

Learning when effort matters: Neural dynamics underlying updating and adaptation to changes in
performance efficacy

Ivan Grahek^{1*}, Romy Frömer¹, Mahalia Prater Fahey¹, & Amitai Shenhav¹

1. *Department of Cognitive, Linguistic, & Psychological Sciences and Carney Institute for
Brain Science, Brown University, Providence, RI 02912, USA*

* Corresponding author. E-mail address: ivan_grahek@brown.edu

Conflicts of interest: The authors have no conflicts of interest to declare.

Abstract

To determine how much cognitive control to invest in a task, people need to consider whether exerting control *matters* for obtaining rewards. In particular, they need to account for the efficacy of their performance – the degree to which rewards are determined by performance or by independent factors. Yet it remains unclear how people learn about their *performance efficacy* in an environment. Here we combined computational modeling with measures of task performance and EEG, to provide a mechanistic account of how people (a) learn and update efficacy expectations in a changing environment, and (b) proactively adjust control allocation based on current efficacy expectations. Across two studies subjects performed an incentivized cognitive control task while their performance efficacy (the likelihood that rewards are performance-contingent or random) varied over time. We show that people update their efficacy beliefs based on prediction errors – leveraging similar neural and computational substrates as those that underpin reward learning – and adjust how much control they allocate according to these beliefs. Using computational modeling, we show that these control adjustments reflect changes in information processing, rather than the speed-accuracy tradeoff. These findings demonstrate the neurocomputational mechanism through which people learn how worthwhile their cognitive control is.

Keywords: Motivation; Cognitive control; Learning; Performance efficacy; Reward rate; EEG

1 Cognitive control is critical for achieving most goals, but it is effortful (Botvinick and
2 Cohen 2014; Shenhav et al. 2017). To decide how to invest control into a task (e.g., writing an
3 essay for a competition), a person must weigh these effort costs against the potential benefits of a
4 given type and amount of control (Manohar et al. 2015; Verguts et al. 2015; Kool and Botvinick
5 2018). One aspect of these benefits is the significance of the expected outcomes, both positive
6 (e.g., a monetary prize, social acclaim) and negative (e.g., missed revenue, social derision)
7 (Atkinson 1966; Leng et al., 2021). An equally important aspect of the expected benefits of
8 control is the extent to which control *matters* for achieving good outcomes and avoiding bad
9 ones (Frömer et al., 2021a; Shenhav et al., 2021). This can in turn be decomposed into the extent
10 to which higher levels of control translate into better performance (e.g., whether writing a good
11 essay will require substantial or only minimal control resources; *control efficacy*) and the extent
12 to which better performance translates into better outcomes (e.g., whether prizes are determined
13 by the strength of an essay or by arbitrary or even biased metrics unrelated to essay-writing
14 performance; *performance efficacy*). Whereas studies have increasingly characterized the ways
15 in which control allocation is influenced by expected outcomes (e.g., Parro et al., 2018; Leng et
16 al., 2021) and the expected efficacy of control (e.g., as a function of task difficulty; Krebs et al.,
17 2012; Vassena et al., 2014; Chiu and Egnér 2019), much less is known about how people
18 estimate and adjust to the perceived efficacy of their performance in a given environment.

19 We recently showed that when participants are explicitly instructed about how
20 efficacious their performance will be on an upcoming trial, they exhibit behavioral and neural
21 responses consistent with increased control (Frömer et al., 2021a). We had participants perform a
22 standard cognitive control task (Stroop) for potential monetary rewards, and we varied whether
23 obtaining those rewards was contingent on performing well on the task (high performance

24 efficacy) or whether those rewards were determined at random (low efficacy). We showed that
25 people allocate more control when they expect to have high compared to low efficacy, reflected
26 in higher amplitudes of an EEG index of proactive control (the contingent negative variation
27 [CNV]) and in improved behavioral performance. These results demonstrate that participants
28 leverage expectations about the extent to which their performance matters when deciding how
29 much cognitive effort to invest in a task. However, in this work, participants were explicitly cued
30 with the level of performance efficacy to expect on a given trial, and those predictive cues
31 retained the same meaning across the session. Thus, how it is that people learn these efficacy
32 expectations in environments where contingencies are not instructed, and how they dynamically
33 update their expectations as contingencies change, remains unanswered.

34 Outside of the domain of cognitive control, a relevant line of work has examined how
35 people learn about the factors that determine future outcomes when selecting between potential
36 courses of action. In particular, work in this area has shown that people are able to learn about
37 and update their expectations of the likelihood that a given action will generate a given outcome
38 (action-outcome contingency; Dickinson and Balleine 1995; Moscarello and Hartley 2017; Ly et
39 al. 2019). People preferentially, and more vigorously, select actions that reliably lead to desired
40 outcomes (i.e., the more contingent those outcomes are on the action in question; Liljeholm et al.
41 2011; Manohar et al. 2017), and work in both animals (Balleine & O’Doherty, 2010) and humans
42 (Norton & Liljeholm, 2020; Dorfman et al., 2021; Ligneul et al., 2022; Morris et al., 2022) has
43 helped to characterize the neural systems that support this process of learning and action
44 selection. However, given the focus on discrete actions and their immediate relationship with
45 outcomes, research into these action-outcome contingencies is unable to capture key aspects that
46 are unique to selection of control states. Most notably, cognitive control signals (e.g., attention to

47 one or more features of the environment) are multidimensional, not immediately observable by
48 either the participant or experimenter, and their relationship with potential outcomes is
49 intermediated by the many-to-many relationship between control states and task performance
50 (Ritz et al., 2022). While there has been research into how people learn to adjust cognitive
51 control signals based on changes in their task environment, here again work has focused on how
52 people adapt to changes in outcomes (e.g., Otto and Daw 2019; Bustamante et al. 2021) and
53 changes in the relationship between control and performance (with increasing task difficulty;
54 Bugg et al. 2011; Nigbur et al. 2015; Bejjani et al. 2018; Jiang et al. 2020). The mechanisms by
55 which people learn about the relationship between performance and outcomes (performance
56 efficacy), and how they adjust their control allocation accordingly, remain largely unexplored.

57 Here, we seek to fill this gap by studying the mechanisms through which expectations of
58 performance efficacy are formed, updated, and used to guide control allocation. To do so, we
59 extend our previous approach (Frömer et al., 2021a) – which studied how behavioral and neural
60 correlates of control allocation vary when performance efficacy is explicitly cued – to examine
61 how participants learn and adapt their control under conditions where efficacy was un-cued
62 (having to instead be learned from feedback) and gradually varied across a wide range of
63 potential efficacy values over the course of the session. We use computational reinforcement
64 learning models to show that expected efficacy can be learned from feedback through iterative
65 updating based on weighted prediction errors (Sutton and Barto 2018), and model-based single-
66 trial EEG analyses to show that these efficacy prediction errors modulate a canonical neural
67 marker of reward-based learning and behavioral adjustment (Fischer and Ullsperger 2013). We
68 further provide evidence that efficacy estimates learned in this way are used to guide the
69 allocation of control. In our EEG study and a second behavioral study, participants tended to

70 perform better when efficacy was higher. We also provide evidence that a neural marker of
71 control allocation (Schevernels et al. 2014a) tends to increase with increasing model-based
72 efficacy estimates. Using a drift diffusion model (Ratcliff & McKoon, 2008; Wiecki, Sofer, &
73 Frank, 2013), in Study 2 we further show that the performance improvements related to
74 increased performance-efficacy reflect facilitation of task-related information processing
75 (reflected in increased drift rates), rather than changes in the speed-accuracy tradeoff (i.e.,
76 thresholds). Taken together, these results show that efficacy estimates can be learned and
77 updated based on feedback, leveraging general cognitive and neural mechanisms of predictive
78 inference.

79

80

Materials and Methods

Study 1

Participants

83 We recruited forty-one participants with normal or corrected-to-normal vision from the
84 Brown University subject pool. One participant was excluded due to technical issues. The final
85 data set included 40 participants (24 females, 16 males; median age = 19). Participants gave
86 informed consent and were compensated with course credits or a fixed payoff of \$20. In
87 addition, they received up to \$5 bonus that depended on their task performance (\$3.25 on
88 average). The research protocol was approved by Brown University's Institutional Review
89 Board.

Experimental design

91 In the main task, taking approximately 45 minutes, participants performed 288 Stroop
92 trials (Figure 1A). Each trial started with the presentation of a cue (grey circle) that remained on

93 the screen throughout the trial. After a period of 1500 ms, a Stroop stimulus was superimposed
94 until a response was made or 1000 ms elapsed, at which time it was sequentially replaced with
95 two types of feedback presented for 1000 ms each. Each trial onset was preceded by a fixation
96 cross (randomly jittered between 1000 and 1500ms). Participants responded to the ink color of
97 the Stroop stimulus (red, green, blue, or yellow) by pressing one of four keyboard keys (D, F, J,
98 and K). Stimuli were either color words same as the ink color (congruent, $n = 108$) or different
99 (incongruent, $n = 108$), or a string of letters “XXXXX” (neutral, $n = 72$). Feedback informed
100 them whether they obtained a reward (reward, “\$50c” or no reward, “\$0c”) and whether the
101 reward they received depended on their performance (performance-based feedback, a button
102 graphic), or not (random feedback, a dice graphic). In order to earn rewards in the performance-
103 based case, participants had to be both accurate and respond within an individually calibrated
104 response deadline (see details below). The order of the two types of feedback was pseudo-
105 randomized with half of the trials showing reward feedback first and the other half efficacy
106 feedback. Every 2-4 trials the feedback was followed by a probe of efficacy (“How much do you
107 think your rewards currently depend on your performance?”) or reward rate (“How often do you
108 think you are currently being rewarded?”) to which participants responded on a visual analog
109 scale ranging from 0 to 100. The number and timing of the probes was randomized per subject
110 resulting in a median of 45 efficacy probes ($SD=3.38$) and 47 reward probes ($SD=2.81$).

111 Efficacy (performance-based or random rewards) on each trial was sampled from a
112 binomial distribution with probabilities ranging between 0.1 and 0.9 that drifted over the course
113 of the experiment and were predetermined (Figure 1B). In order to ensure that the performance-
114 based and random trials did not differ in reward rate, reward feedback for the random trials was
115 sampled from the moving window of the reward feedback of the previous 10 performance-based

116 trials. At the beginning of the experiment a window with 8 rewards and 2 no rewards was used to
117 reflect the pre-calibrated reward rate (details below), and this moving window was then updated
118 after every trial. Thus, reward rate was not experimentally manipulated in the experiment and
119 remained constant. We confirmed that reward rate was matched across performance-based and
120 random trials by comparing reward probability between these trial types ($b = 0.01$; 95% CrI [-
121 0.07, 0.10]; $p_{b > 0} = 0.38$).

122 Prior to the main task, participants performed several practice phases of the Stroop task
123 (approximately 15 minutes). First, they practiced the mappings between colors and keyboard
124 keys (80 trials). Then they completed a short practice of the Stroop task with written feedback
125 (“correct” or “incorrect”) on each trial (30 trials). Participants then completed 100 more of such
126 trials during which we individually calibrated the reaction time deadline such that participants
127 yielded approximately 80% reward rate. The reaction time calibration started from a fixed
128 deadline of 750ms and increased or decreased this threshold in order to ensure that participants
129 earn rewards on 80% of trials (i.e., that they are both accurate and below the deadline). The
130 deadline obtained in this way ($M = 796\text{ms}$; $SD = 73\text{ms}$) was used in the main experiment and
131 was not further adjusted. In the final practice phase participants were introduced to the two types
132 of feedback which they would see in the main experiment (30 trials).

133 The experimental task was implemented in Psychophysics Toolbox (Brainard 1997; Pelli
134 1997; Kleiner et al. 2007) for Matlab (MathWorks Inc.) and presented on a 23 inch screen with a
135 1920 x 1080 resolution. All of the stimuli were presented centrally while the participants were
136 seated 80 cm away from the screen.

137 ***Psychophysiological recording and preprocessing***

138 EEG data were recorded at a sampling rate of 500 Hz from 64 Ag/AgCl electrodes
139 mounted in an electrode cap (ECI Inc.), referenced against Cz, using Brain Vision Recorder
140 (Brain Products, München, Germany). Vertical and horizontal ocular activity was recorded from
141 below both eyes (IO1, IO2) and the outer canthi (LO1, LO2), respectively. Impedances were
142 kept below 10 k Ω . Offline, data were processed using custom made Matlab scripts (Frömer et al.
143 2018) employing EEGLab functions (Delorme and Makeig 2004). Data were re-referenced to
144 average reference, ocular artifacts were corrected using brain electric source analyses (Ille et al.
145 2002) based on separately recorded prototypical eye movements. The cleaned continuous EEG
146 was then low pass filtered at 40 Hz and segmented into epochs around cue onset (-200 to 1500
147 ms), stimulus onset, and both efficacy and reward feedback (-200 to 800 ms). Baselines were
148 corrected to the average of each 200 ms pre-stimulus interval. Segments containing artifacts,
149 values exceeding $\pm 150 \mu\text{V}$ or gradients larger than $50 \mu\text{V}$, were excluded from further analyses.

150 Single trial ERPs were then exported for further analyses in R (R Core Team 2017). The
151 late CNV was quantified between 1000 and 1500 ms post neutral cue onset (Schevernels et al.
152 2014b; Frömer et al. 2016; Frömer et al. 2021a) as the average activity over 9 fronto-central
153 electrodes (Fz, F1, F2, FCz, FC1, FC2, Cz, C1, and C2). The P3b was quantified between 350
154 and 500 ms (Fischer and Ullsperger 2013) for both reward and efficacy feedback and calculated
155 as the averaged activity over 9 centro-parietal electrodes (Pz, P1, P2, POz, PO1, PO2, CPz, CP1,
156 and CP2).

157 ***Learning models and statistical analyses***

158 **Learning models.** Participants provided their subjective estimates of efficacy and reward every
159 4-8 trials (a total of 45 estimates), and we sought to fit a learning model to these estimates to be

160 able to predict trial-by-trial adjustments in performance and neural markers of learning and
161 cognitive control allocation. In order to obtain trial-by-trial estimates of efficacy and reward rate,
162 we fitted two temporal difference learning models (Gläscher et al., 2010; Sutton & Barto, 2018)
163 to the continuous subjective estimates of efficacy and reward rate (Rutledge et al., 2014; Eldar et
164 al., 2016; Nagase et al., 2018). The first model (“1 learning rate efficacy model”) assumed that
165 the estimate of efficacy for the next trial (E_{t+1}) depended on the current efficacy estimate (E_t)
166 and the prediction error (δ_t) weighted by a constant learning rate (α):

$$167 \quad E_{t+1} = E_t + \alpha * \delta_t$$

168 Where $0 \leq \alpha \leq 1$, and the prediction error is calculated as the difference between the contingency
169 feedback on the current trial (e_t) and the efficacy estimate on that trial: $\delta_t = e_t - E_t$. The model
170 started from an initial value (free parameter) and updated the model-based efficacy estimate
171 based on the binary efficacy feedback on each trial. For example, assuming a learning rate of 0.5
172 and the initial value of 0.5, the model would update the initial estimate following efficacy
173 feedback signaling “performance-based” ($e_t = 1$) to 0.75. If on the next trial contingency
174 feedback was “random” ($e_{t+1} = 0$) the model-based efficacy estimate would drop to 0.6. The
175 model was fitted separately to the subjective estimates of efficacy with only the learning rate as a
176 free parameter. The second model (“2 learning rates efficacy model”) was the same as the first
177 model, but it included two learning rates: one learning rate for learning from the “performance-
178 based” feedback, and another for learning from the “random” feedback. Finally, as a baseline, we
179 also included the “intercept model” which did not update the efficacy estimate throughout the
180 experiment, but just assumed that the estimate took one constant value. Importantly, the same
181 models were fitted to obtain the model-based estimates of reward on each trial (“1 learning rate
182 reward model” and the “2 learning rate reward model”). These models were fitted using trial-by-

183 trial reward feedback and the subjective estimates of reward. The models were fit hierarchically
184 to the data using maximum likelihood estimation (using mfit (<https://github.com/sjgershm/mfit>)).
185 To calculate the likelihood of each data point, model-based estimates (0-1 range) were compared
186 to the subjective efficacy estimates (range normalized to 0-1 range for each participant).
187 Likelihood was evaluated on trials which included a subjective estimate, as the likelihood that
188 the difference between the model-based and the empirical estimate comes from a Gaussian
189 distribution centered on 0 with a variance which was fitted as a free parameter each subject. This
190 variance parameter served as the noise in the estimates. Likelihoods were log transformed,
191 summed, and then maximized using the fmincon function in MATLAB.

192 We performed a parameter recovery study to show that the most complex model (the 2
193 learning rates model) can be successfully recovered. We simulated a dataset with the same
194 number of trials and subjective efficacy or reward probes as in the actual experiment. We used
195 the efficacy drifts presented to the actual subjects (half of the simulated subjects saw one drift,
196 and half its inverse), and we used the reward feedback sequences of two actual subjects from our
197 experiment. We simulated 200 agents which learned both efficacy and reward with the noise
198 parameter fixed to 0.2, intercept fixed to 0.5, and the positive and negative learning rates
199 sampled from a uniform distribution ranging from 0.001 to 0.5. These parameters were matched
200 based on the range of values obtained from the empirical fits to our data. As shown in Figure S1
201 we were able to very reliably recover the simulated parameters for both efficacy and reward rate
202 learning.

203 **Statistical analyses.** The efficacy and reward rate estimates obtained through fitting the learning
204 model were then used to analyze the behavioral and EEG data. To this end, we fitted Bayesian
205 multilevel regressions to predict subjective estimates of efficacy and reward rates, reaction times,

206 accuracy, as well as the CNV and P3b amplitudes. Subjective estimates of efficacy and reward
207 rate were regressed onto efficacy or reward feedback. Reaction times and accuracies were
208 regressed onto trial-by-trial model-based estimates of efficacy and reward rate, as well as trial-
209 by-trial CNV amplitude, while controlling for congruency. The P3b amplitudes were analyzed in
210 two ways: with trial-by-trial model-based estimates of efficacy and reward rate and current
211 feedback as predictors, and with model-based prediction errors and learning rates for each
212 feedback type. CNV amplitudes were regressed onto trial-by-trial model-based estimates of
213 efficacy. All of the fitted models controlled for the influence of the reward rate estimates.
214 Parallel analyses were done to predict the P3b in response to reward feedback, while controlling
215 for the efficacy estimates.

