggestive evidence that spontaneous transitions between network states occur rapidly.”

(R1.2b) Line 148: "whether fluctuations between high- and low-PE networks occur sufficiently fast" - this implies spatial localization to networks that is not supported by the EEG analysis.

Building on our changes described in our immediately prior response, we adjusted our text here to say our analyses searched for evidence consistent with the idea that the network fluctuations occur quickly rather than searching for decisive evidence favoring this idea (p. 7-8):

“Finally, we examined rs-fMRI-EEG data to assess whether we find parallels consistent with the high/low-PE network fluctuations occurring at fast timescales suitable for the type of cognitive operations typically targeted by PE theories.”

(R1.2c) Line 480: "how underlying neural oscillators can produce BOLD and EEG measurements" - no evidence is provided that the same neural sources underlie both modalities.

As described above, these claims are based on the simulation study demonstrating that this is a possibility, and we have revised the manuscript overall to be clearer that this is our interpretation while providing alternative explanations.

Reviewer #2 (Public review):

Strengths:

Clearly, a lot of work and data went into this paper, including 2 task-based fMRI experiments and the resting state data for the same participants, as well as a third EEG-fMRI dataset. Overall, well written with a couple of exceptions on clarity, as per below, and the methodology appears overall sound, with a couple of exceptions listed below that require further justification. It does a good job of acknowledging its own weakness.

Weaknesses:

(R2.1) The paper does a good job of acknowledging its greatest weakness, the fact that it relies heavily on reverse inference, but cannot quite resolve it. As the authors put it, "finding the same networks during a prediction error task and during rest does not mean that the networks' engagement during rest reflects prediction error processing". Again, the authors acknowledge the speculative nature of their claims in the discussion, but given that this is the key claim and essence of the paper, it is hard to see how the evidence is compelling to support that claim.

We thank the reviewer for this comment. We agree that reverse inference is a fundamental challenge and that our central claim requires a particularly high bar of evidence. While no single analysis resolves this issue, our goal was to build a cumulative case that is compelling by converging on the same conclusion from multiple, independent lines of evidence.

For our investigation, we initially established a task-general signature of prediction error (PE). By showing the same neural pattern represents PE in different contexts, we constrain the reverse inference, making it less likely that our findings are a task-specific artifact and more likely that they reflect the core, underlying process of PE. Building on this, our most compelling evidence comes from linking task and rest at the individual level. We didn't just find the same general network at rest; we showed that an individual's unique anatomical pattern of PE-related connectivity during the task specifically predicts their own brain's fluctuation patterns at rest. This highly specific, person-by-person correspondence provides a direct bridge between an individual's task-evoked PE processing and their intrinsic, resting-state dynamics. Furthermore, these resting-state fluctuations correlate specifically with the 3-6 Hz theta rhythm—a well-established neural marker for PE.

While reverse inference remains a fundamental limitation for many studies on resting-state cognition, the aspects mentioned above, we believe, provide suggestive evidence, favoring our PE interpretation. Nonetheless, we have made changes throughout the manuscript to be more conservative in the language we use to describe our results, to make it clear what claims are based on correlative/observational evidence, and to put forth alternative explanations for the identified effects. Please find our numerous changes detailed in our response to comment R3.1.

(R2.2) Given how uncontrolled cognition is during "resting-state" experiments, the parallel made with prediction errors elicited during a task designed for that effect is a little difficult to make. How often are people really surprised when their brains are "at rest", likely replaying a previously experienced event or planning future actions under their control? It seems to be more likely a very low prediction error scenario, if at all surprising.

We (and some others) take a broad interpretation of PE and believe it is often more intuitive to think about PE minimization in terms of uncertainty rather than “surprise”; the word “surprise” usually implies a sudden emotive reaction from the violation of expectations, which is not useful here.

When planning future actions, each step of the plan is spurred by the uncertainty of what is the appropriate action given the scenario set up by prior steps. Each planned step erases some of that uncertainty. For example, you may be mentally simulating a conversation, what you will say, and what another person will say. Each step of this creates uncertainty of “what is the appropriate response?” Each reasoning step addresses contingencies. While planning, you may also uncover more obvious forms of uncertainty, sparking memory retrieval to finish it. A resting-state participant may think to cook a frozen pizza when they arrive home, but be uncertain about whether they have any frozen pizzas left, prompting episodic memory retrieval to address this uncertainty. We argue that every planning step or memory retrieval can be productively understood as being sparked by uncertainty/surprise (PE), and the subsequent cognitive response minimizes this uncertainty.

We updated the Introduction to include a paragraph near the start providing this explanation (p. 3-4):

“PE minimization may broadly coordinate brain functions of all sorts, including abstract cognitive functions. This includes the types of cognitive processes at play even in the absence of stimuli (e.g., while daydreaming). While it may seem counterintuitive to associate this type of cognition with PE – a concept often tied to external surprises – it has been proposed that the brain's internal generative model is continuously active.[12–14] Spontaneous thought, such as planning a future event or replaying a memory, is not a passive, low-PE process. Rather, it can be seen as a dynamic cycle of generating and resolving internal uncertainty. While daydreaming, you may be reminded of a past conversation, where you wish you had

said something different. This situation contains uncertainty about what would have been the best thing to say. Wondering about what you wish you said can be viewed as resolving this uncertainty, in principle, forming a plan if the same situation ever arises again in the future. Each iteration of the simulated conversation repeatedly sparks and then resolves this type of uncertainty.”

(R2.3) The quantitative comparison between networks under task and rest was done on a small subset of the ROIs rather than on the full network - why? Noting how small the correlation between task and rest is ($r=0.021$) and that's only for part of the networks, the evidence is a little tenuous. Running the analysis for the full networks could strengthen the argument.

