UNIVERSITY OF CALIFORNIA SAN DIEGO

Role of Retrosplenial Cortex in Value-Based Decision Making

A thesis submitted in partial satisfaction of the requirements for degree Master of Science

in

Biology

by

Zeljana Babic

Committee in charge:

Professor Takaki Komiyama, Chair
Professor Cory Root
Professor Stefan Leutgeb

2020

The Thesis of Zeljana Babic is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

Chair

University of California San Diego

2020

TABLE OF CONTENTS

Signature Page…….. iii

Table of Contents…….. iv

List of Figures…… v

Acknowledgments…… vi

Abstract of the Thesis ……...……... vii

Introduction…….. 1

Methods…….. 8

Experiment I…….....…….... 8

Experiment II…….…….…… 14

Experiment III…….…….…… 15

Results……...…… 16

Experiment I…….……..…….. 16

Experiment II…….……...…… 17

Experiment III…….……..…….. 23

Discussion……..……..…… 24

References……27

LIST OF FIGURES

Figure 1: Acute Inactivation of RSC, but Not Its Chronic Lesion, Impairs Reward History-Based Strategy……..……18

Figure 2: Acute Inactivation of RSC and ACC Projections to RSC Impair Reward History-Based Strategy, while inactivation of V1 projections to RSC show no effect…… 19

Figure 3: Immunostaining of pAAV-CaMKlla-eNpHR 3.0-EYFP Viral Expression in Hippocampal Areas.……..…….……23

ACKNOWLEDGEMENTS

I would like to thank Dr. Takaki Komiyama for his support and guidance over the past three years - through all the highs and lows.

I would also like to acknowledge my immediate mentor Bethanny Dunskin whose day-to-day support proved invaluable too many times to count, amplified my understanding and love for science, and ultimately inspired me to take the Ph.D. neuroscience route as the next step in my academic career.

I would like to thank Ryoma Hattori for his mentorship and help throughout the years, as well as Nicole Mynaryk for being a kind and compassionate colleague and collaborator.

Parts of the introduction that describe Experiment I is a rewrite of the material as it appears in Area-Specificity and Plasticity of History-Dependent Value Coding During Learning, 2019, Ryoma Hattori, Bethanny Danskin, Zeljana Babic, Nicole Mlynaryk, and Takaki Komiyama published in *Cell* 2019. The thesis author was one of the authors of this publication.

Parts of the method section that describe Experiment I is a rewrite of the material as it appears in Area-Specificity and Plasticity of History-Dependent Value Coding During Learning, 2019, Ryoma Hattori, Bethanny Danskin, Zeljana Babic, Nicole Mlynaryk, and Takaki Komiyama published in *Cell* 2019. The thesis author was one of the authors of this publication.

Parts of the discussion section that describe Experiment I is a rewrite of the material as it appears in Area-Specificity and Plasticity of History-Dependent Value Coding During Learning, 2019, Ryoma Hattori, Bethanny Danskin, Zeljana Babic, Nicole Mlynaryk, and Takaki Komiyama published in *Cell* 2019. The thesis author was one of the authors of this publication.

ABSTRACT OF THE THESIS

Role of Retrosplenial Cortex in Value-Based Decision Making

by

Zeljana Babic

Master of Science in Biology

University of California San Diego, 2020

Professor Takaki Komiyama, Chair

This Master's thesis attempts to assess the question of what the role of the retrosplenial cortex (RSC) is in foraging task - a task in which a mouse is supposed to make a value-based decision that uses the internal representation of value, rather than an external sensory cue or perceptual input. This paper also attempts to find out what the effect of inactivation of the surrounding areas' projections to RSC is on the same task. To test these questions, three experiments have been conducted - Experiment I utilized both optogenetic inactivation and NMDA lesions of RSC, and the other two experiments examined the effects of inactivation of

the projections from RSC, ACC, and V1 to RSC (Experiment II) and projections from CA1 and subiculum to RSC (Experiment III).

