








Stress-sensitive inference of task controllability

Romain Ligneul^{1,2}  , Zachary F. Mainen¹, Verena Ly^{2,3,5}  and Roshan Cools^{2,4,5}  

Estimating the controllability of the environment enables agents to better predict upcoming events and decide when to engage controlled action selection. How does the human brain estimate controllability? Trial-by-trial analysis of choices, decision times and neural activity in an explore-and-predict task demonstrate that humans solve this problem by comparing the predictions of an ‘actor’ model with those of a reduced ‘spectator’ model of their environment. Neural blood oxygen level-dependent responses within striatal and medial prefrontal areas tracked the instantaneous difference in the prediction errors generated by these two statistical learning models. Blood oxygen level-dependent activity in the posterior cingulate, temporoparietal and prefrontal cortices covaried with changes in estimated controllability. Exposure to inescapable stressors biased controllability estimates downward and increased reliance on the spectator model in an anxiety-dependent fashion. Taken together, these findings provide a mechanistic account of controllability inference and its distortion by stress exposure.

Influential theories suggest that the human brain navigates its environment by building predictive models of the world, which in turn fuel cognitive processes such as directed exploration, goal-directed decisions and forward planning^{1–3}. While these internal models can take diverse mathematical forms, their efficiency always depends on the use of task-relevant and cost-efficient state spaces^{4–6}. Most often, these state spaces are state–action spaces in which the actions of the agent actively contribute to the prediction of upcoming events. For example, a driver must take into account the movement of their hands to predict the future position of their car. By contrast, a passenger worried for their safety should ignore their own hands and instead focus on the hands of the driver to anticipate potential hazards.

Determining whether an environment is controllable or not is key to deciding the extent to which one’s actions should influence the prediction process since only controllable environments afford causal influence over state transitions. Controllable contexts thus prompt the use of ‘actor’ models including one’s own actions as predictors, whereas uncontrollable contexts prompt the use of simpler, ‘spectator’ models linking past and future states of the environment. By gating the causal influence of action selection, controllability probably plays a central role in the engagement of elaborate action selection mechanisms. Supporting this idea, it is well established that prior exposure to controllable contexts promotes proactive and goal-directed strategies in a variety of cognitive tasks^{7,8}. Conversely, the lack of perceived control over events, especially stressful ones, constitutes a well-established correlate and a potential predictor of prevalent psychiatric disorders involving an increased influence of reactive and habitual behaviours, such as depression, anxiety, post-traumatic stress or obsessive–compulsive disorders^{9–13}.

Numerous studies have shown that exposure to uncontrollable stressors can induce a state of learned helplessness characterized by the generalization of passive reactions to subsequent challenges^{7,8}. Evidence indicates that this maladaptive state largely depends on functional changes within the medial prefrontal cortex (mPFC) and the serotonin system^{14–18}. In humans, a handful of neuroimaging experiments have further suggested that the anterior insula and

cingulate cortex contribute to the detrimental effects of uncontrollable stressors^{19,20}. Beyond stress induction studies, the sense of being in control of one’s own actions and their outcomes is known to modulate haemodynamic responses parietal and prefrontal cortices^{21–23}, and the right temporoparietal junction (TPJ) was found to track the divergence of action–outcome transitions, a feature of controllable environments²⁴.

Yet, little is known about the algorithms by which the brain estimates dynamically the degree to which a task is controllable. A general strategy is to estimate controllability by computing the causal effect the agent’s own actions have over the environment. Formally, a task can thus be deemed controllable when the transfer entropy (TE)—a generalization of Granger causality to nonlinear and discrete systems—linking state and action time series is positive^{25,26}. By comparing the entropy of observed states given previous states and actions $[H(S'|S,A)]$ with the entropy of observed states given previous states only $[H(S'|S)]$, this information-theoretic quantity isolates the effective causal influence of actions over state transitions (Fig. 1a and Supplementary Note 1). In the vocabulary of causal mediation analysis²⁷, positive TE values entail the existence of a natural direct effect linking actions to future states of the environment (Supplementary Note 2). Here, we develop a computational model that tracks a dynamic approximation of TE and use it to shed light on the cognitive and neural mechanisms supporting the ability to infer task controllability and adapt behaviour accordingly.

Based on this information-theoretic formalism, we designed an explore-and-predict task that allowed us to manipulate controllability and assess the resulting changes in terms of subjective controllability and prediction accuracy. This new task was first used in behavioural ($n=50$) and functional magnetic resonance imaging (fMRI) ($n=32$) experiments aimed at (1) demonstrating that humans infer task controllability by estimating an approximation of TE, (2) establishing the dissociation of spectator and actor models predicted by the TE hypothesis at the behavioural and neural levels and (3) unravelling the neural substrates underlying the representation of controllability itself and its influences on behaviour. In a subsequent stress experiment ($n=54$), we exposed participants to

¹Champalimaud Research, Champalimaud Centre for the Unknown, Lisbon, Portugal. ²Donders Institute for Brain, Cognition and Behaviour, Centre for Cognitive Neuroimaging, Radboud University, Nijmegen, the Netherlands. ³Department of Clinical Psychology, Institute of Psychology Office, Leiden Institute for Brain and Cognition, Leiden University, Leiden, the Netherlands. ⁴Department of Psychiatry, Radboud University Medical Centre, Nijmegen, the Netherlands. ⁵These authors contributed equally: Verena Ly, Roshan Cools. ✉e-mail: romain.ligneul@research.fchampalimaud.org; roshan.cools@fcdonders.ru.nl




Fig. 1 | Theoretical framework and experimental protocol. **a**, Controllability can be inferred from TE, an information-theoretic measure quantifying the extent to which a time series causally influences another one. **b**, Time course of a novel explore-and-predict task divided in short mini-blocks. Each mini-block consists of a series of exploratory trials (violet) followed by two counterfactual prediction trials (green) used to assess learning and subjective controllability. **c**, Representation of the two uncontrollable rules (U1, U2) and the two controllable (C1, C2) rules, which alternate covertly to govern the evolution of the environment. Note that rule C2 was the only rule allowing state repetition, a feature that was taken into account in our analyses. **d**, Simulations showing the dissociation of controllability, as indexed by TE, and predictability, as indexed by the mutual information (MI) shared between successive state-action pairs (random exploration policy). **e**, Simulations under controllable rule C2 showing that TE requires exploration to be used as a proxy for controllability. In the complete absence of exploration, the conditional entropies $H(S'|S)$ and $H(S'|S,A)$ are both null because the agent maintains itself indefinitely in a single preferred state (Supplementary Fig. 1c). **f**, Synthetic overview of the algorithm able to derive an online approximation of TE (termed Ω) by comparing on each trial the transition probabilities of an actor (SAS') and a spectator (SS') model of the world. By thresholding Ω , the algorithm could in turn arbitrate between spectator and actor models when making predictions depending on current controllability estimates (ω). This architecture was compared with other controllability estimation schemes and a standard model-based learning model tracking only SAS' transitions.

either uncontrollable or controllable electric shocks before administering the explore-and-predict task to (4) provide causal evidence supporting a dissociation of the spectator and actor models and (5) test whether learned helplessness can be characterized by an increased reliance on the former relative to the latter.

Results

Experimental paradigm and computational model. Healthy human participants were invited to explore an abstract environment composed of three states (square, circle and triangle) and three actions (yellow, blue and magenta). A hidden transition rule always determined upcoming states, dependent on either the action of the participants (controllable rules, C) or the previous state only (uncontrollable rules, U) (Fig. 1b). The transition rules were probabilistic and reversed covertly, so participants needed to explore and accumulate evidence to tell which rule was operative. From time to time, the participant was asked to predict the most likely upcoming state given a state-action pair (for example, ‘blue’ action in ‘circle’ state), and its counterfactual (for example, ‘yellow’ action in ‘circle’ state). This procedure yielded a direct yet implicit assessment of their subjective sense of controllability because counterfactual predictions should only differ in controllable contexts, where selected actions determine upcoming states (Fig. 1c). A novel and distinguishing feature of our task is that controllability varied independently of uncertainty (Fig. 1d), a methodological improvement over earlier paradigms where the two constructs

covary systematically^{13,24,28}. Another key difference from previous studies is that we did not include any reinforcers: participants were merely instructed to explore their environment to perform accurate predictions when asked to. Here, controllability estimation can interact with but does not depend on reward and punishment processing^{7,29,30}, the only requirement being to maintain a minimal level of exploration, or noise in the action selection process (Fig. 1e and Supplementary Note 1).

To untangle the mechanisms of controllability estimation in this task, we designed a computational architecture for dynamically tracking an approximation of TE (see Supplementary Fig. 1a and Methods for a detailed description). Paralleling the standard computation of TE, two sets of transition probabilities were monitored, one corresponding to an ‘actor’ model (tracking state-action-state transition, SAS') and the other to a ‘spectator’ model (tracking state-state transition, SS' , Fig. 1f.i). Following each transition, an approximation of TE (termed Ω) was updated in proportion to the difference between ‘actor’ and ‘spectator’ transition probabilities $p_{SAS'} - p_{SS'}$ (Fig. 1f.ii). Intuitively, this difference term can be understood as an instantaneous causality signal, reflecting how likely the last state transition towards S' was due to the influence of action A rather than state S . By integrating $p_{SAS'} - p_{SS'}$ over time, Ω thus reflects the causal influence of actions on recent state transitions (Supplementary Fig. 1b).

This causality signal Ω is at the core of the proposed algorithm, which arbitrates between the actor and the spectator model when




Fig. 2 | Behavioural performance. **a**, Accuracy in the prediction trials was above chance for both conditions in each of the experiments (behavioural: $n=50$, $t(49)=11.31$, $P=2.9 \times 10^{-15}$, $d=1.60$, 95% CI 0.20–0.29; fMRI: $n=32$, $t(31)=12.67$, $P=8.5 \times 10^{-14}$, $d=2.24$, 95% CI 0.24–0.33; stress: $n=54$, $t(49)=11.64$, $P=3.1 \times 10^{-16}$, $d=1$, 95% CI 0.20–0.28; Supplementary Table 1). Perm., permutation. **b**, Rapid recovery of predictive accuracy for all reversal types (first pair of predictions after reversal: $n=50$, $t(49)=6.64$, $P=2.4 \times 10^{-9}$, $d=0.94$, 95% CI 0.11–0.20). Accuracies were split by reversal type for visual purpose (U denotes uncontrollable rules, C denotes controllable rules). **c**, Positive correlation between working memory (WM) capacity indexed by a two-back task (Supplementary Methods) and predictive accuracy in the explore-and-predict task for controllable (blue: $n=46$, $\rho=0.52$, $P=1.9 \times 10^{-4}$, 95% CI 0.18–0.67) and uncontrollable (orange, $n=46$, $\rho=0.40$, $P=0.006$, 95% CI 0.13–0.62) contexts. Shaded areas represent s.e.m.

making predictions about upcoming states. Specifically, the relative weight of the actor versus spectator model is set by an arbitrator (termed ω) whose value can be interpreted as an estimate of controllability. Two parameters influence the mapping between Ω and ω : a threshold determining how much causal evidence is required to infer controllability and a slope determining how fast controllability estimates change around that threshold (Fig. 1f,iii). This SAS'–SS'– Ω algorithm was contrasted with a conventional model-based learning algorithm^{2,31} and with two models estimating controllability based on the uncertainty (as indexed by the conditional entropy $H(S'|A,S)$) or the divergence of SAS' transition probabilities bound to different actions (as indexed by the Jensen–Shannon divergence). Importantly, this simpler algorithm could still learn transition probabilities from both uncontrollable and controllable conditions in stable environments, but the lack of controllability-dependent arbitration makes it less efficient in volatile environments alternating rapidly between controllable and uncontrollable rules.

