Posterior parietal cortex mediates rarity-induced decision bias and learning under uncertainty

Weihao Sheng[1], Xinrui Huang[1], Yang Xie[2], Manolis C. Tsakiris[3], Yang Yang[1*]

1. School of Life Science and Technology, ShanghaiTech University, Shanghai, China.
2. Lingang Laboratory, Shanghai, China
3. Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China

* To whom the correspondence should be addressed (E-mail: yangyang2@shanghaitech.edu.cn)

ABSTRACT

Making decisions when outcomes are uncertain requires judgment of outcome probability, yet such judgments are often inaccurate, owing to reliance on heuristics that introduce systematic errors like overweighting of low probabilities. Here, using a decision-making task in which the participants were unaware of outcome probabilities, we discovered that both humans and mice exhibit a rarity-induced decision bias (RIDB), *i.e.*, a preference towards rare rewards, which persists throughout task performance. Optogenetics experiments demonstrated that activity in the posterior parietal cortex (PPC) is required for the RIDB. Using *in vivo* electrophysiology, we found that rare rewards bidirectionally modulate choice-encoding PPC neurons to bias subsequent decisions towards rare rewards. Learning dynamically recruits PPC neurons and enhances their stimulus encoding, and the activity of these neurons is essential for stimulus-guided decisions. We then developed a dual-agent model that successfully recapitulates the decision-making and learning behaviours, and corroborates the specific functions of PPC neurons in mediating decision-making and learning. This study not only extends our understanding of rare probability overweighting to a context where the outcome probability is unknown and characterizes the neural basis for RIDB in the PPC, but moreover reveals an evolutionarily conserved heuristic that persistently impacts decision-making and learning under uncertainty.

INTRODUCTION

In a world that is probabilistic in nature, the ability to learn from uncertain outcomes and make optimal decisions is essential for survival[1-3]. This behaviour is an active research topic in economics, psychology, and neuroscience [4-6]. Decision making requires accurate judgment of event probabilities, yet these judgments are often biased, as people tend to rely on heuristics and therefore make systematic errors when thinking under uncertainty [7]. One well-established judgement error is the overweighting of low probabilities, *i.e.*, people tend to give more decision weights to rare events, and act as if these events are more likely to occur than they really are [8-11]. This phenomenon results in a concave probability weighting function and constitutes a key component of prospect theory, a fundamental theory in economics [8, 9]. Studies have been done to explore the neural mechanisms underlying such probability distortion when outcome probabilities are explicitly shown to the subjects [12-14].

Previous work in humans and non-human primates has demonstrated that the posterior parietal cortex (PPC) is activated by outcome uncertainty [15-17], and that PPC neurons represent the degree of uncertainty [18-21]. Studies in rodents have demonstrated a variety of functions of the PPC during decision-making: PPC neurons encode choices [22-26] as well as visual [22, 24, 27, 28] and auditory stimuli [29, 30]. Additionally, the PPC encodes short-term memory of stimulus, choice, and outcome [31-34], which are causally related to decision-making behaviour. PPC neurons have also been implicated in learning [24, 29, 35].

Here, using an auditory stimulus-guided decision-making task with probabilistic outcomes, we demonstrate that when the outcome probability is unknown to the subjects, both humans and mice exhibit a rarity-induced decision bias (RIDB) affecting subsequent decisions, which persists throughout task performance. Optogenetics experiments establish that PPC activity is required for RIDB, and *in vivo* recordings show that rare rewards bidirectionally modulate the activity of choice-selective PPC neurons to induce decision bias towards rarely rewarded choices. Moreover, as mice learned to make stimulus-guided correct choices, PPC neurons were dynamically recruited and are essential for stimulus-guided decision-making. We built a data-inspired dual-agent model that accurately reproduces the decision-making and learning behaviours observed for normal and PPC-inhibited conditions, and reveals essential impacts of PPC on the decision bias and learning.

RESULTS

Humans exhibit decision bias towards rarely rewarded choices

To assess the impact of rarity on decision-making under uncertainty, we designed a task with probabilistic outcomes (**Figure 1a**). Human subjects were instructed to make decisions based on auditory stimuli to collect maximal monetary rewards, without being informed of the probabilistic task structure. Pressing the "S" key triggers one of four auditory stimuli: duck, cat, canary, or frog sounds. After a duck or cat sound, pressing the "L" key leads to a monetary reward at 80% probability (common) and no reward at 20% probability (rare); pressing the "R" key leads to a reward of the same amount at 20% probability and no reward at 80% probability. Choice-reward associations were reversed for canary and frog sounds. All stimuli and common/rare outcomes are pseudorandomly interleaved (Methods), so the optimal strategy is to always choose the key with 80% reward probability under the current stimulus, *i.e*., press "L" for duck and cat, and "R" for canary and frog (termed stimulus-guided correct choices; **Extended Data Fig. 1a, b)**. The outcome was shown to the subjects as a money symbol or a no-money symbol upon each choice, but whether it was common or rare was not explicitly shown.

Previous studies have established that uncertain outcomes have a strong influence on decision-making [3]. To quantitatively evaluate the impact of previous outcomes on future decisions, we tracked the subjects' probability of repeating the preceding choices upon receiving 4 different outcomes: 80% rewarded (*comm*+), 20% rewarded (*rare*+), 80% unrewarded (*comm*-), and 20% unrewarded (*rare*-). We found that before subjects learned to reliably make stimulus-guided correct choices (termed naïve stage, defined as performance<70%, **Figure 1b**), a rewarded choice leads to a higher choice repeating probability (RP) in the immediately next trial (termed next-trial RP, nRP) compared to an unrewarded choice (**Figure 1c**, empty bars). This indicates a win-stay, lose-switch (WSLS) tendency, a decision heuristic that promotes repeating of rewarded choices and switching from unrewarded ones. To assess the impact of rare outcomes on future decisions, we examined the difference in nRP when the previous outcome was common *vs*. rare. Rare rewards (*rare*+) triggered a slight increase in win-stay tendency (*i.e.,* increased nRP) and unrewarded rare trials (*rare*-) triggered a significant increase in lose-switch tendency (*i.e.,* reduced nRP) (**Figure 1c**, empty bars), together showing that rare outcomes have a stronger impact on subsequent decisions than common outcomes.

Because this task is stimulus-guided by design, we also inspected the RP for the next same-stimulus trials (termed sRP, **Figure 1c**, filled bars). We observed a much stronger WSLS tendency for the next same-stimulus trial than the immediately next trial (**Figure 1c**), suggesting that choices were more strongly influenced by same-stimulus trials. Focusing on same-stimulus trials, we found a strong decision bias caused by rare rewards: subjects exhibited a significantly higher sRP following *rare*+ than *comm*+ (**Figure 1c**, filled bars). We should emphasize that because *rare*+ choices are opposite to the correct choices, repeating *rare*+ choices resulted in much lower reward probability than repeating *comm*+ choices (**Extended Data Fig. 1c**). In spite of that, subjects still exhibited a significantly higher sRP_{rare+} than sRP_{comm+}, indicating a rarity-induced decision bias (RIDB) induced by preceding rare rewards (**Figure 1c**, filled bars). Thus, in an uncertain environment with probabilistic outcomes, rare rewards strongly influence subsequent decisions, revealing a rarity-induced heuristic that can contribute to overweighting of low probabilities.

A mouse model recapitulates the rarity-induced decision bias

To investigate the neural mechanisms underlying the RIDB observed in humans, we developed a behavioural paradigm in mice by adding probabilistic outcomes to a classical auditory two-alternative forced choice (2AFC) task [36] (**Figure 1d**). The mouse was allowed to freely poke into the centre port, triggering one of two auditory stimuli (A: 5k Hz pure tone; B: 13 kHz). Stimulus A was associated with 80% probability of water reward for left choices and 20% probability for right choices, and vice versa for stimulus B. All mice trained in this task showed a sigmoidal learning curve across sessions (performance calculated as the proportion of correct choices in a session; **Figure 1e, Extended Data Fig. 2a**), similar to an auditory discrimination task with certain outcomes [37].

Like humans, mice exhibited a WSLS tendency both in the immediately next and next same-stimulus trials in the naïve stage (defined as performance < 60%, **Figure 1e, Extended Data Fig. 2b**). Also, like humans, mice exhibited significantly higher sRP_{rare+} than sRP_{comm+} in the naïve stage (**Figure 1f**). Therefore, our behavioural paradigm for mice successfully recapitulates the RIDB found in humans, pointing to an apparently conserved rarity-induced heuristic manifest as a tendency to repeat rarely rewarded choices despite suboptimal outcomes.

We next investigated whether the strength of the RIDB was associated with the rarity levels. We modified the common:rare ratio (0.9:0.1, n = 4 mice; 0.95:0.05, n = 5 mice, **Figure 1f,**

Extended Data Fig. 2c), and quantified ΔsRP (sRP_{rare+} - sRP_{comm+}). We observed larger ΔsRP with higher common:rare ratios (**Figure 1g**), indicating that the rarer the rewards, the stronger the influence. Thus, beyond detecting a rarity-induced heuristic in both humans and mice, we established an empirical system for studying the neural basis of this conserved heuristic.

RIDB persists across all performance levels

Another informative trend we observed was that in both humans and mice, as the performance improves, ΔsRP gradually declines (**Figure 1h, i**). Although it is obvious that ΔsRP > 0 at low performance levels reflects the existence of RIDB, it is important to clarify that ΔsRP < 0 does not indicate the disappearance of RIDB. This is because, in addition to the impact of preceding trials, through which RIDB manifests, there also exists a stimulus-guided factor that influences choices. This factor is related only to the current stimulus and thus independent of RIDB, yet modulating sRP_{comm+} and sRP_{rare+} in opposite directions: it pushes choices towards *comm*+ and away from *rare*+, thus increasing sRP_{comm+} and reducing sRP_{rare+}. As performance improves, more stimulus-guided choices are made, and the modulation by the stimulus-guided factor grows stronger.

