Dopamine reports reward prediction errors, but does not update policy, during inference-guided choice

Marta Blanco-Pozo[1#], Thomas Akam[1#], Mark Walton[1,2*]

Affiliations:

1. Department of Experimental Psychology, Oxford University, Oxford, UK

2. Wellcome Centre for Integrative Neuroimaging, Oxford University, Oxford, U.K.

# Equal contribution

mark.walton@psy.ox.ac.uk

thomas.akam@psy.ox.ac.uk

marta.blancopozo@psy.ox.ac.uk

Abstract

Dopamine is thought to carry reward prediction errors (RPEs), which update values and hence modify future behaviour. However, updating values is not always the most efficient way of adapting to change. If previously encountered situations will be revisited in future, inferring that the *state* of the world has changed allows prior experience to be reused when situations are reencountered. To probe dopamine's involvement in such inference-based behavioural flexibility, we measured and manipulated dopamine while mice solved a sequential decision task using state inference. Dopamine was strongly influenced by the value of states and actions, consistent with RPE signalling, using value information that respected task structure. However, though dopamine responded strongly to rewards, stimulating dopamine at the time of trial outcome had no effect on subsequent choice. Therefore, when inference guides choice, rewards have a dopamine-independent influence on policy through the information they carry about the world's state.

Introduction

Accurately estimating future reward associated with different actions is critical to adaptive behaviour. Reinforcement learning (RL) has provided an influential mechanistic account in which reward prediction errors (RPEs) update values (i.e., predictions of long run reward). Dopamine activity consistent with RPEs is observed in diverse behaviours[1–4], and causal manipulations of dopamine can reinforce or suppress behaviours[5–9].

However, RPE-driven value learning is not always the most efficient way of updating preferences in dynamic environments. In many situations, changes in observable events reflect underlying dynamics which are themselves highly structured. For example, the cyclical change between week and weekend constitutes a *latent* or not directly observable feature of our environment, which affects observable events, with consequences for behaviour. It would be profoundly inefficient to relearn values each Monday and Saturday to cope with this change; instead, by *inferring* that the *state* of the world has changed, we can just redeploy values and policies from a previous encounter. Such structured change is also common in the laboratory, and state inference could in principle underlie behavioural adaptation in many reward-guided tasks[10,11]. Indeed, direct evidence for inference based behavioural flexibility has been demonstrated in several tasks[12,13], and is also suggested by meta-learning, where subjects get faster at adapting to changes with repeated experience[14–17].

Therefore, a pressing issue is to understand how dopamine reflects and shapes behaviour during inference-guided choice. Dopamine activity is informed by inferences about latent states[18–22]. However, little is known about activity in more complex tasks where latent states and causal relationships interact during sequential decision making. This is ubiquitous in the real world; for instance in the earlier example, when the best route home depends on the day of the week. Nor do we know how the observed heterogeneity in dopamine projections to different targets represents such behaviours[6,23–28]. Most critically, we do not know whether dopamine is responsible for the reinforcing effects of natural rewards in situations where inference not value learning underlies behavioural flexibility.

To address these questions, we measured and manipulated dopamine in mice performing a two-step decision task. The task had two important features. First, reward probabilities were anti-correlated and reversed periodically, thus introducing an inferable latent variable. Second, the action-state transition structure meant that under an optimal policy, rewards did not simply reinforce the preceding choice (Fig 1A). Both behaviour and dopamine activity showed evidence that mice inferred a single latent variable that controlled both reward probabilities, and assigned value to actions consistent with the task's structure. Dopamine activity was strongly modulated at specific time points

across the trial by the values of their chosen actions, the states they encountered, and the outcome they obtained. This coexisted with a separate influence of recent reward rate throughout the trial, and signals correlated with lateralised movements. However, stimulating midbrain dopamine neurons at outcome time, using parameters sufficient to reinforce preceding actions, did not recapitulate the effect of natural rewards, and had no effect on subsequent choice. These data demonstrate that when state inference underlies flexible behaviour, dopamine has a limited causal influence on choices, despite prominently reporting a set of task-related variables.

Results

Behaviour reflects task causal and correlation structure

We trained DAT-cre mice (n=18) on a sequential decision task, which required them to choose between two options – left and right – to gain access to one of two reward ports – top or bottom (Figure 1A). Each trial started with a central port lighting up, which mice poked to initiate the trial. This triggered either a choice state, where both left and right ports lit up (75% of trials), or a forced choice where either left or right lit up (25% of trials). Poking an illuminated side port triggered a transition to one of two possible second step states, signalled by an auditory cue, in which either the top or bottom port was illuminated. Each first-step choice (left or right) led commonly (80% of trials) to one second-step state (up or down) and rarely (20% of trials) to the other (Figure 1B). Poking the active reward port caused a 500-ms auditory cue indicating the trial outcome (reward or not), with reward delivered at cue offset. Reward delivery in each port was probabilistic, with probabilities that changed in blocks between *0.8/0.2*, *0.5/0.5* and *0.2/0.8* on the up/down ports respectively. The reward probabilities were therefore anticorrelated at the up and down ports (Figure 1A, B).

Subjects tracked which option was currently best, performing 361 $\pm$ 102 (mean $\pm$ SD) trials and completing 8.4 $\pm$ 4.2 blocks per session (Figure 1C). Rewards did not simply reinforce the choice made at the start of the trial, but rather the trial outcome (rewarded or not) interacted with the state transition (common or rare) to determine the probability of repeating choice (Figure 1D-F). Rewards following common transitions strongly promoted staying while rewards following rare transitions strongly promoted switching ($P < 0.001$, mixed effects logistic regression) (Figure 1D), and these effects persisted over multiple trials (Figure 1F).

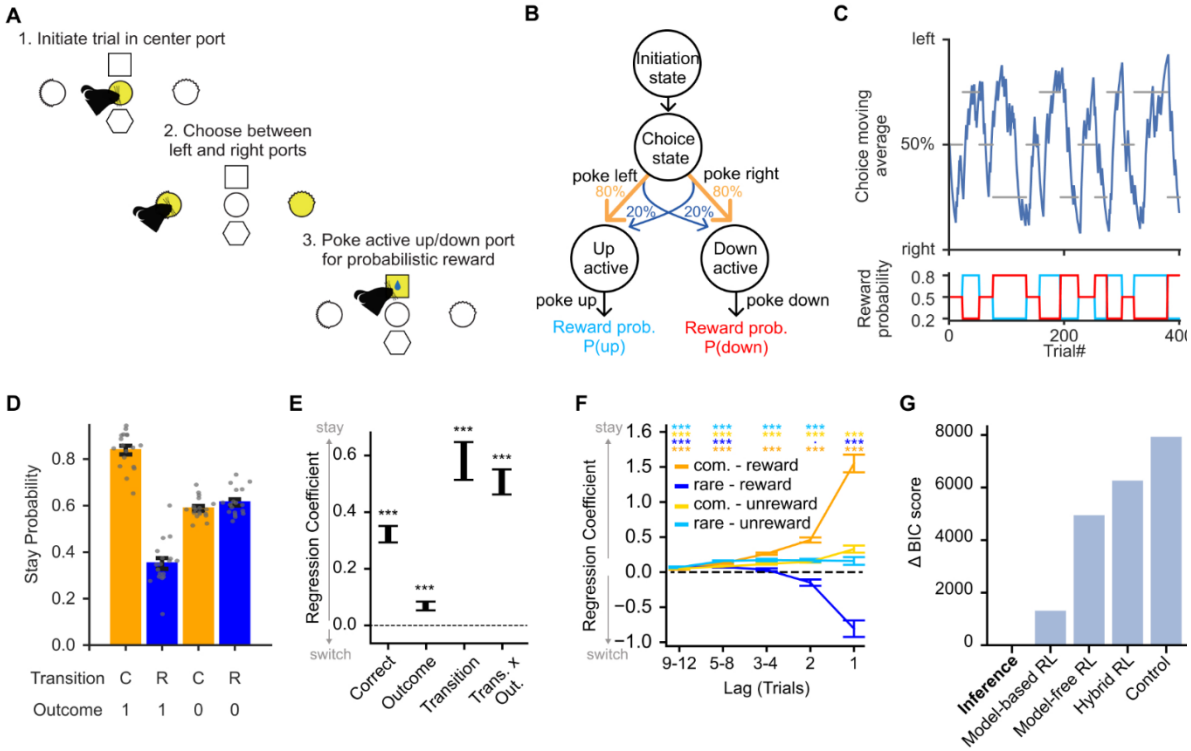

Figure 1. Two step task behaviour. A) Diagram of trial events. **B)** Diagram of the task state space, reward and transition probabilities. **C)** Example behavioural session. Top panel: Moving average of choices (blue trace) with reward probability blocks shown by grey bars, indicating by their vertical position whether the reward probability was higher for the state commonly reached for the left or right choice or a neutral block. Bottom panel: Reward probability at the up (blue) and down (red) ports. **D)** Probability of a choice being repeated as a function of the subsequent state transition (common or rare) and trial outcome (rewarded or unrewarded). Dots show individual subjects, error bars cross subject SEM. **E)** Mixed effects logistic regression predicting repeating choice as a function of the trial outcome (rewarded or not), state transition (common or rare), the transition-outcome interaction, and whether the choice was to the correct (high reward probability) option. **F)** Lagged logistic predicting choice as a function of the history of different trial types defined by the transition and outcome. **G)** Model comparison showing difference in BIC score for different models relative to the best model. Only the best model from each class of models considered is shown, the full set of models compared is shown in Figure S1A. *p<0.05, **p<0.01, ***p<0.001

This interaction between outcome and transition is consistent with model-based RL – where subjects evaluate first-step actions by combining the transition probabilities linking them to second-step states with the values of those states[24]. However, it is also consistent with a strategy of inferring the latent state of the reward probabilities and learning separate model-free values for observable actions and trial events conditioned on this[11]. Model-free values learnt under the latent state where reward probabilities are high at the up port will be high for the "up-active" second-step state, and for the first-step action that commonly leads to this state, and low for the "down-active" state and the alternative first-step choice. These values will be reversed for the other latent state. Under this account, rewards change subjects' *beliefs* about the current latent state, changing which set of values is in play, and hence the preferred first-step action. Subjects can then adapt to reversals by changing their beliefs about the current state, without needing to relearn values.

Model comparison supported an inference-based account over standard model-based RL (Figure 1G, S1A). The best fitting model used a single latent variable updated in opposite directions by reward at the up and down port, which determined the relative value of the two first-step actions - i.e., the model tracked whether it had received more reward recently at the up or down port, and chose the first step action which commonly led to the more rewarded port. It additionally incorporated a simpler strategy component, termed 'reward-as-cue' in[11], which uses the most recent outcome (e.g., reward at the up-port) as a discriminative stimulus for which first-step action to choose on the next trial. Both strategy components therefore use the trial outcome as information about the task's state, although they differ with respect to whether they maintain this information beyond the subsequent choice.

Unexpectedly, model-comparison indicated that subjects only updated their preferences following reward but not reward omission. This is consistent with both the stay probability and lagged regression analysis (Figure 1D,F), which both show that the state transition, and hence *where* the outcome was obtained, strongly affected future choices on rewarded but not non-rewarded trials. This inference-based model fit the data better than standard RL models which learn values independently for the two second-step states, irrespective of whether these were model-free, model-based or a mixture of both, including models with asymmetric learning rates for positive and negative outcomes. Behaviour simulated from the best model closely recapitulated subjects' data (Figure S1B).

Value information in dopamine activity respects task causal and correlation structure

To assess how task structure influenced dopamine activity, we used fibre photometry to record calcium signals from GCaMP6f expressing dopamine neuron cell bodies in the ventral tegmental area (VTA) and axons in nucleus accumbens (NAc) and dorsomedial striatum (DMS), and dopamine signals in NAc and DMS using dLight1.1 expressed pan-neuronally in striatum.

Dopamine activity fluctuated dynamically across the trial, as the mice made their initial choice, received information about the second-step state reached, and discovered the trial outcome (Figure 2). Reward responses were prominent in all signals, though relatively weaker in DMS calcium – in part as average DMS calcium activity masked a strong medio-lateral gradient in reward responsivity, with larger reward responses more laterally in DMS (Figure S7).

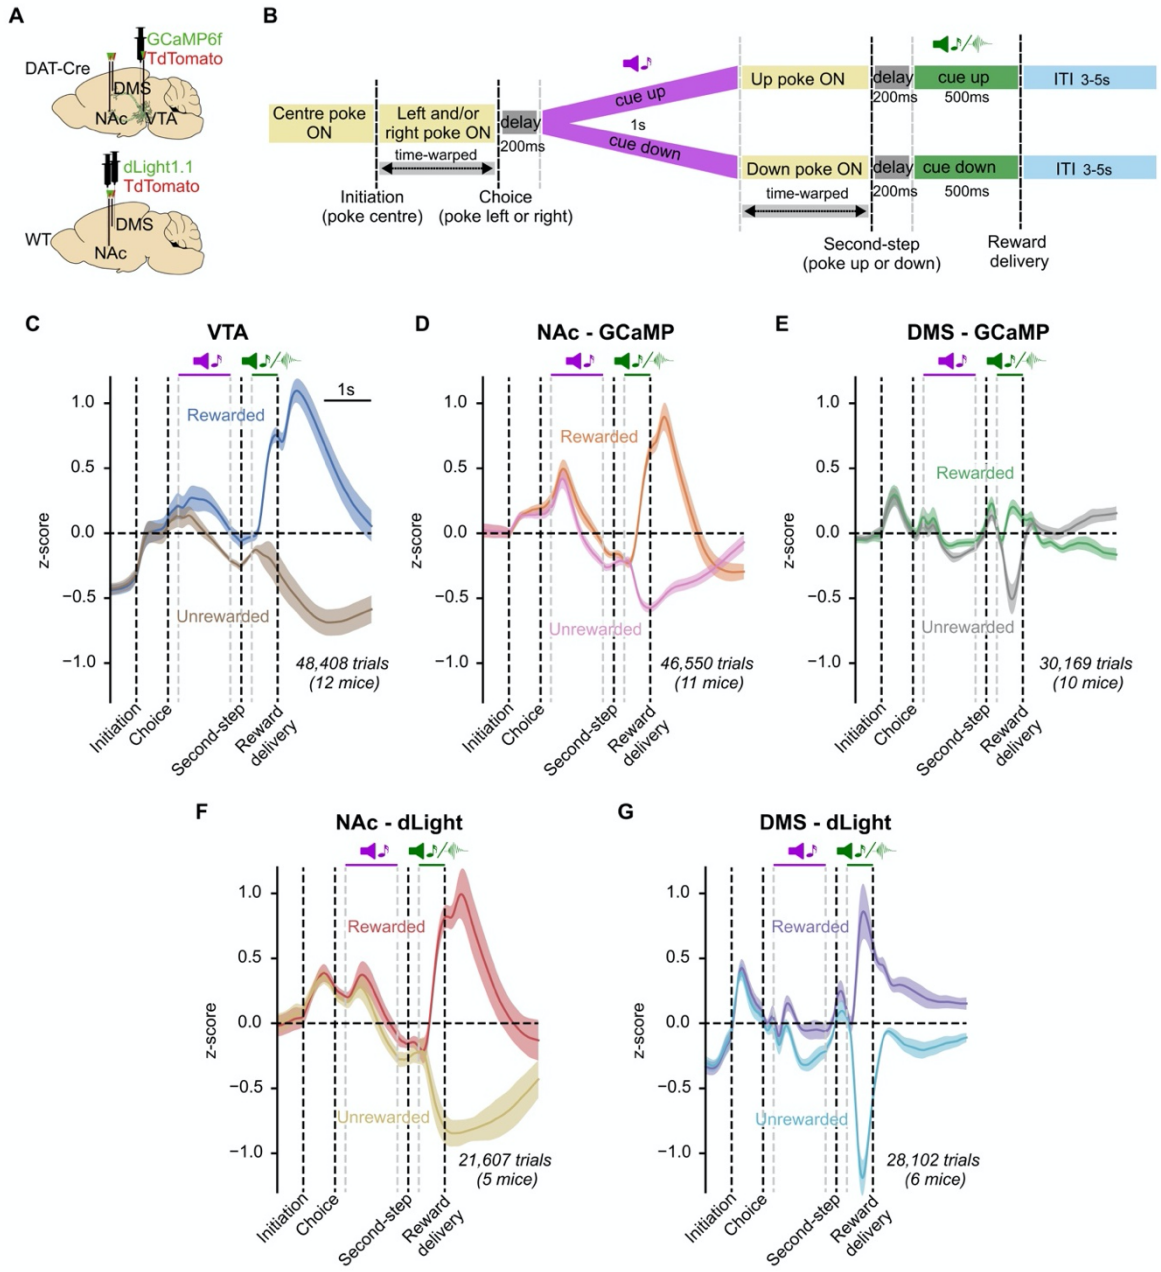

Figure 2. Dopamine activity and release are modulated across the trial. A) Injection and implant schematic. For calcium photometry, DAT-Cre mice were injected bilaterally with GCaMP6f and TdTomato in VTA, and optic fibres were implanted in VTA and NAc in one hemisphere, and DMS in the other hemisphere. For dLight photometry, WT mice were injected with dLight1.1 and TdTomato in NAc and DMS in different hemispheres, and optic fibres implanted at the same sites. **B)** Trial timeline. Trials were aligned by warping the intervals between the trial initiation and choice, and between the second-step port lighting up and being poked. **C-E)** Mean z-scored calcium signal on rewarded and unrewarded trials with shaded areas indicating cross-subject standard error in **C)** VTA (n: 12 mice, 125 sessions, 48,408 trials), **D)** NAc (n: 11 mice, 116 sessions, 46,550 trials), **E)** DMS (n: 10 mice, 100 sessions, 30,169 trials), **F-G)** As C-E but for NAc dopamine concentrations in **E)** NAc (n: 5 mice, 69 sessions, 21,607 trials), and **G)** DMS (n: 6 mice, 108 sessions, 28,102 trials.