Controllability drives learning and predictive decisions.

Participants performed well on the task. In all experiments, the average prediction accuracy was substantially above chance (Fig. 2a). In the fMRI experiment, for which participants received more training, accuracy was also stable across conditions and time (Supplementary Table 1). Prediction accuracy dropped and then rapidly recovered after covert reversals in transition rules, already exceeding chance levels on the first pair of prediction trials after

reversal (Fig. 2b). Prediction accuracy also correlated positively with working memory capacity as indexed by d' values in a standard two-back task (Fig. 2c), consistent with the engagement of a model-based learning process³².

In line with our prediction that humans solve the task by estimating Ω , Bayesian model comparisons demonstrated that SAS'–SS'– Ω schemes outperformed the conventional model-based learning algorithm (SAS' alone) in all experiments (Fig. 3a and Supplementary Fig. 2a). Simulation analyses confirmed that the model was identifiable and that most of its parameters could be recovered accurately (Supplementary Fig. 2b–d). As expected, the arbitrator ω captured quantitative changes in subjective controllability, indexed by the proneness of participants to predict that different actions would lead to different states in counterfactual prediction trials (Fig. 3b). Critically, the SAS'–SS'– Ω scheme that included an arbitration mechanism accounted better for the dynamics of subjective controllability changes around reversals than did the SAS' model alone (Fig. 3c, correct prediction of subjective controllability: 72.1% versus 66.5%, $N=50$, $z(49)=3.90$, $d=0.64$, $P=9.7 \times 10^{-5}$, 95% confidence interval (CI) 0.004–0.028). As compared with other controllability estimation schemes (that is, SAS'–SS'–H and SAS'–SS'–JS), its arbitrator variable ω also predicted variations in subjective controllability more accurately (Supplementary Fig. 2e,f). The benefits of monitoring controllability are further illustrated by the finding that the likelihood of using the SAS'–SS'– Ω scheme over the SAS scheme increased with accuracy across subjects (Supplementary Fig. 2g). Importantly, since only one controllable rule (and none of the uncontrollable rules) allowed immediate state repetitions (rule C2), state repetition events provided a salient psychological cue that contributed to controllability detection. Accordingly, rule C2 was associated with a higher frequency of subjective controllability responses than rule C1 across the three experiments (behaviour: $N=50$, $82.3 \pm 13.3\%$ versus $75.6 \pm 14.0\%$, $t(49)=4.1$, $P=1.5 \times 10^{-4}$, $d=0.58$, 95% CI 0.03–0.10; fMRI: $N=32$, $88.7 \pm 14.5\%$ versus $69.3 \pm 15.4\%$; $t(31)=6.29$, $P=5.5 \times 10^{-7}$, $d=0.55$, 95% CI 0.13–0.26; stress: $N=27$, $t(53)=4.03$, $P=1.8 \times 10^{-4}$, $d=0.55$, 95% CI 0.04–0.13). On average, participants also chose more frequently the actions that could lead to state repetition whenever rule C2 was active, thereby indicating that they leveraged this feature of our task to refine controllability inferences in all experiments (behavioural: $53 \pm 8\%$, $t(49)=2.6$, $P=0.012$, $d=0.37$, 95% CI 0.007–0.05; fMRI: $51.5 \pm 4.6\%$, $t(31)=1.85$, $P=0.07$, $d=0.33$, 95% CI -0.001 to 0.03; stress: $53.2 \pm 10.3\%$, $t(53)=2.29$, $P=0.025$, $d=0.31$, 95% CI 0.004–0.06). Computational models took this factor into account by allowing prior knowledge about transition rules—as derived from the instruction phase—to constrain the update of actor and spectator models (see Supplementary Information, modelling subsection for more details).

Dissociation of actor and spectator models. Model comparison results are consistent with our proposal that subjects estimate the subjective controllability of an environment by separately tracking and comparing an actor and a spectator model. To further test the dissociation of the actor and spectator models, we used subject-level general linear models (GLMs) to assess trial-by-trial fluctuations of decision times. It is known that decision times slow down following state prediction errors^{33,34}. The large amount of exploratory trials per participant thus allowed us to analyse decision times as a proxy of model updating and to evaluate the extent to which controllability per se influences the speed of action selection (behaviour: 562 ± 163 trials; fMRI: 550 ± 115 ; stress: 519 ± 84). To do so, we extracted the prediction errors derived from both the actor and spectator models (termed $\delta_{SAS'}$ and $\delta_{SS'}$, average correlation 0.22 ± 0.17). We found that both types of prediction errors slowed responding (Fig. 3d) and independently explained variance in decision times (Supplementary Fig. 3). We also observed that in periods of higher estimated

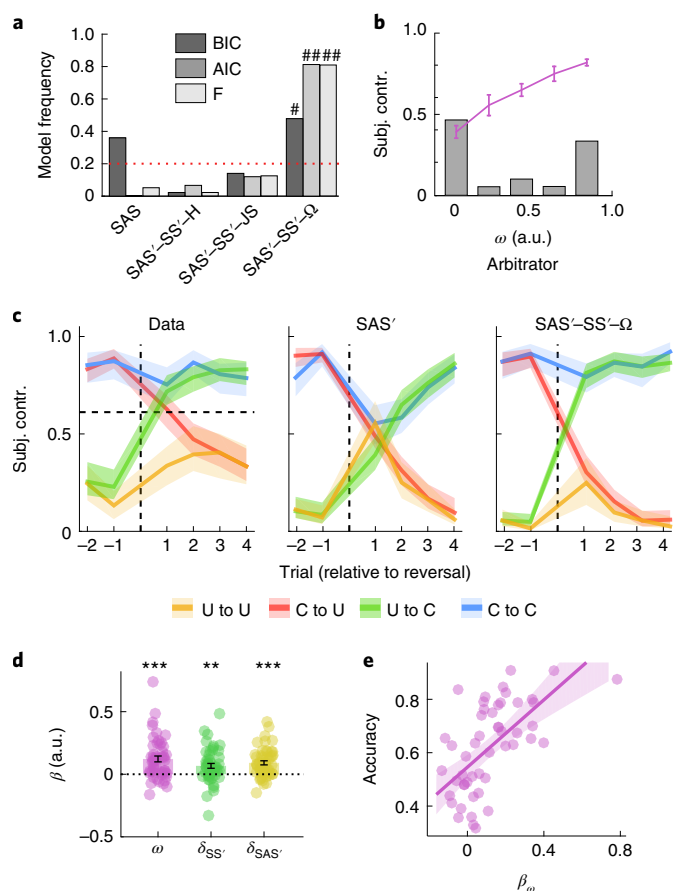


Fig. 3 | Computational modelling and decision times. **a**, Bayesian group model comparison pooled over three experiments showed the advantage of the model estimating controllability using an approximation of TE (SS'-SAS'-Ω) over the actor model (SAS') alone and two models estimating controllability based on the uncertainty (SS'-SAS'-H) or the Jensen-Shannon divergence (SS'-SAS'-JS) of SAS' transitions (Methods and Supplementary Figs. 1 and 2). **b**, Normalized distribution of the arbitrator variable ω (grey bar) and its linear relationship with subjective controllability (Subj. contr.; pink line). Pairs of prediction trials are labelled as 'subjectively controllable' when counterfactual predictions differed (for example, different responses for blue and yellow actions in the circle state). **c**, The SS'-SAS'-Ω scheme captured the dynamics of subjective controllability around reversals better than the SAS' model. **d**, Coefficients of the regression predicting log-transformed reaction times in the behavioural experiment using actor and spectator prediction errors as well as ω ($N=50$; $\delta_{SAS'}: t(49)=6.17, P=1.1 \times 10^{-6}, d=0.79, CI\ 0.6-0.12$; $\delta_{SS'}: t(49)=4.39, P=0.001, d=0.49, CI\ 0.3-0.11$; $\omega: z=4.10, P=4.1 \times 10^{-5}, d=0.74, CI\ 0.08-0.18$). **e**, Positive correlation between the effect of ω on reaction times and predictive accuracy across participants ($N=50, \rho=0.67, P=1.1 \times 10^{-7}, CI\ 0.45-0.81$). Error bars and shaded areas represent s.e.m. $***P < 0.005$, $***P < 0.001$, $####P_{\text{exceedance}} > 0.999$, $\#P_{\text{exceedance}} > 0.95$. BIC, Bayesian information criterion; AIC, Akaike information criterion; F, free energy.

environmental controllability (that is, higher ω), decision times were slower. This controllability-dependent slowing correlated positively with predictive accuracy (Fig. 3e), suggesting that learning in controllable contexts is supported by a more controlled action selection process even when no reinforcement is at stake.

Separable neural correlates should therefore exist for the prediction errors generated by the actor and spectator probability tracking processes, $\delta_{SAS'}$ and $\delta_{SS'}$. A conjunction analysis first revealed that both types of prediction errors activated the typical set of bilateral

brain areas commonly associated with state prediction errors^{2,31}, such as the frontoparietal network and the pre-supplementary motor area (Fig. 4a and Supplementary Table 2). Testing directly the effect of $\delta_{SS'} - \delta_{SAS'}$ (mathematically equivalent to $p_{SAS'} - p_{SS'}$) using a conventional parametric analysis at the whole-brain level showed that the mPFC and the nucleus accumbens encoded negatively this signal required for the update of controllability. In both cases, region-of-interest analyses indicated a positive response to $\delta_{SAS'}$ and an absence of a relationship with $\delta_{SS'}$ (Fig. 4b,c and Supplementary Table 3). A similar pattern was observed in the right TPJ ($P^{\text{FWE}}=0.09$, corrected for family-wise error (FWE)) and in the dopaminergic nuclei of the brainstem at a more lenient threshold (Supplementary Fig. 4c). To ascertain this dissociation, we performed two additional analyses that fully circumvented the collinearity issues that might arise due to the correlation of prediction errors (average $r=0.56 \pm 0.057$). First, we contrasted events where only $\delta_{SS'}$ or only $\delta_{SAS'}$ was above the respective 66th percentile. Second, we contrasted the parametric effects of $\delta_{SS'}$ and $\delta_{SAS'}$ derived from two separate first-level GLMs. These analyses confirmed that the nucleus accumbens and the mPFC significantly dissociated the two prediction error terms, although only the mPFC systematically survived correction for multiple comparisons (see Supplementary Fig. S4a,b and Supplementary Table 3 for robustness checks).

Neural correlates of dynamic controllability. Having established the dissociation of $\delta_{SAS'}$ and $\delta_{SS'}$ at the behavioural and neural levels, we next probed the correlates of the prediction error δ_{Ω} governing changes in estimated controllability ($\delta_{\Omega} = \delta_{SS'} - \delta_{SAS'} - \Omega_{t-1}$). This second-order learning mechanism is key to accumulate, over time, evidence in favour or against the controllability of the ongoing rule. Whole-brain analyses revealed a significant negative relationship between δ_{Ω} and neurovascular responses in the posterior cingulate (PCC), the right dorsal anterior insula (dAI), the right TPJ and the mPFC (Fig. 4d and Supplementary Table 4). Mixed-effects region-of-interest analyses including decision times and Ω confirmed that these effects reflected a genuine response to controllability prediction error, peaking 4–8 s after trial onset.

