


Temporal Dynamics Underlying Prelimbic Prefrontal Cortical Regulation of Action Selection and Outcome Evaluation during Risk/Reward Decision-Making

Debra A. Bercovici, Oren Prinz-Lebel, Jackson D. Schumacher, Valerie M. Lo, and  Stan B. Floresco

Department of Psychology and Djavad Mowafaghian Centre for Brain Health, University of British Columbia, Vancouver, British Columbia V6T1Z4, Canada

Risk/reward decision-making is a dynamic process that includes periods of deliberation before action selection and evaluation of the action outcomes that bias subsequent choices. Inactivation of the prelimbic (PL) cortex has revealed its integral role in updating decision biases in the face of changes in probabilistic reward contingencies, yet how phasic PL signals during different phases of the decision process influence choice remains unclear. We used temporally specific optogenetic inhibition to selectively disrupt PL activity coinciding with action selection and outcome phases to examine how these signals influence choice. Male rats expressing the inhibitory opsin eArchT within PL excitatory neurons were well trained on a probabilistic discounting task, entailing choice between small/certain versus large/risky rewards, the probability of which varied over a session (50–12.5%). During testing, brief light pulses suppressed PL activity before choice or after different outcomes. Prechoice suppression reduced bias toward more preferred/higher utility options and disrupted how recent outcomes influenced subsequent choice. Inhibition during risky losses induced a similar profile, but here, the impact of reward omissions were either amplified or diminished, relative to the context of the estimated profitability of the risky option. Inhibition during large or small reward receipt reduced risky choice when this option was more profitable, suggesting these signals can both reinforce rewarded risky choices and also act as a relative value comparator signal that augments incentive for larger rewards. These findings reveal multifaceted contributions by the PL in implementing decisions and integrating action–outcome feedback to assign context to the decision space.

Key words: anterior cingulate; choice; decision-making; optogenetics; prediction error

Significance Statement

The PL prefrontal cortex plays an integral role in guiding risk/reward decisions, but how activity in this region during different phases of the decision process influences choice is unclear. By using temporally specific optogenetic manipulations of this activity, the present study unveiled previously uncharacterized and differential contributions by PL in implementing decision policies and how evaluation of decision outcomes shape subsequent choice. These findings provide novel insight into the dynamic processes engaged by the PL that underlie action selection in situations involving reward uncertainty that may aid in understanding the mechanism underlying normal and aberrant decision-making processes.

Introduction

Value-based decisions involving reward uncertainty are guided by distributed neural circuits linking different regions of the

striatum, temporal, and frontal lobe. Among the subdivisions of the prefrontal cortex (PFC) known to refine action selection, the anterior cingulate cortex (ACC) has been heavily implicated in biasing risk/reward decision-making. Lesion studies in humans have implicated the ACC in promoting optimal decisions on a variety of tests involving choices between options that may or may not yield different rewards (Clark et al., 2008; Camille et al., 2011; Gläscher et al., 2012; Pujara et al., 2015). Imaging studies have further revealed distinct profiles of ACC activation during choice or outcome evaluation phases of the decision-making process (Rogers et al., 2004; Kuhnen and Knutson, 2005; Christopoulos et al., 2009; Gläscher et al., 2009; Kolling et al., 2014, 2016). These findings imply that neural activity within the

Received Apr. 24, 2022; revised Dec. 23, 2022; accepted Dec. 30, 2022.

Author contributions: D.A.B. and S.B.F. designed research; D.A.B., O.P.-L., and V.M.L. performed research; D.A.B., J.D.S., and S.B.F. analyzed data; and D.A.B., J.D.S., and S.B.F. wrote the paper.

This work was supported by Canadian Institutes of Health Research Grant PJT-162444 to S.B.F. and a Natural Sciences and Engineering Research Council of Canada Fellowship to D.A.B. We thank Jin H. Wen and Patrick Klaiber for insight and advice on the multilevel modeling analysis.

The authors declare no competing financial interests.

Correspondence should be addressed to Stan B. Floresco at floresco@psych.ubc.ca.

<https://doi.org/10.1523/JNEUROSCI.0802-22.2022>

Copyright © 2023 the authors

ACC during different phases of the decision process may aid in tracking changes in the long-term value of different courses of action. This promotes updating of choice behavior as a decision maker progresses through a sequence of decisions to update choice behavior accordingly (Kolling et al., 2014).

The prelimbic (PL) region of the rodent medial PFC (mPFC) shares similar anatomic connectivity to Area 32 of the primate ACC (van Eden et al., 1992; Heilbronner et al., 2016), and this region has been implicated in guiding risk/reward decisions using a variety of different assays (St Onge and Floresco, 2010; de Visser et al., 2011; Paine et al., 2015; Zeeb et al., 2015; Orsini et al., 2018). Notably, the manner in which disruption of PL function alters these decisions can vary based on the specific procedures used. For example, lesions/inactivation of the PL cause a disadvantageous increase in risky choice using assays patterned after the human Iowa gambling task, where rats choose between different options with fixed probabilities of rewards and time-out punishments (de Visser et al., 2011; Paine et al., 2015; Zeeb et al., 2015). Conversely, when decisions are guided by external cues, PL inactivations reduce risky choices when they would be potentially more profitable (van Holstein and Floresco, 2020). On the other hand, previous work by our group has used a probabilistic discounting task entailing choices between small/certain versus large/risky rewards, the probability of which changes over the course of a session. This task requires animals to track choice-outcome history to enable more profitable decisions. Inactivation of the PL increased or decreased risky choice depending on whether reward probabilities decreased or increased over a session, indicating that this region keeps track of changes in the profitability of different actions to facilitate flexible shifts in decision biases (St Onge and Floresco, 2010). Subsequent findings revealed that different networks of PL neurons interfacing with the basolateral amygdala or nucleus accumbens play distinct roles in refining choice. The PL→amygdala pathway tracks changes in rewarded and nonrewarded choices, whereas the PL→accumbens circuit reinforces risky wins (Jenni et al., 2017; St Onge et al., 2012).

Like the human ACC, the rodent PL displays distinct changes in phasic activity that are temporally linked to deliberation (pre-choice) or outcome evaluation phases of the decision-making process. For example, neurophysiological and fiber photometry recordings during risk/reward decision-making revealed some PL cells show phasic increases in activity before action selection that appear to represent preferred choices (Braunscheidel et al., 2019; Sackett et al., 2019). Other cells show brief changes in firing following either nonrewarded or rewarded choices that track outcomes over multiple trials and predict subsequent choice or shifts in action selection (Sul et al., 2010; Del Arco et al., 2017; Braunscheidel et al., 2019; Passecker et al., 2019). Yet, conventional lesion/inactivations abolish all these temporally discrete neural signals associated with different phases of the decision process. What remains to be clarified is how temporally precise changes in PL activity that occur specifically during deliberation phases before action selection, or during evaluation of different choice outcomes, may bias current choices and influence subsequent ones, as few studies have addressed this directly (Passecker et al., 2019). Here, we used temporally discrete optogenetic suppression of activity in PL neurons expressing the inhibitory opsin eArchT to clarify how their phasic activity during different phases of the decision-making sequence shapes choice during probabilistic discounting.

Materials and Method

Subjects

Forty-six male Long–Evans rats (Charles River Laboratories) were used across the different experiments. These animals weighed ~250–275 g on arrival, were group housed, provided food *ad libitum*, and handled daily for 1 week. After acclimatization, rats underwent stereotaxic surgery to infuse virus and implant optic fiber ferrules into the PL. Following surgery, animals were single housed for the remainder of the experiment. Before beginning behavioral training, rats were food restricted to ~85% of their free-feeding weight. Their weights were monitored daily, and food was adjusted to maintain a weight gain of ~5 g per week. All procedures were conducted in accordance with the Canadian Council of Animal Care and the Animal Care Committee at the University of British Columbia.

Stereotaxic surgery

Rats were given a subanesthetic intraperitoneal dose of a ketamine (50 mg/kg) and xylazine (5 mg/kg) cocktail for initial sedation and analgesia and were maintained on isoflurane for the full procedure. They were placed into a stereotaxic frame secured with ear bars (flat skull), and analgesia was administered subcutaneously (Anafen, 10 mg/kg). Burr holes were drilled into the skull and a 0.6 μ l solution containing a virus encoding for an inhibitory opsin (rAAV5-CaMKII α -eArchT3.0-eYFP; concentration, 5×10^{13} particles/ml) or a control vector (rAAV5-CaMKII α -eYFP, University of North Carolina Vector Core) was infused bilaterally at a 10° angle into the PL via microinfusion pumps (coordinates from bregma, +3.4 mm anteroposterior; \pm 1.6 mm mediolateral; –3.5 mm dorsoventral from dura) at a flow rate of 0.1 μ l per minute. Injectors were left in place for 10 min following the infusion to allow for virus diffusion in tissue. Subsequently, optic fibers consisting of 400 μ m cores and 0.50 NA (Thorlabs) threaded through 2.5-mm-wide metal ferrules (Precision Fiber Products) were implanted into the PL at the above coordinates. A head assembly secured the fibers in place using six screws and dental cement. Animals received postoperative treatment and monitoring for 5 d following surgery ahead of beginning food restriction and behavioral training.

Apparatus

Behavioral testing was conducted in operant chambers (30.5 \times 24 \times 21 cm; Med Associates) enclosed in sound-attenuating boxes. Each box was equipped with a fan that provided ventilation and limited extraneous sounds. The chamber was fitted with a central food receptacle where sucrose food reward pellets (45 mg; Bio-Serv) were dispensed. Two retractable levers were located on either side of the food receptacle. The chamber was illuminated by a 100 mA house light located on the top center of the box opposite the food receptacle that delivered ~25 mW (~5 lux) illuminance in the area around the levers. All data were recorded by a personal computer connected to the operant chambers via an interface. Lasers were controlled by Med PC software, which delivered a transistor-transistor logic (TTL^{+/-}) pulse to lasers to initiate/terminate light delivery.

Lever press training

The initial training protocols described below were identical to those described in previous studies conducted by our group (Stopper et al., 2014; Bercovici et al., 2018). Rats were food restricted for 3 d. One day before beginning operant training, rats

were given ~30 reward pellets in their cage. The first day of training began with two pellets placed in the food receptacle, with either the right or left lever extended, and crushed sugar pellets sprinkled on the extended lever. Animals were trained to lever press for pellets under a fixed ratio-1 schedule until a criterion of 60 presses in 30 min was met for both levers. During the next phase of training, rats were trained on a simplified version of the full task. This consisted of 90 trials where rats were presented with one of the levers, which, if pressed within 10 s, would deliver one pellet with a 50% probability. If the lever was not pressed within this time, it was retracted, and the trial was scored as an omission. Trials occurred every 40 s. Rats trained for ~4 d until a criterion of <10 omissions for a minimum of 2 consecutive days. Next, rats learned to choose between one lever associated with a larger, four-pellet reward (delivered with a 50% probability) and another lever that always delivered a one-pellet reward. Assignment of the large-reward lever was counterbalanced across animals. Sessions consisted of 72 trials portioned into four blocks of 18 trials. The first 8 trials of each block were forced choice, where only one lever was inserted (randomized in pairs), and the remaining 10 trials were free choice where both levers were inserted. Rats were trained until they chose the large lever on >60% of the free choices (~3 d).

Probabilistic discounting training

This task was adapted from Bercovici et al. (2018) in which we previously used optogenetic suppression of neural activity during discrete phases of the probabilistic discounting task and is diagrammed in Figure 1A. Animals were trained 5–7 d per week. During the 40 min session, each ferrule was connected to a fiber optic patch cable encased in stainless steel spring coils that were tethered to a rotary joint that permitted free movement through the chamber. Each session consisted of 60 trials separated into two blocks of 30 (10 forced-choice followed by 20 free-choice trials). One lever was designated the small/certain lever and the other designated the large/risky lever, which were the same as in the last phase of pretraining. Every 40 s, a trial began with the illumination of the house light. Four seconds later, one (forced choice) or both (free choice) levers were inserted into the chamber. Rats were given 10 s to press a lever, otherwise the lever(s) was/were retracted, and the trial scored as an omission. Selection of either lever caused both to retract. Choice of the small/certain option always delivered one pellet. Choice of the large/risky option delivered four pellets at changing probabilities. When probability of large reward delivery was set at 50%, the large/risky option was the optimal choice. Conversely, when the probability of reward delivery was set at 12.5% the small/certain option had greater objective utility. On rewarded trials, the house light remained illuminated for another 3 s, whereas after nonrewarded choices or trial omissions, the light was extinguished coincidentally with lever retraction. Pellet delivery was initiated immediately after a press, and multiple pellets were delivered 0.5 s apart.

Animals were trained on one of two variants of the task. In the descending variant, the probability of reward delivery on the large/risky lever is initially set to 50% for the first 30 trials and then subsequently set to 12.5% for the remaining 30 trials. In the ascending variant, the order of probabilities was reversed, starting with the 12.5%, followed by the 50% probability block. Rats were trained until the group demonstrated stable patterns of optimal risky choice (~30 d). Optimal choice was defined based on the following criteria: First, rats showed a bias toward the risky option (>50% risky choice) in the 50% probability block

when the large/risky option had greater objective utility (i.e., would yield more reward in the long-term vs the small certain option). Second, rats showed a bias away from the risky option (<50% risky choice) in the 12.5% block, where the small/certain option had greater objective utility. Choice stability was evaluated by analyzing data from 3 consecutive days using a two-way repeated-measures ANOVA, with day and probability block as the two factors. Behavior was deemed stable when there was no main effect of day and no day × block interaction (at $p > 0.10$). Once stable patterns of choice were displayed, optogenetic test sessions commenced.