To resolve other influences on dopamine activity, we used linear regression to model dopamine fluctuations across trials and time-points. We aligned activity across trials (Figure 2B) and ran a separate regression analysis for each timepoint in the aligned data, using predictors which varied from

trial-to-trial but took the same value for all timepoints in each trial. The time courses of predictor loadings across the trial therefore reflect when each predictor explained variance in the activity. We focus on the GCaMP data in the main figures, but results from dLight were closely comparable, except where noted in the text (Figure S4).

Our first analysis was designed to test specific hypotheses about how the previous trial's events influenced value estimates for actions and states experienced on the current trial. These hypotheses were expressed as a set of regressors detailed below, but the analysis also included regressors modelling trial outcome and other possible sources of variance (see methods, Figure S3,4). As expected, dopamine activity in all regions was strongly positively modulated by trial outcome (rewarded or not) from when this information became available (Figure 3A). However, while all regions responded to the (fully predictive) outcome cue, responses to subsequent reward delivery were strong in VTA, weaker in NAc and absent in DMS.

We next looked for evidence of subjects inferring a single latent variable that controlled both reward probabilities, using two regressors that coded the previous trial's outcome (rewarded or not) either on trials where the second-step state was (i) the same as that reached on the previous trial, or (ii) different from the previous trial. When the second-step state was the same, reward on the previous trial *increased* dopamine at the time the second-step was revealed on the current trial (Figure 3B). This is predicted by both independent value learning at the two second-step states and by inference about a single latent variable that controls both reward probabilities. However, when the second-step state was *different* from the previous trial, reward on the previous trial had a *negative* influence on dopamine at the time the second-step state was revealed on the current trial. This is not predicted by independent value learning – because reward at one second-step state does not affect the value of the other second-step state – but is predicted by inference of a single latent variable: evidence one state is good is also evidence the other is bad. In NAc, but unexpectedly in neither other region, loading on these regressors reversed sign at outcome time. This is exactly as predicted for an RPE signal: at a state transition, the value of the new state contributes positively to the RPE, while the value of the old state contributes negatively[29–31].

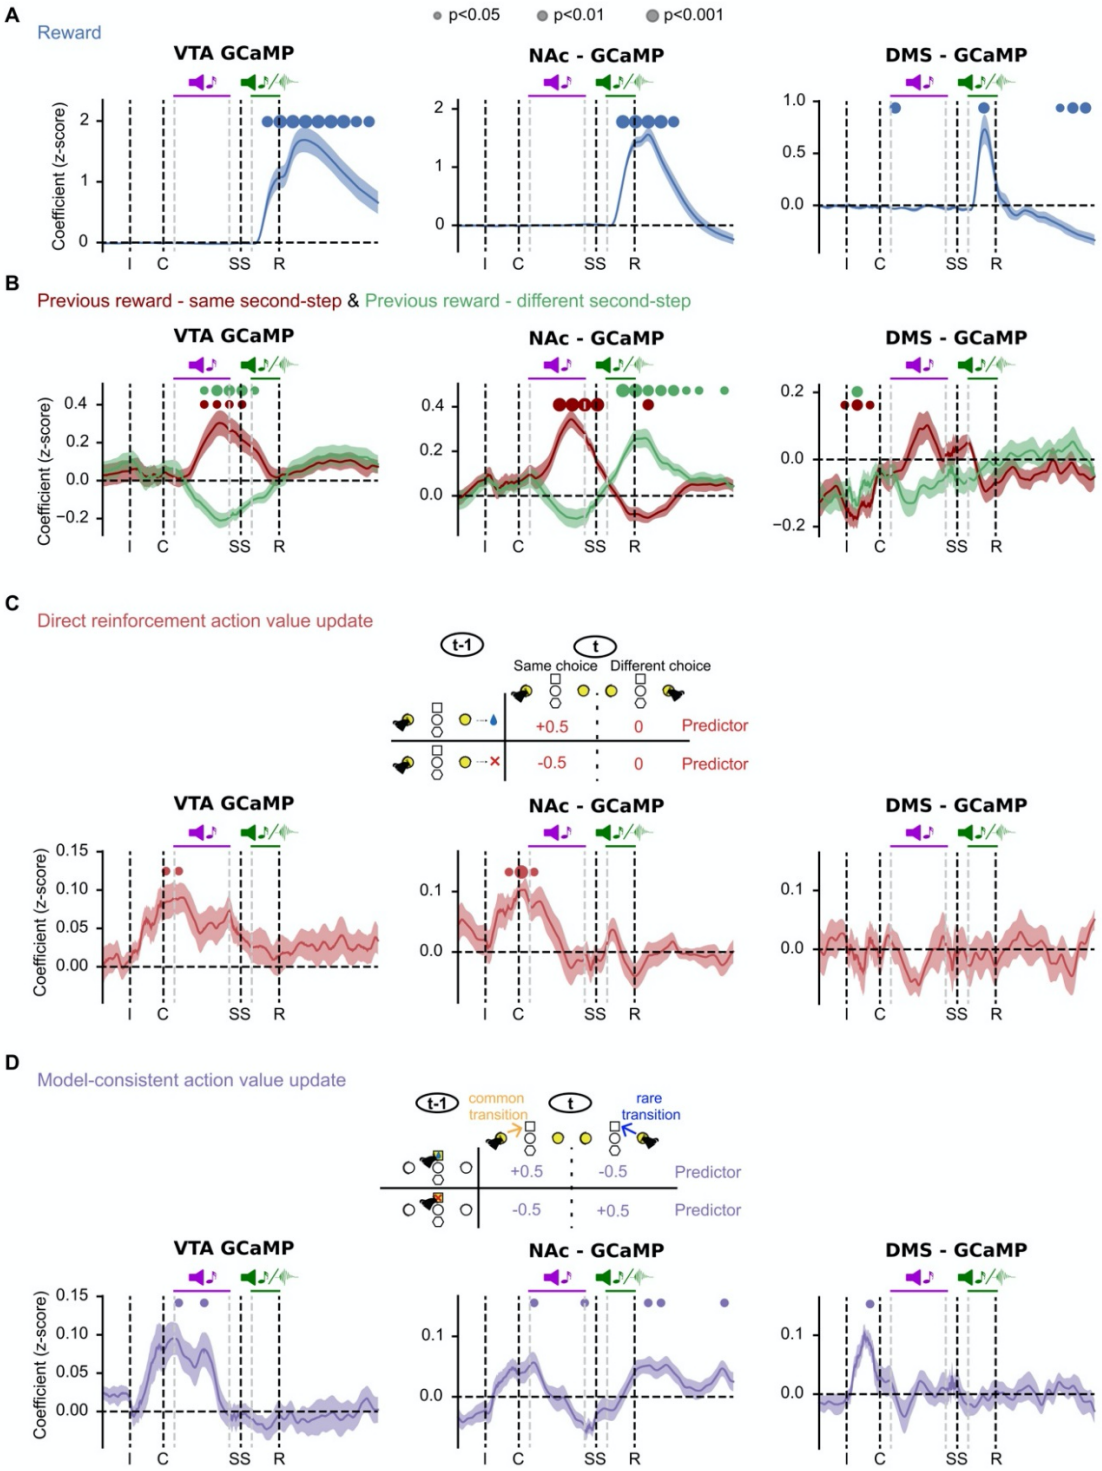

Figure 3. Dopaminergic value information respects task structure **A-D)** Coefficients in linear regression predicting dopamine activity for each timepoint in the trial for each region. **A)** Current trial outcome **B)** Previous trial outcome when second-step state is the same (red) or different (green) from the current trial. **C)** Direction of chosen value update by model-free direction reinforcement of previous choice by previous outcome. **D**) model-based action value update regressor, computed as positive if the choice taken on the current trial commonly leads to the previous second-step that was rewarded, or rarely leads to the previous step that was not rewarded, and negative for the opposite. Mean and cross-subject standard error (shaded area). Dots represents statistical significance after Benjamini-Hochberg multiple comparison correction, p<0.05, p<0.001, p<0.001.

We next assessed what value information informed the dopamine response to subject's first-step choice, using two additional regressors. A *direct-reinforcement* regressor coded the previous trial outcome (rewarded or not) only on trials where the current choice was the same as the previous choice (Figure 3C) – i.e., the direction in which model-free direct reinforcement of the previous choice by trial outcome would update the value of the first-step action chosen on the current trial. This regressor loaded positively around choice time in VTA and NAc calcium but not DMS, consistent with action value estimates used by mesolimbic dopamine incorporating a simple model-free update over observable events.

A second *model-consistent* predictor coded the direction in which the previous trial's outcome would affect the value of the action chosen on the current trial, given a correct model of both the anticorrelated reward probabilities and transition structure linking first-step actions to second-step states (Figure 3D). This loaded positively around choice time in VTA, NAc and DMS calcium activity, consistent with the dopamine system having access to value information about the chosen action that respected the task structure. Both the second-step value updates and model-consistent first-step action value update were primarily driven by rewarded trials and not reward omissions (Figure S5), consistent with the asymmetric effect of reward and omission on behaviour.

These data demonstrate that dopamine signals were shaped by value information that respected two aspects of task structure: the anticorrelated reward probabilities and the action-state transitions. Mesolimbic dopamine also carried a signature of simpler model-free value updating of actions by subsequent outcomes.

In addition to these value-related signals, dopamine was also modulated by lateralised movement (Figure S3, 4). Consistent with previous reports[6], activity in DMS but not VTA or NAc increased at the time of contralateral choices. However, each trial also required a second lateralised movement from choice port back to the centrally located second-step port, which did not evoke a response in DMS, suggesting the lateralised activity may reflect cognitive aspects of choice, or action sequence initiation[32], rather than the movement itself. Unexpectedly, VTA and NAc showed a lateralized response during contralateral movements from choice to second-step ports, despite not responding to the initial choice.

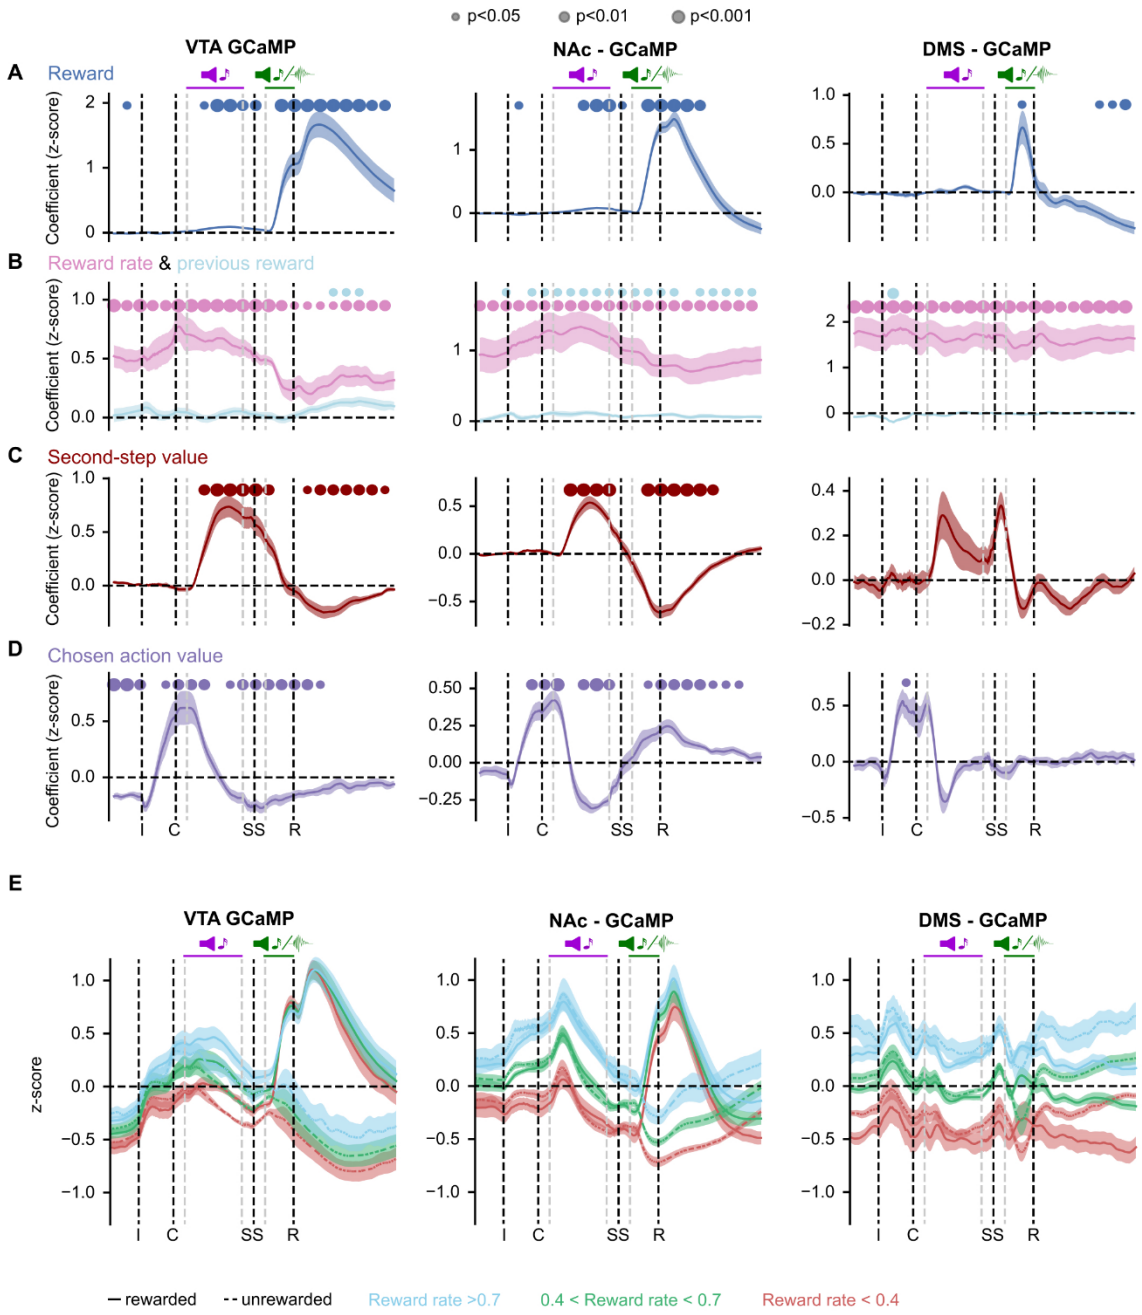

Figure 4. Dissociable influence of value and reward rate on dopamine. **A-D)** Coefficients from linear regression predicting dopamine activity using decision variables from RL model, showing cross-subject mean and standard error (shaded area), dots indicate statistical significance using Benjamini-Hochberg multiple comparison correction. For predictors: **A)** Current trial outcome. **B)** Recent reward rate (exponential moving average with tau = 10 trials) (pink) and previous trial outcome (blue). **C)** Behavioural model value of second-step reached on current trial. **D)** Behavioural model value of chosen first-step action. **E)** Mean-z-score activity split by rewarded/unrewarded trials and recent reward rate. Blue – high reward rate > 0.7 rewards/trial (exponential moving average with tau = 10 trials). Green, medium reward rate (between 0.4 and 0.7). Red, low reward rate (less than 0.4).

