

Running head: AVERAGE REWARD RATE

The average reward rate modulates behavioral and neural indices of effortful control allocation

Hause Lin^{1,2}, Jelena Ristic³, Michael Inzlicht^{4,5}, & A. Ross Otto³

¹Hill/Levene Schools of Business, University of Regina, Regina, Canada

²Sloan School of Management, Massachusetts Institute of Technology, Cambridge, USA

³Department of Psychology, McGill University, Montreal, Canada

⁴Department of Psychology, University of Toronto Scarborough, Toronto, Canada

⁵Rotman School of Management, University of Toronto, Toronto, Canada

Correspondence to: ross.otto@mcgill.ca

Abstract

People tend to avoid exerting cognitive effort, and findings from recent behavioral studies suggest that effort allocation is in part determined by the opportunity cost of slothful responding—operationalized as the average reward rate per unit time. When the average rate of reward is high, individuals make more errors in cognitive control tasks, presumably owing to reduction in cognitive control. An open question remains whether the presumed modulations of cognitively effortful control processes are observable at the neural level. Here, we measured EEG while participants completed the Simon task, a well-known response conflict task, while the experienced average reward rate fluctuated across trials. We examined neural activity associated with the opportunity cost of time by applying generalized eigendecomposition, a hypothesis-driven source separation technique, to identify a midfrontal component associated with the average reward rate. Fluctuations in average reward rate modulated not only component amplitude, but most importantly, component theta power (4-8 Hz). Higher average reward rate was associated with reduced theta power, suggesting that the average reward rate modulates effort allocation. These neural results provide evidence for the idea that individuals modulate the amount of cognitive effort they exert based on their experienced average rate of reward.

Introduction

Owing to the limited capacity nature of human information processing, we tend to expend cognitive effort only when it is worthwhile. In the influential framework of cost-benefit cognitive effort decision-making, people allocate cognitive resources to a particular task when the benefits of effort exertion—for example, reward incentives tied to performance—outweigh their perceived costs (Kool & Botvinick, 2018; Shenhav et al., 2017). On this view, a growing body of work demonstrates that individuals dynamically allocate their level of effort investment in accordance with shifting costs and benefits (Botvinick & Braver, 2015; Kool et al., 2017; Otto et al., 2022; Otto & Vassena, 2021; Sandra & Otto, 2018; Westbrook & Braver, 2015).

One important source of costs is the opportunity cost of time (Dora et al., 2021; Kurzban et al., 2013)—which has previously been formalized as the average reward rate per unit time (Niv et al., 2007). On this view, studies examining motor vigor—the costly outlay of effort assumed to be required for fast responding in free-operant tasks—find that individuals make faster responses in a simple vigilance task when the average rate of reward per unit time is high—suggesting that they balanced the costs of the harder work to emit faster actions (“vigor”) against the rewards foregone by responding slowly (Beierholm et al., 2013; Griffiths & Beierholm, 2017; Guitart-Masip et al., 2011). More recently, extending this idea to the domain of effortful cognitive tasks in a classic cognitive control task (Simon, 1990), we found that the average rate of reward per second time modulated the level of cognitive control that individuals applied towards inhibiting inappropriate, prepotent responses (Otto & Daw, 2019). During periods where participants experienced a higher rate of reward receipt per second, participants made more errors on difficult, incongruent trials, while participants made fewer errors on these trials when the experienced average reward rate was low. That is, when time was “expensive,”

individuals appeared to modulate (presumably effortful) controlled cognitive processing, which resulted in more errors on difficult, incongruent trials in the Simon task. More recently, this effect was replicated in a structurally similar Flanker task (Eriksen & Eriksen, 1974), which also requires cognitive control to inhibit distracting information (Devine et al., 2021).

In short, these behavioral studies suggest that moment-to-moment varying average rate of reward influences individual's strategic allocation of cognitive resources, which could be similar to the observed adjustments to cognitive control in accordance with, for example, recently experienced response conflict (Ridderinkhof, 2002) or cues signaling upcoming conflict (Gratton et al., 1992). However, it remains unclear whether this behavioral signature of average reward rate-evoked effort modulation (i.e., increased error rates on incongruent trials) is also accompanied by well-characterized neural signatures of cognitive control modulation—in particular, midfrontal theta oscillations (4-8 Hz) measured using electroencephalography (EEG; Cavanagh & Frank, 2014; Cooper et al., 2017, 2019; De Loof et al., 2019; Lin et al., 2018; Umemoto, Inzlicht, & Holroyd, 2019). Revealing such a relationship would demonstrate that shifts in the average reward rate prompts modulations of (effortful) cognitive control. Supporting this idea, a body of past work has observed increases in midfrontal theta power in response to cued conflict (van Driel et al., 2015), recently experienced conflict (Cohen & Cavanagh, 2011) in Simon-like tasks, deciding between similarly valued options in a delay discounting task (Lin et al., 2018), and strategic control allocation during decision making (Bogdanov et al., 2022; Cavanagh, Figueroa, et al., 2012).

Accordingly, we recorded EEG from participants while they performed a Simon response conflict task (Forstmann et al., 2008; Simon, 1990). This design allowed us to not only replicate existing behavioral findings (Devine et al., 2021; Otto & Daw, 2019), but to also extend them by

examining the average reward rate per unit time modulates the neural processes underlying cognitive control allocation. Following previous work, we manipulated the (experienced) average reward rate of the environment and investigated whether a high average reward rate brings about modulations in midfrontal theta power—thought to reflect adjustments to effortful cognitive control processes originating in the anterior cingulate cortex (Cavanagh & Frank, 2014; Cohen, 2014).

In this task (Figure 1A), participants are required to respond to a green circle using a right-hand response, and to a blue circle with a left-hand response. Because the stimulus can appear either on the left or right side of the display, on most trials (“congruent” trials) participants can effectively use the location of the stimulus to guide responses, but on incongruent trials, participants must ignore the location of the stimulus in order to make a correct, color-based response. As most trials (75%) are congruent, incongruent trials require overriding a prepotent, stimulus-driven response established by congruent trials, and as a result, responses are markedly slower and more error-prone. To manipulate the average reward rate, we induced random fluctuations in the reward available for making a correct response (Figure 1B), which we used to compute a time-varying “experienced” average rate of reward per second (Figure 1C). As participants had 5 minutes per block to complete as many trials as they could, the average reward rate effectively imposed an opportunity cost for slow responding (Beierholm et al., 2013; Guitart-Masip et al., 2011), which we have previously observed (Devine et al., 2021; Otto & Daw, 2019) increases error rates on difficult, incongruent trials. In other words, because this task is self-paced (with a fixed amount of time per block), a reward-maximizing strategy is

to make responses as quickly as possible to ensure as many trials as possible—and this is especially true when the average rate of reward is high (versus low).

If increases in the average reward rate prompts individuals to reduce their outlay of cognitive control, as suggested by our previous behavioral observations (e.g., Otto & Daw, 2019), we should observe neural activity that reflects preparatory adjustments to effortful control in accordance with this average reward rate. Since neural activity originating in the anterior cingulate cortex and surrounding midfrontal regions (i.e., 4-8 Hz oscillations) are thought to underlie control-related processes (Cavanagh & Frank, 2014; Cohen, 2014), we therefore predicted that higher average reward should be associated with reduced midfrontal theta power, reflecting trial-to-trial reductions in cognitive control prompted by the average reward rate per second. Because the modulatory effects of trial-to-trial fluctuations in average reward rate on neural activity are likely to be weak (especially relative to task-unrelated neural activity), we leveraged generalized eigendecomposition method (GED; Cohen, 2017, 2022), a hypothesis-driven source separation technique that can be used to increase signal-to-noise ratio, reduce data dimensionality, and identify task-relevant statistical neural sources. Specifically, following recent work that used GED to investigate the midfrontal neural components underlying cognitive control processing (Cohen, 2018; Zuure et al., 2020), we also used GED to investigate whether activity in midfrontal components is also modulated by trial-to-trial fluctuations in average reward rate.

Method

Participants

Twenty-six participants were recruited through McGill University's classified ads system and gave written consent in accordance with the McGill Research Ethics Board. Participants were compensated \$20 CAD for the session, and a performance bonus ranging from \$1-5. We excluded the data of one participant who missed over 40 response deadlines in the Simon Task, yielding 25 participants for the behavioral analyses. We further excluded the data of 3 participants for whom technical issues precluded the collection of usable EEG data, resulting a sample size of 22 participants in the EEG analyses. A sensitivity analysis indicated that we could detect EEG-based effect sizes of approximately $r = .29$ ($d = 0.60$) or greater with at least 80% statistical power.

Simon Task

Our version of the Simon task used blue and green circles as stimuli (Figure 1A). The blue or green color was either associated with a left- or right-hand response (the 'Z' or '/' buttons on the keyboard). Stimulus presentation sequence, triggers, and response timing were controlled by the Psychophysics Toolbox (Brainard, 1997). On each trial, a green or blue circle was presented on the left or right side of the screen (Figure 1A). In the version of the Simon task used here, 75% percent of trials were congruent—that is, the side on which the stimulus was presented matched the correct response hand. On the remaining 25% of trials, participants needed to use stimulus color and fully ignore the stimulus side in order to respond correctly.

Following a short practice task phase to gain familiarity with the task, participants then completed ten 5-minute blocks of the Simon task, completing as many trials as they could in

blocks of 5 minutes, mirroring our previous Simon task study (Otto & Daw, 2019; Experiment 2). During these blocks, the rewards available on each trial were determined using a Gaussian random walk with standard deviation 30 and reflecting boundaries at 5 and 95 cents (Figure 1B). At the outset of each trial, participants were presented visually with a number representing the reward on offer that trial, ranging from 1–100 cents (Figure 1A), which lasted from 750-1250ms, after which the Simon stimulus was displayed, and participants had 600ms to make a response. After the 500ms ISI following the response, reward obtained (e.g., “+9” for a correct response; or “0” otherwise) was displayed for 1000 ms. After each self-paced 5-minute block, participants were given a short break, and following the 10 five-minute blocks, participants were then paid a bonus proportional to their earnings in the task.