Optogenetic inhibition

Separate experiments were conducted in groups of rats expressing eArchT3.0 and those expressing eYFP as a control. The eArchT opsin was chosen for these experiments as it has been shown to be effective at reducing evoked activity for short durations, particularly when targeting cell bodies (Wiegert et al., 2017). Although rebound excitation from sustained photoinhibition of presynaptic terminals using eArchT has been reported, this is only observed following prolonged 5 min periods of light application, whereas shorter millisecond-to-second-range periods suppress neural activity (Mahn et al., 2016). Moreover, these rebound excitatory effects are more likely to occur when using this approach for terminal inhibition on account of the relatively small intracellular volume of axons, whereas they are less of a concern when targeting cell bodies as was done in the present study as the larger soma is better suited to buffering against significant changes in pH and ionic composition (Wiegert et al., 2017; Lafferty and Britt, 2020). Additionally, we have shown that 4 s durations of light delivery to basolateral amygdala (BLA) terminals expressing eArchT within the nucleus accumbens were effective at suppressing evoked firing in the absence of an increase in spontaneous or evoked firing compared with baseline (Bercovici et al., 2018). Moreover, others have validated the alterations in reward-related behaviors seen following the use of Arch-mediated cell body inhibition (Lafferty and Britt, 2020). Viewed collectively, it is highly likely that the behavioral effects reported here were attributable to suppression of PL neural activity.

Green (532 nm) diode-pumped solid-state lasers (Laserglow Technologies) were coupled to a 200- μ m-core patch cable (Thorlabs) followed by a dual-channel optical rotary joint (Doric Lenses) that split the light so that each channel emitted 50% of the light intensity output directly emitted from the laser. The rotary joint was attached to optic fiber patch cables (Thorlabs), which were then plugged into ferrules on the animal heads. Before each test day, lasers were turned on using a TTL pulse delivered from the Med Associates control system, and light intensity from each patch cable was measured to be between 20–30 mW. Before surgical implantation, all optic fibers were measured to emit 85–95% of the light emitted through each patch cable. and it is estimated that throughout testing, between 17 and 28 mW of light reached the tip of the fibers in PL with each TTL pulse.

We used an approach similar to that used by our group and others to dissect how activity during different phases of the decision-making sequence influences choice (Orsini et al., 2017; Bercovici et al., 2018). Rats received multiple optogenetic tests, each consisting of a 3 d sequence: The first 2 d were baseline days, where the animal was connected to the fiber optic cables, but no light was delivered. On the subsequent test day, animals received brief pulses of light to suppress PL activity during

Probabilistic Discounting

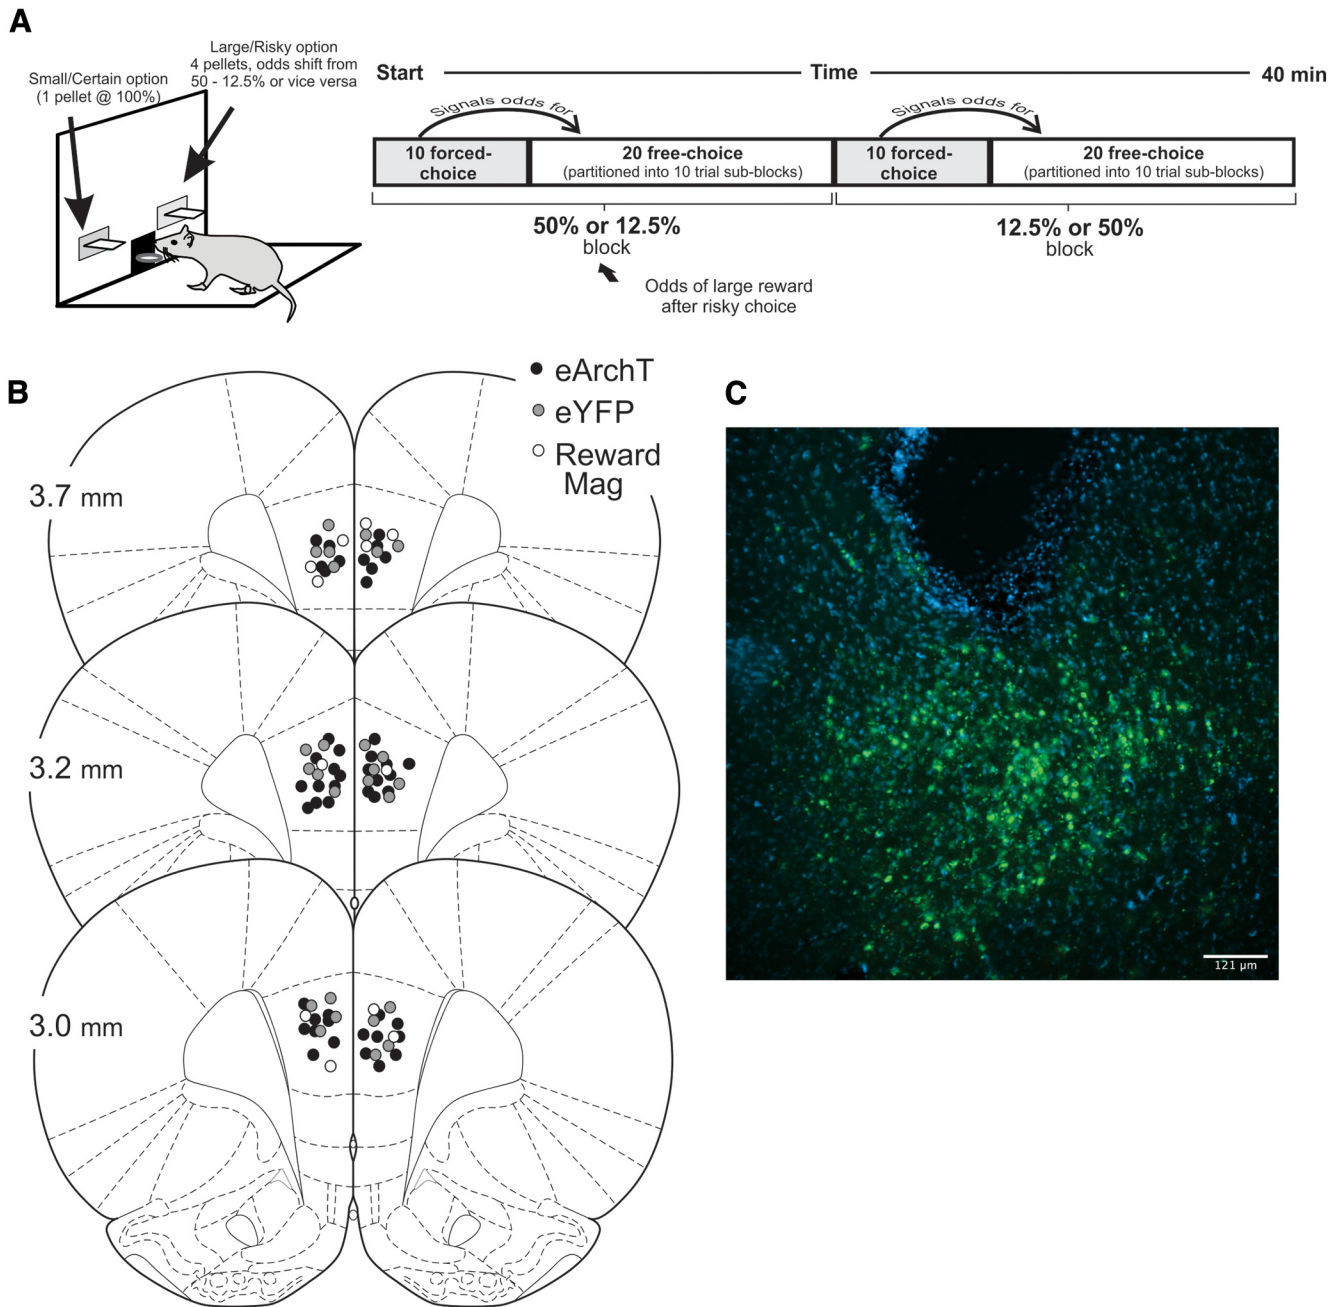


Figure 1. Probabilistic discounting task and histology. **A**, Probabilistic discounting task diagram. **B**, Location of fiber optic placements in PL corresponding to animals receiving the inhibitory eArchT opsin or control eYFP infusions tested on the probabilistic task and those receiving eArchT tested on a reward magnitude discrimination (Reward Mag). Numbers correspond to millimeters from bregma. **C**, Representative slice of PL expression in cell bodies (blue is DAPI; green is eYFP) with optic fiber placement.

discrete task events. Rats received two test sequences for each type of optogenetic manipulation, separated by at least 2 d of retraining. Behavioral data from the baseline days were averaged and compared with those obtained on the silencing test days. There were no differences in percentage choice of the risky option, our primary dependent variable, from the baseline days across all manipulations (all p values > 0.32). The order in which animals received the different optogenetic tests was counterbalanced, and rats completed both tests for a particular manipulation before being retrained for 3–5 d and receiving the next series of tests. Some rats did not receive all tests because of damage to headcaps. This resulted in a different number of subjects in each analysis.

Silencing before choice. In this experiment, laser light was initiated at the start of each trial, 4 s before lever extension, and was terminated either when a choice was made or after 10 s elapsed following lever extension (omission). Under these conditions, light was delivered for 4–10 s each trial depending on response latency. Light was delivered only during free-choice trials, as we were primarily interested in how PL activity influences action selection when choosing between both options.

Silencing during reward omissions. Here, PL activity was inhibited during the outcome of trials where animals chose the risky option and did not receive the larger reward (a risky loss). For this and all other outcome silencing experiments, laser light

was delivered during the outcomes of both free- and forced-choice trials. This is because even during forced-choice trials, the contingency between the probability of reward delivery and the outcome of lever press remains the same, and as such, outcome-related PL activity during these trials is still relevant as it informs the animal of the relative likelihood that a choice will or will not be rewarded. In this particular experiment, a reward omission following a forced choice is still a loss and therefore can still have an impact on future choices. Lasers were left on for 7 s after lever press, which would have overlapped the moment when pellets would have been delivered on a rewarded trial.

Silencing during large rewards. Another experiment inhibited activity in the PL following rewarded risky choices (risky wins). On these test days, light was delivered on all free- and forced-choice trials after a rat selected the large/risky lever and received the larger reward. Laser light was initiated immediately after these choices and was terminated 7 s after lever press, overlapping with pellet delivery and consumption.

Silencing during small/certain rewards. In these experiments, PL activity was suppressed after small/certain choices (small wins), wherein light was delivered on all free- and forced-choice trials immediately after a rat selected the small/certain option. Lasers were left on for 7 s after lever press, which included the time it took for pellet delivery and consumption.

Intertrial interval. To verify that the outcome-associated effects of silencing the PL were attributable to inhibiting neural activity temporally linked to these events, a control experiment was conducted where activity in the PL was inhibited during a random 4 s interval starting 6–14 s after the start of the 40 s intertrial interval (ITI) for all free- and forced-choice trials.

Reward magnitude discrimination

A separate cohort of rats expressing eArchT3.0 was trained for ~25 d on a control reward magnitude discrimination task. In this task both the large (four pellets) and small (one pellet) reward choice options were set to 100% probability of reward delivery for four blocks of two forced-choice followed by 10 free-choice trials. On separate tests, animals received optogenetic silencing during the period before choice and during the large reward delivery as described above.

Histology

Rats were killed via transcardial perfusion with 4% paraformaldehyde. Brains were fixed in 4% paraformaldehyde for 24 h and then stored in 30% sucrose in 1 M PBS. Each brain was flash frozen on dry ice and sliced in 50 μ m sections using a cryostat. Sections were mounted onto slides, counterstained, and coverslipped using Fluoromount-G with DAPI (eBioscience). Viral expression and ferrule placements (Fig. 1B,C) were verified in the PL using a 1 \times objective on an Axio Zoom microscope (Zeiss). Three rats whose placements were found to be outside the borders of the PL were determined referencing a neuroanatomical atlas (Paxinos and Watson, 2005) and were subsequently removed from data analysis (one in the eArchT group and two from the eYFP control group).

Data analysis

The primary dependent variable was the proportion of choices of the large/risky option after each behavioral manipulation. The total number of large/risky choices made within a probability block was divided by the total number of choices made in that block, thereby factoring out trial omissions. Although this index of choice can sometimes be skewed if an animal makes a large

number of omissions during a free-choice block, as described in the results, omission rates were relatively low in this study. Across all experiments, only two animals that received PL inhibition before choice showed a relatively large increase in omissions over 20 trials of a free-choice block (8 and 11), and eliminating the data from these animals from the analysis did not qualitatively change the results. Choice data were analyzed with three-way ANOVAs, with treatment (optogenetic inhibition vs baseline) and probability block (50 vs 12.5%) as within-subject factors and task variant (descending vs ascending probabilities) as a between-subjects factor. The main effect of probability block for all these analyses was always significant ($p < 0.001$). Simple main effects analyses partitioning significant two-way interactions consisted of one-way ANOVAs (Bonferroni corrected).

Additional choice-by-choice analyses examined how suppression of PL activity during different task events influenced action selection after different risky choice outcomes as indices of reward and negative feedback sensitivity (win-stay and lose-shift behavior, respectively). Each free choice was compared with the outcome of the preceding choice of the risky option. Win-stay ratios were calculated from the proportion of trials where rats chose the risky option following a risky win (receipt of the large reward), divided by the total number of free-choice risky wins. As many animals did not experience a risky win in the low probability block, win-stay values were pooled across both blocks and analyzed with one-way ANOVAs. Conversely, lose-shift ratios were calculated for each probability block separately and were based on the proportion of trials where rats chose the small/certain option following a risky loss (nonrewarded choice) over the total number of risky losses within each block. For these data, we were able to obtain enough risky losses in both the 50 and 12.5% probability blocks. As we were particularly interested in examining how reward omissions differentially influence subsequent choice as a function of reward probability, these data were analyzed using two-way ANOVAs, with treatment and probability block as two within-subject factors. For all significant interaction effects, simple main effects analyses were conducted using one-way ANOVAs where appropriate. Other performance measures included the number of trial omissions and average choice latencies (time between lever extension and lever press) and were analyzed with one-way repeated-measures ANOVAs. In instances where suppression of PL activity increased choice latencies, subsequent analyses probed if these effects varied based on either probability block or choice type (risky vs certain) using two-way repeated-measures ANOVAs, although as discussed below, none of these analyses revealed any differential effects.