To assess RIDB regardless of performance levels, we define M_{rare+} and M_{comm+} as a measure of the influence of *rare*+ and *comm*+ on subsequent choices, respectively (Methods). Thus, RIDB can be evaluated as ΔM (M_{rare+} - M_{comm+}). To compute M_{rare+}, we measured the difference between the conditional probability of making a *rare*+ choice following a *rare*+ trial, and the probability of making a *rare*+ choice following any trial:

$$\mathrm{M}_{rare+} = \mathrm{P}(rare+ | \text{ preceding } rare+) - \mathrm{P}(rare+)$$

Note that P(*rare*+|preceding *rare*+) is the definition of sRP_{rare+}, and P(*rare*+) equals to (1 – performance), performance being the average performance of the session. Similarly,

$$\mathrm{M}_{comm+} = \mathrm{P}(comm+ | \text{ preceding } comm+) - \mathrm{P}(comm+)$$

Here, P(*comm*+|preceding *comm*+) is sRP_{comm+}, and P(*comm*+) equals to performance.

Results show that M_{rare+} is consistently larger than M_{comm+} across all performance levels for humans and mice (**Figure 1j-m, Extended Data Fig. 2d**), indicating that RIDB persisted throughout task performance. Thus, rare rewards have a strong impact on decision-making even after animals were extensively exposed to the task and had learned to predominantly make stimulus-guided decisions.

PPC activity is necessary for generating RIDB

Previous studies in rodents have shown that the posterior parietal cortex (PPC) encodes behavioural history [31-34], which is ostensibly related to the RIDB. To investigate whether the PPC is involved in mediating RIDB, we used optogenetics (eNpHR3.0) to bilaterally inhibit the PPC in rare trials, and evaluated sRP_{rare+} and sRP_{comm+} in the PPC-inhibited mice (**Figure 2a, Extended Data Fig. 3a**). We first inhibited the PPC during the outcome epoch of the rare trials, as this is the epoch during which mice apparently perceive rarity (n=5, *outcome-opto*, Methods). A masking yellow light was delivered in all trials to ensure that mice did not treat the stimulating light as an additional cue. To our disappointment, sRP_{rare+} is still higher than sRP_{comm+} in the naïve stage of *outcome-opto* mice (**Figure 2b, c**, yellow; **Extended Data Fig. 3b**). However, for a separate group of mice in which we inhibited the PPC during the stimulus epoch of rare trials (n=5, *stim-opto*), we found a significant reduction in ΔsRP in the naïve stage as compared to non-inhibited controls (**Figure 2b, c,** orange). Further, we observed a significant reduction of ΔsRP as well as ΔM (representing RIDB) in *opto-stim* mice compared to non-inhibited controls, across a wide range of performance levels (**Figure 2d-g, Extended Data Fig. 3c, d**), corroborating the notion that RIDB persisted even at high performance levels. These results indicate that PPC activity is required for generating the RIDB.

Rare rewards bidirectionally modulate choice-selective PPC neurons to induce RIDB

To investigate the neural mechanisms underlying the RIDB in the PPC, we made extracellular recordings in the mouse PPC during task performance (5 mice, 100 behavioural sessions, 3135 units, **Figure 3a, Extended Data Fig. 4a-c, 5a**). Previous rodent studies have shown that PPC activity can drive the animal's choices [12, 19, 20]. Thus, it is plausible that preceding *rare*+ trials somehow modulate PPC activity to influence subsequent choices. To investigate the rarity-induced modulation, we first examined how a *rare*+ trial influences PPC neuronal activity in the next same-stimulus trial. We separated trials into 4 groups based on the outcome of the preceding trial (**Figure 3a, b**). Then, we used ROC analysis to calculate a *rare*+ modulation index (RMI) that measures how strongly a neuron's firing rates following *rare*+ trials differ from that neuron's firing rates following other trials. We observed that ~15% of recorded PPC neurons were differentially modulated by preceding *rare*+ trials compared to other trials (**Figure 3b, c**), among which about half were significantly enhanced by preceding *rare*+ and the others were significantly suppressed

(**Figure 3c**). These results establish that *rare*+ modulates PPC neuronal activity in the next same-stimulus trial. Further, and consistent with the aforementioned persistence of RIDB across performance levels, we found that *rare*+ modulates PPC activity across all performance levels (**Extended Data Fig. 6a, b**).

We subsequently investigated how this rarity-induced modulation of PPC activity influences choices (**Figure 3d**). Using ROC analysis to compute the choice preference index (CPI) of individual PPC neurons, we found that a majority of PPC neurons either preferred right choices (*i.e.,* higher firing rates for right than left choices) or left choices (**Extended Data Fig. 5b-g, 6c**). Upon combining CPI with RMI for each neuron, we found that most right-preferring neurons were enhanced by preceding right *rare*+ (*i.e.*, *rare*+ for stimulus A), whereas most left-preferring neurons were suppressed by preceding right *rare*+ (**Figure 3e, f, Extended Data Fig. 6d**). These trends are both consistent with the right bias induced by right *rare*+. Similarly, preceding left *rare*+ enhanced left-preferring neurons and suppressed right-preferring neurons (**Figure 3d-f, Extended Data Fig. 6d**), leading to a left bias. These results establish a neural basis for how *rare*+ trials influence subsequent decisions (*i.e.*, a neural basis for RIDB): by bi-directionally modulating choice-encoding PPC neurons to induce bias towards *rare*+ choices.

Learning under uncertainty enhances stimulus-encoding of PPC neurons to drive decisions

Despite the persistence of RIDB, most humans and all mice performing this task eventually learned to make decisions based on the current stimulus, manifesting successful learning under uncertainty. Although it has been established that the PPC participates in learning under uncertainty [17, 38], whether PPC activity is required remains unclear. In *stim-opto* mice, in which we inhibited PPC during the stimulus epoch in rare trials, we observed an increasing impairment on performance in PPC-inhibited trials as performance improves (**Figure 4a, b**, orange). The impairment became highly significant when mice reached the expert stage (defined as performance>85%, **Figure 4c, d**, orange). No difference was observed for *outcome-opto* mice (**Figure 4a, c, d**, yellow**)**. For some of the *outcome-opto* mice (n=4), after they reached the expert stage, we switched the inhibition window to the stimulus epoch (*stim-opto acute*), and immediately observed performance impairment in PPC-inhibited trials (**Figure 4a, c, d**). These findings demonstrate that the PPC activity during the stimulus epoch plays a causal role in stimulus-guided decision-making.

To analyse how PPC neurons evolve during learning, we conducted a two-way ANOVA considering stimulus and choice [22-26], on neurons recorded at different performance levels. Whereas the proportion of choice-encoding neurons remained consistently high throughout learning (**Figure 4e**), the proportion of stimulus-encoding neurons varied significantly, marked by a quick rise at 65% performance, and gradual decline after 85% (**Figure 4e**). Stimulus and choice dual-encoding neurons show similar trend as stimulus-encoding neurons (**Figure 4e**). These results suggest that learning dynamically recruits PPC neurons, especially when performance is rising.

The increase followed by a decrease of stimulus-encoding neurons (**Figure 4e**) seems to contradict the monotonically increasing impairment caused by PPC inhibition as performance improves (**Figure 4b**). A possible explanation is that although the number of stimulus-encoding neurons fluctuates, the magnitude of stimulus-encoding increases. To test this, we computed the stimulus discriminability (d') by comparing each neuron's responses to stimulus A *vs*. B [39] (**Figure 4f**), and subsequently calculated the AUC (area under curve) of d' during the stimulus epoch (**Figure 4g**). We found that the d' increased during learning (**Figure 4f, g**). We then used demixed PCA [40] to assess the variance explained by the stimulus at the populational level (**Figure 4h**). This analysis demonstrated an extremely low variance at performance<65%, followed by a continuous increase (**Figure 4h, i**). Thus, as learning progresses, an increasing magnitude of stimulus-encoding by the PPC neurons underlies their causal impact on stimulus-guided decision-making.

A dual-agent model reveals impact of PPC activity on learning under uncertainty

We have demonstrated that two factors contribute to decision-making in this task: a WSLS factor influenced by RIDB, and a stimulus-guided factor independent of RIDB. We have also shown that both decision-making factors rely on PPC activity. However, it is unclear how PPC impacts learning. To conceptualize the functional role of the PPC on learning, we trained a dual-agent model, each agent representing one decision-making factor (Methods, **Figure 5a**). The first agent is a rarity-state-based WSLS agent [41], based on the same-stimulus WSLS tendency (**Figure 1f**). The existence of rarity state throughout learning is inspired by the behavioural and neural data showing the persistent RIDB across performance (**Figure 1k, m, Extended Data Fig. 6a, b**). The WSLS probabilities are updated trial-by-trial, depending on the internal estimate of the rarity state (**Figure 5a, Supp Fig. 7a**) [2, 42, 43], which can effectively track the occurrence of rare trials

(**Supp Fig. 7b, c**). The second agent is a stimulus-guided reinforcement learning (RL) model [44] (**Figure 5a**), which uses a multiplicative update rule [37] and different learning rates for rewarded and unrewarded trials [45]. In light of the data showing persistent RIDB (**Figure 1k, m, Extended Data Fig. 6a, b**) and increasing stimulus-dependency (**Figure 4b, f, g, i**), we used normalized decision weights derived from entropy-based metrics in the RL model to combine the two agents [46] (Methods).