Collected data were analyzed in three ways: (1) basic behavioral metrics such as reward rate and choice probability; (2) logistic regression behavioral model that estimates how the mouse integrates recent experience to make a decision on the upcoming trials; (3) fitting behavior data to a reinforcement learning algorithm which models the internal decision making process. The first two experiments have been completed - showing that acute optogenetic inactivation of RSC selectively impaired the reward-history-based strategy, while also revealing an involvement of ACC projections to RSC in this cognitive process. Experiment III has not been completed due to the Covid-19 pandemic.

INTRODUCTION

Decision making represents assigning subjective value to given options - a value formed through past experiences (Hattori et al., 2018). In a value-based choice task, mice show normal sensitivity to changes in reward value and use that reward value to guide their actions (Fila al., 2018). Furthermore, value-based decision making is sometimes also called economic decision making because it involves weighing up different beneficial alternatives to maximize payoff. For instance, when weighing different costs associated with different rewards, decisions require the integration of information about reward size and action/outcome history for guessing the probability of obtaining a reward (Krasheninnikova 2018).

One of the areas implicated in many cognitive functions across the literature, including value-based decision making, is retrosplenial cortex (RSC). RSC is an essential node in the systemic integration network. For example, studies point to its role in learning that involves spatial stimuli and navigation (Radwanska al., 2010), while it has also emerged as a critical brain area involved in memory, including episodic and topographical memory in humans, as well as spatial memory in rodents (Milczarek et al., 2018). Thus, in Hattori et al.'s study we published in 2019, we decided to assess the extent to which RSC is involved in value-based decision making.

In Hattori et al.'s study (2019), we asked two main questions about the neural encoding of value: how does value encoding differ across cortical areas, and does their encoding of history information change throughout the learning of the task? First, the potential differences across different brain areas were tested in terms of their value coding. We reasoned that the brain must stably maintain value information to be able to retrieve it anytime as needed. We also asked

whether a particular area encodes value as a persistent population activity pattern that spans the entire period between one choice and the next.

Second, we addressed the dynamics of history encoding over weeks of task learning. We hypothesized that, whereas animals learn to perform value-based decision making task, specific areas preferentially enhance their encoding of choice-outcome history. We addressed these questions in mice performing a foraging task based on history-dependent value.

Two-photon calcium imaging was conducted during the performance of foraging task in 6 dorsal cortical areas: anterior-lateral motor cortex (ALM), posterior premotor cortex (pM2), posterior parietal cortex (PPC), retrosplenial cortex (RSC), the primary somatosensory cortex (S1), and primary visual cortex (V1). Although we found that all six cortical areas significantly encoded history- and value-related information, direct comparisons of population activity across areas revealed area specificity in the encoding of history- and value-related information, and in the plasticity of history coding during task learning. In particular, we found that RSC uniquely and potently encodes value-related information in a persistent population activity pattern, and preferentially extends history coding during learning. These results highlight RSC as a critical region for decision making based on history-dependent value.

This study also sought to compare neural ensemble activity across cortical areas during decision making driven by history-dependent value. To quantify the history dependency of their choices, we fit the behavior with a logistic regression model:

$$logit(P_L(t)) = \sum_{i=1}^{5} \beta_{RewC(t-i)} * RewC(t-i) + \sum_{i=1}^{5} \beta_{UnrC(t-i)} * UnrC(t-i) + \sum_{i=1}^{5} \beta_{C(t-i)} * C(t-i) + \beta_0$$
$$+ \sum_{i=1}^{5} o\beta_{RewC(t-i)} * oRewC(t-i) + \sum_{i=1}^{5} o\beta_{UnrC(t-i)} * oUnrC(t-i) + \sum_{i=1}^{5} o\beta_{C(t-i)} * oC(t-i) + o\beta_0$$

Information about past rewarded choices and unrewarded choices are essential for adaptive behavior in this task. This regression analysis revealed a dominant increase in the influence of reward choice history from previous trials on decision making during learning. This result indicates that mice improved the reward rate in the task by learning to adjust their strategy such that they preferentially chose the option that was more frequently rewarded in the recent trials.