We thank the reviewer for this opportunity to clarify our method. A single correlation between the full, aggregated networks would be conceptually misaligned with what we aimed to assess. To test for a person-specific anatomical correspondence, it is necessary to examine the link between task and rest at a granular level. We therefore asked whether the specific parts of an individual's network most responsive to PE during the task are the same parts that show the strongest fluctuations at rest. Our analysis, performed iteratively across all 3,432 possible ROI subsets, was designed specifically to answer this question, which would be obscured by an aggregated network measure.

We appreciate the reviewer's concern about the modest effect size ($r = .021$). However, this must be contextualized, as the short task scan has very low reliability (.08), which imposes a severe statistical ceiling on any possible task-rest correlation. Finding a highly significant effect ($p < .001$) in the face of such noisy data, therefore, provides robust evidence for a genuine task-rest correspondence.

We updated the Discussion to discuss this point (p. 22-23):

“A key finding supporting our interpretation is the significant link between individual differences in task-evoked PE responses and resting-state fluctuations. One might initially view the effect size of this correspondence ($r = .021$) as modest. However, this interpretation must be contextualized by the considerable measurement noise inherent in short task-fMRI scans; the split-half reliability of the task contrast was only .08. This low reliability imposes a severe statistical ceiling on any possible task-rest correlation. Therefore, detecting a highly significant ($p < .001$) relationship despite this constraint provides robust evidence for a genuine link. Furthermore, our analytical approach, which iteratively examined thousands of ROI subsets rather than one aggregated network, was intentionally granular. The goal was not simply to correlate two global measures, but to test for a person-specific anatomical correspondence – that is, whether the specific parts of an individual's network most sensitive to PE during the task are the same parts that fluctuate most strongly at rest. An aggregate analysis would obscure this critical spatial specificity. Taken together, this granular analysis provides compelling evidence for an anatomically consistent fingerprint of PE processing that bridges task-evoked activity and spontaneous resting-state dynamics, strengthening our central claim.”

(R2.4) Looking at the results in Figure 2C, the four-quadrant description of the networks labelled for low and high PE appears a little simplistic. The authors state that this four-quadrant description omits some ROIs as motivated by prior knowledge. This would benefit from a more comprehensive justification. Which ROIs are excluded, and what is the evidence for exclusion?

Our four-quadrant model is a principled simplification designed to distill the dominant, large-scale connectivity patterns from the complex modularity results. This approach focuses on coherent, well-documented anatomical streams while setting aside a few anatomically distant

and disjoint ROIs that were less central to the main modules. This heuristic additionally unlocks more robust and novel analyses.

The two low-PE posterior-anterior (PA) pathways are grounded in canonical processing streams. (i) The OCATL connection mirrors the ventral visual stream (the “what” pathway), which is fundamental for object recognition and is upregulated during the smooth processing of expected stimuli. (ii) The IPL-LPFC connection represents a core axis of the dorsal attention stream and the Fronto-Parietal Control Network (FPCN), reflecting the maintenance of top-down cognitive control when information is predictable; the IPL-LPFC module excludes ROIs in the middle temporal gyrus, which are often associated with the FPCN but are not covered here.

In contrast, the two high-PE ventral-dorsal (VD) pathways reflect processes for resolving surprise and conflict. (i) The OC-IPL connection is a classic signature of attentional reorienting, where unexpected sensory input (high PE) triggers a necessary shift in attention; the OC-IPL module excludes some ROIs that are anterior to the occipital lobe and enter the fusiform gyrus and inferior temporal lobe. (ii) The ATL-LPFC connection aligns with mechanisms for semantic re-evaluation, engaging prefrontal control regions to update a mental model in the face of incongruent information.

Beyond its functional/anatomical grounding, this simplification provides powerful methodological and statistical advantages. It establishes a symmetrical framework that makes our dynamic connectivity analyses tractable, such as our “cube” analysis of state transitions, which required overlapping modules. Critically, this model also offers a statistical safeguard. By ensuring each quadrant contributes to both low- and high-PE connectivity patterns, we eliminate confounds like region-specific signal variance or global connectivity. This design choice isolates the phenomenon to the pattern of connectivity itself (posterior-anterior vs. ventral-dorsal), making our interpretation more robust.

We updated the end of the Study 1A results (p. 10-11):

“Some ROIs appear in Figure 2C but are excluded from the four targeted quadrants (Figures 2C & 2D) – e.g., posterior inferior temporal lobe and fusiform ROIs are excluded from the OC-IPL module, and middle temporal gyrus ROIs are excluded from the IPL-LPFC modules. These exclusions, in favor of a four-quadrant interpretation, are motivated by existing knowledge of prominent structural pathways among these quadrants. This interpretation is also supported by classifier-based analyses showing connectivity within each quadrant is significantly influenced by PE (Supplemental Materials 2.2), along with analyses of single-region activity showing that these areas also respond to PE independently (Supplemental Materials 3). Hence, we proceeded with further analyses of these quadrants’ connections, which summarize PE’s global brain effects.

“This four-quadrant setup also imparts analytical benefits. First, this simplified structure may better generalize across PE tasks, and Study 1B would aim to replicate these results with a different design. Second, the four quadrants mean that each ROI contributes to both the posterior-anterior and ventral-dorsal modules, which would benefit later analyses and rules out confounds such as PE eliciting increased/decreased connectivity between an ROI and the rest of the brain. An additional, less key benefit is that this setup allows more easily evaluating whether the same phenomena arise using a different atlas (Supplemental Materials Y).”

(R2.5) The EEG-fMRI analysis claiming 3-6Hz fluctuations for PE is hard to reconcile with the fact that fMRI captures activity that is a lot slower, while some PEs are as fast as 150 ms. The discussion acknowledges this but doesn't seem to resolve it - would benefit from a more comprehensive argument.

