Context-dependent goal-directed control in psychosis: Prioritizing loss avoidance over reward pursuit

Wolfgang Omlor[1]*, Giacomo Cecere[1]*, Akhil Ratan Misra[1]*, Gao-Yang Huang[2], Philipp Homan[1,3]

[1]Psychiatric Hospital, University of Zurich, Lenggstrasse 31, Zurich, 8032, Switzerland
[2]Swiss Federal Institute of Technology Zurich, Zurich, Switzerland
[3]Neuroscience Center Zurich, University of Zurich & Swiss Federal Institute of Technology Zurich, Zurich, Switzerland
*These authors contributed equally

Correspondence concerning this article should be addressed to Philipp Homan, MD, PhD, Psychiatric Hospital, University of Zurich, Lenggstrasse 31, 8032 Zurich, Switzerland. E-mail: philipp.homan@bli.uzh.ch

Abstract: While individuals with psychosis struggle to pursue rewards, their ability to avoid negative outcomes appears relatively intact. The neural mechanisms of this dissociation and its relation to negative symptoms remain poorly understood. We hypothesized that in early psychosis, cognitive resources are preferentially directed toward avoiding losses rather than pursuing rewards, potentially limiting reward processing and contributing to negative symptoms. Fifty medicated and stable individuals with early psychosis and 48 healthy controls completed a two-step sequential decision task, which distinguishes between goal-directed and habitual strategies, under reward and loss conditions during functional magnetic resonance imaging (fMRI). Compared to the loss context, healthy participants showed more goal-directed (model-based) control in the reward setting. In contrast, individuals with early psychosis showed the opposite pattern, engaging model-based control more strongly in the loss context. These behavioral shifts were reflected in the prefrontal cognitive network. In the anterior cingulate, orbitofrontal, and dorsolateral prefrontal cortex, individuals with early psychosis showed increased activation during model-based learning specifically in the loss condition. Within this group, loss-related activation in the anterior cingulate was associated with anhedonia, but no other consistent links to negative symptom domains were found. Our findings suggest that early psychosis is associated with a reallocation of cognitive resources from reward pursuit towards loss avoidance. Loss-focused processing in the anterior cingulate may help explain the reduced capacity for pleasure seen in early psychosis.

Introduction:

How does the mind decide when it is worth investing cognitive effort to chase a reward versus avoid a loss? Everyday choices - double-checking a tax return to avoid a penalty, or staying late at work to secure a bonus - force the brain to weigh potential gains against possible setbacks. In schizophrenia-spectrum disorders, which are marked by profound disturbances in cognition and goal-directed behavior [1,2], this balance between reward pursuit and loss avoidance appears to be altered. Strikingly, the drive to pursue rewards is blunted, whereas the capacity to avoid losses remains intact or is even heightened [3,4]. This asymmetry raises the possibility that cognitive resources for goal-directed behavior are preferentially channeled into loss avoidance at the expense of reward pursuit [5,6]. Yet this cognitive resource-reallocation hypothesis has never been tested within a computational framework.

The two-stage reinforcement-learning task, coupled with computational modeling, allows researchers to quantify the balance between two principal decision processes that place different demands on cognitive resources: model-based control is goal-directed, prospective, cognitively demanding, and flexible, whereas model-free control is habitual, history-based, resource-efficient, and rapid [7-28]. The two-stage decision task has been used to demonstrate that compulsive conditions, including obsessive–compulsive disorder and substance dependence, are associated with reduced model-based control in reward contexts, but not in those involving loss [29,30]. Because model-based control hinges on cognitive resources [12,13] and dopaminergic signaling [31-34], both often compromised in psychotic disorders [2], individuals on the schizophrenia spectrum are expected to exhibit impaired model-based learning patterns [35]. In reward contexts of the two-stage decision task, individuals with schizophrenia have shown impairments in model-based control [36].

It remains unclear, however, how reward and loss contexts modulate the balance between model-based and model-free control in individuals with psychosis. Resolving this gap is essential to determine whether psychosis shifts goal-directed cognitive control toward loss avoidance at the expense of reward pursuit, and whether this distortion drives reward devaluation and negative symptoms. We therefore aimed to test whether model-based control and its neural implementation are redirected from reward pursuit to loss avoidance in early psychosis, and how this reallocation relates to distinct symptom domains. Guided by evidence that schizophrenia-spectrum disorders entail restricted cognitive capacity [1,2,37], value representation is dependent on cognition [5,6,38], and reward pursuit is attenuated relative to loss avoidance [3], we derived a set of specific expectations for behavior, neural engagement and their clinical correlates. We hypothesized that limited cognitive resources in psychotic disorders are preferentially allocated to goal-directed avoidance of negative

outcomes rather than reward pursuit. We therefore predicted (i) preserved model-based control in individuals with psychosis when choices prevent losses but diminished goal-directed behavior when they secure gains. We further predicted (ii) that this goal-directed control asymmetry in psychosis would be mirrored by a corresponding neural signature: stronger loss-than-reward activation in the three prefrontal hubs that instantiate model-based computations—the anterior cingulate cortex (ACC) for motivational–cognitive integration [39,40], the orbitofrontal cortex (OFC) for prospective value representation [3,41,42], and the dorsolateral prefrontal cortex (dlPFC) for planning and working memory [43,44]. We further expected (iii) that the magnitude of this loss-biased recruitment would scale particularly with apathy-related negative symptoms (avolition, asociality, and anhedonia), which have been associated with a motivational asymmetry characterized by impaired reward seeking relative to preserved loss avoidance [4,45].

To test these predictions, we combined functional MRI with the canonical two-step sequential decision task [7] in 50 medicated, clinically stable individuals with early psychosis and 48 healthy controls. Separate reward and loss versions of the task enabled us to quantify context-dependent shifts between model-based and model-free control while simultaneously mapping prefrontal activity. Such mechanistic insights may inform interventions aimed at restoring adaptive motivational processing in schizophrenia spectrum disorders.

Methods

Participants

Fifty individuals with early psychosis were recruited from the University Hospital of Psychiatry, Zurich, and 48 healthy controls were recruited through local advertisements. Psychosis diagnoses were established with a structured clinical interview according to ICD-10 criteria. Eligible patients met criteria for schizophrenia, schizoaffective disorder, bipolar disorder with psychotic features, or acute polymorphic psychotic disorder, and were within 5 years of their first psychotic episode. In addition, eligible patients showed no evidence of substance-induced psychosis. All patients were clinically stable, in remission, and receiving antipsychotic medication at the time of testing. Healthy controls had no lifetime history of psychiatric or neurological disorders. The study was approved by the Cantonal Ethics Committee of Zurich, and all procedures adhered to its guidelines.