We also attempted to study neural correlates of history and value signals during this history-dependent, value-based decision making task, and we performed two-photon calcium imaging in task-performing mice. We used CaMKIIa-tTA::tetO-GCaMP6s double transgenic mice that express GCaMP6s widely in cortical excitatory neurons. We adopted a surgical preparation that exposed most dorsal cortical areas for optical access, which allowed us to image multiple cortical areas in each mouse. We focused on neurons in layer 2/3, and fluorescence signals from individual neurons were deconvolved to remove fluorescence decay and estimate spiking activity before all analyses. We first focused on expert sessions (defined as training day > 15 and RL index > 0.08) and imaged 6 dorsal cortical areas: ALM (n = 6,721 neurons from 13 fields in 7 mice), pM2 (n = 9,759 neurons from 17 fields in 10 mice), PPC (n = 7,703 neurons from 16 fields in 11 mice), RSC (n = 10,481 neurons from 17 fields in 9 mice), S1 (n = 7,576 neurons from 14 fields in 7 mice), and V1 (n = 2,767 neurons from 6 fields in 4 mice). We included only one image field for each cortical area per hemisphere per animal for analysis. In each behavioral session, Dr. Hattori simultaneously imaged hundreds of neurons (542.3 ± 122.0

cells; mean ± SD) in one of the six cortical areas. Behavioral performance was equivalent in the sessions used to image each area. Each area contained many neurons with activity aligned to task epochs. Individual neurons showed heterogeneous activity patterns, and their peak timing tiled the entire trial period and inter-trial intervals in all six areas. Many of these task-related neurons were tuned to specific task events. In addition to such activity modulations based on ongoing task events, the activity of many neurons was modulated by past events. To analyze how these historical events modulated population activity in each area, decoding analysis was performed to classify various history and current events via population activity of 200 neurons from each area.

The results above indicated that history information is differentially represented across cortical areas in expert animals. This raises the possibility that these areas differentially encode value related variables, given that value in this task is a specific computational product of various types of history. We focused on three types of value-related variables:

1. Value difference between the two options (DQ) which is the main driver of choice in this task and directly updated by rewarded and unrewarded choice history.
2. Value of upcoming choice (Qch), which could contribute to whether the mouse chooses the high-value option ("exploitive" behavior) or the low-value option ("explorative" behavior) in each trial.
3. The sum of the values for the two options (SQ), which might relate to the motivational state of the mouse.

We performed a multiple linear regression for the activity of each neuron with these three value-related variables. In particular, RSC stood out as the area with the highest fractions of neurons encoding value related variables among the six cortical areas.

Consistent with the population analysis above, we were able to identify individual RSC neurons whose activity was persistently modulated according to the DQ and Qch values. In contrast, neurons in the other areas that encoded DQ and Qch in the ready period did not consistently encode the value signals outside the ready period, in line. These results highlight RSC as a unique cortical area where value-related signals are potently and encoded explicitly in a persistent population activity pattern across trials.

Furthermore, the other cortical areas did not show a correlation between behavioral strategy and population encoding. Altogether, these results indicate that RSC flexibly encodes history information depending on the ongoing behavioral strategy and strongly encodes the specific history information required for that ongoing behavior.

In addition to findings presented above, literature shows that the retrosplenial cortex is also involved in visuospatial integration and spatial learning and that RSC neurons exhibit discrete, place cell-like sequential activity that resembles the population code of space in the hippocampus (Mao al., 2017). Using calcium imaging, Mao and colleagues (2017) found that a subpopulation of neurons in the mouse RSC expresses "a sparse, orthogonal and continuous representation of a linear environment," very similar to CA1 place cell population. This study inspired Experiment III in this thesis. Although our project did not assess spatial memory in the foraging task - because there is evidence of functional similarities between these hippocampus and RSC, we wanted to further assess if hippocampal projections to RSC - specifically CA1 and subiculum - would also extend to value-based decision making.