The simulated results successfully recapitulated the RIDB we observed at 20%, 10% and 5% rarity levels (**Figure 5b, Extended Data Fig. 7d**), as well as the relationship between ΔsRP and performance (**Figure 5c**). This model also replicated two phenomena observed in the optogenetics experiments: 1, throughout learning, invalidating the rarity state in the model reduced ΔsRP and ΔM (**Figure 5d Extended Data Fig. 7e**), reproducing the observed reduction of ΔsRP and ΔM by PPC inhibition (**Figure 2e, g**); 2, in the expert stage, invalidating the RL agent impaired performance (**Figure 5e**), reproducing the observed impairment of stimulus-guided decisions (**Figure 4c, d**).

Having recapitulated the decision-making results, we used this model to simulate the learning process. The model predicts faster learning speeds for lower rarity levels, closely resembling the learning speeds we observed for 20%, 10%, and 5% rarity levels in mice (**Figure 5f**). Interestingly, this model also predicts that learning can be accelerated by PPC inhibition during rare trials (**Figure 5g**), which was indeed observed in *stim-opto* mice (**Figure 5g**). Thus, in addition to reproducing RIDB, our model established an essential impact of PPC activity on learning under uncertainty, as validated by optogenetic experiments.

Discussion

To navigate in a world where uncertainty is universal, one needs to make good assessments of the probabilities of upcoming events. Here, using a behavioural paradigm with probabilistic outcomes, we show that in both humans and mice, rarity causes a persistent decision bias. We demonstrate that the PPC is required for generating this decision bias, and that PPC neurons are modulated by preceding rare rewards to induce the decision bias (**Extended Data Fig. 8**). By tracking PPC neuronal activity throughout learning, we observed a dynamic change in stimulus-encoding of PPC neurons, in both the number of neurons and the magnitude of encoding (**Extended Data Fig. 8**). This stimulus-encoding by the PPC is behaviourally relevant, as PPC inhibition increasingly

deteriorates performance as learning progresses. Finally, we built a dual-agent behavioural model which successfully predicted that PPC inhibition during rare trials would accelerate learning, further demonstrating a causal influence of the PPC on decision-making and learning under uncertainty.

The ability to adjust one's behaviour in response to uncertain outcomes is essential for survival, and it is well-established that the subjective assessment of probability in uncertain environments is often biased and results in systematic errors [7, 47]. These errors are caused by many factors such as risk-seeking for rewards, risk-aversion for losses, ease of remembering familiar or extreme events, and emotions [8, 9]. In the present study, we found that in a task where both humans and mice were motivated to obtain maximal rewards, they exhibited a systematic decision bias: preference for rarely rewarded over commonly rewarded choices. As this task excludes influences from risk, reward size, and emotions, our results support that rarity *per se* induces a heuristic that causes humans and mice to inaccurately estimate the reward probability associated with their choices. We also show that this decision bias persists even when the subjects have been extensively exposed to the probabilistic outcomes and are capable of performing the task at an expert level, findings consistent with the notion that "the reliance on heuristics and the prevalence of biases are not restricted to laymen" [7]. Further, this decision bias aligns with commonly reported overweighting of small probabilities in studies related to prospect theory, in which the outcome probabilities are explicit and evaluated consciously by the subjects [8, 9, 11]. Our study has thus uncovered a conserved, intrinsic preference for rare rewards, extending this well-known phenomenon to a condition in which the probability is not obviously revealed, a common real-world scenario [2]. This behavioural bias, though seemingly counterproductive, may reflect an evolutionary advantage of exploration over exploitation [4].

Previous studies have speculated that animals adopt different learning strategies depending on the uncertainty type: fast learning under unexpected uncertainty (changing environment) and slow learning under expected uncertainty (stable environment) [3, 48]. Our dual-agent behavioural model complies with this hypothesis: 1) the fast-updating WSLS agent represents RIDB in decisions, coinciding with the unexpected uncertainty in the early phase of the task, as subjects were not informed of the probabilities *a priori*; 2) the slow-updating RL agent represents stimulus-guided learning, achieved by the subjects after extensive task exposure, consistent with slow learning under expected uncertainty. Our data indicate that the PPC is able to represent both fast-

updating WSLS and slow-updating RL agents, and it is possible that the neuronal populations representing these agents communicate with distinct brain regions and neuromodulatory systems [29, 32, 49]. Future studies are needed to further understand the neural circuit mechanisms.

Many studies in primates and rodents have established that the PPC is involved in decision-making [23-26, 35, 50]. Specifically, the PPC has been linked to decision-making under uncertainty in primates [18-20], but manipulation experiments probing the causal role of the PPC in uncertainty processing remain limited. Our experiments demonstrate two separate roles of the PPC in decision-making under uncertainty: mediating the decision bias towards rare rewards, and the stimulus-guided decisions at the expert stage. In addition, our study filled in the gap in understanding the PPC function in learning under uncertainty, by showing that PPC activity is dynamically modulated by learning, and can influence learning speed.

Author Contributions

W.H.S. and Y.Y. designed the experiments. W.H.S. and X.R.H. performed the experiments. W.H.S, M.C.T and Y.Y. analysed the data. W.H.S. and Y.X. built the model. W.H.S, M.C.T and Y.Y. wrote the manuscript.

Acknowledgements

We thank Drs. Dan-Qian Liu and Ju Lu for comments on the manuscript. We thank the animal facility of ShanghaiTech for their excellent care of mice. This work was supported by grants from the Ministry of Science and Technology of China (2022ZD0204900), Central Guidance on Local Science and Technology Development Fund (YDZX20233100001002), and ShanghaiTech University Startup Fund to Y.Y., National Key R&D Program of China (2023YFA1009402) and of the CAS Project for Young Scientists in Basic Research (YSBR-034) to M.C.T.

References

1. Platt, M.L. and S.A. Huettel, *Risky business: the neuroeconomics of decision making under uncertainty.* Nature Neuroscience, 2008. **11**(4): p. 398-403.
2. Soltani, A. and A. Izquierdo, *Adaptive learning under expected and unexpected uncertainty.* Nature Reviews Neuroscience, 2019. **20**(10): p. 635-644.
3. Monosov, I.E., *How Outcome Uncertainty Mediates Attention, Learning, and Decision-Making.* Trends in Neurosciences, 2020. **43**(10): p. 795-809.
4. Steiner, J. and C. Stewart, *Perceiving Prospects Properly.* American Economic Review, 2016. **106**(7): p. 1601–31.
5. Fiser, J., et al., *Statistically optimal perception and learning: from behavior to neural representations.* Trends in Cognitive Sciences, 2010. **14**(3): p. 119-130.
6. Lee, D., H. Seo, and M.W. Jung, *Neural Basis of Reinforcement Learning and Decision Making.* Annual Review of Neuroscience, Vol 35, 2012. **35**: p. 287-308.
7. Tversky, A. and D. Kahneman, *Judgment under Uncertainty: Heuristics and Biases.* Science, 1974. **185**(4157): p. 1124-1131.
8. Kahneman, D. and A. Tversky, *Prospect Theory: An Analysis of Decision under Risk.* Econometrica, 1979. **47**(2): p. 263-291.
9. Tversky, A. and D. Kahneman, *Advances in prospect theory: Cumulative representation of uncertainty.* Journal of Risk and Uncertainty, 1992. **5**(4): p. 297-323.
10. Hsu, M., et al., *Neural Response to Reward Anticipation under Risk Is Nonlinear in Probabilities.* Journal of Neuroscience, 2009. **29**(7): p. 2231-2237.
11. Elston, T.W., et al., *Extreme Outcomes Accentuate Overweighting of Low Probabilities and Underweighting of High Probabilities in Experience-Based Choice.* Journal of Behavioral Decision Making, 2024. **37**(2): p. e2380.
12. Constantinople, C.M., A.T. Piet, and C.D. Brody, *An Analysis of Decision under Risk in Rats.* Current Biology, 2019. **29**(12): p. 2066-+.
13. Imaizumi, Y., et al., *A neuronal prospect theory model in the brain reward circuitry.* Nature Communications, 2022. **13**(1).
14. Tymula, A., et al., *Dynamic prospect theory: Two core decision theories coexist in the gambling behavior of monkeys and humans.* Science Advances, 2023. **9**(20).
15. Huettel, S.A., A.W. Song, and G. McCarthy, *Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices.* Journal of Neuroscience, 2005. **25**(13): p. 3304-3311.
16. Huettel, S.A., et al., *Neural signatures of economic preferences for risk and ambiguity.* Neuron, 2006. **49**(5): p. 765-775.
17. Bach, D.R., et al., *The Known Unknowns: Neural Representation of Second-Order Uncertainty, and Ambiguity.* Journal of Neuroscience, 2011. **31**(13): p. 4811-4820.
18. Horan, M., N. Daddaoua, and J. Gottlieb, *Parietal neurons encode information sampling based on decision uncertainty.* Nature Neuroscience, 2019. **22**(8): p. 1327-+.
19. Yang, T. and M.N. Shadlen, *Probabilistic reasoning by neurons.* Nature, 2007. **447**(7148): p. 1075-U2.
20. Kiani, R. and M.N. Shadlen, *Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex.* Science, 2009. **324**(5928): p. 759-764.