216 The regression models were fitted in R with the *brms* package (Bürkner 2016) which
217 relies on the probabilistic programming language *Stan* (Carpenter et al. 2016) to implement
218 Markov Chain Monte Carlo (MCMC) algorithms and estimate posterior distributions of model
219 parameters. The analyses were done based on the recommendations for Bayesian multilevel
220 modeling using *brms* (Bürkner 2016; Bürkner 2017; Nalborczyk and Bürkner 2019). The fitted
221 models included constant and varying effects (also known as fixed and random) with weakly
222 informative priors (except for the behavioral and CNV analyses, see below for details) for the
223 intercept and the slopes of fixed effects and the likelihoods appropriate for the modeled data (Ex-
224 Gaussian for reaction times, Bernoulli for accuracy, and Gaussian for the subjective estimates
225 and the EEG amplitudes). The fitted models included all of the fixed effects as varying effects.
226 All of the continuous predictors in the model were centered and the categorical predictors were
227 contrast coded. Four MCMC simulations (“chains”; 20,000 iterations; 19,000 warmup) were run
228 to estimate the parameters of each of the fitted models. The convergence of the models was

229 confirmed by examining trace plots, autocorrelation, and variance between chains (Gelman and
230 Rubin 1992). After convergence was confirmed, we analyzed the posterior distributions of the
231 parameters of interest and summarized them by reporting the means of the distribution for the
232 given parameter (b) and the 95% credible intervals (95% CrI) of the posterior distributions of
233 that model. We report the proportion of the posterior samples on the relevant side of 0 (e.g., $p_{b < 0}$
234 = 0.9), which represents the probability that the estimate is below or above 0. We also report
235 Bayes factors (BF) calculated using the Savage-Dickey method (Wagenmakers et al. 2010). We
236 report the BFs in support of the alternative hypothesis against the null (BF_{10}), except for the
237 analyses of accurate RT, accuracies, and CNV amplitude in which we have informative priors
238 based on our previous study (Frömer et al. 2021a), and in which case we support the evidence in
239 favor of the null (BF_{01}).

240 To compare the positive and negative learning rates we fitted a model in which the
241 learning rates were predicted by the learning rate type (Kruschke 2013). In this model we used
242 Gaussian distributions (mean, standard deviation) as priors (intercept: (0.5,0.5); slopes: (0,0.5)).

243 We fitted two separate models to predict the subjective estimates of efficacy and reward
244 rate based on previous feedbacks. At each timepoint the estimates were predicted by the current,
245 and previous 4 feedbacks. The feedback on each of the trials (performance-based vs. random or
246 reward vs. no reward) was entered as a constant effect and the models also included the intercept
247 as a varying effect. As the subjective estimates could vary between 0 and 1, we used Gaussian
248 distributions (intercept: (0.5,0.2); slopes: (0,0.2)) as priors.

249 For predicting the P3b amplitude in response to the onset of the efficacy feedback, we
250 fitted two models. First, we fitted a model which included the model-based estimate of efficacy
251 (prior to observing the current feedback), the actual feedback, and the interaction between the

252 expected efficacy and the observed efficacy feedback. Additionally, we controlled for the reward
253 rate estimate. Second, we fitted separate models which included the model-based prediction
254 errors, the influence of the between-subject learning rates, and their interaction with the
255 prediction errors, while controlling for the estimate of the reward rate. In this analysis the
256 learning rates (one for feedback type for each subject) were mean-centered within subjects and
257 thus any effect of the learning rates is driven by the difference between the random and the
258 performance-based learning rate. For these models we selected wide Gaussian priors (intercept:
259 (5,3); slopes (0,3)). The same logic in building models was applied for the analyses of the reward
260 feedback. In these analyses we focused on the reward feedback processing and how it interacted
261 with the model-based estimates of reward rates, while controlling for the model-based estimates
262 of efficacy. We analyzed only the trials with correct responses for both the efficacy and the
263 reward feedback analyses.

264 To test the influence of efficacy on the late CNV, we fitted a model which predicted the
265 CNV based on the model-based efficacy estimates, while controlling for the effect of the reward
266 rate estimates. Drawing on the results of our previous study (Frömer et al. 2021a), this model
267 included Gaussian priors for the intercept (-0.16, 1.59) and the efficacy (-0.30, 0.73) and reward
268 (0,0.73) slopes.

269 For predicting reaction times and accuracy we fitted models which included congruency
270 (Facilitation: difference between neutral and congruent trials; Interference: difference between
271 incongruent and neutral trials) and the model-based efficacy estimates, while controlling for the
272 reward rate estimates. We used Gaussian distributions as informative priors based on our
273 previous study (Frömer et al. 2021a), for both the reaction times (intercept: (624, 58.69);
274 facilitation (15.54, 21.93), interference (61.12, 37.49); efficacy (-10.26, 19.51); reward (0,

275 19.51)) and accuracy¹ analyses (intercept: (2.11, 0.81); facilitation (-0.45, 0.64), interference (-
276 0.53, 0.81); efficacy (0.09, 0.32); reward (0, 0.32)).

277 To investigate how the late CNV influences the behavior, we fitted two models in which
278 we predicted the reaction times and accuracy based on the CNV amplitude. The prior
279 distributions for these models were weakly informative Gaussian distributions for predicting
280 both the reaction times (intercept: (650, 200); slope: (0, 50)) and accuracy (intercept: (0.7, 0.2);
281 slope: (0, 0.2)).

282 To visualize the topographies of the relevant ERP effects, we fitted the relevant models to
283 all 64 channels and then plotted the posterior estimates of the effects of interest at each electrode
284 (cf. Frömer et al., 2021b).

285 **Study 2**

286 *Participants*

287 We recruited eighty-seven participants residing in the United States from Prolific – an
288 online platform for data collection. Participants had normal or corrected-to-normal vision and
289 gave informed consents. They were compensated with a fixed payoff of \$8 per hour (median
290 completion time of 74 minutes) plus a monetary bonus based on points earned during the task
291 (\$1 on average). The research protocol was approved by Brown University’s Institutional
292 Review Board.

293 We a priori excluded participants who did not pass attention checks (N=8) or who took
294 substantially longer than the average participant to complete the study (N=2 participants who
295 took over 130 minutes), suggesting that they did not sustain attention to the experiment over that

¹ Note that the prior distributions are set in log-odds.

296 time. We fit our learning models to data from the remaining 77 subjects, and then excluded
297 participants whose performance suggested inattention to the overall task (based on accuracies
298 less than 70% across all trials – including the trials in which performance efficacy was low, N=6)
299 or inattention to the task feedbacks and efficacy probes (based on low learning rates (N=19), and
300 one subject with no variance in responses to reward probes). To identify participants with
301 exceedingly low learning rates, we submitted all positive and negative efficacy learning rates to
302 unsupervised Gaussian mixture models (as implemented in the Mclust package; (Scrucca et al.,
303 2016) to determine the best fitting number and shape of clusters (model comparison via BIC).
304 This procedure identified four clusters of subjects with different overall learning rates (Figure
305 S2B-C) and we excluded subjects from the first cluster as they all had very low learning rates
306 relative to the other participants (both learning rates < 0.03). The subjects excluded based on low
307 learning rates were most likely not paying attention to efficacy feedback, or were always giving
308 very similar responses to the efficacy probes (Figure S3). The final sample included 51
309 participants (31 females, 20 males; median age = 29).

310 *Experimental design*

311 In order to better understand the computational mechanisms that lead to improved
312 behavioral performance in high efficacy states (Study 1), we wanted to fit a Drift Diffusion
313 model (DDM; Ratcliff & McKoon, 2008) to our behavioral data. If people allocate more
314 attention when they think they have high performance efficacy, this should be observed as an
315 increase in the drift rate (speed of evidence accumulation). However, Study 1 included a tight
316 respond deadline for earning a reward, making it more challenging to fit the DDM. To avoid this
317 issue, in Study 2 we adjusted the task to a free response paradigm which allowed us to

318 investigate drift rate and threshold adjustments (cf. Leng et al., 2021). We used this design to test
319 the hypothesis that higher efficacy estimates should predict increased drift rates.

320 Instead of single trials, participants now completed 288 intervals during which they could
321 respond to as many trials (congruent and incongruent, removing the neutral condition) as they
322 wished within a fixed time window (randomly selected between 2000, 3000, or 4000ms). Apart
323 from this, the structure of the task remained the same: participants saw a fixation cross (1000,
324 1500, or 2000ms), then completed as many trials as they wished during a fixed interval, followed
325 by the feedback (1500 ms) on how many points they earned and whether this was based on their
326 performance or awarded to them based on random chance. Note that participants now received
327 continuous reward feedback (10 points per correct response instead of the binary reward-no
328 reward in Study 1). For example, if participants completed 4 trials correctly and 2 trials
329 incorrectly within a performance-based interval they would receive 40 points and see feedback
330 informing them that the points were based on their performance. To determine the number of
331 points on random intervals, the same yoking procedure as in Study 1 was employed ensuring that
332 the amount of reward was matched between performance-based and random intervals (reward
333 amounts on random intervals were sampled from the moving average window of the past 10
334 performance-based intervals). We confirmed that the yoking procedure was successful by
335 comparing the reward amounts on the two interval types ($b = 0.00$; 95% CrI [-0.00, 0.02]; $p_{b>0} =$
336 0.16). As in Study 1, participants were probed every 2-4 intervals to estimate either how much
337 they thought their rewards depended on their performance, or how often they were rewarded. We
338 again implemented an efficacy drift (modified, but comparable to the drift in Study 1; Figure
339 S2A), now across the 288 intervals of the task.

340 We gamified the task in order to make it more appealing for the participants. Instead of
341 the Stroop task, we used a picture-word interference task in which four grey-scaled images of
342 fruit (apple, pear, lemon, and peach) were overlaid by those fruit words. Participants responded
343 to the image, while ignoring the word, by pressing one of the 4 corresponding keys. They first
344 practiced this task, and then were introduced with a cover story telling them that they are in the
345 garden and need to water the fruit in little patches by pressing the correct keys. They were
346 instructed that they will be moving through a garden and that in some patches watering will
347 directly translate into how many points they will be earning, while in the others that will not be
348 the case (the efficacy drift). The experiment was implemented in Psiturk (Gureckis et al., 2016)
349 and the participants performed the task on their own computers and were required to have a
350 keyboard.

351 *Learning models, Drift Diffusion model, and statistical analyses*

352 **Learning models.** We fitted the same set of learning models as in Study 1, performed model
353 comparison, and got the interval-by-interval model-based estimates of performance efficacy and
354 reward. Note that in this version of the task participants earned points in each interval, unlike the
355 binary rewards (reward vs. no reward) in Study 1. This meant that the reward learning model
356 learned reward magnitudes rather than reward probabilities. However, model fitting and the
357 further analyses were the same as in Study 1.

358 **Statistical analyses.** We fitted the same Bayesian multilevel models as in Study 1 to predict the
359 influence of previous efficacy feedbacks on the subjective efficacy estimates, as well as the
360 influence of the model-based efficacy estimates on reaction times and accuracies. For the
361 analyses of the subjective efficacy estimates we used the same priors as in Study 1. For the
362 analyses of the reaction times and accuracies, we used the posterior distributions obtained in

363 Study 1 as the informative priors for the congruency, efficacy, and reward effects. The reaction
364 times and accuracy models also controlled for the effects of the average congruency level in the
365 interval and the interval length.

366 **Drift Diffusion model.** This model decomposes participant's behavior into drift rate (the speed
367 of evidence accumulation) and response threshold (the level of caution), allowing us to
368 investigate which of these two components is affected by the efficacy estimates. We fitted the
369 model using Bayesian hierarchical estimation as implemented in the HDDM package (Wiecki,
370 Sofer, & Frank, 2013). The fitted model included the main effects of efficacy and reward rate
371 estimates onto both drift rate and threshold, and included the effect of congruency on the drift
372 rate. The responses were coded as correct or incorrect, and trials with reaction times below
373 250ms were excluded. All of the effects were allowed to vary across subjects, and we ran five
374 MCMC chains (12,000 iterations; 10,000 warmup). We confirmed convergence by examining
375 trace plots and variance between chains.

376

377

Results

378 Study 1

379 To investigate how efficacy estimates are learned, and how they affect control allocation,
380 in Study 1 we recorded EEG while 40 participants performed a modified version of the Stroop
381 task (Figure 1A). Across trials we varied whether reward outcomes (\$0.50 vs. \$0.00) were
382 determined by a participant's performance on a given trial (responding accurately and below a
383 pre-determined response time criterion; *performance-based* trials) or whether those outcomes
384 were determined independent of performance (based on a weighted coin flip; *random* trials).
385 Over the course of the session, we gradually varied the likelihood that a given trial would be

386 performance-based or random such that, at some points in the experiment, most trials were
387 performance-based (high efficacy level), and at other points most trials had random outcomes
388 (low efficacy level) (Figure 1B). Importantly, unlike in our previous study (Frömer et al., 2021a),
389 participants were not told whether a given trial would be performance-based (maximal efficacy)
390 or random (minimal efficacy), but instead had to estimate their current efficacy level based on
391 recent trial feedback, which indicated both the reward outcome (\$0.50 vs. \$0.00) and how that
392 outcome was determined (performance-based [button symbol] vs. random [dice symbol]). We
393 held reward rate constant across both feedback types by yoking reward rate on random trials to
394 the reward rate on performance-based trials, and counter-balanced the time-course of the gradual
395 change in efficacy (see Methods for details). To capture changes in efficacy expectations over
396 the course of the session, we probed participants every 4-8 trials (averaging 44.3 probes per
397 participant) to report their current estimates of efficacy. These efficacy probes were interleaved
398 with probes asking participants to estimate the current reward rate, serving as foils and for
399 control analyses.

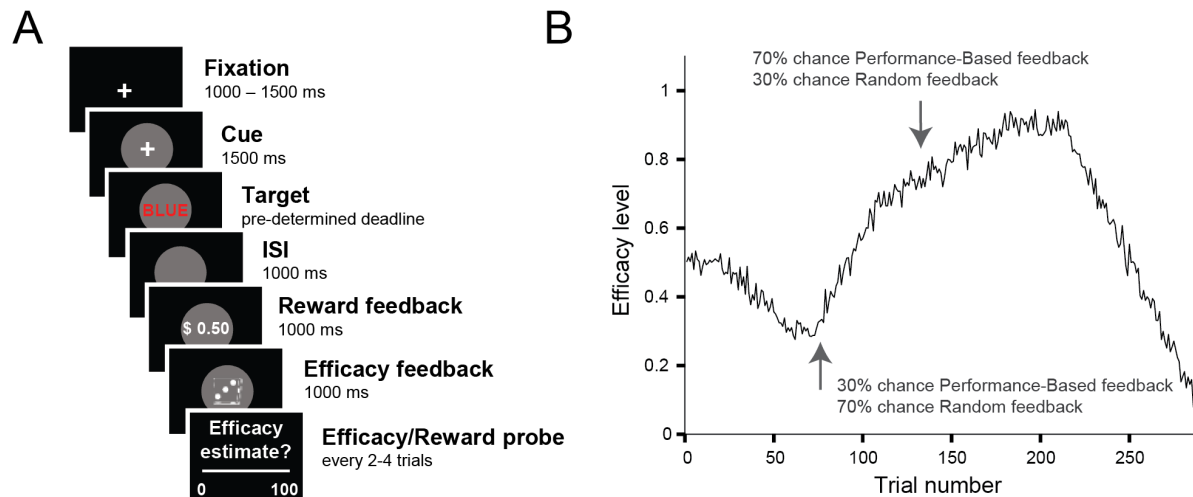


Figure 1. Manipulating expected efficacy and assessing learning in Study 1. A. Trial Schematic. On each trial participants saw a cue (gray circle), predicting the onset of a Stroop stimulus (target), and were then sequentially presented with reward and efficacy feedback. On half of the trials, efficacy feedback was presented first, and on the other half reward feedback was presented first. Every 2-4 trials participants were subsequently asked to estimate their current efficacy level (“How much do you think your rewards currently depend on your performance?”) or reward rate (“How often do you think you are currently being rewarded?”). **B. Efficacy manipulation.** We let the probability of performance-based vs. random feedback continuously drift over the course of the experiment (inversed for one half of the sample). Arrows mark time points with low and high efficacy, respectively. When efficacy was low, rewards were more likely to be random, whereas when efficacy was high, rewards were more likely to be performance-based.

400

401 *Participants dynamically update efficacy expectations based on feedback*

402 To determine whether and how participants learned their current efficacy levels, we first
403 analyzed the influence of previous efficacy feedback (whether outcomes on a given previous trial
404 had been performance-based or random) on one’s current subjective estimates of efficacy. We
405 found very strong evidence that participants adjusted their subjective efficacy upward or
406 downward depending on whether the most recent trial was Performance-Based or Random ($b =$
407 0.14 ; 95% CrI [$0.12, 0.16$]; $p_{b < 0} = 0$; $BF_{10} > 100$), and that this remained true (but to diminishing
408 degrees) up to five trials back (all $p_{b < 0} < 0.01$; Figure 2A; Table S1). This effect of feedback was
409 present only for the efficacy feedback, while reward feedback did not predict the subjective

410 estimates of efficacy (Figure 2A). These results suggest that participants were dynamically
 411 updating their efficacy estimates based on efficacy feedback.