Clinical Assessments

Symptoms were assessed using several rating scales. The Brief Negative Symptom Scale (BNSS) was used to assess the different negative symptom domains of the dimensions apathy (avolition, asociality, anhedonia) and diminished expression (blunted affect, alogia). To complement this domain-specific measure, overall psychopathology was evaluated with the Brief Psychiatric Rating Scale (BPRS), which provides a concise, clinician-rated index of positive, affective, activation and hostility symptoms in schizophrenia. Detailed cognitive functioning was additionally characterized with the MATRICS Consensus Cognitive Battery (MCCB). All participants were required to pass a urine drug screen and a Breathalyzer test.

Task Design

We used the 2-stage task as described in previous research [7]. In total, participants completed 200 trials with 3 seconds allocated for each choice. The interval between first and second-stage stimuli was 1 second, followed by 1 second of feedback presentation after the second-stage choice. The sequence of events in each trial is illustrated in **Fig. 1a**. We employed the simplified version of the two-stage decision task described by Decker, et al. [46], which replaces abstract fractal cues with easily discriminable "spaceship" and "alien" stimuli. This adaptation with an embedded story line lowers perceptual and working-memory demands, accommodating the cognitive impairments often observed in psychosis, and provides a more intuitive, engaging environment for participants. Each trial included two stages: In the first stage, participants used an MRI-compatible button box to select one of two spaceships. Choices had to be made within 2 seconds; otherwise, the trial was aborted. Once a choice was made, the selected option moved to the top of the screen, while the unselected option disappeared. In the second stage, participants were presented with one of two new states (two red or two purple aliens) each containing two additional options. The participants selected one option, and the outcome was displayed. In the reward setting, the outcome was either a monetary reward (represented by a coin) or no reward (represented by a red cross). In the loss setting, the outcome was either a monetary loss (represented by a broken coin) or no loss (represented by a green cross). Trials were separated by an inter-trial interval of randomized duration, averaging approximately one repetition time (TR). The second-stage state presented was probabilistically determined by the first-stage choice, based on the transition scheme shown in **Fig. 1a**. The position (left or right) of the options at each state was randomly varied between trials. Reward probabilities for each bottom-stage option were dynamic, fluctuating with independent Gaussian noise [7]. This dynamic structure encouraged ongoing learning. Before the fMRI task, participants completed a computerized training session to understand the task structure. They were informed that reward probabilities would vary throughout the task, but the transition probabilities between the first

and second stages would remain fixed. Participants were also told that each first-stage choice was predominantly associated with one of the two second-stage states, and the specific mapping was disclosed. To familiarize themselves with the task, participants completed 50 practice trials prior to the scanning session. Participants were informed that they would receive 50 CHF for task completion. Four participants with early psychosis did not engage with the task (always selected the same response option/pressed the same button during one stage)—and were therefore excluded from further analysis.

Behavior analysis

To examine learning effects without presupposing a particular reinforcement-learning model, we first asked how events on the previous trial shaped choice on the subsequent trial. We computed a stay-probability (StP) index based on the combinations of transition type (common/rare) and outcome (reward/loss). From this, we derived straightforward and intuitive "raw" metrics of model-based (MB_{raw}) and model-free (MF_{raw}) control (**Supplementary Figure 1**). Model-based and model-free raw scores were calculated separately for reward and loss contexts as follows:

Reward context:

$$MB_{raw} = StP_{reward, common} - StP_{reward, rare} - StP_{no\ reward, common} + StP_{no\ reward, rare}$$

$$MF_{raw} = StP_{reward, common} + StP_{reward, rare} - StP_{no\ reward, common} - StP_{no\ reward, rare}$$

Loss context:

$$MB_{raw} = StP_{no\ loss, common} - StP_{no\ loss, rare} - StP_{loss, common} + StP_{loss, rare}$$

$$MF_{raw} = StP_{no\ loss, common} + StP_{no\ loss, rare} - StP_{loss, common} - StP_{loss, rare}$$

Separately for each context, we defined the balance between model-based and model-free control (B_{raw}) as the difference between the respective raw scores, providing a straightforward and intuitive metric to quantify their relative contributions:

$$B_{raw} = MB_{raw} - MF_{raw}$$

To quantify the context-dependent difference in the deployment of model-based versus model-free control, we defined the raw context bias (CB_{raw}) as the difference between the B_{raw} scores obtained in the reward and loss contexts, computed as follows:

$$CB_{raw} = B_{raw, reward} - B_{raw, loss}$$

We additionally deployed a hierarchical Bayesian reinforcement-learning model that incorporates the full choice history and individual learning rates, providing a generative, mechanistically explicit estimate of model-based versus model-free control. We fitted a hybrid reinforcement-learning model that integrates two algorithms: (i) model-based control, which computes action values through prospective (recursive) planning over the task's transition structure; and (ii) model-free control, which updates action values via reward-prediction errors, strengthening actions that are rewarded and weakening those followed by loss [7,47]. A key parameter, denoted as ω, determines the balance between these systems in shaping first-stage choices—where their predicted values diverge most. This weighting parameter is fixed across trials, with $\omega = 1$ reflecting fully model-based control, and $\omega = 0$ indicating purely model-free influence. Further model specifications and fitting procedures are detailed in the supplementary materials. Similar to ω, B_{raw} quantifies the relative contribution of model-based versus model-free learning, with higher values indicating greater reliance on model-based control. To assess the correspondence between the two metrics, we computed their Spearman correlation. To quantify the context-dependent difference in the deployment of model-based relative to model-free control, we defined the context bias (CB) as the difference between the ω scores obtained in the reward and loss contexts (**Fig. 1**):

$$CB = \omega_{reward} - \omega_{loss}$$

Hierarchical Bayesian parameter estimation was implemented through Stan (v. 2.18.2). Prior research has demonstrated that hierarchical parameter estimation provides more stable and reliable estimates compared to individual-level estimation [48]. Posterior distributions were approximated using the Markov Chain Monte Carlo (MCMC) algorithm. Each model was run with four chains, comprising 1000 burn-in samples and 5000 post-burn-in samples. Mixing and convergence of the chains to stationary values were assessed by computing the convergence index ($\Box$) and through visual inspection [48]. Both in the reward and loss setting, MCMC samples of all model parameters were well-mixed and converged to stationary distributions (**Supplementary Fig. 2 and 3**). Convergence to stationary distributions was additionally indicated by the index of convergence ($\Box$ across parameters: 1.00 to 1.04 in the reward and loss setting). For the two-stage decision task, we utilized the *ts_par7* model that is implemented in *hBayesDM package in R* and incorporated both reward and loss contexts [48]. This model quantifies trial-by-trial decision-making dynamics, using 7 parameters, $\alpha 1$ (learning rate in stage 1), $\beta 1$ (inverse temperature in stage 1), $\alpha 2$ (learning rate in stage 2), $\beta 2$ (inverse temperature in stage 2), pi (perseverance), ω (balance between model-based and model-free learning) and lambda (eligibility trace). To determine the optimal specification, we also fitted 4- and 6-parameter variants of the two-stage reinforcement-

learning model and evaluated their performance using the leave-one-out information criterion (LOOIC) [48]. LOOIC favoured the 7-parameter model over both the 4- and 6-parameter models (**Supplementary Table 1**).