21. Payzan-LeNestour, E., et al., *The neural representation of unexpected uncertainty during value-based decision making.* Neuron, 2013. **79**(1): p. 191-201.
22. Raposo, D., M.T. Kaufman, and A.K. Churchland, *A category-free neural population supports evolving demands during decision-making.* Nature Neuroscience, 2014. **17**(12): p. 1784-1792.
23. Najafi, F., et al., *Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning.* Neuron, 2020. **105**(1): p. 165-+.
24. Licata, A.M., et al., *Posterior Parietal Cortex Guides Visual Decisions in Rats.* Journal of Neuroscience, 2017. **37**(19): p. 4954-4966.
25. Harvey, C.D., P. Coen, and D.W. Tank, *Choice-specific sequences in parietal cortex during a virtual-navigation decision task.* Nature, 2012. **484**(7392): p. 62-68.
26. Tseng, S.Y., et al., *Shared and specialized coding across posterior cortical areas for dynamic navigation decisions.* Neuron, 2022. **110**(15).
27. Pho, G.N., et al., *Task-dependent representations of stimulus and choice in mouse parietal cortex.* Nature Communications, 2018. **9**.
28. Goard, M.J., et al., *Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions.* Elife, 2016. **5**.
29. Zhong, L., et al., *Causal contributions of parietal cortex to perceptual decision-making during stimulus categorization.* Nature Neuroscience, 2019. **22**(6): p. 963-+.
30. Yao, J.D., et al., *Parietal Cortex Is Required for the Integration of Acoustic Evidence.* Curr Biol, 2020. **30**(17): p. 3293-3303.e4.
31. Hwang, E.J., et al., *History-based action selection bias in posterior parietal cortex.* Nat Commun, 2017. **8**(1): p. 1242.
32. Hwang, E.J., et al., *Corticostriatal Flow of Action Selection Bias.* Neuron, 2019. **104**(6): p. 1126-+.
33. Morcos, A.S. and C.D. Harvey, *History-dependent variability in population dynamics during evidence accumulation in cortex.* Nature Neuroscience, 2016. **19**(12): p. 1672-1681.
34. Akrami, A., et al., *Posterior parietal cortex represents sensory history and mediates its effects on behaviour.* Nature, 2018. **554**(7692): p. 368-+.
35. Najafi, F., et al., *Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning.* Neuron, 2020. **105**(1): p. 165-179.e8.
36. Otazu, G.H., et al., *Engaging in an auditory task suppresses responses in auditory cortex.* Nat Neurosci, 2009. **12**(5): p. 646-54.
37. Bathellier, B., et al., *A multiplicative reinforcement learning model capturing learning dynamics and interindividual variability in mice.* Proceedings of the National Academy of Sciences, 2013. **110**(49): p. 19950-19955.
38. Lyamzin, D. and A. Benucci, *The mouse posterior parietal cortex: Anatomy and functions.* Neuroscience Research, 2019. **140**: p. 14-22.
39. Constantinople, C.M., et al., *Lateral orbitofrontal cortex promotes trial-by-trial learning of risky, but not spatial, biases.* Elife, 2019. **8**.
40. Kobak, D., et al., *Demixed principal component analysis of neural population data.* eLife, 2016. **5**: p. e10989.

41. Worthy, D.A. and W. Todd Maddox, *A comparison model of reinforcement-learning and win-stay-lose-shift decision-making processes: A tribute to W.K. Estes.* Journal of Mathematical Psychology, 2014. **59**: p. 41-49.
42. Jang, A.I., et al., *Positive reward prediction errors during decision-making strengthen memory encoding.* Nature Human Behaviour, 2019. **3**(7): p. 719-732.
43. Rouhani, N. and Y. Niv, *Signed and unsigned reward prediction errors dynamically enhance learning and memory.* eLife, 2021. **10**: p. e61077.
44. Sutton, R.S. and A.G. Barto, *Reinforcement learning: an introduction*. Second edition ed. Adaptive computation and machine learning series. 2018, Cambridge, Massachusetts: The MIT Press. 526.
45. Fusi, S., et al., *A Neural Circuit Model of Flexible Sensorimotor Mapping: Learning and Forgetting on Multiple Timescales.* Neuron, 2007. **54**(2): p. 319-333.
46. Trepka, E., et al., *Entropy-based metrics for predicting choice behavior based on local response to reward.* Nature Communications, 2021. **12**(1): p. 6567.
47. Tversky, A. and D. Kahneman, *Availability: A heuristic for judging frequency and probability.* Cognitive Psychology, 1973. **5**(2): p. 207-232.
48. Behrens, T.E.J., et al., *Learning the value of information in an uncertain world.* Nature Neuroscience, 2007. **10**(9): p. 1214-1221.
49. Yu, A.J. and P. Dayan, *Uncertainty, neuromodulation, and attention.* Neuron, 2005. **46**(4): p. 681-692.
50. Zhou, Y. and D.J. Freedman, *Posterior parietal cortex plays a causal role in perceptual and categorical decisions.* Science, 2019. **365**(6449): p. 180-185.

Figure 1

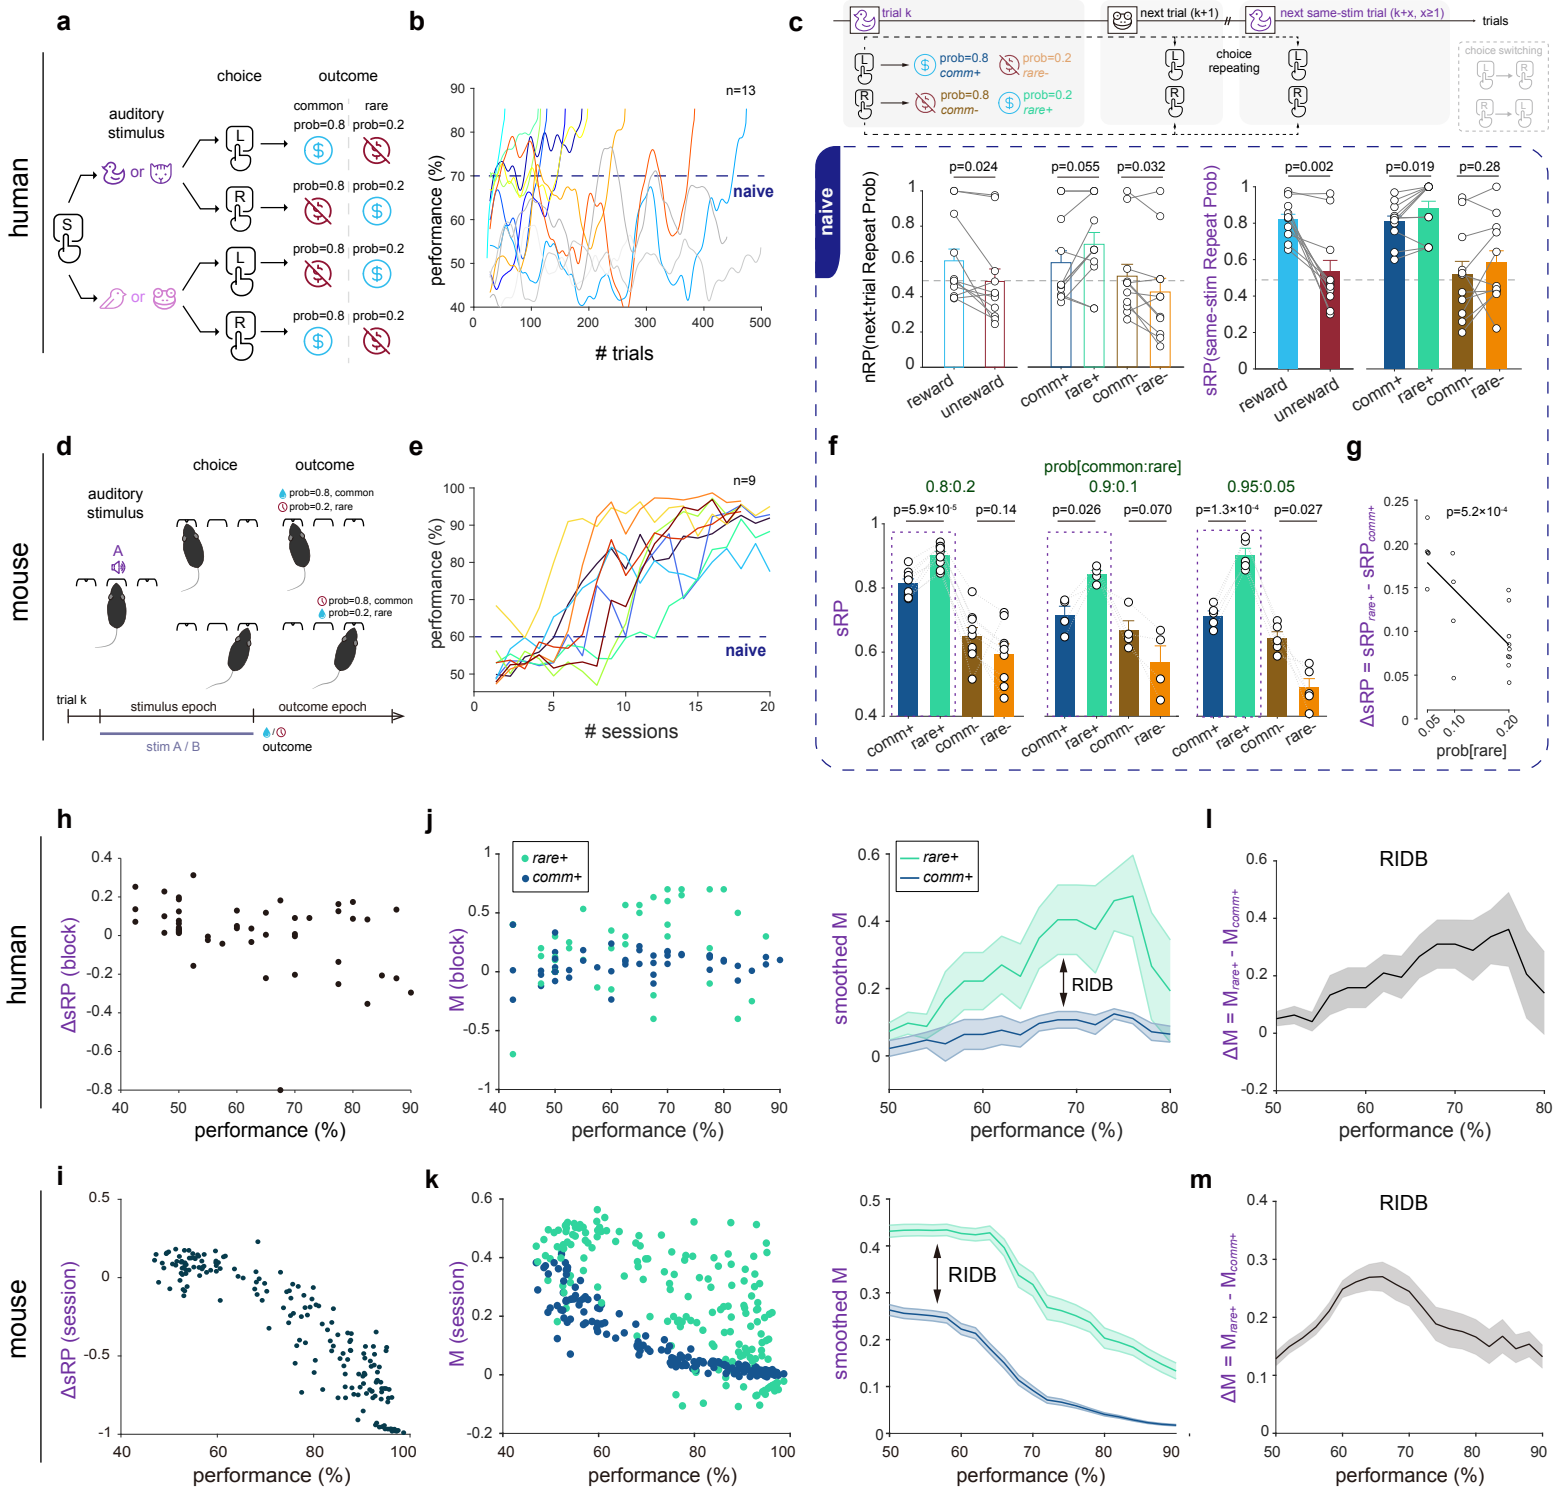
a
b
c
human
auditory stimulus
choice
outcome
common
rare
prob=0.8
prob=0.2
n=13
naive
performance (%)
trials
trial k
next trial (k+1)
next same-stim trial (k+x, x≥1)
trials
choice repeating
choice switching
prob=0.8 comm+
prob=0.2 rare-
prob=0.8 comm-
prob=0.2 rare+
nRP(next-trial Repeat Prob)
sRP(same-stim Repeat Prob)
p=0.024
p=0.055
p=0.032
p=0.002
p=0.019
p=0.28
reward
unreward
comm+
rare+
comm-
rare-
d
e
f
g
mouse
A
prob=0.8, common
prob=0.2, rare
trial k
stimulus epoch
outcome epoch
stim A / B
outcome
n=9
# sessions
prob[common:rare]
0.8:0.2
0.9:0.1
0.95:0.05
p=5.9×10⁻⁵
p=0.14
p=0.026
p=0.070
p=1.3×10⁻⁴
p=0.027
sRP
p=5.2×10⁻⁴
ΔsRP = sRP_rare+ - sRP_comm+
prob[rare]
h
j
l
i
k
m
ΔsRP (block)
M (block)
rare+
comm+
smoothed M
RIDB
ΔM = M_rare+ - M_comm+
ΔsRP (session)
M (session)