To unravel the neural correlates of controllability with maximal sensitivity, we performed a multivoxel pattern analysis. A support vector machine classifier was trained to predict whether streaks of consecutive exploratory trials would result in an implicit report of subjective controllability in the upcoming pair of prediction trials (that is, different or identical responses for each counterfactual). Whole-brain maps of classification accuracy were obtained using the searchlight method (leave-one-run-out cross-validation). Local patterns of activity in the precuneus, the right TPJ, the supplementary motor area, the left premotor cortex and the left dorsolateral prefrontal cortex (dlPFC) contained information relative to subjective controllability (Fig. 4e and Supplementary Table 5). Interestingly, the decoding performance in the ACC and the right dlPFC scaled with the lengthening of decision times in controllable contexts from one participant to another. The dorsal bank of the mPFC and the right TPJ were the only two regions whose activity was simultaneously sensitive to subjective controllability (as probed by multivoxel pattern analysis), to controllability prediction errors and to the instantaneous difference between $\delta_{SAS'}$ and $\delta_{SS'}$ (Fig. 4f).

Uncontrollable stressors promote the spectator model. We applied this new paradigm to better understand the computational mechanisms underlying learned helplessness. More precisely, we hypothesized that exposure to uncontrollable stressors might bias controllability estimation mechanisms to promote reliance on the spectator model relative to the actor model. We invited participants to perform an active avoidance task exposing them to mild electric shocks before completing the explore-and-predict task (Fig. 5a). Participants in the controllable group learned to avoid the shock

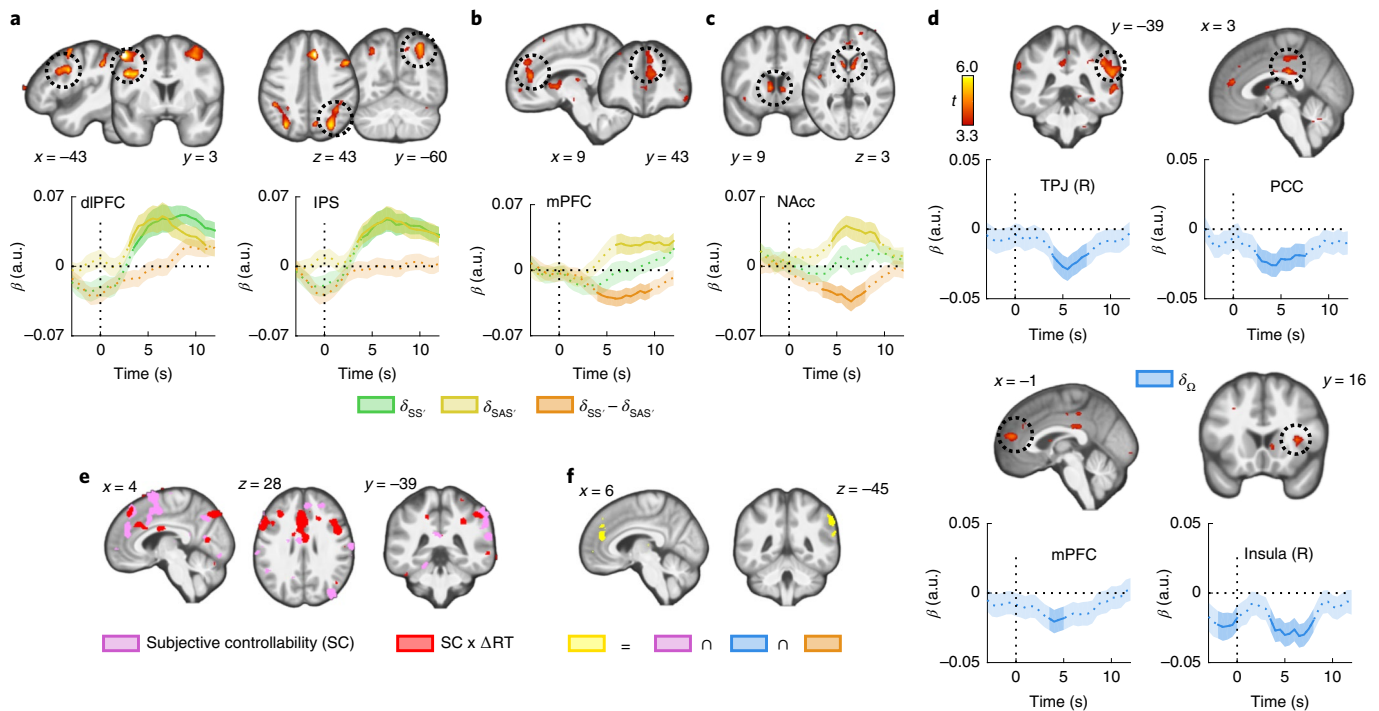


Fig. 4 | Neuroimaging. **a**, A conjunction analysis revealed brain regions whose activity encoded both $\delta_{SS'}$ and $\delta_{SAS'}$ positively. **b,c**, Parametric analysis of BOLD responses showed that the mPFC (**b**) and the nucleus accumbens (**c**) encoded negatively the difference term $\delta_{SS'} - \delta_{SAS'}$ used to update controllability. Both regions encoded $\delta_{SAS'}$ positively, but showed no clear-cut modulation by $\delta_{SS'}$. At a more lenient voxel-wise threshold ($P < 0.005$ uncorrected), the right TPJ also survived correction for multiple comparisons (Supplementary Table 3 and Fig. 4e). **d**, Brain regions encoding signed the second-order prediction errors δ_{Ω} . All areas surviving correction for multiple comparisons showed a negative effect, implying greater activity when actions appeared less causal than expected. **e**, Decoding of subjective controllability from brain data. A searchlight analysis based on the six exploratory trials preceding a prediction pair revealed that the mPFC, the posterior dmPFC, the right TPJ and the precuneus were sensitive to upcoming reports of controllability (pink). Decodability extended to the diPFC and ACC in participants who displayed slower reaction time (RT) in controllable contexts (red). **f**, Spatial overlap (yellow) between the decodability of subjective controllability (**e**), controllability prediction errors (**d**) and difference term $\delta_{SS'} - \delta_{SAS'}$ (**b** and **c**). The right TPJ and the mPFC were the only regions highlighted by each of these analyses (threshold of each map: $P < 0.005$ uncorrected). The time courses are shown below (**a-d**) were only used for robustness checks and visualization. Statistical inferences were based on whole-brain effects at standard thresholds (voxel-wise: $P < 0.001$, uncorrected; cluster-wise: $P < 0.05^{FWE}$). Shaded areas represent s.e.m. Brain co-ordinates are based on the Montreal Neurological Institute system. NAcc, nucleus accumbens; R, right hemisphere.

following one of the three possible cues by pressing the correct response button (out of six alternatives). Shocks received by participants in the uncontrollable group were yoked to the former so that their decisions did not influence shock probability. As expected, this procedure induced a dissociation between actual shock frequency, matched across groups by design, and reported shock expectancy (Fig. 5b), so that shock expectancy remained high until the end of the induction phase in the uncontrollable group.

Despite the absence of aversive reinforcers in the explore-and-predict task, we observed clear carry-over effects from the shock experiment when analysing the model parameters governing controllability estimation (Fig. 5c and Supplementary Table 6). In particular, the threshold parameter increased significantly in the uncontrollable group compared with the controllable group (Fig. 5c, left). This parameter determines how much causal evidence is required before controllability is inferred. Therefore, when making predictions, participants exposed to uncontrollable stressors relied more on the spectator model, demonstrated by the reduced average value of the arbitrator ω (Fig. 5c, right) as well as the direct analysis of subjective controllability estimates, revealing that counterfactual predictions were more often identical in the uncontrollable group ($z = 1.69$, $P = 0.045$, $d = 0.48$, 95% CI 0.006–0.10, one-tailed).

We found no statistically significant evidence that the stress manipulation affected the overall motivation of ability to perform

the task. Accuracies (U, 0.61 ± 0.16 ; C, $0.65 \pm 0.16\%$; $t(52) = 0.81$, $P = 0.43$, $d = 0.22$, 95% CI -0.05 to 0.13) and decision times (U, 2.07 ± 0.73 s; C, 1.97 ± 0.65 s; $t(52) = 0.56$, $P = 0.59$, $d = 0.14$, 95% CI -0.48 to 0.27) did not differ significantly across groups in prediction trials. Furthermore, when restricting our analysis to the uncontrollable group, the SAS'–SS'– Ω model still outperformed a variety of simpler models, including the standard SAS' model (actor only), a SS' model (spectator only) and a reinforcement learning model that only learned through the feedback delivered on predictive trials (Supplementary Fig. 5a). Exposure to uncontrollable stressors thus elicits an imbalance between actor and spectator mechanisms for transition probability learning consistent with a sustained shift in controllability expectations. This finding provides a parsimonious account of the cross-contextual generalization of passive strategies, a core feature of helpless states^{7,8}. Interestingly, state anxiety, as assessed before the experiment, moderated the induction of controllability estimation biases. It predicted the average value of the arbitrator only in participants exposed to uncontrollable stressors (Fig. 5d).

Since uncontrollable stressors promoted increased reliance on the spectator model and decreased reliance on the actor model, we expected prediction error-dependent slowing effects to follow a similar pattern. Confirming this prediction, the type of stress induction profoundly altered the slowing of decision times by actor and spectator prediction errors (Fig. 5e). Post hoc tests showed that the effect of

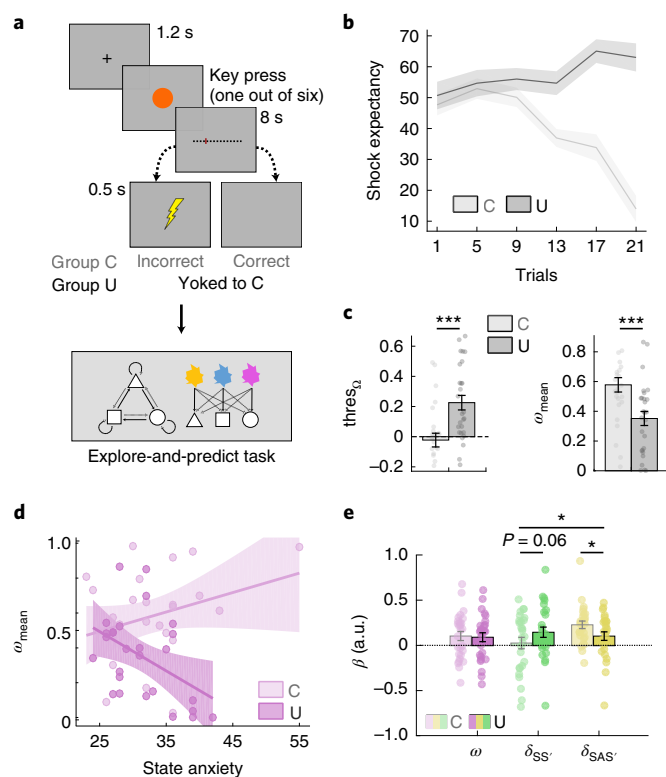


Fig. 5 | Stress experiment. **a**, Induction of controllable and yoked uncontrollable stress followed by the explore-and-predict task. **b**, Temporal evolution of shock expectancy during the induction phase, split by condition. **c**, Impact of induction type on the parameter thres_{Ω} and the average value of the arbitrator ω . The value above which Ω was treated as evidence for a controllable environment increased significantly following uncontrollable stressors ($N=27$ per group, $t(52)=4.56$, $P=2.7 \times 10^{-4}$, $d=1.1$, 95% CI 0.12–0.39), resulting in increased reliance on the spectator model when making predictions, as indexed by the reduction of ω at the group level ($N=27$ per group, $z=2.60$, $P=0.009$, $d=0.77$, 95% CI 0.08–0.29) (Supplementary Fig. 5 and Supplementary Table 6). **d**, State anxiety moderated the effect of induction type on the arbitrator variable (ω) reflecting controllability estimation. Higher state anxiety was associated with greater reliance on the spectator model after exposure to uncontrollable stressors (dark pink, $N=27$, $r=-0.46$, $P=0.015$, 95% CI -0.70 to -0.08) but not after controllable stressors (bright pink, $N=27$, $r=0.30$, $P=0.12$, CI -0.16 to 0.61). Correlation coefficients were significantly different between uncontrollable and controllable group ($z=2.87$, $P=0.002$). **e**, The impact of induction type on the slowing of decision times induced by $\delta_{SS'}$ and $\delta_{SAS'}$ (average correlation of the two first-order prediction errors 0.45 ± 0.17) was consistent with this increased reliance on the spectator model (interaction group by PE type: $F_{(1,52)}=5.97$, $P=0.018$, $\eta^2=0.07$, 95% CI 0.005–0.20). All error bars and shaded areas represent s.e.m. * $P < 0.05$, *** $P < 0.001$.