The win-stay/lose shift analyses were complemented by additional multilevel modeling analyses examining how a combination of outcomes spanning two trials back influenced a particular choice and how PL silencing during different task events altered the impact of these different outcomes. In so doing, multilevel logistic regression analyses were conducted on trial-level data from the different experiments, targeting the probability blocks where PL inhibition induced a significant change in overall risky choice. This yielded six separate models that analyzed choice data from the 50% block, and in some instances the 12.5% blocks, separately for each manipulation (silencing during pre-choice, reward omission (losses), large/risky wins, and small/certain wins). Each model was specified the same way with Trials (level 1) nested within Rats (level 2). Treatment (baseline vs inhibition) was included as a level 1 predictor. To assess how the outcome history from the preceding two trials influenced choice, we included another series of categorical level 1 predictors, with

each level representing one of nine possible combinations of outcomes on the past two trials. This resulted in the following nine levels, herein denoted based on the outcomes 2 and 1 trials back from a choice (denoted n-2 and n-1, respectively); Risky win→Risky win, Risky win→Risky loss, Risky win→Certain win, Risky loss→Risky win, Risky loss→Risky loss, Risky loss→Certain win, Certain win→Risky win, Certain win→Risky loss, and finally Certain win→Certain win. Choice trials that followed two consecutive wins were excluded from the regression models on data from 12.5% probability blocks as this combination of outcomes was extremely rare or nonexistent across the datasets. With these predictors as a backbone of the analyses, how PL inhibition affected choice following different outcomes was examined by including Treatment × Outcome interaction terms across the eight to nine possible outcome combinations in the model. When one or more of these terms were found to be statistically significant, the interactions were partitioned with simple slopes analyses assessing how PL inhibition may have differentially altered the probability of a risky choice versus baseline on trials that followed a particular combination of outcomes. In comparison, a main effect of Treatment in the absence of any interaction with outcomes indicated that PL inhibition altered risky choice in a manner that was not modulated by the recent experienced outcomes. Here, odds ratios that were significantly less than or >1.0 reflected an overall decrease or increase in risky choice within a particular block. These analyses were conducted using the lmerTest and interactions packages in R software.

Animals trained on the reward magnitude discrimination task were analyzed on each manipulation for the proportion of choices of the large reward option, using a two-way ANOVA with treatment and trial block (four blocks of 10 free-choice trials) as two within-subject factors. Other performance measures were analyzed using one-way ANOVAs.

Results

Each of the probabilistic discounting experiments discussed below included animals trained on either the descending or ascending variant of the task. However, in each of these analyses, there were no main effects of task variant or variant × treatment interaction effects (all p values > 0.09), indicating that alterations in choice induced by suppression of PL activity were comparable regardless of the order in which reward probabilities changed over a session. These findings will not be mentioned further. The graphical presentation of the choice data in each probability block is partitioned over sub-blocks of 10 trials to display the relative consistency of these effects within each probability block.

PL inhibition before choice reduced choice bias toward more preferred options

One of our primary interests was to examine how activity in the PL occurring during epochs immediately before initiation of a choice guides decision-making. To this end, PL activity was suppressed during the deliberation period before action selection. During baseline tests, rats displayed optimal choice patterns, displaying a strong preference toward the large/risky option during the 50% block and a bias away from this option during the 12.5% block. In animals in the eArchT group ($n = 26$), optogenetic suppression during the prechoice period markedly perturbed decision-making (Fig. 2A,B). Analysis of these data revealed a significant main effect of treatment ($F_{(1,24)} = 8.67$,

$p = 0.007$) and more pertinently, a significant treatment × block interaction ($F_{(1,24)} = 35.48$, $p < 0.0001$). This interaction was driven by a decrease in risky choice in the 50% block ($p < 0.0001$) and an increase in risky choice in the 12.5% block ($p = 0.04$). In comparison, laser light administered before choice in control rats expressing eYFP ($n = 13$) caused no significant changes in behavior (all F values < 1, all p values > 0.40; Fig. 2C).

In addition to disrupting choice behavior, suppression of pre-choice PL activity slowed decision latencies in rats expressing eArchT ($F_{(1,24)} = 23.02$, $p = 0.0001$; Fig. 2A, inset) in a manner that did not vary across probability block or risky versus certain choices (all p values > 0.20; Table 1). This was not observed in the eYFP group (mean ± SEM, baseline = 0.59 ± 0.09 s; laser = 0.63 ± 0.07 s; $F_{(1,11)} = 0.13$, $p = 0.72$). Likewise, there was a slight but significant increase in trial omissions for the eArchT group ($F_{(1,24)} = 12.04$, $p = 0.002$; Table 1), with these occurring with comparable frequency in the 50 and 12.5% blocks. However, laser light did not alter omissions in the eYFP group (baseline = 0.6 ± 0.2 ; laser = 0.7 ± 0.3 ; $F_{(1,11)} = 0.20$, $p = 0.67$).

Further analyses examined how these perturbations in decision biases related to changes in how the most recent rewarded and nonrewarded outcomes influenced subsequent action selection. Under baseline conditions, animals had a strong tendency to follow a rewarded risky choice with another risky choice, displaying win-stay behavior on >80% of these trials; the vast majority of these occurred in the 50% block. Suppression of PL activity during the deliberation period caused a robust reduction in win-stay behavior ($F_{(1,24)} = 24.74$, $p < 0.0001$; Fig. 2D, left bars). Indeed, during these tests, animals were seemingly indifferent to previous rewarded risky choices as they were just as likely to make a risky choice or shift to the small/certain option on the next choice trial (one-sample t test vs 50%; $t_{(25)} = 1.95$, $p = 0.06$). On the other hand, lose-shift behavior was strongly dependent on the context of the probability block. Under baseline conditions, rats rarely shifted to the small/certain option after a nonrewarded risky choice when reward probabilities were relatively high (50%), but they did so on the majority of these trials when reward probabilities were low (12.5%; Fig. 2D, right bars). Here, suppression of PL activity differentially altered lose-shift behavior (treatment × block interaction, $F_{(1,25)} = 48.99$, $p < 0.001$), increasing it during the 50% block ($p < 0.001$; one-sample t test vs 50%, $t_{(25)} = 1.86$, $p = 0.07$), whereas in the 12.5% block, lose-shift behavior was also reduced to chance levels ($p = 0.03$; one-sample t test vs 50%, $t_{(25)} = 1.03$, $p = 0.31$). Figure 2E displays trial-by-trial choice and outcome patterns of an exemplar rat in this experiment. This average risky choice of the rat during the 50 and 12.5% blocks was comparable to the group mean under baseline and test conditions. Thus, suppression of PL neural activity before action selection disrupted choice of more preferred, higher utility options, and this appeared to be linked to an inability to incorporate information about the outcomes of preceding decisions to guide subsequent action selection.

PL inhibition during reward omissions differentially alters risky choice

It has previously been shown that activity of PL neurons associated with reward omissions are predictive of subsequent choices during risk/reward decision-making (Passecker et al., 2019). Thus, in this experiment, PL activity was suppressed on all trials in which the large/risky option was selected, but no reward was delivered. For rats in the eArchT group ($n = 26$), suppression of

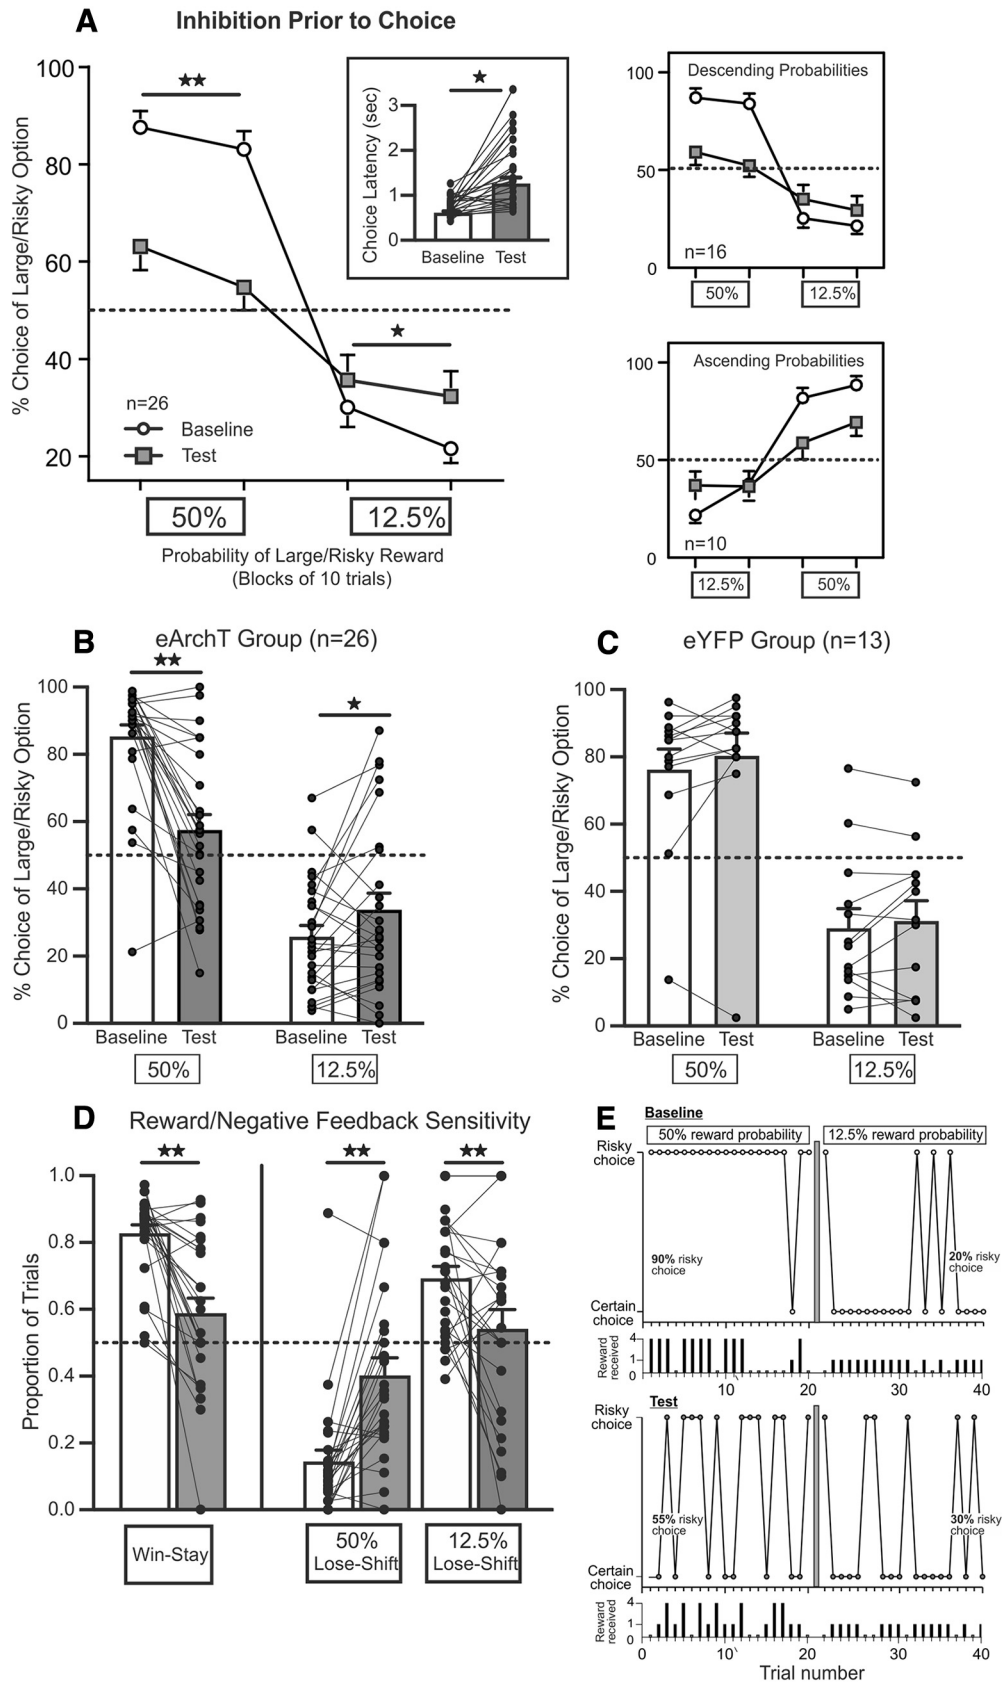


Figure 2. Inhibition of PL activity before choice disrupts bias toward more preferred, higher-utility options. **A**, Left, percentage choice of the large/risky option under baseline conditions and during optogenetic inhibition tests for rats in the eArchT group. Inhibition before choice reduced risky choice during the higher 50% probability block and increased risky choice in the lower 12.5% block. For this and all other figures, data are partitioned over blocks of 10 trials. Inset, Plot of average decision latencies, suppressing PL activity before choice increased deliberation times. Right, Data from animals trained on descending (top, $n = 14$) and ascending (bottom, $n = 9$) task variants. **B**, Individual risky choice data for the eArchT group, averaged over 20 trials in the 50 and 12.5% probability blocks. **C**, Individual risky choice data for eYFP group. **D**, PL inhibition before choice induced near-random patterns of reward/negative feedback sensitivity.

Table 1. Mean ± SEM number of omissions (over 60 trials) and choice latency data partitioned by probability block and choice type for rats expressing eArchT under baseline conditions and following PL silencing (laser) across the difference experiments where alterations in choice were observed

Experiment	Condition	Omissions	Latency (by block)		Latency (by choice type)	
			50%	12/5%	Risky	Certain
Prior to choice	Baseline	0.5 ± 0.2	0.53 ± 0.02	0.70 ± 0.05	0.58 ± 0.04	0.63 ± 0.03
	Laser	4.7 ± 1.1*	1.29 ± 0.20*	1.25 ± 0.15*	1.25 ± 0.12*	1.26 ± 0.16*
Reward omission	Baseline	0.4 ± 0.2	0.54 ± 0.04	0.58 ± 0.04	0.55 ± 0.03	0.56 ± 0.04
	Laser	2.5 ± 0.8*	0.81 ± 0.11*	0.85 ± 0.09*	0.71 ± 0.11*	0.68 ± 0.11*
Large/risky reward	Baseline	0.4 ± 0.1;	0.59 ± 0.04	0.65 ± 0.06	0.62 ± 0.05	0.62 ± 0.05
	Laser	1.4 ± 0.3*	0.60 ± 0.05	0.74 ± 0.07	0.67 ± 0.05	0.65 ± 0.04
Small/certain reward	Baseline	0.4 ± 0.2	0.49 ± 0.04	0.57 ± 0.05	0.55 ± 0.04	0.55 ± 0.03
	Laser	0.5 ± 0.1	0.47 ± 0.03	0.58 ± 0.05	0.56 ± 0.03	0.55 ± 0.04

* $p < 0.05$ laser versus baseline.