Figure 1. Rare rewards induce a persistent decision bias in humans and mice

(a) Schematic showing an auditory decision-making task with probabilistic outcomes for humans. Pressing "S" key triggers the delivery of 1 of 4 animal sounds (duck, cat, canary or frog). Upon hearing duck or cat sounds, pressing "L" key leads to money reward at 80% probability, and no money at 20% probability; pressing "R" key leads to money reward at 20% probability, and no money at 80% probability. The choice-reward contingency is the opposite for canary and frog sounds.

(b) Performance of all subjects, calculated as a 20-trial running average of the proportion of correct choices. Each line represents one subject. Naïve stage is defined as performance < 70%.

(c) Top, schematic showing choice repeating for the immediately next trial (stimulus may be the same or different) and the next same-stimulus trial after each possible outcome: common reward (*comm*+), rare no-reward (*rare*-), common no-reward (*comm*-), and rare reward (*rare*+). Bottom left panel, empty bars showing choice repeating probabilities of rewarded & unrewarded; *comm*+, *rare*+, *comm*-, and *rare*- trials in the immediately next trial (nRP) in the naïve stage. Bottom right panel, filled bars showing choice repeating probabilities of rewarded & unrewarded; *comm*+, *rare*+, *comm*-, and *rare*- trials in the next same-stimulus trial (sRP) in the naïve stage. Each data point represents one subject. Error bar, s.e.m. n=11. Wilcoxon signed-rank test.

(d) Schematic showing an auditory 2-alternative-forced choice (2AFC) task with probabilistic outcomes for mice. Poking in the centre port triggers stimulus A (5kHz pure tone) or B (13kHz pure tone). For stimulus A, left choice is associated with water reward at 80% probability and timeout punishment at 20%; right choice is associated with timeout at 80% probability and water at 20%. The choice-reward contingency is the opposite for stimulus B.

(e) Performance of all mice across sessions, calculated as the average proportion of correct choices in each session. Each line represents one mouse. Naïve stage is defined as performance < 60%.

(f) Choice repeating probabilities of *comm*+, *rare*+, *comm*-, or *rare*- trials in the next same-stimulus trial (sRP) for 3 rarity levels: 20%, 10%, and 5% rare trials, in the naïve stage. Each data point represents one mouse. Error bar, s.e.m. n=9. Wilcoxon signed-rank test.

(g) Difference in sRP for *rare*+ vs. *comm*+ trials (ΔsRP) for 3 rarity levels in (f). One-way ANOVA.

(h) ΔsRP across performance for human subjects. Each behavioural session is divided into 40-trial blocks *post hoc* for calculating ΔsRP and performance. Each data point represents one block. n=64 blocks.

(i) ΔsRP across performance for mice. Each data point represents one session. n=183 sessions.

(j) Left, M_{rare+} and M_{comm+} for 40-trial blocks of all human subjects, plotted against performance. Right, smoothed M_{rare+} and M_{comm+} across performance. Error bar, s.e.m.

(k) Left, M_{rare+} and M_{comm+} for all sessions of all mice, plotted against performance. Right, smoothed M_{rare+} and M_{comm+} across performance. Error bar, s.e.m.

(l) ΔM representing RIDB across performance for human subjects. Error bar, s.e.m.

(m) ΔM representing RIDB across performance for mice. Error bar, s.e.m.