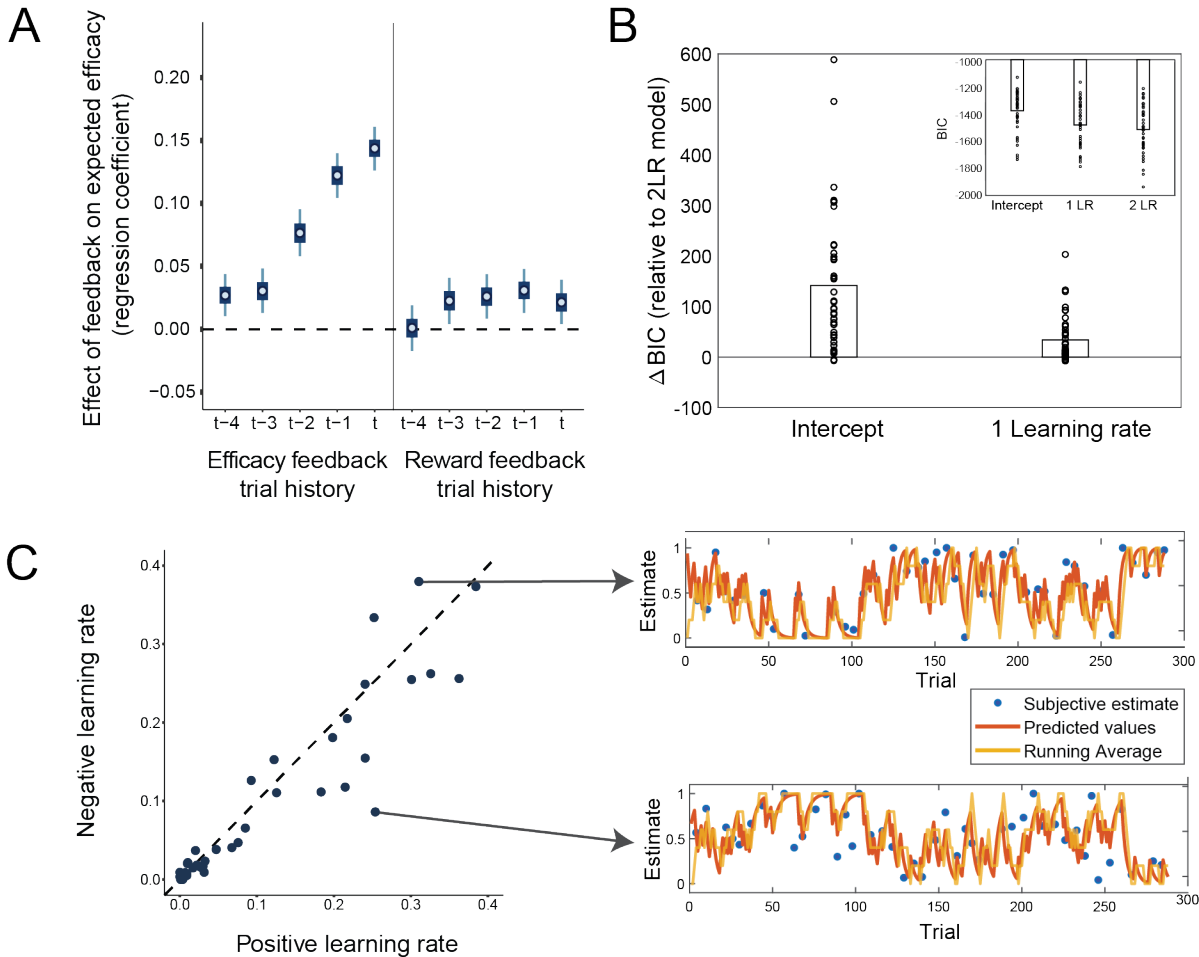


Figure 2. Efficacy learning is captured by a reinforcement learning model with separate learning rates for performance-based and random feedback in Study 1. A. Efficacy estimates track recent efficacy feedback. Regression weights for the influence of the current (t) and previous contingent vs. random feedback, as well as reward feedback, on the subjective efficacy estimate. Error bars represent 50% and 95% highest density intervals. **B. A separate learning rate model captures efficacy learning best.** BICs of fitted intercept only and one learning rate model relative to two learning rate models are plotted for each participant and favor the two learning rate model. **C. Learning rate biases vary between participants.** Positive and negative learning rate estimates are plotted for all participants. Points below the diagonal indicate higher learning rates for performance-based compared to random feedback, and points above the diagonal indicate the opposite. **D. Example learning trajectories.** Subjective and model-based efficacy estimates, and a running average of the previous 5 efficacy feedbacks, for a participant with a higher learning rate for the contingent compared to random feedback (upper) and a participant with the reverse bias (lower).

412 This pattern of learning was accounted for by a standard reinforcement learning (RL)
413 algorithm, the delta-rule, according to which efficacy estimates are updated in proportion to the
414 prediction error between the expected efficacy and the efficacy feedback on a given trial (i.e.,
415 whether a given outcome was determined by performance or not) (Sutton and Barto 2018).
416 Interestingly, consistent with studies of reward learning (Niv et al. 2012; Collins and Frank 2014;
417 Lefebvre et al. 2017; Chambon et al. 2020; Garrett and Daw 2020), we found that the RL model
418 that best accounted for our data allowed efficacy estimates to be updated differently from trials
419 that were more efficacious than expected (Performance-Based feedback) than from trials that
420 were less efficacious than expected (Random feedback), by having separate *learning rates*
421 scaling prediction errors in the two cases. Even when accounting for model complexity, this Two
422 Learning Rate Efficacy model outperformed a One Learning Rate Efficacy model as well as a
423 baseline model that fits a single constant efficacy estimate and no learning (Intercept Model)
424 (Figure 2B). In addition, we were able to successfully recover the parameters of this model from
425 a simulated dataset matched to our data (see Methods and Figure S1). We found that the two
426 learning rates for this best-fit model varied across the group (Figure 2E), but did not find that one
427 was consistently higher than the other ($b = 0.02$; 95% CrI [-0.04, 0.08]; $p_{b < 0} = 0.260$;
428 $BF_{01} = 13.55$; Figure S4), despite most participants (80%) tending to learn more from
429 Performance-Based than Random trials. Finally, model-based efficacy estimates were strongly
430 related to the raw subjective estimates on trials on which participants reported efficacy ($b = 0.77$;
431 95% CrI [0.62, 0.91]; $p_{b < 0} < 0.001$; $BF_{01} > 100$; $R^2 = 0.50$), demonstrating that the model
432 successfully captured the raw estimates. Taken together, these results show that participants
433 dynamically updated their expected efficacy based on trial-by-trial feedback, and that they did so
434 differentially based on whether the trial was more or less efficacious than expected. The fitted

435 models further enable us to generate trial-by-trial estimates of expected efficacy and efficacy
436 prediction errors, which we use in model-based analyses of behavior and neural activity below.
437 Note that in all the following analyses we control for reward estimates obtained from models fit
438 to reward feedback (for details see below).

439 ***The feedback-related P3b indexes updating of efficacy expectations***

440 To investigate the neural mechanism underlying feedback-based learning of efficacy, we
441 probed the centroparietal P3b ERP component (Figure 3A), an established neural correlate of
442 prediction-error based learning (Fischer and Ullsperger 2013; Nassar et al. 2019). If the P3b
443 indexes feedback-based updating of *efficacy* predictions, as it does for reward predictions, we
444 would expect this ERP to demonstrate several key markers of such a learning signal. First, we
445 would expect the P3b to reflect the extent to which efficacy feedback (Performance-Based vs.
446 Random) deviates from the current level of expected efficacy. In other words, the P3b should
447 track the magnitude of the *unsigned efficacy prediction error (PE)* on a given trial. We tested
448 this by examining how the amplitude of the P3b to a given type of efficacy feedback varied with
449 model-based estimates of the participant's efficacy expectation on that trial, while holding the
450 expected reward rate constant (see the Methods and Materials sections for the details of the
451 experimental design and the statistical models). If the P3b signaled efficacy PE then its
452 amplitude should scale inversely with the expected probability of a given type of feedback (i.e.,
453 how *unexpected* that feedback is), thus correlating *negatively* with expected efficacy on trials
454 providing Performance-Based feedback and correlating *positively* with expected efficacy on
455 trials providing Random feedback. In addition to overall higher P3b to performance-based
456 compared to random feedback ($b = 0.86$; 95% CrI [0.42, 1.31]; $p_{b < 0} = 0$; $BF_{10} = 30.86$; Figure
457 2B), we found exactly this predicted interaction between feedback type and expected efficacy (b

458 = -2.40; 95% CrI [-4.07, -0.74]; $p_{b>0} = 0$; $BF_{10} = 24.40$; Figure 2B; Figure S5A; Table S2), with
459 the P3b amplitude decreasing with model-based estimates of expected efficacy on Performance-
460 Based trials ($b = -1.49$; 95% CrI [-2.72, -0.27]; $p_{b>0} = 0.03$; $BF_{10} = 1.52$) and increasing
461 numerically, but not robustly with expected efficacy on Random trials ($b = 0.91$; 95% CrI [-0.34,
462 2.21]; $p_{b<0} = 0.12$; $BF_{10} = 0.44$). Accordingly, when we regressed P3b amplitude on our model-
463 based estimates of trial-to-trial efficacy PE, we found a positive relationship ($b = 1.25$; 95% CrI
464 [0.35, 2.15]; $p_{b<0} = 0.01$; $BF_{10} = 5.84$; Figure 3C-left; Table S2).

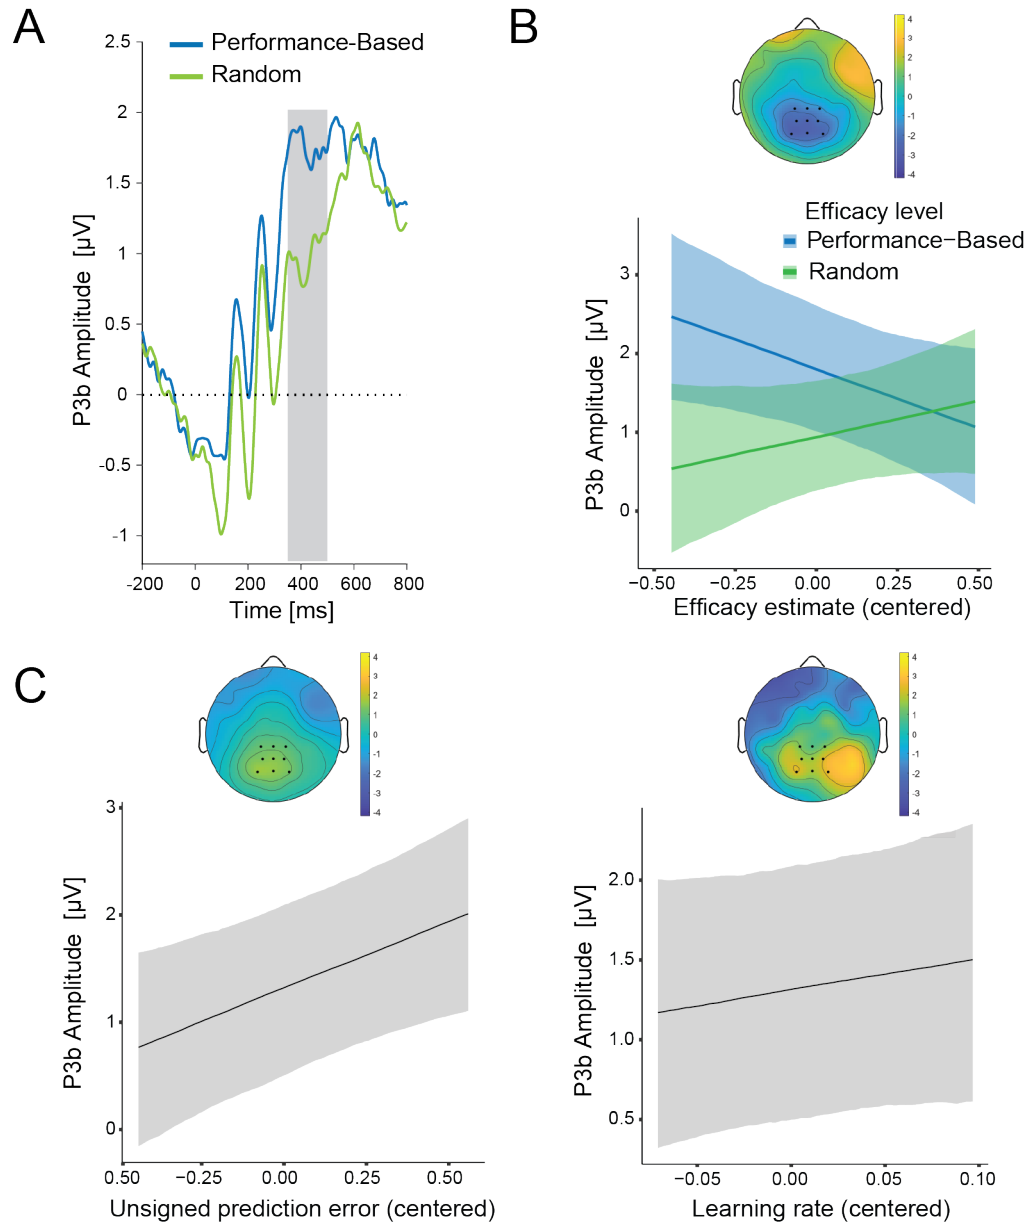


Figure 3. P3b reflects dynamically changing efficacy estimates during processing of efficacy feedback in Study 1. **A.** ERP average for the P3b locked to the onset of the efficacy feedback separately for performance-based and random feedback. The grey area shows the time window used for quantifying the P3b. **B.** LMM predicted P3b amplitudes are plotted for performance-based and random feedback as a function of efficacy estimates. The topography shows the interaction of efficacy estimate with efficacy feedback in the P3b time window. **C.** Predicted (centered) effects of unsigned prediction errors (left) and model-based learning rates (right) on the P3b. Shaded error bars represent 95% confidence intervals. Topographies display fixed-effects estimates.

466 In addition to tracking efficacy PEs, the second key prediction for the P3b if it indexes
467 efficacy learning, is that it should track the extent to which PEs are used to update estimates of
468 efficacy (i.e., the learning rate). In the current study, we found that participants differed in their
469 learning rates for the two forms of efficacy feedback (Performance-Based vs. Random),
470 providing us with a unique opportunity to perform a within-subject test of whether P3b tracked
471 differences in learning rate across these two conditions. Specifically, we could test whether a
472 given subject's P3b was greater in the feedback condition for which they demonstrated a higher
473 learning rate. However, we have not found conclusive evidence for the increase in P3b for the
474 within-subject feedback condition with the higher learning rate ($b = 2.00$; 95% CrI [-2.21, 6.04];
475 $p_{b<0} = 0.17$; BF = 1.08; Figure 3C-right). While, theoretically, prediction error and learning rate
476 would interact in predicting the P3b amplitude, we did not observe such an interaction here. This
477 finding is in line with previous work on reward processing (Fischer and Ullsperger 2013), which
478 has found additive effects of prediction errors and learning rate on P3b. While we found the
479 effect of prediction errors, our learning rate effect was going in the expected direction, but was
480 not conclusive.

481 ***The CNV indexes control allocation based on updated expectations of efficacy***

482 Thus far, our findings suggest that participants dynamically updated expectations of their
483 performance efficacy based on feedback, and that the P3b played a role in prediction error-based
484 updating of these expectations. Next, we tested the prediction that learned efficacy estimates
485 determine the expected benefits of control, and thus influence how much control is allocated
486 (Shenhav et al. 2013). We have previously shown that people exert more control when expecting
487 their performance to be more rather than less efficacious on the upcoming trial (Frömer et al.,
488 2021a). This was reflected in better behavioral performance and higher amplitudes of the

489 contingent negative variation (CNV; Figure 4B-left) - a slow negative fronto-central waveform
490 preceding target onset, which is increasingly negative as the amount of control allocated in
491 preparation for the task increases (Grent-'t-Jong and Woldorff 2007; Schevernels et al. 2014a).
492 Here we used the same marker of control, but, unlike in our previous study, efficacy expectations
493 were (a) learned rather than explicitly instructed; (b) varied over time rather than having a fixed
494 association with a cue; and (c) varied continuously across the range of efficacy probabilities
495 rather than alternating between maximal and minimal efficacy. We were therefore interested in
496 testing whether these dynamically varying expectations of efficacy, as estimated by our model,
497 would still exert the same influence on behavior and neural activity.

498 Consistent with our predictions and our previous findings, participants tended to perform
499 better when they expected performance to be more efficacious, responding faster on correct trials
500 with increasing model-based estimates of efficacy (Figure 4A; Figure S6A; Table S3). This
501 finding replicates the performance effects we observed using instructed cues, albeit with only
502 modest evidence ($b = -16.00$; 95% CrI [-34.91, 2.96]; $p_{b > 0} = 0.05$; $BF_{01} = 1.68$). Like in our
503 previous studies, faster responding was not explained by a change in speed-accuracy trade-off, as
504 accuracy did not decrease with increasing efficacy ($b = 0.12$; 95% CrI [-0.20, 0.44]; $p_{b > 0} = 0.23$;
505 $BF_{01} = 1.99$; Figure S6A). These analyses controlled for the standard behavioral effects related to
506 Stroop congruency (i.e., slower and less accurate responding for incongruent relative to
507 congruent trials; Figure S7), as well as for the reward rate estimates.

508 If the CNV provides an index of control allocation based on current incentive
509 expectations, it should both reflect one's latest efficacy estimate and should predict performance
510 on the upcoming trial. Our findings support both predictions. Regressing single-trial CNV
511 amplitude onto model-based efficacy estimates, and controlling for expectations of reward

512 (discussed below), we found that the CNV amplitudes had a positive relationship² with the
513 current efficacy expectations ($b = -0.35$; 95% CrI [-0.85, 0.16]; $p_{b>0} = 0.09$; $BF_{10} = 2.73$; Figure
514 4B-right; Figure S5B; Table S4). However, this effect was weaker than in the previous
515 experiment with cued efficacy levels, which is to be expected given that in this experiment
516 participants had to learn their efficacy levels. As with the behavioral finding above, this result
517 provides evidence consistent with our previous CNV finding using instructed cues. We further
518 replicate our earlier finding (Frömer et al., 2021a) that larger CNV amplitude in turn predicted
519 better performance on the upcoming trial, with participants responding faster ($b = 11.41$; 95%
520 CrI [8.09, 14.72]; $p_{b<0} = 0$; $BF_{10} > 100$; Figure 4C; Table S5) and more accurately ($b = -0.07$;
521 95% CrI [-0.12, -0.01]; $p_{b>0} = 0.01$; $BF_{10} = 3.14$; Figure 4D; Table S5) as CNV increased.

² Note that the CNV is a negative component, thus higher CNV amplitudes (i.e., more control allocation) are more negative.