fMRI Scanning protocol

Two consecutive fMRI scans were recorded. During the first (resting state), participants were instructed to relax and keep their eyes open. In the second scan, participants performed the modified version of the two-stage task in a "reward" and in a "loss" setting. Participants were monitored to ensure they remained awake and relaxed, wearing a respiratory belt for tracking breathing, and observed with a camera. MRI was conducted on a Philips Achieva 3.0 T magnetic resonance scanner with a 32-channel SENSE head coil. During the 22-minute resting state sequence, 900 echo-planar imaging volumes were recorded (TR = 2000 ms, TE = 30 ms, matrix = 64 x 64, FOV = 240 mm, slice thickness = 3 mm, 40 continuous axial oblique slices). Participants were provided with earplugs, headphones, and head and back support to minimize motion. Each session began with the acquisition of structural images, including a sagittal localizer, a high-resolution T1-weighted SPGR 3D spoiled gradient sequence (TR = 7.5 ms, TE = 3 ms, matrix = 256 × 256, FOV = 240 mm), and a T2 FSE sequence (TR = 6434 ms, TE = 102 ms, matrix = 256 × 256, slice thickness = 2.5 mm). Wakefulness was monitored throughout the scan, and customized head molds were used to reduce head motion.

Anatomical Data Processing

The T1w image was corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection [49], distributed with ANTs 2.5.0 [50], and used as T1w-reference throughout the workflow. The T1w-reference was then skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as the target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM), and gray-matter (GM) was performed on the brain-extracted T1w using fast [51]. Volume-based spatial normalization to one standard space (MNI152NLin2009cAsym) was performed through nonlinear registration with antsRegistration (ANTs 2.5.0) [50], using brain-extracted versions of both T1w reference and the T1w template. The following template was selected for spatial normalization and accessed with TemplateFlow: ICBM 152 Nonlinear Asymmetrical template version 2009c [52].

Functional Data Processing

There was one BOLD run per subject. For pre-processing, a reference volume was first generated, using a custom methodology of *fMRIPrep*, for use in head motion correction. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt [53]. The BOLD reference was then co-registered to the T1w reference using mri_coreg (FreeSurfer) followed by flirt with the boundary-based registration cost-function. Co-registration was configured with six degrees of freedom. Several confounding time series were calculated based on the *preprocessed BOLD*: framewise displacement (FD), DVARS, and three region-wise global signals. FD was computed following Power, et al. [54] and Jenkinson, et al. [53]. FD and DVARS are calculated for each functional run, both using their implementations in *Nipype* [54]. The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow component-based noise correction [55]. Principal components are estimated after high-pass filtering the *preprocessed BOLD* time-series (using a discrete cosine filter with 128s cut-off) for the two *CompCor* variants: temporal (tCompCor) and anatomical (aCompCor). tCompCor components are then calculated from the top 2% variable voxels within the brain mask. For aCompCor, three probabilistic masks (CSF, WM and combined CSF + WM) are generated in anatomical space. The implementation differs from that of Behzadi, et al. [55] in that instead of eroding the masks by 2 pixels on BOLD space, a mask of pixels that likely contain a volume fraction of GM is subtracted from the aCompCor masks. This mask is obtained by thresholding the corresponding partial volume map at 0.05, and it ensures components are not extracted from voxels containing a minimal fraction of GM. Finally, these masks are resampled into BOLD space and binarized by thresholding at 0.99 (as in the original implementation). Components are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the *k* components with the largest singular values are retained, such that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration. The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each [56]. Frames that exceeded a threshold of 0.5 mm FD or 1.5 standardized DVARS were annotated as motion outliers. Additional nuisance time series are calculated by means of principal components analysis of the signal found within a thin band (*crown*) of voxels around the edge of the brain. All resamplings can be performed

with *a single interpolation step* by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using nitransforms, configured with cubic B-spline interpolation. After images were preprocessed with fMRIPrep, preprocessed images were smoothed using a Gaussian kernel with full width at half-maximum of 8 mm and were high-pass filtered with a cutoff of 128s. Four participants with schizophrenia were omitted from further fMRI analysis: two owing to excessive head motion during fMRI acquisition, and two because scanner malfunction produced incomplete imaging data.

Analysis of fMRI Data.

To investigate neural activity associated with task conditions, an event-related analysis was implemented using generalized linear modeling in FSL-FEAT. The analysis was based around the timeseries of Reward Prediction Errors (RPEs) as generated from the simulation of the model over each subject's task choices. This approach followed a framework inspired by Daw, et al. [7], concentrating on prediction errors occurring at two critical moments within each trial: the onset of second-stage stimuli and the onset of reward delivery.