Dissociable influence of value and reward rate on dopamine activity

The above analysis used predictors expressing specific hypotheses about how the previous trial's events affected value information in the dopamine system. A complementary picture can be obtained

using trial-by-trial estimates of decision variables from the best fitting behavioural model to predict dopamine fluctuations. We used the same linear regression framework but with behavioural model values for the first-step choice and second-step state as predictors. As detailed above, the best fitting model used a single latent variable to track which option is best, updated in opposite directions by rewards at the two second-step states. The second-step value is just this variable interacted with the second-step state reached on the trial, while the chosen action value is this variable interacted with the chosen action. Both values therefore respect the anticorrelated reward probabilities, and the chosen action value combines this with the action-state transition structure.

Second-step state value had a strong positive influence on VTA and NAc calcium at the time the second-step was revealed, which subsequently reversed at outcome time, consistent with RPE signalling (Figure 4C). The time-course was more complex in DMS, with peaks following both the second-step cue and second-step port entry, though these only reached significance when multiple comparison corrected for multiple timepoints in the dLight data (Figure S6). Chosen action value had a strong positive influence on activity in all 3 regions around the time of the choice, which then reversed when the second-step was revealed, again consistent with RPE. These data confirm that value information that respecting the task causal and correlation structure strongly influence dopamine responses to both the first-step action and second-step choice.

As both reward rate and value have been shown to correlate with dopamine activity[5,33], we additionally included recent reward rate as a predictor, calculated as an exponential moving average with time constant of 10 trials. Reward rate had a strong positive influence on dopamine neuron calcium in all 3 regions (Figure 4B,E). Unlike the influence of value and reward, which were tightly time-locked to when information became available, reward rate influenced activity at all timepoints, with only weak modulation across the trial. Reward rate also strongly influenced dopamine concentrations in NAc (Figure S6), again accounting for separate variance from the prominent influence of action and state values. However, reward rate had limited influence on DMS dopamine concentrations (Figure S6), despite strongly influencing DMS dopamine axonal calcium.

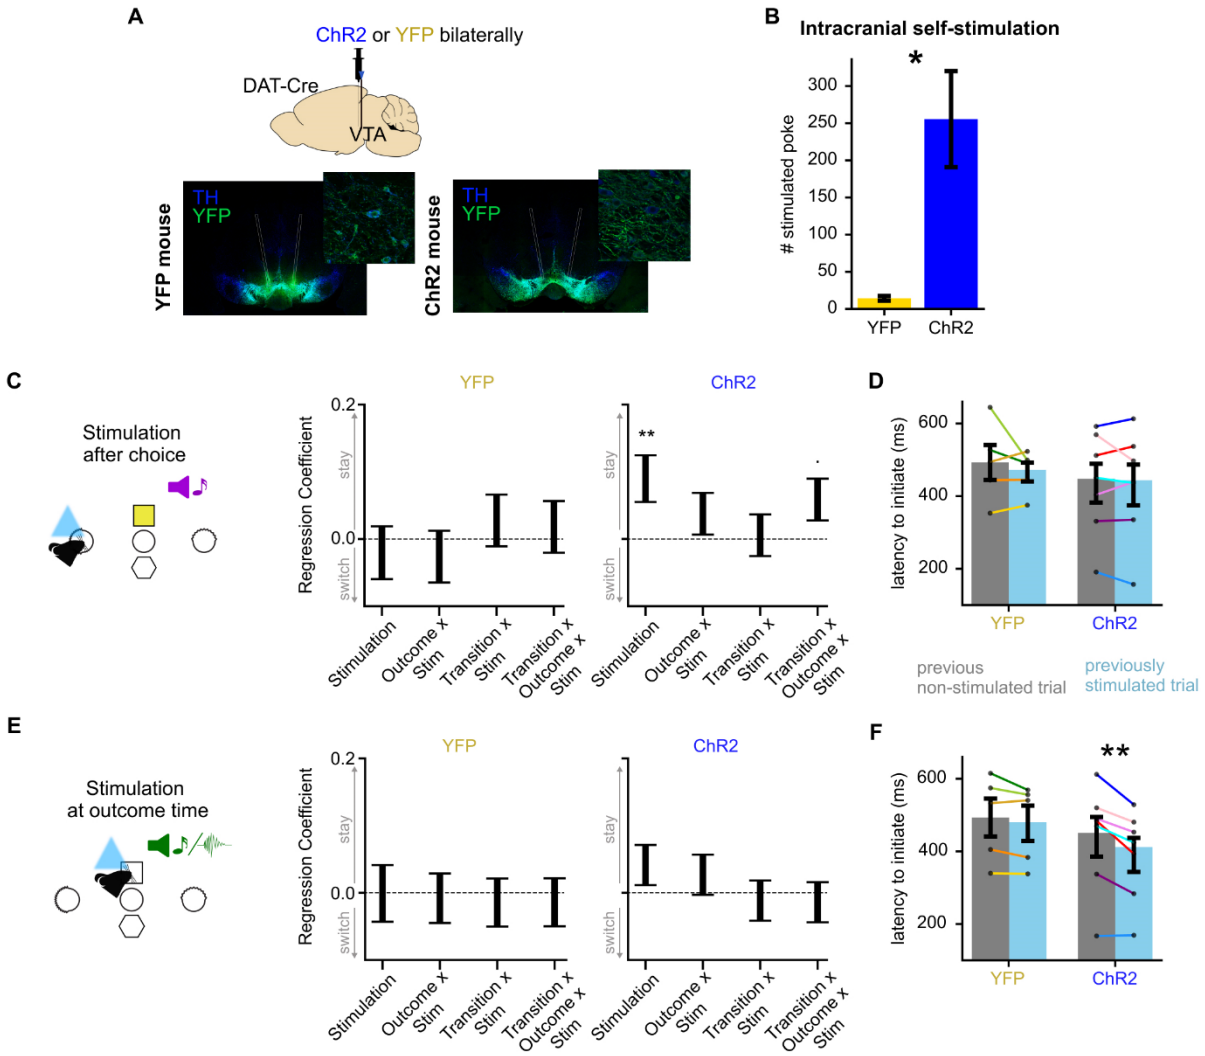

Figure 5. Dopamine stimulation does not recapitulate natural reward effects A) Top: Injection and fibre placement schematics. Bottom: Histology from a YFP and a ChR2 mouse, stained for TH (Tyrosine Hydroxylase) and YFP. **B)** Intracranial self-stimulation. Mean number of entries to a nose-poke port delivering stimulation in the YFP vs the ChR2 group. Error bars show cross subject standard error. *p<0.05 t-test. **C-D)** Effect of optogenetic stimulation 200ms after choice in the the control (YFP) and ChR2 group. **E-F)** Effect of optogenetic stimulation at time of the trial outcome cue. **C,E)** Mixed effects logistic regression predicting repeating the same choice on the next trial using the regression model from Figure 2C, with additional regressors modelling the effect of stimulation and its interaction with trial events. Only regressors showing stimulation effects are shown here, see Figure S8 for the full model. •p<0.1, **p<0.005 **D-F)** Mean latency to initiate trial following stimulated and non-stimulated trials. Dots indicate individual animals, error bars show standard error. **p<0.01 two-sided paired t-test. N: YFP: 5 animals, 13,079 trials (choice-time stimulation sessions), 15,419 trials (outcome-time stimulation sessions); ChR2: 7 animals, 20,817 trials (choice-time stimulation sessions), 23,051 trials (outcome-time stimulation sessions).

Dopamine stimulation does not recapitulate natural rewards despite reinforcing preceding actions

The above data demonstrate that dopamine report RPEs and reward information, but cannot establish whether this signal was causally responsible for updating mice's choice policy. To assess this, we

therefore manipulated dopamine activity in a new cohort of mice expressing either channelrhodopsin (ChR2, N=7) or a control fluorophore (EYFP, N=5) (Figure 5A) in VTA dopamine neurons.

We first verified that our stimulation parameters (5 pulses, 25 Hz, ~8-10 mW power) were sufficient to promote and maintain intracranial self-stimulation in the ChR2 but not YFP groups ($p = 0.011$, t-test) using an assay where optical stimulation was delivered contingent on nose-poking in a different context from the two-step task (Figure 5B). We then examined the effect on two-step task behaviour of optogenetic stimulation at two different timepoints – (i) after the first-step choice, at the time of second-step cue presentation, or (ii) at outcome cue presentation (Figure 5C) (25% stimulated trials, stimulation timepoint fixed for each session, order counterbalanced across sessions and mice, N=6-8 sessions per mouse and stimulation type).

Stimulating dopamine neurons after the first-step choice was reinforcing; it significantly increased the probability of repeating that choice on the next trial ($p=0.0095$, mixed effect logistic regression, Figure 5C). This stimulation did not significantly interact with either the transition or outcome in its effect on next trial choice (all $p>0.059$). Strikingly, stimulating dopamine neurons at the time of trial outcome had no significant influence the subsequent choice (Figure 5E), neither reinforcing the preceding first-step choice ($p=0.17$), nor acting like a natural reward by interacting with the state transition ($p=0.71$). No effect of stimulation was found in the YFP group for either of the stimulation types (all $p>0.47$). By contrast, while stimulation after the first step choice had no effect on the latency to initiate the next trial ($p=0.74$), stimulation at outcome time significantly reduced the latency to initiate the next trial (Figure 5F, $p=0.0055$, two-sided t-test). Again, we did not find an effect on the latency to initiate on the YFP group ($p=0.46$, stimulation after choice; $p=0.16$, stimulation at outcome time).

Therefore, while stimulation of VTA dopamine neurons was sufficient to reinforce immediately preceding actions, it completely failed to recapitulate the behavioural consequences of natural rewards at outcome, despite natural rewards driving the largest dopamine signals observed in the task.

Discussion

We recorded and manipulated dopamine activity in a sequential decision task where rewards do not simply reinforce preceding actions, but rather affect subsequent choices in a manner that respects the task structure. This revealed a striking asymmetry between reporting of values and updating of policy by dopamine. Moment-by-moment dopamine fluctuations were strongly influenced by both the value of actions taken and states encountered, and reward delivery or omission. However, while natural rewards strongly influenced subsequent choices, stimulating dopamine neurons at the time of

trial outcome had no effect on animal's next choice. This was despite positive controls verifying that, in the same subjects, the same stimulation was sufficient to (a) reinforce choice when stimulating after a first-step choice, (b) induce robust self-stimulation in a different context, and (c) speed subsequent trial re-initiation when stimulating at outcome.

These data are surprising given the tight interaction posited by RL theory between RPE and policy; whereby rewards evoke RPEs – putatively carried by dopamine neurons, which update the values of states and actions, and hence shape future policy.

Dopamine in inference guided choice

Our findings can be reconciled with RL theory by recognizing that the task is only partially observable: there are aspects of its state-space that cannot be directly observed, but which nonetheless influence observable events. Rewards on individual trials were determined by underlying reward probabilities, which underwent occasional reversals. The current configuration of the reward probabilities constitutes a latent or unobservable feature of the task's state, which can be inferred by observing where rewards are obtained[11,12]. Task rewards therefore carry *information* about the current *state* of the task, in addition to acting as a reward in the RL sense. A subject who can infer the latent state of the reward probabilities can learn separate values for actions and trial events conditioned on it. State information provided by task rewards can then influence future choices by changing subjects' beliefs about the latent state, and hence which values are in play[11]. In this account, subjects adapt to reversals by changing their beliefs about the current state of the task, not by changing the values associated with each state.

Several lines of evidence support this. First, the best fitting behavioural model used a single variable updated in opposite directions by rewards at the up and down port, rather than learning values for the two second-step states independently. This can be thought of as representing subject's belief about whether the reward probabilities are currently high at the up or down port. It is consistent with state inference because it incorporates knowledge that rewards at the two locations are not independent but rather controlled by a single latent variable. Unexpectedly, only rewards but not reward omissions appear to update this variable, a conclusion supported both by model comparison (Figure 1G) and the lagged regression analysis (Figure 1F). This is not optimal as both rewards and reward omissions were informative, and might reflect either cognitive limitations or divergence between task statistics and those of ethological foraging.

Second, rewards obtained on one trial decreased the dopamine response evoked by reaching the other second-step state on the next trial, consistent with the inference that reward obtained at one second-step state makes obtaining reward at the other less likely.

Third, our manipulation results are consistent under the assumption that dopamine neurons carry RPEs evoked by rewards, used to update the values of states and actions, but not the state information used to update beliefs. Stimulating dopamine immediately after choice will reinforce that choice by increasing the value of the chosen action. Stimulating dopamine at the time of trial outcome will increase the value of that second-step, but this change in second-step value would only influence the next trial choice if subjects were doing prospective model-based evaluation at the first-step. However, if subjects learn values conditioned on the inferred latent state, there is no advantage to doing costly model-based evaluation, because model-free values will generate very similar behaviour[11]. Specifically, model-free values learnt conditioned on the latent state will be high for the first-step action that commonly leads to the second-step port with high reward probability.

These data build on previous evidence for inference underlying behavioural flexibility[12,13,15], and informing dopamine responses[18,20,34,35], but extend them by both recording and manipulating dopamine on fast timescales during inference-guided adaptive decisions, revealing a disconnect between value reporting and policy updating. An inference-based perspective helps explain previous observations that are surprising under a pure value learning account of flexible behaviour. Two studies using probabilistic reward-guided tasks, found that optogenetic inhibition of dopamine neurons on single trials biased animals against the chosen option on the subsequent trial[5,6], interpreted as evidence that task behaviour was dopamine dependent and dopamine inhibition acted as a negative RPE. However, in both cases, the reinforcing effect of natural rewards was unaffected by inhibiting dopamine neurons, inconsistent with dopaminergic RPEs mediating reward effects, but consistent with the inference-based account above. Similarly, pharmacological manipulation of dopamine signalling in reward guided decision tasks often has limited or no effect on learning[12,36–38], suggesting dopamine independent influence of rewards on policy.

This does not imply dopamine is not involved in such tasks. Pharmacological dopamine manipulations show effects on choice consistency and task engagement even where they show limited or no effect on trial-by-trial learning[37–39]. Dopaminergic RPEs likely also drive value learning during initial task acquisition. As representations of observable events are available immediately, but those of latent states will emerge only gradually with task experience, dynamic arbitration must occur between values learnt over observable and inferred features. Normatively, this should be determined by their predictive accuracy, perhaps implemented via synaptic learning rules which track uncertainty[40]. As values conditioned on correctly inferred latent states are stable across changes in those states, they

can presumably be learnt with higher accuracy than values over observable features, which change when latent states change. This suggests that over task acquisition, control will transfer from values learnt over observable features to those over latent features, and hence the mechanism of behavioural flexibility from value-learning to inference.