PL activity during these outcomes altered decision-making in a manner dependent on the likelihood that a risky choice would pay off. Analysis of the choice data yielded a significant treatment × probability block interaction ($F_{(1,24)} = 32.94$, $p < 0.0001$; Fig. 3A,B) with no main effect of treatment ($F_{(1,24)} = 0.99$, $p = 0.33$). This interaction reflected a significant increase in risky choice during the 12.5% block ($p = 0.008$) but, unexpectedly, a decrease in risky choice in the higher 50% probability block ($p = 0.01$). Additionally, silencing PL activity during reward omissions increased the time to make a choice on subsequent trials ($F_{(1,24)} = 7.67$, $p = 0.01$; Fig. 3A, inset; Table 1), but this did not vary across probability block or risky versus certain choices (all p values > 0.50 ; Table 1). PL silencing during these epochs also increased trial omissions versus baseline by slightly more than two omissions over the 60 trials ($F_{(1,24)} = 5.85$, $p = 0.02$; Table 1) with about one each occurring during forced- and free-choice trials. On the other hand, this manipulation did not affect choice behavior in the eYFP control group ($n = 13$; all $F_{(1,11)}$ values < 1.0 , all p values > 0.50 ; Fig. 3C). These animals also showed no change in trial omissions (baseline = 0.9 ± 0.5 ; laser = 0.4 ± 0.2 ; $F_{(1,11)} = 3.63$, $p = 0.08$) and actually a slight 80 ms decrease in choice latency (baseline = 0.58 ± 0.05 s; laser = 0.50 ± 0.05 s; $F_{(1,11)} = 6.20$, $p = 0.03$) during this manipulation.

Silencing PL activity during reward omissions also altered sensitivity to losses occurring on a preceding trial that were dependent on the context of the probability block (treatment × block interaction: $F_{(1,25)} = 15.85$, $p < 0.001$; Fig. 3D, right bars). Thus, in the 12.5% block, when risky choices were less likely to be rewarded, suppressing PL activity rendered animals less sensitive to losses, reducing lose-shift behavior versus baseline ($p = 0.025$). Conversely, in the higher 50% probability block, inhibition of PL neural firing had the opposite effect, increasing the tendency to shift choice after a risky loss ($p = 0.02$). In comparison, win-stay behavior was also reduced ($F_{(1,24)} = 7.05$, $p = 0.01$; Fig. 3D, left bars), although this effect was less robust than that observed following prechoice silencing. Figure 3E displays trial-by-trial choice and outcome patterns of an exemplar rat in this experiment under baseline and test conditions. Collectively, these data show that suppressing PL activity after nonrewarded actions

does not unidirectionally alter subsequent choice. Instead, it appears that silencing this activity alters responses to losses within the context of reward history, amplifying or diminishing their impact relative to the estimated likelihood that risky choices will be more profitable than the alternative.

PL inhibition during large/risky rewards biases choice away from these options

A separate experimental series examined how activity associated with rewarded choices influences decision-making. In one such experiment, laser light was delivered only on trials when animals selected the large/risky option and received the larger reward (eArchT group, $n = 25$). Analysis of the choice data yielded a significant treatment × probability block interaction ($F_{(1,23)} = 15.80$, $p < 0.001$), reflecting that this manipulation reduced risky choice when reward probabilities were relatively high (50%, $p = 0.007$), but not when they were low (12.5%, $p = 0.31$; Fig. 4A,B). These effects on choice occurred in the absence of any change in decision latencies ($F_{(1,23)} = 1.22$, $p = 0.28$; Fig. 4A, inset; Table 1), although there was a slight but statistically reliable increase in trial omissions as, on average, PL silencing led to about one additional omission versus baseline ($F_{(1,23)} = 13.53$, $p = 0.001$; Table 1). eYFP-expressing rats that received the same manipulation ($n = 13$) showed no change in choice (Fig. 4C), choice latency (baseline = 0.59 ± 0.08 s; laser = 0.60 ± 0.09 s), or trial omissions (baseline = 0.7 ± 0.4 ; laser = 0.8 ± 0.4 ; all $F_{(1,11)}$ values < 1.6 , all p values > 0.20).

The reduction in risky choice induced by silencing PL activity during receipt of larger/risky rewards was associated with changes in how animals behaved after both rewarded and nonrewarded decisions. Thus, PL inhibition appeared to reduce the reinforcing properties of rewarded risky choices as indexed by a reduction in win-stay behavior ($F_{(1,24)} = 6.98$, $p = 0.014$; Fig. 4D, left bars). Interestingly, these treatments also altered how risky losses influenced subsequent choice, even though PL activity was not perturbed during these outcomes. Specifically, animals were more likely to shift to the small/certain option after a nonrewarded risky choice in the 50% block (treatment × block interaction, $F_{(1,24)} = 5.42$, $p = 0.03$; simple main effects, $p = 0.03$; Fig. 4D, right bars). Yet, lose-shift behavior was unaffected in the 12.5% block ($p = 0.99$), presumably because rats received very few large rewards (and corresponding laser pulses) during this part of the test session. Figure 4E displays trial-by-trial choice and outcome patterns of an exemplar rat in this experiment under baseline and test conditions. Together, these data show that suppressing PL activity during receipt of larger/risky rewards reduces biases toward these options when their objective utility is relatively high.

←

These treatments decreased win-stay behavior, and caused opposing changes in lose-shift behavior, increasing it in the 50% block and decreasing it in the 12.5% block. E, Individual, trial-by-trial free-choice and outcome data displayed by one rat during a baseline training session (top) and a subsequent test session when PL silencing occurred before choice (bottom). Missing symbols on the curves represent trial omissions. For this and all other figures, error bars indicate SEM. Stars and double stars denote $p < 0.05$ and $p < 0.001$ compared with baseline.

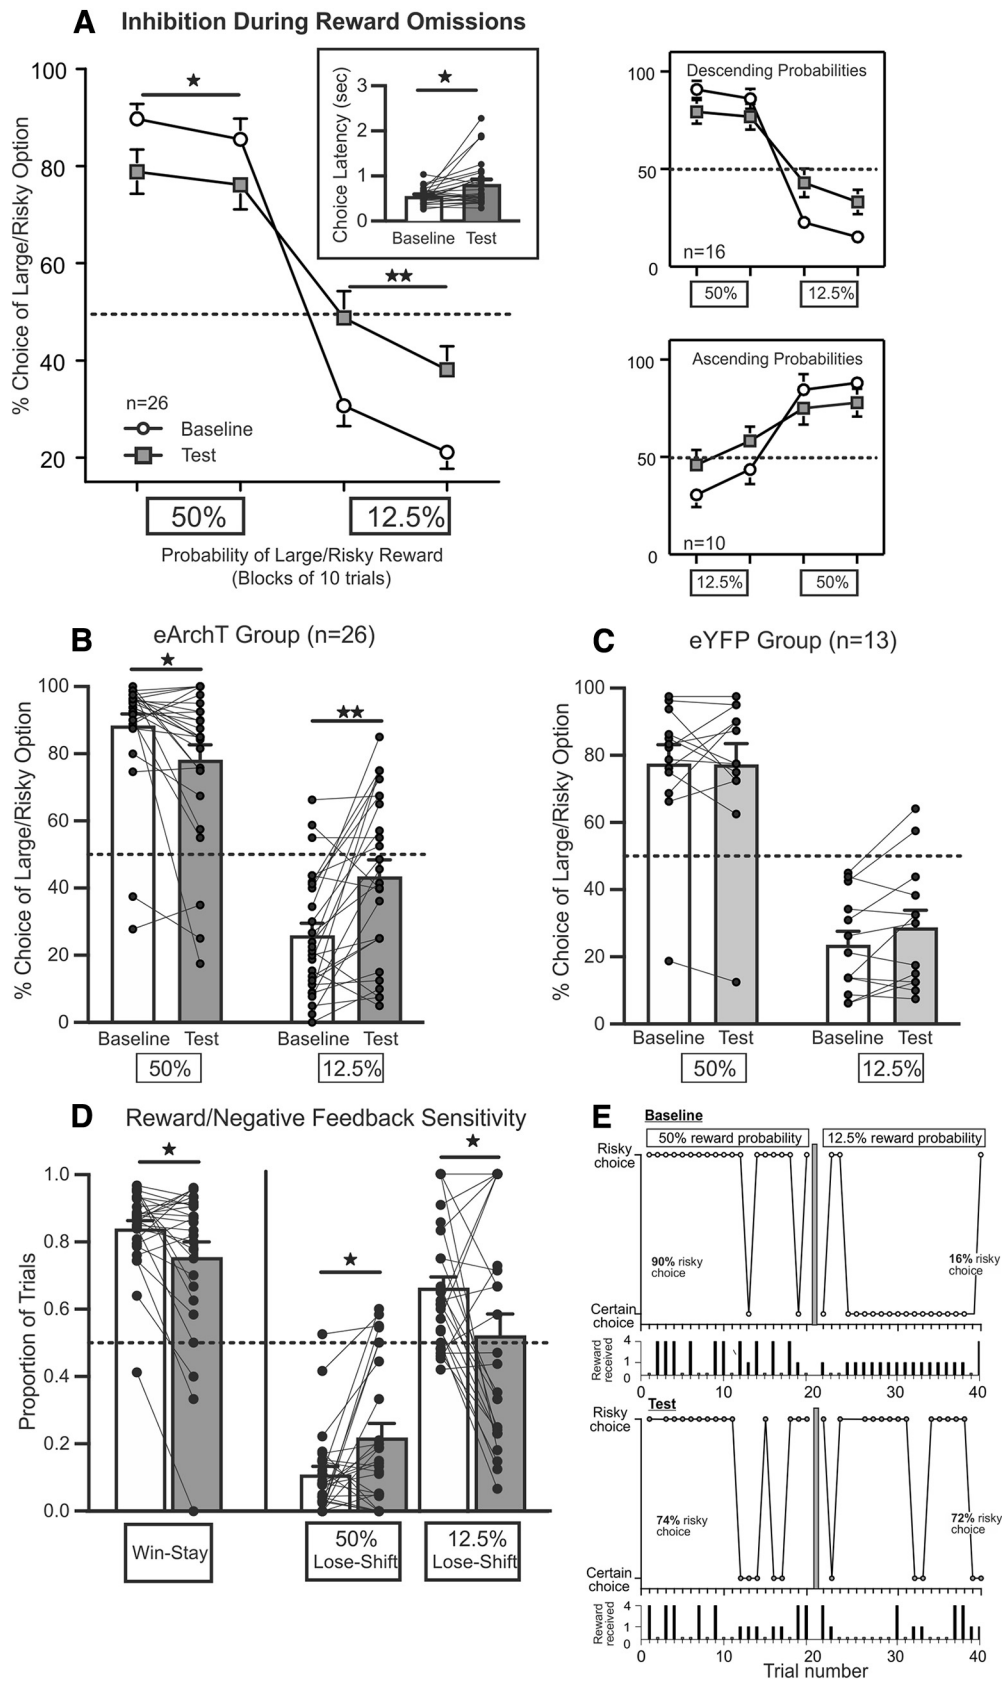


Figure 3. PL inhibition during reward omissions differentially alters risky choice depending on the context of reward probability. **A**, Left, Percentage choice of the large/risky option under baseline conditions and during optogenetic tests for rats in the eArchT group. Inhibition decreased risky choice in the 50% block and increased it in the 12.5% block. Inset, Plot of average decision latencies showing increased latency to make a choice following PL inhibition during reward omissions. Right, Data from rats trained on the descending (top, $n = 15$) and ascending (bottom, $n = 7$) task variants. **B**, Individual average risky choice data for eArchT group across probability blocks. **C**, Individual risky choice data for the eYFP group. **D**, Reward/negative feedback sensitivity of eArchT group. Suppressing PL activity during reward omissions increased lose-shift behavior during the 50% block but decreased it during the 12.5% block, without affecting win-

PL inhibition during small/certain rewards also reduces preference for large/risky rewards

Another experiment suppressed PL activity after selection of the small/certain option during delivery of the small reward. Somewhat surprisingly, this manipulation also caused a slight but reliable reduction in preference for the large/risky option in the high probability block. Analysis of the choice data from the eArchT group ($n = 26$) again yielded a significant treatment \times block interaction ($F_{(1,24)} = 8.86$, $p = 0.007$; Fig. 5A,B) with no main effect of treatment ($F_{(1,24)} = 1.42$, $p = 0.25$). Partitioning this interaction revealed a significant decrease in choice of the risky option during the 50% probability block ($p = 0.02$) but not the 12.5% block ($p = 0.25$). These effects were not associated with alterations in choice latencies ($F_{(1,24)} = 2.94$, $p = 0.10$; Fig. 5A, inset; Table 1) or trial omissions ($F_{(1,24)} = 0.80$, $p = 0.38$; Table 1). Curiously, analysis of the choice data from the eYFP group in this experiment ($n = 13$) yielded a significant main effect of treatment ($F_{(1,11)} = 11.07$, $p = 0.007$; Fig. 5C). However, this reflected a slight increase in risky choice, an effect opposite to what was observed in animals expressing eArchT. No latency (baseline = 0.67 ± 0.13 s; laser = 0.66 ± 0.12 s) or omission effects (baseline = 0.9 ± 0.7 ; laser = 0.6 ± 0.3) were observed in control animals in this experiment (all F values < 3.0 , all p values > 0.11).