For the first-level analysis, two onset regressors were included in the model (i) first-stage onset (ii) second stage and feedback onset time incorporated into the model as a single regressor. The first onset regressor was parametrically modulated by the value of the chosen first-stage action and its partial derivative with respect to w (see also Supplementary Information). On the other hand, the second onset regressor was modulated by model-free prediction error and by the difference of model-free and model-based reward prediction error and its partial derivative with respect to w. Note that this difference regressor is zero at reward delivery as both algorithms converge at this point. To avoid any confound of the neural results due to activity differences between these two time points per se, the difference regressor was mean-centered within subject. Further, the analysis also includes nuisance onset at feedback. Additionally, to correct for artefacts caused by physiological noise we included 6 motion parameters as estimated by fmriprep, a regressor for framewise displacement, a regressor to censor frames in the fmri volumes with FD > 3mm, and additional five noise regressors calculated using anatomical CompCor. To quantify context-dependent differences in neural activation associated with model-based versus model-free control, we defined the neural context bias (nCB) as the difference between β coefficients for model-based reward prediction errors estimated in reward versus loss contexts (**Fig. 1**):

$$nCB = \beta_{reward} - \beta_{loss}$$

Region-of-interest and whole brain analysis

The three prefrontal areas which are arguably most essential for the computation of model-based control in mammals, were selected as primary regions of interest (ROIs) to explore their activation patterns during model-based control in different contexts. These ROIs included the ACC, dlPFC and OFC [40,41,43]. While it can be argued that additional regions such as the hippocampus are important for model-based control [57], a quantitative whole-brain meta-analysis in human subjects did not confirm this for standard two-stage decision tasks and concluded that the hippocampus plays a substantial role primarily when tasks include explicit spatial relations between decision stages [41]. We excluded striatal regions as primary ROIs because they encode a mixture of model-based and model-free signals, whereas our focus was on regions selective for model-based control [7,41]. ROI masks were taken from the Harvard-Oxford cortical atlas to ensure anatomical precision (see also Supplementary Information). Due to the absence of explicit labels for certain prefrontal subregions in the Harvard–Oxford atlas, the middle frontal gyrus was used as a proxy for the dlPFC [58]. Each ROI was further divided into left and right hemispheres to allow for lateralized analysis. These resulting binary masks were applied to extract the mean BOLD activity within each region for subsequent analyses. This anatomically guided approach ensured consistency across subjects and enhanced interpretability of ROI-based neural activation patterns during task performance. Using the Glasser parcellation, we additionally examined the spatial distribution of activation within each ROI. This analysis provided a finer-grained view of how each region supports model-based control in the reward and loss contexts. Because our ROI analysis was hypothesis-driven—and therefore might overlook regions that contribute to model-based control—we also conducted an exploratory whole-brain analysis.

Results

Characteristic	HC	SSD	p-value
Demographics			
Age (years), mean (SD)	28.04 (5.52)	26.92 (6.39)	0.350
Sex (% female)	21 (42.9)	12 (23.1)	0.057
Ethnicity, n (% Caucasian)	39 (79.6)		
% Asian	6 (12.2)		0.989
% African	4 (8.2)		
Schooling (years), mean (SD)	2.27 (1.30)	1.62 (1.17)	0.010
Illness Characteristics			
Disease Duration (months), mean (SD)	NA	19.21 (29.41)	
Hospitalizations, mean (SD)	NA	7.18 (5.11)	
Medication Status			
Dosage, antipsychotic (CPZ eq.)	NA	271.35 (160.57)	
Clinical Ratings			
BPRS Total Score, mean (SD)	NA	39.10 (10.40)	
Suspicion	NA	1.94 (1.09)	
Hallucination	NA	1.88 (1.15)	
BNSS Total Score, mean (SD)	NA	40.67 (15.25)	
Asociality	NA	6.08 (2.75)	
Avolition	NA	6.60 (2.75)	
Anhedonia	NA	10.88 (4.67)	
Distress	NA	3.31 (1.52)	
Blunted Affect	NA	9.40 (4.53)	
Alogia	NA	4.37 (2.51)	
Cognitive Performance			
MATRICS - Attention, T-score	NA	0.97 (1.06)	
MATRICS - Working Memory, T-score	NA	1.47 (0.84)	
MATRICS - Verbal Learning, T-score	NA	1.61 (1.13)	
MATRICS - Reasoning, T-score	NA	1.47 (1.03)	
MATRICS - Visual Learning, T-score	NA	1.53 (0.88)	

Table 1. Summary of participant characteristics. Clinical symptoms were assessed using the Brief Psychiatric Rating Scale (BPRS) for general psychopathology and the Brief Negative Symptom Scale

(BNSS) for negative symptoms. Cognitive performance was measured using the MATRICS Consensus Cognitive Battery (MCCB) and is reported as standardized T-scores. Antipsychotic medication dosage is given in chlorpromazine equivalents (CPZ eq.). Values are expressed as mean (standard deviation) or count (percentage). *Schooling years* are coded as follows: 0 = mandatory school, 1 = SEK/Gymnasium, 2 = Lehre, 3 = Hochschulen, 4 = university. HC: Healthy controls; SSD: Individuals with schizophrenia spectrum disorders.

The early psychosis cohort (n = 50) and healthy controls (n = 48) showed no significant differences in age, sex, or ethnicity. However, participants with early psychosis had completed fewer years of post-secondary education compared to controls (psychosis: 1.62 ± 1.17; controls: 2.27 ± 1.30 years, p = 0.01). All patients received antipsychotic medication. Symptom severity was mild to moderate (BPRS total score: 39.1 ± 10.4). Negative symptom ratings on the BNSS showed both apathy-related symptoms (anhedonia: 10.9 ± 4.7; avolition: 6.6 ± 2.8; asociality: 6.1 ± 2.8) and diminished-expression components (blunted affect: 9.4 ± 4.5; alogia: 4.4 ± 2.5). Cognitive screening using the MATRICS battery indicated deficits in attention (mean T-score: 1.0 ± 1.1) and working memory (mean T-score: 1.5 ± 0.8).

Figure 1: Experimental setting and modeling hypothesis. (a) *Two-stage decision-making task.* Participants make two consecutive choices to either earn play coins (*reward setting*) or avoid losing play coins (*loss setting*). Probabilistic transitions between the first and second stages (common vs. rare) allow for the quantification of model-based and model-free learning. Adapted from Decker, et al. [46]. **(b)** Schematic of the human pre-frontal network supporting model-based reinforcement learning. Dorsolateral pre-frontal cortex (dlPFC) supports planning and working-memory processes; orbitofrontal cortex (OFC) encodes prospective option values; anterior cingulate cortex (ACC) integrates these signals to allocate cognitive control. **(c)** Schematic showing the computation of context biases in goal-directed control. For each participant, the model-based weight ω and the task-evoked BOLD parameter β are estimated separately in *reward* and *loss* contexts (grey bars). Context bias is defined as the difference $\Delta\omega = \omega_{reward} - \omega_{loss}$ for behavior and $\Delta\beta = \beta_{reward} - \beta_{loss}$ for neural activation. Positive Δ indicates a *reward bias* (back bar with arrow pointing up), whereas negative Δ indicates a *loss bias* (back bar with arrow pointing down).

Across all contexts and groups, the hybrid model's weight parameter ω showed a strong correlation with our straightforward metric B_{raw} which similarly quantifies the relative contributions of model-based and model-free control (r = 0.60–0.70, p ≤ 1*10⁻☐). The correspondence supports that B_{raw} of our simplified index also captures the model-based/model-free balance quantified by the computational model.