Recording and manipulation studies implicate medial and orbitofrontal cortex in learning and representing latent structure[13,41–46]. This may reflect their position near the top of hierarchically organised predictive models in cortex, which learn to minimise sensory prediction errors by inferring their latent causes. Another possible learning mechanism is meta RL, where plasticity in a recurrent network on a slow timescale, driven by reward prediction errors, implements network dynamics that adapt to environment changes on a fast timescale[47].

RPE and value and reward rate

Recent work has questioned the relative influence of value and RPE on dopamine activity and striatal concentrations[4,5,33]. A key signature of RPE is that the value of encountered states contributes at two timepoints with opposite signs – first positively when the state is entered, then negatively on transition to the next state[29–31]. We observed this biphasic influences of the values of both first-step actions and second-step states in both calcium activity in VTA, NAc and DMS (Figure 4C,D), and in dopamine concentrations in DMS and NAc (Figure S6). This biphasic pattern was most prominent in NAc, with VTA and DMS signals both showing substantially larger initial positive modulation than later negative modulation. Intriguingly, when evaluating only the influence of the most recent outcome on second-step state value, rather than the extended history, the biphasic pattern was prominent in NAc but absent in VTA (Figure 3B). Differences between the VTA and striatal signals could reflect local modulation of terminals, for example by cholinergic neurons[48], or alternatively the VTA signal may be influenced by calcium activity in dendrites that is partially decoupled from spiking due to somatic inhibition.

In addition, we observed the average reward rate over recent trials had a strong positive influence on dopamine activity in all regions (Figure 4B, S6). Importantly, this accounted for an entirely separate portion of variance from the value of trial events. Value-related signals were tightly time-locked to specific events, whereas reward rate had sustained effects across the trial from before initiation to after choice. This appears broadly consistent with theoretical proposals that tonic dopamine represents average reward rate, which acts as the opportunity cost of time in average-reward RL[49]. Correlation between reward rate and NAc dopamine concentrations is consistent with recent

reports[5,33], but that with VTA calcium is unexpected given previous reports of no correlation with VTA spikes[33].

Conclusions

Our findings emphasise that flexible behaviour can involve two processes operating in parallel; inference about the current state of the world, and evaluation of those states and actions taken in them. The involvement of dopamine in updating values has rightly been a major focus of accounts of flexible decision making. However, in the structured environments common both in the lab and real world, it is only half the picture, as subjects can adapt to change by updating their beliefs about the world's state. We showed here that although such belief updates strongly influence dopamine activity, this relationship is not reciprocal, as stimulating dopamine failed to reproduce natural reward effects. A pressing direction for future work is to understand the interplay between these two processes in flexibly adapting to a changing world.

Acknowledgements

This research was funded by the Wellcome (202831/Z/16/Z: MEW; 214314/Z/18/Z: MEW, TA; WT096193AIA: TA; 215198/Z/19/Z: MBP). For the purpose of open access, the author has applied a CC-BY public copyright licence to any Author Accepted Manuscript version arising from this submission. We thank Marios Panayi for assistance with the mixed effects modelling and Tim Behrens for discussions about the work.

Author Contributions

TA and MEW conceived the project and, with MBP, designed the experiments. MBP and TA collected the data. MBP and TA analyzed the data. All authors wrote the manuscript.

Conflict of Interest

The authors have no competing interests to declare.

References

1. Montague, P. R., Dayan, P. & Sejnowski, T. J. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* **16**, 1936–47 (1996).
2. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–9 (1997).
3. Eshel, N. *et al.* Arithmetic and local circuitry underlying dopamine prediction errors. *Nature* **525**, 243–6 (2015).
4. Kim, H. R. *et al.* A Unified Framework for Dopamine Signals across Timescales. *Cell* **183**, 1600-1616.e25 (2020).
5. Hamid, A. A. *et al.* Mesolimbic dopamine signals the value of work. *Nat. Neurosci.* **19**, 117–26 (2016).
6. Parker, N. F. *et al.* Reward and choice encoding in terminals of midbrain dopamine neurons depends on striatal target. *Nat. Neurosci.* **19**, 845–54 (2016).
7. Kravitz, A. V, Tye, L. D. & Kreitzer, A. C. Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nat. Neurosci.* **15**, 816–8 (2012).
8. Kim, K. M. *et al.* Optogenetic mimicry of the transient activation of dopamine neurons by natural reward is sufficient for operant reinforcement. *PLoS One* **7**, e33612 (2012).
9. Steinberg, E. E. *et al.* A causal link between prediction errors, dopamine neurons and learning. *Nat. Neurosci.* **16**, 966–73 (2013).
10. Bartolo, R. & Averbeck, B. B. Inference as a fundamental process in behavior. *Curr. Opin. Behav. Sci.* **38**, 8–13 (2021).
11. Akam, T., Costa, R. & Dayan, P. Simple Plans or Sophisticated Habits? State, Transition and Learning Interactions in the Two-Step Task. *PLoS Comput. Biol.* **11**, e1004648 (2015).
12. Costa, V. D., Tran, V. L., Turchi, J. & Averbeck, B. B. Reversal learning and dopamine: a bayesian perspective. *J. Neurosci.* **35**, 2407–16 (2015).
13. Vertechi, P. *et al.* Inference-Based Decisions in a Hidden State Foraging Task: Differential Contributions of Prefrontal Cortical Areas. *Neuron* **106**, 166-176.e6 (2020).
14. Izquierdo, A., Suda, R. K. & Murray, E. A. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *J. Neurosci.* **24**, 7540–8 (2004).

15. Samborska, V., Butler, J. L., Walton, M. E., Behrens, T. E. & Akam, T. Complementary Task Representations in Hippocampus and Prefrontal Cortex for Generalising the Structure of Problems. *bioRxiv* 2021.03.05.433967 (2021). doi:10.1101/2021.03.05.433967

16. Mackintosh, N. J., McGonigle, B., Holgate, V. & Vanderver, V. Factors underlying improvement in serial reversal learning. *Can. J. Psychol.* **22**, 85–95 (1968).

17. Bond, A. B., Kamil, A. C. & Balda, R. P. Serial reversal learning and the evolution of behavioral flexibility in three species of North American corvids (Gymnorhinus cyanocephalus, Nucifraga columbiana, Aphelocoma californica). *J. Comp. Psychol.* **121**, 372–9 (2007).

18. Bromberg-Martin, E. S., Matsumoto, M., Hong, S. & Hikosaka, O. A pallidus-habenula-dopamine pathway signals inferred stimulus values. *J. Neurophysiol.* **104**, 1068–76 (2010).

19. Nakahara, H., Itoh, H., Kawagoe, R., Takikawa, Y. & Hikosaka, O. Dopamine neurons can represent context-dependent prediction error. *Neuron* **41**, 269–80 (2004).

20. Starkweather, C. K., Babayan, B. M., Uchida, N. & Gershman, S. J. Dopamine reward prediction errors reflect hidden-state inference across time. *Nat. Neurosci.* **20**, 581–589 (2017).

21. Engelhard, B. *et al.* Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. *Nature* **570**, 509–513 (2019).

22. Lak, A. *et al.* Neural basis of learning guided by sensory confidence and reward value. *bioRxiv* 411413 (2019). doi:10.1101/411413

23. Tsutsui-Kimura, I. *et al.* Distinct temporal difference error signals in dopamine axons in three regions of the striatum in a decision-making task. *Elife* **9**, (2020).

24. 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**, 1204–15 (2011).

25. 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**, 1595–600 (2015).

26. Colas, J. T., Pauli, W. M., Larsen, T., Tyszka, J. M. & O'Doherty, J. P. Distinct prediction errors in mesostriatal circuits of the human brain mediate learning about the values of both states and actions: evidence from high-resolution fMRI. *PLoS Comput. Biol.* **13**, e1005810 (2017).

27. Brown, H. D., McCutcheon, J. E., Cone, J. J., Ragozzino, M. E. & Roitman, M. F. Primary food reward and reward-predictive stimuli evoke different patterns of phasic dopamine signaling

437 throughout the striatum. *Eur. J. Neurosci.* **34**, 1997–2006 (2011).

438 28. Howe, M. W. & Dombeck, D. A. Rapid signalling in distinct dopaminergic axons during
439 locomotion and reward. *Nature* **535**, 505–10 (2016).

440 29. Behrens, T. E. J., Hunt, L. T., Woolrich, M. W. & Rushworth, M. F. S. Associative learning of
441 social value. *Nature* **456**, 245–9 (2008).

442 30. Rutledge, R. B., Dean, M., Caplin, A. & Glimcher, P. W. Testing the reward prediction error
443 hypothesis with an axiomatic model. *J. Neurosci.* **30**, 13525–36 (2010).

444 31. Niv, Y., Edlund, J. A., Dayan, P. & O'Doherty, J. P. Neural prediction errors reveal a risk-
445 sensitive reinforcement-learning process in the human brain. *J. Neurosci.* **32**, 551–62 (2012).

446 32. Jin, X. & Costa, R. M. Start/stop signals emerge in nigrostriatal circuits during sequence
447 learning. *Nature* **466**, 457–62 (2010).

448 33. Mohebi, A. *et al.* Dissociable dopamine dynamics for learning and motivation. *Nature* **570**,
449 65–70 (2019).

450 34. Babayan, B. M., Uchida, N. & Gershman, S. J. Belief state representation in the dopamine
451 system. *Nat. Commun.* **9**, 1891 (2018).

452 35. Gershman, S. J. & Uchida, N. Believing in dopamine. *Nat. Rev. Neurosci.* **20**, 703–714 (2019).

453 36. Grogan, J. P. *et al.* Effects of dopamine on reinforcement learning and consolidation in
454 Parkinson's disease. *Elife* **6**, (2017).

455 37. Verharen, J. P. H., Adan, R. A. H. & Vanderschuren, L. J. M. J. Differential contributions of
456 striatal dopamine D1 and D2 receptors to component processes of value-based decision
457 making. *Neuropsychopharmacology* **44**, 2195–2204 (2019).

458 38. Korn, C. *et al.* Distinct roles for dopamine clearance mechanisms in regulating behavioral
459 flexibility. *bioRxiv* 823401 (2021). doi:10.1101/823401

460 39. Lee, E., Seo, M., Dal Monte, O. & Averbeck, B. B. Injection of a dopamine type 2 receptor
461 antagonist into the dorsal striatum disrupts choices driven by previous outcomes, but not
462 perceptual inference. *J. Neurosci.* **35**, 6298–306 (2015).

463 40. Aitchison, L. *et al.* Synaptic plasticity as Bayesian inference. *Nat. Neurosci.* **24**, 565–571
464 (2021).

465 41. Schuck, N. W., Cai, M. B., Wilson, R. C. & Niv, Y. Human Orbitofrontal Cortex Represents a
466 Cognitive Map of State Space. *Neuron* **91**, 1402–1412 (2016).

42. Klein-Flügge, M. C., Wittmann, M. K., Shpektor, A., Jensen, D. E. A. & Rushworth, M. F. S. Multiple associative structures created by reinforcement and incidental statistical learning mechanisms. *Nat. Commun.* **10**, 4835 (2019).

43. Bradfield, L. A., Dezfouli, A., van Holstein, M., Chieng, B. & Balleine, B. W. Medial Orbitofrontal Cortex Mediates Outcome Retrieval in Partially Observable Task Situations. *Neuron* **88**, 1268–1280 (2015).

44. Starkweather, C. K., Gershman, S. J. & Uchida, N. The Medial Prefrontal Cortex Shapes Dopamine Reward Prediction Errors under State Uncertainty. *Neuron* **98**, 616-629.e6 (2018).

45. Jocham, G. *et al.* Reward-Guided Learning with and without Causal Attribution. *Neuron* **90**, 177–90 (2016).

46. Walton, M. E., Behrens, T. E. J., Buckley, M. J., Rudebeck, P. H. & Rushworth, M. F. S. Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* **65**, 927–39 (2010).

47. Wang, J. X. *et al.* Prefrontal cortex as a meta-reinforcement learning system. *Nat. Neurosci.* **21**, 860–868 (2018).

48. Threlfell, S. & Cragg, S. J. Dopamine signaling in dorsal versus ventral striatum: the dynamic role of cholinergic interneurons. *Front. Syst. Neurosci.* **5**, 11 (2011).

49. Niv, Y., Daw, N. D., Joel, D. & Dayan, P. Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl).* **191**, 507–20 (2007).

50. Akam, T. *et al.* pyControl: Open source, Python based, hardware and software for controlling behavioural neuroscience experiments. *bioRxiv* 2021.02.22.432227 (2021). doi:10.1101/2021.02.22.432227

51. Lopes, G. *et al.* Bonsai: an event-based framework for processing and controlling data streams. *Front. Neuroinform.* **9**, 7 (2015).

52. Akam, T. & Walton, M. E. pyPhotometry: Open source Python based hardware and software for fiber photometry data acquisition. *Sci. Rep.* **9**, 3521 (2019).

53. Singmann, H., Bolker, B., Westfall, J. & Aust, F. afex': analysis of factorial experiments. R package. (2018).

METHODS

Subjects

All experiments were performed in line with the UK Animal (Scientific Procedure) Act 1986 and in accordance with the University of Oxford animal use guidelines. 12 DAT-Cre mice (7 females and 5 males) were used for the GCaMP photometry recordings, 6 WT C57BL6 mice (3 females and 3 males) for the dLight recordings, and 12 DAT-Cre mice (YFP: 2 females and 3 males; ChR2: 4 females and 3 males) for the optogenetic manipulation experiment. Mice were aged 8-16 weeks at the start of behavioural training.

Data and code availability

All data and code used in the manuscript will be made available in a public repository from acceptance of the manuscript.

Behavioural setup

The task was run in custom-built 12x12 cm operant boxes (design files at https://github.com/pyControl/hardware/tree/master/Behaviour_box_small) controlled using pyControl[50]. Five nose poke ports were located on the back wall of the boxes; a central initiation port flanked by two choice ports 4 cm to the left and right and two second-step ports located 1.6 cm above and below the central poke. The second-step ports each had a solenoid to deliver water rewards. A speaker located above the ports was used to deliver auditory stimuli. Video was acquired from an FLIR Flea 3 camera positioned above each setup using a Bonsai based workflow (https://github.com/ThomasAkam/Point_Grey_Bonsai_multi_camera_acquisition)[51].

Behavioural task and training

The behavioural task was adapted from the human two-step task[24]. Each trial started with the central initiation port lighting up. Subjects initiated the trial by poking the illuminated port, which caused the choice ports to illuminate. On free-choice trials (75% of trials), both the left and right port lit up, allowing subjects to choose, while on forced-choice trials only one randomly selected choice port lit up, forcing animals to select that specific port. Poking a choice port was followed, after a 200ms delay, by the second-step port lighting up and 1s presentation of one of two auditory cues (either a 5- or 12-kHz tone depending on whether the top or bottom second-step became active). Each choice port was

commonly (80% trials) associated with transitioning to one second-step state (up or down) and rarely (20% trials) to the other. The second-step port only became responsive to pokes after cue offset. Poking the second-step port triggered a 200ms delay, after which a 500ms auditory cue signalled whether the trial was rewarded or not (same 5- or 12-kHz tone pulsed with at 10Hz on rewarded trials, white noise on unrewarded trials). Reward was delivered at the offset of this cue. To ensure mice knew when they had made a nose poke, a click sound was presented whenever the subject poked a port that was active (e.g., the initiation port during the initiation state).

Reward probabilities for the two second-step ports changed in blocks. In non-neutral blocks, one second-step port had 80% reward probability and the other 20%, while in neutral blocks both second-step ports were rewarded with 50% probability. Block transitions from non-neutral blocks were triggered 5 to 15 trials after mice crossed a threshold of 75% "correct" choices (i.e., choosing the higher reward probability), computed as the exponential moving average with time constant of 8 free choice trials. Transitions from neutral blocks were triggered after a 20-30 trials. An inter-trial interval of 2-4 s duration started once the subject remained out of the second-step port for 250ms after the trial outcome.