Another study that supported the idea that hippocampal projections to RSC could possibly affect value-based decision making in our task was done by Masuda et al. (2020). This

lab wanted to explore the exact mechanisms the hippocampus uses to support decision making when the length of a delay and size of a rewarding change. Because previous goal-directed behavior studies indicated that some CA1 neurons fire when animals "approach, wait for or acquire rewards", Masuda and colleagues recorded CA1 neuronal activity in mice during delay-based decision making in an automated T-maze task while independently manipulating delay length and reward size across a session. They found that distinct populations of CA1 neurons increased or decreased their firing during the delay of a reward. Additionally, they found that dissociable subpopulations of hippocampal neurons represent delay and reward information in opposing ways (Masuda, 2020). These findings could be relevant to our foraging task - as we introduce a delay of ~2 seconds between the signal and a potential reward.

We also thought hippocampal projections to RSC could affect foraging behavior because the hippocampus is in charge of the retrieval of short term memory. Studies also show that RSC appears to be a significant consolidation area for hippocampal-dependent memories, particularly contextual memories (Vedder al., 2016). Although in our task, we were primarily interested in value encoding of RSC as opposed to the acquisition of the task, it would be hard to imagine that history guided decision making processes would not be affected by the ability to retrieve short-term memory quickly. An interesting study that explains how hippocampal retrieval of short term memories work was done by Roy and colleagues (2017). They posed a question of whether hippocampus retrieves episodic memory by reinstating cortical activity patterns that occur during learning. This assumption is not new in the neuroscience field - it has been one of the main theories regarding hippocampal function for decades. Previous studies also indicate that hippocampal and cortical neurons are reactivated after learning during memory retrieval and

sleep. However, previous studies failed to explain if the hippocampus is required to induce reactivation in the cortex. Roy and colleagues attempted to test this idea and found that when CA1 neurons were silenced during retrieval, reactivation was disrupted in areas such as subiculum and retrosplenial cortex. They found that while disruption was selective, activity appeared normal in cortical neurons not engaged during learning. Thus, they found that when dorsal CA1 cannot retrieve a specific context memory, representations in anatomically connected cortical regions cannot be reactivated. This finding indicates that hippocampal inputs to these cortical areas - most importantly, RSC in terms of our study - fail to completely activate representations established during learning (Tanaka al., 2014).

Additionally, another study that presented evidence that hippocampal projections could affect the performance on the foraging task was done by Roy and colleagues in 2017. They found subiculum to be a crucial brain area for the formation and retrieval of episodic memory. They also found that the dorsal subiculum and CA1-to-dorsal subiculum circuit play an essential role for the selective retrieval of episodic memories. In other words, Roy et al. (2017) found evidence that the subiculum-containing loop is necessary for the recall of rapid memory updating, providing another reason for this thesis to also assess the effects of subiculum projections to RSC, in addition to CA1-RSC projections.

Given how RSC is well suited to contribute to the consolidation process, being connected to both the hippocampus and neocortical areas, and is necessary for both recent and remote memory retrieval (Sousa al., 2018), the retrosplenial cortex (RSC) appears to be a major component of the brain's memory and navigation systems. One of the areas that are very densely connected with RSC is the anterior cingulate cortex (ACC). For this reason, we decided to also

examine the effect that inactivation of the projections from ACC to RSC would have on the foraging task. Additionally, given that the RSC is an important hub for visual-spatial information, we decided to also inactivate the primary visual area (V1) of the cerebral cortex and examine what effects on the value-based decision making would that inactivation show.

Therefore, this thesis will present three separate experiments. Experiment I is essentially data published in Hattori et al. (2019), and it examines what effects channelrhodopsin2 guided optogenetic inactivation of RSC, as well as lesions of RSC, have on value-based decision making. The other two experiments will examine the effects of inactivating ACC and V1 projections to RSC (Experiment II) and CA1 and subiculum projections to RSC (Experiment III) are.