Figure 2: Analysis of decision behavior: (a) *Simplified raw behavior scores:* Regarding the relative contribution of model-based learning, healthy controls (blue) exhibited a trend toward a reward bias, whereas individuals with schizophrenia spectrum disorders (SSD) showed a trend toward a loss bias. When the context bias of raw behavior scores was compared between groups the difference was not significant. **(b)** *ω-scores from the modelling analysis:* Regarding the relative contribution of model-based learning, the ω-scores indicated a significant reward bias in healthy controls and a significant loss bias in individuals with SSD. In addition, the context bias of ω-values differed significantly between groups. Mean and 95% confidence intervals are shown. Asterisk indicates *p* < 0.05.

In the raw behavioral analysis, individuals with psychosis showed reduced model-based control in both the reward and loss contexts (reward: p = 0.004; loss: p = 0.014). Healthy controls showed a trend toward a reward bias and individuals with psychosis exhibited a trend toward a loss bias, but the group difference was not significant (**Fig. 2a**). When decision behavior was analyzed using the more sophisticated hybrid model, the weight parameter (ω) indicated that individuals with psychosis exhibited reduced model-based control in the reward context (p = 7.9*10⁻¹³), but not in the loss context (**Supplementary Fig. 4 and 5**). Moreover, the hybrid model revealed a reward bias for healthy controls (p = 3.5*10⁻⁴) and a loss bias for individuals with psychosis (p = 2.9*10⁻⁹), and the context bias in controls and individuals with psychosis differed significantly between groups (p = 2.8*10⁻¹⁵; **Fig. 2b**).

Context-dependent neural activations: In the ACC, healthy controls showed a significant reward bias in both hemispheres (*signrank test*; left: $p = 4*10^{-\square}$; right: $p = 6*10^{-\square}$), while individuals with schizophrenia exhibited a significant loss bias (*signrank test*; left: $p = 0.048$; right: $p = 0.04$). Group comparisons revealed a significant difference in context bias for both the left and right ACC (*ranksum test*; both $p = 2*10^{-\square}$, **Fig. 3**). In the dlPFC, a significant reward bias was found in healthy controls (*signrank test*; left: $p = 0.001$; right: $p = 0.008$).

Figure 3: Context-dependent BOLD activation during model-based control.
The context bias of activation in the anterior cingulate cortex (ACC), orbitofrontal cortex (OFC) and dorsolateral prefrontal cortex (dlPFC) for healthy controls (blue squares) and individuals with schizophrenia spectrum disorders (orange triangles) is shown, plotted separately for the left and right hemispheres. Symbols indicate group means and 95 % confidence intervals. Asterisk indicates p < 0.05.

This reward bias differed significantly from the response in individuals with psychosis in the left dlPFC (**Fig. 3**), where no reward or loss bias was detected (*ranksum test*; $p = 0.02$). In the OFC, healthy controls showed a significant reward bias in both hemispheres (*signrank test*; left: $p = 6*10^{-\square}$; right: $p = 0.01$), while individuals with psychosis demonstrated a significant shift toward a loss bias (*ranksum test*; left: $p = 6*10^{-\square}$; right: $p = 0.001$, **Fig. 3**).

Figure 4: Right-ACC context bias and negative-symptom domains. Spearman correlations between the right-ACC context bias (reward minus loss activation during model-based learning) and each BNSS negative-symptom domain are displayed. Within the *apathy* dimension, only anhedonia showed a significant association with right-ACC bias; avolition and asociality did not reach significance. No domain in the *diminished-expression* dimension (blunted affect, alogia) correlated with the ACC measure. Bars represent 95% confidence intervals; asterisk indicates $p < 0.05$.

We next examined whether context bias in any of these regions was associated with negative symptom domains. Only the context bias in the right ACC showed a significant negative correlation with the symptom domain of anhedonia (Spearman correlation $r_s = 0.4$; $p = 0.008$; **Fig. 4**; **Supplementary Fig. 6**). This relationship indicates that greater loss-biased activation in rACC is associated with greater anhedonia. As additional control, we investigated the association between right ACC context bias and anhedonia using a linear model with age, gender, disease duration and chlorpromazine equivalents as covariates. This analysis confirmed a negative association between right ACC context bias and anhedonia ($p = 0.005$).

Outside the a-priori ROIs (ACC, OFC, dlPFC), healthy controls showed extensive reward-biased activation along the medial cortical wall in the exploratory brain analysis (**Fig. 5**). Large clusters spanned the dorsomedial prefrontal cortex, pre-supplementary motor area (pre-SMA) and posterior cingulate/precuneus regions. In the early-psychosis group the pattern flipped but was markedly sparser: only small foci in the anterior and posterior cingulate cortex and adjacent pre-SMA expressed a significant loss bias. Thus, whereas healthy participants recruit a broad medial fronto-parietal network more strongly for reward-based control, patients exhibit a selective shift toward loss-biased engagement confined to the anterior cingulate–pre-SMA axis.

Figure 5: Whole-brain analysis reveals dissociable valence effects on model-based learning networks in individuals with psychosis and healthy controls. Surface projections of significant z-scores are displayed for healthy controls (**a**) and individuals with psychosis (**b**). Warm colors indicate brain regions with greater model-based signals for loss outcomes compared to reward outcomes, while cool colors indicate the reverse pattern. For visualization, Z-scores are projected onto the pial surface of a template brain.

Discussion

In the present study we provide converging behavioral and neural evidence that early psychosis is characterized by a systematic re-allocation of goal-directed cognitive resources away from reward pursuit and toward loss avoidance. Relative to the reward-biased deployment of model-based strategies in healthy controls, individuals with psychosis showed model-based control primarily when potential losses were at stake. This loss-biased strategy was mirrored at the neural level by preferential engagement of anterior cingulate (ACC), orbitofrontal (OFC) and dorsolateral prefrontal cortex (dlPFC) during model-based learning in the loss condition, and the degree of loss-biased ACC activation tracked the severity of anhedonia. Together, these results offer a mechanistic explanation for the clinical observation that individuals with schizophrenia show impaired pursuit of positive outcomes yet intact loss avoidance [3,4], while linking this asymmetry to the negative symptom domain anhedonia. It is worth mentioning that the raw behavioral analysis showed only trend-level evidence for a context effect, while the hierarchical reinforcement learning model revealed a robust difference. By fitting the full choice history and jointly estimating learning-rate, perseveration and lapse noise, the hierarchical computational model isolates the underlying goal-directed signal that may be masked by excess variability in the raw behavioral analysis.