Training

Behavioural training took 4-6 weeks. Animals were put on water restriction 48 hours before starting training, and received 1 hour ad lib water access in their home cage 24 hours before starting training. On training days (typically 6 days/week) animals usually received all their water from the task, but were topped up outside the task as necessary to maintain a body weight >85% of their pre restriction baseline. On days off from training mice received 1 hour ad lib water access in their home cage. Water reward size was decreased from 15ul to 4ul across training in order to increase the number of trials performed and blocks experienced on each session.

Training consisted of multiple stages with increasing complexity (Table 1). At the first training stage 1.1, only the second-step ports were exposed, with all other ports covered. Second-step ports were illuminated in a pseudorandom order with inter-trial interval of 2-4 seconds. Poking an illuminated port delivered reward with 100% probability, with no auditory cues. When animals obtained >50 rewards on this stage, transitioned to stage 1.2, where the auditory cues for second-step state and reward were introduced. Once animals obtaining >70 rewards in a session, they were switched to stage 2 on the next session. At stage 2, the choice ports were introduced, but all trials were forced choice, such that only one choice port lit up on each trial. Mice were switched to stage 3 when they obtained >70 rewards on a single session. At stage 3, the initiation poke was introduced, and when

561 they obtained >70 rewards on a single session, they transitioned to stage 4 on the next training
562 session. Stage 4 consisted on multiple substages where progressively more free choice trials were
563 introduced and the reward probabilities gradually changed until reaching the final task parameters.
564 Mice were transitioned to next substage after two training sessions of 45min or a single 90min session,
565 until they reached substage 4.6. Subjects were only transition to the final stage (full task) when they
566 completed at least 5 blocks in a single session.

STAGE	START POKE	CHOICE POKES	REWARD POKES	TASK STRUCTURE
1.1			X	Reward probability 100%
1.2				Introduce auditory cues Reward probability 100%
2		X	X	Transition probability 80/20% Reward probability 100%
3	X	X	X	Transition probability 80/20% Reward probability 100%
	x	x	x	Start introducing free choice trials and reward probabilities
SUBSTAGE	**Free choice trials**	**Non-neutral blocks reward probabilities** (high/low)	**Neutral block reward probability**	**Block change** (independent/ dependent on behaviour)
4.1	0%	0.9/0.7	0.8	Independent – 20-30 trials blocks
4.2	25%	0.9/0.5	0.7	Independent – 20-30 trials blocks
4.3	25%	0.9/0.3	0.6	Independent – 20-30 trials blocks
4.4	25%	0.9/0.1	0.5	Independent – 20-30 trials blocks
4.5	50%	0.9/0.1	0.5	Dependent
4.6	75%	0.9/0.1	0.5	Dependent
4.7 – FINAL	75%	0.8/0.2	0.5	Dependent

567
568 **Table 1. Description of training stages and task contingencies**

569

570 **Surgery**

571 Mice were anaesthetised with isoflurane (3% induction, 0.5-1% maintenance), and injected with
572 buprenorphine (0.08 mg/kg), meloxicam (5 mg/kg) and glucose saline (0.5ml). Marcaine (max 2
573 mg/kg) was injected into the scalp and before placing mice into the stereotaxic frame. Mice were
574 maintained at ~37 degrees with rectal probe and heating blanket (Harvard Apparatus). Surgery
575 proceeded as described below for the different experiments. After surgery, mice were given
576 additional doses of meloxicam each day for 3 days after surgery, and were monitored carefully for 7
577 days post-surgery.

578

579 *GCaMP photometry*

580 Mice were intracranially injected with 1ul of saline containing AAV1.Syn.Flex.GCaMP6f.WPRE.SV40
581 (titre: 6.22 x 10^{12} vg/ml, Penn Vector Core) and AAV1.CAG.Flex.tdTomato.WPRE.bGH

582 (AllenInstitute864) (titre: 1.535 x 10^{12} vg/ml, Penn Vector Core) at 2nl/s in VTA (antero-posterior (AP):
583 -3.3, medio-lateral (ML): ±0.4, dorso-ventral (DV): -4.3 from bregma) in one hemisphere and VTA/SNc
584 (AP: -3.1, ML: ±0.9, DV: -4.2 from bregma) in the other hemisphere. Three 200 µm diameter ceramic
585 optic fibres were implanted chronically in each animal in VTA (AP: -3.3, ML: ±0.4, DV: -4.3 from
586 bregma) and NAc (AP: +1.4, ML: ±0.8, DV: -4.1 from bregma) in the same hemisphere, and DMS (AP:
587 +0.5, ML: ±1.5-1.7, DV: -2.6 from bregma) in the contralateral hemisphere.

588

589 *dLight photometry*

590 Mice were intracranially injected with 500 nl of saline containing pAAV5-CAG-dLight1.1 (titre: 1.4 x
591 10^{12} vg/ml, addgene) and pssAAV-2/5-hSyn1-chI-tdTomato-WPRE-SV40p(A) (titre: 4.9 x 10^{11} vg/ml,
592 ETH Zurich) at 2nl/s in NAc (AP: +1.4, ML: ±0.8, DV: -4.1 from bregma) and DMS (AP: +0.5, ML: ±1.5/1.7,
593 DV: -2.6 from bregma) in opposite hemispheres. Two 200 µm diameter ceramic optic fibres were
594 implanted chronically in the injection sites.

595

596 *Optogenetic stimulation*

597 Mice were injected bilaterally with 500 nl/hemisphere of saline containing either AAV2-EF1a-DIO-EYFP
598 (titre: >1 x 10^{12} vg/ml, UNC Vector Core) (YFP group) or rAAV2/Ef1a-DIO-hchR2(E123t/T159C)-EYFP
599 (titre: 5.2 x 10^{12} vg/ml, UNC Vector Core) (ChR2 group) at 2 nl/s in VTA (AP: -3.3, ML: ±0.4, DV: -4.3
600 from bregma). Two 200 µm diameter ceramic optic fibres were implanted chronically in the injection
601 sites at a 10 degree angle.

602

603 **Histology**

604 Mice were terminally anesthetised with sodium pentobarbital and transcranial perfused with saline
605 then 4% PFA solution. 50 µm coronal brain slices were cut, covering striatum and VTA, and
606 immunostained with anti-GFP and anti-TH primary antibodies, and 488 and Cy5 secondary antibodies.
607 An Olympus FV3000 microscope was used to image the slices.

608

609 **Photometry recordings**

610 Dopamine calcium activity (GCaMP) and release (dLight) were recorded at 130Hz sampling rate using
611 pyPhotometry[52]. The optical system comprised a 465 nm and a 560 nm LED, a 5 port minicube and
612 fibre optic rotary joint (Doric Lenses), and two Newport 2151 photoreceivers. Time division

613 illumination with background subtraction was used to prevent crosstalk between fluorophores due to
614 the overlap of their emission spectra, and changes in ambient light from affecting the signal[52].

615 Photometry signals were pre-processed using custom Python code. A median filter (width 5 samples)
616 was first used to remove any spikes due to electrical noise picked up by the photodetectors. To correct
617 for bleaching of the fibre and fluorophores, signals were detrended using a 3rd order polynomial fit.
618 Afterwards, a 5Hz low-pass filter was used to denoise the signal. To motion-correct the signals, we
619 band-passed the detrended signals between 0.2 Hz and 5 Hz (motion band), and used linear regression
620 to predict the GCaMP or dLight signal using the control fluorophore (tdTomato) signal. The predicted
621 signal due to motion was subtracted from the detrended signal. Finally, the pre-processed dopamine
622 signal was z-scored for each session to allow comparison across sessions and animals with different
623 signal intensities.

624 For the GCaMP data, we recorded from 12, 11 and 12 mice in VTA, NAc and DMS, respectively (in one
625 animal, the fibre targeting NAc did not presented GCaMP modulation, which later histological analysis
626 confirmed was caused by the fibre being in the anterior commissure). In DMS, 2 animals were
627 excluded from the main analysis as they presented a negative modulation to reward, however, they
628 were included for the M-L gradient analysis of DMS signal (Figure S7).

629 For the dLight data, we recorded from 5 and 6 mice in NAc and DMS, respectively (1 subject in NAc
630 did not presented any dLight modulation, posterior histological analysis confirmed that the fibre was
631 misplaced into the ventricle).

632

633 **Optogenetic manipulation**

634 Dopamine population activity in VTA was stimulated bilaterally using two 465 nm LEDs (Plexon
635 Plexbright) connected to 200 µm optic fiber patch cords. All stimulation experiments used stimulation
636 parameters of 5 pulses at 25 Hz, 25% duty cycle, ~8-10mW optical power at the fibre tip.

637 First, as a positive control, we performed an intracranial self-stimulation (ICSS) assay in which mice
638 were presented with either 4 or 2 nose poke ports, one of which triggered optical stimulation when
639 poked. A minimum 1-second delay was imposed between stimulations. Mice were tested on ICSS on
640 40-60min sessions for 4 days.

641 We then tested the effect of optogenetic manipulation during the two-step task. On each stimulation
642 session, we either stimulated (i) 200ms after the first-step choice at the time of second-step cue onset,
643 or (ii) at outcome-cue onset. Stimulation occurred on 25% of the trials, under the constraints that (i)
644 the trial after stimulation was always a free choice trial, and (ii) there were always at least 2 non-

stimulated trials after each stimulation. The stimulation sessions were interspersed with baseline no stimulation sessions (data not shown). The timing of stimulation was fixed within a session, with the order counterbalanced across animals (e.g., no stimulation → second-step cue stimulation → outcome stimulation → no stimulation → outcome stimulation → second-step cue stimulation, etc).

Analysis

All behavioural and photometry analysis was performed using custom Python code.

Behavioural logistic regression model

The logistic regression analysis shown in Figure 1E predicted repeating choices (or "staying") as a function of the subsequent trial events, considering only free choice trials, implemented as a mixed effects model using the afex package[53] in R programming language. The model formula was:

$$stay = correct + bias + transition * outcome + (random\ effects|\ subject)$$

For the analysis of two-step optogenetic manipulation (Figure 5C,E), we added stimulation and its interaction with trials events as addition predictors, giving the formula:

$$stay = correct + bias + transition * outcome * stim + (random\ effects||\ subject)$$

We used orthogonal sum-to-zero contrasts, and likelihood ratio test to calculate p-values. The random effects structure used subject as a grouping factor and all random slopes and intercepts that could be included without having a singular fit. This was achieved by removing any random effects fitting to close to zero variance.

The predictors were coded as:

- *Correct*: 3 level categorical variable indicating whether the previous choice was correct, incorrect or in a neutral block.
- *Bias*: Binary categorical variable explaining whether previous choice was left or right.
- *Outcome*: Binary categorical variable indicating whether the previous trial was rewarded or not.
- *Transition*: Binary categorical variable indicating whether previous trial had a common or a rare transition.
- *Stimulation*: Binary categorical variable indicating whether the previous trial was stimulated or not. This regressor was only used for the optogenetic analysis.

For the optogenetic manipulation analysis, we performed separate regressions per group (YFP & ChR2) and stimulation time (after choice & outcome). To ensure stimulated and non-stimulated trials had matching histories, we only included trials where stimulation could have potentially been delivered, i.e., excluding the 2 trials following each stimulation where stimulation was never delivered.

The lagged logistic regression analysis (Figure 1F) assessed how subjects' choices were affected by the history of trial events over the last 12 trials. The regression predicted subjects' probability of choosing left, using the following predictors at lags 1, 2, 3-4, 5-8, 9-12 (where lag 3-4 etc. means the sum of the individual trial predictors over the specified range of lags).

- *Common Transition*: Rewarded at lag n: +0.5/-0.5 if the n'th previous trial was a left/right choice followed by a common transition and reward, 0 otherwise.
- *Rare Transition*: Rewarded at lag n: +0.5/-0.5 if the n'th previous trial was a left/right choice followed by a rare transition and reward, 0 otherwise.
- *Common Transition*: Unrewarded at lag n: +0.5/-0.5 if the n'th previous trial was a left/right choice followed by a common transition and no reward, 0 otherwise.
- *Rare Transition*: Unrewarded at lag n: +0.5/-0.5 if the n'th previous trial was a left/right choice followed by a rare transition and no reward, 0 otherwise.

The lagged regression was fitted separately for each subject as a fixed effects model. Cross subjects mean and standard error for each predictor coefficient were plotted, significance of coefficients was assessed using a t-test comparing the distribution of the individual subjects' coefficients against zero, and performing Bonferroni multiple comparison correction.

Behavioural models

We evaluated goodness-of-fit to subjects' choices for a set of different reinforcement learning (RL) agents created by combining one or more of the following learning strategies:

<u>Model-free</u>

The model-free strategy updated value $Q_{MF}(c)$ of the chosen action and value $V(s)$ of second-step state as:

$$Q_{MF}(c) \leftarrow (1 - \alpha)Q_{MF}(c) + \alpha((1 - \lambda)V(s) + \lambda r)$$

$$V(s) \leftarrow (1 - \alpha)V(s) + \alpha r$$

Where α is the learning rate, λ is the eligibility trace parameter and r is the outcome.

Model-based

The model-based strategy updated value $V(s)$ of the second-step reached, and both first-step action values $Q_{MB}(a)$ as:

$$V(s) \leftarrow (1 - \alpha)V(s) + \alpha r$$

$$Q_{MB}(a) = \sum_{s} P(s|a)V(s)$$

Where α is the learning rate, r is the outcome, $P(s|a)$ is the transition probability of reaching second-step s after taking action a (i.e. 0.8 or 0.2). We also considered variants on this strategy which updated their estimates of the transition probabilities based on the experienced transitions.

Inference

The inference strategy used a single variable B to track the anticorrelated reward probabilities at the two second-step states, representing subjects' belief about which second-step state currently had high reward probability. This was updated as:

$$B \leftarrow (1 - \alpha_r)B + \alpha_r S \qquad \text{if previous trial rewarded}$$

$$B \leftarrow (1 - \alpha_n)B + \alpha_r(1 - S) \qquad \text{if previous trial not rewarded}$$

Where α_r is a learning rate from rewarded trials, α_n is a learning rate from non-rewarded trials, and S is a variable which takes value 1 if the up second-step state was reached on the trial and 0 if the down second-step state was reached.

The variable B determined the value of both the second-step states (used only for predicting dopamine activity in Figures 4 & S6) and first-step actions as:

$$V_{Inf}(up) = B$$

$$V_{Inf}(down) = (1 - B)$$

$$Q_{Inf}(a \rightarrow up) = B$$

$$Q_{Inf}(a \rightarrow down) = (1 - B)$$

Where $V_{Inf}(up)$ and$V_{Inf}(down)$ are the values of the up and down second-step states, and $Q_{Inf}(a \rightarrow up)$ and $Q_{Inf}(a \rightarrow down)$ are the values of the actions which commonly lead to the up and down state.

Reward-as-cue

We considered another strategy in which the most recent trial's outcome acts as a discriminative stimulus for the state of the task, termed 'reward-as-cue' in [11]. This strategy has fixed habits (here expressed as first-step action values) deployed depending on the previous trial's outcome and second-step state where it was obtained:

If previous trial rewarded:

$$Q_{RC}(a \rightarrow s) = z$$

$$Q_{RC}(a \rightarrow s') = 0$$

If previous trial not rewarded:

$$Q_{RC}(a \rightarrow s) = 0$$

$$Q_{RC}(a \rightarrow s') = (1 - z)$$

Where $Q_{RC}(a \rightarrow s)$ is the first-step action that commonly leads to the second-step state reached on the previous trial, $Q_{RC}(a \rightarrow s')$ is the first-step action that commonly leads to the second-step state not reached on the previous trial, and z is a variable between 0 and 1 which determines the relative strength with which reward and reward omission affect the next trial choice.