Behavioral Data Analysis

Following previous work, we calculated the average reward in units of reward per second, using the following update rule (Constantino & Daw, 2015; Otto & Daw, 2019):

$$\bar{r}_{t+1} = (1 - \alpha)^\tau \bar{r}_t + (1 - (1 - \alpha)^\tau) \frac{r}{\tau}$$

where r is the obtained reward on trial t , τ is the time elapsed since the last update (which depends, critically, on each trial’s RT and ITI), and α is a learning rate parameter. We fit a single learning rate to congruent trial RTs of the entire sample of participants, by running estimating a separate regression for each participant finding the learning rate that minimizes the total error across the group. Specifically, the participant-level RT regression included the following terms:

$$RT = \bar{r} + R + prev_error + ITI + trial_num + resp_side + stim_rep + prev_type$$

where RTs were log-transformed and z-scored response times, \bar{r} is the average reward rate, R is the reward available that trial, $prev_error$ and $prev_type$ are binary variables specifying whether an error response or incongruent stimulus (respectively) occurred on the previous trial, $trial_num$

is a linear term representing trial number (to capture practice effects), and *resp_side* represents whether a left or right-hand response was made (to capture simple response bias). Our estimation procedure yielded a best-fitting α of 0.0144.

To assess the influence of this inferred average reward rate (Figure 1C) upon log-transformed Simon RTs at the group level, we conducted mixed-effects regressions using the lme4 package (Pinheiro & Bates, 2000) for R, using the following formula:

$$\begin{aligned}
 RT \sim 0 &+ congruent + incongruent + congruent: (prv_errors + run_num + iti \\
 &+ trial_in_run + key_rep + prev_type + \bar{r} + R + resp_side) \\
 &+ incongruent: (prev_errors + run_num + iti + trial_in_run \\
 &+ key_rep + prev_type + \bar{r} + R + resp_side)
 \end{aligned}$$

As we examined congruent and incongruent trials separately, our regression models jointly analyzed behavior using two dummy variables (*congruent* and *incongruent*) specifying the trial type. Similarly, the effect of the average reward rate upon error rates was estimated using a logistic regression, taking incorrect (versus correct) responses as the outcome variable:

$$\begin{aligned}
 error \sim 0 &+ congruent + incongruent + congruent: (prv_errors + run_num \\
 &+ trial_in_run + key_rep + prev_type + \bar{r} + R + resp_side) \\
 &+ incongruent: (prev_errors + run_num + trial_in_run + key_rep \\
 &+ prev_type + \bar{r} + R + resp_side)
 \end{aligned}$$

All terms estimated at the fixed-effects level and as random effects at the participant level, taking all continuously-valued predictor variables as within-participant z-scores. Two-tailed probability values and degrees of freedom associated with each statistic was determined using the Satterthwaite approximation implemented in lmerTest (Kuznetsova et al., 2017).

Electroencephalography (EEG) Data and Processing

EEG data were acquired using an ActiveTwo system (BioSemi, Amsterdam, Netherlands) from 64 Ag/AgCl electrodes positioned according to a 10-20 international system using a typical 64-channel montage. The data were recorded in an electrically shielded room. Stimuli were presented on a 16-inch CRT monitor and viewed from an approximate distance of 120cm. Horizontal and vertical electrooculograms (EOG) were recorded from the electrodes placed above and below the left eye and 1cm lateral to the left and right canthi.

Offline, the EEG data were re-referenced to the average of electrodes placed on the two earlobes. The continuous data were high-pass filtered at 0.10 Hz (12 dB/oct, zero phase-shift Butterworth filter) and decomposed into independent components using the infomax independent component analysis algorithm implemented in the MATLAB toolbox EEGLAB (Delorme & Makeig, 2004). We inspected the independent components and used ICLabel, an extension for EEGLAB (Pion-Tonachini et al., 2019), to identify components that were classified as eye or muscle components. The algorithm assigns probabilities to 7 categories: brain, muscle, eye, heart, line noise, channel noise, other. The extension also provides an interface (see <https://sccn.ucsd.edu/wiki/ICLabel>) that shows the topography, time course, power spectrum, and ERP-image (sorted by trial number) of each component. Guided by ICLabel's classification algorithm, for each participant, we excluded, on average, 2 to 3 eye frontal components (e.g., blinks, vertical/horizontal eye movements) and 1 to 3 muscle components (usually components that showed maximal activity at temporal channels). Components were considered blinks or eye movement components and were excluded if (1) there was a high probability ($> 85\%$ & $< 1\%$ brain) of them being classified as an eye-related IC; (2) the IC time course activity resembled blinks or vertical/horizontal eye movements (e.g., activity that looks like step-functions); and (3) the topography showed maximal activity at frontal channels. Components were considered as

muscle components and were excluded if (1) there was a high probability ($> 95\%$ muscle & $< 1\%$ brain) of them being classified as a muscle component and (2) the power spectrum resembled noise or muscle activity more than neural activity (i.e., power peaks at higher frequencies rather than lower frequencies).

To prepare the preprocessed EEG data for generalized eigendecomposition (GED; see next section Statistical Source Separation of EEG Data), we epoched the single-trial data relative to cue onset and Simon stimulus onset (see Figure 1A). The single-trial cue-locked epochs (time-locked to reward onset; -2.5 to 2.5 s) were used to test the hypothesis that trial-to-trial fluctuations in average reward rate would modulate preparatory adjustments to effortful control. The single-trial stimulus-locked epochs (-0.2 to 0.8 s) were used to demonstrate, as a proof of principle, that the GED approach can effectively identify components or statistical sources associated with conflict and control-related neural processes.

The single-trial cue- and stimulus-locked epochs were baseline-corrected by subtracting the mean amplitude (-0.2 to 0 s) before their respective event onsets. The single-trial epochs were then lowpass filtered at 30 Hz with a finite impulse response (FIR) filter (Hamming window with 0.0194 passband ripple and 53 dB stopband attenuation; upper passband edge of 30 Hz; upper transition bandwidth of 7.50 Hz and -6 dB cutoff frequency of 33.75 Hz; filter length of 227 samples). Epochs containing artifacts, with amplitudes exceeding ± 150 μ V or gradients larger than 50 μ V, were excluded from further analysis.

For the stimulus-locked analysis, we focused on the N2 conflict-related ERP component (Yeung et al., 2004) and determined the component window (peak: 0.28 s) by inspecting the grand average event-related potential (ERP) waveform at three midfrontal channels (Fz, FCz, Cz; (see Luck & Gaspelin, 2017)). Then, for each single-trial epoch, we extracted the amplitude

(averaged across channels and all time points in a 0.06 s window [-0.25 to 0.31 s] around the peak) and fitted statistical models (see Statistical Analyses of EEG Data section) to these single-trial data to examine the effects of congruency on component amplitudes.

Statistical Source Separation of EEG Data

EEG data recorded at the scalp are a mixture of electric fields, produced by separate underlying neural sources, that propagate simultaneously via volume conduction to all EEG scalp channels. Whereas standard univariate channel-level analytic approaches do not directly address this source mixing problem, recent work has proposed generalized eigendecomposition-based (GED) multivariate methods that leverage the rich spatiotemporal dynamics of EEG data to decompose multi-channel time-series voltage data into independent but non-orthogonal statistical sources (Blankertz et al., 2007; Cohen, 2017, 2022; Parra et al., 2005; Parra & Sajda, 2003).

To validate the GED approach before applying it to test our main hypothesis, we first applied GED to the stimulus-locked single-trial epochs to examine classic conflict-related neural processes in the Simon conflict task (Yeung et al., 2004). After providing a proof-of-principle demonstration, we then applied GED to the cue-locked data to test our main hypothesis (i.e., average reward rate modulates midfrontal activity). For each of the two sets of analyses, we performed GED by first identifying two time-windows (Figure 2): One contained task-related neural signals (stimulus-locked analysis: 0.24 to 0.32 s post-stimulus onset; cue-locked analysis: 0.10 to 0.75 s post-cue onset) and the other task-unrelated or reference signals (-0.20 to 0 s pre-stimulus or pre-cue onset). Then, we computed the S (using task-related signals) and R (using reference or baseline signals) channel-by-channel covariance matrices separately for activity in the two windows. The covariance matrices were computed separately for each single-trial epoch

before averaging over epochs. GED was implemented using the *numpy* Python library and the method *numpy.linalg.eig(S, R)*.

Note that the stimulus-locked window (0.24 to 0.32 s) was partly informed by previous work (e.g., Cavanagh et al., 2012; Yeung et al., 2004). On the one hand, stimulus congruency often modulates the relatively transient midfrontal N2 ERP component (often observed about 0.2 to 0.3 s after stimulus onset). On the other hand, the cue-locked window was much longer (0.10 to 0.75 s) in the present study because we did not have specific hypotheses about when these effects would occur and how long they would last. However, as the average reward rate depends on the integration of multiple variables (e.g., reward on offer, time elapsed), it is likely that the underlying neurocognitive processes involved in its computation would last longer.

Since regularization can reduce noise and improve the quality of the decomposition (Wong et al., 2018), we added 1% shrinkage regularization to the R matrix as follows (Cohen, 2018; Zuure et al., 2020): $\tilde{R} = (1 - \gamma)R + \gamma\alpha I$, where $\alpha = n^{-1} \sum_{i=1}^n \lambda_i$, and $\gamma \in [0,1]$, where \tilde{R} is the regularized R covariance matrix, R is the original "reference" covariance matrix, I is the identity matrix, γ is the amount of regularization (can take only any value between 0 and 1), λ_i is the eigenvalue for eigenvector i , n is the number of eigenvectors, and α is the mean eigenvalue across all eigenvectors. To apply 1% shrinkage regularization, γ was set to 0.01. Note that if $\gamma = 0$, 0% or no regularization is applied, whereas $\gamma = 1$ is equivalent to 100% regularization, which essentially turns GED into principal component analysis (PCA).