Hierarchical pooling further reduces within-group noise through shrinkage, increasing power to detect true between-group differences [48].

Canonical reinforcement-learning theory assumes that model-based control is deployed when the expected benefits of flexible planning outweigh its cognitive costs [7]. In healthy adults, appetitive contexts typically increase the perceived benefit of planning, leading to the reward-biased pattern observed in our control sample and in prior work [41,43]. The reversed, loss-biased pattern in schizophrenia suggests that for these individuals the subjective cost–benefit calculus is distorted: avoiding losses is deemed more valuable than securing gains. One parsimonious interpretation is that patients discount prospective rewards because they anticipate diminished pleasure from realizing potential gains [59], whereas losses remain fully salient. In line with this idea, low "decision acuity" — a trait-like reduction in overall planning capacity marked by noisier, short-sighted choices [60] — could further restrict cognitive resources and channel them chiefly into preventing losses, thereby producing the loss-biased goal-directed control we observe. A complementary line of evidence shows that people with schizophrenia or bipolar psychosis are less willing to expend cognitive effort even for large rewards, pointing to a domain-general undervaluation of appetitive incentives [61,62]. These findings align with evidence for impaired value representation during reward processing relative to preserved loss avoidance in schizophrenia [3], and our results frame this as the consequence of a strategic shift toward loss-oriented planning. Our fMRI results reveal how this motivational bias is implemented in the prefrontal cognitive network. In controls, reward-related increases in model-based control were accompanied by greater activation of ACC, OFC and dlPFC - regions known to orchestrate cost-dependent allocation of cognitive resources, value representation, and prospective planning according to goal-representation [35,44]. In schizophrenia spectrum disorders, the same circuit was preferentially engaged for model-based control when potential losses, rather than rewards, were at stake. These findings extend prior observations of blunted OFC and dlPFC responses to positive outcomes [6] by demonstrating that computations underpinning flexible decision-making are context-dependent and retuned toward loss avoidance in psychotic disorders. The ACC emerged as a particularly pivotal hub: its loss-biased engagement was the only neural measure that covaried with a negative symptom domain. The ACC is ideally positioned to integrate motivational signals with task-control demands and to evaluate whether engaging the additional cognitive effort required for model-based learning is worthwhile [40,63]. A loss-dominated encoding of action value within ACC could therefore funnel limited cognitive resources into preventing negative outcomes, leaving insufficient capacity for goal-directed reward processing and manifesting clinically as anhedonia. Our data reinforce converging evidence that ACC dysfunction is central to schizophrenia, underlying failures of proactive

control and showing uniform structural features in schizophrenia [35,39,64]. Identifying a loss-biased planning mode in schizophrenia has potential translational relevance. First, it provides a neurocomputational phenotype that could be quantified and used to stratify patients or monitor treatment responses. Second, the link between anhedonia and ACC loss bias supports the ACC as a target for circuit-based treatments. Non-invasive brain stimulation targeting the ACC or novel pharmacological interventions may re-balance resource allocation and thereby ameliorate motivational deficits. Our findings also extend earlier active-inference results showing elevated forgetting and reduced policy precision during action selection in early psychosis [65], and place that computational profile in a motivational framework that distinguishes loss from reward contexts.

This study has several limitations. First, the study is cross-sectional and cannot establish whether loss-biased planning predates illness onset, emerges progressively or reflects a secondary adaptation to chronic disability. Longitudinal studies in high-risk and first-episode cohorts are needed to pinpoint when loss-biased cognitive processing emerges in schizophrenia. Secondly, participants were medicated, and modulation of cognitive processing and behavioral strategies by the complex receptor modulation of antipsychotic agents cannot be excluded. Thirdly, a recent study showed that modelling parameters in reinforcement tasks can vary over time, potentially limiting their capacity for precision psychiatry [66]. Moreover, the monetary outcomes used here do not capture the rich spectrum of rewards and losses encountered in daily life, and virtual-reality tasks embedding social, vocational and other incentives could enhance ecological validity.

Conclusions

By integrating computational modelling with neuroimaging, we demonstrate that schizophrenia spectrum disorders are not marked by a uniform deficit in goal-directed cognition but by a selective redirection of cognitive resources toward loss avoidance. This behavioural shift is instantiated in a prefrontal network comprising ACC, OFC and dlPFC, and loss-biased engagement of the ACC is coupled to the negative symptom domain anhedonia. Our findings provide a mechanistic explanation for the preservation of avoidance learning alongside impaired reward pursuit in schizophrenia and outline circuit-level targets for interventions aimed at restoring adaptive motivation.

References

Marder, S. R. & Cannon, T. D. Schizophrenia. *N Engl J Med* **381**, 1753-1761, doi:10.1056/NEJMra1808803 (2019).

McCutcheon, R. A., Reis Marques, T. & Howes, O. D. Schizophrenia-An Overview. *JAMA Psychiatry* **77**, 201-210, doi:10.1001/jamapsychiatry.2019.3360 (2020).

Gold, J. M. *et al.* Negative symptoms and the failure to represent the expected reward value of actions: behavioral and computational modeling evidence. *Arch Gen Psychiatry* **69**, 129-138, doi:10.1001/archgenpsychiatry.2011.1269 (2012).

Marder, S. R. & Umbricht, D. Negative symptoms in schizophrenia: Newly emerging measurements, pathways, and treatments. *Schizophr Res* **258**, 71-77, doi:10.1016/j.schres.2023.07.010 (2023).

Kring, A. M. & Elis, O. Emotion deficits in people with schizophrenia. *Annu Rev Clin Psychol* **9**, 409-433, doi:10.1146/annurev-clinpsy-050212-185538 (2013).

Strauss, G. P., Waltz, J. A. & Gold, J. M. A review of reward processing and motivational impairment in schizophrenia. *Schizophr Bull* **40 Suppl 2**, S107-116, doi:10.1093/schbul/sbt197 (2014).

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, doi:10.1016/j.neuron.2011.02.027 (2011).

Lee, S. W., Shimojo, S. & O'Doherty, J. P. Neural computations underlying arbitration between model-based and model-free learning. *Neuron* **81**, 687-699, doi:10.1016/j.neuron.2013.11.028 (2014).

Shahar, N. *et al.* Credit assignment to state-independent task representations and its relationship with model-based decision making. *Proc Natl Acad Sci U S A* **116**, 15871-15876, doi:10.1073/pnas.1821647116 (2019).