We thank the reviewer for raising this important point, which allows us to clarify the logic of our multimodal analysis. Our analysis does not claim that the fMRI BOLD signal itself oscillates at 3-6 Hz. Instead, it is based on the principle that the intensity of a fast neural process can be reflected in the magnitude of the slow BOLD response. It's akin to using a long-exposure photograph to capture a fast-moving object; while the individual movements are blurred, the intensity of the blur in the photo serves as a proxy for the intensity of the underlying motion. In our case, the magnitude of the fMRI network difference ($|PA - VD|$) acts as the "blur," reflecting the intensity of the rapid fluctuations between states within that time window.

Following this logic, we correlated this slow-moving fMRI metric with the power of the fast EEG rhythms, which reflects their amplitude. To bridge the different timescales, we averaged the EEG power over each fMRI time window and convolved it with the standard hemodynamic response function (HRF) – a crucial step to align the timing of the neural and metabolic signals. The resulting significant correlation specifically in the 3-6 Hz band demonstrates that when this rhythm is stronger, the fMRI data shows a greater divergence between network states. This allows us to infer the characteristic frequency of the underlying neural fluctuations without directly measuring them at that speed with fMRI, thus reconciling the two timescales.

Reviewer #3 (Public review):

Bogdan et al. present an intriguing and timely investigation into the intrinsic dynamics of prediction error (PE)-related brain states. The manuscript is grounded in an intuitive and compelling theoretical idea: that the brain alternates between high and low PE states even at rest, potentially reflecting an intrinsic drive toward predictive minimization. The authors employ a creative analytic framework combining different prediction tasks and imaging modalities. They shared open code, which will be valuable for future work.

(R3.1) Consistency in Theoretical Framing

The title, abstract, and introduction suggest inconsistent theoretical goals of the study.

The title suggests that the goal is to test whether there are intrinsic fluctuations in high and low PE states at rest. The abstract and introduction suggest that the goal is to test whether the brain intrinsically minimizes PE and whether this minimization recruits global brain networks. My comments here are that a) these are fundamentally different claims, and b) both are challenging to falsify. For one, task-like recurrence of PE states during resting might reflect the wiring and geometry of the functional organization of the brain emerging from neurobiological constraints or developmental processes (e.g., experience), but showing that mirroring exists because of the need to minimize PE requires establishing a robust relationship with behavior or showing a causal effect (e.g., that interrupting intrinsic PE state fluctuations affects prediction).

The global PE hypothesis-"PE minimization is a principle that broadly coordinates brain functions of all sorts, including abstract cognitive functions"-is more suitable for discussion rather than the main claim in the abstract, introduction, and all throughout the paper.

Given the above, I recommend that the authors clarify and align their core theoretical goals across the title, abstract, introduction, and results. If the focus is on identifying fluctuations that resemble taskdefined PE states at rest, the language should reflect that more narrowly, and save broader claims about global PE minimization for the discussion. This hypothesis also needs to be contextualized within prior work. I'd like to see if there is similar evidence in the literature using animal models.

Thank you for bringing up this issue. We have made changes throughout the paper to address these points. First, we have omitted reference to a “global PE hypothesis” from the Abstract and Introduction, in favor of structuring the Introduction in terms of a falsifiable question (p. 4):

“We pursued this goal using three studies (Figure 1) that collectively targeted a specific question: Do the taskdefined connectivity signatures of high vs. low PE also recur during rest, and if so, how does the brain transition between exhibiting high/low signatures?”

We made changes later in the Introduction to clarify that the investigation is based on correlative evidence and requires interpretations that may be debated (p. 5-7):

“Although this does not entirely address the reverse inference dilemma and can only produce correlative evidence, the present research nonetheless investigates these widely speculated upon PE ideas more directly than any prior work.

Although such speed outpaces the temporal resolution of fMRI, correlating fluctuations in dynamic connectivity measured from fMRI data with EEG oscillations can provide an estimate of the fluctuations’ speed. This interpretation of a correlation again runs up against issues related to reverse inference but would nonetheless serve as initial suggestive evidence that spontaneous transitions between network states occur rapidly.

Second, we examined the recruitment of these networks during rs-fMRI, and although the problems related to reverse inference are impossible to overcome fully, we engage with this issue by linking rs-fMRI data directly to task-fMRI data of the same participants, which can provide suggestive evidence that the same neural mechanisms are at play in both.”

We made changes throughout the Results now better describing the results as consistent with a hypothesis rather than demonstrating it (p. 12-19):

“In other words, we essentially asked whether resting-state participants are sometimes in low PE states and sometimes in high PE states, which would be consistent with spontaneous PE processing in the absence of stimuli.

These emerging states overlap strikingly with the previous task effects of PE, suggesting that rs-fMRI scans exhibit fluctuations that resemble the signatures of low- and high-PE states.

To be clear, this does not entirely dissuade concerns about reverse inference, which would require a type of causal manipulation that is difficult (if not impossible) to perform in a resting state scan. Nonetheless, these results provide further evidence consistent with our interpretation that the resting brain spontaneously fluctuates between high/low PE network states.

These patterns are most consistent with a characteristic timescale near 3–6 Hz for the amplitude of the putative high/low-PE fluctuations. This is notably consistent with established links between PE and Delta/Theta and is further consistent with an interpretation in which these fluctuations relate to PE-related processing during rest.”