Parts of the introduction that describe Experiment I is a rewrite of the material as it appears in Area-Specificity and Plasticity of History-Dependent Value Coding During Learning, 2019, Ryoma Hattori, Bethanny Danskin, Zeljana Babic, Nicole Mlynaryk, and Takaki Komiyama published in *Cell* 2019. The thesis author was one of the authors of this publication.

METHODS

Experiment I - Inactivating RSC using channelrhodopsin2 (ChR2)

Surgery for imaging and optogenetics

Before surgery, adult mice were injected with dexamethasone (2 mg/kg) subcutaneously. Animals were continuously anesthetized with 1%–2% isoflurane during the procedure. After cleaning the surface of the dorsal skull with a razor blade, saline was applied, which made the skull become transparent enough to visualize vasculature patterns. Stereotactic coordinates of

vasculature patterns were recorded through an intact skull. Collected information was used to identify imaging regions under the two-photon microscope. A craniotomy with a changeable window size (ranging from a small circular window with 2 mm diameter for imaging a single cortical area/mouse to a large hexagonal window for imaging multiple cortical areas/mouse) was performed on each mouse. The dura was left unimpaired. A glass window was secured on the edges of the remaining skull using 3M Vetbond (WPI), followed by cyanoacrylate glue and dental acrylic cement (Lang Dental). The largest glass windows used in this study were made by cutting a coverslip rectangular glass and further cutting two frontal edges to make a hexagonal window. After the glass implantation, a custom-built metal head-bar was secured on the skull above the cerebellum with cyanoacrylate glue and dental acrylic cement. Buprenorphine (0.1 mg/kg of body weight) and Baytril (10 mg/kg of body weight) were subcutaneously injected after surgery, and mice were monitored until they recovered from anesthesia.

Behavior

Following a minimum of five days of recovery upon the procedure, mice were water-restricted at 1-2 mL/day. Following a period of at least a week of water restriction, behavioral training began. Behavioral control was automated with a real-time system running on Linux, communicating with MATLAB. During behavioral sessions, two lickports were placed to the left and right sides of the head-fixed mouse, and IR beams tracked licking. An amber LED was used as a ready cue (5mm diameter, placed 5cm away from nose), and a speaker was used as a go cue (10 kHz tone). Each test began with a ready period (2 or 2.5 s with LED light), followed by an answer period with an auditory go cue. The go cue was terminated when mice made a choice (the first lick to one of the two ports) or when the answer period exceeded the maximum

duration of 2 seconds. Each decision triggered a 50 ms feedback tone (left, 5 kHz; right, 15 kHz). In rewarded trials, 2.5 ml of water reward was given immediately after the choice. Before starting the dynamic foraging task, mice went through three phases of pre-training under head-fixation. In the primary stage of pre-training (2-3 days), mice were compensated for every left or right choice during the answer period, with a 100% reward probability. At this stage, licking during the ready period was not punished. The mean intertrial interval (ITI) steadily increased from 1 s to 6 s. In the following phase of pre-training (1-3 days), mice were trained in another task. This time, the reward was delivered alternating from left and right lickports following either choice (first lick during the answer period). Beginning with this training stage, licking during the ready period was punished by 500 ms white noise alarm and trial abort with an extra 2 s ITI.