Like PCA, GED decomposes the data into eigenvectors and eigenvalues (the number of eigenvectors is equivalent to the number of EEG channels), and each eigenvector, w , is a new basis vector that acts as a spatial filter. When applied to the data, a spatial filter up- or down-weighs different channels' activities to amplify or suppress activity before summing across all

weighted channels to derive the component time series, $w^T X$. The component's topography (also known as forward model of the filter or the activation pattern; Haufe et al., 2014) is computed by pre-multiplying the signal covariance matrix by the eigenvector, $w^T S$ (see component pattern in Figure 2).

We used GED to reduce the multi-dimensional EEG dataset into a single dimension or component. This component is a statistical source that maximizes differences between the post-event S and pre-event R activities, and it often has higher signal-to-noise ratio than the unfiltered EEG data (Cohen, 2017; Parra & Sajda, 2003). To identify this component, we visually inspected the component activation patterns and selected the midfrontal component with the largest eigenvalue for each participant (eigenvectors with larger eigenvalues explain more variance). For the stimulus-locked analysis, we selected a midfrontal component because the N2 component is usually localized to midfrontal channels (Yeung et al., 2004). We expected to find a GED component that has the characteristics of an N2 component—that is, the component amplitudes should be more negative on incongruent (vs. congruent) Simon-task trials. Similarly, we also selected a midfrontal component for the cue-locked analysis because previous work has consistently shown that that cognitive control processes (e.g., 4-8 Hz theta-band activity) are most apparent over midfrontal channels (Cavanagh, Zambrano-Vazquez, et al., 2012; Cavanagh & Frank, 2014). Thus, we expected the average reward rate to influence the amount of cognitive effort allocated to the upcoming Simon task trial, and that neural activity associated with effort allocation would be reflected in a GED component with a midfrontal topography.

Time-Frequency Analysis of Cue-Locked GED Component Activity

We performed time-frequency analysis on single-trial GED cue-locked component data using custom Python scripts and the MNE Python library (Gramfort et al., 2014, 2014). Each

trial consisted of 5 s (-2.5 to 2.5 s relative to cue onset) to avoid potential time-frequency decomposition edge artifacts. Time-frequency measures were computed by multiplying the fast Fourier transformed (FFT) power spectrum of single-trial EEG data with the FFT power spectrum of a set of complex Morlet wavelets, and taking the inverse FFT. The wavelet family is defined as a set of Gaussian-windowed complex sine waves, $e^{-i2\pi t f} e^{-t^2/2\sigma^2}$, where t is time, f is frequency (increased from 1 to 25 Hz), and σ is the width or cycles of each frequency band (increased from 3 to 10 in logarithmically spaced steps). Time-frequency power was normalized by converting to a decibel (dB) scale, $10 * \log_{10}\left(\frac{power_t}{power_{baseline}}\right)$, allowing different frequency bands to be directly compared, and $power_{baseline}$ is the mean power from -0.30 to -0.10 s pre-cue onset. Two windows were selected for further analysis based on examination of peak time-frequency points in the grand average time-frequency representations: The first window was 4-6 Hz at 0.15 to 0.25 s (captured the positive 5 Hz theta peak at 0.20s) and the second window was 6-8 Hz at 0.60 to 0.70 s (captured the negative 7 Hz theta peak at 0.65s). For each participant, single-trial power values in these two windows were exported for statistical analyses.

Statistical Analyses of EEG Data

Stimulus-locked analysis. To validate the GED approach, we fitted linear mixed-effects single-trial regression models to examine the effects of congruency (coded -0.5 [congruent] and 0.5 [incongruent], then within- participant z-scored) on the N2 ERP and GED component amplitudes. The models were fitted using (Pinheiro & Bates, 2000) [syntax: $lmer(amplitude \sim congruency + (1 + congruency | participant))$]. After validating the GED approach by showing that GED component amplitudes were sensitive to stimulus congruency, we then explored whether the N2 and GED component amplitudes were also sensitive to reward on offer and average reward rate [syntax: $lmer(amplitude \sim congruency * (reward_on_offer +$

$average_reward_rate) + (1 + congruency * (reward_on_offer + average_reward_rate) | participant)]$ (reward on offer and average reward rate were within-participant z-scored).

Cue-locked analysis. For each participant, we regressed the component amplitude at each time point on reward on offer and average reward rate ($y \sim b_0 + b_1r + b_2\bar{r}$) to obtain the regression coefficients for average reward rate. We then performed non-parametric permutation tests using the MNE library (Gramfort et al., 2013; Gramfort et al., 2014) to determine the temporal clusters where the regression coefficients significantly differed from 0 (5,000 permutations, threshold = 5.00, $p < .01$).

We fitted linear mixed-effects single-trial regression models to examine time-frequency power effects in the two windows described above. Time-frequency power values were regressed on reward on offer and average reward rate; they were nested within participants and the models included varying intercepts and slopes for both regressors, which were within-participant z-scored. The models were also fitted using [syntax: $lmer(power \sim reward_on_offer + average_reward_rate + (1 + reward_on_offer + average_reward_rate | participant))$].

Results

Behavioral Results

Mirroring typical Simon task performance (Forstmann et al., 2008; Simon, 1990), we found that participants made more errors (Figure 3A; mixed-effects logistic regression $\beta=2.835$, $SE=0.299$, $p < .0001$) and were slower to respond on incongruent trials compared to congruent trials (Figure 3B; mixed-effects regression $\beta= 0.1483$, $SE=0.0054$, $p < .0001$).

We then turned to examining the effect of the average reward rate upon Simon task performance. Notably, we found that a high average reward rate engendered a marked decrease

in accuracy on the more difficult incongruent trials (Figure 2A). Statistically, this accuracy effect was confirmed by a mixed-effect logistic regression, which revealed a significant effect of average reward rate upon error rates only on incongruent trials ($\beta=0.2332$, $SE=0.1094$, $p=0.033$; see Table 1 for full coefficient estimates) but not congruent trials ($\beta=-0.0635$, $SE=0.2221$, $p=0.775$). We observed a marginally significant linear contrast between the average reward effects on incongruent versus congruent trials ($p=0.063$). Further, we found no main effects of reward on offer on either trial type ($ps>0.785$), mirroring our previous work (Otto & Daw, 2019) and observations by others (Beierholm et al., 2013; Guitart-Masip et al., 2011). Put another way, the average reward rate of the environment, but not the reward available for making a correct response modulated effortful control.

While a high average reward rate appeared, visually, to slightly speed responses across both trial types (Figure 3B), the mixed-effects RT regression only revealed a non-significant (negative) effect of average reward rate on congruent trials ($\beta=-0.005$, $SE=0.0034$, $p=0.145$; see Table 2) but no apparent effect on incongruent trials ($\beta=-0.0004$, $SE=0.0057$, $p=0.946$).

EEG Results: Stimulus-Locked Conflict-Related Activity

Consistent with previous work (Yeung et al., 2004), the stimulus-locked N2 component was larger (i.e., more negative) for incongruent (vs. congruent) Simon-task trials ($\beta=-0.63$, $SE=0.13$, $p<.001$; Figure 4A). This component peaked at 0.28 s after stimulus onset: Contrasting the activity on incongruent and congruent trials revealed a predominantly central topographical distribution (Figure 4B right panel). Note that this topography is not evident when averaging activity across all trials (Figure 4B, left panel), likely because activity in other regions (e.g., P3 in the parietal region) is stronger and thus masks the relatively weak N2 effects.

We then applied the GED approach to all the stimulus-locked EEG trials to demonstrate, as a proof of principle, that the GED approach not only addresses the volume conduction problem associated with regular ERP analyses, but most importantly, isolates task-relevant relevant statistical sources (Cohen, 2017)—and by extension, cognitive processes of interest (see Figure 2). As a large body of work has linked activity over midfrontal channels to conflict processing and/or cognitive control processes (Cavanagh, Zambrano-Vazquez, et al., 2012; Cavanagh & Frank, 2014), we therefore identified, for each participant, a component with maximally midfrontal spatial distribution (Figure 4D). As with the N2 component (Figure 4A), activity in this GED component (peaked at around 0.27 s) was also larger for incongruent (vs. congruent) Simon-task trials ($\beta=-0.25$, $SE=0.07$, $p<.001$; Figure 4C). Unlike the ERP topography for all trials (Figure 4B left), the GED component spatial distributions (i.e., activation patterns) across all trials and the difference between incongruent and congruent trials were maximal at midfrontal channels (Fig. 4D), which is unsurprising given that we explicitly selected, separately for each participant, a component with a predominantly midfrontal spatial distribution.

Having established the effectiveness of the GED approach for separating sources that are mixed at the scalp via volume conduction and isolating stimulus-locked conflict processes, we then explored whether trial-by-trial fluctuations in reward on offer and average reward rate also modulated these stimulus-locked conflict processes. GED component amplitude was not significantly modulated by reward on offer ($\beta=0.002$, $SE=0.05$, $p=0.968$), average reward ($\beta=0.06$, $SE=0.06$, $p=0.288$), and their interactions with stimulus congruency ($ps>0.151$). Similarly, N2 amplitude was not significantly modulated by ($\beta=0.07$, $SE=0.15$, $p=0.618$), average reward ($\beta=0.09$, $SE=0.14$, $p=0.519$), and their interactions with stimulus congruency

($p > 0.135$). These results suggest that stimulus-locked reactive control was not modulated by reward on offer and average reward rate.