The shift in bias away from the large/risky option in the 50% block was associated with a slight reduction in win-stay behavior, although this effect did not achieve statistical significance ($F_{(1,25)} = 2.83$, $p = 0.10$; Fig. 5D, left bars). Conversely, silencing PL activity during receipt of smaller rewards did not influence lose-shift behavior during any part of the session (treatment \times probability block interaction: $F_{(1,25)} = 2.65$, $p = 0.12$; Fig. 5D, right bars). Figure 5E displays trial-by-trial choice and outcome patterns of an exemplar rat in this experiment under baseline and test conditions. Collectively these findings show that PL neural activity coinciding with receipt of smaller/certain rewards can also influence the allure that larger, higher probability rewards exert over action selection.

Multilevel modeling analyses on broader choice-outcome history

In each of the experiments described above, PL inhibition reduced risky choice in the higher (50%) probability block. Conversely, silencing before choice or during reward omissions increased risky choice in the lower-probability 12.5% block. These perturbations in choice biases of more preferred options were associated with differential alterations in the impact that preceding risky wins or losses exerted on a current choice. It was of further interest to explore how different combinations of recent choice outcomes spanning the past two trials back from a current choice influenced action selection and how PL silencing during different task events altered how these wins/losses shaped choice biases. We used multilevel logistic regression models to analyze data across the four different experiments, focusing on probability blocks where we observed a significant alteration in risky choice (see above, Materials and Methods). This resulted in six separate model analyses, the results of which are displayed in Table 2. In this table, the treatment \times loss \rightarrow loss

outcome factor was set as the reference for calculating odds ratios for the other interaction terms to determine whether they were significantly greater or < 1.0 . From these analyses, we plotted the probability of a large/risky choice under baseline and test conditions across the various outcome combinations along with the 95% confidence intervals for the purposes of data presentation (Fig. 6). A summary of the main findings from each of these follows.

Analysis of data from the prechoice inhibition experiment revealed a main effect of Treatment in the 50% block (odds ratio = 0.31, $p < 0.001$) but no Treatment \times Outcome interactions (all p values < 0.067 ; Table 2; Fig. 6A, left). This indicates that prechoice PL inhibition reduced risky choice in the higher probability block in a manner that was independent of the particular combination of recent outcomes experienced. In contrast, analysis of data from the 12.5% block also revealed significant Treatment \times Outcome interactions (Table 2). Simple slopes analysis revealed that during this block, prechoice PL inhibition reduced the impact of the most recent losses. This manipulation increased risky choice only on trials immediately following a loss (i.e., reduced lose-shift behavior; all p values < 0.05), regardless of the outcome experienced $n-2$ trials back, but did not alter choice when losses occurred two trials back (Fig. 6A, right).

Separate analyses were conducted on data obtained when PL was silenced during reward omissions. Analysis of data from the 50% block revealed a comparable effect to that of prechoice inhibition, yielding a significant main effect of Treatment (odds ratio = 0.37, $p < 0.001$) in the absence of any interactions with the outcome levels (all p values > 0.13 ; Table 2; Fig. 6B). Yet, when the probability of obtaining the larger reward was 12.5%, PL inhibition increased risky choice when rats experienced any combination of the two least favorable outcomes (i.e., a loss or small reward, all p values < 0.01), but not when a risky win was received within the preceding two trials.

Silencing the PL during receipt of either large/risky or small/certain rewards shifted bias away from the risky option in the higher 50% probability block without affecting choice in the 12.5% block. Yet, regression analyses on these data showed this effect was more selectively dependent on recent outcomes experienced, compared with silencing before choice or during reward omissions. For example, analysis of the data from the large/risky reward experiment did not yield a significant main effect of Treatment (odds ratio = 0.58, $p = 0.073$), but did yield significant Treatment \times Outcome interactions (Table 2; Fig. 6C). Specifically, this manipulation tended to reduce risky choice and win-stay behavior after combinations of small and large reward outcomes except for two consecutive risky wins (all p values < 0.01). In addition, lose-shift behavior (i.e., reduced risky choice after a loss) was increased only after recent losses were preceded by a reward (all p values < 0.05) but not when losses occurred two trials back.

Last, silencing PL during small/certain wins did not produce a significant main effect of Treatment (odds ratio = 0.98, $p = 0.95$) but did yield significant Treatment \times Outcome interactions (Table 2; Fig. 6D). Note that this manipulation did not result in significant changes in our conventional analyses of win-stay or lose-shift behavior when only taking into account the most recent choice outcomes. However, simple slopes partitioning of the Treatment \times Outcome interactions showed that PL silencing during receipt of small rewards reduced risky choice after recent losses preceded by a reward or when a loss preceded a risky win (all p values < 0.05). In addition, silencing during these

←

stay behavior. **E**, Individual, trial-by-trial free-choice and outcome data displayed by one rat during a baseline training session (top) and subsequent test session when PL silencing occurred during reward omissions (bottom).

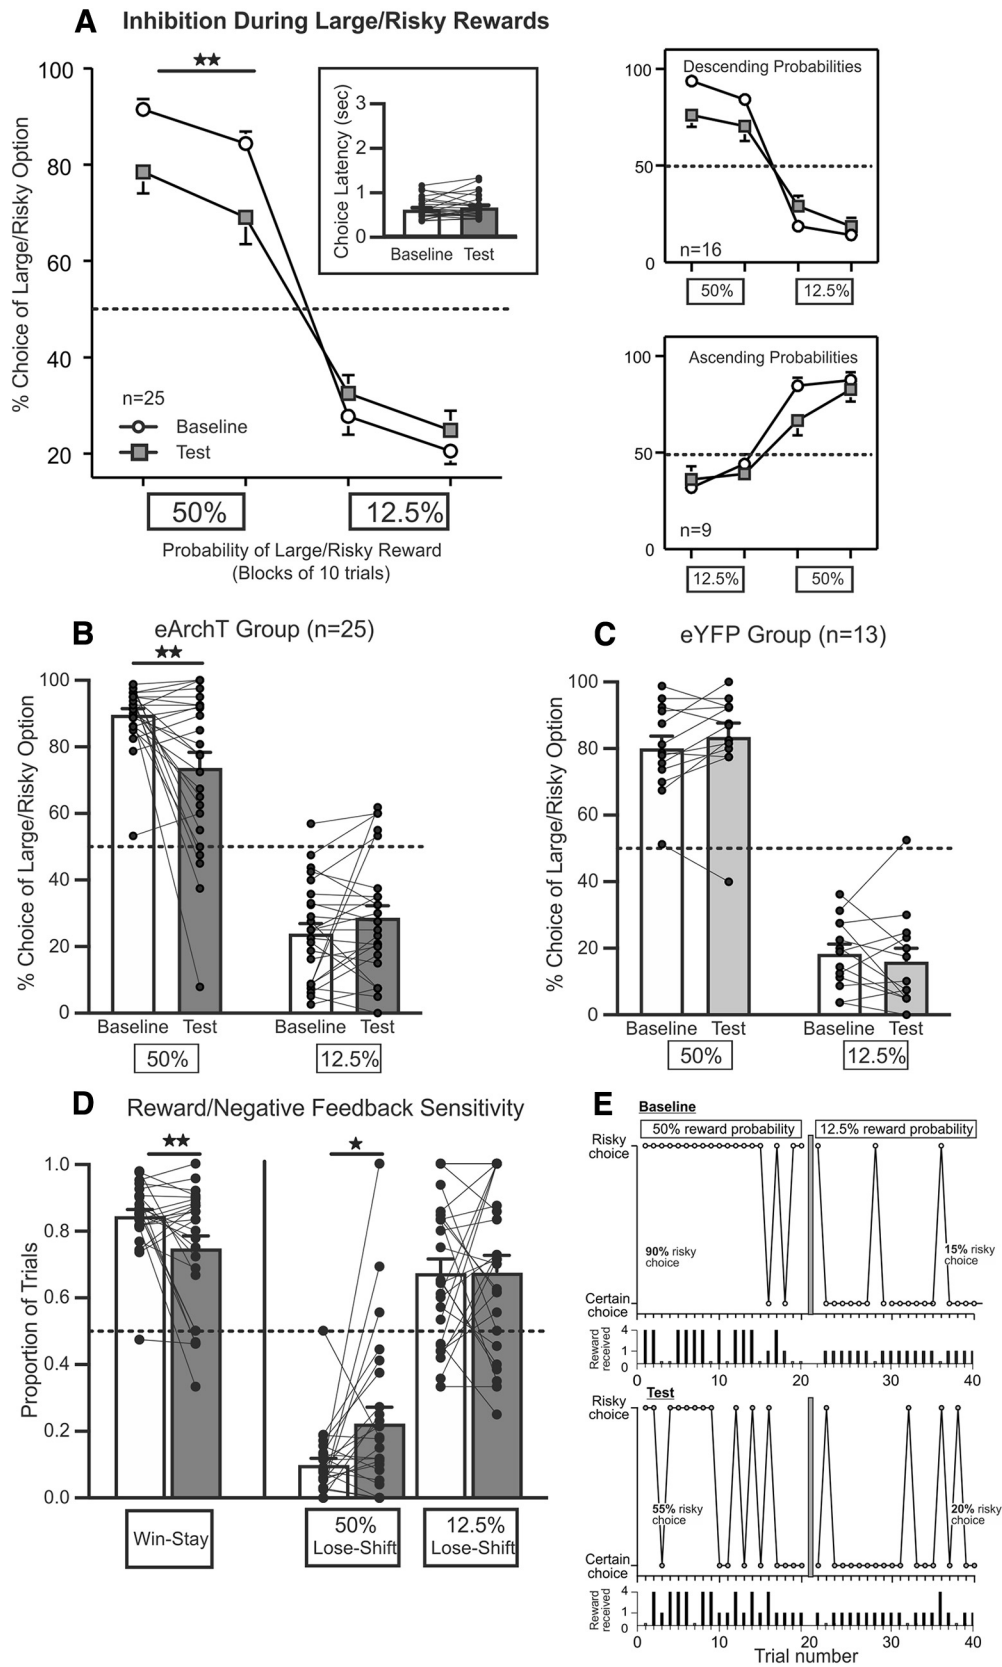


Figure 4. PL inhibition during large rewards decreases risky choice when reward probabilities are high. **A**, Left, Percentage choice of the risky option for eArchT ($n = 24$) group during baseline and on optogenetic tests. Inset, Average choice latency, which was unaffected by PL inhibition during large rewards. Right, Plots of risky choice data for animals trained on the descending (top, $n = 16$) and ascending (bottom, $n = 8$) versions of the task. **B**, Individual risky choice data for the eArchT group showed a pattern of reduction in risky choice in the 50% block. **C**, Individual risky choice data for the eYFP group ($n = 13$), which displayed no reliable change in choice. **D**, Win-stay/lose-shift analyses of eArchT group revealed that PL inhibition during large reward delivery reduced win-stay behavior and also increased lose-shift behavior during the 50% block. **E**, Individual, trial-by-trial free-choice and outcome data displayed by one rat during a baseline training session (top) and subsequent test session when PL silencing occurred during delivery of the large reward (bottom).