In the third stage of pre-training (1-2 weeks), mice were required to alternate the choices between left and right on every choice trial. Through the three phases of pre-training, mice learned the general task structure. Only their first lick during the answer period is associated with outcome, and mice need to withhold licking during the ready period. After the pre-training, mice started training in the dynamic foraging assignment where the reward was probabilistic. Inter-trial intervals ranged randomly between 5-7 s. In a test in which mice made a choice, 2.5 ml of water reward was delivered immediately after a reward had been assigned to the lickport on trial. The reward was assigned at each lickport on every choice trial with a specific reward assignment probability for the lickport. Once a reward was assigned to a lickport, the reward was maintained until it was chosen. The combinations of reward assignment probabilities were either [60%, 10%] or [52.5%, 17.5%] per trial, and reward assignment probability switched randomly

every 60-80 trials in the order of [Left, Right] = [60%, 10%], [10%, 60%], [52.5%, 17.5%], [17.5%, 52.5%], [60%, 10%]. The probability switch was postponed if the fraction of choosing the lickport with higher reward assignment probability was below 50% in the recent 60 trials - until the fraction reached at least 50%. Trials in which mice licked during the ready period ('alarm trials') and the trials in which mice failed to lick during the answer period ('miss trials') were not rewarded. Alarm trials were not included, and neither were the miss trials in neural activity analyses. Expert sessions were defined as the sessions in which the mice have been trained for at least 15 days in the dynamic foraging task, and the reinforcement learning (RL) index was above 0.08 for the session. Expert mice typically completed about 600 trials per session (711.5 ± 203.5 trials / session, 4.97 ± 3.29% alarm rate and 4.51 ± 4.61% miss rate in experts (mean ± SD)).

Optogenetic inactivation

In order to activate PV-positive inhibitory neurons in RSC of PV-Cre:LSL-ChR2 double transgenic mice using optogenetics, elliptical illumination patterns were generated with a DLP projector (Optoma X600 XGA). A single-lens reflex (SLR) lens (Nikon, 50 mm, f/1.4D, AF) was coupled with two achromatic doublets (Thorlabs, AC508-150-A-ML, f = 150 mm; Thorlabs, AC508-075-A-ML, f = 75 mm) as to shrink and focus illumination patterns on RSC. A dichroic mirror (Thorlabs, DMLP490L) and a blue filter (Thorlabs, FESH0450) were installed between the two achromatic pairs and after the second achromatic doublet, respectively, to pass only blue light (400-450 nm). Illumination patterns were generated with Psychtoolbox in MATLAB (http://psychtoolbox.org/). In RSC inactivation trials, a 2 mm-0.5 mm ellipse was concentrated on RSC in each hemisphere (Center at 0.3 mm lateral and 2 mm posterior to bregma). The

cumulative light intensity was equivalent between RSC inactivation trials and control (light over head bar) trials. The patterns at 30 Hz were projected as a sequence of square pulses from the onset of the ready period until the choice was made. There was a linear attenuation in intensity over the last 100 ms. The intensity at the focus ranged between 2.5-6 mW/mm2. The frequency of RSC inactivation trials was set within a session to either 15% (12 sessions) or 5% (3 sessions). A constraint was added - each RSC inactivation trial had to be followed by at least three head bar trials to avoid unnecessary perturbation of reinforcement learning. For four mice, the RSC was inactivated through a glass window, and for one mouse, it was done through the skull. The skull for the through-skull inactivation was made semi-transparent by covering the dorsal skull surface with a cyanoacrylate glue layer.

Lesion experiment

Twelve adult mice were trained to perform the foraging task, and after at least seven days of stable performance, underwent excitotoxic-lesion or sham lesion (saline) surgery. Stable, expert performance for this task was determined to be choice prediction accuracy of 65% with a standard RL model (Equation 1 and Equation 2) in at least six sessions during the seven days; these sessions also met the 0.08 RL index criterion for imaging mice in at least six sessions during the seven days. Mice were anesthetized with 1%–2% isoflurane during surgery. Three burr-hole craniotomies per hemisphere (6 total) were opened on the dorsal skull over RSC. A tapered glass pipette was inserted to perform the cortical microinjection. Injection sites relative to Bregma were the following: AP = 1.6, 2.3, 3.0, ML = ± 0.3, ± 0.35, ± 0.4, and DV = 0.4 from the dura surface in all sites. The injection concentration was of 50 nL/site of either NMDA in sterile saline (20 mg/ml or 10 mg/ml; Sigma) at a rate of 0.05-0.1 ml/min. After injection, the