EEG Results: Cue-Locked Activity Associated with Average Reward Rate Processing

We now turn to our primary analysis of interest, in which we examine whether preparatory (proactive) control neural activity is modulated by fluctuations in average reward rate. We applied GED to the pre-stimulus cue-locked component and identified a component with midfrontal topography in each participant (Figure 5A; see inset) because the behavioral results from the current and previous studies suggest that average reward rate modulates effortful control allocation (Devine et al., 2021; Otto & Daw, 2019; Sharp et al., in preparation), which, in turn, has been associated with changes in theta-band dynamics over midfrontal channels (Cavanagh, Figueroa, et al., 2012). The time-series activity of the midfrontal component is depicted in the top panel of Figure 5A.

To examine whether component amplitude correlated with average reward rate, for each participant, we regressed the amplitude at each timepoint on reward on offer, r , and average reward rate, \bar{r} . We then performed non-parametric permutation tests to determine the time windows or temporal clusters where the regression coefficients for average reward rate were statistically significant and different from 0 ($p < .01$). The t-statistic time series for the significant cluster (0.04 to 0.75 s after cue onset) of regression coefficients are shown Figure 5A (bottom panel). Importantly, we observed a significant negative relationship between average reward rate and component amplitude: a higher average experienced reward rate—which was associated with increased errors on incongruent trials—was associated with reduced amplitude of the component identified by GED (Figure 5B), and this effect was statistically significant for an extended period (at least 0.71 s; Figure 5A; blue shaded regions). Note that this effect may have

lasted even longer, but we had to restrict our analyses from 0 to 0.75 s to prevent overlap with activity related to the processing of the stimulus, which was shown between 0.75 to 1.25 s after cue onset (see Figure 1A).

The above results suggest that the GED approach has effectively isolated a midfrontal component whose source activity reflects the neural activity associated with average reward rate processing across trials. We next turned to time-frequency analysis of the component time-series data to examine whether average reward rate (but not reward on offer) modulated theta power (4-8 Hz) in this midfrontal component. The analysis revealed two distinct time windows where we observed changes in theta power (Figure 5C). Theta power in the earlier window (4-6 Hz, 0.15 to 0.25 s) did not correlate with reward on offer ($\beta=0.02$, $SE=0.02$, $p=.739$) or average reward rate ($\beta=0.004$, $SE=0.01$, $p=.432$). However, theta power in the later window (6-8 Hz, 0.60 to 0.70 s) correlated negatively with average reward rate ($\beta=-0.02$, $SE=0.01$, $p=.043$), whereas reward on offer did not correlate with theta power ($\beta=-0.01$, $SE=0.01$, $p=.256$). To ensure these estimates had not been biased by outliers (because single-trial power estimates can be relatively noisy), we applied a robust outlier-detection approach, which identified 5.88% of the trials with power estimates ± 3 times the median absolute deviation (Leys et al., 2019). After excluding these single-trial theta power outliers, we again found a negative effect for average reward rate ($\beta=-0.03$, $SE=0.01$, $p=.010$) and a non-significant effect for reward on offer ($\beta=-0.01$, $SE=0.01$, $p=.589$). That is, when average reward rate is high, theta power in this midfrontal component was reduced, suggesting reduced cognitive control allocation (Figure 5D). Mirroring the behavioral analyses, the reward on offer did not appear to modulate theta power in either direction.

Discussion

The idea that we invest (versus withhold) cognitive effort in accordance with the costs and benefits of effort exertion has been an influential proposal which has found broad empirical support and stemmed considerable interest in the last decade (Frömer et al., 2021; Inzlicht et al., 2018; Kool & Botvinick, 2018; Kurzban et al., 2013; Otto et al., 2022; Shenhav et al., 2017). Within this framework, here we considered the more specific hypothesis that individuals should reduce their use of cognitively costly processing in accordance with the time-varying, experienced average reward rate of the environment (Otto & Daw, 2019), examining behavioral and EEG signatures of cognitive control modulations in accordance with the average reward rate, in a Simon response conflict task. Behaviorally, we observed that individuals modulated their level of (presumably effortful) controlled processing—manifesting as higher error rates on difficult, incongruent trials—when the average reward rate was high but readily employed cognitive control when the average reward rate was low, replicating our previous performance results in simple response conflict tasks (Devine et al., 2021; Otto & Daw, 2019).

Examining neural activity by applying generalized eigendecomposition (GED) to perform statistical source separation on each participant's single-trial EEG time-series data (Cohen, 2017), we could identify, separately for each participant, one predominantly midfrontal component to examine task-relevant cue-locked neural activity that occurred prior to presentation of the Simon stimulus (i.e., preparatory adjustments to effortful control). We found that the average rate of reward correlated negatively with theta power (6-8 Hz) in a midfrontal component identified via GED. Specifically, on trials in which the average reward rate per second was high, theta power was reduced in the GED-identified midfrontal component; conversely, when this average reward rate was low, component theta power was enhanced. In

other words, echoing the observed modulation of effortful behavior (evinced by incongruent trial accuracy; Figure 3A), the trial-to-trial strength of participants' component theta power—which we take to index preparatory effortful control allocation (Cavanagh & Frank, 2014; Umemoto et al., 2019)—also varied inversely with the experienced trial-to-trial environmental average reward rate. Together, these behavioral and neural results provide evidence for the idea that the average reward rate per unit time modulates adjustments to effortful control.

Indeed, a body of previous work has found that activity in midfrontal regions (in particular, theta-band oscillatory activity) relates to the application of effortful control processes involved in processing conflicting stimulus-response requirements (Cavanagh & Frank, 2014; Cohen, 2014, 2017). Enhancements in the strength of these theta oscillations have been observed following events that suggest an increased need for action monitoring—for example, following an error or the experience of response conflict (Cavanagh, Zambrano-Vazquez, et al., 2012). On this view, these frontal midline theta dynamics are thought to reflect synchronization of goal-relevant information in the service of successful action selection, possibly reflecting the control function of the anterior cingulate cortex (ACC; Cavanagh & Frank, 2014; Holroyd & Umemoto, 2016). Of note, a large body of theoretical and empirical work has implicated the ACC in both cost-benefit valuation of cognitive control and regulation of the level control applied (Shenhav et al., 2017; Fromer et al., 2021). In light of this proposed role for ACC function, the observed trial-to-trial modulations of cue-locked (pre-stimulus) midfrontal component theta power observed here suggest the possibility that these modulations of midfrontal activity reflect reductions in effortful cognitive control prompted by the experienced average reward rate.

Critically, our analyses of average reward-induced modulations in the strength of midfrontal component power focused on the period of each trial *before* the Simon stimulus onset

(i.e., preparatory/proactive effortful control adjustments; see Figure 1A) which suggests that these modulations of component theta power did not stem from the congruence of the current trial (as participants would have no knowledge of stimulus congruence during this period) but rather appear to evidence a calculation of background average reward rate computed on the basis of recent history of rewards obtained per second. Crucially, even though the GED analysis was not informed by the average reward rate or reward on offer on each trial during the pre-stimulus period, activity in the midfrontal component we identified was selectively sensitive to only average reward rate, but not reward on offer. This dissociation mirrors participants' lack of behavioral sensitivity to reward on offer—observed both here and in previous studies examining cognitive control (Devine et al., 2021; Otto & Daw, 2019; Sharp et al., in preparation) and motor vigor (Beierholm et al., 2013; Guitart-Masip et al., 2011), which employ an identical average reward rate manipulation. A feature of this available reward manipulation worth noting is that reward incentives change on a trial-to-trial basis, rather than manipulated in a block-wise fashion, as is often found in studies examining reward-motivated cognitive control (Chiew & Braver, 2014; Locke & Braver, 2008; Massar et al., 2016; Otto & Vassena, 2021). One possibility for the apparent lack of sensitivity to reward on offer levels here—and in our previous work—is that individuals are reluctant to adjust their control levels if they believe these control levels will only be appropriate for very short periods of time. This idea dovetails with a recent theoretical proposal, building upon the cost-benefit model of effort allocation, that a “reconfiguration cost” accompanying adjustments to control levels rides atop a control cost (Grahek et al., 2022). Crucially, average reward rate and reward on offer did not modulate stimulus-locked neural activity, suggesting that they did not influence reactive control. This finding is consistent with recent work showing that people combine information about expected

reward and task efficacy to proactively (but not reactively) adjust control allocation (Frömer et al., 2021). Together, these behavioral and neural results provide evidence for the idea that the average reward rate modulates preparatory but not reactive adjustments to effortful control.

Further, as a proof of principle, we also used GED to investigate a well-characterized ERP component—namely the stimulus-locked N2, a negative wave typically observed between 200 and 350 ms, which has been associated with incongruence in Simon and Simon-like response conflict tasks (Folstein & Van Petten, 2008). As expected, a traditional ERP analysis revealed a stronger negative deflection on incongruent (versus congruent) trials across midline channels (Fz, FCz, and Cz). Mirroring this result, our analysis of midfrontal GED-identified component time-series data revealed a similar pattern of deflections (i.e., more negative deflections on incongruent versus congruent trials) in a nearly identical time window. In other words, our GED analysis of stimulus-locked EEG was able to uncover component deflections that resemble the typical patterns of N2 obtained in traditional ERP analyses.

Our results also highlight the utility of applying multivariate source separation techniques that not only decompose multi-channel EEG data into independent sources, but also simultaneously highlight task-relevant neural activity and de-emphasize task-unrelated background activity (e.g., Cohen, 2017). The midfrontal component we identified clearly captured neural activity associated with trial-to-trial fluctuations in average reward rate but not reward on offer, but this dissociation only indicates that reward on offer was not tracked by this particular midfrontal component. In fact, because GED produces as many statistical components as the number of channels in the EEG data, it is likely another component's activity might covary with reward on offer but not average reward rate. However, this possibility does not

affect our main finding that average reward rate modulates proactive allocation of cognitive resources, and that we found evidence for this modulation in a midfrontal component.