Deserno, L. *et al.* Dopamine enhances model-free credit assignment through boosting of retrospective model-based inference. *Elife* **10**, doi:10.7554/eLife.67778 (2021).

Feher da Silva, C., Lombardi, G., Edelson, M. & Hare, T. A. Rethinking model-based and model-free influences on mental effort and striatal prediction errors. *Nat Hum Behav* **7**, 956-969, doi:10.1038/s41562-023-01573-1 (2023).

Otto, A. R., Gershman, S. J., Markman, A. B. & Daw, N. D. The curse of planning: dissecting multiple reinforcement-learning systems by taxing the central executive. *Psychol Sci* **24**, 751-761, doi:10.1177/0956797612463080 (2013).

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 U S A* **110**, 20941-20946, doi:10.1073/pnas.1312011110 (2013).

Karthik, S. K., Pereira, K. M. K., Culbreth, A. J., Barch, D. M. & Moran, E. K. Higher motivation and pleasure scores predict more reliance on model-free decision making. *Cogn Affect Behav Neurosci*, doi:10.3758/s13415-025-01302-3 (2025).

Daw, N. D. Are we of two minds? *Nat Neurosci* **21**, 1497-1499, doi:10.1038/s41593-018-0258-2 (2018).

Kool, W., Cushman, F. A. & Gershman, S. J. When Does Model-Based Control Pay Off? *PLoS Comput Biol* **12**, e1005090, doi:10.1371/journal.pcbi.1005090 (2016).

Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A. & Daw, N. D. Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *Elife* **5**, doi:10.7554/eLife.11305 (2016).

Seow, T. X. F. *et al.* Model-Based Planning Deficits in Compulsivity Are Linked to Faulty Neural Representations of Task Structure. *J Neurosci* **41**, 6539-6550, doi:10.1523/JNEUROSCI.0031-21.2021 (2021).

Vaghi, M. M. *et al.* Compulsivity is linked to reduced adolescent development of goal-directed control and frontostriatal functional connectivity. *Proc Natl Acad Sci U S A* **117**, 25911-25922, doi:10.1073/pnas.1922273117 (2020).

Grosskurth, E. D., Bach, D. R., Economides, M., Huys, Q. J. M. & Holper, L. No substantial change in the balance between model-free and model-based control via training on the two-step task. *PLoS Comput Biol* **15**, e1007443, doi:10.1371/journal.pcbi.1007443 (2019).

Sebold, M. *et al.* When Habits Are Dangerous: Alcohol Expectancies and Habitual Decision Making Predict Relapse in Alcohol Dependence. *Biol Psychiatry* **82**, 847-856, doi:10.1016/j.biopsych.2017.04.019 (2017).

Drummond, N. & Niv, Y. Model-based decision making and model-free learning. *Curr Biol* **30**, R860-R865, doi:10.1016/j.cub.2020.06.051 (2020).

Doll, B. B., Simon, D. A. & Daw, N. D. The ubiquity of model-based reinforcement learning. *Curr Opin Neurobiol* **22**, 1075-1081, doi:10.1016/j.conb.2012.08.003 (2012).

Collins, A. G. E. & Cockburn, J. Beyond dichotomies in reinforcement learning. *Nat Rev Neurosci* **21**, 576-586, doi:10.1038/s41583-020-0355-6 (2020).

Karagoz, A. B., Moran, E. K., Barch, D. M., Kool, W. & Reagh, Z. M. Evidence for shallow cognitive maps in Schizophrenia. *Cogn Affect Behav Neurosci*, doi:10.3758/s13415-025-01283-3 (2025).

Konovalov, A. & Krajbich, I. Gaze data reveal distinct choice processes underlying model-based and model-free reinforcement learning. *Nat Commun* **7**, 12438, doi:10.1038/ncomms12438 (2016).

Economides, M., Kurth-Nelson, Z., Lubbert, A., Guitart-Masip, M. & Dolan, R. J. Model-Based Reasoning in Humans Becomes Automatic with Training. *PLoS Comput Biol* **11**, e1004463, doi:10.1371/journal.pcbi.1004463 (2015).

Knolle, F. *et al.* Investigating disorder-specific and transdiagnostic alterations in model-based and model-free decision making. *J Psychiatry Neurosci* **49(6)**, E389-E401 (2024).

Voon, V. *et al.* Disorders of compulsivity: a common bias towards learning habits. *Mol Psychiatry* **20**, 345-352, doi:10.1038/mp.2014.44 (2015).

Voon, V. *et al.* Motivation and value influences in the relative balance of goal-directed and habitual behaviours in obsessive-compulsive disorder. *Transl Psychiatry* **5**, e670, doi:10.1038/tp.2015.165 (2015).

Wunderlich, K., Smittenaar, P. & Dolan, R. J. Dopamine enhances model-based over model-free choice behavior. *Neuron* **75**, 418-424, doi:10.1016/j.neuron.2012.03.042 (2012).

Deserno, L. *et al.* Ventral striatal dopamine reflects behavioral and neural signatures of model-based control during sequential decision making. *Proc Natl Acad Sci U S A* **112**, 1595-1600, doi:10.1073/pnas.1417219112 (2015).

Doll, B. B., Bath, K. G., Daw, N. D. & Frank, M. J. Variability in Dopamine Genes Dissociates Model-Based and Model-Free Reinforcement Learning. *J Neurosci* **36**, 1211-1222, doi:10.1523/JNEUROSCI.1901-15.2016 (2016).

Sharp, M. E., Foerde, K., Daw, N. D. & Shohamy, D. Dopamine selectively remediates 'model-based' reward learning: a computational approach. *Brain* **139**, 355-364, doi:10.1093/brain/awv347 (2016).

Barch, D. M., Culbreth, A. J. & Sheffield, J. M. Cognitive Control in Schizophrenia: Advances in Computational Approaches. *Curr Dir Psychol Sci* **33**, 35-42, doi:10.1177/09637214231205220 (2024).

Culbreth, A. J., Westbrook, A., Daw, N. D., Botvinick, M. & Barch, D. M. Reduced model-based decision-making in schizophrenia. *J Abnorm Psychol* **125**, 777-787, doi:10.1037/abn0000164 (2016).

Collins, A. G. E., Albrecht, M. A., Waltz, J. A., Gold, J. M. & Frank, M. J. Interactions Among Working Memory, Reinforcement Learning, and Effort in Value-Based Choice: A New Paradigm and Selective Deficits in Schizophrenia. *Biol Psychiatry* **82**, 431-439, doi:10.1016/j.biopsych.2017.05.017 (2017).