We have also made targeted edits to the Discussion to present the findings in a more cautious way, more clearly state what is our interpretation, and provide alternative explanations (p. 19-26):

“The present research conducted task-fMRI, rs-fMRI, and rs-fMRI-EEG studies to clarify whether PE elicits global connectivity effects and whether the signatures of PE processing arise spontaneously during rest. This investigation carries implications for how PE minimization may characterize abstract task-general cognitive processes. [...] Although there

are different ways to interpret this correlation, it is consistent with high/low PE states generally fluctuating at 3-6 Hz during rest. Below, we discuss these three studies' findings.

Our rs-fMRI investigation examined whether resting dynamics resemble the task-defined connectivity signatures of high vs. low PE, independent of the type of stimulus encountered. The resting-state analyses indeed found that, even at rest, participants' brains fluctuated between strong ventral-dorsal connectivity and strong posterior-anterior connectivity, consistent with shifts between states of high and low PE. This conclusion is based on correlative/observational evidence and so may be controversial as it relies on reverse inference.

These patterns resemble global connectivity signatures seen in resting-state participants, and correlations between fMRI and EEG data yield associations, consistent with participants fluctuating between high-PE (ventral-dorsal) and low-PE (posterior-anterior) states at 3-6 Hz. Although definitively testing these ideas is challenging, given that rs-fMRI is defined by the absence of any causal manipulations, our results provide evidence consistent with PE minimization playing a role beyond stimulus process."

(R3.2) Interpretation of PE-Related Fluctuations at Rest and Its Functional Relevance. It would strengthen the paper to clarify what is meant by "intrinsic" state fluctuations. Intrinsic might mean taskindependent, trait-like, or spontaneously generated. Which do the authors mean here? Is the key prediction that these fluctuations will persist in the absence of a prediction task?

Of the three terms the reviewer mentioned, "spontaneous" and "task-independent" are the most accurate descriptors. We conceptualize these fluctuations as a continuous background process that persists across all facets of cognition, without requiring a task explicitly designed to elicit prediction error – although we, along with other predictive coding papers, would argue that all cognitive tasks are fundamentally rooted in PE mechanisms and thus anything can be seen as a "prediction task" (see our response to comment R2.2 for our changes to the Introduction that provide more intuition for this point). The proposed interactions can be seen as analogous to cortico-basal-thalamic loops, which are engaged across a vast and diverse array of cognitive processes.

The prior submission only used the word "intrinsic" in the title. We have since revised it to "spontaneous," which is more specific than "intrinsic," and we believe clearer for a title than "task-independent" (p. 1): "Spontaneous fluctuations in global connectivity reflect transitions between states of high and low prediction error"

We have also made tweaks throughout the manuscript to now use "spontaneously" throughout (it now appears 8 times in the paper).

Regardless of the intrinsic argument, I find it challenging to interpret the results as evidence of PE fluctuations at rest. What the authors show directly is that the degree to which a subset of regions within a PE network discriminates high vs. low PE during task correlates with the magnitude of separation between high and low PE states during rest. While this is an interesting relationship, it does not establish that the resting-state brain spontaneously alternates between high and low PE states, nor that it does so in a functionally meaningful way that is related to behavior. How can we rule out brain dynamics of other processes, such as arousal, that also rise and fall with PE? I understand the authors' intention to address the reverse inference concern by testing whether "a participant's unique connectivity response to PE in the reward-processing task should match their specific patterns of resting-state fluctuation". However, I'm not fully convinced that this analysis establishes the functional role of the identified modules to PE because of the following:

Theoretically, relating the activities of the identified modules directly to behavior would demonstrate a stronger functional role.

(R3.2a) Across participants: Do individuals who exhibit stronger or more distinct PE-related fluctuations at rest also perform better on tasks that require prediction or inference? This could be assessed using the HCP prediction task, though if individual variability is limited (e.g., due to ceiling effects), I would suggest exploring a dataset with a prediction task that has greater behavioral variance.

This is a good idea, but unfortunately difficult to test with our present data. The HCP gambling task used in our study was not designed to measure individual differences in prediction or inference and likely suffers from ceiling effects. Because the task outcomes are predetermined and not linked to participants' choices, there is very little meaningful behavioral variance in performance to correlate with our resting-state fluctuation measure.

While we agree that exploring a different dataset with a more suitable task would be ideal, given the scope of the existing manuscript, this seems like it would be too much. Although these results would be informative, they would ultimately still not be a panacea for the reverse inference issues.

Or even more broadly, does this variability in resting state PE state fluctuations predict general cognitive abilities like WM and attention (which the HCP dataset also provides)? I appreciate the inclusion of the win-loss control, and I can see the intention to address specificity. This would test whether PE state fluctuations reflect something about general cognition, but also above and beyond these attentional or WM processes that we know are fluctuating.

This is a helpful suggestion, motivating new analyses: We measured the degree of resting-state fluctuation amplitude across participants and correlated it with the different individual differences measures provided with the HCP data (e.g., measures of WM performance). We computed each participant's fluctuation amplitude measure as the average absolute difference between posterior-anterior and ventral-dorsal connectivity; this is the average of the TR-by-TR fMRI amplitude measure from Study 3. We correlated this individual difference score with all of the ~200 individual difference measures provided with the HCP dataset (e.g., measures of intelligence or personality). We measured the Spearman correlation between mean fluctuation amplitude with each of those ~200 measures, while correcting for multiple hypotheses using the False Discovery Rate approach.[18]

We found a robust negative association with age, where older participants tend to display weaker fluctuations ($r = -.16$, $p < .001$). We additionally find a positive association with the age-adjusted score on the picture sequence task ($r = .12$, $p_{\text{corrected}} = .03$) and a negative association with performance in the card sort task ($r = -.12$, $p_{\text{corrected}} = .046$). It is unclear how to interpret these associations, without being speculative, given that fluctuation amplitude shows one positive association with performance and one negative association, albeit across entirely different tasks. We have added these correlation results as Supplemental Materials 8 (SM p. 11):

“(8) Behavioral differences related to fluctuation amplitude

To investigate whether individual differences in the magnitude of resting-state PE-state fluctuations predict general cognitive abilities, we correlated our resting-state fluctuation measure with the cognitive and demographic variables provided in the HCP dataset.