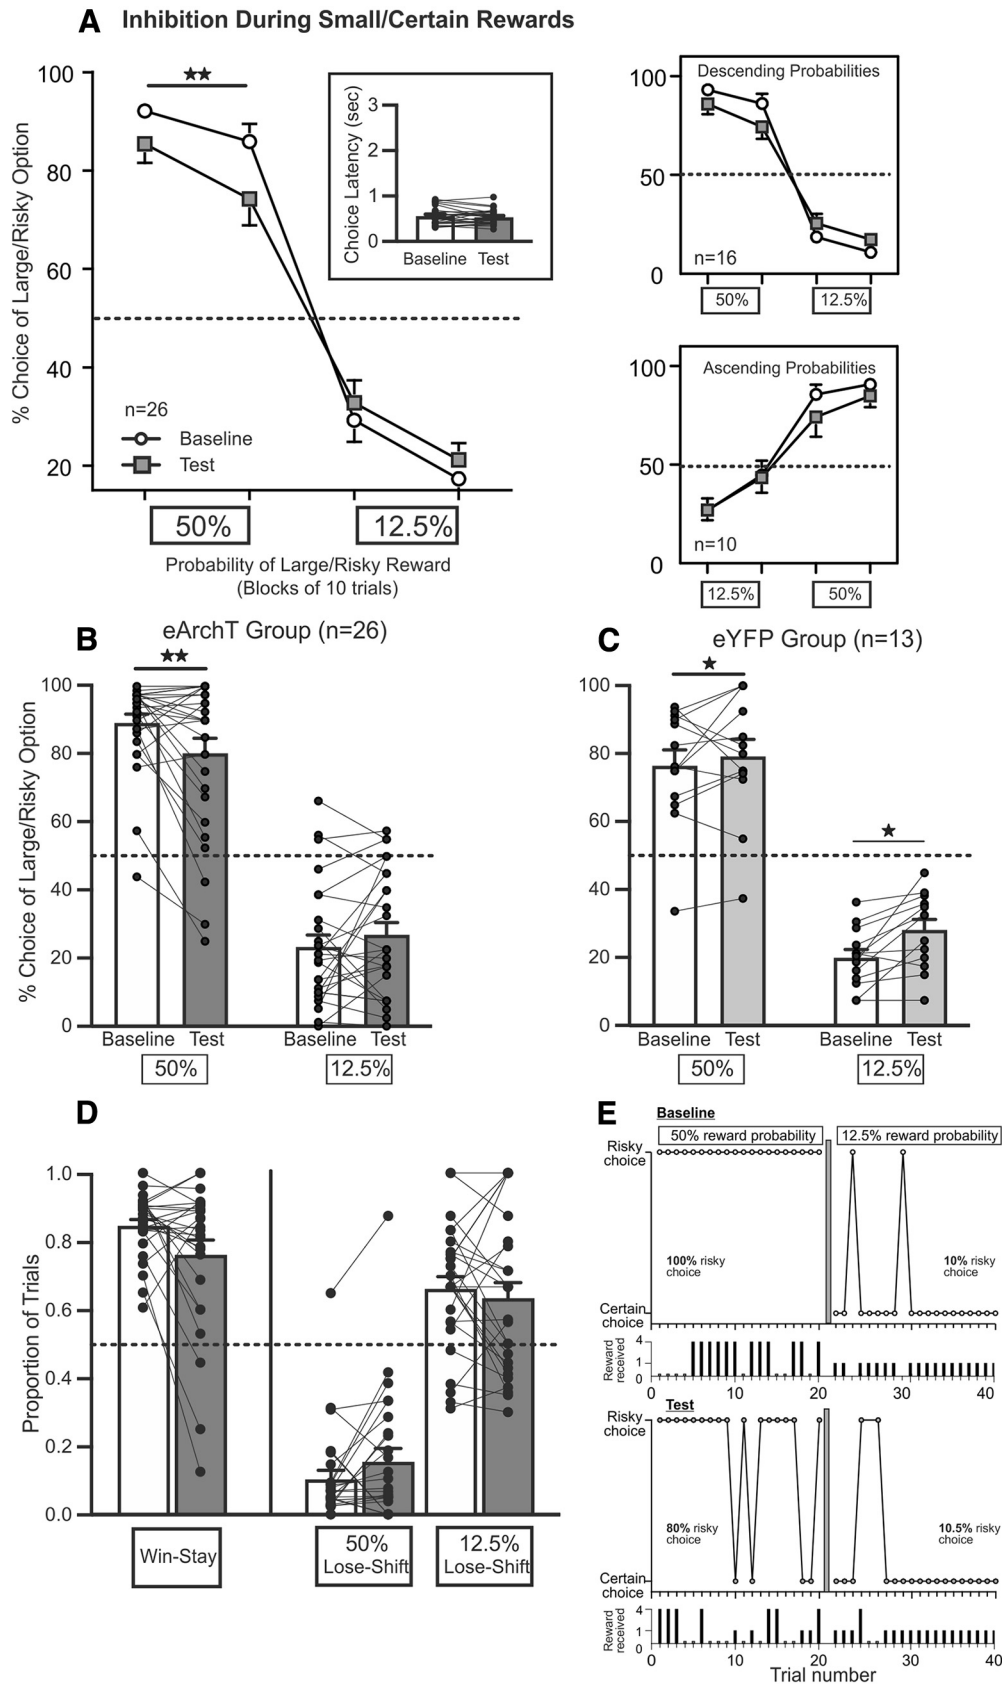


Figure 5. PL inhibition during small reward delivery decreased preference for the large/risky option when it was more advantageous. **A**, Left, Risky choice data for eArchT ($n = 22$) group. PL inhibition decreased risky choice during the 50% block on optogenetic tests relative to baseline. Inset, Choice latency, which was not affected by PL inhibition. Right, Plots of risky choice data for animals trained on the descending (top, $n = 15$) and ascending (bottom, $n = 7$) versions of the task. **B**, Individual risky choice data for the eArchT group. **C**, Individual risky choice data for the eYFP group ($n = 13$). In these animals, laser light delivery actually caused a small increase in risky choice, an effect opposite to that seen in rats expressing eArchT. **D**, Win-stay/lose shift data for the eArchT group. **E**, Individual trial-by-trial free-choice and outcome data displayed by one rat during a baseline training session (top) and subsequent test session when PL silencing occurred during delivery of the small reward (bottom).

Table 2. Summary of odds ratios, 95% confidence intervals, and *p* values from multilevel logistic regression analyses of how PL inhibition during different task events altered risky choice as a function of different outcome histories 2→1 trials back

Predictors	Prechoice inhibition 50% block			Prechoice inhibition 12.5% block			Reward omission inhibition 50% block			Reward omission inhibition 12.5% block			Risky win inhibition 50% block			Certain win inhibition 50% block		
	Odds Ratios	CI	<i>p</i>	Odds Ratios	CI	<i>p</i>	Odds Ratios	CI	<i>p</i>	Odds Ratios	CI	<i>p</i>	Odds Ratios	CI	<i>p</i>	Odds Ratios	CI	<i>p</i>
Treatment [baseline (reference) vs laser inhibition]	0.31	0.17–0.56	<0.001	3.12	1.83–5.33	<0.001	0.37	0.22–0.61	<0.001	4.62	2.82–7.56	<0.001	0.58	0.32–1.05	0.073	0.98	0.54–1.78	0.947
Treatment × outcome (win→win)	0.65	0.26–1.62	0.350				0.59	0.22–1.58	0.293				0.94	0.36–2.50	0.907	0.48	0.18–1.26	0.135
Treatment × outcome (loss→win)	0.49	0.20–1.22	0.125	0.43	0.11–1.67	0.224	1.86	0.73–4.70	0.192	0.65	0.17–2.48	0.530	1.40	0.57–3.42	0.462	0.43	0.17–1.11	0.083
Treatment × outcome (certain→win)	0.97	0.41–2.29	0.948	0.22	0.06–0.76	0.017	1.69	0.61–4.65	0.313	0.45	0.16–1.26	0.129	0.54	0.21–1.40	0.204	0.72	0.28–1.86	0.492
Treatment × outcome (win→loss)	0.46	0.20–1.05	0.067	1.18	0.31–4.48	0.810	1.15	0.49–2.72	0.743	0.31	0.10–0.98	0.047	0.89	0.38–2.09	0.791	0.50	0.20–1.26	0.141
Treatment × outcome (certain→loss)	0.65	0.27–1.60	0.350	0.51	0.26–1.02	0.057	2.08	0.81–5.31	0.127	0.42	0.22–0.80	0.008	0.64	0.26–1.59	0.338	0.35	0.13–0.91	0.032
Treatment × outcome (win→certain)	1.11	0.45–2.71	0.824	0.39	0.11–1.40	0.150	0.83	0.29–2.35	0.722	0.28	0.09–0.87	0.028	0.31	0.11–0.87	0.026	0.60	0.21–1.66	0.323
Treatment × outcome (loss→certain)	1.90	0.81–4.44	0.139	0.37	0.19–0.74	0.005	1.65	0.68–3.97	0.268	0.41	0.22–0.78	0.007	0.92	0.37–2.32	0.865	0.84	0.33–2.18	0.725
Treatment × outcome (certain→certain)	0.92	0.41–2.04	0.836	0.38	0.20–0.71	0.002	1.17	0.50–2.70	0.719	0.37	0.21–0.66	0.001	0.29	0.11–0.79	0.016	0.40	0.16–0.98	0.044
Residual Variance σ^2	3.29			3.29			3.29			3.29			3.29			3.29		
Intercept ₀	0.67	Subject		0.65	Subject		0.95	Subject		0.68	Subject		0.35	Subject		0.83	Subject	
Intraclass Correlation Coefficient	0.17			0.16			0.22			0.17			0.10			0.20		
<i>N</i>	26	Subject		26	Subject		26	Subject		26	Subject		25	Subject		27	Subject	
Observations	3034			3026			3063			3047			2985			3115		
Marginal R^2 /conditional R^2	0.210/0.344			0.049/0.205			0.139/0.332			0.082/0.238			0.153/0.235			0.125/0.301		

Values for main effects of outcome predictors were included in the model, but are not shown for clarity. Boldface *p* values highlight statistically significant interactions terms.

epochs also increased the tendency to follow two small/certain choices with another certain choice ($p = 0.005$).

PL inhibition during ITI

Suppression of PL activity that was time-locked to each of the choice outcome events tested here altered risk/reward decision-making in different ways. To ascertain whether these effects were specifically because of disruption of activity that coincided with these events, another experiment was conducted to inhibit PL activity during a randomized time point within the intertrial interval (6–14 s after the end of each trial). Notably, this did not alter choice in any way (main effect of treatment, $F_{(1,22)} = 1.32$, $p = 0.26$; treatment × probability block interaction, $F_{(1,22)} = 0.008$, $p = 0.93$; Fig. 7A,B), nor were there changes in other performance measures (latency, $F_{(1,22)} = 4.08$, $p = 0.06$; Fig. 7A, inset; omissions, baseline = 0.37 ± 0.12 ; laser = 0.33 ± 0.14 ; $F_{(1,22)} = 0.18$, $p = 0.68$). These lacks of effects are in keeping with reports that temporally specific inhibition of either the dopamine system or BLA→accumbens circuitry during the intertrial interval also did not alter probabilistic discounting or risky decision-making, although these manipulations did affect choice when delivered around the time when rewards/punishments were or were not received (Stopper et al., 2014; Orsini et al., 2017; Bercovici et al., 2018). The null effect in this experiment, combined with the alterations in choice observed when PL silencing coincided with choice outcomes, highlights that activity occurring in close temporal proximity to rewarded and nonrewarded actions exert a more discernable impact on the direction of subsequent choice, compared with activity that occurs some time after outcomes are realized.

PL inhibition during reward magnitude discrimination

Inhibiting PL activity before choice caused near-indiscriminate patterns of choice as reward probabilities changed over the session. One potential interpretation of these findings is that this may reflect nonspecific impairments in discriminating between larger versus smaller rewards or between the two levers. To control for this, a separate group of rats that received PL infusions of eArchT were trained on a simpler reward magnitude discrimination task, where they chose between a larger four-pellet and smaller one-pellet reward, both delivered with 100% certainty.

After 3 d of training, all rats showed a strong bias for the large reward option across blocks of trials and continued to do so for the duration of testing. Previous work by our group has shown that under these conditions, choice behavior remains goal directed rather than habitual as reinforcer and contingency devaluation are effective at reducing preference for the large reward (Stopper et al., 2014). In these animals, we suppressed PL activity before choice in a manner identical to the procedures used in rats trained on the probabilistic task. This had no effects on choice (treatment, $F_{(1,5)} = 2.91$, $p = 0.15$; treatment × block interaction, $F_{(1,5)} = 1.21$, $p = 0.34$; Fig. 7C). In addition, although this manipulation increased choice latencies and trial omissions in rats making risk/reward decisions, under these simpler task conditions, neither of these performance variables was altered (latency, $F_{(1,5)} = 0.77$, $p = 0.48$; Fig. 7C, inset; omissions, baseline = 0.1 ± 0.1 ; laser = 1.3 ± 1.3 ; $F_{(1,5)} = 0.89$, $p = 0.42$).

After this first test, these rats were retrained and then retested, wherein we silenced PL activity during receipt of the large reward, which in this experiment, occurred on nearly every free-choice trial. This also did not alter choice or performance measures (choice, treatment effect, $F_{(1,5)} = 0.0005$, $p = 0.98$; treatment × block interaction, $F_{(1,5)} = 1.25$, $p = 0.33$, Fig. 7D; latency, $F_{(1,5)} = 0.19$, $p = 0.85$; Fig. 7D, inset; omissions, baseline = 0.1 ± 0 ; laser = 0.1 ± 0.1 ; $F_{(1,5)} = 1.00$, $p = 0.36$). Given that administration of laser light into the PL did not affect any aspects of behavior during reward magnitude discrimination in rats expressing eArchT, we did not test the effects of laser light in rats expressing eYFP.

Collectively, the findings from these control experiments indicate that the effects of PL silencing during risk/reward decision-making are unlikely to be attributable to impairments in motivational or discrimination processes. Relatedly, the lack of effect of prechoice silencing indicate that alterations in probabilistic discounting induced by this manipulation were not driven by nonspecific perturbations in movement that may have altered the approach toward the levers. Rather, they suggest that PL activity during these different task events is essential for guiding action selection in situations requiring integration of reward history to bias choice between rewards of different magnitudes and probabilities.