Finally, it is important to note that the strength of conclusions drawn from these EEG-based results is constrained by the exploratory nature of our analysis and the limited sample size of our study (N=22). While we did observe a statistically significant behavioral effect of the average reward rate upon incongruent trial accuracy—thereby replicating a series of previous results obtained using employing Simon or Simon-like tasks (Devine et al., 2021; Otto & Daw, 2019; Sharp et al., in preparation)—further research is needed to bolster our conclusions, especially with regards to the relationship between EEG and average rate of reward. Specifically, we see a need for independent and high-powered studies to 1) directly replicate how the experienced average reward rate correlates with midfrontal activity, and 2) to probe for possible correspondences between behavioral reactivity (i.e., average reward-rate-driven modulations of accuracy) and average reward rate-induced changes in midfrontal oscillatory power, both at the intra- and interindividual levels.

Acknowledgements

This research was supported by NSERC, SSHRC, and William Dawson funds to JR, grants from the Fonds de Research du Quebec – Nature et Technologies, NSERC and a CFI Infrastructure Grant to ARO, and a grant from the Natural Sciences and Engineering Research Council of Canada to MI. The authors are grateful to Jenna Morris for assistance with data collection.

References

- Beierholm, U., Guitart-Masip, M., Economides, M., Chowdhury, R., Düzel, E., Dolan, R., & Dayan, P. (2013). Dopamine Modulates Reward-Related Vigor. *Neuropsychopharmacology*, *38*(8), 1495–1503. <https://doi.org/10.1038/npp.2013.48>
- Blankertz, B., Tomioka, R., Lemm, S., Kawanabe, M., & Müller, K. (2007). Optimizing Spatial filters for Robust EEG Single-Trial Analysis. *IEEE Signal Processing Magazine*, *25*(1), 41–56. <https://doi.org/10.1109/MSP.2008.4408441>
- Bogdanov, M., Renault, H., LoParco, S., Weinberg, A., & Otto, A. R. (2022). Cognitive Effort Exertion Enhances Electrophysiological Responses to Rewarding Outcomes. *Cerebral Cortex*, bhab480. <https://doi.org/10.1093/cercor/bhab480>
- Botvinick, M., & Braver, T. (2015). Motivation and Cognitive Control: From Behavior to Neural Mechanism. *Annual Review of Psychology*, *66*(1), 83–113. <https://doi.org/10.1146/annurev-psych-010814-015044>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. <https://doi.org/10.1163/156856897X00357>
- Cavanagh, J. F., Figueroa, C. M., Cohen, M. X., & Frank, M. J. (2012). Frontal Theta Reflects Uncertainty and Unexpectedness during Exploration and Exploitation. *Cerebral Cortex*, *22*(11), 2575–2586. <https://doi.org/10.1093/cercor/bhr332>
- Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, *18*(8), 414–421. <https://doi.org/10.1016/j.tics.2014.04.012>
- Cavanagh, J. F., Zambrano-Vazquez, L., & Allen, J. J. B. (2012). Theta lingua franca: A common mid-frontal substrate for action monitoring processes. *Psychophysiology*, *49*(2), 220–238. <https://doi.org/10.1111/j.1469-8986.2011.01293.x>

- Chiew, K. S., & Braver, T. S. (2014). Dissociable influences of reward motivation and positive emotion on cognitive control. *Cognitive, Affective, & Behavioral Neuroscience*, *14*(2), 509–529. <https://doi.org/10.3758/s13415-014-0280-0>
- Cohen, M. X. (2014). A neural microcircuit for cognitive conflict detection and signaling. *Trends in Neurosciences*, *37*(9), 480–490. <https://doi.org/10.1016/j.tins.2014.06.004>
- Cohen, M. X. (2017). Multivariate cross-frequency coupling via generalized eigendecomposition. *ELife*, *6*, e21792. <https://doi.org/10.7554/eLife.21792>
- Cohen, M. X. (2018). Using spatiotemporal source separation to identify prominent features in multichannel data without sinusoidal filters. *European Journal of Neuroscience*, *48*(7), 2454–2465. <https://doi.org/10.1111/ejn.13727>
- Cohen, M. X. (2022). A tutorial on generalized eigendecomposition for denoising, contrast enhancement, and dimension reduction in multichannel electrophysiology. *NeuroImage*, *247*, 118809. <https://doi.org/10.1016/j.neuroimage.2021.118809>
- Cohen, M. X., & Cavanagh, J. F. (2011). Single-Trial Regression Elucidates the Role of Prefrontal Theta Oscillations in Response Conflict. *Frontiers in Psychology*, *2*. <https://doi.org/10.3389/fpsyg.2011.00030>
- Constantino, S. M., & Daw, N. D. (2015). Learning the opportunity cost of time in a patch-foraging task. *Cognitive, Affective, & Behavioral Neuroscience*, 1–17. <https://doi.org/10.3758/s13415-015-0350-y>
- Cooper, P. S., Karayanidis, F., McKewen, M., McLellan-Hall, S., Wong, A. S. W., Skippen, P., & Cavanagh, J. F. (2019). Frontal theta predicts specific cognitive control-induced behavioural changes beyond general reaction time slowing. *NeuroImage*, *189*, 130–140. <https://doi.org/10.1016/j.neuroimage.2019.01.022>

- Cooper, P. S., Wong, A. S. W., McKewen, M., Michie, P. T., & Karayanidis, F. (2017). Frontoparietal theta oscillations during proactive control are associated with goal-updating and reduced behavioral variability. *Biological Psychology, 129*, 253–264. <https://doi.org/10.1016/j.biopsycho.2017.09.008>
- De Loof, E., Vassena, E., Janssens, C., De Taeye, L., Meurs, A., Van Roost, D., Boon, P., Raedt, R., & Verguts, T. (2019). Preparing for hard times: Scalp and intracranial physiological signatures of proactive cognitive control. *Psychophysiology, 0(0)*, e13417. <https://doi.org/10.1111/psyp.13417>
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods, 134(1)*, 9–21. <https://doi.org/10.1016/j.jneumeth.2003.10.009>
- Devine, S., Neumann, C., Otto, A. R., Bolenz, F., Reiter, A., & Eppinger, B. (2021). Seizing the opportunity: Lifespan differences in the effects of the opportunity cost of time on cognitive control. *Cognition, 216*, 104863. <https://doi.org/10.1016/j.cognition.2021.104863>
- Dora, J., van Hooff, M. L. M., Geurts, S. A. E., Kompier, M. A. J., & Bijleveld, E. (2021). The effect of opportunity costs on mental fatigue in labor/leisure trade-offs. *Journal of Experimental Psychology: General*, No Pagination Specified-No Pagination Specified. <https://doi.org/10.1037/xge0001095>
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics, 16(1)*, 143–149. <https://doi.org/10.3758/BF03203267>

- Folstein, J. R., & Van Petten, C. (2008). Influence of cognitive control and mismatch on the N2 component of the ERP: A review. *Psychophysiology*, *45*(1), 152–170.
<https://doi.org/10.1111/j.1469-8986.2007.00602.x>
- Forstmann, B. U., van den Wildenberg, W. P. M., & Ridderinkhof, K. R. (2008). Neural Mechanisms, Temporal Dynamics, and Individual Differences in Interference Control. *Journal of Cognitive Neuroscience*, *20*(10), 1854–1865.
<https://doi.org/10.1162/jocn.2008.20122>
- Frömer, R., Lin, H., Dean Wolf, C. K., Inzlicht, M., & Shenhav, A. (2021). Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications*, *12*(1), 1030. <https://doi.org/10.1038/s41467-021-21315-z>
- Grahek, I., Leng, X., Prater Fahey, P., Yee, D. M., & Shenhav, A. (2022). Empirical and Computational Evidence for Reconfiguration Costs During Within-Task Adjustments in Cognitive Control. *Proceedings of the 44th Annual Meeting of the Cognitive Science Society*.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Parkkonen, L., & Hämäläinen, M. S. (2014). MNE software for processing MEG and EEG data. *NeuroImage*, *86*, 446–460. <https://doi.org/10.1016/j.neuroimage.2013.10.027>
- Gramfort, A., Strohmeier, D., Haueisen, J., Hämäläinen, M. S., & Kowalski, M. (2013). Time-frequency mixed-norm estimates: Sparse M/EEG imaging with non-stationary source activations. *NeuroImage*, *70*, 410–422. <https://doi.org/10.1016/j.neuroimage.2012.12.051>
- Gratton, G., Coles, M., & Donchin, E. (1992). Optimizing the use of information: Strategic control of activation of responses. *Journal of Experimental Psychology: General*, *121*(4), 480–506. <https://doi.org/10.1037/0096-3445.121.4.480>

- Griffiths, B., & Beierholm, U. R. (2017). Opposing effects of reward and punishment on human vigor. *Scientific Reports*, 7(1), 42287. <https://doi.org/10.1038/srep42287>
- Guitart-Masip, M., Beierholm, U. R., Dolan, R., Duzel, E., & Dayan, P. (2011). Vigor in the Face of Fluctuating Rates of Reward: An Experimental Examination. *Journal of Cognitive Neuroscience*, 23(12), 3933–3938. https://doi.org/10.1162/jocn_a_00090
- Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J.-D., Blankertz, B., & Bießmann, F. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage*, 87, 96–110. <https://doi.org/10.1016/j.neuroimage.2013.10.067>
- Holroyd, C. B., & Umemoto, A. (2016). The research domain criteria framework: The case for anterior cingulate cortex. *Neuroscience & Biobehavioral Reviews*, 71, 418–443. <https://doi.org/10.1016/j.neubiorev.2016.09.021>
- Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The Effort Paradox: Effort Is Both Costly and Valued. *Trends in Cognitive Sciences*, 22(4), 337–349. <https://doi.org/10.1016/j.tics.2018.01.007>
- Kool, W., & Botvinick, M. (2018). Mental labour. *Nature Human Behaviour*, 1. <https://doi.org/10.1038/s41562-018-0401-9>
- Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost-Benefit Arbitration Between Multiple Reinforcement-Learning Systems. *Psychological Science*, 28(9), 1321–1333. <https://doi.org/10.1177/0956797617708288>
- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, 36(06), 661–679. <https://doi.org/10.1017/S0140525X12003196>

- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(13).
<https://doi.org/10.18637/jss.v082.i13>
- Leys, C., Delacre, M., Mora, Y. L., Lakens, D., & Ley, C. (2019). How to Classify, Detect, and Manage Univariate and Multivariate Outliers, With Emphasis on Pre-Registration. *International Review of Social Psychology*, 32(1), 5. <https://doi.org/10.5334/irsp.289>
- Locke, H. S., & Braver, T. S. (2008). Motivational influences on cognitive control: Behavior, brain activation, and individual differences. *Cognitive, Affective, & Behavioral Neuroscience*, 8(1), 99–112. <https://doi.org/10.3758/CABN.8.1.99>
- Luck, S. J., & Gaspelin, N. (2017). How to get statistically significant effects in any ERP experiment (and why you shouldn't). *Psychophysiology*, 54(1), 146–157.
<https://doi.org/10.1111/psyp.12639>
- Massar, S. A. A., Lim, J., Sasmita, K., & Chee, M. W. L. (2016). Rewards boost sustained attention through higher effort: A value-based decision making approach. *Biological Psychology*, 120, 21–27. <https://doi.org/10.1016/j.biopsycho.2016.07.019>
- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology*, 191(3), 507–520.
<https://doi.org/10.1007/s00213-006-0502-4>
- Otto, A. R., Braem, S., Silvetti, M., & Vassena, E. (2022). Is the juice worth the squeeze? Learning the marginal value of mental effort over time. *Journal of Experimental Psychology: General*, No Pagination Specified-No Pagination Specified.
<https://doi.org/10.1037/xge0001208>

- Otto, A. R., & Daw, N. D. (2019). The opportunity cost of time modulates cognitive effort. *Neuropsychologia*, *123*, 92–105. <https://doi.org/10.1016/j.neuropsychologia.2018.05.006>
- Otto, A. R., & Vassena, E. (2021). It's all relative: Reward-induced cognitive control modulation depends on context. *Journal of Experimental Psychology: General*, *150*(2), 306–313. <https://doi.org/10.1037/xge0000842>
- Parra, L. C., & Sajda, P. (2003). Blind source separation via generalized eigenvalue decomposition. *The Journal of Machine Learning Research*, *4*(null), 1261–1269.
- Parra, L. C., Spence, C. D., Gerson, A. D., & Sajda, P. (2005). Recipes for the linear analysis of EEG. *NeuroImage*, *28*(2), 326–341. <https://doi.org/10.1016/j.neuroimage.2005.05.032>
- Pinheiro, J. C., & Bates, D. M. (2000). *Mixed-Effects Models in S and S-PLUS*. Springer.
- Pion-Tonachini, L., Kreutz-Delgado, K., & Makeig, S. (2019). ICLabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, *198*, 181–197. <https://doi.org/10.1016/j.neuroimage.2019.05.026>
- Ridderinkhof, R. K. (2002). Micro- and macro-adjustments of task set: Activation and suppression in conflict tasks. *Psychological Research*, *66*(4), 312–323. <https://doi.org/10.1007/s00426-002-0104-7>
- Sandra, D. A., & Otto, A. R. (2018). Cognitive capacity limitations and Need for Cognition differentially predict reward-induced cognitive effort expenditure. *Cognition*, *172*, 101–106. <https://doi.org/10.1016/j.cognition.2017.12.004>
- Sharp, M., Beierholm, U., LoParco, S., & Otto, A. R. (in preparation). *Dopamine modulates the effect of opportunity costs on cognitive effort allocation*.

- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a Rational and Mechanistic Account of Mental Effort. *Annual Review of Neuroscience*, *40*(1), 99–124. <https://doi.org/10.1146/annurev-neuro-072116-031526>
- Simon, J. R. (1990). The effects of an irrelevant directional cue on human information processing. In R. W. Proctor & T. G. Reeve (Eds.), *Stimulus-response compatibility: An integrated perspective* (pp. 31–86). North-Holland.
- Umemoto, A., Inzlicht, M., & Holroyd, C. B. (2019). Electrophysiological indices of anterior cingulate cortex function reveal changing levels of cognitive effort and reward valuation that sustain task performance. *Neuropsychologia*, *123*, 67–76. <https://doi.org/10.1016/j.neuropsychologia.2018.06.010>
- van Driel, J., Swart, J. C., Egner, T., Ridderinkhof, K. R., & Cohen, M. X. (2015). (No) time for control: Frontal theta dynamics reveal the cost of temporally guided conflict anticipation. *Cognitive, Affective, & Behavioral Neuroscience*, *15*(4), 787–807. <https://doi.org/10.3758/s13415-015-0367-2>
- Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience*, 1–21. <https://doi.org/10.3758/s13415-015-0334-y>
- Wong, D. D. E., Fuglsang, S. A., Hjortkjær, J., Ceolini, E., Slaney, M., & de Cheveigné, A. (2018). A Comparison of Regularization Methods in Forward and Backward Models for Auditory Attention Decoding. *Frontiers in Neuroscience*, *12*. <https://doi.org/10.3389/fnins.2018.00531>
- Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The Neural Basis of Error Detection: Conflict Monitoring and the Error-Related Negativity. *Psychological Review*, *111*(4), 931–959. <https://doi.org/10.1037/0033-295X.111.4.931>

Zuure, M. B., Hinkley, L. B., Tiesinga, P. H. E., Nagarajan, S. S., & Cohen, M. X. (2020).

Multiple Midfrontal Thetas Revealed by Source Separation of Simultaneous MEG and EEG. *Journal of Neuroscience*, 40(40), 7702–7713.

<https://doi.org/10.1523/JNEUROSCI.0321-20.2020>

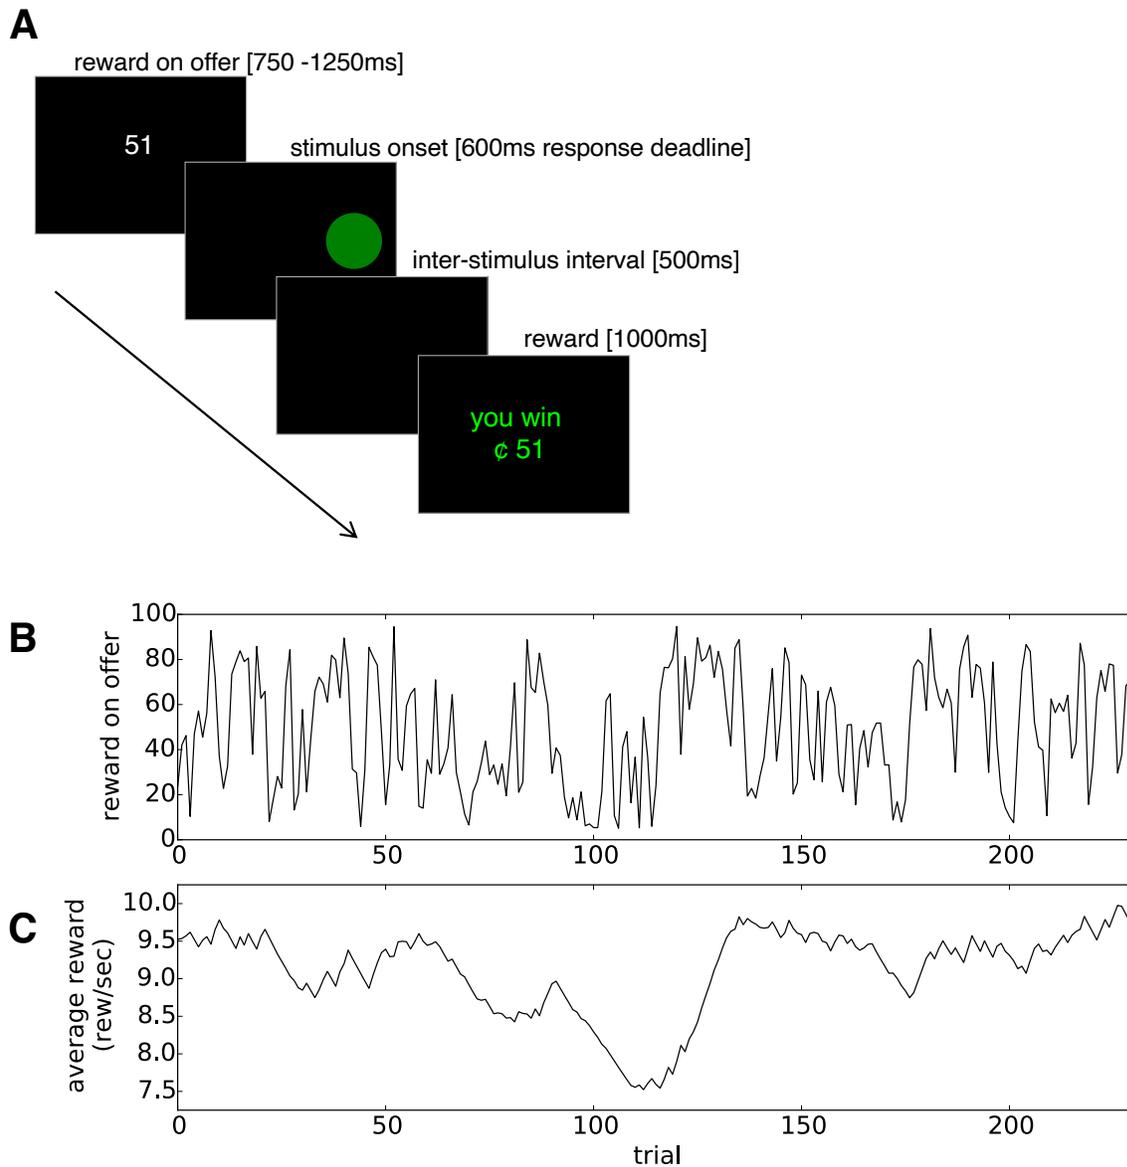


Figure 1. (A) Task flow in the Simon task. Before the stimulus is displayed, participants were shown the potential reward for making a correct response, after which they responded to a circle on the basis of its color, ignoring its location. (B) We induced random fluctuations in trial-to-trial available rewards (C) An example participant's experienced average reward rate, in units of reward per second, computed jointly from the participant's history of rewards and RTs, yielded an experienced, empirical average reward.