(8.1) Methods

For each of the 1,000 participants, we calculated a single fluctuation amplitude score. This score was defined as the average absolute difference between the time-varying posterior-anterior (PA) and ventral-dorsal (VD) connectivity during the resting-state fMRI scan (the average of the TR-by-TR measure used for Study 3). We then computed the Spearman correlation between this score and each of the approximately 200 individual difference measures provided in the HCP dataset. We corrected for multiple comparisons using the False Discovery Rate (FDR) approach.

(8.2) Results

The correlations revealed a robust negative association between fluctuation amplitude and age, indicating that older participants tended to display weaker fluctuations ($r = -.16$, $p_{\text{corrected}} < .001$). After correction, two significant correlations with cognitive performance emerged: (i) a positive association with the age-adjusted score on the Picture Sequence Memory Test ($r = .12$, $p_{\text{corrected}} = .03$), (ii) a negative association with performance on the Card Sort Task ($r = -.12$, $p_{\text{corrected}} = .046$). As greater fluctuation amplitude is linked to better performance on one task but worse performance on another, it is unclear how to interpret these findings.”

We updated the main text Methods to direct readers to this content (p. 39-40):

“(4.4.3) Links between network fluctuations and behavior

We considered whether the extent of PE-related network expression states during resting-state is behaviorally relevant. We specifically investigated whether individual differences in the overall magnitude of resting-state fluctuations could predict individual difference measures, provided with the HCP dataset. This yielded a significant association with age, whereby older participants tended to display weaker fluctuations. However, associations with cognitive measures were limited. A full description of these analyses is provided in Supplemental Materials 8.”

(R3.2b) Within participants: Do momentary increases in PE-network expression during tasks relate to better or faster prediction? In other words, is there evidence that stronger expression of PE-related states is associated with better behavioral outcomes?

This is a good question that probes the direct behavioral relevance of these network states on a trial-by-trial basis. We agree with the reviewer's intuition; in principle, one would expect a stronger expression of the low-PE network state on trials where a participant correctly and quickly gives a high likelihood rating to a predictable stimulus.

Following this suggestion, we performed a new analysis in Study 1A to test this. We found that while network expression was indeed linked to participants' likelihood ratings: higher likelihood ratings correspond to stronger posterior-anterior connectivity, whereas lower ratings correspond to stronger ventral-dorsal connectivity (Connectivity-Direction $\times$ likelihood, β [standardized] = .28, $p = .02$). Yet, this is not a strong test of the reviewer's hypothesis, and different exploratory analyses of response time yield null results ($p > .05$). We suspect that this is due to the effect being too subtle, so we have insufficient statistical power. A comparable analysis was not feasible for Study 1B, as its design does not provide an analogous behavioral measure of trial-by-trial prediction success.

(R3.3) A priori Hypothesis for EEG Frequency Analysis.

It's unclear how to interpret the finding that fMRI fluctuations in the defined modules correlate with frontal Delta/Theta power, specifically in the 3-6 Hz range. However, in the EEG literature, this frequency band is most commonly associated with low arousal, drowsiness, and mind wandering in resting, awake adults, not uniquely with prediction

error processing. An a priori hypothesis is lacking here: what specific frequency band would we expect to track spontaneous PE signals at rest, and why? Without this, it is difficult to separate a PE-based interpretation from more general arousal or vigilance fluctuations.

This point gets to the heart of the challenge with reverse inference in resting-state fMRI. We agree that an interpretation based on general arousal or drowsiness is a potential alternative that must be considered. However, what makes a simple arousal interpretation challenging is the highly specific nature of our fMRI-EEG association. As shown in our confirmatory analyses (Supplemental Materials 6), the correlation with 3-6 Hz power was found exclusively with the absolute difference between our two PE-related network states ($|PA - VD|$)—a measure of fluctuation amplitude. We found no significant relationship with the signed difference (a bias toward one state) or the sum (the overall level of connectivity). This specificity presents a puzzle for a simple drowsiness account; it seems less plausible that drowsiness would manifest specifically as the intensity of fluctuation between two complex cognitive networks, rather than as a more straightforward change in overall connectivity. While we cannot definitively rule out contributions from arousal, the specificity of our finding provides stronger evidence for a structured cognitive process, like PE, than for a general, undifferentiated state.

We updated the Discussion to make the argument above and also to remind readers that alternative explanations, such as ones based on drowsiness, are possible (p. 24):

“We specifically interpret the fMRI-EEG correlation as reflecting fluctuation speed because we correlated EEG oscillatory power with the fluctuation amplitude computed from fMRI data. Simply correlating EEG power with the average connectivity or the signed difference between posterior-anterior and ventral-dorsal connectivity yields null results (Supplemental Materials 6), suggesting that this is a very particular association, and viewing it as capturing fluctuation amplitude provides a parsimonious explanation. Yet, this correlation may be interpreted in other ways. For example, resting-state Theta is also a signature of drowsiness, [2] which may correlate with PE processing, but perhaps should be understood as some other mechanism.”

(R3.4) Significance Assessment

The significance of the correlation above and all other correlation analyses should be assessed through a permutation test rather than a single parametric t-test against zero. There are a few reasons: a) EEG and fMRI time series are autocorrelated, violating the independence assumption of parametric tests;

Standard t-tests can underestimate the true null distribution's variance, because EEG-fMRI correlations often involve shared slow drifts or noise sources, which can yield spurious correlations and inflating false positives unless tested against an appropriate null.