Combined action values:

A set of different candidate models were created by combining one of more of the above strategies in a weighted sum with (optionally) bias and perseveration parameters, to give net action values:

$$Q_{net}(a) = \sum_i w_i Q_i(a) + K(a)$$

Where w_i is the weight assigned to strategy i whose first step action values are given by $Q_i(a)$, and $K(a)$ is the modifier to the value of first-step action a due to any bias or perseveration terms included in the model. In models where bias was included this increased the value of the left action by a parametrised amount on all trials. In models where perseveration was included this increased the value of the first step action chosen on the previous trial by a parametrised amount.

The combined action values determined choice probabilities via a softmax decision rule as:

$$P(a) = \frac{e^{Q_{net}(a)}}{\sum_a e^{Q_{net}(a)}}$$

Model fitting and comparison

We generated a total of 46 different individual models from the following classes:

Model-free*:* These models used only the model-free strategy, and varied with respect to whether they had asymmetric learning rates for positive and negative outcomes, and whether they included perseveration and/or bias.

Model-based*:* These models used only the model-based strategy, and varied with respect to whether they had asymmetric learning rates for positive and negative outcomes, and whether they included perseveration and/or bias.

Hybrid*:* These models used both the model-based and model-free strategy, and varied with respect to whether they had asymmetric learning rates for positive and negative outcomes, and whether they included perseveration and/or bias.

Inference*:* These models used the inference and/or the reward-as-cue strategies, and varied with respect to whether they learned from both rewarded and non-rewarded outcomes or just rewarded outcomes, and whether they included perseveration and/or bias.

Each model was fit separately to data from each session using maximum a posteriori fitting with priors: Beta (α=2, β=2) for unit range parameters, Gamma (α=2, β=0.4) for positive range parameters, and Normal (σ=5) for unconstrained parameters. For each session, the optimisation was repeated 15 times starting with randomised initial parameter values drawn from the prior distributions, and the best of these fits was used. Model comparison used the Bayesian Information criterion, which was computed per session, and summed across sessions. In Figure 1G we show only the best fitting model from each of the 4 above classes, the full set of compared models is shown in Figure S1A.

The overall best fitting model combined the inference and reward-as-cue learning strategies with bias, but without perseveration. Both strategies were only influenced by rewarded but not non-rewarded trials - i.e. in the inference strategy the learning rate for non-rewarded trials $\alpha_n = 0$, and in the reward-as-cue strategy, the parameter determining asymmetry between rewarded and non-rewarded trials $z = 1$.

To compare data simulated from the model with real data (Figure S1), we simulated 540 sessions of 350 trials each, using parameters values drawn from the distribution of each animal's fits across sessions.

Photometry analysis

Photometry signals were aligned across trials by linearly time-warping the signal at the two points in the trial where timings were determined by subject behaviour; between initiation and choice, and between the second-step port illuminating and being poked (Figure 2B). Activity at other time periods was not warped.

For the analyses presented in Figures 3, 4, S3-6, we used linear regression to predict trial-by-trial dopamine activity at each timepoint in a trial as:

$$y(i,t) = \sum_{p} \beta_p(t) X_p(i) + \beta_o(t) + \varepsilon(i,\ t)$$

where $y(i,t)$ is the calcium z-scored activity on trial i at timepoint t, $\beta_p(t)$ is the weight for predictor p at timepoint t, $X_p(i)$ is the value of the predictor p on trial i, $\beta_o(t)$ is the intercept at timepoint t, and $\varepsilon(i,\ t)$ is the residual unexplained variance.

The linear regressions in Figures 3 & S3-4 used the following predictors:

- *Reward*: +0.5 if current trial is rewarded, -0.5 otherwise
- *Good Second-Step*: +0.5 is the second step reached on the current trial has high reward probability , -0.5 if low reward probability, 0 if neutral block.
- *Correct Choice*: +0.5 if subject current trial choice commonly leads to the high reward probability second-step port, -0.5 if to the low reward probability second-step port, 0 if neutral block.
- *Repeat Choice*: +0.5 if same choice as previous choice, -0.5 if different choice to previous trial.
- *Direct Reinforcement Action Value Update*: +0.5/-0.5 if current choice is the same as the previous choice and the previous trial was rewarded/not-rewarded, 0 if different choice from previous trial.
- *Model-Consistent Action Value Update*: +0.5/-0.5 if current choice commonly leads to the previous second-step when it was rewarded/not-rewarded, -0.5/+0.5 if current choice rarely leads to the previous second-step when it was rewarded/not-rewarded.

- *Previous Reward, Same Second-Step*: +0.5 if current second-step is the same as in the previous trial and previous trial was rewarded, -0.5 if same second-step and previous trial was unrewarded, 0 if current second-step different from second-step on the previous trial.
- *Previous Reward, Different Second-Step*: +0.5 if current second-step is different from previous second-step and previous trial was rewarded, -0.5 if different second-step and previous second-step was unrewarded, 0 if current second-step is the same as the second-step on the previous trial.
- *Common Transition*: +0.5 if a common transition occurs on the current trial, -0.5 if a rare transition.
- *Reward Rate*: exponential moving average of the recent reward rate (tau=10 trials).
- *Contralateral Choice*: +0.5 if current choice is in the contralateral side from the recording site, -0.5 if current choice is in the ipsilateral side.
- *Up Second-Step:* +0.5 if current second-step is up, -0.5 if current second-step is down.

The linear regression used in Figure S5 use the same predictors as above (figures 3 & S3-4) except (i) second-step value update predictors (*Previous Reward, Same Second-Step* and *Previous Reward, Different Second-Step*) and (ii) *model-consistent action value update*, which were replaced with predictors split by outcome as:

- *Same vs Different Second-Step, Previously Rewarded*: +0.5/-0.5 if current second-step is the same/different as in the previous trial and the previous trial was rewarded, 0 if the previous trial was not rewarded.
- *Same vs Different Second-Step, Previously Non-Rewarded*: +0.5/-0.5 if current second-step is the same/different as in the previous trial and the previous trial was not rewarded, 0 if the previous trial was rewarded.
- *Model-consistent Action Value Update from Rewarded Trials*: +0.5/-0.5 if current choice commonly/rarely leads to the previous second-step and the previous trial was rewarded, 0 if the previous trial was not rewarded.
- *Model-consistent Action Value Update from Unrewarded Trials*: -0.5/+0.5 if current choice commonly/rarely leads to the previous second-step and the previous trial was not rewarded, 0 if the previous trial was rewarded.

The linear regressions used in figures 4 & S6 used the above-described *Reward* & *Reward Rate* regressors with the following additional regressors:

- *Previous Reward:* +0.5 if previous trial was rewarded, -0.5 otherwise
- *Second-Step Value:* The value $V_{Inf}(s)$ of the second-step reached on the current trial from the best fitting RL model.
- *Chosen Action Value*: The value $Q_{Inf}(c)$ of the first-step action chosen on the current trial from the best fitting RL model.

SUPPLEMENTARY TABLES AND FIGURES

		Fitting parameters								
Best fitting model from each agent type	**BIC**	**alpha**	**iTemp**	**eligibility trace**	**weight MF**	**weight MB**	**reversal prob.**	**bias**	**rew bias.**	**persv.**
MB inference pos. out. + bias + rew. bias	283.73 ± 92.92	0.32 ± 0.17	2.24 ± 1.00					0.03 ± 0.33	1.14 ± 1.01	
MB + bias + rew. bias + persv.	286.11 ± 93.55	0.23 ± 0.15	3.04 ± 1.86					0.03 ± 0.33	0.92 ± 0.88	0.35 ± 0.40
MF + bias + rew. bias + persv.	292.94 ± 94.84	0.29 ± 0.17	2.64 ± 1.44	0.39 ± 0.17				0.04 ± 0.33	1.11 ± 0.90	0.36 ± 0.39
Hybrid + bias + rew. bias + persv.	295.57 ± 94.22	0.27 ± 0.19		0.40 ± 0.15	1.79 ± 1.28	1.70 ± 1.06		0.05 ± 0.64	2.10 ± 1.49	0.74 ± 0.59
Control (bias + rew.bias + persv.)	298.57 ± 98.66		2.64 ± 0.11					0.01 ± 0.14	0.99 ± 0.54	0.30 ± 0.21

Supplementary Table 1. Best RL model fits. Fitting parameters for each model (Mean ± standard deviation) (Models shown in Figure 1G)

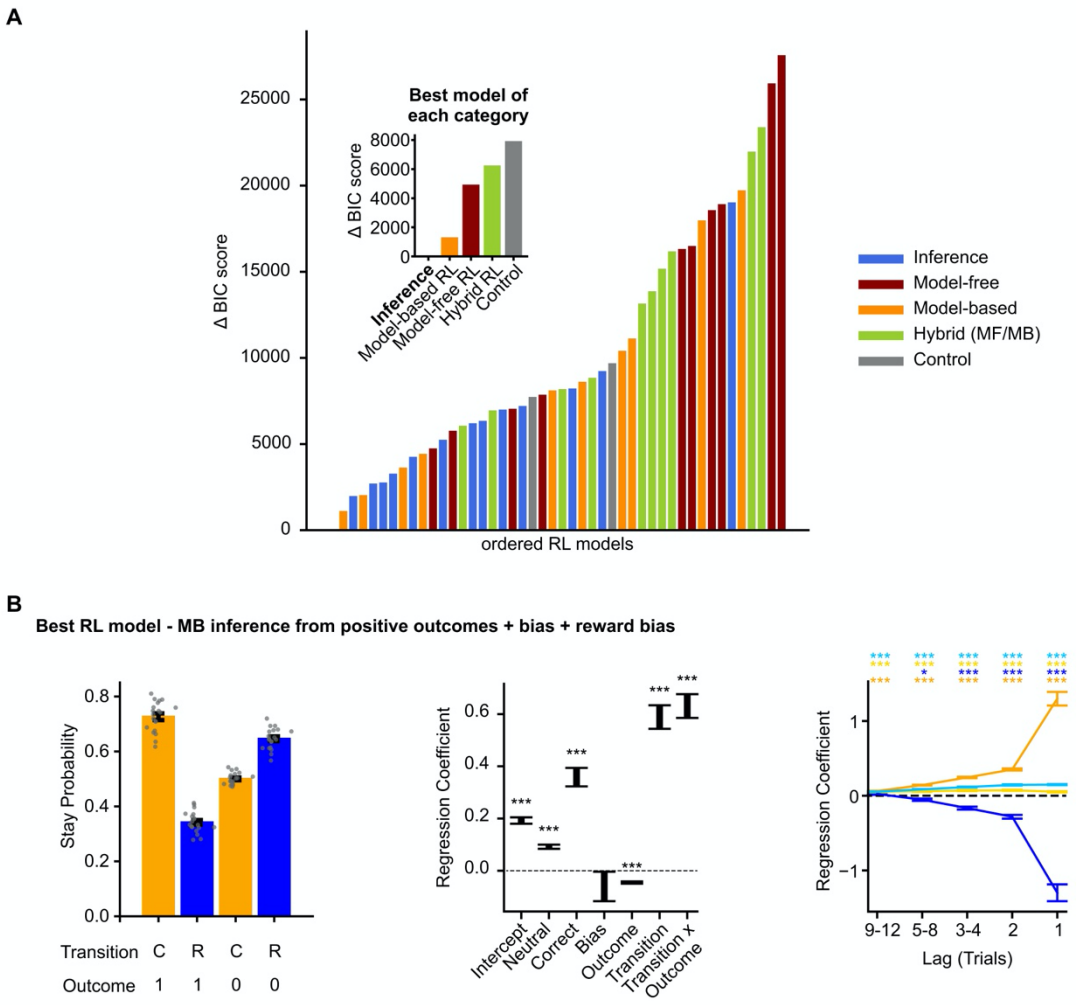

Supplementary Figure 1. Model comparison and simulated behaviour. A) BIC score model-comparison of all the models tested. Colours reflect the RL model type. Each model type was fitted a different combination of the following kernels: choice bias, reward bias (reward-as-cue), choice perseveration. Models were also tested with same or different learning rates for rewarded and unrewarded trials or just value updating from positive outcome (this latter one just for models with inference learning). **B)** Best fitting model: model-based inference learning only from positive outcomes with bias and reward bias. *p<0.05, **p<0.01, ***p<0.001.

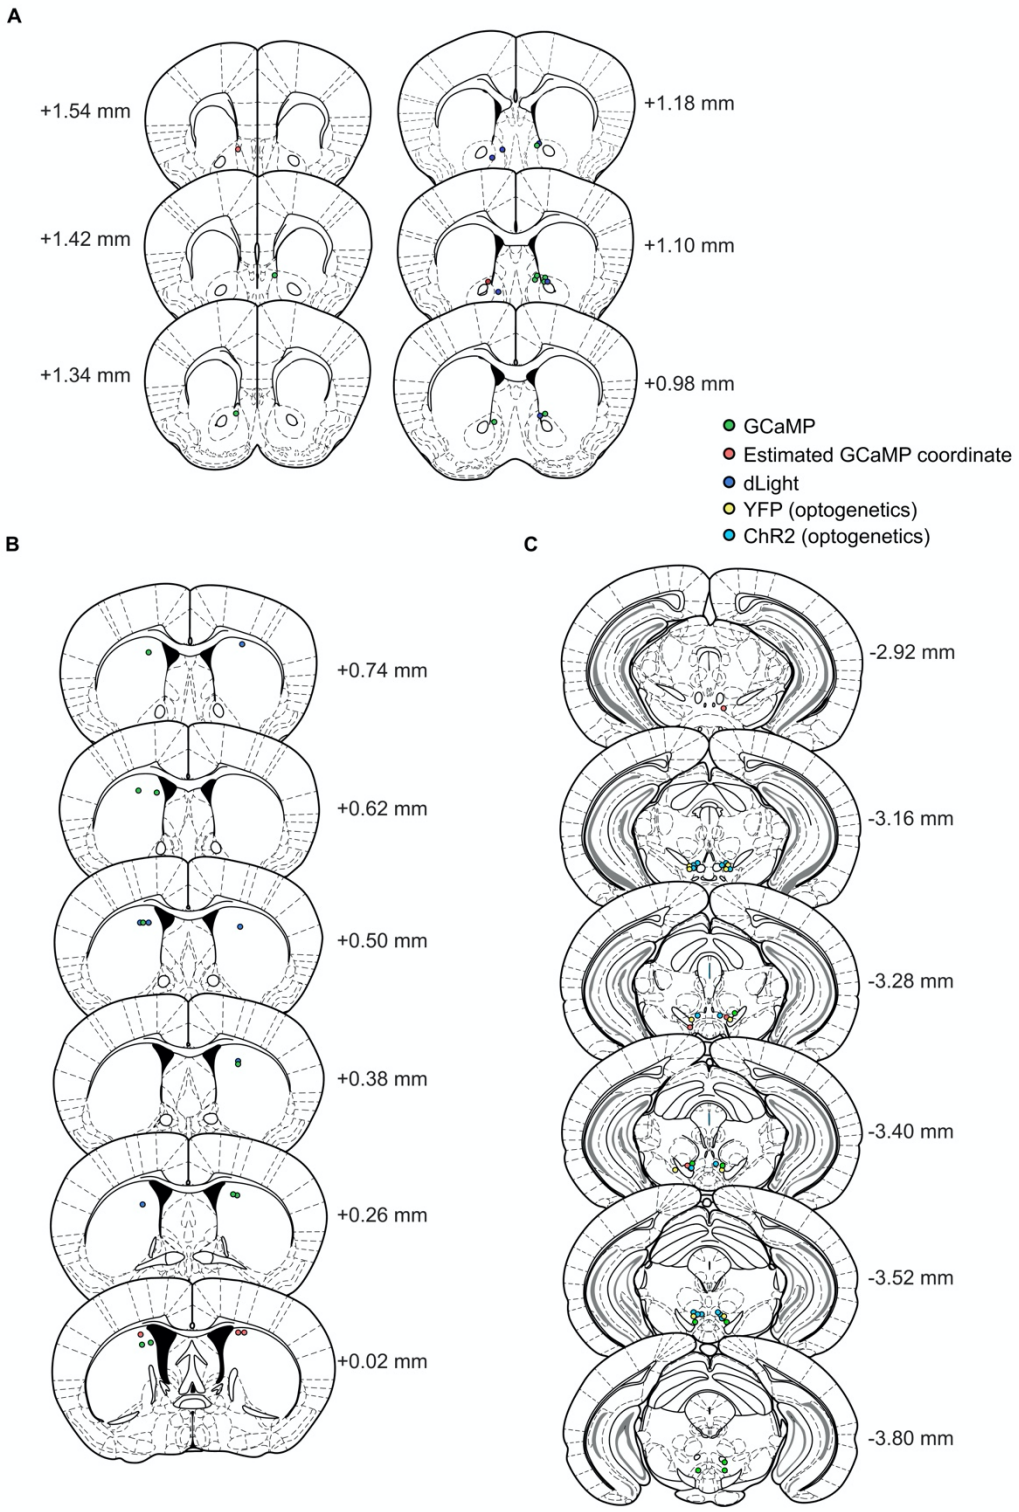

Supplementary Figure 2. Fibre placement in **A)** NAc, **B)** DMS, and **C)** VTA. Green, GCaMP6f; red, GCaMP6f animals where the tip of the fibre placement is estimated; dark blue, dLight1.1; yellow, YFP animals from the optogenetics experiment; light blue, ChR2 animals.