Collins, A. G., Brown, J. K., Gold, J. M., Waltz, J. A. & Frank, M. J. Working memory contributions to reinforcement learning impairments in schizophrenia. *J Neurosci* **34**, 13747-13756, doi:10.1523/JNEUROSCI.0989-14.2014 (2014).

Omlor, W. *et al.* Estimating Multimodal Structural Brain Variability in Schizophrenia Spectrum Disorders: A Worldwide ENIGMA Study. *Am J Psychiatry* **182**, 373-388, doi:10.1176/appi.ajp.20230806 (2025).

Akam, T. *et al.* The Anterior Cingulate Cortex Predicts Future States to Mediate Model-Based Action Selection. *Neuron* **109**, 149-163 e147, doi:10.1016/j.neuron.2020.10.013 (2021).

Huang, Y., Yaple, Z. A. & Yu, R. Goal-oriented and habitual decisions: Neural signatures of model-based and model-free learning. *Neuroimage* **215**, 116834, doi:10.1016/j.neuroimage.2020.116834 (2020).

Wilson, R. C., Takahashi, Y. K., Schoenbaum, G. & Niv, Y. Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267-279, doi:10.1016/j.neuron.2013.11.005 (2014).

Smittenaar, P., FitzGerald, T. H., Romei, V., Wright, N. D. & Dolan, R. J. Disruption of dorsolateral prefrontal cortex decreases model-based in favor of model-free control in humans. *Neuron* **80**, 914-919, doi:10.1016/j.neuron.2013.08.009 (2013).

Stein, D. J. *et al.* Obsessive-compulsive disorder. *Nat Rev Dis Primers* **5**, 52, doi:10.1038/s41572-019-0102-3 (2019).

Begue, I., Kaiser, S. & Kirschner, M. Pathophysiology of negative symptom dimensions of schizophrenia - Current developments and implications for treatment. *Neurosci Biobehav Rev* **116**, 74-88, doi:10.1016/j.neubiorev.2020.06.004 (2020).

Decker, J. H., Otto, A. R., Daw, N. D. & Hartley, C. A. From Creatures of Habit to Goal-Directed Learners: Tracking the Developmental Emergence of Model-Based Reinforcement Learning. *Psychol Sci* **27**, 848-858, doi:10.1177/0956797616639301 (2016).

Worbe, Y. *et al.* Valence-dependent influence of serotonin depletion on model-based choice strategy. *Mol Psychiatry* **21**, 624-629, doi:10.1038/mp.2015.46 (2016).

Ahn, W. Y., Haines, N. & Zhang, L. Revealing Neurocomputational Mechanisms of Reinforcement Learning and Decision-Making With the hBayesDM Package. *Comput Psychiatr* **1**, 24-57, doi:10.1162/CPSY_a_00002 (2017).

Tustison, N. J. *et al.* N4ITK: improved N3 bias correction. *IEEE Trans Med Imaging* **29**, 1310-1320, doi:10.1109/TMI.2010.2046908 (2010).

Avants, B. B., Epstein, C. L., Grossman, M. & Gee, J. C. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med Image Anal* **12**, 26-41, doi:10.1016/j.media.2007.06.004 (2008).

Zhang, Y., Brady, M. & Smith, S. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging* **20**, 45-57, doi:10.1109/42.906424 (2001).

Fonov, V. S., Evans, A. C., McKinstry, R. C., Almli, C. R. & Collins, D. L. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *Neuroimage* **47**, S102 (2009).

Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**, 825-841, doi:10.1016/s1053-8119(02)91132-8 (2002).

Power, J. D. *et al.* Methods to detect, characterize, and remove motion artifact in resting state fMRI. *Neuroimage* **84**, 320-341, doi:10.1016/j.neuroimage.2013.08.048 (2014).

Behzadi, Y., Restom, K., Liau, J. & Liu, T. T. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage* **37**, 90-101, doi:10.1016/j.neuroimage.2007.04.042 (2007).

Satterthwaite, T. D. *et al.* An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage* **64**, 240-256, doi:10.1016/j.neuroimage.2012.08.052 (2013).

Miller, K. J., Botvinick, M. M. & Brody, C. D. Dorsal hippocampus contributes to model-based planning. *Nat Neurosci* **20**, 1269-1276, doi:10.1038/nn.4613 (2017).

Schienle, A., Kogler, W. & Wabnegger, A. A randomized trial that compared brain activity, efficacy and plausibility of open-label placebo treatment and cognitive repappraisal for reducing emotional distress. *Sci Rep* **13**, 13998, doi:10.1038/s41598-023-39806-y (2023).

Kring, A. M. & Barch, D. M. The motivation and pleasure dimension of negative symptoms: neural substrates and behavioral outputs. *Eur Neuropsychopharmacol* **24**, 725-736, doi:10.1016/j.euroneuro.2013.06.007 (2014).

Moutoussis, M. *et al.* Decision-making ability, psychopathology, and brain connectivity. *Neuron* **109**, 2025-2040 e2027, doi:10.1016/j.neuron.2021.04.019 (2021).

Barch, D. M. *et al.* Dissociation of Cognitive Effort-Based Decision Making and Its Associations With Symptoms, Cognition, and Everyday Life Function Across Schizophrenia, Bipolar Disorder, and Depression. *Biol Psychiatry* **94**, 501-510, doi:10.1016/j.biopsych.2023.04.007 (2023).

Treadway, M. T., Peterman, J. S., Zald, D. H. & Park, S. Impaired effort allocation in patients with schizophrenia. *Schizophr Res* **161**, 382-385, doi:10.1016/j.schres.2014.11.024 (2015).

Shenhav, A., Botvinick, M. M. & Cohen, J. D. The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* **79**, 217-240, doi:10.1016/j.neuron.2013.07.007 (2013).

Brugger, S. P. & Howes, O. D. Heterogeneity and Homogeneity of Regional Brain Structure in Schizophrenia: A Meta-analysis. *JAMA Psychiatry* **74**, 1104-1111, doi:10.1001/jamapsychiatry.2017.2663 (2017).

Knolle, F. *et al.* Action selection in early stages of psychosis: an active inference approach. *J Psychiatry Neurosci* **48**, E78-E89, doi:10.1503/jpn.220141 (2023).

Vrizzi, S., Najar, A., Lemogne, C., Palminteri, S. & Lebreton, M. Behavioral, computational and self-reported measures of reward and punishment sensitivity as predictors of mental health characteristics. *Nature Mental Health* (2025).