Figure 2

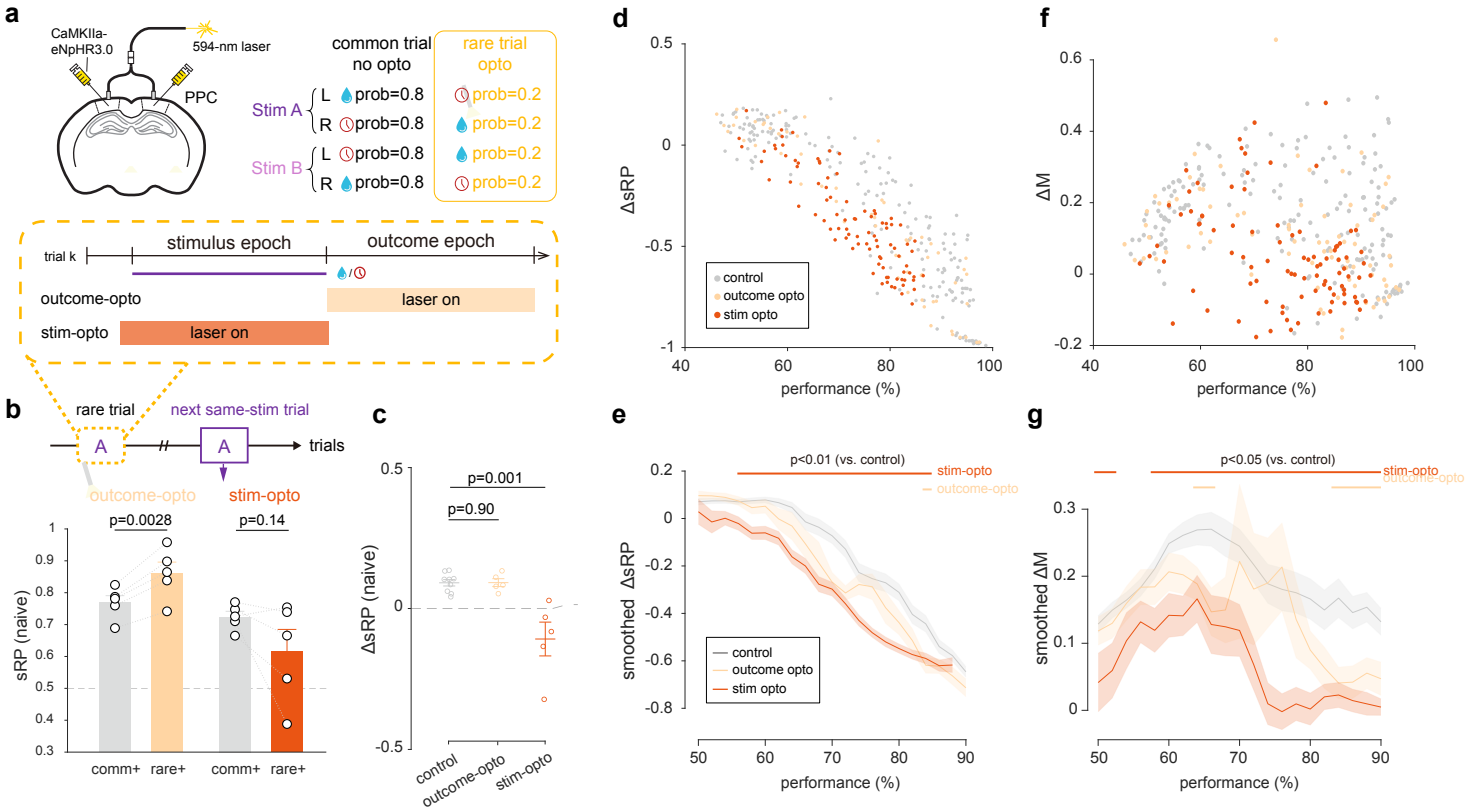
a
CaMKIIa-eNpHR3.0
594-nm laser
PPC
common trial no opto
rare trial opto
Stim A
L prob=0.8
R prob=0.8
prob=0.2
prob=0.2
Stim B
L prob=0.8
R prob=0.8
prob=0.2
prob=0.2
trial k
stimulus epoch
outcome epoch
outcome-opto
laser on
stim-opto
laser on
b
rare trial
next same-stim trial
A
A
trials
outcome-opto
stim-opto
p=0.0028
p=0.14
sRP (naive)
comm+
rare+
comm+
rare+
c
p=0.001
p=0.90
ΔsRP (naive)
control
outcome-opto
stim-opto
d
ΔsRP
control
outcome opto
stim opto
performance (%)
e
p<0.01 (vs. control)
stim-opto
outcome-opto
smoothed ΔsRP
control
outcome opto
stim opto
performance (%)
f
ΔM
performance (%)
g
p<0.05 (vs. control)
stim-opto
outcome-opto
smoothed ΔM
performance (%)

Figure 2. Optogenetic inhibition of the PPC reduces rarity-induced decision bias (RIDB)

(a) Schematic showing optogenetic experiment. Optogenetic inhibition of the PPC was performed either during stimulus epoch (*stim-opto*) or outcome epoch (*outcome-opto*) in rare trials.

(b) Bar graph showing sRP_{rare+} and sRP_{comm+} for *stim-opto* (left, n=5) and *outcome-opto* (right, n=5) mice, in the naïve stage. Each data point represents one mouse. Error bar, s.e.m. Paired t-test.

(c) ΔsRP of *non-opto* (control)*, outcome-opto* and *stim-opto* mice in the naïve stage. Error bar, s.e.m. Wilcoxon ranksum test.

(d) ΔsRP across performance for *non-opto* (control)*, stim-opto,* and *outcome-opto* mice. Each data point represents one session. *Non-opto*, n=183; *stim-opto*, n=61; *outcome-opto*, n=91.

(e) Smoothed ΔsRP across performance for *non-opto* (control)*, stim-opto,* and *non-opto* mice. Wilcoxon ranksum test, between control and *stim-opto,* and between control and *outcome-opto.* Error bar, s.e.m.

(f) ΔM across performance for *non-opto* (control)*, stim-opto,* and *outcome-opto* mice. Each data point represents one session.

(g) Smoothed ΔM across performance for *non-opto* (control)*, stim-opto,* and *non-opto* mice. Wilcoxon ranksum test, between control and *stim-opto,* and between control and *outcome-opto.* Error bar, s.e.m.

Figure 3

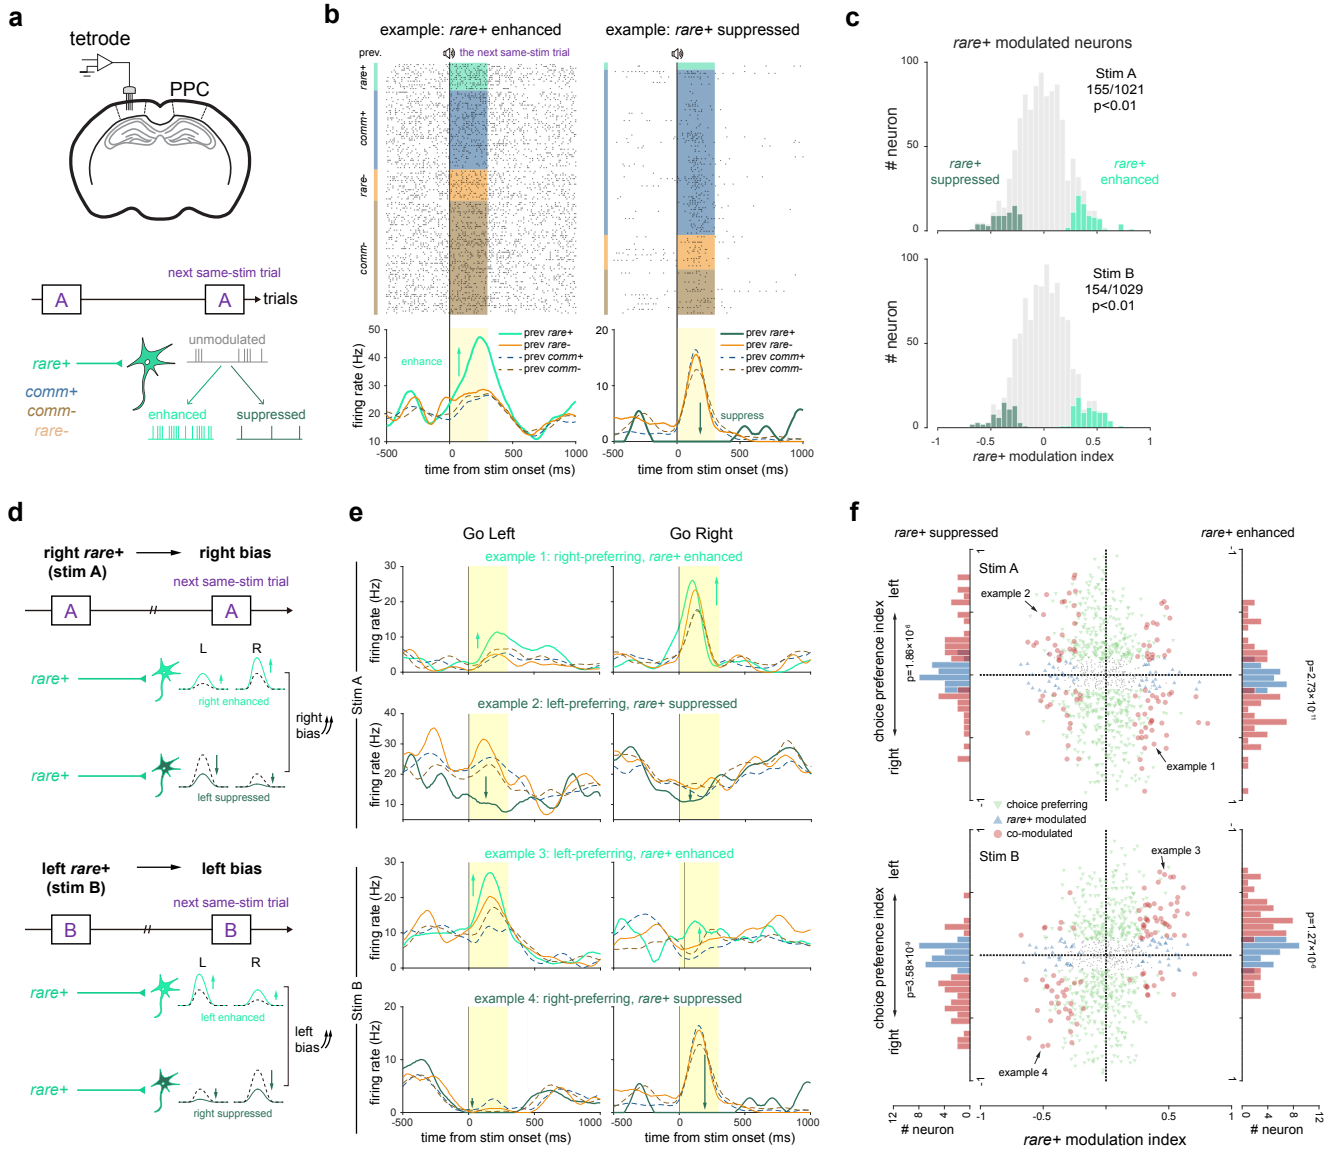
a
tetrode
PPC
next same-stim trial
A
A
trials
rare+
comm+
comm-
rare-
unmodulated
enhanced
suppressed
b
example: rare+ enhanced
example: rare+ suppressed
prev.
the next same-stim trial
rare+
comm+
rare-
comm-
enhance
suppress
prev rare+
prev rare-
prev comm+
prev comm-
firing rate (Hz)
time from stim onset (ms)
c
rare+ modulated neurons
Stim A
155/1021
p<0.01
Stim B
154/1029
p<0.01
neuron
rare+ suppressed
rare+ enhanced
rare+ modulation index
d
right rare+ (stim A)
right bias
next same-stim trial
A
A
L
R
rare+
right enhanced
right bias
rare+
left suppressed
left rare+ (stim B)
left bias
next same-stim trial
B
B
L
R
rare+
left enhanced
left bias
rare+
right suppressed
e
Go Left
Go Right
example 1: right-preferring, rare+ enhanced
example 2: left-preferring, rare+ suppressed
example 3: left-preferring, rare+ enhanced
example 4: right-preferring, rare+ suppressed
Stim A
Stim B
firing rate (Hz)
time from stim onset (ms)
f
rare+ suppressed
rare+ enhanced
Stim A
Stim B
example 1
example 2
example 3
example 4
choice preference index
left
right
choice preferring
rare+ modulated
co-modulated
neuron
rare+ modulation index

Figure 3. Rare rewards modulate choice-selective PPC neurons to induce decision bias

(a) Top panel, schematic showing tetrode recording in PPC. Bottom panel, schematic showing modulation (enhancement or suppression) by a preceding *rare*+ trial on PPC neuronal activity in the next same-stimulus trial.