Building a null distribution that preserves the slow drifts, for example, would help us understand how likely it is for the two time series to be correlated when the slow drifts are still present, and how much better the current correlation is, compared to this more conservative null. You can perform this by phase randomizing one of the two time courses N times (e.g., $N=1000$), which maintains the autocorrelation structure while breaking any true co-occurrence in patterns between the two time series, and compute a non-parametric p-value. I suggest using this approach in all correlation analyses between two time series.

This is an important statistical point to clarify, and the suggested analysis is valuable. The reviewer is correct that the raw fMRI and EEG time series are autocorrelated. However, because our statistical approach is a twolevel analysis, we reasoned that non-independence

at the correlation-level would not invalidate the higher-level t-test. The t-test's assumption of independence applies to the individual participants' coefficients, which are independent across participants. Thus, we believe that our initial approach is broadly appropriate, and its simplicity allows it to be easily communicated.

Nonetheless, the permutation-testing procedure that the Reviewer describes seems like an important analysis to test, given that permutation-testing is the gold standard for evaluating statistical significance, and it could guarantee that our above logic is correct. We thus computed the analysis as the reviewer described. For each participant, we phase-randomized the fMRI fluctuation amplitude time series. Specifically, we randomized the Fourier phases of the |PA-VD| series (within run), while retaining the original amplitude spectrum; inverse transforms yielded real surrogates with the same power spectrum. This was done for each participant once per permutation. Each participant's phase-randomized data was submitted to the analysis of each oscillatory power band as originally, generating one mean correlation for each band. This was done 1,000 times.

Across the five bands, we find that the grand mean correlation is near zero ($M_r = .0006$) and the 97.5th percentile critical value of the null distribution is $r = \sim .025$; this 97.5th percentile corresponds to the upper end of a 95% confidence interval for a band's correlation; the threshold minimally differs across bands ($.024 < r_s < .026$). Our original correlation coefficients for Delta ($M_r = .042$) and Theta ($M_r = .041$), which our conclusions focused on, remained significant ($p \leq .002$); we can perform family-wise error-rate correction by taking the highest correlation across any band for a given permutation, and the Delta and Theta effects remain significant ($p_{\text{FWEcorrected}} \leq .003$); previously Reviewer comment R1.4c requested that we employ family-wise error correction.

These correlations were previously reported in Table 1, and we updated the caption to note what effects remain significant when evaluated using permutation-testing and with family-wise error correction (p. 19):

"The effects for Delta, Theta, Beta, and Gamma remain significant if significance testing is instead performed using permutation-testing and with family-wise error rate correction ($p_{\text{corrected}} < .05$)."

We updated the Methods to describe the permutation-testing analysis (p. 43):

"To confirm the significance of our fMRI-EEG correlations with a non-parametric approach, we performed a group-level permutation-test. For each of 1,000 permutations, we phase-randomized the fMRI fluctuation amplitude time series. Specifically, we randomized the Fourier phases of the |PA-VD| series (within run), while retaining the original amplitude spectrum; inverse transforms yielded real surrogates with the same power spectrum. This procedure breaks the true temporal relationship between the fMRI and EEG data while preserving its structure. We then re-computed the mean Spearman correlation for each frequency band using this phase-randomized data. We evaluated significance using a family-wise error correction approach that accounts for us analyzing five oscillatory power bands. We thus create a null distribution composed of the maximum correlation value observed across all frequency bands from each permutation. Our observed correlations were then tested for significance against this distribution of maximums."

(R3.5) Analysis choices

If I'm understanding correctly, the algorithm used to identify modules does so by assigning nodes to communities, but it does not itself restrict what edges can be formed from these modules. This makes me wonder whether the decision to focus only on connections between adjacent modules, rather than considering the full connectivity, was an analytic choice by the authors. If so, could you clarify the rationale? In particular, what justifies assuming that the gradient of PE states should be captured by edges

formed only between nearby modules (as shown in Figure 2E and Figure 4), rather than by the full connectivity matrix? If this restriction is instead a by-product of the algorithm, please explain why this outcome is appropriate for detecting a global signature of PE states in both task and rest.

We discuss this matter in our response to comment R2.(4).

When assessing the correspondence across task-fMRI and rs-fMRI in section 2.2.2, why was the pattern during task calculated from selecting a pair of bilateral ROIs (resulting in a group of eight ROIs), and the resting state pattern calculated from posterior-anterior/ventral-dorsal fluctuation modules? Doesn't it make more sense to align the two measures? For example, calculating task effects on these same modules during task and rest?

We thank the reviewer for this question, as it highlights a point in our methods that we could have explained more clearly. The reviewer is correct that the two measures must be aligned, and we can confirm that they were indeed perfectly matched.

For the analysis in Section 2.2.2, both the task and resting-state measures were calculated on the exact same anatomical substrate for each comparison. The analysis iteratively selected a symmetrical subset of eight ROIs from our larger four quadrants. For each of these 3,432 iterations, we computed the task-fMRI PE effect (the Connectivity Direction $\times$ PE interaction) and the resting-state fluctuation amplitude ($E[|PA - VD|]$) using the identical set of eight ROIs. The goal of this analysis was precisely to test if the fine-grained anatomical pattern of these effects correlated within an individual across the task and rest states. We will revise the text in Section 2.2.2 to make this direct alignment of the two measures more explicit.