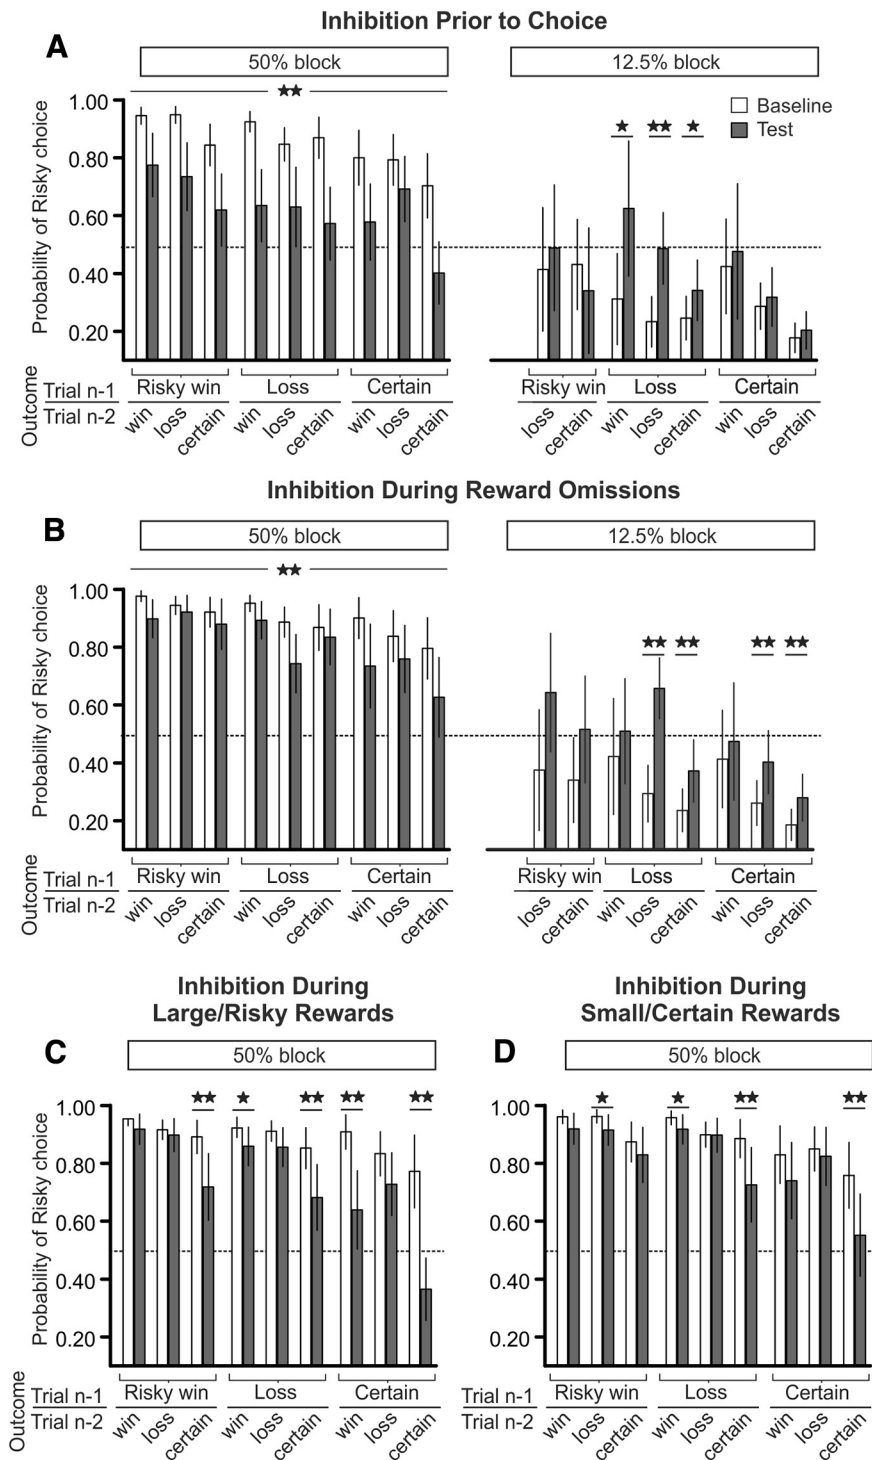


Figure 6. Results of the multilevel logistic regression analyses examining alterations in risky choice induced by PL inhibition during different task events varied as a function of outcome history. Data are presented as means and 95% confidence intervals of the probability of a risky choice after a combination of risky wins, losses, and small/certain outcomes occurring one and two trials before a particular choice. **A**, Reductions in risky choice in the 50% block induced by prechoice PL inhibition did not vary as a function of outcome history (left, stars denote $p < 0.001$ main effect of treatment), whereas increases in risky choice in the 12.5% block occurred after losses on the preceding trial (right). **B**, PL inhibition during reward omissions reduced in risky choice in the 50% block, but this did not vary as a function of outcome history (left, stars denote $p < 0.001$ main effect of treatment), whereas in the 12.5% block, this manipulation increased in risky choice when the large reward was not received within the preceding two trials (right). **C, D**, PL inhibition during **C**, large/risky wins, or small/certain rewards (**D**) reduced risky choice in the 50% block following many but not all the various outcome combinations. Stars and double stars denote $p < 0.05$ and $p < 0.001$ compared with baseline.

Discussion

The present findings provide novel insight into how temporally discrete patterns of PL neural activity, occurring during different phases of the decision-making sequence, shape choice biases during risk/reward decision-making. Activity before choices or during the evaluation of their different outcomes plays a multifaceted role in promoting more profitable decisions. During periods before choice, PL activity promotes choices of higher-utility options. In contrast, activity occurring after nonrewarded actions aids in evaluating losses within the context of reward history. Additionally, activity linked to larger or smaller rewards appears to serve as a value comparator for different reward options, promoting bias toward larger rewards when they are more likely to be received.

PL activity before action selection promotes optimal choice

PL inhibition during deliberation reduced bias for more-preferred/higher-utility options, reducing or increasing risky choice when the odds of winning were comparatively high (50%) or low (12.5%), rendering animals more ambivalent toward either option. This also disrupted how information about the outcomes of preceding decisions guided subsequent choice as animals displayed near-random patterns of reward/negative feedback sensitivity in response to the most recent risky choice outcomes. In addition, prechoice PL activity promotes the timely implementation of decision policies as suppressing it increased both deliberation times and choice omissions. Admittedly, we did not video record behavior, making it difficult to disentangle if these latency effects were attributable specifically to increased hesitation in making decisions or other factors, such as displacement from the levers that may have slowed response times. In this regard, it is important to note that disruptions in risk/reward decision-making or associated choice latencies and omissions were not attributable to nonspecific discrimination/motivational impairments or movement artifacts as PL inhibition before choice on a reward magnitude discrimination had no effect on action selection or other performance measures. Thus, during deliberation phases of decision-making, PL activity is integral to promoting more profitable and timely choices in more cognitively demanding situations requiring integration of recent outcome history and context.

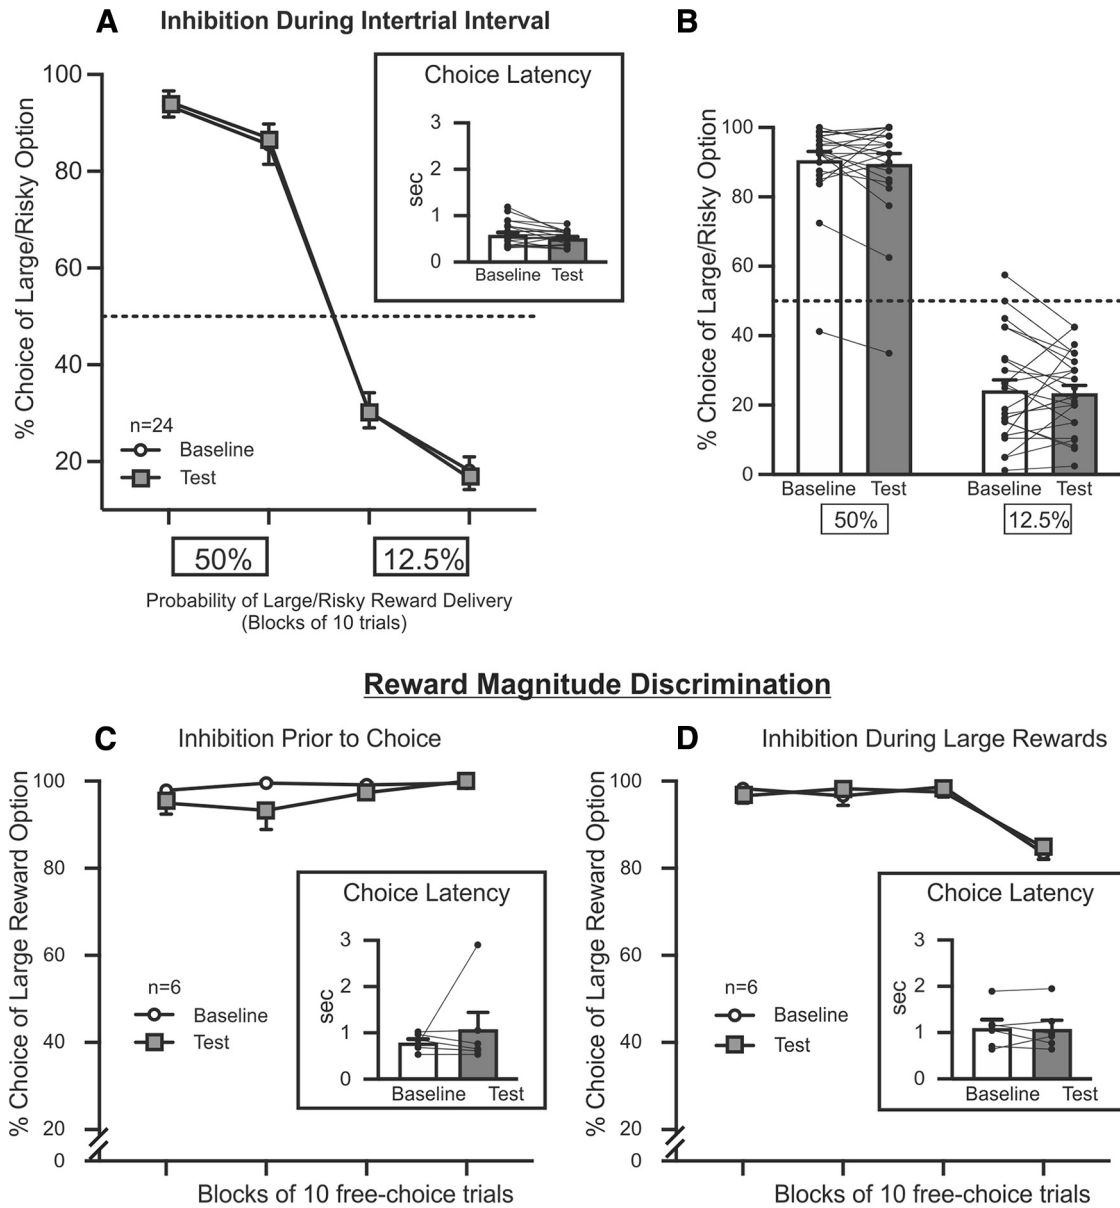


Figure 7. Intertrial interval and reward magnitude discrimination control experiments. **A**, PL inhibition at pseudorandom epochs during the intertrial interval did not alter probabilistic discounting in animals expressing eArchT ($n = 24$) or choice latencies (inset). **B**, Individual risky choice data comparing baseline to optogenetic inhibition during the intertrial interval averaged across probability blocks. **C**, Inhibition of PL before choice on a reward magnitude discrimination ($n = 6$) did not affect preference for the large reward, nor did it affect response latencies (inset). **D**, PL inhibition occurring during receipt of the large reward also did not alter choice preference or latencies (inset).

The present data complement imaging studies demonstrating increased ACC activation during decision-making deliberation that may represent preferred choice (Rogers et al., 2004; Gläscher et al., 2009; Wittmann et al., 2016). Moreover, they align with neurophysiological findings that populations of PL neurons display phasic increases in activity before initiating a particular course of action. These changes in firing tracks preferred choices during cost/benefit decision-making, which may contribute to action monitoring and encoding relative value (Bari et al., 2019; Braunscheidel et al., 2019; Sackett et al., 2019; Choi et al., 2021). In this regard, similar alterations in risk/reward decision-making are induced by suppression of prechoice activity in either the dopamine system (Stopper et al., 2014) or BLA→nucleus accumbens circuitry (Bercovici et al., 2018). The similarities of these effects across different brain systems highlights that the PL operates as part of a broader network to establish choice preferences

and then acts on them. Around the time of decision implementation, near-simultaneous activity within PL, BLA, and dopamine projection pathways work congruently to evaluate which options may be deemed better and increase the likelihood that actions are biased toward them.

PL activity during nonrewarded outcomes frames losses in context

PL activity associated with losses was also integral to guiding subsequent decisions, depending on the estimated profitability of the risky option. Suppressing this activity increased risky choice when reward probabilities were low (12.5%), in keeping with previous reports (Passecker et al., 2019). Yet, this unexpectedly reduced bias toward the larger reward when the objective utility of this option was higher (50%). Unlike prechoice silencing, which led to near-random win-stay and lose-shift patterns,

suppressing activity during loss evaluation had a more prominent effect on how losses affected subsequent choice, which was shaped by the context of the relative profitability of the risky option. Under baseline conditions, animals tended to disregard risky losses in the 50% block as it was still more profitable to choose risky in this context. PL silencing enhanced sensitivity to these occasional losses, increasing lose-shift behavior independent of recent outcome history. Conversely, more frequent losses occurring in the 12.5% block led to more frequent small/certain choices. In this context, PL silencing had the opposite effect, reducing the impact of losses over a broader reward history (i.e., at least one to two choices back). These differential effects provide novel insight into how PL activity integrates information about nonrewarded actions to guide risk/reward decisions, highlighting that these signals do not necessarily encode a more rudimentary negative reward prediction error. Rather, PL activity, coincidental with unrewarded actions, provides top-down control that places losses in context, biasing their impact relative to the estimated profitability of different options, and determining whether they should be attended to (play it safe next time) or disregarded (keep playing risky) given the recent context of reward history.

These findings complement human imaging data demonstrating the importance of ACC activity linked to negative outcomes during other forms of complex decision-making that require monitoring of action-outcome contingencies to support flexible choice preferences (Gläscher et al., 2009; Camille et al., 2011). Likewise, there have been multiple reports that populations of mPFC neurons display differential changes in activity after unrewarded choices during probabilistic reversal or set-shifting tasks (Del Arco et al., 2017; Passecker et al., 2019; Choi et al., 2021; Spellman et al., 2021). As a comparison, reward omission-related activity within other circuits appears to signal a more uniform negative prediction error. For example, interfering with either phasic reductions in dopamine activity or increases in BLA→accumbens activity during reward omissions increases risky choice and perturbs other forms of associative learning (Steinberg et al., 2013; Stopper et al., 2014; Bercovici et al., 2018; Fischbach and Janak, 2019). PL activity during reward omissions appears to subservise a more complex function. Rather than highlighting all unrewarded actions as something to be avoided, it operates more holistically to evaluate losses in context to either minimize their impact or promote modifications in decision biases.

Reward-related PL signals

Reward-associated PL activity also influenced subsequent action selection. Optogenetic inhibition time locked to receipt of larger rewards reduced preference for this option but only when their delivery was probabilistic and their likelihood was high. Here, animals were less likely to follow a rewarded risky choice with another one but also showed amplified lose-shift tendencies after recent losses, although here, activity was unperturbed after nonrewarded choices. These results complement primate neurophysiological findings revealing that ACC encoding of rewarded actions is maintained across multiple trials, and its computation of reward value is updated following each subsequent rewarded action (Seo and Lee, 2007; Donahue et al., 2013). This dovetails with the results of our multilevel modeling, unveiling that PL activity during these epochs promotes choice persistence toward larger, higher probability rewards when risky choices paid off one to two trials back. Thus, PL activity linked to uncertain rewards exerts a broad impact on their perceived value, which

spans a longer choice history, augmenting the impact of rewarded actions while also minimizing the impact of the recent nonrewarded ones. These signals may work in concert with reward-related dopamine activity as inhibiting these signals during receipt of larger/risky rewards induced similar changes in choice and feedback sensitivity (Stopper et al., 2014). It is possible that in situations where rewards are uncertain, neural computations by the PL may provide top-down influence that shape dopamine neuron reward-prediction errors (Jo et al., 2013).