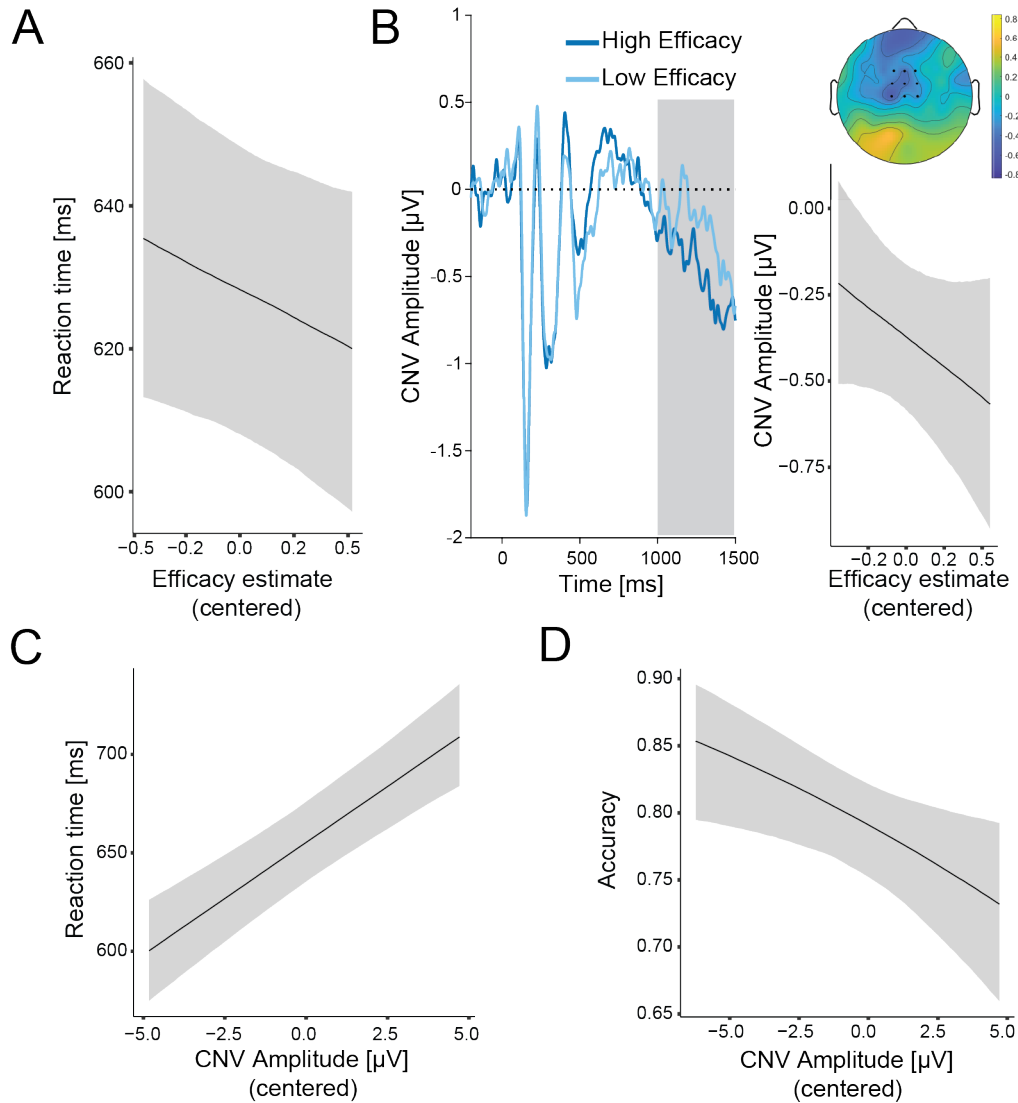


Figure 4. Efficacy estimates influence allocation of control and behavior in Study 1. A. Higher efficacy predicts faster accurate responses. **B. CNV increases with higher efficacy.** Left: Grand-average ERP for high and low efficacy estimates (median split used for plotting). The shaded area marks the time window used for CNV quantification. Time 0 corresponds to the onset of the neutral cue. Right: Predicted CNV amplitudes as a function of efficacy estimates. The topography shows the fixed effect of the efficacy estimate from the fitted linear model. **C. -D. Larger CNV amplitude predicts better performance.** Predicted accurate RT (**C.**) and accuracy (**D.**) as a function of efficacy estimates. Shaded error bars indicate 95% confidence intervals.

523 ***Parallel learning of efficacy and reward rate***

524 We held the amount of reward at stake constant over the course of the experiment, but the
525 frequency of reward (reward rate) varied over the course of the session based on a participant's
526 recent performance, and participants were sporadically probed for estimates of their reward rate
527 (interleaved with trials that were followed by efficacy probes). Our efficacy manipulation
528 explicitly controlled for this variability by yoking Random-Outcome feedback to a participant's
529 recent reward rate (see methods for details). However, this additional source of variability also
530 provided an opportunity to examine additional mechanisms of learning and adaptation in our
531 task. As in the case of efficacy estimates, reward rate estimates were robustly predicted by
532 reward feedback on previous trials (Table S1; Figure S8A), and this reward rate learning process
533 was well captured by a two learning rate reward rate model (Garrett and Daw 2020; S8B-C),
534 with the model-based estimates successfully predicting the reported subjective estimates ($b =$
535 0.82 ; 95% CrI[0.60, 1.02]; $p_{b < 0} < 0.001$; $BF_{01} > 100$; $R^2 = 0.58$; Figure S8B-C). Updates to these
536 reward rate estimates were reflected in P3b modulations of (unsigned) reward prediction errors
537 and associated learning rates (Fischer and Ullsperger 2013; Figure S9; Table S6). This pattern of
538 results provides additional evidence that efficacy learning involves similar neural and
539 computational mechanisms as reward-based learning.

540 **Study 2**

541 ***Learned efficacy modulates control over information processing***

542 Our findings suggest that people rely on domain-general mechanisms to learn about their
543 performance efficacy in a given environment, and use these learned estimates of efficacy to
544 optimize performance on their task. Specifically, in Study 1 we found that higher levels of
545 learned efficacy are associated with faster responses, albeit with modest evidence ($b = -16.00$;

546 95% CrI [-34.91, 2.96]; $p_{b > 0} = 0.05$; $BF_{01} = 1.68$). We also found that this speeding occurred on
547 correct but not incorrect trials, suggesting that these performance adjustments reflected
548 attentional control rather than adjustments to speed-accuracy tradeoffs. However, these findings
549 remain only suggestive in the absence of a formal model, and the presence of a stringent
550 response deadline in this study (individually calibrated for each subject during the practice phase
551 to ensure the reward rate of 80%; $M = 796\text{ms}$; $SD = 73\text{ms}$) presented an obstacle to fitting our
552 behavioral data to such a model without additional assumptions (e.g., regarding the form of a
553 collapsing threshold).

554 To provide further support for our proposal that learned efficacy influences control over
555 information processing, we ran an additional behavioral study (Study 2). Participants in this
556 study ($N = 51$) performed a web-based version of the task in Study 1, with the biggest
557 modification being that the Stroop trials (now using picture-word rather than color-word
558 interference) were performed over the course of short self-paced time intervals rather than trial
559 by trial as in Study 1. Specifically, participants were given limited time windows (2-4s) to
560 complete as many Stroop trials as they wanted to and were rewarded at the end of each interval
561 (cf. Leng et al., 2021). When rewards were performance-based, participants received a number of
562 points exactly proportional to the number of correct responses they gave during that window.
563 When rewards were random, the number of points was unrelated to performance on that interval
564 but (as in Study 1) was yoked to their performance in previous performance-contingent intervals.
565 Following our approach in Study 1, we varied the likelihood of a given interval being
566 performance-based or random over the course of the session (Figure S2A), and sporadically
567 probed participants for their subjective estimates of expected efficacy and reward rate. While in
568 most respects very similar to the paradigm in Study 1, one noteworthy feature of this self-paced

569 design is that it resulted in a much less stringent deadline for responding, thus producing reaction
570 time patterns more typical of free-response paradigms for which the traditional (fixed-threshold)
571 DDM was designed. Note that because this was an online sample we also employed additional
572 cutoff criteria to exclude inattentive participants, as detailed in the Methods section and Figures
573 S2B-C and S3.

574 Replicating the learning patterns observed in Study 1, subjective estimates of efficacy in
575 Study 2 reflected a running average over recent efficacy feedback ($b = 0.18$; 95% CrI [0.16,
576 0.20]; $p_{b < 0} = 0$; $BF_{10} > 100$). This effect was again weighted towards the most recent feedback
577 but still present up to five intervals back (all $p_{b < 0} < 0.01$; Figure 5A, Figure S10, and Table S7).
578 As in Study 1, this learning pattern was best captured by an RL algorithm with two learning rates
579 (Figure 5B), though positive and negative efficacy learning rates did not significantly differ from
580 one another on average ($b = 0.03$; 95% CrI [-0.03, 0.08]; $p_{b < 0} = 0.260$; $BF_{01} = 11.74$; Figure 5C).
581 As in Study 1, the model-based efficacy estimates successfully predicted the raw subjective
582 estimates on intervals on which participants reported their efficacy beliefs ($b = 0.86$; 95% CrI
583 [0.81, 0.90]; $p_{b < 0} < 0.001$; $BF_{01} > 100$; $R^2 = 0.60$), and the same was true for the model-based
584 reward estimates predicting the subjective reward estimates ($b = 1.00$; 95% CrI [0.95, 1.04]; $p_{b < 0}$
585 < 0.001 ; $BF_{01} > 100$; $R^2 = 0.66$).

586 Critically, we once again found that higher model-based estimates of efficacy predicted
587 better performance on the upcoming interval. Participants responded faster ($b = -10.25$; 95% CrI
588 [-19.86, -0.20]; $p_{b > 0} = 0.02$; $BF_{01} = 2.33$) and more accurately ($b = 0.23$; 95% CrI [0.03, 0.43];
589 $p_{b > 0} = 0.01$; $BF_{01} > 100$) with increasing model-based efficacy estimates (Figure 6B, Figure
590 S6B, and Table S8). To formally test whether these behavioral patterns reflected adjustments in
591 information processing (i.e., the rate of evidence accumulation once the stimuli appeared) or

592 instead reflected adjustments in speed-accuracy tradeoffs (i.e., the threshold for responding), we
593 fit these data with the hierarchical drift diffusion model (HDDM; Wiecki, Sofer, & Frank, 2013).
594 We tested whether model-based efficacy estimates predicted trial-by-changes in drift rate,
595 threshold, or both, while controlling for influences of expected reward rate on those same DDM
596 parameters. We found that higher levels of expected efficacy were associated with higher drift
597 rates ($b = 0.07$; 95% CrI [0.14, 0.43]; $p_{b < 0} = 0.00$) but were uncorrelated with threshold levels (b
598 $= -0.00$; 95% CrI [-0.04, 0.04]; $p_{b < 0} = 0.74$). Expected reward rate was not correlated with either
599 drift rate or threshold (Table S9). These results suggest that participants responded to changes in
600 performance efficacy by adjusting their attention to the task, rather than simply adjusting their
601 response threshold (i.e., becoming more or less impulsive).

602

603

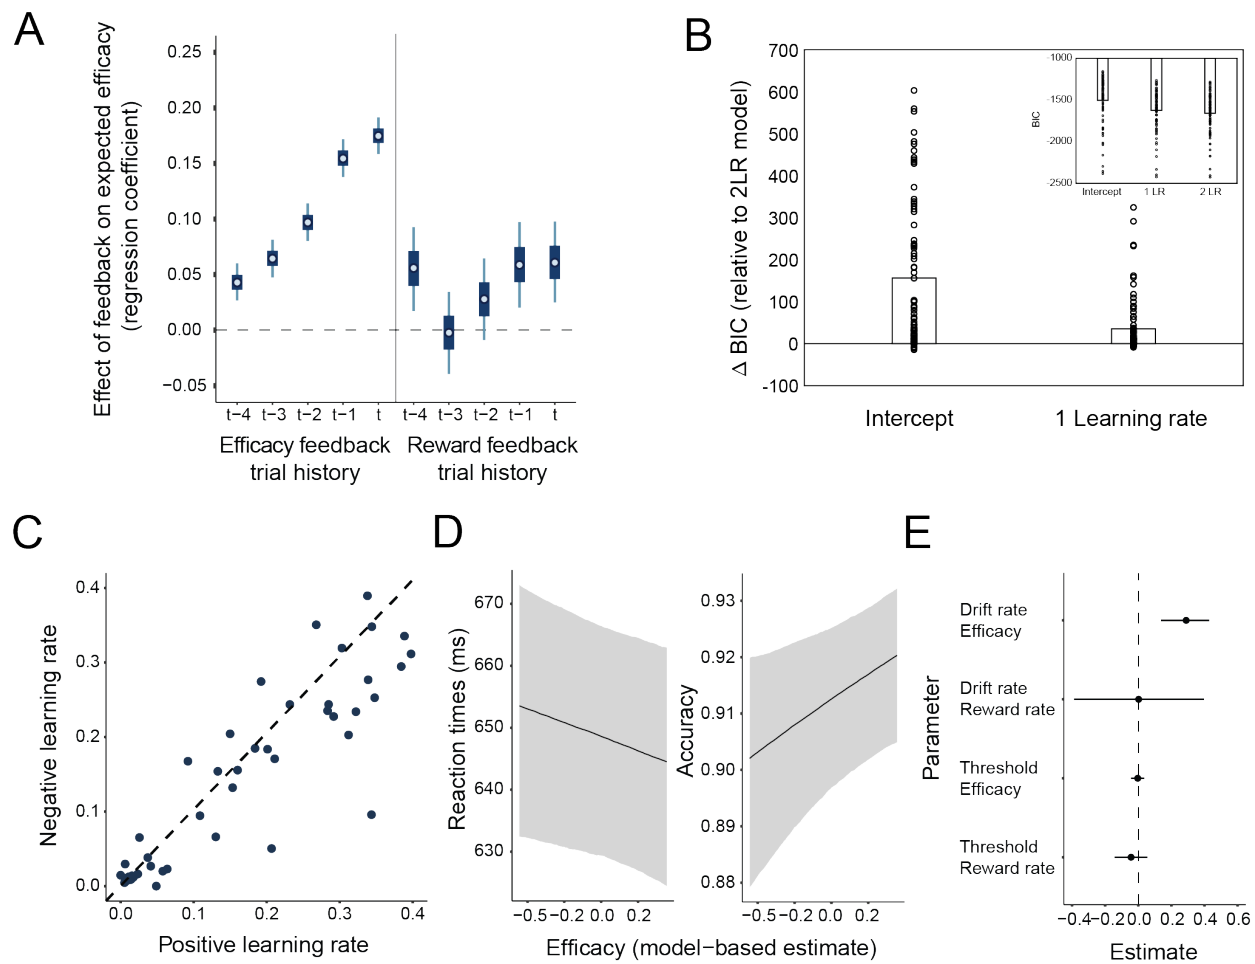


Figure 5. Efficacy is learned in the same way in a modified task and influences a behavioral marker of control allocation. **A.** Regression weights for the influence of previous feedback type (efficacy and reward) on the subjective efficacy estimate. Error bars represent 50% and 95% highest density intervals. **B.** **Two learning rate model captures efficacy learning best.** Differences in BICs between the two learning rate model, and the intercept-only and one learning rate models. **C.** **Efficacy learning rates.** Positive and negative efficacy learning rates for each participant. **D.** **Higher model-based efficacy estimates predict better behavioral performance.** Higher efficacy estimates reduce reaction times (left) and improve accuracy (right). **E.** **Higher model-based efficacy estimates predict increased allocation of attention.** Parameter estimates from the drift diffusion model. Higher efficacy estimates increase drift rates, but not response caution (thresholds).

604

Discussion

605

To evaluate the expected benefits of investing cognitive control into a task, people need

606

to consider whether this control is necessary and/or sufficient for achieving their desired

607

outcome (i.e., whether these efforts are worthwhile). A critical determinant of the

608 worthwhileness of control is performance efficacy, the extent to which performance on a control-
609 demanding task matters for outcome attainment versus those outcomes being determined by
610 performance-unrelated factors. However, the mechanisms through which people estimate the
611 efficacy of their performance based on previous experience are largely unknown. Here, we
612 identified the computational and neural mechanism through which efficacy is learned and used to
613 guide the allocation of cognitive control. Across two experiments, we found that participants
614 dynamically updated expectations of the efficacy of their task performance (i.e., the likelihood
615 that this performance will determine reward attainment), and used those expectations to adjust
616 how much control they allocated. The feedback-based updating of efficacy was well-captured by
617 a prediction error-based learning algorithm. Model-based estimates of efficacy and efficacy
618 prediction errors were encoded by canonical neural signatures of effort allocation and belief
619 updating, respectively. Importantly, these findings cannot be explained by variability in reward,
620 as reward rate was held constant across efficacy levels, and the subjective reward rate was
621 controlled for statistically. Further, our model-based analysis revealed that people allocated more
622 control when they learned that they had more efficacy, extending our previous findings on
623 instructed efficacy (Frömer et al., 2021a). Taken together, our results uncover the mechanism
624 through which efficacy estimates are learned and used to inform mental effort investment within
625 a given task environment.

626 Previous research has characterized the learning algorithms people use to learn the reward
627 value of different states and actions in their environment (Gläscher et al. 2010; Sutton and Barto
628 2018). Recent theoretical (Jiang et al. 2014; Lieder et al. 2018; Verbeke and Verguts 2019) and
629 empirical (Bejjani et al. 2018; Otto and Daw 2019; Jiang et al. 2020; Bustamante et al. 2021 Jan
630 6) work has extended this research to show how similar algorithms guide learning and adaptation

631 of cognitive control under varying rewards and task demands within a given task environment.
632 Our findings extend this work further in several ways. First, we show that people leverage
633 weighted prediction errors when learning about the efficacy of task performance, independently
634 of potential reward and task difficulty. Second, we show that they update their efficacy
635 expectations differently depending on whether performance was more efficacious or less
636 efficacious than they expected, demonstrating a striking parallel with dual learning rate models
637 that have been found to prevail in research on reward learning (Collins and Frank 2013; Lefebvre
638 et al. 2017; Garrett and Daw 2020), including in our own reward rate data. Third, we show that
639 participants dynamically adjust their control allocation based on learned efficacy, just as they do
640 for learned rewards and task demands (Bugg et al. 2011; Jiang et al. 2014; Lieder et al. 2018).

641 Our neural findings build further on past research on learning and control adaptation. The
642 P3b component has been shown to track prediction-error based learning from action-relevant
643 outcome values (Fischer and Ullsperger 2013; Nassar et al. 2019; Lohse et al. 2020). Here we
644 show that this neural marker tracks learning about efficacy in the same way as it tracks learning
645 about rewards. We found increased P3b amplitudes when people experienced feedback about
646 outcome contingencies that was less expected given their current estimate of efficacy (e.g.,
647 expecting low efficacy, but getting performance-contingent feedback), relative to when these
648 contingencies were more expected (e.g., expecting low efficacy and getting random feedback).
649 Our additional finding that P3b amplitude was overall larger for efficacy compared to no efficacy
650 feedback demonstrates that our participants were not just tracking the frequency of the two types
651 of feedback, as predicted by an oddball account. Instead this finding suggests that they were
652 actively learning from the feedback.

653 Extending previous findings on cueing efficacy and/or reward (Schevernels et al. 2014b;
654 Frömer et al. 2021a), our CNV and behavioral results further show that participants used these
655 learned efficacy estimates to calibrate their effort and their performance. Notably, unlike in
656 previous work, our study shows effort-related changes in CNV amplitude entirely divorced from
657 perceptual cues, providing evidence that this activity truly reflects adjustments in control, rather
658 than reactive processing of features associated with the cue. Taken together, our findings suggest
659 that similar neural mechanisms underlie learning and control adaptation in response to variability
660 in one's efficacy in a task environment, as they do in response to variability in expected rewards
661 (Leng et al., 2021; Otto & Daw, 2019). By fitting behavioral data from this task to a drift
662 diffusion model (Study 2), we were able to further demonstrate that participants were adapting to
663 changes in expected efficacy by enhancing the processing of stimuli (i.e., increasing their rate of
664 evidence accumulation) - potentially via attentional control mechanisms – rather than by
665 adjusting their threshold for responding. This particular pattern of control adjustments was
666 predicted for the current task because performance-contingent rewards depended on responses
667 being *both* fast and accurate (as in Frömer et al., 2021a), but future work should test the
668 prediction that different control adjustments should emerge under different performance
669 contingencies (cf. Leng et al., 2021; Ritz et al., 2021).