$\delta_{SAS'}$ was significantly lower ($t(52)=-2.45$, $P=0.018$, $d=0.65$, 95% CI -0.13 to -0.1) and that the impact of $\delta_{SS'}$ was marginally higher ($\delta_{SS'}$: $t(52)=1.86$, $P=0.068$, $d=0.51$, 95% CI -0.006 to 0.16) in the uncontrollable group compared with the controllable group (Fig. 5e).

Discussion

Taken together, these findings shed light on one of the most fundamental aspects of human experience, that is, the ability to estimate the extent to which our actions affect our environment and to adjust our decisions accordingly. Our results demonstrate that this ability involves the comparison of actor and spectator models

of the ongoing task, which are dissociable computationally, behaviourally and neurally. In turn, controllability estimates can be used to arbitrate between these models when making predictions about future events. The mPFC and the striatum encode the difference between the prediction errors generated by each model, while signals related to the update of controllability estimates are found in a more posterior brain network encompassing the TPJ and the PCC. Furthermore, exposure to uncontrollable stressors biases this process assessed by the explore-and-predict task, hence establishing its relevance for the study of neuropsychiatric disorders involving altered perceptions of controllability^{9,12,13,35}.

Historically, the concept of task controllability has been heavily influenced by learned helplessness studies in which animals granted the ability to actively terminate stressors are compared with yoked animals exposed to the exact sequence of stressors but whose actions are made independent from stressor termination^{8,14,15}. In this line of research, focused on the long-lasting consequences of stress exposure, more controllable contexts were defined as those in which the mutual information linking the timings of actions and stressor offsets is higher³⁶. However, a positive mutual information linking an organism's actions and upcoming states of the environment is a necessary but not sufficient condition to declare a task controllable. For example, the highly positive mutual information linking the statements of a weather forecaster with the occurrence of rain should obviously not be interpreted as a sign that the forecaster controls the weather because the statements of the forecaster and the occurrence of rain are both conditioned by past meteorological states. Moreover, following this incomplete definition, variations of task controllability were often obtained by manipulating uncertainty about future states^{17,20,28}, hence leading to ambiguous conclusions regarding the mechanisms underlying the estimation of controllability per se and its downstream influence on behaviour.

Formalizing controllability using TE rather than mutual information allowed us to design a task in which controllability varied independently from uncertainty. In addition, this approach provided an algorithm for detecting genuine changes in task controllability. Model comparisons showed that, across the three experiments, algorithms monitoring controllability using an approximation of TE (that is, the SAS'–SS'– Ω scheme) accounted better for participants' choices than a simpler SAS' learning algorithm. While the neural correlates of state prediction errors elicited by the explore-and-predict task were remarkably similar to those reported in two-step decision tasks analysed using SAS' (or actor) learning^{2,37,38}, model comparisons revealed the concomitant engagement of SS' (or spectator) learning. Crucially, the SAS'–SS'– Ω model also outperformed two architectures in which controllability was derived from statistical features of the SAS' transition probabilities only. Namely, they exploited the uncertainty (SAS'–SS'–H) or the divergence of counterfactual SAS' transitions (SAS'–SS'–JS) to discriminate between controllable and uncontrollable contexts. While they predicted qualitatively similar behaviours (for example dynamics of subjective controllability change around reversals), these alternative models were less efficient at capturing the dynamic fluctuations of subjective controllability during the task. Finally, the analysis of reaction times confirmed that participants were sensitive to the prediction errors generated by the spectator and actor models, whose comparison governed the update of controllability estimates in the SAS'–SS'– Ω scheme only.

The actor and the spectator can be viewed as two state spaces competing to structure the learning of statistical contingencies. When controllability estimates are low, the spectator model representing only the successive states of the environment dominates. In contrast, when controllability estimates are high, the actor model representing both states and actions takes over. By defining the most appropriate state space dynamically, controllability estimation improves predictions about the future states of one's environment

and can therefore contribute to maximizing utility when reward or punishment rates depend on such predictions. And by promoting reliance on a simpler spectator model when the environment is deemed uncontrollable, it can also minimize the metabolic cost and subjective effort associated with controlled action selection^{39,40}. These hypotheses could be tested directly by introducing reinforcers in the explore-and-predict task, but it is already worth noting that the controllability-dependent arbitration logics can readily explain why Pavlovian (equivalent to *SS'*) and instrumental (equivalent to *SAS'*) learning mechanisms are respectively promoted in uncontrollable and controllable contexts²⁸. Furthermore, the preferential encoding of actor prediction errors by the mPFC, the striatum and the dopaminergic midbrain is consistent with earlier findings showing that the mesolimbic pathway preferentially encodes reward prediction errors in instrumental learning tasks^{6,41–43}. The finding that the actor and spectator models only dissociated in these deep structures close to the midline is consistent with a recent magnetoencephalography study showing that the human brain exploits shared 'neural codes' to address sensorimotor and perceptual demands in controllable and uncontrollable contexts⁴⁴.

By comparing the predictions emanating from the actor and spectator models, one can derive an instantaneous causality signal (that is, how likely it is that the last action caused the last state transition). Encoded by mPFC and striatal blood oxygen level-dependent (BOLD) responses, this instantaneous signal can then be integrated over time, hence reflecting the causal influence of actions over recent transitions. A signature of the second-order prediction errors supporting this integration was found in the right TPJ, the dorsal mPFC, the right insula and the PCC. The right TPJ and the mPFC were the only regions sensitive to the difference of the two first-order prediction errors, to these second-order prediction errors used to update controllability and to subjective controllability as assessed by the decoding analysis. They are thus strong candidates for the implementation of controllability monitoring in our task. Supporting this view, the right (but not left) TPJ has previously been found to encode the divergence in action–outcome distributions²⁴ and the discrepancy between expected and actual outcome timings in a simple sensorimotor task alternating controllable and uncontrollable trials²³. It is also consistent with a study showing that mPFC lesions can alter the perception of controllability in simple instrumental learning tasks⁴⁵. Given that the mPFC and the uncertainty of *SAS'* transitions are involved in the trade-off between model-based and model-free decision-making systems⁴⁶, these results support the emerging hypothesis^{7,47} that perceived or expected controllability may play a role in the relative influence of these systems as assessed by two-step tasks^{31,48}.

Other prefrontal areas were sensitive to variations in subjective controllability according to the decoding analysis, probably reflecting adaptations of brain networks to task controllability⁴⁹ or the contribution to controllability estimation of cognitive processes that were not captured by the *SAS'*–*SS'*– Ω scheme. For example, the ACC and the anterior insula, which encoded controllability prediction errors, play an established role in the signalling of state uncertainty or task volatility, both of which may participate in controllability estimation^{28,50,51}. Strikingly, a higher sensitivity of the dlPFC and the dorsomedial prefrontal cortex (dmPFC, extending ventrally to the dorsal ACC) to controllability was also associated with a stronger influence of controllability on decision times. This finding suggests that controllability detection may foster a form of proactive response inhibition previously linked with dlPFC and dmPFC activity^{52,53}. The engagement of more elaborate action selection processes may ultimately depend on the valuation of control itself, which is known to involve dorsal ACC activity⁵⁴. Indeed, while this process is usually studied by varying task difficulty, it is clear that by gating the causal influence of actions, variations of task controllability moderate the expected benefits of exerting cognitive control.

Having described the computational principles and outlined neural correlates of controllability estimation, we sought to test whether an experimental manipulation could alter this process and simultaneously contribute to a better understanding of the learned helplessness phenomenon. Indeed, exposure to uncontrollable stressors is known to induce passive responses to subsequent controllable stressors, but the origins of this maladaptive strategy remain poorly understood. In particular, it is unclear whether prior exposure to uncontrollable stressors induces an increased sensitivity to future aversive events, reduces the expectation of control with respect to future stressors or reduces expectations of control in general^{8,29}. Our results support the latter hypothesis by showing sustained alterations of controllability estimation in human participants previously exposed to uncontrollable versus controllable stressors. More precisely, the specific increase observed for the threshold parameter implies that the former group needed to integrate more causal evidence before considering a given rule as controllable in the explore-and-predict task. The dAI is involved in the modulation of pain perception by controllability²⁰ and was found to encode controllability prediction errors in our fMRI experiment. Therefore, prior exposure to uncontrollable stressors may have altered dAI excitability to distort subsequent controllability estimation mechanisms. Supporting this idea, a study showed that a lower perception of control mediates the exacerbation of dAI responses to physical threats in more anxious individuals, who also displayed lower controllability estimates following uncontrollable stressors in our data¹⁹. Yet, this increased reliance on the spectator relative to the actor model following uncontrollable stressors probably involves several other brain areas, including the mPFC and the dorsal raphe nucleus, both of which are highly sensitive to stressor controllability^{14–16}.

In sum, the explore-and-predict task allowed us to isolate the core computations supporting the inference of task controllability by excluding reinforcers and matching uncertainty across contexts. The mPFC and the right TPJ emerged as the two most promising candidates for the neural implementation of controllability inference, and it would therefore be interesting to confirm their causal contribution using brain stimulation methods in the future. By showing that the human brain can compute an approximation of TE, our study may help to bridge the gap between neuroscience and artificial intelligence research, where TE plays an important role in solving unsupervised learning problems^{55,56}. Investigating in greater detail the interactions between controllability estimation and model-based reinforcement learning mechanisms will constitute an important step in this direction. More invasive techniques will also be required to understand how these computations are implemented within local neural circuits and how neuromodulators such as dopamine or serotonin mediate their broad impact on stress responses and mental health^{7,14,15,57}.

Methods

Participants. For the behavioural experiment, 50 young adult participants (mean age 24.7 years, range 18–43 years, 27 women) were recruited via the Sona system (a human subject pool management system) of the Radboud University (the Netherlands). All participants were included in the data analysis. For the fMRI experiment, 32 young adult participants (mean age 25.1 years, range 20–43 years, 18 women) were recruited through the same system. For the stress experiment, a total of 62 participants (mean age 21.8 years, range 18–27 years, 52 women) were recruited via the Sona system of Leiden University. One additional participant was excluded a posteriori from the fMRI experiment, and four participants were excluded from the stress experiments, together with their yoked counterparts (see Supplementary Methods for details on exclusion and inclusion criteria). The behavioural and fMRI experiments were approved by the local ethics committee (CMO region Arnhem/Nijmegen, the Netherlands, CMO2001/095). The stress experiment was approved by the Psychology Research Ethics Committee (CEP17-0905/282) at Leiden University. All participants provided written informed consent, in line with the Declaration of Helsinki, and were compensated for their participation in the study (€10 per hour for the behavioural and fMRI experiments, €7.5 per hour for the stress experiment).

Explore-and-predict task. In the three experiments, the overall structure of the task was identical. Participants performed six (fMRI and stress experiment) or seven (behavioural experiment) exploratory trials before a pair of predictions were required. Pairs of predictions always probed the two actions available for a given state (for example, blue followed by yellow in the circle state), to derive subjective controllability from counterfactual responses. Participants received feedback about their predictions in 50% (fMRI and stress) or 100% (behavioural experiment) of the trials. In the fMRI and stress tasks, feedback was delivered only after one of the two counterfactual predictions, to prevent participants from inferring whether the rule was controllable or not based on feedback.