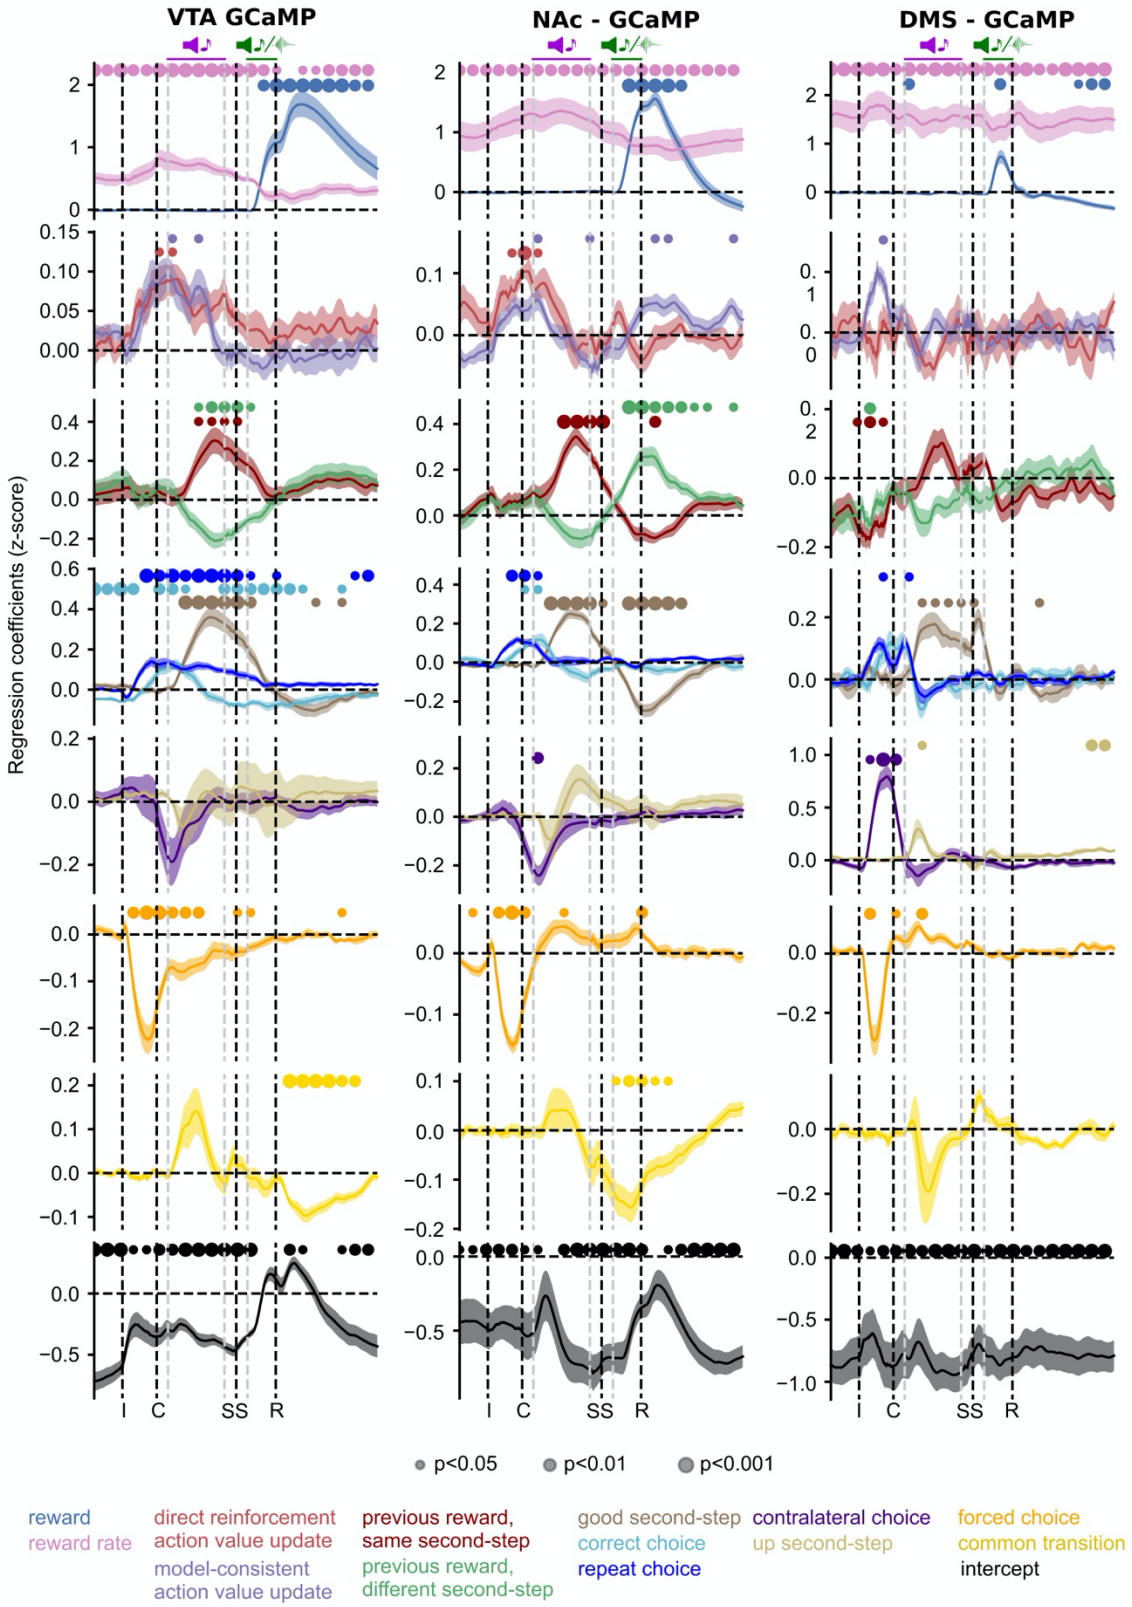

Supplementary Figure 3. Regression coefficients from all the predictors used in the behavioural logistic regression model predicting VTA, NAc and DMS dopamine activity (GCaMP recordings). Shaded area indicates cross subject standard error. Shaded area indicates cross subject standard error. Dots represents statistical significance after Benjamini-Hochberg multiple comparison correction, p<0.05, p<0.001, p<0.001.

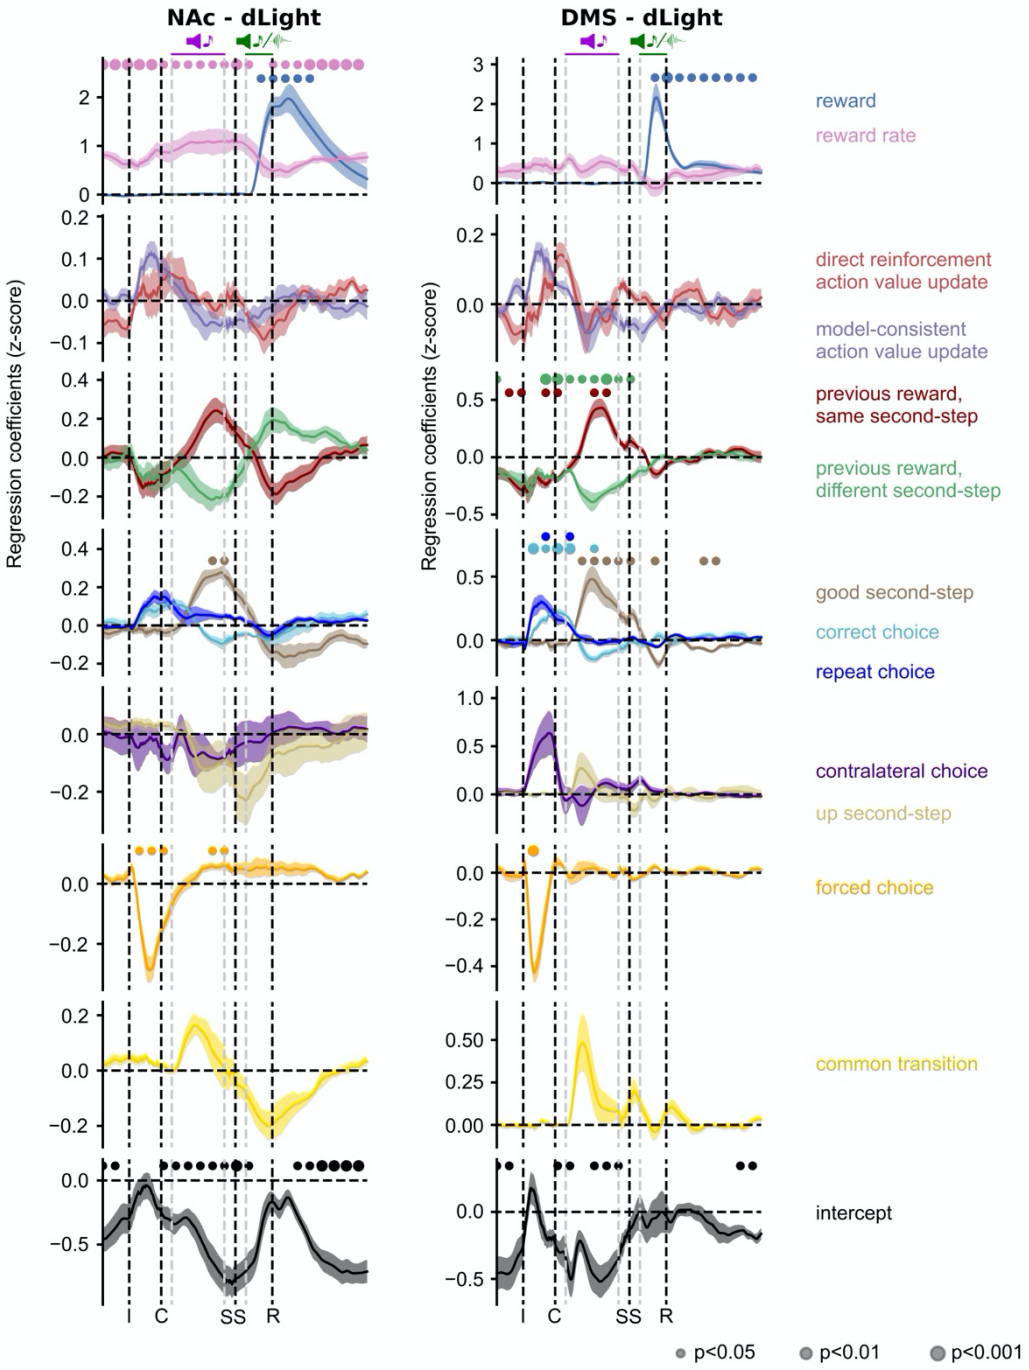

Supplementary Figure 4. Regression coefficients from all the predictors used in the behavioural logistic regression model predicting NAc and DMS dopamine receptor binding (dLight recordings). Shaded area indicates cross subject standard error. Dots represents statistical significance after Benjamini-Hochberg multiple comparison correction, p<0.05, p<0.001, p<0.001.

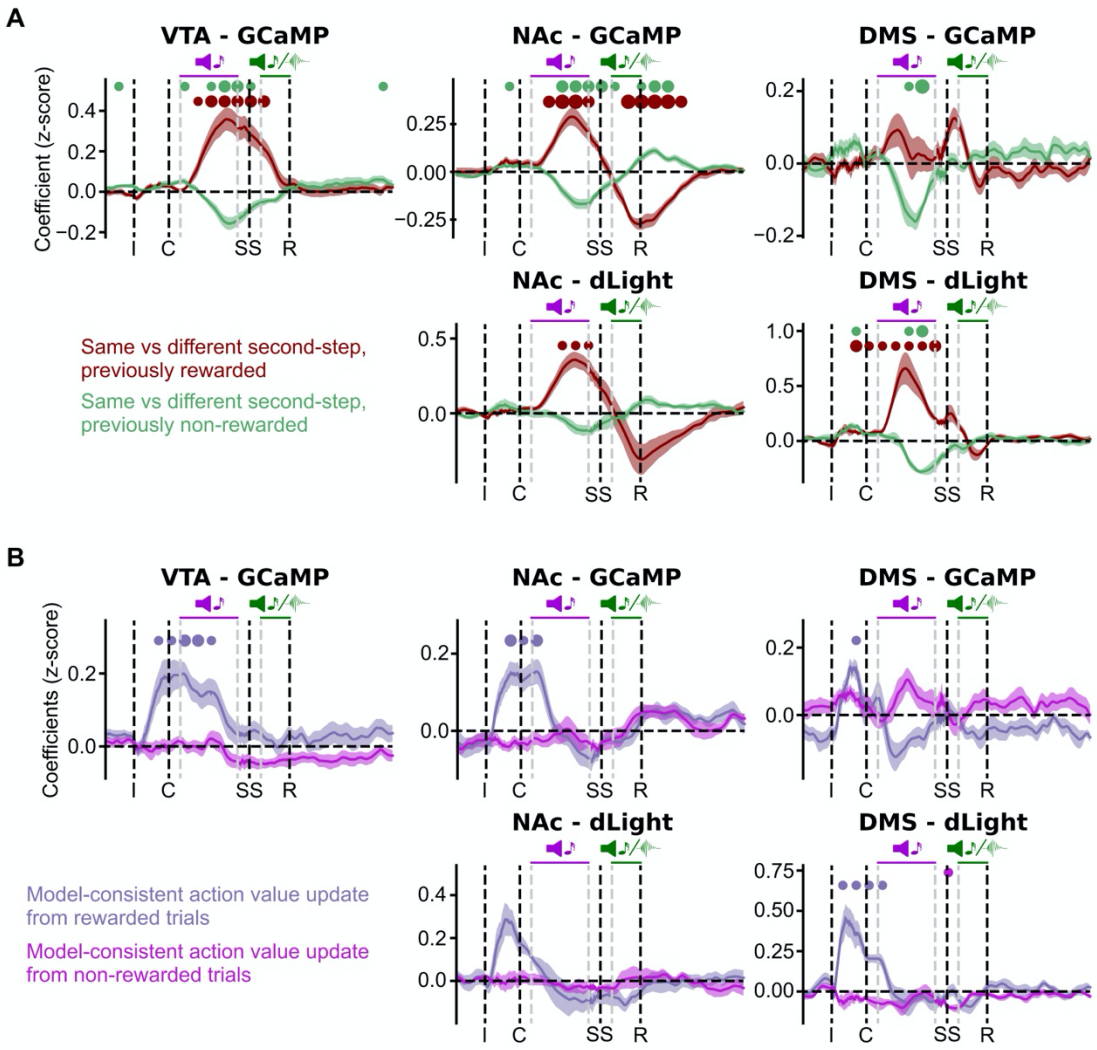

Supplementary Figure 5. Value update regressors split by rewarded/not rewarded outcome. A) Second-step value update based on whether previous trial was rewarded or not. **B)** Model-based action value update split by outcome on the previous trial. Shaded area indicates cross subject standard error. Dots represents statistical significance after Benjamini-Hochberg multiple comparison correction, $p<0.05$, $p<0.001$, $p<0.001$.