670 Our findings build on previous research on how people learn about controllability of their
671 environment. Studies have examined the neural and computational mechanisms by which
672 humans and other animals learn about the contingencies between an action and its associated
673 outcome, and demonstrated that these learned action-outcomes contingencies influence which
674 actions are selected and how vigorously they are enacted (Dickinson and Balleine 1995;
675 Liljeholm et al. 2011; Manohar et al. 2017; Moscarello and Hartley 2017; Ly et al. 2019). Our

676 work extends this research into the domain of cognitive control, where the contingencies
677 between actions (i.e., control adjustments) and outcomes (e.g., reward) depend both on whether
678 control adjustments predicts better performance and whether better performance predicts better
679 outcomes (Shenhav et al., 2021). Learning control-outcome contingencies therefore requires
680 learning about how control states map onto performance (control efficacy) as well as how
681 performance maps onto outcomes (performance efficacy). By describing the mechanisms by
682 which people solve the latter part of this learning problem, and demonstrating that these are
683 comparable to those engaged during action-outcome learning, our study lays critical groundwork
684 for better understanding the links between selection of actions and control states.

685 Our efficacy-updating results are a reminder that many aspects of feedback are reflected in
686 prediction error signals (Langdon et al. 2018; Frömer et al. 2021b). In the present study, we
687 intentionally separated feedback about reward and efficacy to isolate the cognitive and neural
688 learning mechanisms associated with each. In doing so, we have taken an important first step
689 towards understanding the updating mechanisms underlying each. Further work is needed to
690 understand how they are inferred in more naturalistic settings, in which different forms of
691 feedback are often multiplexed, containing information about the values of actions that were
692 taken as well as about the features and structure of the environment (cf. Dorfman et al., 2019;
693 2021).

694 Another distinct element of more complex naturalistic environments is that the same
695 feedback can be used to evaluate multiple targets, internal ones, such as the selected response
696 and its predicted outcome, or external ones, such as the source of feedback/environment
697 (Carlebach and Yeung 2020). Such multi-level prediction error signals might for instance
698 explain, why despite close links between P3b and behavioral adaptation (Yeung and Sanfey

699 2004; Chase et al. 2011; Fischer and Ullsperger 2013), this link is context dependent and
700 attempts to link the P3b uniquely to behavioral updating have failed (Nassar et al., 2019).
701 Reinforcement learning, and predictive inference more generally, have been proposed to not only
702 support the selection of individual actions, but also extended sequences of actions and control
703 signals (Holroyd and Yeung 2012; Lieder et al. 2018). Alongside evaluations of actions and
704 environmental states, neural signatures of feedback-based learning could thus further reflect
705 evaluations of control signals, their quality or intensity. Given the many potential causes a given
706 outcome can have, and the flexibility that people have in how they use the feedback, it is easy to
707 see how feedback could be misattributed and lead to inaccurate beliefs about performance
708 efficacy. Such beliefs about environmental statistics can drive changes in feedback-processing
709 and behavioral adaptation, above and beyond the true statistics (Schiffer et al. 2017), and are thus
710 of particular importance for understanding some of the cognitive symptoms of mental disorders.

711 Understanding how efficacy estimates develop based on previous experiences is crucial for
712 understanding why people differ in how much cognitive control they allocate in different
713 situations (Shenhav et al.). People differ in their beliefs about how much control they have over
714 their environments (Leotti et al. 2010; Moscarello and Hartley 2017), and in their ability to
715 estimate their efficacy (Cohen et al. 2020). Further, many mental disorders, including depression
716 and schizophrenia, have been linked with one's estimates of their ability to control potential
717 outcomes in their environment, including through allocation of control (Huys and Dayan 2009;
718 Maier and Seligman 2016), and we have recently proposed that such changes can drive
719 impairments of motivation and control in those populations (Grahek et al. 2019). As we show in
720 this study, when people have learned to expect low efficacy, they will allocate less cognitive
721 control, which can manifest as apparent control deficits. The experimental and modeling

722 approach taken in our study helps uncover a more fine-grained view of how components of
723 motivation are learned and used to support the allocation of cognitive resources. In this way, our
724 study takes a first step toward a better computational and neural account of efficacy learning,
725 which can aid the understanding of individual differences in the willingness to exert mental
726 effort, as well as the development of interventions aimed at teaching individuals when these
727 efforts truly matter.

728

729

730

Supplementary materials

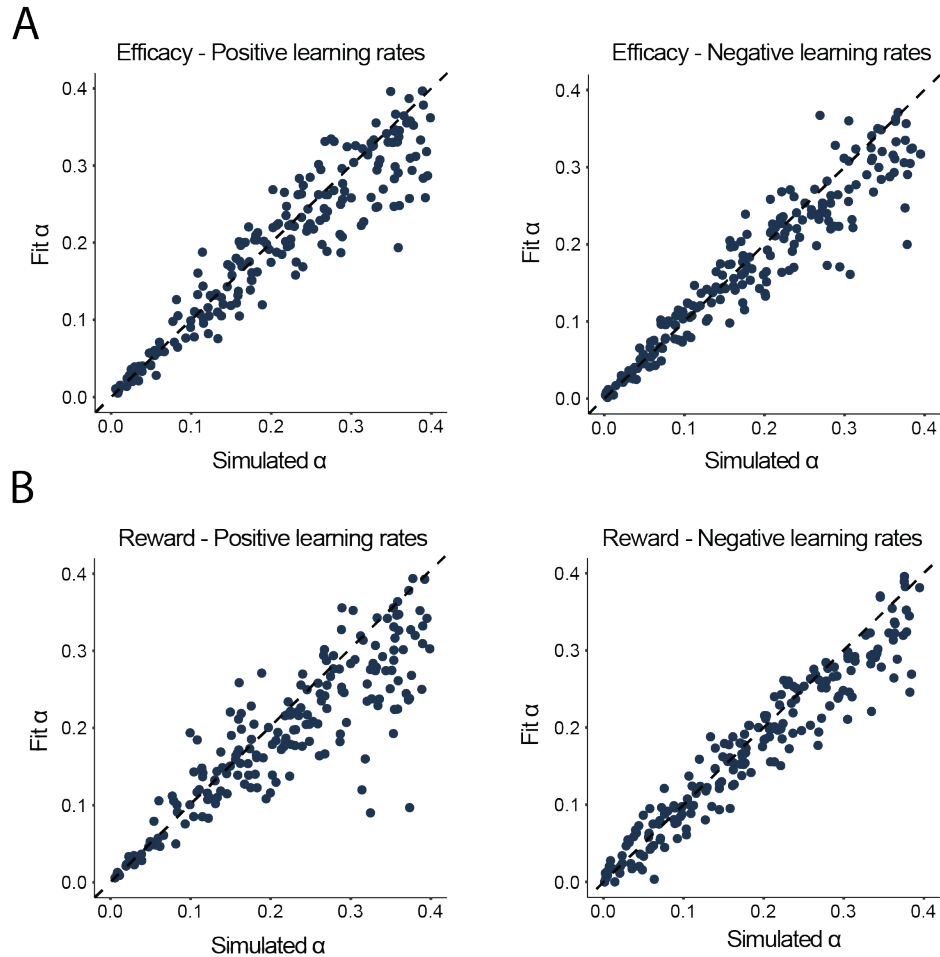


Figure S1. Parameter recovery study for the two learning rate models for the efficacy and reward learning. We simulated 200 subjects with the range of learning rates matched to the empirically observed range. The noise and the intercept parameters were fixed to match the empirically observed mean value. The number of trials and subjective estimates, as well as the sequence of efficacy and reward feedbacks were based on Experiment 1. **A. Learning rates for efficacy.** Simulated and recovered learning rates were highly correlated both in the case of positive ($r = 0.93$, $p < 0.001$) and negative learning rates ($r = 0.94$, $p < 0.001$). **B. Learning rates for reward.** Simulated and recovered learning rates were highly correlated both in the case of positive ($r = 0.88$, $p < 0.001$) and negative learning rates ($r = 0.96$, $p < 0.001$).

731

732

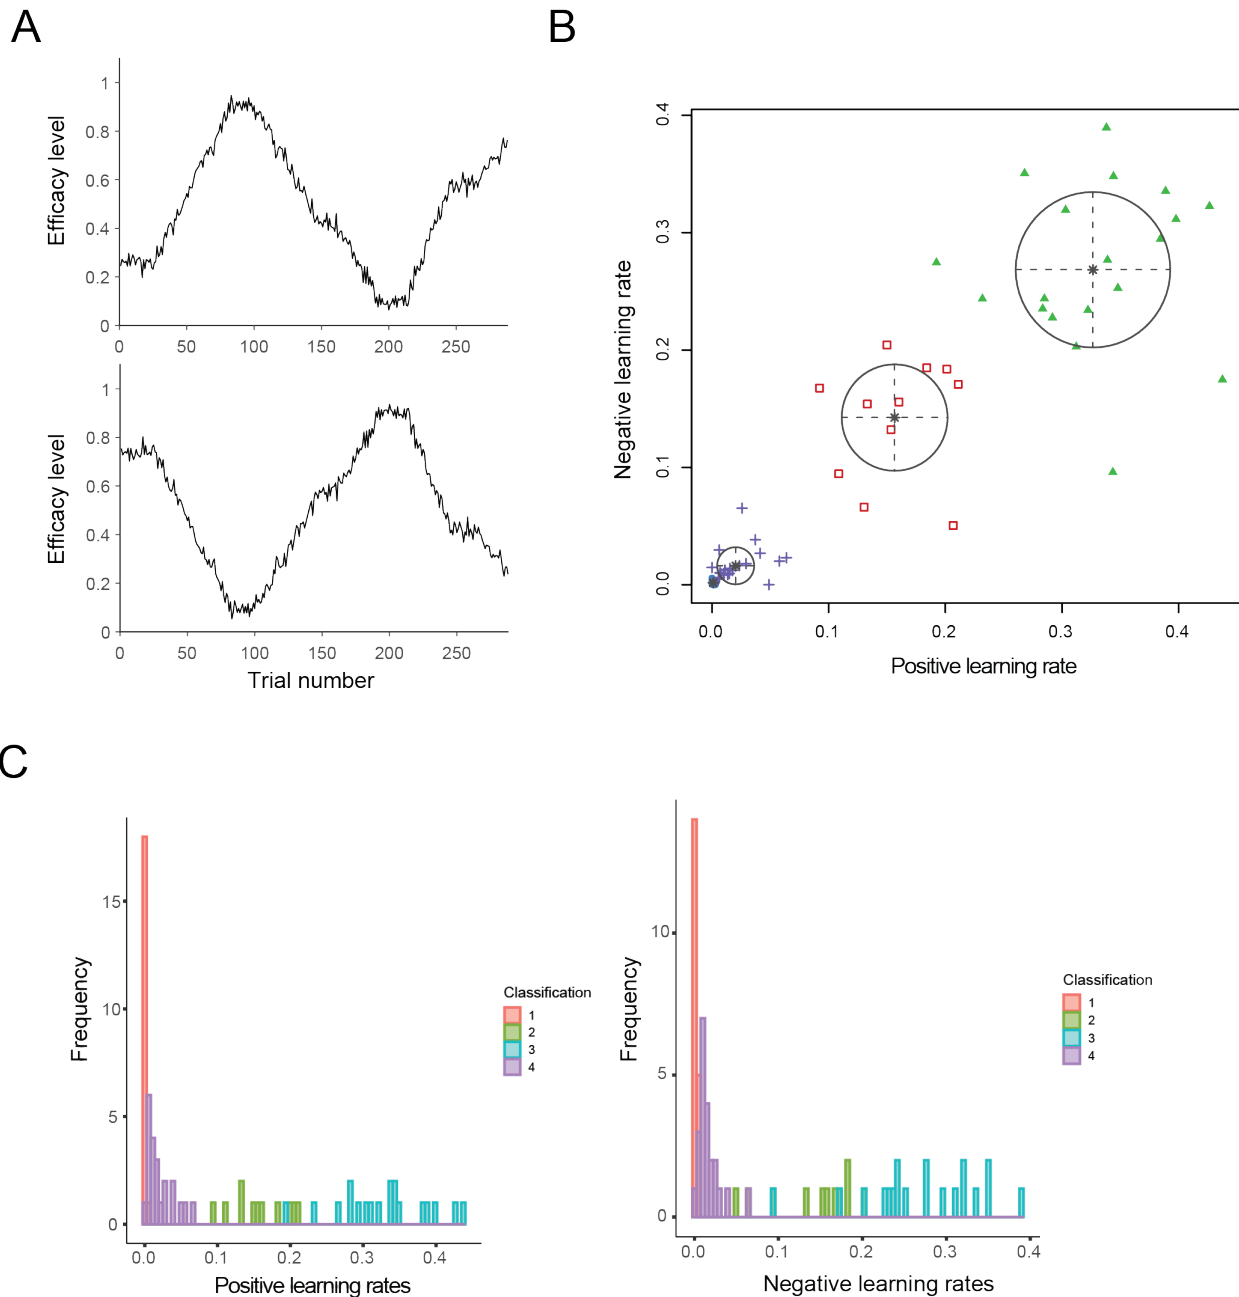


Figure S2. Efficacy drifts and the unsupervised clustering of the efficacy learning rates in Study 2. A. Efficacy drifts used in Study 2. The probability drifts of contingent feedback presented to the first (top) and second (bottom) half of the participants. **B. The results of the winning Gaussian mixture model for clustering of the efficacy learning rates.** One of the clusters (blue dots) included only the subjects with extremely small learning rates (all learning rates < 0.03; N=20; group 1 – blue dots). These subjects were excluded from the further analyses with the assumption that they did not pay attention to the feedbacks or that they were giving random answers to the probes. **C. Histograms of the positive and negative learning rates and their clusters.** For both then negative (left) and positive (right) learning rates cluster 1 included only the participants with very low learning rates.

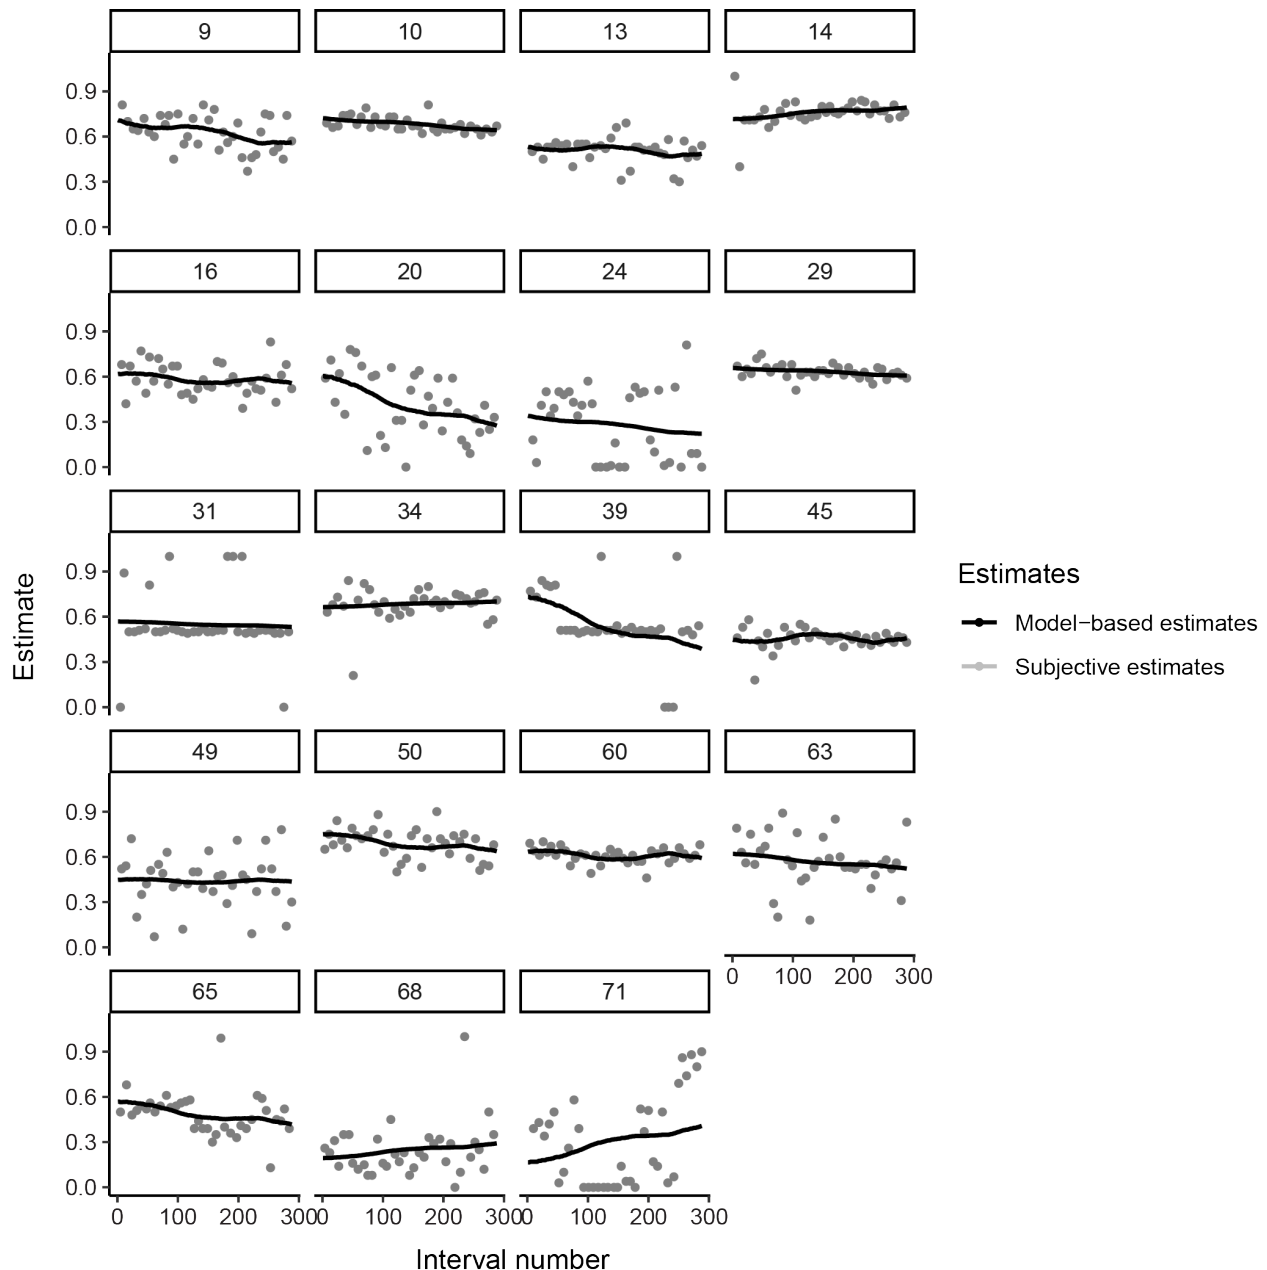


Figure S3. Self-reported and model-based efficacy estimates for subjects excluded from Study 2 based on very low learning rates. Subjective efficacy estimates (grey) and model-based efficacy estimates for each of the 19 subjects identified to form a cluster due to very low learning rates (both learning rates < 0.03) based on the unsupervised clustering algorithm using Gaussian mixture models. These subjects had low variance in their self-reported efficacy estimates, or did not appear to update their efficacy estimates based on feedbacks, suggesting that they were not paying attention to the feedbacks, or always providing very similar efficacy estimates.