On each exploratory trial, two identical geometrical shapes were displayed side by side. The colour of each shape determined the action corresponding to left and right button presses (side randomly assigned in each trial). An urgency signal was displayed after 1.5 s. Transitions to the next state were always governed by one of the four rules (Fig. 1c). To maximize the variation of prediction errors, the transitions were stochastic (noise 0.05–0.2).

The first prediction trial of each pair was simply displayed at the end of the intertrial interval (ITI) of the previous exploratory trial. An urgency signal was displayed after 4 s. The hypothetical state–action pair was displayed at the centre of the screen, just below a question mark, and the three possible next states were displayed as white geometrical shapes at the top of the screen. The selected state was then highlighted, and the feedback was displayed when applicable.

The ongoing rule was never changed before four pairs of predictions were completed. In the behavioural experiment, the rule changed from then as soon as five correct responses were provided in the last six predictions or if the last four predictions were accurate. In the fMRI experiment, the rule was changed as soon as the P value of a binomial test indicated that accuracy was significantly below chance ($P < 0.05$, one-tailed, chance level 1/3), hence making the accuracy threshold more lenient as the number of predictions made for a given rule increased. In all experiments, the rule changed after ten pairs of predictions, even if performance did not meet the learning criterion. Prediction trials were pseudo-randomly ordered with the constraint that each state would be tested a similar number of times. An exhaustive description of instructions, counterbalancing, reversal schedules, transition noises and timings is available in Supplementary Methods.

Stress induction task. To test the impact of prior controllability over stress on subsequent controllability estimations, participants underwent a stressor controllability manipulation before the explore-and-predict task in a between-subjects design. Critically, we employed a between-subjects yoked control procedure to match the amount and order of aversive outcome stimuli between the controllable and uncontrollable conditions. We randomized participants in blocks of four where the controllable condition of a yoked pair was always administered first to create the schedule for the yoked counterpart in the uncontrollable condition.

Electric stimuli served as stressors in the manipulation task and were delivered by a Digitimer DS7 stimulator. First, individual levels of intensity of the electric stimulus for the manipulation task were determined using a stepwise procedure in which the intensity of the stimulus was gradually increased until participants reported a ‘just bearable, but not yet painful’ experience of shock. A yoked control design with pre-programmed pseudo-randomized schedule enabled us to match the amount and order of electric stimuli between the conditions.

In the controllable condition, a total of four cues (different in shape and colour) were presented for at least six repetitions each following a pre-programmed pseudo-randomized schedule. Participants could learn by trial and error the correct response corresponding to the cue (a key between ‘1’ and ‘6’) to avoid the electric stimulus. Critical trials on which participants would be able to prevent the electric stimulus for the first time according to the schedule were repeated until the participants arrived at a correct response. As such, all participants underwent the whole schedule with a minimum of 24 trials and were able to acquire the correct response for each cue.

The uncontrollable condition was yoked to the controllable condition, such that participants experienced a comparable pattern of events across conditions. However, in the uncontrollable condition, participants were not able to acquire these action–outcome contingencies to prevent the shocks, whose sequences were merely replayed from the yoked participants performing the controllable condition. Data collection and analysis were not performed blind to the conditions of the experiments. An exhaustive description of counterbalancing, instructions and procedures is available in Supplementary Methods.

Computational modelling. The main purpose of all SAS’–SS’– Ω algorithms is to provide a way to dynamically estimate the causal influence of actions over state transitions by updating a variable termed Ω . In all models, S represents the previous state of the environment, A represents the previous action and S' represents the current state of the environment. The local causality estimate Ω can only be used as a proxy for controllability, which is a property not of actions but of the environment. It is this ‘inferred controllability’ variable, termed ω , which can then be used to decide (arbitrate) whether one should make predictions using learned S – S' transitions or learned SA – S' transitions. Ω is homologous to

TE, which is itself a generalization of Granger causality to discrete and nonlinear domains. See Supplementary Methods for a detailed explanation of the differences between TE and Ω .

To demonstrate that participants used a dynamic estimate of TE to solve the task, we compared the SAS’–SS’– Ω algorithm with a standard model-based architecture tracking SAS’ transitions². This latter algorithm corresponds to the actor model alone. Its asymptotic performance in stable environments is identical to that of SAS’–SS’– Ω algorithms. We also compared the SAS’–SS’– Ω algorithm with alternative controllability estimation schemes based solely on statistical features of SAS’ transitions, namely the uncertainty of SAS’ transitions (SAS’–SS’– H , lower during periods of high controllability) and the Jensen–Shannon divergence of counterfactual SAS’ transitions (SAS’–SS’– J , higher during periods of high controllability). These two alternatives schemes are fully described in Supplementary Information.

The actor model tracks transitions linking state–action pairs to newly encountered states (that is, SAS’). It updates transition probabilities in the following fashion:

Realized transitions:

$$p(s'|a, s) \leftarrow p(s'|a, s) + \alpha(1 - p(s'|a, s)).$$

Unrealized transitions:

$$p(s'|a, s) \leftarrow p(s'|a, s)(1 - \alpha),$$

Where $\alpha \in [0, 1]$ controls the extent to which learned transition probabilities are determined by the most recent transitions. The prediction error $1 - P(s'|a, s)$ is noted δ_{SAS} in the main text.

The spectator model tracks transitions linking states to newly encountered states (that is, SS’). Therefore, it updates transition probabilities exactly like the actor model, except that only states are represented: $p(s'|s, a)$ is simply replaced by $p(s'|s)$ in the two equations above, and the prediction error $1 - p(s'|s)$ is noted $\delta_{SS'}$ in the main text.

Following the update of the actor and spectator models, we allowed prior transition probabilities derived from the instruction phase to constrain the update of each model. This was done by multiplying (element-wise) the relevant vector of probabilities by the corresponding vector of prior probabilities. For example, after a transition from a blue circle (state 2, action 2), the transition probabilities of the spectator model were multiplied by $[0.5 \lambda \ 0.5]$ (reflecting the fact that states did not repeat under uncontrollable rules) and those of the actor model were multiplied by $[\lambda \ 0.5 \ 0.5]$ (reflecting the fact that the triangle state never appeared after choosing blue under uncontrollable rules). Thus, for any $\lambda < 0.5$, this prior injection step constrained the update of the spectator and actor model in a way that reflected the transitions a priori possible under uncontrollable and controllable rules, respectively. By altering the prediction errors elicited by each model, prior injection had an indirect influence on controllability estimation. For example, a lower value of λ leads to a greater increase of estimated controllability following state repetition events, by reducing the prediction error generated by the actor model relative to the spectator model (see below).

The variable Ω supports the controllability estimation process by tracking the expected difference $p(s'|a, s) - p(s'|s)$ dynamically (or, equivalently, $\delta_{SAS} - \delta_{SAS}$). The logic of this process is that, in a controllable environment, actions contribute to predicting the upcoming states and therefore $p(s'|a, s) > p(s'|s)$. Higher values of Ω therefore imply higher evidence that the environment is controllable. The update of Ω is governed by the following equation:

$$\Omega \leftarrow \Omega + \alpha_{\Omega}(p(s'|a, s) - p(s'|s) - \Omega),$$

where $\alpha_{\Omega} \in [0, 1]$ is the learning rate controlling the extent to which Ω is determined by the most recent observations.

Since Ω reflects the causal influence of one’s action over state transition, it can be used as a proxy to infer whether the environment is probably controllable or uncontrollable. To form the arbitration term reflecting this inference and accommodate inter-individual differences at this step, Ω is thus transformed using a parametrized sigmoid function,

$$\omega = 1/(1 + \exp(-\beta_{\Omega}(\Omega - \text{thres}_{\Omega}))),$$

where $\text{thres}_{\Omega} \in [-1, 1]$ corresponds to the threshold above which Ω is interpreted as evidence that the environment is controllable and where $\beta_{\Omega} \in [0, \infty]$ determines the extent to which evidence that the environment is controllable (that is, $\Omega - \text{thres}_{\Omega} > 0$) favours reliance on learned SAS’ transitions when making predictions (and vice versa for SS’ transitions when $\Omega - \text{thres}_{\Omega} < 0$). Thus, the variable ω implements the arbitration between the ‘actor’ and the ‘spectator’ model.

When only SAS’ learning is considered, the probability that a given state $S' = i$ will be observed given s and a is directly given by

$$p(S' = i) = p(S' = i|a, s).$$

When the $SS'-SAS'-\Omega$ architecture is used, the probability that a given state $S' = i$ will be observed given s , a and ω is directly given by

$$p(S' = i) = \omega p(S' = i|a, s) + (1 - \omega) p(S' = i|s).$$

In turn, the probability that the participant will predict that the next state would be i (for example, a square) when confronted with the hypothetical state-action pair s , a (for example, circle state–blue action) is given by

$$p(\text{prediction} = i) = \exp^{\beta_{\text{choice}}(S'=i)} / \sum_{j=1}^3 \exp^{\beta_{\text{choice}}(S'=j)},$$

where $\beta_{\text{choice}} \in [-\infty, \infty]$ determines the extent to which the participants will systematically select the most likely transition (that is, the highest $p(S' = i)$, according to what has been learned) to make their predictions. A very positive β_{choice} implies that the participant systematically selects this most likely transition. A β_{choice} around zero implies that the participant mostly makes random guesses. And a very negative β_{choice} would imply that the participant mostly goes against what he/she has learned.

The full model space was composed of SAS' alone (two parameters), SAS'-SS'- Ω (six parameters), SAS'-SS'-H (five parameters) and SAS'-SS'-J (five parameters). An exhaustive description of these models is available in Supplementary Information.

Model fitting procedures. Model fitting was performed using a variational Bayesian estimation procedure using the well-validated VBA toolbox⁵⁸. The fitting procedure only attempted to explain decisions made in prediction trials. In other words, the decisions made in exploratory trials only indirectly constrained the fit by determining the information gleaned between pairs of prediction trials. For the behavioural experiments, the prior distributions of the various learning rates and the threshold parameter were innately defined as Gaussian distributions of mean 0 and variance 3, which approximates the uniform distribution over the interval of interest after sigmoid transformations. The prior distributions of β_{choice} and β_{ω} parameters were defined as Gaussians of mean 0 and variance 10. For the fMRI and the stress experiments, the prior distributions of every parameter were defined using the posterior mean and variance obtained from the 50 participants who passed the behavioural experiment. Hidden states corresponding to transition probabilities were systematically initialized at 1/3 (equiprobability prior), while Ω was initialized at 0. Detailed information about parameter transformation, model fitting, model comparison and simulation procedures is available in Supplementary Information.

fMRI acquisition. All images were collected using a 3-T Siemens Magnetom Prismafit MRI scanner (Erlangen, Germany) with a 32-channel head coil. A T2*-weighted multiband echo-planar imaging sequence with acceleration factor 8 (MB8) was used to acquire BOLD fMRI whole-brain covered images (repetition time 700 ms, echo time 39 ms, flip angle 52°, voxel size 2.4 × 2.4 × 2.4 mm³, slice gap 0 mm, field of view 210 mm). This state-of-the-art sequencing protocol was optimized from the recommended imaging guidelines of the Human Connectome Project, with the fast acquisition speed facilitating the detection and removal of non-neuronal contributions to BOLD changes (<http://protocols.humanconnectome.org/HCP/3T/imaging-protocols.html>).