(b) Example *rare*+ modulated neurons. Top panel, raster plot showing spikes of a PPC neuron in next same-stimulus trials of *comm*+, *rare*+, *comm*-, and *rare*-. The left example neuron fires more spikes after *rare*+, and the right example neuron fires fewer spikes after *rare*+. Bottom panel, averaged firing rates of the example neurons above, showing enhanced (left example) or suppressed (right example) firing after *rare*+.

(c) *Rare*+ modulation index (RMI) for PPC neurons during the stimulus epoch of next same-stimulus trials, for stimulus A (top panel) and stimulus B (bottom panel). RMI is calculated using ROC analysis (see Methods). Colour bars represent significant modulation. Statistical significance was assessed via bootstrapping (200 iterations, $p<0.01$).

(d) Top panel, schematic showing a preceding right *rare*+ trial (stimulus A *rare*+) enhances the firing of a right-preferred PPC neuron, and suppresses the firing of a left-preferred PPC neuron in the next same-stimulus trial. Bottom panel, schematic showing a preceding left *rare*+ trial (stimulus B) enhances a left-preferred PPC neuron and suppresses a right-preferred PPC neuron in the next same-stimulus trial.

(e) Example *rare*+ modulated neurons. Top panel, a right-preferred neuron (example 1) is enhanced by preceding right *rare*+; a left-preferred neuron (example 2) is suppressed by preceding right *rare*+. Bottom panel, a left-preferred neuron (example 3) is enhanced by preceding left *rare*+; a right-preferred neuron (example 4) is suppressed by preceding left *rare*+.

(f) RMI plotted against choice preference index (CPI) for all neurons. CPI is calculated using ROC analysis. Top panel, modulation by right *rare*+; bottom panel, modulation by left *rare*+. Each data point represents one neuron. Red, significant choice preference and *rare*+ modulation. Green, significant choice preference. Blue, significant *rare*+ modulation. Significance was assessed via bootstrapping (200 iterations, $p<0.01$). Histograms show CPI for *rare*+ enhanced (right bar graphs) and *rare*+ suppressed (left bar graphs) neurons. Red bars represent significant *rare*+ modulation. Kolmogorov-Smirnov test.

Figure 4

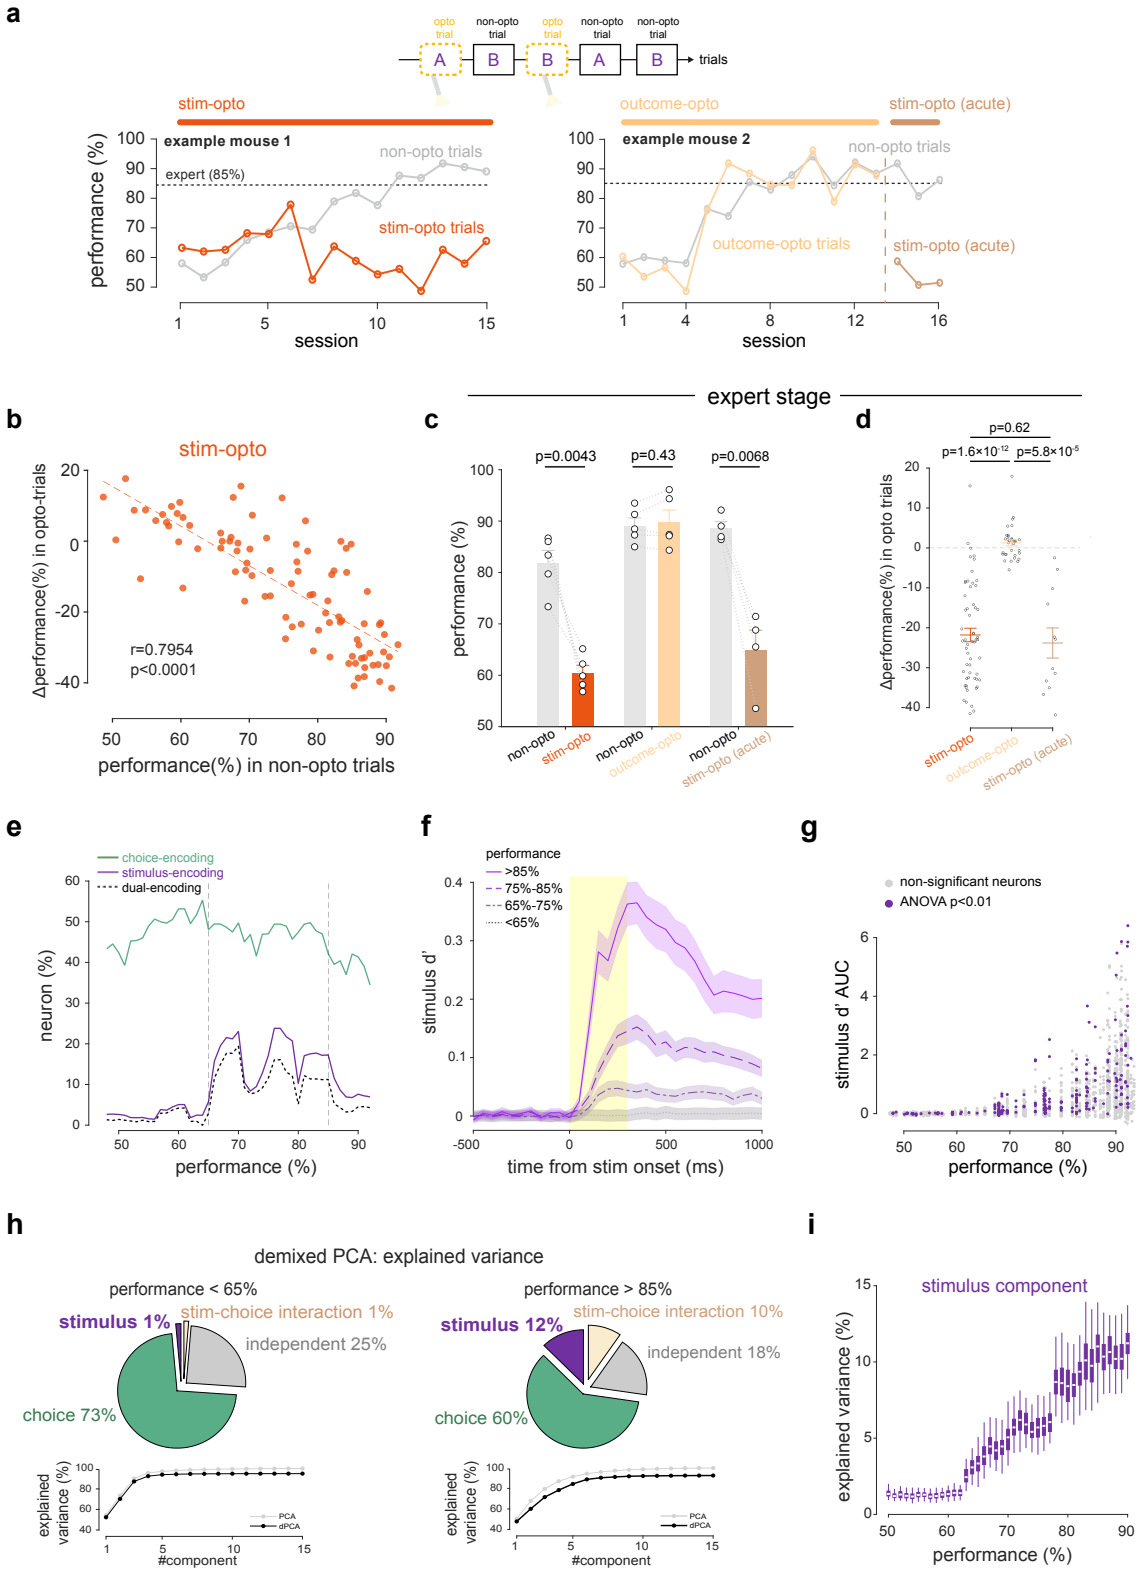
a
opto trial
non-opto trial
A
B
trials
stim-opto
outcome-opto
stim-opto (acute)
example mouse 1
example mouse 2
expert (85%)
non-opto trials
stim-opto trials
outcome-opto trials
performance (%)
session
b
stim-opto
Δperformance(%) in opto-trials
r=0.7954
p<0.0001
performance(%) in non-opto trials
c
expert stage
p=0.0043
p=0.43
p=0.0068
performance (%)
non-opto
stim-opto
outcome-opto
stim-opto (acute)
d
p=0.62
p=1.6×10-12
p=5.8×10-5
Δperformance(%) in opto trials
e
choice-encoding
stimulus-encoding
dual-encoding
neuron (%)
performance (%)
f
performance
>85%
75%-85%
65%-75%
<65%
stimulus d'
time from stim onset (ms)
g
non-significant neurons
ANOVA p<0.01
stimulus d' AUC
performance (%)
h
demixed PCA: explained variance
performance < 65%
stimulus 1%
stim-choice interaction 1%
independent 25%
choice 73%
performance > 85%
stimulus 12%
stim-choice interaction 10%
independent 18%
choice 60%
explained variance (%)
PCA
dPCA
#component
i
stimulus component
explained variance (%)
performance (%)

Figure 4. PPC inhibition impairs stimulus-guided decision-making and learning enhances stimulus-encoding of PPC neurons

(a) Performance across sessions of an example *stim-opto* (left) and *outcome-opto* (right) mouse, showing averaged performance of non-inhibited (non-opto) and inhibited (opto) trials. The *outcome-opto* mice was switched to *stim-opto* protocol for 3 sessions after reaching expert performance level (>85%). These sessions were termed "*stim-opto (acute)*".