Recommendations for authors:

Reviewer #1 (Recommendations for authors):

(R1.3) Several prior studies have described co-activation or connectivity "templates" that spontaneously alternate during rest and task states, and are linked to behavioral variability. While they are interpreted differently in terms of cognitive function (e.g., in terms of sustained attention: Monica Rosenberg; alertness: Catie Chang), the relationship between these previously reported templates and those identified in the current study warrants discussion. Are the current templates spatially compatible with prior findings while offering new functional interpretations beyond those already proposed in the literature? Or do they represent spatially novel patterns?

Thank you for this suggestion. Broadly, we do not mean to propose spatially novel patterns but rather focus on how these are repurposed for PE processing. In the Discussion, we link our identified connectivity states to established networks (e.g., the FPCN). We updated this paragraph to mention that these patterns are largely not spatially novel (p. 20):

"The connectivity patterns put forth are, for the most part, not spatially novel and instead overlap heavily with prior functional and anatomical findings."

Regarding the specific networks covered in the prior work by Rosenberg and Chang that the reviewer seems to be referring to, [7,8] this research has emphasized networks anchored heavily in sensorimotor, subcortical– cerebellar, and medial frontal circuits, and so mostly do not overlap with the connectivity effects we put forth.

(R1.4) Additional points:

(R1.4a) I do not think that the logic for taking the absolute difference of fMRI connectivity is convincing. What happens if the sign of the difference is maintained ?

Thank you for pointing out this area that requires clarification. Our analysis targets the amplitude of the fluctuation between brain states, not the direction. We define high fluctuation amplitude as moments when the brain is strongly in either the PA state ($PA > VD$) or the VD state ($VD > PA$). The absolute difference $|PA - VD|$ correctly quantifies this intensity, whereas a signed difference would conflate these two distinct high-amplitude moments. Our simulation study (Supplemental Materials, Section 5) provides the theoretical validation for this logic, showing how this absolute difference measure in slow fMRI data can track the amplitude of a fast underlying neural oscillator.

When the analysis is tested in terms of the signed difference, as suggested by the Reviewer, the association between the fMRI data and EEG power is insignificant for each power band ($p_{\text{uncorrected}} \geq .47$). We updated Supplemental Materials 6 to include these results. Previously, this section included the fluctuation amplitude (fMRI) $\times$ EEG power results while controlling for: (i) the signed difference between posterior-anterior and ventral-dorsal connectivity, (ii) the sum of posterior-anterior and ventral-dorsal connectivity, and (iii) the absolute value of the sum of posterior-anterior and ventral-dorsal connectivity. For completeness, we also now report the correlation between each EEG power band and each of those other three measures (SM, p. 9)

“We additionally tested the relationship between each of those three measures and the five EEG oscillation bands. Across the 15 tests, there were no associations ($p_{\text{uncorrected}} \geq .04$); one uncorrected p-value was at $p = .044$, although this was expected given that there were 15 tests. Thus, the association between EEG oscillations and the fMRI measure is specific to the absolute difference (i.e., amplitude) measure.”

(R1.4b) Reasoning of focus on frontal and theta band is weak, and described as "typical" (line 359) based on a single study.

Sorry about this. There is a rich literature on the link between frontal theta and prediction error,[3,9–11] and we updated the Introduction to include more references to this work (p. 18): “The analysis was first done using power averaged across frontal electrodes, as these are the typical focus of PE research on oscillations.[3,9–11]”

We have also updated the Methods to cite more studies that motivate our electrode choice (p. 41): “The analyses first targeted five midline frontal electrodes (F3, F1, Fz, F2, F4; BioSemi64 layout), given that this frontal row is typically the focus of executive-function PE research on oscillations.[9–11]”

(R1.4c) No correction appears to have been applied for the association between EEG power and fMRI connectivity. Given that 100 frequency bins were collapsed into 5 canonical bands, a correction for 5 comparisons seems appropriate. Notably, the strongest effects in the delta and theta bands (particularly at fronto-central electrodes) may still survive correction, but this should be explicitly tested and reported.

Thanks for this suggestion. We updated the Table 1 caption to mention what results survive family-wise error rate correction – as the reviewer suggests, the Delta/Theta effects would survive Bonferroni correction for five tests, although per a later comment suggesting that we evaluate statistical significance with a permutation-testing approach (comment R3.4), we instead report family-wise error correction based on that. The revised caption is as follows (p. 19):

“The effects for Delta, Theta, Beta, and Gamma remain significant if significance testing is instead performed using permutation-testing and with family-wise error rate correction ($p_{\text{corrected}} < .05$).”

(R1.4d) Line 135. Not sure I understand what you mean by "moods". What is the overall point here?

The overall argument is that the fluctuations occur rapidly rather than slowly. By slow "moods" we refer to how a participant could enter a high anxiety state of >10 seconds, linked to high PE fluctuations, and then shift into a low anxiety state, linked to low PE fluctuations. We argue that this is not occurring. Regardless, we recognize that referring to lengths of time as short as 10 seconds or so is not a typical use of the word "mood" and is potentially ambiguous, so we have omitted this statement, which was originally on page 6: "Identifying subsecond fluctuations would broaden the relevance of the present results, as they rule out that the PE states derive from various moods."

(R1.4e) Line 100. "Few prior PE studies have targeted PE, contrasting the hundreds that have targeted BOLD". I don't understand this sentence. It's presumably about connectivity vs activity?

Yes, sorry about this typo. The reviewer is correct, and that sentence was meant to mention connectivity. We corrected (p. 5): "Few prior PE studies have targeted connectivity, contrasting the hundreds that have targeted BOLD."

(R1.4f) Line 373: "0-0.5Hz" in the caption is probably "0-50Hz".

Yes, this was another typo, thank you. We have corrected it (p. 19): "... every 0.5 Hz interval from 0-50 Hz."

Reviewer #2 (Recommendations for authors):

(R2.6) (Page 3) When referring to the "limited" hypothesis of local PE, please clarify in what sense is it limited. That statement is unclear.