Somewhat more surprising was the finding that inhibiting PL activity associated with small/certain rewards also made animals less likely to chase larger/riskier ones and more likely to choose the less-profitable certain option repeatedly. Thus, PL activity associated with smaller rewards in this context may reduce the tendency to direct actions toward these options when riskier ones may be more profitable. In this regard, subpopulations of PL cells display changes in firing linked to small/certain rewards that predict subsequent risky choices, and inhibiting this activity decreased risk seeking (Passecker et al., 2019). This combination of findings demonstrates that PL activity associated with smaller/certain rewards exerts a discernable influence on the incentive for pursuing larger/uncertain ones. One interpretation for these seemingly counterintuitive findings is that smaller rewards trigger patterns of activity within a population of PL neurons that encode a form of comparator signal used by other PFC-related networks to evaluate how much better alternative rewards may be. When larger rewards are more likely, these small reward signals may act as a reminder that the alternative is preferable, augmenting its attractiveness in future choice situations. Conversely, by perturbing comparator signals related to the stable objective value of small/certain reward, decisions may become more stochastic and less likely to be directed toward larger ones. In essence, the absence of information of the value of one option may hinder evaluation of how much better the other one may be. This idea is in keeping with findings that fluctuations in persistent, postoutcome PL activity encode the relative value of different options in mice performing a probabilistic foraging task (Bari et al., 2019). Viewed from a broader perspective, these data reveal that PL reward-related activity exerts a multimodal influence that promotes optimal risky choices, both increasing the likelihood that rewarded risky choices are made again as well as providing a relative value signal that further enhances incentive for larger/risky rewards.

There have been reports that termination of photoinhibition with Arch can be associated with brief (~1 s) and moderate (~20%, 1 Hz) increases in cortical neuron firing above baseline (Li et al., 2019). Thus, it is possible that the effects of PL silencing during choice outcomes may have been driven partially by rebound excitation, which in our experiments would have occurred ~7 s after outcomes were experienced. That said, photoinhibition during the ITI, at time points further removed from choice outcomes, had no effect on action selection. This indicates that perturbations in neural activity coincidental with choice outcomes have a much greater influence on subsequent choices, and any rebound excitation that may occur when photoinhibition was terminated would have less, if any, impact.

The present findings that both prechoice and reward-related PL activity pivotally influence choice direction differs from those of a recent examination of how PL suppression alters decisions using a seemingly similar task (Passecker et al., 2019). In that study, stimulating PL GABA cells before choice or on receipt of large/risky rewards did not significantly alter behavior. Aside

from certain methodological differences (inactivation approach, within- vs between-subject designs, statistical power), two key analytical and procedural differences are of particular relevance. First, our analyses focused more on isolating risky choice as a function of probability, making it better suited to identify more nuanced roles for PL during different phases of the decision process. A second issue pertains to different task parameters and how reward probabilities were varied. The task structure used by Passecker et al. (2019) varied reward probabilities across three blocks within a session, with block order randomized across sessions. During a particular session, rats inferred reward probability within a block based on outcome history on that given day. In contrast, our task used a more stable structure so that our well-trained rats had a firmer representation of the reward probability context at the start of a session that could aid in guiding choice. In this regard, the ACC has been implicated in using reward context to shape decision policies (Kolling et al., 2014). What is novel here is that when contrasting these two experiments, our data reveal that the PL may play different roles in guiding choice in stochastic situations versus more stable ones, where context more heavily guides goal-directed actions.

Summary and conclusions

Collectively, the present findings provide novel insight into the temporal dynamics through which PL shapes risk/reward decision-making, highlighting how activity during deliberation and outcome-evaluation phases differentially influence action selection. It is likely that the distinct functions that emerge during different phases of the decision process are mediated by separate populations of PL neurons, which may be distinguished in part by their projection targets (St Onge et al., 2012; Jenni et al., 2017). For example, subpopulations of PL cells projecting to anterior versus posterior striatum appear to encode negative versus positive outcomes (Choi et al., 2021). Thus, clarifying how distinct PL ensembles guide choice during different phases of the decision sequence remains an important direction for future studies. Additionally, a limitation of the present study is that only male subjects were used. Previous studies have shown mixed results in terms of baseline differences in probabilistic discounting in male versus female rats (Braunscheidel et al., 2019; Islas-Preciado et al., 2020), but sex differences have been shown in rats assessed on other decision-making tasks (Orsini and Setlow, 2017; Orsini et al., 2017; Pellman et al., 2017). Given this, it will be important to assess how PL activity associated with different decision phases influences choice in females. It is also notable that the constellation of effects reported here differed from how pharmacological inactivation of this region alters this form of decision-making. Complete suppression of all PL activity during probabilistic discounting impaired flexible updating of choice biases (St Onge and Floresco, 2010). In comparison, the use of more targeted perturbations in PL signaling revealed a plethora of functions embedded in different phasic events that promote advantageous choices before decision implementation and provide important context when evaluating decision outcomes that shape future choices. As such, these findings underscore how this approach can yield a more comprehensive understanding of the neural underpinnings of complex forms of cognition and behavior (Orsini et al., 2019).

References

Bari BA, Grossman CD, Lubin EE, Rajagopalan AE, Cressy JI, Cohen JY (2019) Stable representations of decision variables for flexible behavior. *Neuron* 103:922–933.e7.

- Bercovici DA, Princz-Lebel O, Moorman DE, Floresco SB, Tse MT (2018) Optogenetic dissection of temporal dynamics of amygdala-striatal interplay during risk/reward decision making. *eNeuro* 5:ENEURO.0422-18.2018.
- Braunscheidel KM, Okas MP, Hoffman M, Mulholland PJ, Floresco SB, Woodward JJ (2019) The abused inhalant toluene impairs medial prefrontal cortex activity and risk/reward decision-making during a probabilistic discounting task. *J Neurosci* 39:9207–9220.
- Camille N, Griffiths CA, Vo K, Fellows LK, Kable JW (2011) Ventromedial frontal lobe damage disrupts value maximization in humans. *J Neurosci* 31:7527–7532.
- Choi K, Piasini E, Cifuentes-Vargas L, Díaz-Hernández E, Henderson NT, Subramanian M, Gerfen CR, Fuccillo MV (2021) Distributed processing for action control by prelimbic circuits targeting anterior-posterior dorsal striatal subregions. *bioRxiv* 469698. <https://doi.org/10.1101/2021.12.01.469698>.
- Christopoulos GI, Tobler PN, Bossaerts P, Dolan RJ, Schultz W (2009) Neural correlates of value, risk, and risk aversion contributing to decision making under risk. *J Neurosci* 29:12574–12583.
- Clark L, Bechara A, Damasio H, Aitken MRF, Sahakian BJ, Robbins TW (2008) Differential effects of insular and ventromedial prefrontal cortex lesions on risky decision-making. *Brain* 131:1311–1322.
- de Visser L, Homberg JR, Mitsogiannis M, Zeeb FD, Rivalan M, Fitoussi A, Galhardo V, van den Bos R, Winstanley CA, Dellsu-Hagedorn F (2011) Rodent versions of the Iowa gambling task: opportunities and challenges for the understanding of decision-making. *Front Neurosci* 5:109.
- Del Arco A, Park J, Wood J, Kim Y, Moghaddam B (2017) Adaptive encoding of outcome prediction by prefrontal cortex ensembles supports behavioral flexibility. *J Neurosci* 37:8363–8373.
- Donahue CH, Seo H, Lee D (2013) Cortical signals for rewarded actions and strategic exploration. *Neuron* 80:223–234.
- Fischbach S, Janak PH (2019) Decreases in cue reward seeking after reward-paired inhibition of mesolimbic dopamine. *Neuroscience* 412:259–269.
- Gläscher J, Hampton AN, O'Doherty JP (2009) Determining a role for ventromedial prefrontal cortex in encoding action-based value signals during reward-related decision making. *Cereb Cortex* 19:483–495.
- Gläscher J, Adolphs R, Damasio H, Bechara A, Rudrauf D, Calamia M, Paul LK, Tranel D (2012) Lesion mapping of cognitive control and value-based decision making in the prefrontal cortex. *Proc Natl Acad Sci U S A* 109:14681–14686.
- Heilbronner SR, Rodriguez-Romaguera J, Quirk GJ, Groenewegen HJ, Haber SN (2016) Circuit-based corticostriatal homologies between rat and primate. *Biol Psychiatry* 80:509–521.
- Islas-Preciado D, Wainwright SR, Sniegocki J, Lieblich SE, Yagi S, Floresco SB, Galea LA (2020) Risk-based decision making in rats: modulation by sex and amphetamine. *Horm Behav* 125:104815.
- Jo YS, Lee J, Mizumori SJ (2013) Effects of prefrontal cortical inactivation on neural activity in the ventral tegmental area. *J Neurosci* 33:8159–8171.
- Jenni NL, Larkin JD, Floresco SB (2017) Prefrontal dopamine D1 and D2 receptors regulate dissociable aspects of decision-making via distinct ventral striatal and amygdalar circuits. *J Neurosci* 37:6200–6213.
- Kolling N, Wittmann M, Rushworth MFS (2014) Multiple neural mechanisms of decision making and their competition under changing risk pressure. *Neuron* 81:1190–1202.
- Kolling N, Wittmann MK, Behrens TEJ, Boorman ED, Mars RB, Rushworth MFS (2016) Value, search, persistence, and model updating in anterior cingulate cortex. *Nat Neurosci* 19:1280–1285.
- Kuhnen CM, Knutson B (2005) The neural basis of financial risk taking. *Neuron* 47:763–770.
- Lafferty CK, Britt JP (2020) Off-target influences of arch-mediated axon terminal inhibition on network activity and behavior. *Front Neural Circuits* 14:10.
- Li N, Chen S, Guo ZV, Chen H, Huo Y, Inagaki HK, Chen G, Davis C, Hansel D, Guo C, Svoboda K (2019) Spatiotemporal constraints on optogenetic inactivation in cortical circuits. *Elife* 18:8:e48622.
- Mahn M, Prigge M, Ron S, Levy R, Yizhar O (2016) Biophysical constraints of optogenetic inhibition at presynaptic terminals. *Nat Neurosci* 19:554–556.
- Orsini CA, Setlow B (2017) Sex differences in animal models of decision making. *J Neurosci Res* 95:260–269.
- Orsini CA, Hernandez CM, Singhal S, Kelly KB, Frazier CJ, Bizon JL, Setlow B (2017) Optogenetic inhibition reveals distinct roles for basolateral

- amygdala activity at discrete time points during risky decision making. *J Neurosci* 37:11537–11548.
- Orsini CA, Heshmati SC, Garman TS, Wall SC, Bizon JL, Setlow B (2018) Contributions of medial prefrontal cortex to decision making involving risk of punishment. *Neuropharmacology* 139:205–216.
- Orsini CA, Hernandez CM, Bizon JL, Setlow B (2019) Deconstructing value-based decision making via temporally selective manipulation of neural activity: insights from rodent models. *Cogn Affect Behav Neurosci* 19:459–476.
- Paine TA, Hara AO, Plaut B, Lowes DC (2015) Effects of disrupting medial prefrontal cortex GABA transmission on decision-making in a rodent gambling task. *Psychopharmacology (Berl)* 232:1755–1765.
- Passecker J, Mikus N, Malagon-Vina H, Anner P, Dimidschstein J, Fishell G, Dorffner G, Klausberger T (2019) Activity of prefrontal neurons predict future choices during gambling. *Neuron* 101:152–164.e7.
- Paxinos G, Watson C (2005) *The rat brain in stereotaxic coordinates*. Amsterdam: Academic.
- Pellman BA, Schuessler BP, Tellakat M, Kim JJ (2017) Sexually dimorphic risk mitigation strategies in rats. *eNeuro* 4:ENEURO.0288-16.2017.
- Pujara MS, Wolf RC, Baskaya MK, Koenigs M (2015) Ventromedial prefrontal cortex damage alters relative risk tolerance for prospective gains and losses. *Neuropsychologia* 79:70–75.
- Rogers RD, Ramnani N, Mackay C, Wilson JL, Jezzard P, Carter CS, Smith SM (2004) Distinct portions of anterior cingulate cortex and medial prefrontal cortex are activated by reward processing in separable phases of decision-making cognition. *Biol Psychiatry* 55:594–602.
- Sackett DA, Moschak TM, Carelli RM (2019) Prelimbic cortical neurons track preferred reward value and reflect impulsive choice during delay discounting behavior. *J Neurosci* 39:3108–3118.
- Seo H, Lee D (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci* 27:8366–8377.
- Spellman T, Svei M, Kaminsky J, Manzano-Nieves G, Liston C (2021) Prefrontal deep projection neurons enable cognitive flexibility via persistent feedback monitoring. *Cell* 184:2750–2766.e17.
- Steinberg EE, Keiflin R, Boivin JR, Witten IB, Deisseroth K, Janak PH (2013) A causal link between prediction errors, dopamine neurons and learning. *Nat Neurosci* 16:966–973.
- St Onge JR, Stopper CM, Zahm DS, Floresco SB (2012) Separate prefrontal-subcortical circuits mediate different components of risk-based decision making. *J Neurosci* 32:2886–2899.
- St Onge JR, Floresco SB (2010) Prefrontal cortical contribution to risk-based decision making. *Cereb Cortex* 20:1816–1828.
- Stopper CM, Tse MTL, Montes DR, Wiedman CR, Floresco SB (2014) Overriding phasic dopamine signals redirects action selection during risk/reward decision making. *Neuron* 84:177–189.
- Sul JH, Kim H, Huh N, Lee D, Jung MW (2010) Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron* 66:449–460.
- Van Eden CG, Lamme VAF, Uylings HBM (1992) Heterotopic cortical afferents to the medial prefrontal cortex in the rat. A combined retrograde and anterograde tracer Study. *Eur J Neurosci* 4:77–97.
- van Holstein M, Floresco SB (2020) Dissociable roles for the ventral and dorsal medial prefrontal cortex in cue-guided risk/reward decision making. *Neuropsychopharmacology* 45:683–693.
- Wiegert JS, Mahn M, Prigge M, Printz Y, Yizhar O (2017) Silencing neurons: tools, applications, and experimental constraints. *Neuron* 95:504–529.
- Wittmann MK, Kolling N, Akaishi R, Chau BKH, Brown JW, Nelissen N, Rushworth MFS (2016) Predictive decision making driven by multiple time-linked reward representations in the anterior cingulate cortex. *Nat Commun* 7:12327.
- Zeeb FD, Baarendse PJJ, Vanderschuren LJMJ, Winstanley CA (2015) Inactivation of the prelimbic or infralimbic cortex impairs decision-making in the rat gambling task. *Psychopharmacology (Berl)* 232:4481–4491.