(b) The reduction in performance (Δperformance) of *stim-opto* trials compared to non-opto trials, plotted against the performance of non-opto trials in the same session. Each data point represents one session. n=91.

(c) Performance of non-opto and opto trials for *stim-opto* (n=5), *outcome-opto* (n=5), and *stim-opto (acute)* mice (n=4) after mice reached expert stage. Each data point represents one mouse. Error bar, s.e.m. Wilcoxon signed-rank test.

(d) Δperformance of *stim-opto* trials in *stim-opto* (n=5), *outcome-opto* (n=5), and *stim-opto (acute)* (n=4) mice. Each data point represents one session. *stim-opto*, n=57 sessions; *outcome-opto*, n=33; *stim-opto (acute)*, n=12. Error bar, s.e.m. Wilcoxon ranksum test.

(e) Fraction of significant choice-encoding (green), stimulus-encoding (purple), and dual-encoding (black dotted) neurons across performance. Significance is determined using two-way ANOVA ($p<0.01$).

(f) Mean stimulus discriminability (d') for stimulus-encoding neurons recorded in sessions during which performance<65%, 65%<performance<75%, 75%<performance<85%, and performance>85%. Error bar, s.e.m.

(g) AUC (area under curve) of stimulus d' during the stimulus epoch of stimulus-encoding neurons. Each data point indicates one neuron. Purple points indicate significant stimulus-encoding neurons. Significance is determined using two-way ANOVA ($p<0.01$).

(h) Demixed PCA (dPCA) is performed on neurons recorded in sessions in which performance <65% (left), and performance>85%. Pie graph shows explained variance for stimulus (purple), choice (green), stimulus-choice interaction (orange), and independent (grey) components.

(i) Explained variance by stimulus across performance. Each box represents the explained variance calculated using dPCA on 200 neurons, bootstrapped 100 times from sessions within the range of [performance -5%, performance +5%]. The boxes show 25th and 75th percentiles.

Figure 5

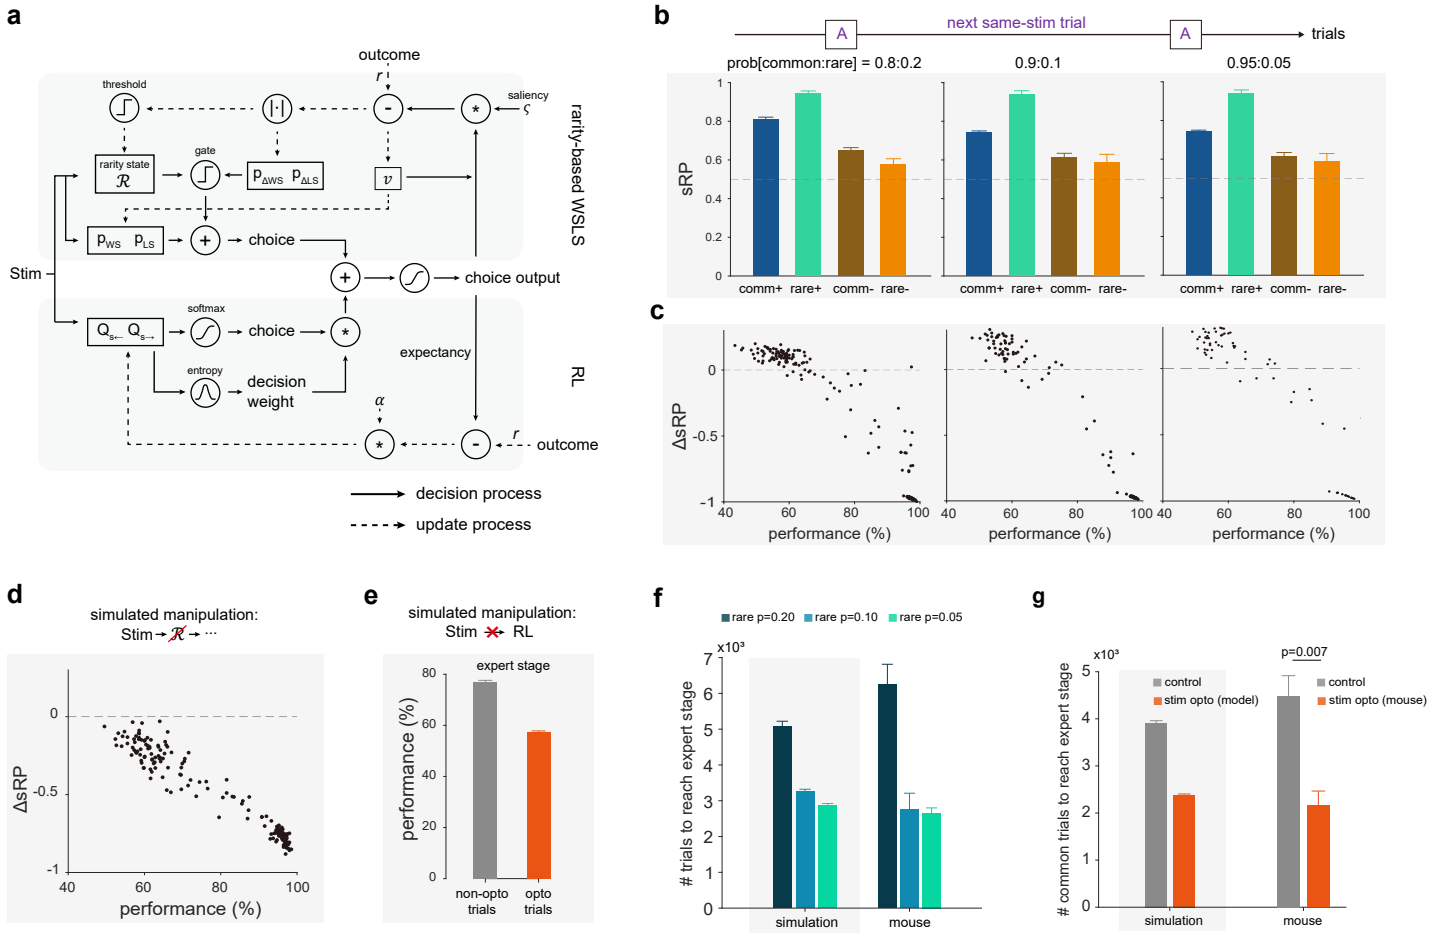
a
outcome
threshold
saliency
rarity-based WSLS
rarity state
gate
choice
Stim
choice output
softmax
choice
expectancy
RL
entropy
decision weight
outcome
decision process
update process
b
next same-stim trial
trials
prob[common:rare] = 0.8:0.2
0.9:0.1
0.95:0.05
sRP
comm+ rare+ comm- rare-
c
ΔsRP
performance (%)
d
simulated manipulation:
Stim → ℛ → ···
ΔsRP
performance (%)
e
simulated manipulation:
Stim ✗ RL
expert stage
performance (%)
non-opto trials
opto trials
f
rare p=0.20
rare p=0.10
rare p=0.05
x10³
trials to reach expert stage
simulation
mouse
g
x10³
p=0.007
control
stim opto (model)
control
stim opto (mouse)
common trials to reach expert stage
simulation
mouse

Figure 5. A dual-agent model shows the impact of PPC on decision-making and learning under uncertainty

(a) Schematic showing the dual-agent behavioural model, consisting of a rarity-based WSLS agent (top) and an RL agent (bottom). See Methods for details.

(b) Simulated repeat probability for next same-stimulus trials (sRP) at 3 rarity levels. Error bar, s.e.m. n=10 iterations.

(c) Simulated ΔsRP across performance for 3 rarity levels. 10 iterations, 20%: n=159 pseudo-sessions (each pseudo-session contains 500 trials); 10%: n=112; 5%: n=75.

(d) Simulated ΔsRP across performance at 20% rarity level, upon elimination of rarity state in the WSLS agent of the model. 10 iterations, n=159 pseudo-sessions.

(e) Simulated performance impairment by inactivation of the RL agent at expert stage. Error bar, s.e.m. n=10 iterations.

(f) Trials to reach expert stage at 20%, 10% and 5% rarity levels, for the model (left) and mouse (right). Error bar, s.e.m.

(j) Trials to reach expert stage at 20% rarity level upon simulated inactivation of the PPC in the model (left), reproduced by optogenetic PPC inhibition in mice (right). Error bar, s.e.m. Wilcoxon ranksum test.