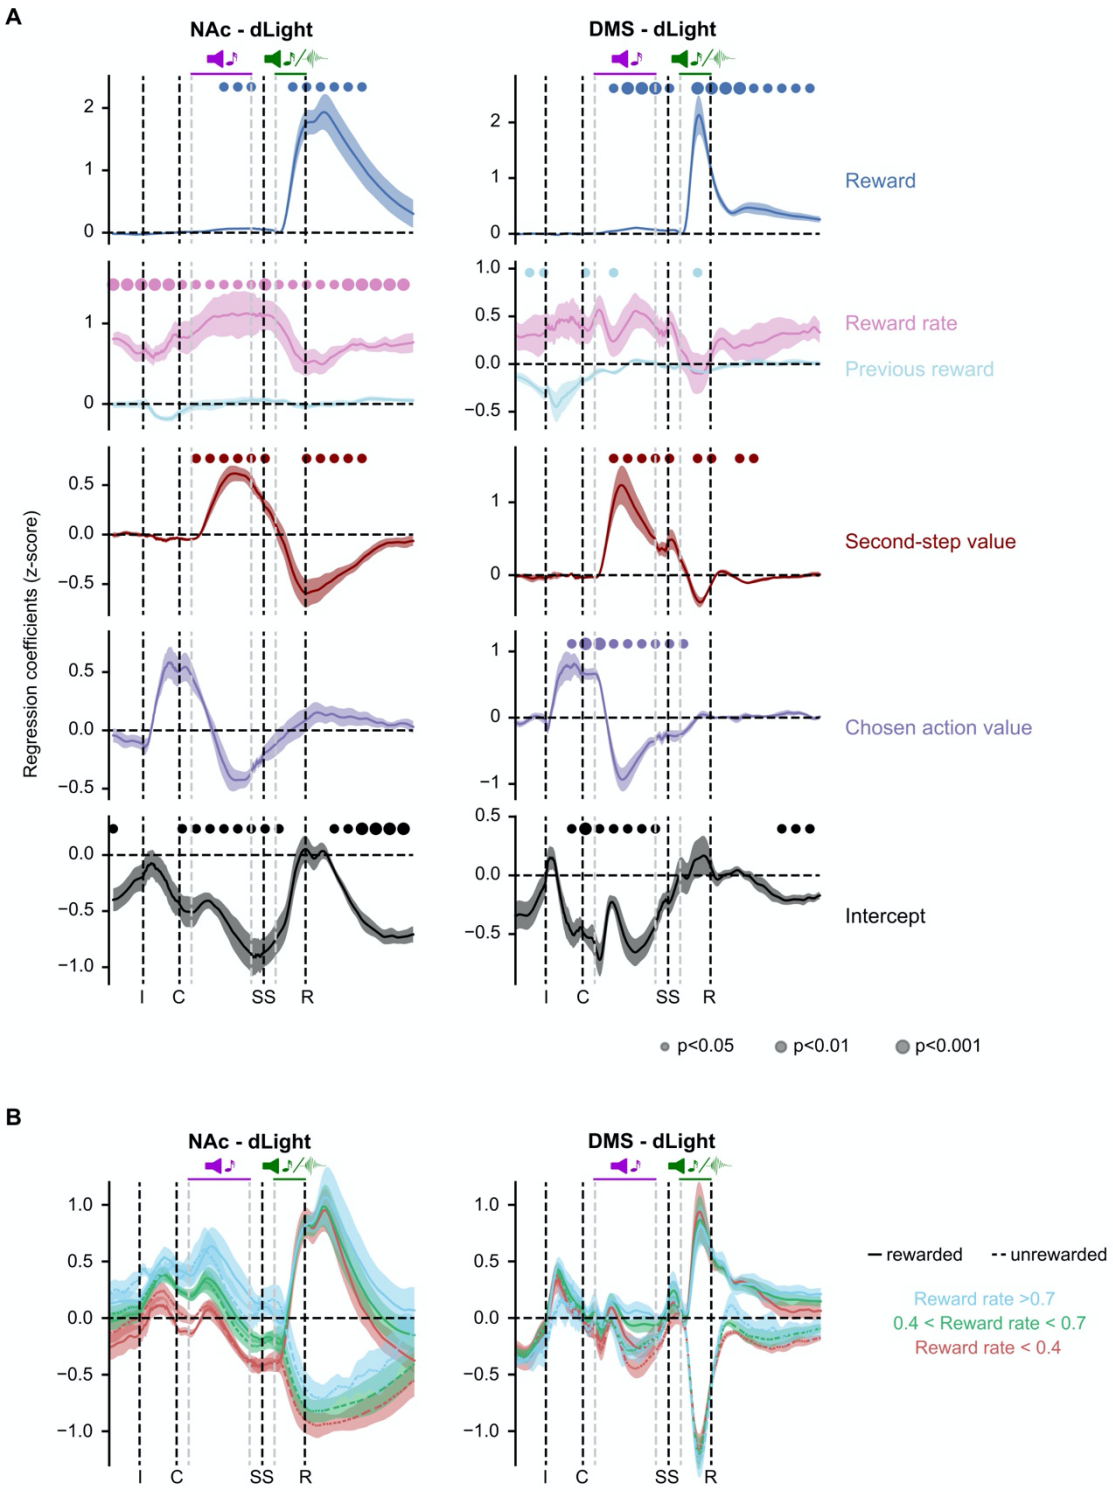

Supplementary Figure 6. A) Regression model using second-step and action values derived from the best RL fitting model predicting dopamine concentrations (dLight) in NAc (left) and DMS (right). Shaded area indicates cross subject standard error. Dots represents statistical significance after Benjamini-Hochberg multiple comparison correction, p<0.05, p<0.001, p<0.001. **B)** Mean z-scored signal across a trial split between rewarded (continuous trace) and unrewarded (dashed trace) trials, and split between high reward rate (>0.7 exponential moving average of rewards, tau=10) in blue, medium reward rate (between 0.4 and 0.7) in green, and low reward rate (<0.4) in red.

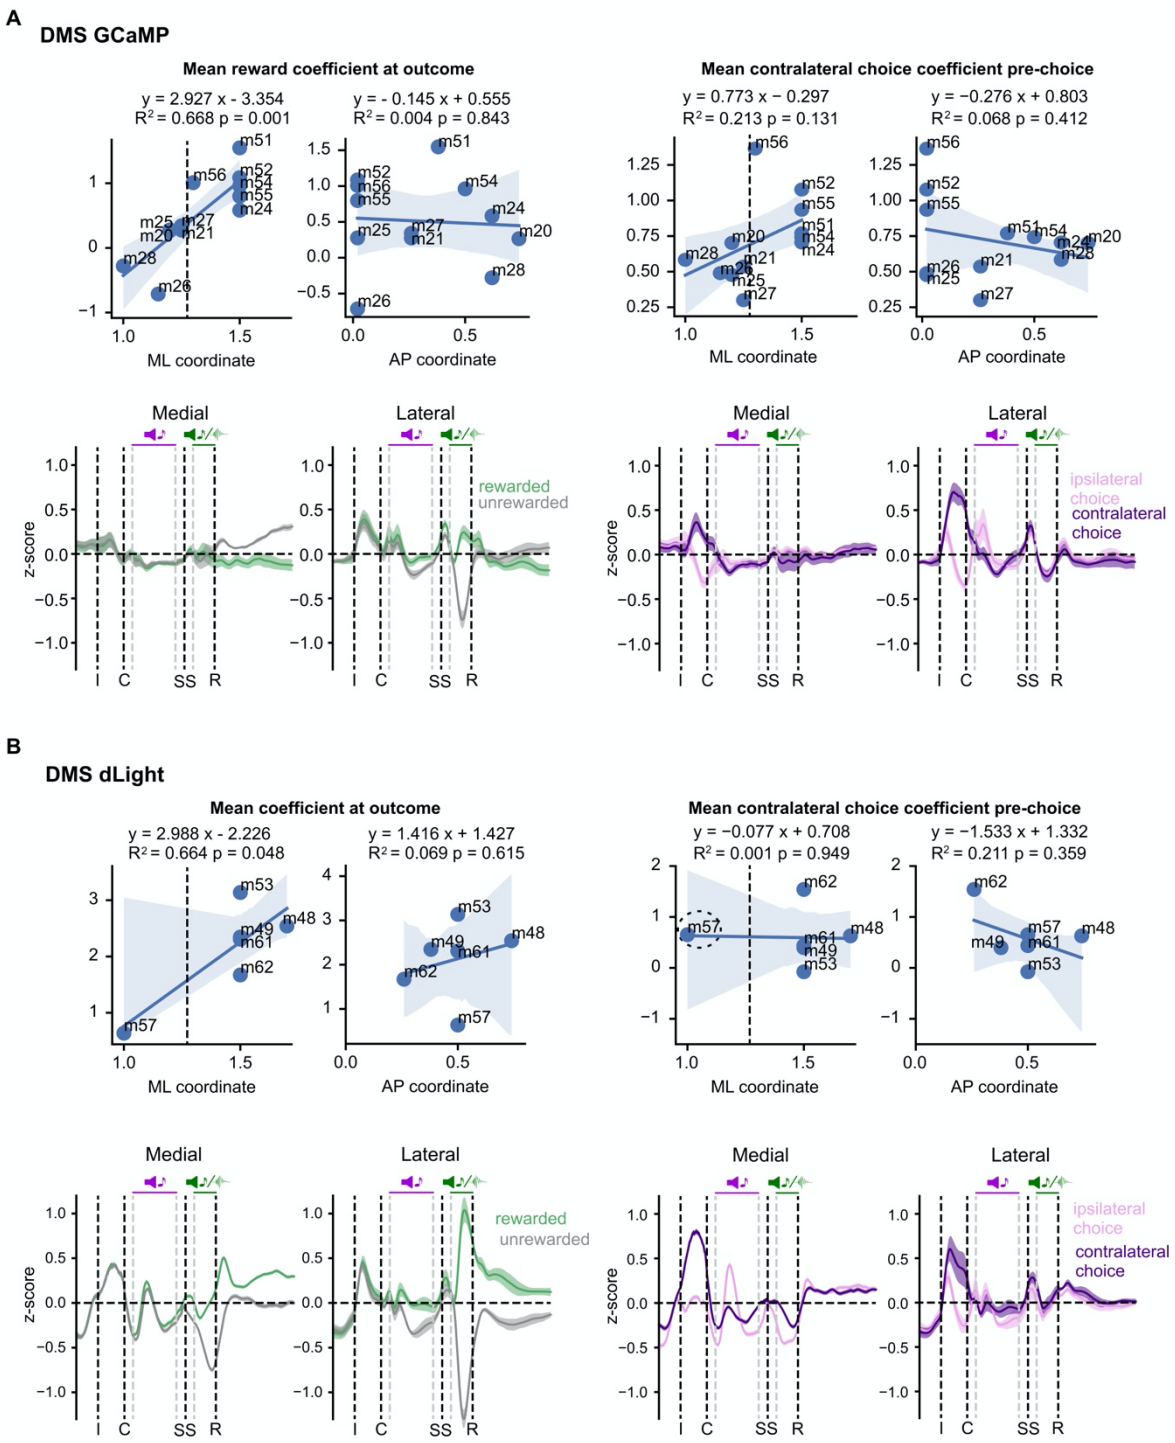

Supplementary Figure 7. Medio-lateral gradient in reward modulation in DMS. A) Dopamine activity measured through GCaMP. **B)** Dopamine receptor binding measured through dLight. *Left*, mean reward coefficient (from the linear regression model) at the time of outcome cue; *Right*, mean contralateral choice coefficient pre-choice. *Top*, correlation between mean coefficient and medio-lateral (ML) or antero-posterior (AP) location of the optic fibre. *Bottom*, mean z-scored activity split by outcome or lateralised choice for animals with medial or lateral placement in DMS. Dashed line in the correlation plots show how animals where divided in two groups (medial and lateral).

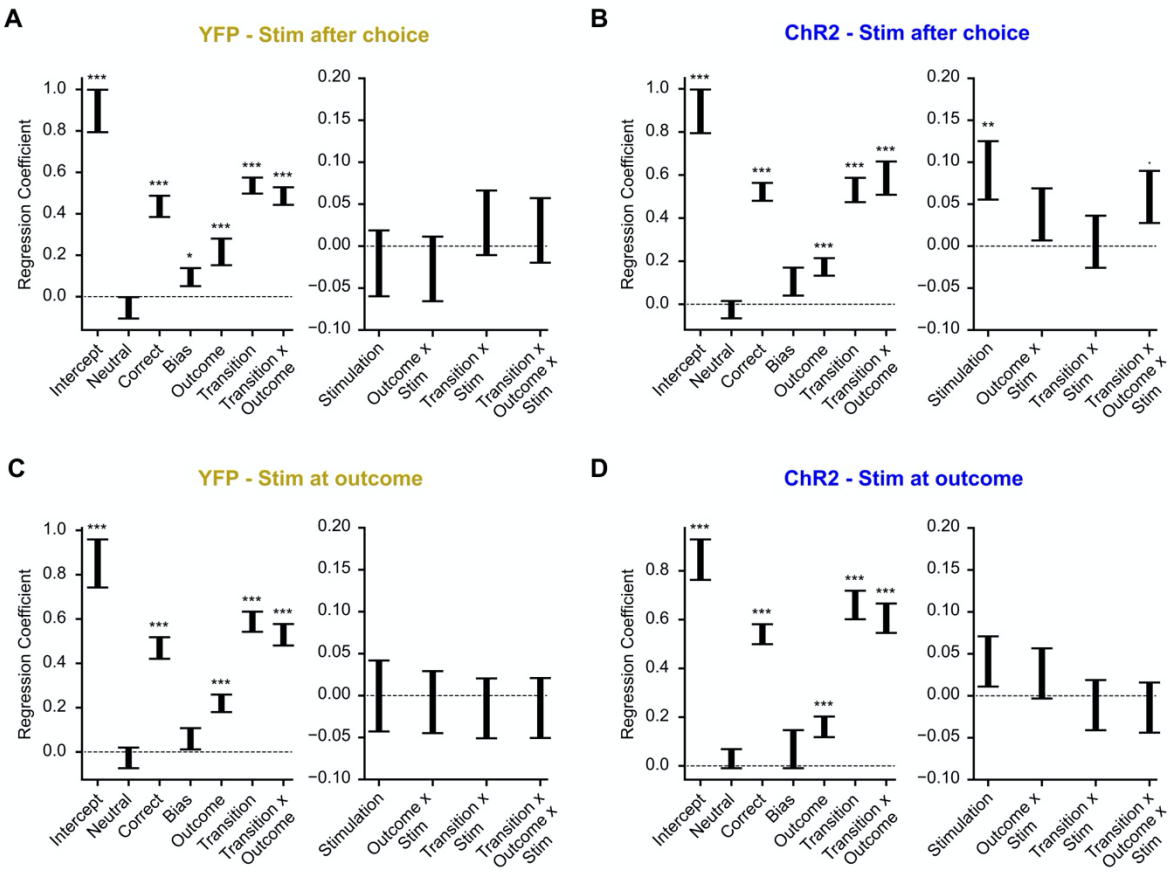

Supplementary Figure 8. Mixed effects logistic regression model predicting repeat choice after optogenetic stimulation in the YFP controls (**A, C**) and ChR2 group (**B, D**) on sessions when stimulation occurred after choice (**A, B**) or at outcome (**C, D**). •p<0.1, *p<0.05, **p<0.01, ***p<0.001.