Thank you for pointing out this text, which we now see is ambiguous. We originally use "limited" to refer to the hypothesis's constrained scope – namely, that PE is relevant to various low-level operations (e.g., sensory processing or rewards) but the minimization of PE does not guide more abstract cognitive processes. We edited this part of the Introduction to be clearer (p. 3)

"It is generally agreed that the brain uses PE mechanisms at neuronal or regional levels, [15,16] and this idea has been useful in various low-level functional domains, including early vision [15] and dopaminergic reward processing.[17] Some theorists have further argued that PE propagates through perceptual pathways and can elicit downstream cognitive processes to minimize PE."

(R2.7) (Page 5) "Few prior PE have targeted PE"... this statement appears contradictory. Please clarify.

Sorry about this typo, which we have corrected (p. 5):

"Few prior PE studies have targeted connectivity, contrasting the hundreds that have targeted BOLD."

(R2.8) What happened to the data of the medium PE condition in Study 1A?

The medium PE condition data were not excluded. We modeled the effect of prediction error on connectivity using a linear regression across the three conditions, coding them as a continuous variable (Low = -1, Medium = 0, High = +1). This approach allowed us to identify brain connections that showed a linear increase or decrease in strength as a function of increasing PE. This linear contrast is a more specific and powerful way to isolate PRelated

effects than a High vs. Low contrast. We updated the Results slightly to make this clearer (p. 8-9):

“In the fMRI data, we compared the three PE conditions’ beta-series functional connectivity, aiming to identify network-level signatures of PE processing, from low to high. [...] For the modularity analysis, we first defined a connectome matrix of beta values, wherein each edge’s value was the slope of a regression predicting that edge’s strength from PE (coded as Low = -1, Medium = 0, High = +1; Figure 2A).”

(R2.9) (Page 15) The point about how the dots in 6H follow those in 6J better than those in 6I is a little subjective - can the authors provide an objective measure?

Thank you for pointing out this issue. The visual comparison using Figure 6 was not meant as a formal analysis but rather to provide intuition. However, as the reviewer describes, this is difficult to convey. Our formal analysis is provided in Supplemental Materials 5, where we report correlation coefficients between a very large number of simulated fMRI data points and EEG data points corresponding to different frequencies. We updated this part of the Results to convey this (p. 16-17):

“Notice how the dots in Figure 6H follow the dots in Figure 6J (3 Hz) better than the dots in Figure 6I (0.5 Hz) or Figure 6K (10 Hz); this visual comparison is intended for illustrative purposes only, and quantitative analyses are provided in Supplemental Materials 5.”

References

- (1) Zalesky, A., Fornito, A. & Bullmore, E. T. Network-based statistic: identifying differences in brain networks. *Neuroimage* 53, 1197–1207 (2010)
- (2) Strijkstra, A. M., Beersma, D. G., Drayer, B., Halbesma, N. & Daan, S. Subjective sleepiness correlates negatively with global alpha (8–12 Hz) and positively with central frontal theta (4–8 Hz) frequencies in the human resting awake electroencephalogram. *Neuroscience letters* 340, 17–20 (2003).
- (3) Cavanagh, J. F. & Frank, M. J. Frontal theta as a mechanism for cognitive control. *Trends in cognitive sciences* 18, 414–421 (2014).
- (4) Grech, R. et al. Review on solving the inverse problem in EEG source analysis. *Journal of neuroengineering and rehabilitation* 5, 25 (2008)
- (5) Palva, J. M. et al. Ghost interactions in MEG/EEG source space: A note of caution on inter-areal coupling measures. *Neuroimage* 173, 632–643 (2018).
- (6) Koles, Z. J. Trends in EEG source localization. *Electroencephalography and clinical Neurophysiology* 106, 127–137 (1998).
- (7) Rosenberg, M. D. et al. A neuromarker of sustained attention from whole-brain functional connectivity. *Nature neuroscience* 19, 165–171 (2016).
- (8) Goodale, S. E. et al. fMRI-based detection of alertness predicts behavioral response variability. *elife* 10, e62376 (2021).
- (9) Cavanagh, J. F. Cortical delta activity reflects reward prediction error and related behavioral adjustments, but at different times. *NeuroImage* 110, 205–216 (2015)
- (10) Hoy, C. W., Steiner, S. C. & Knight, R. T. Single-trial modeling separates multiple overlapping prediction errors during reward processing in human EEG. *Communications Biology* 4, 910 (2021).

- (11) Neo, P. S.-H., Shadli, S. M., McNaughton, N. & Sellbom, M. Midfrontal theta reactivity to conflict and error are linked to externalizing and internalizing respectively. *Personality neuroscience* 7, e8 (2024).
- (12) Friston, K. J. The free-energy principle: a unified brain theory? *Nature reviews neuroscience* 11, 127–138 (2010)
- (13) Feldman, H. & Friston, K. J. Attention, uncertainty, and free-energy. *Frontiers in human neuroscience* 4, 215 (2010).
- (14) Friston, K. J. et al. Active inference and epistemic value. *Cognitive neuroscience* 6, 187–214 (2015).
- (15) Rao, R. P. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extraclassical receptive-field effects. *Nature neuroscience* 2, 79–87 (1999)
- (16) Walsh, K. S., McGovern, D. P., Clark, A. & O’Connell, R. G. Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the new York Academy of Sciences* 1464, 242– 268 (2020)
- (17) Niv, Y. & Schoenbaum, G. Dialogues on prediction errors. *Trends in cognitive sciences* 12, 265–272 (2008).
- (18) Benjamini, Y. & Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 289–300 (1995).

<https://doi.org/10.7554/eLife.107723.2.sa0>