Table S1

Study 1 regression weights for the models predicting subjective efficacy and reward rate estimates based on efficacy and reward feedback 5 trials back.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability ($p < 0$)</i>	<i>BF₁₀</i>
<i>Subjective efficacy</i>				
Intercept	0.52	0.49, 0.55	0.00	
T1 Performance-Based - Random	0.14	0.12, 0.16	0.00	>100
T2 Performance-Based - Random	0.12	0.10, 0.14	0.00	>100
T3 Performance-Based - Random	0.08	0.05, 0.10	0.00	>100
T4 Performance-Based - Random	0.03	0.01, 0.05	0.00	2.20
T5 Performance-Based - Random	0.03	0.01, 0.05	0.00	1.72
T1 Reward - No Reward	0.02	0.00, 0.04	0.02	0.37
T2 Reward - No Reward	0.03	0.01, 0.05	0.00	4.23
T3 Reward - No Reward	0.03	0.01, 0.05	0.01	0.95
T4 Reward - No Reward	0.02	0.00, 0.04	0.02	0.39
T5 Reward - No Reward	0.00	-0.02, 0.02	0.46	0.05
<i>Subjective reward rate</i>				
Intercept	0.50	0.46, 0.54	0.00	
T1 Reward - No Reward	0.19	0.17, 0.21	0.00	>100
T2 Reward - No Reward	0.16	0.14, 0.18	0.00	>100
T3 Reward - No Reward	0.09	0.07, 0.11	0.00	>100
T4 Reward - No Reward	0.05	0.03, 0.06	0.00	30.86
T5 Reward - No Reward	0.05	0.03, 0.07	0.00	>100
T1 Performance-Based - Random	-0.01	-0.03, 0.01	0.73	0.06
T2 Performance-Based - Random	0.00	-0.01, 0.02	0.33	0.05
T3 Performance-Based - Random	0.01	-0.01, 0.03	0.00	0.05
T4 Performance-Based - Random	0.02	-0.01, 0.03	0.20	0.08
T5 Performance-Based - Random	0.02	0.00, 0.03	0.04	0.20

734

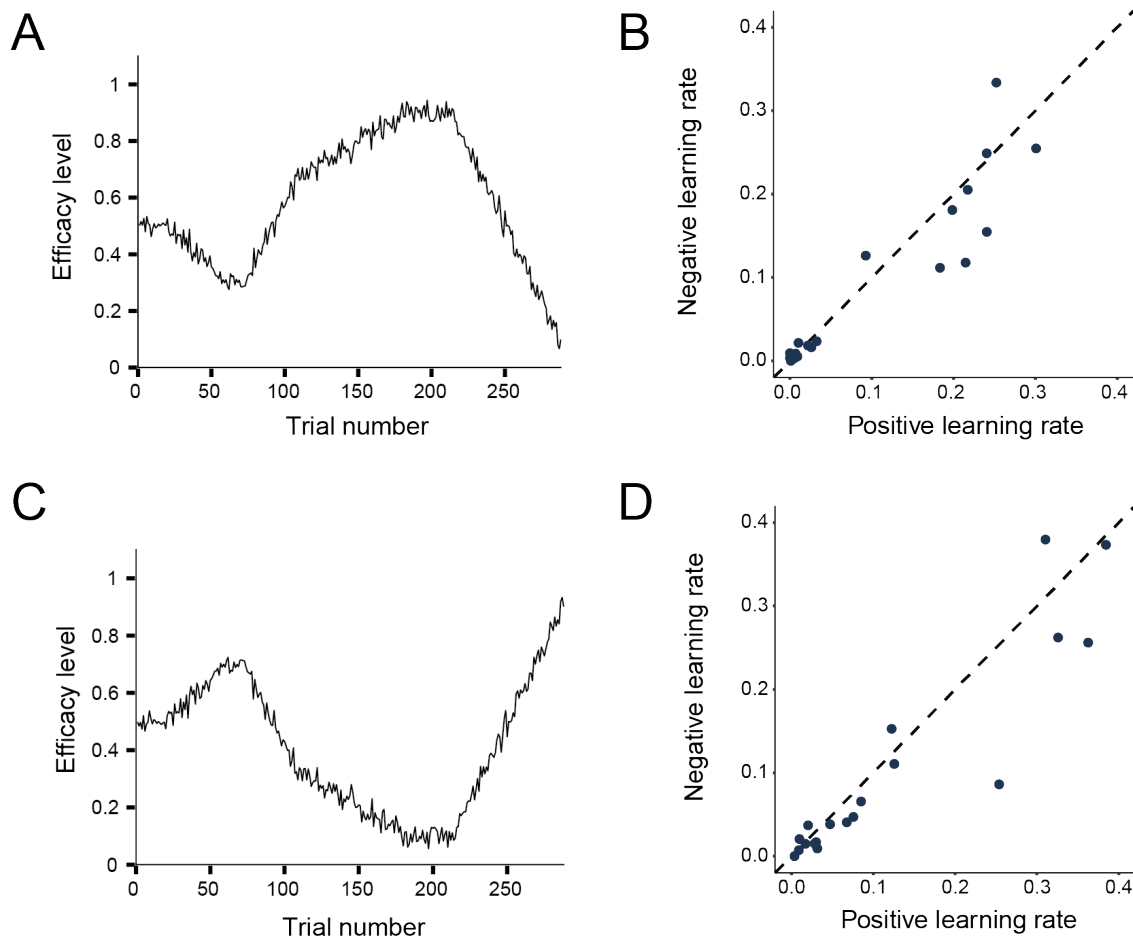


Figure S4. Efficacy learning in Study 1 – the effect of different feedback probability drifts on learning. **A.** The probability drift of contingent feedback presented to the first half of subjects. **B.** Learning rates for contingent and random feedback for the first half of subjects. Learning rates for the contingent and random feedback did not differ ($b = 0.01$; 95% CrI [-0.06, 0.08]; $p_{b < 0} = 0.38$; $BF_{10} = 0.09$). **C.** The inversed probability drift of contingent feedback presented to the second half of participants. **D.** Learning rates for contingent and random feedback did not differ ($b = 0.02$; 95% CrI [-0.06, 0.11]; $p_{b < 0} = 0.31$; $BF_{10} = 0.10$).

735

736

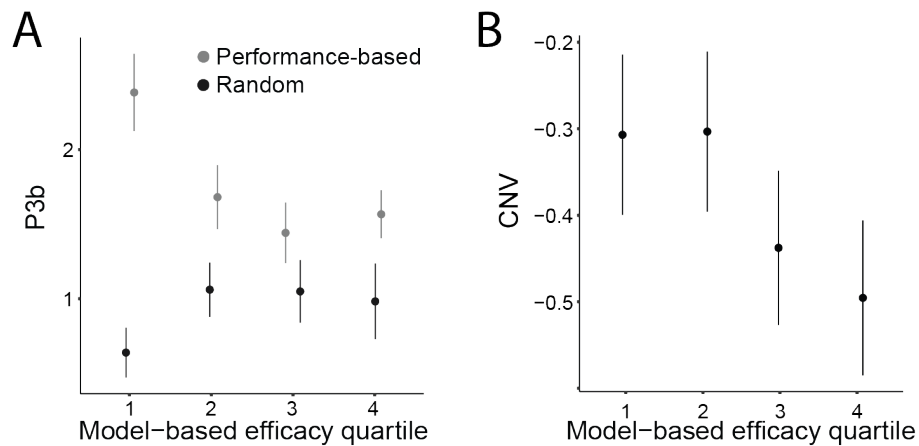


Figure S5. The influence of efficacy estimates on efficacy feedback processing and proactive cognitive control allocation. **A.** Means for the P3b amplitude in response to performance-based and random feedbacks across each quartile of the model-based efficacy estimates calculated for each subject. The P3b amplitude in response to performance-based feedback is higher when efficacy estimates are lower, suggesting more updating based on feedback. Error bars represent standard errors of the mean. **B.** Means for the CNV amplitude pre target onset across quartiles of the model-based efficacy estimates calculated for each subject. The CNV amplitude is more negative (more proactive control allocation) when efficacy estimates are higher. Error bars represent standard errors of the mean.

737

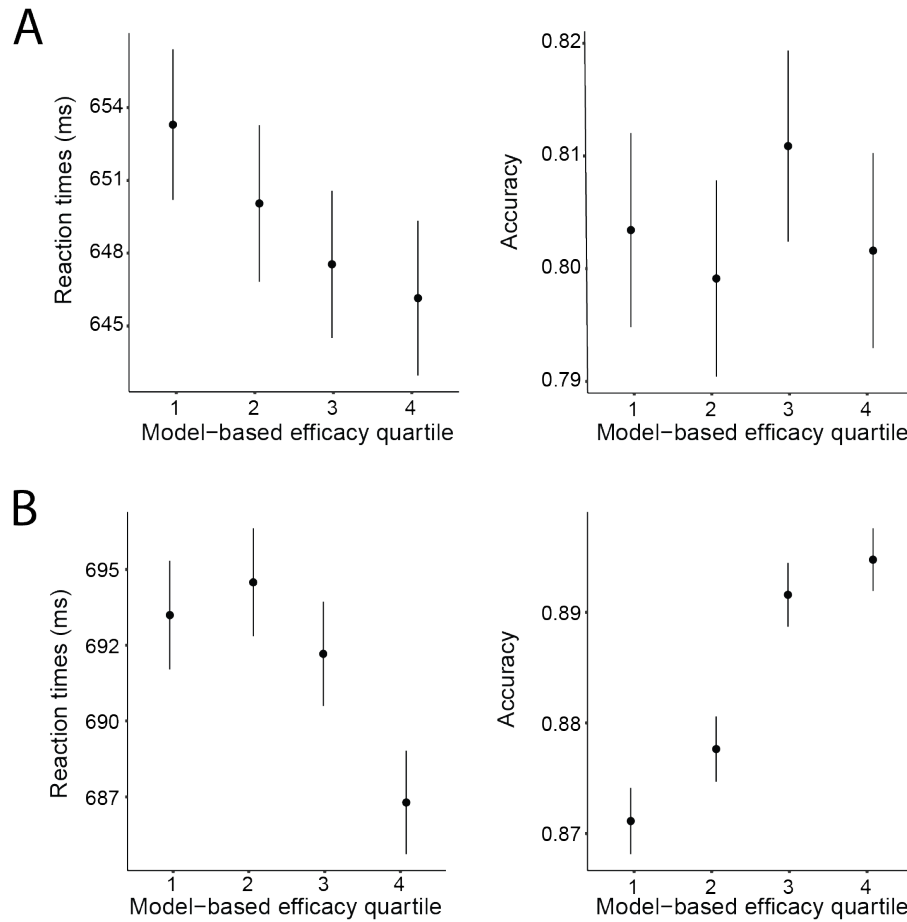


Figure S6. The influence of efficacy estimates on behavioral performance. A. Means for the reaction times (left) and accuracy (right) across quartiles of model-based efficacy estimates in Experiment 1. Reaction times are faster when efficacy estimates are higher, while there is no consistent pattern in accuracy data. Error bars represent standard errors of the mean. **B.** Means for the reaction times (left) and accuracy (right) across quartiles of model-based efficacy estimates in Experiment 2. Participants are faster to respond and more accurate when efficacy estimates are high relative to low. Error bars represent standard errors of the mean.

738

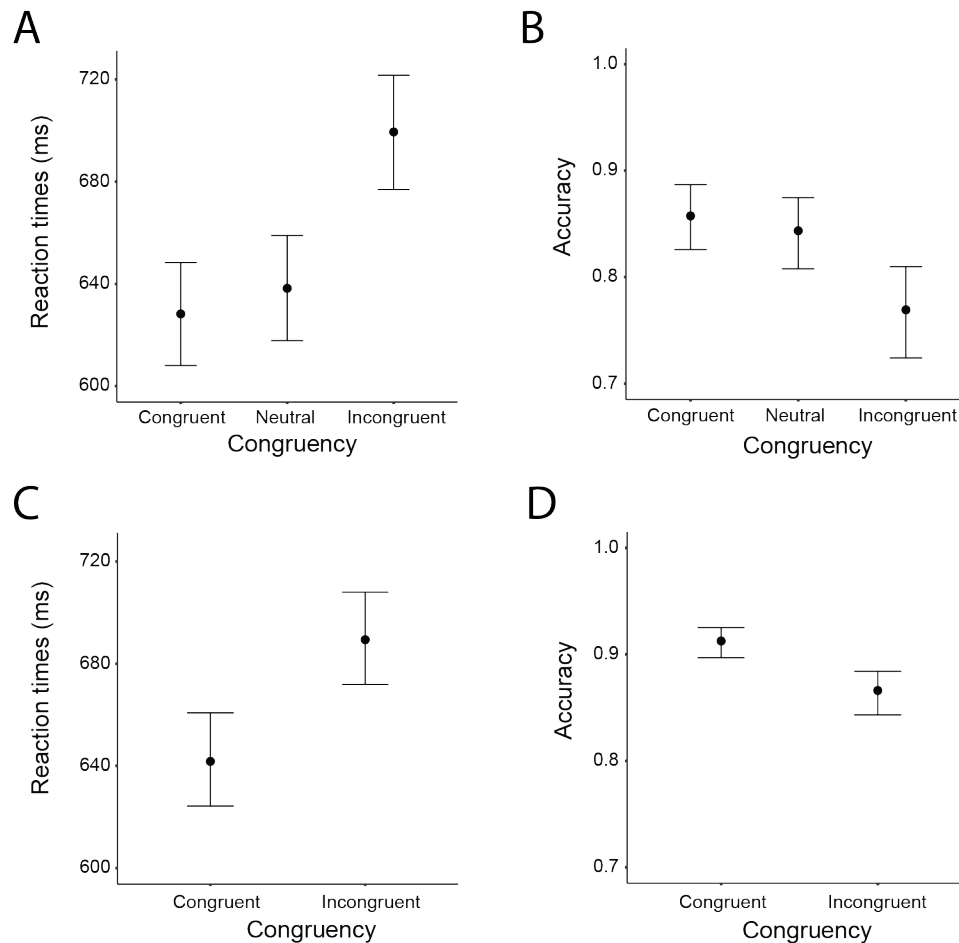


Figure S7. Behavioral effects of congruency in Study 1 and Study 2. **A.** In comparison to the neutral trials in Study 1, participants were faster to respond to congruent ($b = 9.95$; 95% CrI [3.33, 16.59]; $p_{b < 0} = 0.00$; BF = 1.60) and slower to respond to incongruent trials ($b = 61.24$; 95% CrI [49.50, 72.59]; $p_{b < 0} = 0.00$; BF = 6.73). **B.** In comparison to the neutral trials in Study 1, participants were equally likely to be correct when responding to congruent ($b = -0.11$; 95% CrI [-0.29, 0.06]; $p_{b > 0} = 0.091$; BF = 0.01) and less likely to be correct when responding to incongruent trials ($b = -0.48$; 95% CrI [-0.64, -0.33]; $p_{b > 0} = 0.00$; BF > 8.22). The regression weights for the accuracy analysis are in log odds. **C.** In Study 2 participants were slower to respond to incongruent relative to congruent trials ($b = 47.48$; 95% CrI [40.69, 52.22]; $p_{b < 0} = 0.001$; BF = 50). **D.** In Study 2 participants were less accurate when responding to incongruent compared to congruent trials ($b = -0.48$; 95% CrI [-0.57, -0.39]; $p_{b < 0} = 1$; BF > 100).

739

Table S2

Study 1 regression weights for the models predicting the P3b to efficacy feedback.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability ($p < 0$)</i>	<i>BF₁₀</i>
<i>Efficacy feedback</i>				
Intercept	1.38	0.67, 2.09	0.00	
Model-based efficacy	-0.29	-1.56, 0.97	0.68	0.23
Efficacy feedback type (PB-R)	0.86	0.42, 1.31	0.00	30.86
Model-based reward rate	-0.19	-1.50, 1.18	0.62	0.23
Model-based efficacy × Efficacy feedback type	-2.40	-4.07, -0.74	1.00	24.40
Model-based efficacy slope for Performance-Based feedback	-1.49	-2.72, -0.27	0.97	1.52
Model-based efficacy slope for Random feedback	0.91	-0.34, 2.21	0.12	0.44
<i>Predictions errors and learning rates</i>				
Intercept	1.31	0.51, 2.09	0.00	
Unsigned prediction error	1.25	0.35, 2.15	0.01	5.84
Learning rate	2.00	-2.21, 6.04	0.17	1.08
Model-based reward rate	-0.09	-1.48, 1.34	0.55	0.23
Unsigned prediction error × Learning rate	0.57	-4.87, 6.14	0.41	0.91

740

Table S3

Study 1 regression weights for the models predicting reaction times and accuracy based on efficacy and reward rate estimates.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability (p<0)</i>	<i>BF₀₁</i>
<i>Reaction times</i>				
Intercept	655.37	635.09, 675.51	0.00	
Facilitation	9.95	3.33, 16.59	0.00	1.60
Interference	61.24	49.50, 72.59	0.00	6.73
Model-based efficacy	-16.00	-34.91, 2.96	0.95	1.68
Model-based reward rate	-1.56	-27.66, 24.04	0.54	1.36
<i>Accuracy</i>				
Intercept	1.56	1.33, 1.80	0.00	
Facilitation	-0.11	-0.29, 0.06	0.91	0.01
Interference	-0.48	-0.64, -0.33	1.00	8.22
Model-based efficacy	0.12	-0.20, 0.44	0.23	1.99
Model-based reward rate	-0.36	-0.73, 0.03	0.97	0.29

741

Table S4

Study 1 regression weights for the models predicting the CNV based on efficacy and reward rate.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability (p<0)</i>	<i>BF₀₁</i>
<i>CNV</i>				
Intercept	-0.37	-0.58, -0.16	0.00	
Model-based efficacy	-0.35	-0.85, 0.16	0.09	2.73
Model-based reward rate	-0.09	-0.72, 0.55	0.38	2.19

742

743

Table S5 744

Study 1 regression weights for the models predicting reaction times and accuracy based on the CNV. 745

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability ($p < 0$)</i>	<i>BF₁₀</i>
<i>Reaction times</i>				
Intercept	655.28	635.44, 675.91	0.00	
CNV amplitude	11.41	8.09, 14.72	0.00	>100
<i>Accuracy</i>				
Intercept	1.33	1.11, 1.53	0.00	
CNV amplitude	-0.07	-0.12, -0.01	0.99	3.14