The experiment was divided into four blocks lasting on average 7.7 ± 2.1 min (662 ± 179 volumes). We recorded participants' heartbeats using the scanner's built-in photoplethysmograph, placed on the right index finger. Respiration was measured with a pneumatic belt positioned at the level of the abdomen. Anatomical images were acquired using a T1-weighted magnetization-prepared rapid gradient-echo sequence with a generalized autocalibrating partial parallel acquisition acceleration factor of 2 (repetition time 2,300 ms, echo time 3.03 ms, voxel size 1 × 1 × 1 mm³, 192 transversal slices, flip angle 8°). Field magnitude and phase maps were also acquired.

fMRI pre-processing. fMRI data processing and statistical analyses were performed using statistical parametric mapping (SPM12; Wellcome Trust Centre for Neuroimaging). For each session, the first four volumes were automatically discarded by the scanner. Functional images were slice-time corrected, unwarped using the field maps and realigned to the mean functional image using a rigid-body registration. Functional images were then coregistered to the anatomical T1. Next, the anatomical images were segmented based on tissue prior probability maps for spatial normalization using the diffeomorphic anatomical registration through exponentiated Lie algebra algorithm, and the resulting normalization matrix was applied to all functional images. Finally, all images were spatially smoothed with a 6 mm Gaussian kernel, except in the decoding analysis for which unsmoothed images were used.

fMRI analyses. Statistical analyses of fMRI signals were performed using a conventional two-level random-effects approach in SPM12. All GLMs described below included the six unconvolved motion parameters from the realignment step. We also included the eigenvariate of signals from cerebrospinal fluid in our GLM

(fourth and lateral ventricular). Moreover, we used a retrospective image correction (RETROICOR) method to regress out physiological noise, using ten cardiac phase regressors and ten respiratory phase regressors obtained by expanding cosines and sines of each signal phase to fifth order. We also included time-shifted cardiac rates (lag +6, +10 and +12 s) and respiratory volume (-1 and +5 s) as nuisance regressors. All regressors of interest were convolved with the canonical haemodynamic response function. All GLM models included a high-pass filter to remove low-frequency artefacts from the data (cut-off 96 s) as well as a run-specific intercept. Temporal autocorrelation was modelled using an AR(1) process. All motor responses recorded were modelled using a zero-duration Dirac function. We used standard voxel-wise threshold to generate SPM maps ($P < 0.001$ uncorrected), unless otherwise indicated. All statistical inferences based on whole-brain analyses satisfied the standard multiple comparison threshold ($P^{\text{FWE}} < 0.05$) at the cluster level unless otherwise indicated. The cluster-size correction was based on random field theory. Prediction error and (log-transformed) decision time regressors were systematically z-scored within individual blocks to exclude scaling effects.

All GLM models included separate onset regressors for motor responses, for prediction trials and for the first trial of each exploratory sequence (where no prediction error was elicited). All models also included parametric regressors for reaction time and ω (reflecting controllability estimates) on prediction trials. A detailed description of the GLMs used to analyse neuroimaging data is available in Supplementary Methods. These GLMs only differ in the way exploratory trials were treated.

To verify the robustness of our whole-brain results and inspect the time course of our parametric effects of interest, we performed mixed-effects analyses on BOLD signals filtered and adjusted for nuisance regressors. This adjusted signal was extracted from the functional clusters uncovered by whole-brain analyses and segmented into trial epochs from -3 to +16 s around the onset of each exploration trial (excluding the first of each streak). We then estimated the effect of each regressor of interest, at each time point, for all subjects simultaneously. Subject identity was included as a random effect, and a subject-specific intercept was included. Parametric regressors were z-scored in the same way as in the mass univariate analyses. Importantly, this approach was not used for statistical inference (since doing so would constitute double dipping) but merely for visualization purposes.

Decoding analyses were performed using the TDT toolbox⁵⁹. Each mini-block of six exploratory trials was arbitrarily coded as +1 (controllable) or -1 (uncontrollable) based on the responses given in the upcoming prediction pair (uncontrollable for identical responses, controllable for different responses). We used a leave-one-run-out cross-validation scheme with 100 permutations per subject, so that classes remained balanced for training. The training was performed on the beta values associated with each mini-block using a support vector machine classifier (L2-loss function, cost parameter set to 1; Liblinear, version 1.94), without feature selection or feature transformation. Since we did not constrain the testing sets to have balanced classes, balanced accuracies were used when reporting the results of the searchlight analysis (performed within an 8 mm sphere) at the whole-brain level.

Statistical procedures. Model selections relied on Bayesian model comparisons and exceedance probabilities, as implemented by the VBA toolbox⁵⁸. The analysis of predictive accuracies over time and across conditions relied on two-way repeated-measure analyses of variance (ANOVAs) or one-sample t tests, assuming normal distribution of the data following arcsin transformation. To assess whether predictive performance was significantly superior to chance, we permuted correct responses for each hypothetical state independently and compared these permuted responses with actual predictions (1,000 permutation per participants). The resulting empirical chance levels were indeed higher than the theoretical level of 1/3 in all experiments (behavioural: 38.0 ± 2.5%; fMRI: 39.4 ± 1.6%; stress: 38.6 ± 2.6%), reflecting the fact that participants knew that three transition rules out of four did not allow state repetitions.

The analysis of decision times was performed in two steps. First, a logistic regression was performed on binarized decision times (median-split). Second, group-level significance was assessed by means of one-sample t tests. For the analysis of decision times, we excluded trials in which decision times were three standard deviations above the mean. Comparison between conditions relied on paired t -tests, and comparison between groups (stress experiments) relied on two-sample t tests, unless normality assumptions were violated, in which case non-parametric equivalents were used (Wilcoxon signed-rank and rank-sum tests, respectively). All t tests were two sided unless otherwise indicated. Correlations were based on Pearson coefficients unless normality assumptions were violated, in which case Spearman rank coefficients were used. Confidence intervals were computed using a bootstrapping approach (2,000 permutations).

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The experimental paradigm and the code used to generate the figures are available at the following address: <https://github.com/romainligneu/NHBControllability>.

Second-level SPM images are available at <https://neurovault.org/collections/8810/>. Anonymized data can be accessed using ORCID identification at https://data.donders.ru.nl/collections/di/dccn/DSC_3017049.01_905.

Code availability

The scripts used to collect and analyse data are available upon publication at <https://github.com/romainligneu/NHBcontrollability>.

Received: 25 November 2020; Accepted: 17 January 2022;

Published online: 10 March 2022

References

- Bastos, A. M. et al. Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
- Gläscher, J., Daw, N., Dayan, P. & O’Doherty, J. P. States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* **66**, 585–595 (2010).
- Rao, R. P. N. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).
- Niv, Y. Learning task-state representations. *Nat. Neurosci.* **22**, 1544–1553 (2019).
- Kim, D., Park, G. Y., O’Doherty, J. P. & Lee, S. W. Task complexity interacts with state-space uncertainty in the arbitration between model-based and model-free learning. *Nat. Commun.* **10**, 5738 (2019).
- Hamid, A. A., Frank, M. J. & Moore, C. I. Wave-like dopamine dynamics as a mechanism for spatiotemporal credit assignment. *Cell* **184**, 2733–2749 (2021).
- Moscarello, J. M. & Hartley, C. A. Agency and the calibration of motivated behavior. *Trends Cogn. Sci.* **21**, 725–735 (2017).
- Maier, S. F. & Seligman, M. E. P. Learned helplessness at fifty: insights from neuroscience. *Psychol. Rev.* **123**, 349–367 (2016).
- Cheng, C., Cheung, S. F., Chio, J. H. & Chan, M.-P. S. Cultural meaning of perceived control: a meta-analysis of locus of control and psychological symptoms across 18 cultural regions. *Psychol. Bull.* **139**, 152–188 (2013).
- Hammack, S. E., Cooper, M. A. & Lezak, K. R. Overlapping neurobiology of learned helplessness and conditioned defeat: implications for PTSD and mood disorders. *Neuropharmacology* **62**, 565–575 (2012).
- Harrow, M., Hansford, B. G. & Astrachan-Fletcher, E. B. Locus of control: relation to schizophrenia, to recovery, and to depression and psychosis — A 15-year longitudinal study. *Psychiatry Res.* **168**, 186–192 (2009).
- Gillan, C. M. et al. Obsessive-compulsive disorder patients have a reduced sense of control on the illusion of control task. *Front. Psychol.* **5**, 204 (2014).
- Diener, C., Kuehner, C., Brusniak, W., Struve, M. & Flor, H. Effects of stressor controllability on psychophysiological, cognitive and behavioural responses in patients with major depression and dysthymia. *Psychol. Med.* **39**, 77–86 (2009).
- Amat, J. et al. Medial prefrontal cortex determines how stressor controllability affects behavior and dorsal raphe nucleus. *Nat. Neurosci.* **8**, 365–371 (2005).
- Bland, S. T. et al. Stressor controllability modulates stress-induced dopamine and serotonin efflux and morphine-induced serotonin efflux in the medial prefrontal cortex. *Neuropsychopharmacology* **28**, 1589–1596 (2003).
- Challis, C. et al. Raphe GABAergic neurons mediate the acquisition of avoidance after social defeat. *J. Neurosci.* **33**, 13978–13988 (2013).
- Kerr, D. L., McLaren, D. G., Mathy, R. M. & Nitschke, J. B. Controllability modulates the anticipatory response in the human ventromedial prefrontal cortex. *Front. Psychol.* **3**, 557 (2012).
- Wood, K. H. et al. Controllability modulates the neural response to predictable but not unpredictable threat in humans. *NeuroImage* **119**, 371–381 (2015).
- Alvarez, R. P. et al. Increased anterior insula activity in anxious individuals is linked to diminished perceived control. *Transl. Psychiatry* **5**, e591–e591 (2015).
- Bräscher, A.-K., Becker, S., Hoeppli, M.-E. & Schweinhardt, P. Different brain circuitries mediating controllable and uncontrollable pain. *J. Neurosci.* **36**, 5013–5025 (2016).
- Haggard, P. Sense of agency in the human brain. *Nat. Rev. Neurosci.* **18**, 196–207 (2017).
- Kühn, S., Brass, M. & Haggard, P. Feeling in control: neural correlates of experience of agency. *Cortex* **49**, 1935–1942 (2013).
- Spengler, S., von Cramon, D. Y. & Brass, M. Was it me or was it you? How the sense of agency originates from ideomotor learning revealed by fMRI. *NeuroImage* **46**, 290–298 (2009).
- Liljeholm, M., Wang, S., Zhang, J. & O’Doherty, J. P. Neural correlates of the divergence of instrumental probability distributions. *J. Neurosci.* **33**, 12519–12527 (2013).
- Vejmelka, M. & Paluš, M. Inferring the directionality of coupling with conditional mutual information. *Phys. Rev. E* **77**, 026214 (2008).
- Barnett, L., Barrett, A. B. & Seth, A. K. Granger causality and transfer entropy are equivalent for Gaussian variables. *Phys. Rev. Lett.* **103**, 238701 (2009).
- Pearl, J. Interpretation and identification of causal mediation. *Psychol. Methods* **19**, 459–481 (2014).
- Dorfman, H. M. & Gershman, S. J. Controllability governs the balance between Pavlovian and instrumental action selection. *Nat. Commun.* **10**, 5826 (2019).
- Ly, V., Wang, K. S., Bhanji, J. & Delgado, M. R. A reward-based framework of perceived control. *Front. Neurosci.* **13**, 65 (2019).
- Huys, Q. J. M. & Dayan, P. A Bayesian formulation of behavioral control. *Cognition* **113**, 314–328 (2009).
- Lee, S. W., Shimojo, S. & O’Doherty, J. P. Neural computations underlying arbitration between model-based and model-free learning. *Neuron* **81**, 687–699 (2014).
- Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A. & Daw, N. D. Working-memory capacity protects model-based learning from stress. *Proc. Natl Acad. Sci. USA* **110**, 20941–20946 (2013).
- Shahar, N. et al. Improving the reliability of model-based decision-making estimates in the two-stage decision task with reaction-times and drift-diffusion modeling. *PLoS Comput. Biol.* **15**, e1006803 (2019).
- Shahar, N. et al. Credit assignment to state-independent task representations and its relationship with model-based decision making. *Proc. Natl Acad. Sci. USA* **116**, 15871–15876 (2019).
- Voss, M., Chambon, V., Wenke, D., Kühn, S. & Haggard, P. In and out of control: brain mechanisms linking fluency of action selection to self-agency in patients with schizophrenia. *Brain* **140**, 2226–2239 (2017).
- Maier, S. F. & Seligman, M. E. Learned helplessness: theory and evidence. *J. Exp. Psychol. Gen.* **105**, 3–46 (1976).
- Miller, K. J., Botvinick, M. M. & Brody, C. D. Dorsal hippocampus contributes to model-based planning. *Nat. Neurosci.* **20**, 1269–1276 (2017).
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-based influences on humans’ choices and striatal prediction errors. *Neuron* **69**, 1204–1215 (2011).
- Westbrook, A. et al. Dopamine promotes cognitive effort by biasing the benefits versus costs of cognitive work. *Science* **367**, 1362–1366 (2020).
- Hahn, A. et al. Reconfiguration of functional brain networks and metabolic cost converge during task performance. *eLife* **9**, e52443 (2020).
- O’Doherty, J. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* **304**, 452–454 (2004).
- Garrison, J., Erdeniz, B. & Done, J. Prediction error in reinforcement learning: a meta-analysis of neuroimaging studies. *Neurosci. Biobehav. Rev.* **37**, 1297–1310 (2013).
- Grogan, J. P., Sandhu, T. R., Hu, M. T. & Manohar, S. G. Dopamine promotes instrumental motivation, but reduces reward-related vigour. *eLife* **9**, e58321 (2020).
- Weiss, A., Chambon, V., Lee, J. K., Drugowitsch, J. & Wyart, V. Interacting with volatile environments stabilizes hidden-state inference and its brain signatures. *Nat. Commun.* **12**, 2228 (2019).
- O’Callaghan, C., Vaghi, M. M., Brummerloh, B., Cardinal, R. N. & Robbins, T. W. Impaired awareness of action-outcome contingency and causality during healthy ageing and following ventromedial prefrontal cortex lesions. *Neurophysiology* **128**, 282–289 (2019).
- Kim, D., Park, G. J., O’Doherty, J.-D. & Lee, S. W. Task complexity interacts with state-space uncertainty in the arbitration between model-based and model-free learning. *Nat. Commun.* **10**, 5738 (2019).
- Boureau, Y.-L., Sokol-Hessner, P. & Daw, N. D. Deciding how to decide: self-control and meta-decision making. *Trends Cogn. Sci.* **19**, 700–710 (2015).
- Piray, P., Toni, I. & Cools, R. Human choice strategy varies with anatomical projections from ventromedial prefrontal cortex to medial striatum. *J. Neurosci.* **36**, 2857–2867 (2016).
- Wanke, N. & Schwabe, L. Subjective uncontrollability over aversive events reduces working memory performance and related large-scale network interactions. *Cereb. Cortex* **30**, 3116–3129 (2019).
- Preuschoff, K., Quartz, S. R. & Bossaerts, P. Human insula activation reflects risk prediction errors as well as risk. *J. Neurosci.* **28**, 2745–2752 (2008).
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E. & Rushworth, M. F. S. Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–1221 (2007).
- Albares, M. et al. The dorsal medial frontal cortex mediates automatic motor inhibition in uncertain contexts: evidence from combined fMRI and EEG studies. *Hum. Brain Mapp.* **35**, 5517–5531 (2014).
- van Belle, J., Vink, M., Durston, S. & Zandbelt, B. B. Common and unique neural networks for proactive and reactive response inhibition revealed by independent component analysis of functional MRI data. *NeuroImage* **103**, 65–74 (2014).
- Shenhav, A., Cohen, J. D. & Botvinick, M. M. Dorsal anterior cingulate cortex and the value of control. *Nat. Neurosci.* **19**, 1286–1291 (2016).
- Leibfried, F., Pascual-Diaz, S. & Grau-Moya, J. A unified Bellman optimality principle combining reward maximization and empowerment. Preprint at <https://arxiv.org/abs/1907.12392> (2020).