pipette was left for 5 min to ensure diffusion of the solution. Buprenorphine (0.1 mg/kg of body weight) and Baytril (10 mg/kg of body weight) were subcutaneously injected after surgery. Mice resumed the behavioral task following a day after the injections. The surgeon and the experimenter for the behavior were blinded to the identity of the substance that was injected. They became unblinded only after the last day of data collection. Of the 12 mice, 5 received saline, seven received NMDA. One of the NMDA mice was excluded due to small and off-target lesions, as quantified by histology. Brains of lesion and saline mice were collected at 21-25 days post-injection. In order to quantify the lesion size, 50 mm-thick coronal sections were prepared with a microtome (Thermo Fisher Scientific) and blocked with 10% goat serum, 1% bovine serum albumin, and 0.3% Triton X-100 in PBS. Immunostaining was then performed with anti-NeuN primary antibody (1:400; Mouse, Millipore) and anti-mouse Alexa Fluor 488 secondary antibody (1:1000; Goat, Thermo Fisher Scientific). Both missing areas and areas that lacked NeuN-positive neurons were considered lesioned. Images of coronal sections with RSC and the corresponding brain atlas were superimposed to quantify the % of the lesion within RSC.

Reinforcement learning model

"In a standard RL model, Equation 1 describes action value for the chosen option and is updated:

$$Q_{ch}(t+1) = Q_{ch}(t) + \alpha * (R(t) - Q_{ch}(t))$$

where Qch(t) is the subjective action value of the chosen option on trial t, R(t) is the reward outcome on trial t (1 if rewarded, 0 if unrewarded), and alpha is the learning rate.

The probability of choosing left on trial t is estimated by a softmax function as follows in Equation 2:

$$P_L(t) = \frac{1}{1 + e^{-\beta_{\Delta Q}(Q_L(t) - Q_R(t))}}$$

where $Q_L(t)$ and $Q_R(t)$ are the action values of the left and right choices on trial t respectively, and $\beta_{\Delta Q}$ defines the sensitivity of decision making to the value difference" (Hattori et al., 2019).

Experiment II - Inactivating projections to RSC using halorhodopsin (RSC, ACC, and V1)

Surgery

Surgeries that followed the steps from Experiment I surgeries were performed on six adult 6 wild-type (WT) mice. However, in Experiment II, animals were also injected with pAAV-CaMKlla-eNpHR 3.0-EYFP virus containing halorhodopsin. Each animal was injected bilaterally with a virus in one target area only - RSC, ACC, or V1 cortex (two animals per site). Injection sites relative to Bregma chosen based on Allen Brain Atlas were the following: for RSC (AP = -1.7, -2.2, -2.7; ML = ± 0.3; DV 0.4.) and for ACC (AP = +0.4, +0; ML = ± 0.35; DV = 1.0, 0.8 (two injection depths per site)). The injection concentration was of 40 nL/site of either at a rate of 0.05-0.1 ml/min. V1 coordinates were relative to Lambda (AP/ML: 0/+-2.4, +.4/+-2.4, 0/+-2.75. DV 0.4 for all sites), with the concentration of 50 nl per site. Post-surgery recovery, water restriction, and foraging training were the same as described in Experiment I.

Behavior

After animals reached the expert stage, mice were subjected to RSC optogenetic inactivation - alternating between the intertrial interval (ITI) and the ready-answer period (RA). Control sessions were performed for both ITI and RA - with the light shining directly over the headbar, and with a black tape placed over the cranial window to prevent any cortical

inactivation. Laser for Experiment two was 532 nm (green) laser. Intensity at the surface of RSC was between 8-10 mW/mm2. Unfortunately, due to the pandemic, only three of the six animals completed the experiment.