746

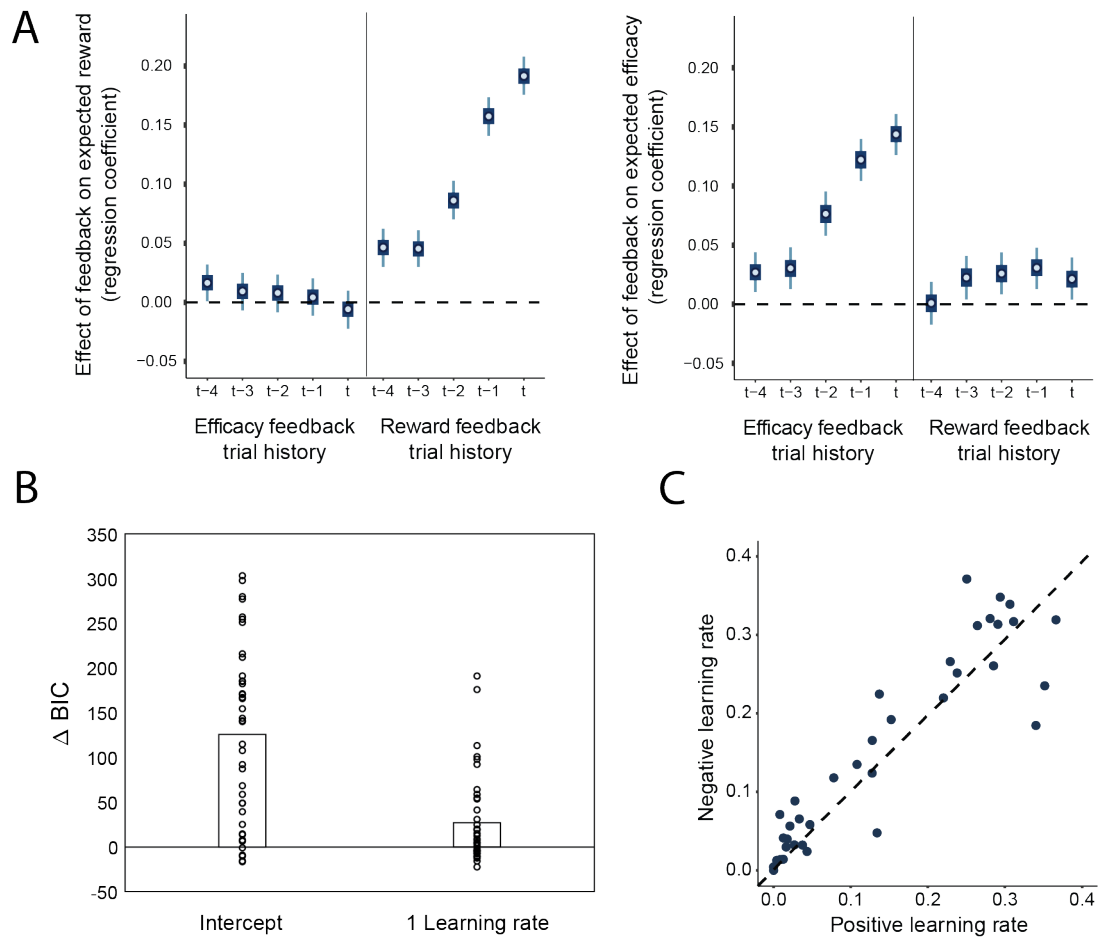


Figure S8. Reward rate learning – model-based analyses for Study 1. **A.** Regression weights for the influence of the current (t) and previous feedbacks on the subjective estimates of reward rate (left) and efficacy (right). Reward rate estimates are strongly predicted by the previous reward feedbacks, but not efficacy feedbacks. The reverse is true for the subjective efficacy estimates. Error bars represent 50% and 95% highest density intervals. **B.** Model comparison between the fitted learning models for the reward rate learning model. **C.** Positive and negative learning rate estimates for all subjects for the reward rate learning model. Negative learning rates were numerically, but not statistically larger than positive ones ($b = 0.01$; 95% CrI [-0.07, 0.09]; $p_b < 0 = 0.38$; $BF_{10} = 0.09$)

747 **ERP analyses of the reward rate feedback in Study 1**

748 To investigate the effect of the model-based reward rate estimates on the processing of
749 reward, we performed complementary analyses on the P3b related to the processing of the
750 reward feedback (Figure S9A). We expected that a similar learning mechanism should operate
751 on both reward and efficacy feedback, but that the neural markers of feedback processing should
752 be sensitive to different model-based estimates. Thus, we expected the reward-locked P3b to be
753 sensitive to reward-rate estimates, but not to efficacy estimates.

754 Our results showed larger P3b amplitudes to no reward compared to reward feedback
755 (note that reward rate was approximately .80 and negative feedback less expected overall), and
756 the processing of the reward feedback was influenced by reward-rate estimates, although to
757 lesser extent than for the P3b (Figure S9B; Table S6). Probing the effects of learning directly, we
758 found robust effects of the prediction errors and learning rates (Figure S9C). The P3b to reward
759 feedback was not influenced by the model-based efficacy estimates in either of these analyses.

760

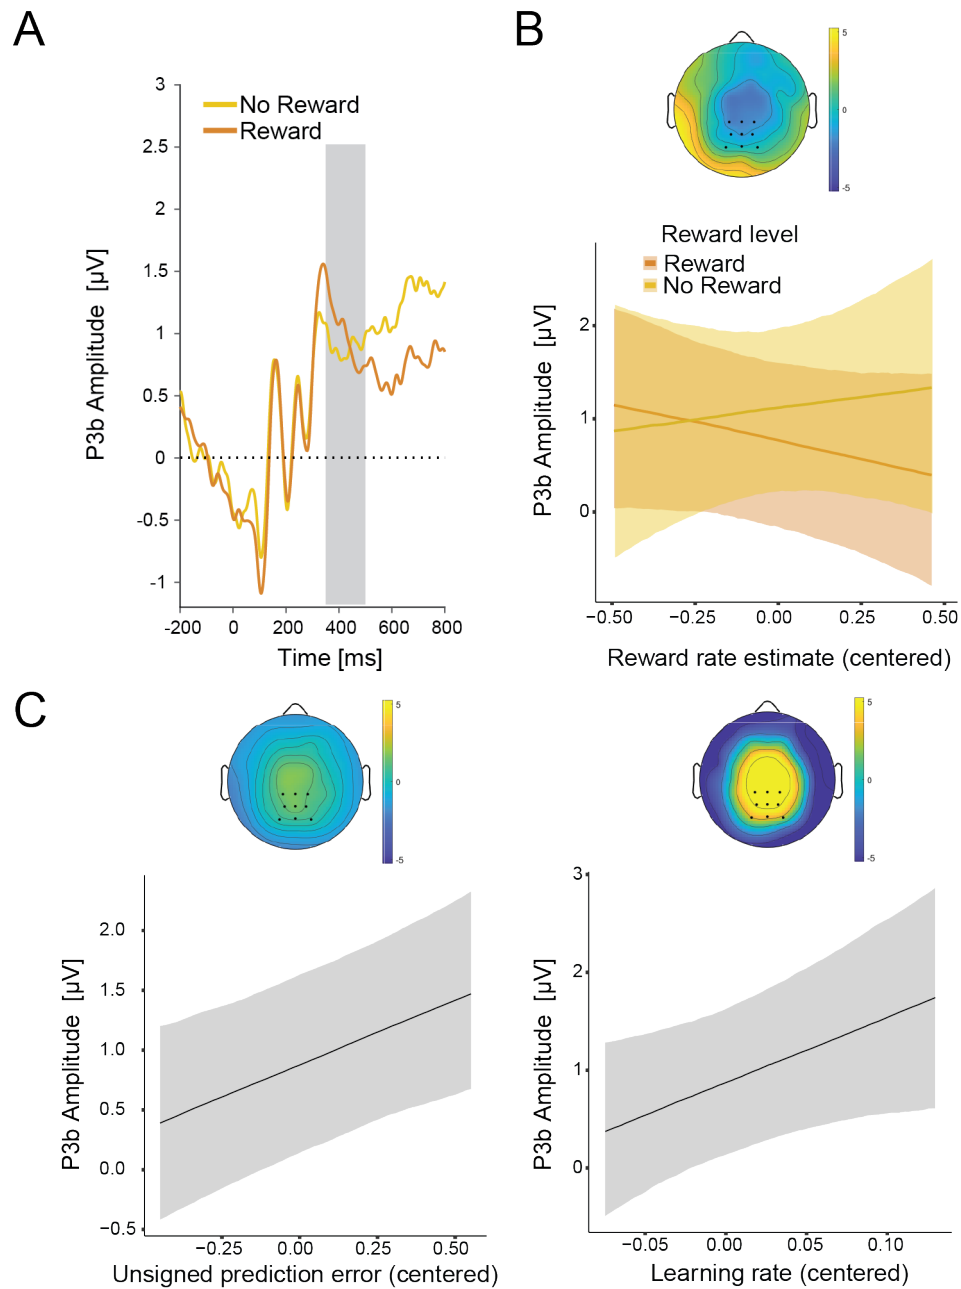


Figure S9. Learning reward rate estimate from reward feedback in Study 1. **A.** ERP average for the P3b locked to the onset of reward feedback. **B.** Interaction between reward feedback and the model-based reward rate estimate. **C.** The effects of unsigned prediction errors (left) and learning rates (right) on the P3b.

Table S6

Study 1 regression weights for the models predicting the P3b to reward feedback.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability (p<0)</i>	<i>BF₁₀</i>
<i>Efficacy feedback</i>				
Intercept	0.93	0.08, 1.74	0.02	
Model-based efficacy	0.45	-0.45, 1.32	0.16	0.15
Reward feedback type (R-NR)	-0.35	-0.90, 0.18	0.90	0.12
Model-based reward rate	-0.15	-1.55, 1.23	0.59	0.14
Model-based reward × Reward feedback type	-1.26	-3.73, 1.19	0.85	0.41
Model-based efficacy slope for No Reward feedback	0.48	-1.34, 2.34	0.87	0.23
Model-based efficacy slope for Reward feedback	-0.78	-1.96, 0.39	0.32	0.21
<i>Predictions errors and learning rates</i>				
Intercept	0.89	0.15, 1.64	0.01	
Unsigned prediction error	1.09	0.38, 1.81	0.00	4.82
Learning rate	6.68	-0.07, 12.99	0.03	4.72
Model-based efficacy	0.39	-0.47, 1.24	0.19	0.13
Unsigned prediction error × Learning rate	0.94	-8.06, 9.71	0.42	0.92

763

764

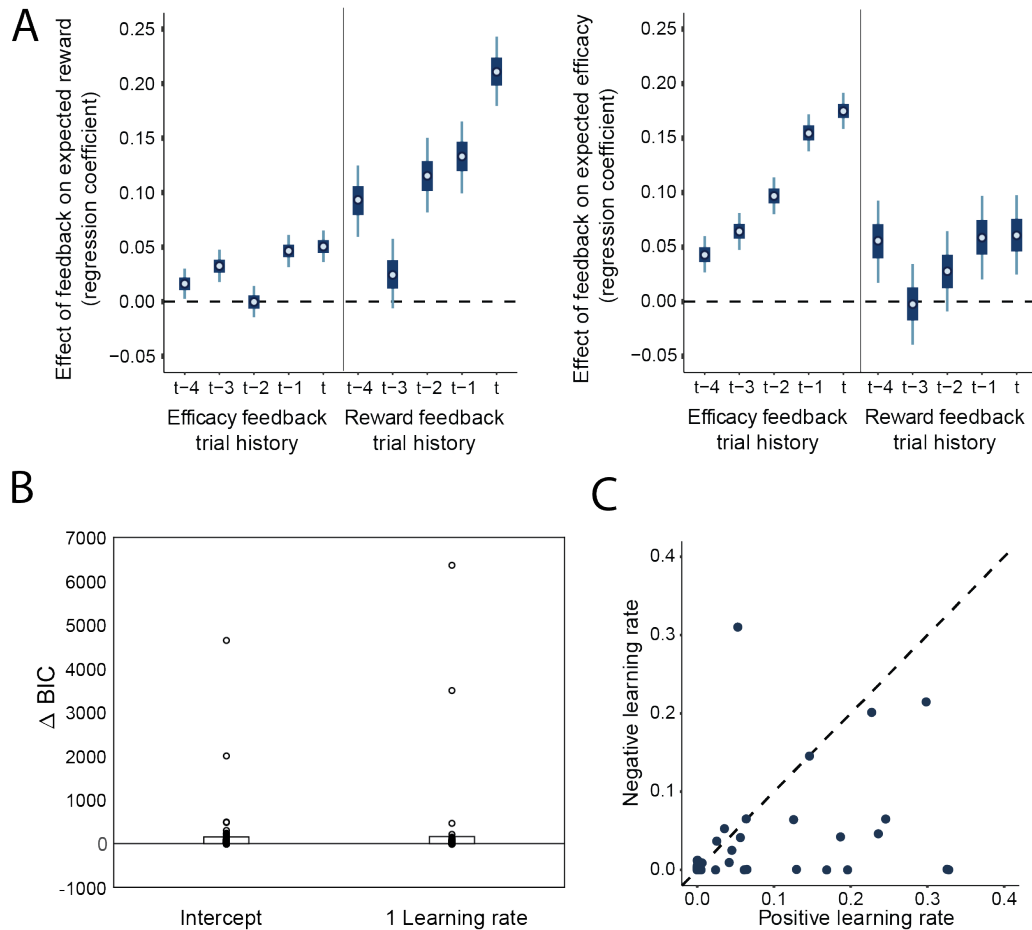


Figure S10. Efficacy and reward learning – model-based analyses for Study 2. **A.** Regression weights for the influence of the current (t) and previous feedbacks on the subjective estimates of reward rate (left) and efficacy (right). Reward rate estimates are strongly influenced by the previous reward feedbacks, and the reverse is true for the subjective efficacy estimates. Error bars represent 50% and 95% highest density intervals. **B.** Model comparison between the fitted learning models for the reward learning model. **C.** Positive and negative learning rate estimates for all subjects for the reward learning model. Positive learning rates were numerically, but not statistically larger than negative ones ($b = 0.02$; 95% CrI [-0.01, 0.05]; $p_{b < 0} = 0.11$; $BF_{10} = 0.06$)

Table S7

Study 2 regression weights for the models predicting subjective efficacy and reward rate estimates based on efficacy and reward feedback 5 trials back.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability ($p < 0$)</i>	<i>BF₁₀</i>
<i>Subjective efficacy</i>				
Intercept	0.43	0.37, 0.49	0.00	
T1 Performance-Based - Random	0.17	0.16, 0.19	0.00	>100
T2 Performance-Based - Random	0.15	0.13, 0.18	0.00	>100
T3 Performance-Based - Random	0.10	0.08, 0.12	0.00	>100
T4 Performance-Based - Random	0.06	0.04, 0.08	0.00	>100
T5 Performance-Based - Random	0.04	0.02, 0.06	0.00	>100
T1 Reward - No Reward	0.06	0.02, 0.10	0.00	5.42
T2 Reward - No Reward	0.06	0.01, 0.11	0.01	2.23
T3 Reward - No Reward	0.03	-0.02, 0.07	0.11	0.23
T4 Reward - No Reward	0.00	-0.05, 0.04	0.54	0.11
T5 Reward - No Reward	0.06	0.01, 0.10	0.01	1.69
<i>Subjective reward rate</i>				
Intercept	0.36	0.29, 0.43	0.00	
T1 Reward - No Reward	0.21	0.17, 0.25	0.00	>100
T2 Reward - No Reward	0.13	0.09, 0.17	0.00	>100
T3 Reward - No Reward	0.12	0.07, 0.16	0.00	>100
T4 Reward - No Reward	0.03	-0.01, 0.06	0.09	0.21
T5 Reward - No Reward	0.09	0.05, 0.13	0.00	>100
T1 Performance-Based - Random	0.05	0.03, 0.07	0.00	0.06
T2 Performance-Based - Random	0.05	0.03, 0.06	0.00	0.05
T3 Performance-Based - Random	0.00	-0.02, 0.02	0.51	0.05
T4 Performance-Based - Random	0.03	0.02, 0.05	0.00	0.08
T5 Performance-Based - Random	0.02	0.00, 0.03	0.03	0.20

Table S8

Study 2 regression weights for the models predicting reaction times and accuracy based on efficacy and reward rate estimates. Bayes Factors calculated only for the parameters with informative priors. Posteriors from Study 1 were used as priors for the effects of congruency, efficacy, and reward.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability (p<0)</i>	<i>BF₀₁</i>
<i>Reaction times</i>				
Intercept	643.02	625.42, 661.81	0.00	
Incongruent-Congruent	47.48	40.69, 55.22	0.00	0.02
Model-based efficacy	-10.25	-19.86, -0.20	0.98	2.33
Model-based reward	-0.17	-21.15, 20.34	0.51	1.24
Interval length	-1.76	-2.73, -0.83	1.00	
Interval congruency	-5.34	-6.69, -4.00	1.00	
<i>Accuracy</i>				
Intercept	2.35	2.16, 2.52	0.00	
Incongruent-Congruent	-0.48	-0.57, -0.39	1.00	>100
Model-based efficacy	0.23	0.03, 0.43	0.01	>100
Model-based reward	-0.29	-0.57, -0.01	0.98	1.07
Interval length	-0.13	-0.16, -0.10	1.00	
Interval congruency	0.04	0.00, 0.08	0.03	

Table S9
Drift Diffusion model estimates.

<i>Parameter</i>	<i>Estimate</i>	<i>95% Credible interval</i>	<i>Posterior probability (p<0)</i>
<hr/>			
<i>Drift rate</i>			
Congruent-Incongruent	0.82	0.74, 0.92	0.00
Model-based efficacy	0.29	0.14, 0.40	0.00
Model-based reward	0.20	-0.39, 0.04	0.50
<hr/>			
<i>Threshold</i>			
Model-based efficacy	-0.00	-0.04, 0.04	0.57
Model-based reward rate	-0.04	-0.14, 0.05	0.81
<hr/>			

766

767

768

Acknowledgements

769 This work was supported by the Special Research Fund (BOF) of Ghent University [grant
770 #01D02415] (I.G.), the Research Foundation Flanders (FWO) travel grant [grant #V432718N]
771 (I.G.), a Center of Biomedical Research Excellence grant P20GM103645 from the National
772 Institute of General Medical Sciences (A.S.), the Alfred P. Sloan Foundation Research Fellowship
773 in Neuroscience (A.S.), and an NSF Graduate Research Fellowship (M.P.F). The funding sources
774 were not involved in the study design; collection, analysis, and interpretation of data; writing of
775 the report; and decision to submit the article for publication. We would like to thank Natalie
776 Knowles and Hattie Xu for help with data collection for Study 1, Peyton Strong for help with data
777 collection for Study 2, Carolyn Dean Wolf and Liz Cory for help with programming the Study 1
778 task, Harrison Ritz for help with constructing the efficacy drifts, and Xiamin Leng for advice on
779 fitting the drift-diffusion models.