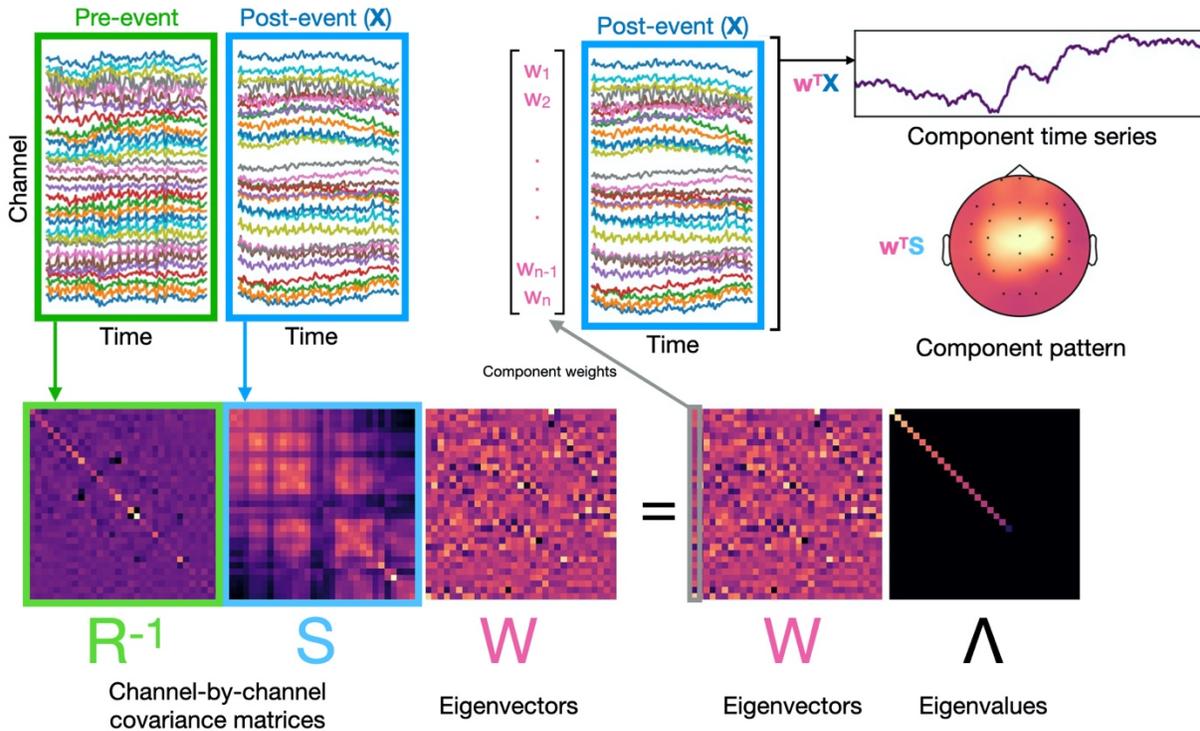


Figure 2. Generalized eigendecomposition (GED) on EEG data. From the recorded EEG data, we defined two periods of activity: pre-event (green) and post-event (blue) periods. We then compute the respective channel-by-channel covariance matrices, R and S (typically, separately for each epoch and then compute the mean covariance across epochs). GED decomposes multi-channel EEG data into independent but non-orthogonal sources by finding weighted combinations of activities across different channels. The weights, W , are defined by some criteria specified by the contrast between the S and R covariance matrices. To generate a single component time series, we select a column in W —i.e., one eigenvector—and use it to compute the linear weighted sum of activity across all channels. The resulting component maximizes the difference between activities in the post-event (blue) and pre-event (green) periods, resulting in hypothesis-driven source separation and dimension reduction.

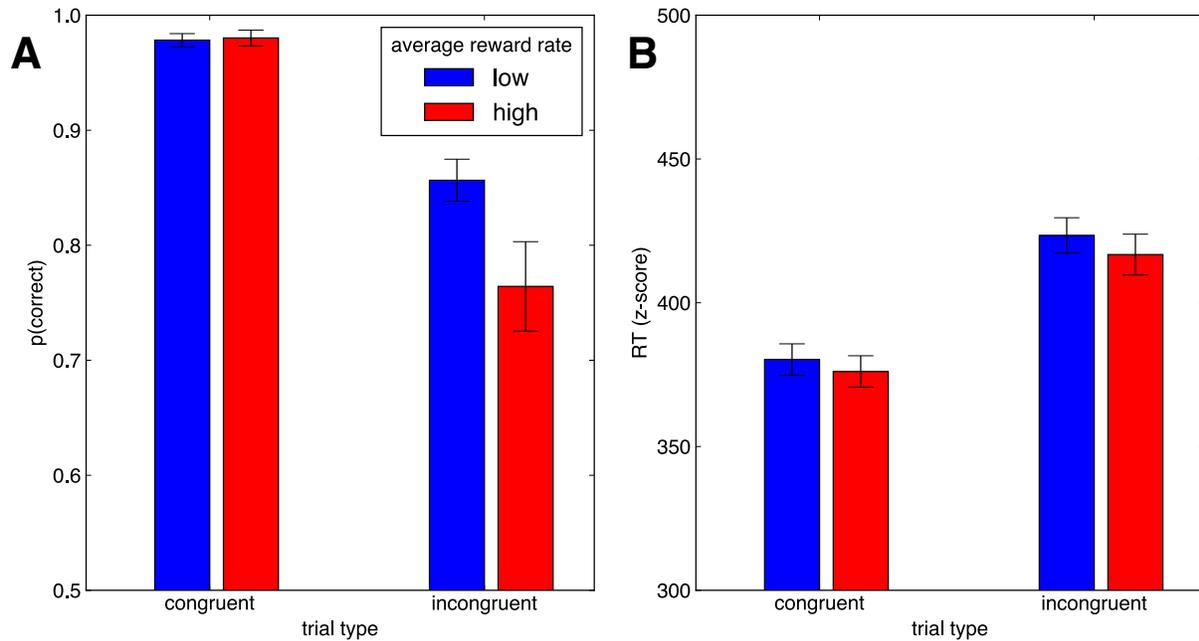


Figure 3. (A) When the average reward rate was high, participants made more errors on incongruent Simon trials, where they needed to override inappropriate, prepotent responses. (B) Participants did not make significantly faster responses, in either congruent or incongruent trials, when the average reward rate was high versus low. High and low average reward rate conditions were computed with a tertile split on participants' experienced average reward rate. Error bars indicate standard error of the mean.