56. Mohamed, S. & Jimenez Rezende, D. Variational information maximisation for intrinsically motivated reinforcement learning. *Adv. Neural Inf. Process. Syst.* **28**, 2125–2133 (2015).
57. Dickerson, S. S. & Kemeny, M. E. Acute stressors and cortisol responses: a theoretical integration and synthesis of laboratory research. *Psychol. Bull.* **130**, 355–391 (2004).
58. Daunizeau, J., Adam, V. & Rigoux, L. VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioural data. *PLoS Comput. Biol.* **10**, e1003441 (2014).
59. Hebart, M. N., Gorgen, K. & Haynes, J.-D. The Decoding Toolbox (TDT): a versatile software package for multivariate analyses of functional imaging data. *Front. Neuroinform.* **8**, 88 (2015).

Acknowledgements

We thank M. Frank for his constructive comments on the manuscript and computational models. We thank P. Gaalman for help with fMRI data acquisition. We thank K. Eleftheriadou, F. Dianellou, K. Koelbel and J. Breen and SOLO lab support Leiden University for their help and support with data acquisition for the stress experiment. This work was supported by grants from the Fyssen Foundation and the Behaviour and Brain Research Foundation awarded to R.L. (Young Investigator 2017) and a Vici award from the Netherlands Organisation for Scientific Research to R.C. (NWO 453-14-005). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

Conceptualization: R.L., R.C. and V.L. Methodology: R.L. and V.L. Software and formal analysis: R.L. Investigation: R.L. and V.L. Resources: R.C. and Z.M. Data curation: R.L. and V.L. Writing original draft: R.L. Writing review and editing: R.L., Z.M., V.L. and R.C. Visualization: R.L. Funding acquisition: R.L., Z.M., V.L. and R.C.

Competing interests

The authors declare no competing financial interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41562-022-01306-w>.

Correspondence and requests for materials should be addressed to Romain Ligneul or Roshan Cools.

Peer review information *Nature Human Behaviour* thanks Sang Wan Lee and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2022

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection Behavioural experiments used Psychtoolbox 3 running on Matlab 2014b (Mathworks), except for the stress induction task which used E-Prime 2.

Data analysis The data was analyzed using Matlab 2014b or 2016a (Mathworks). The fMRI data analyses were done using the SPM12 toolbox. The computational modeling analyses were done using the VBA toolbox. The decoding analyses were done using the TDT toolbox. The psychophysiological interaction analysis was performed using the GPPI toolbox.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

Behavioural data related to the three experiments will be made available upon publication at the following address: <https://github.com/romainligneul/controllability>. Second-level SPM images will be made available upon publication at the following address: <https://identifiers.org/neurovault.collection:8810>. Raw fMRI will be made available upon request to the corresponding authors.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	The sample sizes are explicitly reported in the main manuscript. Overall, we tested 50 subjects in the behavioral experiment, 33 subjects in the fMRI experiment in 62 subjects in the stress experiment. The explore-and-predict task being a new experimental paradigm with no equivalent in the literature, no explicit sample size determination was performed before the behavioral experiment. For the fMRI experiment, we included 32 participants. This exceeds the average sample size of fMRI task-based studies (N=28.5 in 2015, Turner et al., 2018, Communication Biology). Such sample size is sufficient to achieve brain-wide correlation of 2-nd level SPM maps superior or equal to 0.75 for working memory tasks, social inference tasks, motor tasks and emotion recognition tasks (Turner et al., 2018). In order to maximize the accuracy of individual parameter estimation, we maximized the number of trials without exceeding the threshold of 40 minutes often considered as a point of diminishing returns for task-based fMRI studies (Evan Lee, 2019, Communication Biology). For the stress experiment, we included 54 participants in total. In their study investigating the effects of subjective controllability on working memory performance following a stress induction protocol, Wanke and Schwabe reports a effect size of 0.931 (Cohen's d). Post-hoc power calculation indicates that this effect size requires a sample size of 26 participants per group to achieve a power of 0.95 (alpha: 0.05).
Data exclusions	No subjects were excluded from the behavioral experiment (final N=50). One subject was excluded from the fMRI experiment (final N=32), and 6 additional subjects were excluded for the analysis of brainstem responses. Eight subjects were excluded from the stress experiments (final N=54). No data points were excluded, except for the analysis decision time in which trials were decision times superior to mean + 3SD were excluded. These exclusions are justified in the main manuscript.
Replication	The main behavioral results were replicated across the behavioral, fMRI and stress experiments (i.e. Bayesian model selection and reaction time analyses). The conclusions drawn based on fMRI data were robust to qualitative variations in statistical methods.
Randomization	The ordering of blocks was counterbalanced across subjects in the explore-and-predict task, so that controllability and rule reversal effects could not be explained by time. The lateral ordering of action colors (exploration trials) and possible upcoming states (prediction trials) was pseudorandomized so as to counterbalance lateralization effects within subjects.
Blinding	No blinding procedure was used

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	<i>Briefly describe the study type including whether data are quantitative, qualitative, or mixed-methods (e.g. qualitative cross-sectional, quantitative experimental, mixed-methods case study).</i>
Research sample	<i>State the research sample (e.g. Harvard university undergraduates, villagers in rural India) and provide relevant demographic information (e.g. age, sex) and indicate whether the sample is representative. Provide a rationale for the study sample chosen. For studies involving existing datasets, please describe the dataset and source.</i>
Sampling strategy	<i>Describe the sampling procedure (e.g. random, snowball, stratified, convenience). Describe the statistical methods that were used to predetermine sample size OR if no sample-size calculation was performed, describe how sample sizes were chosen and provide a rationale for why these sample sizes are sufficient. For qualitative data, please indicate whether data saturation was considered, and what criteria were used to decide that no further sampling was needed.</i>
Data collection	<i>Provide details about the data collection procedure, including the instruments or devices used to record the data (e.g. pen and paper, computer, eye tracker, video or audio equipment) whether anyone was present besides the participant(s) and the researcher, and whether the researcher was blind to experimental condition and/or the study hypothesis during data collection.</i>
Timing	<i>Indicate the start and stop dates of data collection. If there is a gap between collection periods, state the dates for each sample cohort.</i>

Data exclusions	<i>If no data were excluded from the analyses, state so OR if data were excluded, provide the exact number of exclusions and the rationale behind them, indicating whether exclusion criteria were pre-established.</i>
Non-participation	<i>State how many participants dropped out/declined participation and the reason(s) given OR provide response rate OR state that no participants dropped out/declined participation.</i>
Randomization	<i>If participants were not allocated into experimental groups, state so OR describe how participants were allocated to groups, and if allocation was not random, describe how covariates were controlled.</i>