780

781

Data availability

782 Analysis scripts are available on this link: https://github.com/igrahek/LFXC_EEG_2022.git Please
783 contact the authors for raw data.

784

References

- 785 Atkinson JW, Feather NT, Others. 1966. A theory of achievement motivation. Wiley New York.
- 786 Balleine, B. W., & O'doherty, J. P. 2010. Human and rodent homologies in action control:
787 corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology*,
788 35(1), 48-69.
- 789 Bejjani C, Zhang Z, Egner T. 2018. Control by association: Transfer of implicitly primed
790 attentional states across linked stimuli. *Psychon Bull Rev.* 25(2):617–626. doi:10.3758/s13423-
791 018-1445-6.
- 792 Botvinick MM, Cohen JD. 2014. The computational and neural basis of cognitive control:
793 Charted territory and new frontiers. *Cogn Sci.* 38(6):1249–1285. doi:10.1111/cogs.12126.
- 794 Brainard DH. 1997. The Psychophysics Toolbox. *Spat Vis.* 10(4):433–436.
- 795 Bugg JM, Jacoby LL, Chanani S. 2011. Why it is too early to lose control in accounts of item-
796 specific proportion congruency effects. *J Exp Psychol Hum Percept Perform.* 37(3):844–859.
797 doi:10.1037/a0019957.
- 798 Bürkner P-C. 2016. brms: An R package for Bayesian multilevel models using Stan. *J Stat*
799 *Softw.* 80(1):1–28.
- 800 Bürkner P-C. 2017. Advanced Bayesian Multilevel Modeling with the R Package brms.
801 arXiv:170511123.
- 802 Bustamante L, Lieder F, Musslick S, Shenhav A, Cohen J. 2021 Jan 6. Learning to Overexert
803 Cognitive Control in a Stroop Task. *Cogn Affect Behav Neurosci.* doi:10.3758/s13415-020-
804 00845-x.
- 805 Carlebach N, Yeung N. 2020. Flexible use of confidence to guide advice requests. *PsyArXiv.*
- 806 Carpenter B, Gelman A, Hoffman M, Lee D, Goodrich B, Betancourt M, Brubaker MA, Guo J,
807 Li P, Riddell A. 2016. Stan: A probabilistic programming language. *J Stat Softw.* 2(20):1–37.
- 808 Chambon V, Théro H, Vidal M, Vandendriessche H, Haggard P, Palminteri S. 2020. Information
809 about action outcomes differentially affects learning from self-determined versus imposed
810 choices. *Nat Hum Behav.* doi:10.1038/s41562-020-0919-5.
- 811 Chase HW, Swainson R, Durham L, Benham L, Cools R. 2011. Feedback-related negativity
812 codes prediction error but not behavioral adjustment during probabilistic reversal learning. *J*
813 *Cogn Neurosci.* 23(4):936–946. doi:10.1162/jocn.2010.21456.
- 814 Chiu, Y. C., & Egner, T. (2019). Cortical and subcortical contributions to context-control
815 learning. *Neurosc & Biobeh Rev*, 99: 33-41.
- 816 Cohen A, Nussenbaum K, Dorfman HM, Gershman SJ, Hartley CA. 2020. The rational use of
817 causal inference to guide reinforcement learning strengthens with age. *Npj Sci Learn.*

- 818 Collins AGE, Frank MJ. 2013. Cognitive control over learning: creating, clustering, and
819 generalizing task-set structure. *Psychol Rev.* 120(1):190–229. doi:10.1037/a0030852.
- 820 Collins AGE, Frank MJ. 2014. Opponent actor learning (OpAL): Modeling interactive effects of
821 striatal dopamine on reinforcement learning and choice incentive. *Psychol Rev.* 121(3):337–66.
822 doi:10.1037/a0037015.
- 823 Delorme A, Makeig S. 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG
824 dynamics including independent component analysis. *J Neurosci Methods.* 134:9–21.
825 doi:10.1016/j.jneumeth.2003.10.009.
- 826 Dickinson A, Balleine B. 1995. Motivational Control of Instrumental Action. *Curr Dir Psychol*
827 *Sci.* 4(5):162–167. doi:10.1111/1467-8721.ep11512272.
- 828 Dorfman, H. M., Bhui, R., Hughes, B. L., & Gershman, S. J. 2019. Causal inference about good
829 and bad outcomes. *Psychological science*, 30(4): 516-525.
- 830 Dorfman, H. M., Tomov, M. S., Cheung, B., Clarke, D., Gershman, S. J., & Hughes, B. L.
831 (2021). Causal inference gates corticostriatal learning. *Journal of Neuroscience*, 41(32): 6892-
832 6904.
- 833 Eldar E, Hauser TU, Dayan P, Dolan RJ. 2016. Striatal structure and function predict individual
834 biases in learning to avoid pain. *Proc Natl Acad Sci U S A.* 113(17):4812–4817.
835 doi:10.1073/pnas.1519829113.
- 836 Fischer AG, Ullsperger M. 2013. Real and fictive outcomes are processed differently but
837 converge on a common adaptive mechanism. *Neuron.* 79(6):1243–1255.
838 doi:10.1016/j.neuron.2013.07.006.
- 839 Frömer R, Lin H, Dean Wolf CK, Inzlicht M, Shenhav A. 2021a. Expectations of reward and
840 efficacy guide cognitive control allocation. *Nat Commun.* 12(1). doi:10.1038/s41467-021-21315-
841 z.
- 842 Frömer, R., Nassar, M. R., Bruckner, R., Stürmer, B., Sommer, W., & Yeung, N. 2021b.
843 Response-based outcome predictions and confidence regulate feedback processing and learning.
844 *Elife*, 10, e62825.
- 845 Frömer R, Maier M, Rahman RA. 2018. Group-level EEG-processing pipeline for flexible single
846 trial-based analyses including linear mixed models. *Front Neurosci.* 12(48).
847 doi:10.3389/fnins.2018.00048.
- 848 Frömer R, Stürmer B, Sommer W. 2016. (Don't) Mind the effort: Effects of contextual
849 interference on ERP indicators of motor preparation. *Psychophysiology.* 53(10):1577–1586.
850 doi:10.1111/psyp.12703.
- 851 Garrett N, Daw ND. 2020. Biased belief updating and suboptimal choice in foraging decisions.
852 *Nat Commun.* 11(1):1–12. doi:10.1038/s41467-020-16964-5.
- 853 Gelman A, Rubin DB. 1992. Inference from Iterative Simulation Using Multiple Sequences. *Stat*
854 *Sci.* 7(4):457–472. doi:10.1214/ss/1177011136.

- 855 Gläscher J, Daw N, Dayan P, O’Doherty JP. 2010. States versus rewards: Dissociable neural
856 prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*.
857 66(4):585–595. doi:10.1016/j.neuron.2010.04.016.
- 858 Grahek I, Shenhav A, Musslick S, Krebs RM, Koster EHW. 2019. Motivation and cognitive
859 control in depression. *Neurosci Biobehav Rev*. 102(March):371–381.
860 doi:10.1016/j.neubiorev.2019.04.011.
- 861 Grent-’t-Jong T, Woldorff MG. 2007. Timing and sequence of brain activity in top-down control
862 of visual-spatial attention. *PLoS Biol*. 5(1):0114–0126. doi:10.1371/journal.pbio.0050012.
- 863 Holroyd CB, Yeung N. 2012. Motivation of extended behaviors by anterior cingulate cortex.
864 *Trends Cogn Sci*. 16(2):122–128. doi:10.1016/j.tics.2011.12.008.
- 865 Huys QJM, Dayan P. 2009. A Bayesian formulation of behavioral control. *Cognition*.
866 113(3):314–328. doi:10.1016/j.cognition.2009.01.008.
- 867 Ille N, Berg P, Scherg M. 2002. Artifact Correction of the Ongoing EEG Using Spatial Filters
868 Based on Artifact and Brain Signal Topographies. *J Clin Neurophysiol*. 19(2):113–124.
869 doi:10.1097/00004691-200203000-00002.
- 870 Jiang J, Bramão I, Khazenzon A, Wang SF, Johansson M, Wagner AD. 2020. Temporal
871 Dynamics of Memory-guided Cognitive Control and Generalization of Control via Overlapping
872 Associative Memories. *J Neurosci*. 40(10):2343–2356. doi:10.1523/JNEUROSCI.1869-19.2020.
- 873 Jiang J, Heller K, Egner T. 2014. Bayesian modeling of flexible cognitive control. *Neurosci*
874 *Biobehav Rev*. 46(P1):30–43. doi:10.1016/j.neubiorev.2014.06.001.
- 875 Kleiner M, Brainard D, Pelli D. 2007. What’s new in Psychtoolbox-3. In: *Perception 36 ECVF*
876 *Abstract Supplement*.
- 877 Kool W, Botvinick M. 2018. Mental labour. *Nat Hum Behav*. doi:10.1038/s41562-018-0401-9.
- 878 Krebs, R. M., Boehler, C. N., Roberts, K. C., Song, A. W., & Woldorff, M. G. 2012. The
879 involvement of the dopaminergic midbrain and cortico-striatal-thalamic circuits in the integration
880 of reward prospect and attentional task demands. *Cereb cortex*, 22(3): 607-615.
881
- 882 Kruschke JK. 2013. Bayesian estimation supersedes the T test. *J Exp Psychol Gen*. 142(2):573–
883 588. doi:10.1037/a0029177.
- 884 Langdon AJ, Sharpe MJ, Schoenbaum G, Niv Y. 2018. Model-based predictions for dopamine.
885 *Curr Opin Neurobiol*. 49:1–7. doi:10.1016/j.conb.2017.10.006.
- 886 Lefebvre G, Lebreton M, Meyniel F, Bourgeois-Gironde S, Palminteri S. 2017. Behavioural and
887 neural characterization of optimistic reinforcement learning. *Nat Hum Behav*. 1(4):1–9.
888 doi:10.1038/s41562-017-0067.
- 889 Leng, X., Yee, D., Ritz, H., & Shenhav, A. 2021. Dissociable influences of reward and
890 punishment on adaptive cognitive control. *PLoS computational biology*, 17(12), e1009737.

- 891 Leotti LA, Iyengar SS, Ochsner KN. 2010. Born to choose: The origins and value of the need for
892 control. *Trends Cogn Sci.* 14(10):457–463. doi:10.1016/j.tics.2010.08.001.
- 893 Lieder F, Shenhav A, Musslick S, Griffiths TL. 2018. Rational metareasoning and the plasticity
894 of cognitive control. *Plos Comput Biol.* 14(4). doi:10.13140/RG.2.2.24500.14721.
- 895 Ligneul, R., Mainen, Z. F., Ly, V., & Cools, R. 2022. Stress-sensitive inference of task
896 controllability. *Nat Hum Beh.* 1-11.
- 897 Liljeholm M, Tricomi E, O’Doherty JP, Balleine BW. 2011. Neural Correlates of Instrumental
898 Contingency Learning: Differential Effects of Action-Reward Conjunction and Disjunction. *J*
899 *Neurosci.* 31(7):2474–2480. doi:10.1523/JNEUROSCI.3354-10.2011.
- 900 Lohse KR, Miller MW, Daou M, Valerius W, Jones M. 2020. Dissociating the contributions of
901 reward-prediction errors to trial-level adaptation and long-term learning. *Biol Psychol.*
902 149(January 2019). doi:10.1016/j.biopsycho.2019.107775.
- 903 Ly V, Wang KS, Bhanji J, Delgado MR. 2019. A Reward-Based Framework of Perceived
904 Control. *Front Neurosci.* 13(February):1–11. doi:10.3389/fnins.2019.00065.
- 905 Maier SF, Seligman MEP. 2016. Learned helplessness at fifty: Insights from neuroscience.
906 *Psychol Rev.* 123(4):1–19. doi:10.1037/rev0000033.
- 907 Manohar SG, Chong TTJ, Apps MAJ, Batla A, Stamelou M, Jarman PR, Bhatia KP, Husain M.
908 2015. Reward Pays the Cost of Noise Reduction in Motor and Cognitive Control. *Curr Biol.*
909 25(13):1707–1716. doi:10.1016/j.cub.2015.05.038.
- 910 Manohar SG, Finzi RD, Drew D, Husain M. 2017. Distinct Motivational Effects of Contingent
911 and Noncontingent Rewards. *Psychol Sci.* 28(7):1016–1026. doi:10.1177/0956797617693326.
- 912 Morris, R. W., Dezfouli, A., Griffiths, K. R., Le Pelley, M. E., & Balleine, B. W. 2022. The
913 neural bases of action-outcome learning in humans. *J Neuro*, 42(17), 3636-3647.
- 914 Moscarello JM, Hartley CA. 2017. Agency and the Calibration of Motivated Behavior. *Trends*
915 *Cogn Sci.* 21(10):725–735. doi:10.1016/j.tics.2017.06.008.
- 916 Nagase AM, Onoda K, Foo JC, Haji T, Akaishi R, Yamaguchi S, Sakai K, Morita K. 2018.
917 Neural mechanisms for adaptive learned avoidance of mental effort. *J Neurosci.* 38(10):2631–
918 2651. doi:10.1523/JNEUROSCI.1995-17.2018.
- 919 Nalborczyk L, Bürkner P-C. 2019. An Introduction to Bayesian Multilevel Models Using brms:
920 A Case Study of Gender Effects on Vowel Variability in Standard Indonesian. *J Speech Lang*
921 *Hear Res.*
- 922 Nassar MR, Bruckner R, Frank MJ. 2019. Statistical context dictates the relationship between
923 feedback-related EEG signals and learning. *eLife.* 8:1–26. doi:10.7554/eLife.46975.
- 924 Nigbur, R., Schneider, J., Sommer, W., Dimigen, O., & Stürmer, B. (2015). Ad-hoc and context-
925 dependent adjustments of selective attention in conflict control: An ERP study with visual
926 probes. *NeuroImage*, 107: 76-84.

- 927 Niv Y, Edlund JA, Dayan P, O’Doherty JP. 2012. Neural prediction errors reveal a risk-sensitive
928 reinforcement-learning process in the human brain. *J Neurosci.* 32(2):551–562.
929 doi:10.1523/JNEUROSCI.5498-10.2012.
- 930 Norton, K. G., & Liljeholm, M. 2020. The Rostrolateral Prefrontal Cortex Mediates a Preference
931 for High-Agency Environments. *J Neurosci*, 40(22): 4401-4409.
- 932 Otto AR, Daw ND. 2019. The opportunity cost of time modulates cognitive effort.
933 *Neuropsychologia.* 123:92–105. doi:10.1016/j.neuropsychologia.2018.05.006.
- 934 Parro, C., Dixon, M. L., & Christoff, K. 2018. The neural basis of motivational influences on
935 cognitive control. *Human Brain Mapping*, 39(12), 5097–5111
- 936 Pelli DD. 1997. The VideoToolbox software for visual psychophysics: transforming numbers
937 into movies. *Spat Vis.* 10(4):437–442.
- 938 R Core Team. 2017. R: A Language and Environment for Statistical Computing.
- 939 Ritz, H., Leng, X., & Shenhav, A. (2022). Cognitive control as a multivariate optimization
940 problem. *Journal of Cognitive Neuroscience*, 1-23.
- 941 Rutledge RB, Skandali N, Dayan P, Dolan RJ. 2014. A computational and neural model of
942 momentary subjective well-being. *Proc Natl Acad Sci.* 111(33):12252–12257.
943 doi:10.1073/pnas.1407535111.
- 944 Schevernels H, Krebs RM, Santens P, Woldorff MG, Boehler CN. 2014a. Task preparation
945 processes related to reward prediction precede those related to task-difficulty expectation.
946 *NeuroImage.* 84:639–647. doi:10.1016/j.neuroimage.2013.09.039.
- 947 Schevernels H, Krebs RM, Santens P, Woldorff MG, Boehler CN. 2014b. Task preparation
948 processes related to reward prediction precede those related to task-difficulty expectation.
949 *NeuroImage.* 84:639–647. doi:10.1016/j.neuroimage.2013.09.039.
- 950 Schiffer AM, Siletti K, Waszak F, Yeung N. 2017. Adaptive behaviour and feedback processing
951 integrate experience and instruction in reinforcement learning. *NeuroImage.* 146:626–641.
952 doi:10.1016/j.neuroimage.2016.08.057.
- 953 Shenhav A, Botvinick M, Cohen J. 2013. The expected value of control: An integrative theory of
954 anterior cingulate cortex function. *Neuron.* 79(2):217–240. doi:10.1016/j.neuron.2013.07.007.
- 955 Shenhav A, Musslick S, Lieder F, Kool W, Griffiths TL, Cohen JD, Botvinick MM. 2017.
956 Toward a Rational and Mechanistic Account of Mental Effort. *Annu Rev Neurosci.* 40(1):99–
957 124. doi:10.1146/annurev-neuro-072116-031526.
- 958 Shenhav A, Prater Fahey M, Grahek I. 2021. Decomposing the motivation to exert mental effort.
959 *Curr Dir Psych Sci.* 30(4), 307-314.
- 960 Sutton RS, Barto AG. 2018. Reinforcement Learning: An Introduction. MIT Press.

- 961 Vassena, E., Silvetti, M., Boehler, C. N., Achten, E., Fias, W., & Verguts, T. 2014. Overlapping
962 neural systems represent cognitive effort and reward anticipation. *PLoS One*, 9(3).
- 963 Verbeke P, Verguts T. 2019. Learning to synchronize: How biological agents can couple neural
964 task modules for dealing with the stability-plasticity dilemma. *PLoS Comput Biol*. 15(8):1–25.
965 doi:10.1371/journal.pcbi.1006604.
- 966 Verguts T, Vassena E, Silvetti M. 2015. Adaptive effort investment in cognitive and physical
967 tasks: a neurocomputational model. *Front Behav Neurosci*. 9(March).
968 doi:10.3389/fnbeh.2015.00057.
- 969 Wagenmakers EJ, Lodewyckx T, Kuriyal H, Grasman R. 2010. Bayesian hypothesis testing for
970 psychologists: A tutorial on the Savage-Dickey method. *Cognit Psychol*. 60(3):158–189.
971 doi:10.1016/j.cogpsych.2009.12.001.
- 972 Yeung N, Sanfey AG. 2004. Independent coding of reward magnitude and valence in the human
973 brain. *J Neurosci*. 24(28):6258–6264. doi:10.1523/JNEUROSCI.4537-03.2004.
- 974
- 975