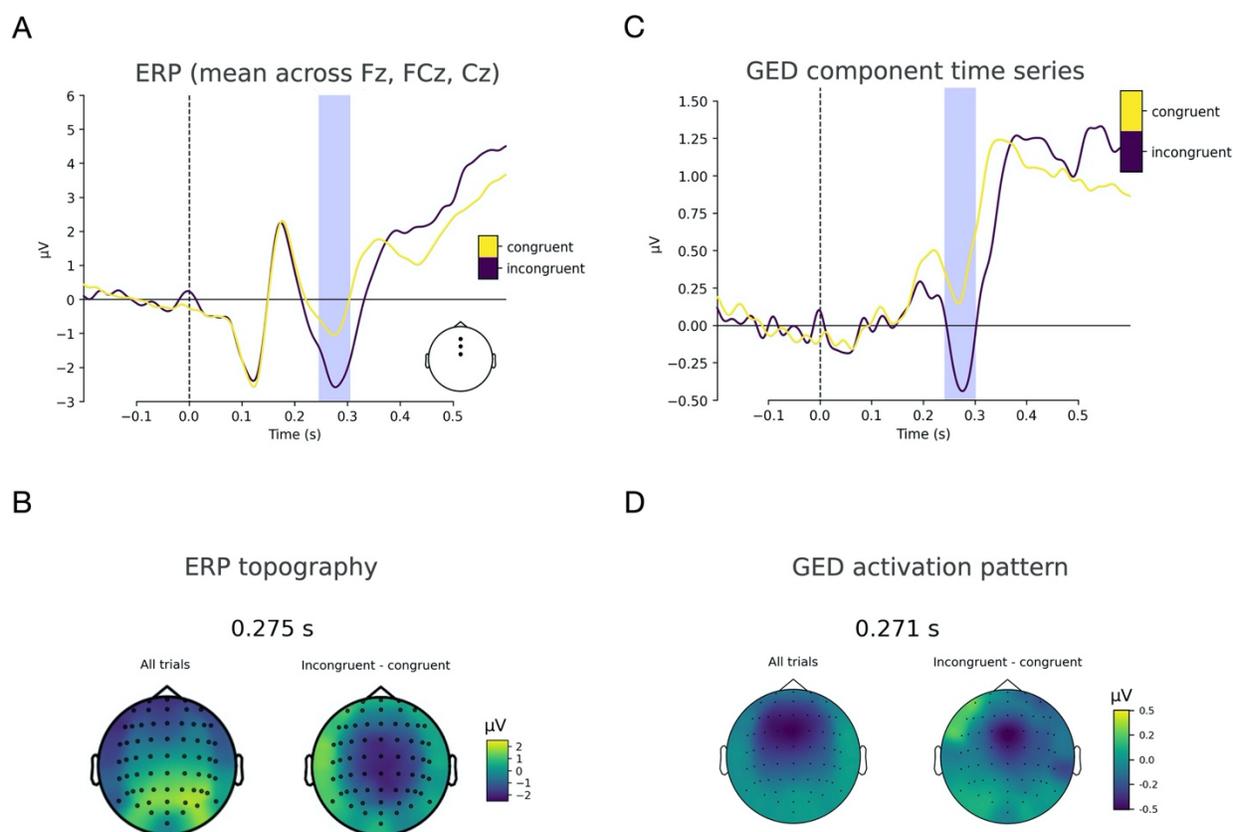


Figure 4. Event-related potentials (ERPs) and generalized eigendecomposition (GED) of EEG Simon-task stimulus-locked data. (A) The N2 component (0.25 to 0.31s) amplitude was more negative on incongruent (vs. congruent) Simon conflict task trials. (B) Spatial topographies at peak N2 amplitude (0.28 s) for grand average (all trials; left) and the difference between incongruent and congruent trials (right). (C). A GED component time series also differentiated between congruent and incongruent trials (0.24 to 0.30 s). (D). We identified a GED component (see C) with midfrontal topography because the N2 ERP component is usually observed in midfrontal channels (see A). Statistical analyses were performed on the mean activity in the blue shaded regions.

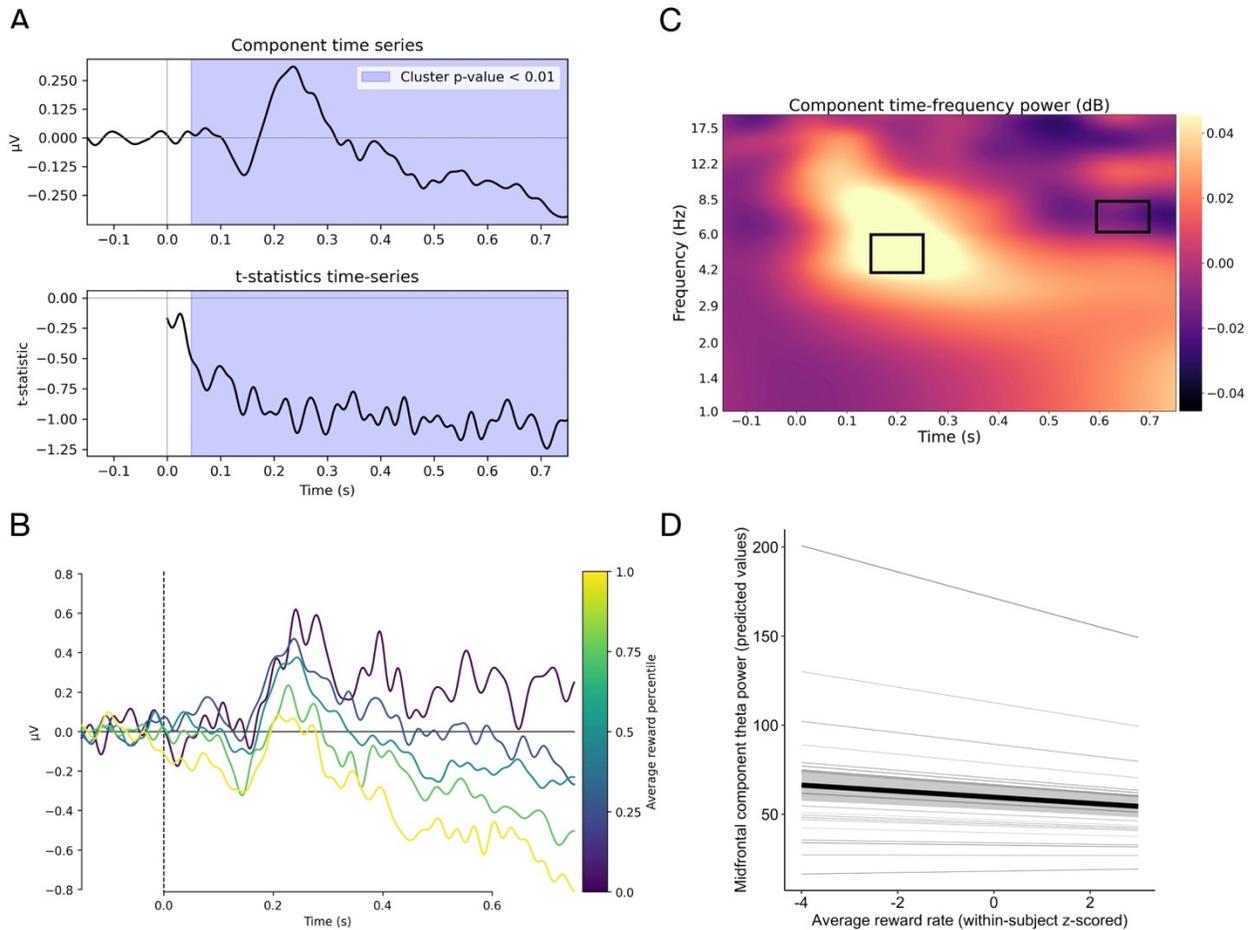


Figure 5. Generalized eigendecomposition of EEG cue-locked data identified a component associated with average reward rate processing. We identified a component with midfrontal topography (see inset in panel A) because theta (4-8 Hz) activity in midfrontal channels (FCz, Cz) has been associated with cognitive control and effort processes. (A) Component time series (top) and t-statistic time series indicating that from 0.04 to 0.75 s, higher average reward rate was associated with reduced component amplitude. Blue shaded regions indicate statistically significant clusters identified using non-parametric permutation tests (5000 permutations). (B) Component time-series activity for average reward rate (a continuous variable) split into four equally-sized bins. (C) Time-frequency representation of component time-series activity. The two highlighted windows showed changes in component theta power (window 1: 4-6 Hz at 0.15 to 0.25 s; window 2: 6-8 Hz at 0.60 to 0.70 s). Only theta power in the second window correlated negatively and significantly with average reward rate, such that higher average reward rate was associated with reduced component theta power. (D) Model-predicted component theta power (6-8 Hz at 0.60 to 0.70 s) as a function of average reward rate (± 1 SE). The gray lines are the average reward rate effects for individual participants.

Table 1. Mixed-effects logistic regression coefficients indicating the influence of the average reward rate and other trial-by-trial covariates upon accuracy in the Simon Task. Asterisks denote significance at the .05 level.

<i>Coefficient</i>	<i>Estimate (SE)</i>	<i>p-value</i>
congruent	-4.8502 (0.4712)	<0.0001*
incongruent	-1.4626 (0.2027)	<0.0001*
congruent:prev_errors	0.1202 (0.3364)	0.721
congruent:run_num	-0.1256 (0.2422)	0.604
congruent:trial_in_run	-0.0476 (0.1869)	0.799
congruent:key_rep	0.0462 (0.2078)	0.824
congruent:prev_type	1.0666 (0.2664)	<0.0001*
congruent:avg_reward	-0.0635 (0.2221)	0.775
congruent:reward	-0.0499 (0.1832)	0.785
congruent:resp_side	-0.2232 (0.2149)	0.299
incongruent:prev_errors	0.8538 (0.1868)	<0.0001*
incongruent:run_num	0.1667 (0.1046)	0.111
incongruent:trial_in_run	0.003 (0.0912)	0.973
incongruent:key_rep	-0.1661 (0.0857)	0.053
incongruent:prev_type	-0.7975 (0.1407)	<0.0001*
incongruent:avg_reward	0.2332 (0.1094)	0.033*
incongruent:reward	-0.0068 (0.0872)	0.938
incongruent:resp_side	-0.0393 (0.0844)	0.642

Table 2. Mixed-effects Regression coefficients indicating the influence of the average reward rate and a number of other trial-by-trial covariates upon RTs in the Simon Task. Asterisks denote significance at the .05 level.

<i>Coefficient</i>	<i>Estimate (SE)</i>	<i>p-value</i>
congruent	5.9015 (0.0157)	<0.0001*
incongruent	6.0783 (0.015)	<0.0001*
congruent:prev_errors	0.0655 (0.0143)	<0.0001*
congruent:run_num	-0.0063 (0.0038)	0.102
congruent:iti	-0.0166 (0.0028)	<0.0001*
congruent:trial_in_run	0.0015 (0.0032)	0.626
congruent:key_rep	0.0337 (0.0072)	<0.0001*
congruent:prev_type	0.0426 (0.0064)	<0.0001*
congruent:avg_reward	-0.005 (0.0034)	0.145
congruent:reward	-0.0039 (0.0028)	0.161
congruent:resp_side	-0.0225 (0.0077)	0.004*
incongruent:prev_errors	0.0271 (0.0245)	0.269
incongruent:run_num	-0.0099 (0.0061)	0.104
incongruent:iti	-0.0126 (0.0048)	0.009*
incongruent:trial_in_run	-0.0012 (0.0054)	0.827
incongruent:key_rep	0.0105 (0.0096)	0.274
incongruent:prev_type	-0.0395 (0.0107)	0.0*
incongruent:avg_reward	-0.0004 (0.0057)	0.946
incongruent:reward	0.0049 (0.005)	0.332
incongruent:resp_side	-0.0209 (0.011)	0.059