Ecological, evolutionary & environmental sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	<i>Briefly describe the study. For quantitative data include treatment factors and interactions, design structure (e.g. factorial, nested, hierarchical), nature and number of experimental units and replicates.</i>
Research sample	<i>Describe the research sample (e.g. a group of tagged <i>Passer domesticus</i>, all <i>Stenocereus thurberi</i> within Organ Pipe Cactus National Monument), and provide a rationale for the sample choice. When relevant, describe the organism taxa, source, sex, age range and any manipulations. State what population the sample is meant to represent when applicable. For studies involving existing datasets, describe the data and its source.</i>
Sampling strategy	<i>Note the sampling procedure. Describe the statistical methods that were used to predetermine sample size OR if no sample-size calculation was performed, describe how sample sizes were chosen and provide a rationale for why these sample sizes are sufficient.</i>
Data collection	<i>Describe the data collection procedure, including who recorded the data and how.</i>
Timing and spatial scale	<i>Indicate the start and stop dates of data collection, noting the frequency and periodicity of sampling and providing a rationale for these choices. If there is a gap between collection periods, state the dates for each sample cohort. Specify the spatial scale from which the data are taken</i>
Data exclusions	<i>If no data were excluded from the analyses, state so OR if data were excluded, describe the exclusions and the rationale behind them, indicating whether exclusion criteria were pre-established.</i>
Reproducibility	<i>Describe the measures taken to verify the reproducibility of experimental findings. For each experiment, note whether any attempts to repeat the experiment failed OR state that all attempts to repeat the experiment were successful.</i>
Randomization	<i>Describe how samples/organisms/participants were allocated into groups. If allocation was not random, describe how covariates were controlled. If this is not relevant to your study, explain why.</i>
Blinding	<i>Describe the extent of blinding used during data acquisition and analysis. If blinding was not possible, describe why OR explain why blinding was not relevant to your study.</i>
Did the study involve field work?	<input type="checkbox"/> Yes <input checked="" type="checkbox"/> No

Field work, collection and transport

Field conditions	<i>Describe the study conditions for field work, providing relevant parameters (e.g. temperature, rainfall).</i>
Location	<i>State the location of the sampling or experiment, providing relevant parameters (e.g. latitude and longitude, elevation, water depth).</i>
Access & import/export	<i>Describe the efforts you have made to access habitats and to collect and import/export your samples in a responsible manner and in compliance with local, national and international laws, noting any permits that were obtained (give the name of the issuing authority, the date of issue, and any identifying information).</i>
Disturbance	<i>Describe any disturbance caused by the study and how it was minimized.</i>

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

- n/a Involved in the study
- Antibodies
- Eukaryotic cell lines
- Palaeontology and archaeology
- Animals and other organisms
- Human research participants
- Clinical data
- Dual use research of concern

Methods

- n/a Involved in the study
- ChIP-seq
- Flow cytometry
- MRI-based neuroimaging

Antibodies

- Antibodies used
- Validation

Eukaryotic cell lines

Policy information about [cell lines](#)

- Cell line source(s)
- Authentication
- Mycoplasma contamination
- Commonly misidentified lines (See [ICLAC](#) register)

Palaeontology and Archaeology

- Specimen provenance
- Specimen deposition
- Dating methods
- Tick this box to confirm that the raw and calibrated dates are available in the paper or in Supplementary Information.
- Ethics oversight

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Animals and other organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research

- Laboratory animals
- Wild animals
- Field-collected samples
- Ethics oversight

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Human research participants

Policy information about [studies involving human research participants](#)

Population characteristics	Behavioral experiment: mean age: 24.7, range: 18–43, 27 women fMRI experiment: mean age: 25.1, range: 20–43, 18 women Stress experiment: mean age = 21.8; range: 18-27, 52 women
Recruitment	For the behavioural experiment, fifty young adult participants were recruited via the Sona system (human subject pool management system) of the Radboud University (The Netherlands). For the fMRI experiment, thirty-two young adult participants were recruited through the same system. For the stress experiment, a total of 62 participants were recruited via the Sona system of Leiden University. All participants provided written informed consent, in line with the declaration of Helsinki and were compensated for their participation in the study (10€/hour for the behavioural and fMRI experiments, 7.5€/hour for the stress experiment)
Ethics oversight	The behavioural and fMRI experiments were approved by the local ethics committee (CMO region Arnhem/Nijmegen, The Netherlands, CMO2001/095). The stress experiment was approved by the Psychology Research Ethics Committee (CEP17-0905/282) at Leiden University. All participants provided written informed consent, in line with the declaration of Helsinki.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Clinical data

Policy information about [clinical studies](#)

All manuscripts should comply with the ICMJE [guidelines for publication of clinical research](#) and a completed [CONSORT checklist](#) must be included with all submissions.

Clinical trial registration	<i>Provide the trial registration number from ClinicalTrials.gov or an equivalent agency.</i>
Study protocol	<i>Note where the full trial protocol can be accessed OR if not available, explain why.</i>
Data collection	<i>Describe the settings and locales of data collection, noting the time periods of recruitment and data collection.</i>
Outcomes	<i>Describe how you pre-defined primary and secondary outcome measures and how you assessed these measures.</i>

Dual use research of concern

Policy information about [dual use research of concern](#)

Hazards

Could the accidental, deliberate or reckless misuse of agents or technologies generated in the work, or the application of information presented in the manuscript, pose a threat to:

No	Yes
<input checked="" type="checkbox"/>	<input type="checkbox"/> Public health
<input checked="" type="checkbox"/>	<input type="checkbox"/> National security
<input checked="" type="checkbox"/>	<input type="checkbox"/> Crops and/or livestock
<input checked="" type="checkbox"/>	<input type="checkbox"/> Ecosystems
<input checked="" type="checkbox"/>	<input type="checkbox"/> Any other significant area

Experiments of concern

Does the work involve any of these experiments of concern:

No	Yes
<input type="checkbox"/>	<input type="checkbox"/> Demonstrate how to render a vaccine ineffective
<input type="checkbox"/>	<input type="checkbox"/> Confer resistance to therapeutically useful antibiotics or antiviral agents
<input type="checkbox"/>	<input type="checkbox"/> Enhance the virulence of a pathogen or render a nonpathogen virulent
<input type="checkbox"/>	<input type="checkbox"/> Increase transmissibility of a pathogen
<input type="checkbox"/>	<input type="checkbox"/> Alter the host range of a pathogen
<input type="checkbox"/>	<input type="checkbox"/> Enable evasion of diagnostic/detection modalities
<input type="checkbox"/>	<input type="checkbox"/> Enable the weaponization of a biological agent or toxin
<input type="checkbox"/>	<input type="checkbox"/> Any other potentially harmful combination of experiments and agents

ChIP-seq

Data deposition

- Confirm that both raw and final processed data have been deposited in a public database such as [GEO](#).
- Confirm that you have deposited or provided access to graph files (e.g. BED files) for the called peaks.

Data access links

May remain private before publication.

For "Initial submission" or "Revised version" documents, provide reviewer access links. For your "Final submission" document, provide a link to the deposited data.

Files in database submission

Provide a list of all files available in the database submission.

Genome browser session

(e.g. [UCSC](#))

Provide a link to an anonymized genome browser session for "Initial submission" and "Revised version" documents only, to enable peer review. Write "no longer applicable" for "Final submission" documents.

Methodology

Replicates

Describe the experimental replicates, specifying number, type and replicate agreement.

Sequencing depth

Describe the sequencing depth for each experiment, providing the total number of reads, uniquely mapped reads, length of reads and whether they were paired- or single-end.

Antibodies

Describe the antibodies used for the ChIP-seq experiments; as applicable, provide supplier name, catalog number, clone name, and lot number.

Peak calling parameters

Specify the command line program and parameters used for read mapping and peak calling, including the ChIP, control and index files used.

Data quality

Describe the methods used to ensure data quality in full detail, including how many peaks are at FDR 5% and above 5-fold enrichment.

Software

Describe the software used to collect and analyze the ChIP-seq data. For custom code that has been deposited into a community repository, provide accession details.

Flow Cytometry

Plots

Confirm that:

- The axis labels state the marker and fluorochrome used (e.g. CD4-FITC).
- The axis scales are clearly visible. Include numbers along axes only for bottom left plot of group (a 'group' is an analysis of identical markers).
- All plots are contour plots with outliers or pseudocolor plots.
- A numerical value for number of cells or percentage (with statistics) is provided.

Methodology

Sample preparation

Describe the sample preparation, detailing the biological source of the cells and any tissue processing steps used.

Instrument

Identify the instrument used for data collection, specifying make and model number.

Software

Describe the software used to collect and analyze the flow cytometry data. For custom code that has been deposited into a community repository, provide accession details.

Cell population abundance

Describe the abundance of the relevant cell populations within post-sort fractions, providing details on the purity of the samples and how it was determined.

Gating strategy

Describe the gating strategy used for all relevant experiments, specifying the preliminary FSC/SSC gates of the starting cell population, indicating where boundaries between "positive" and "negative" staining cell populations are defined.

- Tick this box to confirm that a figure exemplifying the gating strategy is provided in the Supplementary Information.

Magnetic resonance imaging

Experimental design

Design type

Indicate task or resting state; event-related or block design.

Design specifications	: Task
Behavioral performance measures	: During the task, we recorded responses and response times. Behavioral performance was assessed mainly using the accuracy of the choices made in prediction trials reported Fig. 2a

Acquisition

Imaging type(s)	: Anatomical, functional, magnitude and phase field maps. A DTI scan was collected at the very end of the experiment (data not analyzed)
Field strength	: 3T Siemens Magnetom Prismafit MRI scanner (Erlangen, Germany)
Sequence & imaging parameters	: All images were collected using a 32-channel head coil. A T2*-weighted multiband echo planar imaging sequence with acceleration factor 8 (MB8) was used to acquire BOLD-fMRI whole-brain covered images (TR = 700 ms, TE = 39 ms, flip angle = 52, voxel size = 2.4 × 2.4 × 2.4 mm ³ , slice gap = 0 mm, and FOV = 210 mm).
Area of acquisition	<i>State whether a whole brain scan was used OR define the area of acquisition, describing how the region was determined.</i>
Diffusion MRI	<input type="checkbox"/> Used <input checked="" type="checkbox"/> Not used

Preprocessing

Preprocessing software	: SPM12
Normalization	DARTEL normalization
Normalization template	: Template IXI555 MNI152 (1.5mm isotropic) from the CAT12 toolbox
Noise and artifact removal	: Visual inspection, no removal needed
Volume censoring	: No volume censoring needed. Head motion was minimized using tight cushions (<1.5mm)

Statistical modeling & inference

Model type and settings	<i>Specify type (mass univariate, multivariate, RSA, predictive, etc.) and describe essential details of the model at the first and second levels (e.g. fixed, random or mixed effects; drift or auto-correlation).</i>
Effect(s) tested	<i>Define precise effect in terms of the task or stimulus conditions instead of psychological concepts and indicate whether ANOVA or factorial designs were used.</i>
Specify type of analysis:	<input checked="" type="checkbox"/> Whole brain <input type="checkbox"/> ROI-based <input type="checkbox"/> Both
Statistic type for inference (See Eklund et al. 2016)	Classical 2-step inference. Voxel-wise cluster forming threshold: $p < 0.001$ uncorrected Cluster-wise: $p < 0.05$ FWE Minimal cluster extent: 10 voxels
Correction	: FWE

Models & analysis

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Functional and/or effective connectivity
<input checked="" type="checkbox"/>	<input type="checkbox"/> Graph analysis
<input type="checkbox"/>	<input checked="" type="checkbox"/> Multivariate modeling and predictive analysis
Functional and/or effective connectivity	<i>Report the measures of dependence used and the model details (e.g. Pearson correlation, partial correlation, mutual information).</i>
Graph analysis	<i>Report the dependent variable and connectivity measure, specifying weighted graph or binarized graph, subject- or group-level, and the global and/or node summaries used (e.g. clustering coefficient, efficiency, etc.).</i>
Multivariate modeling and predictive analysis	: Independent variable: whether a given streak of 6 exploratory trials would be followed by different counterfactual predictions (+1) or identical counterfactual predictions (-1) Feature extraction: all voxels within the searchlight radius were used Dimensionality reduction: none Model: Support Vector Machine with L2 cost function (cost=1)

Training: A complete leave-one-run out cross-validation procedure was used. Within each fold, controllable and uncontrollable targets were in equalized using a bootstrapping approach (10 draw per fold). Evaluation metrics: balanced classification accuracy was used, since we did not constrain the test sets to have equal number of controllable and uncontrollable targets