Experiment III - Inactivating projections to RSC using halorhodopsin (CA1 and subiculum)

Experiment III was meant to use the exact same methods as Experiment II, except that the target of the viral injections was the CA1 region of hypothalamus in two mice and the subiculum in the other two mice (n = 4). Bilateral injection sites relative to Bregma were chosen based on Allen Brain Atlas: CA1 (AP = -1.7, -2.2, -2.5; ML = ± 1.3, ±2.1, ±2.3; and DV = 1.25, 1.4, 1.4) and subiculum (AP = -3.3; ML = ± 1.6; and DV = 1.6). The injection concentration was 50 nL/site at a rate of 0.05-0.1 ml/min. However, due to the pandemic, this experiment did not pass the surgery stage. The four adult male mice were left to recover for two months, after which they were sacrificed, and the brains were obtained for the histology processing.

Parts of the method section that describe Experiment I is a rewrite of the material as it appears in Area-Specificity and Plasticity of History-Dependent Value Coding During Learning, 2019, Ryoma Hattori, Bethanny Danskin, Zeljana Babic, Nicole Mlynaryk, and Takaki Komiyama published in *Cell* 2019. The thesis author was one of the authors of this publication.

RESULTS

Experiment I

The retrosplenial cortex is required for maintaining a reward-history-based strategy. Optogenetic inactivation of RSC was performed to examine whether RSC is necessary for task performance. Mice used were PV-Cre:LSL-ChR2 double transgenic mice that express channelrhodopsin2 in parvalbumin-positive inhibitory neurons. RSC was bilaterally inactivated from the onset of the ready period until the choice in a subset of trials, by delivering blue light to activate inhibitory neurons. Given that RSC is rostrocaudally elongated, elliptical illumination patterns were generated with a DLP projector (Figures 1A and 1B). In control trials, the light was directed onto the head bar, away from RSC. It was discovered that RSC inactivation decreased the probability of repeating the same action after rewarded trials (‘‘win-stay’’).

Furthermore, the RSC inactivation reduced the likelihood of changing action after unrewarded trials (‘‘lose-switch’’) (Figure 1C). Therefore, it is suggested that the RL strategy is impaired in tests with the RSC inactivation. Regression analysis revealed that the RSC inactivation selectively impaired behavioral dependency on rewarded and unrewarded choice history while preserving the dependence on the outcome- independent choice history and choice bias (Figures 1D–1F). These results indicate that the activity of RSC is required for the reward-history-based behavioral strategy. The imaging results and previous studies suggested widespread encoding of history- and value-related information across areas. To test the possibility that other regions can compensate for the function of RSC when it is removed altogether, bilateral lesions of RSC were performed by injecting N- Methyl-D-aspartate (NMDA) to induce excitotoxicity (Figure 1G). In distinction, to acute optogenetic inactivation, the RSC

lesion did not affect behavioral performance in subsequent sessions (Figures 1H and 1I). These results suggest that other areas can compensate for the chronic loss of RSC, despite RSC being involved in the reward history-based strategy in control mice. It is worth noting that the results do not exclude the possibility that the remaining neurons in RSC that endured the lesion were enough to support the behavior.

Experiment II

Experiment II assessed the effect of ACC, V1 and RSC projections on RSC. Immunostaining is demonstrated in Figure 2A. Two mice with virus injected in ACC showed a similar ITI effect to RSC, but not a similar RA effect. This result indicated that inputs from ACC to RSC mediated the maintenance of reward information, especially distant history reward information during the ITI. However, RSC is downstream of ACC. This could explain why ACC inputs to RSC are not necessary during RA. Nonetheless, these are only preliminary results given a small number of sessions and two mice only (Figure 2A).

Behavioral data of mice with injection in RSC were generally consistent with other RSC effects from Experiment I, but with a smaller effect size, which is expected with limited number of mice and sessions (Figure 2B). Finally, we also tested projections of V1 region to RSC because of its relevance for the behavior due to the visual stimulus in foraging task, but we saw no effect on value-based decision making (Figure 2D).

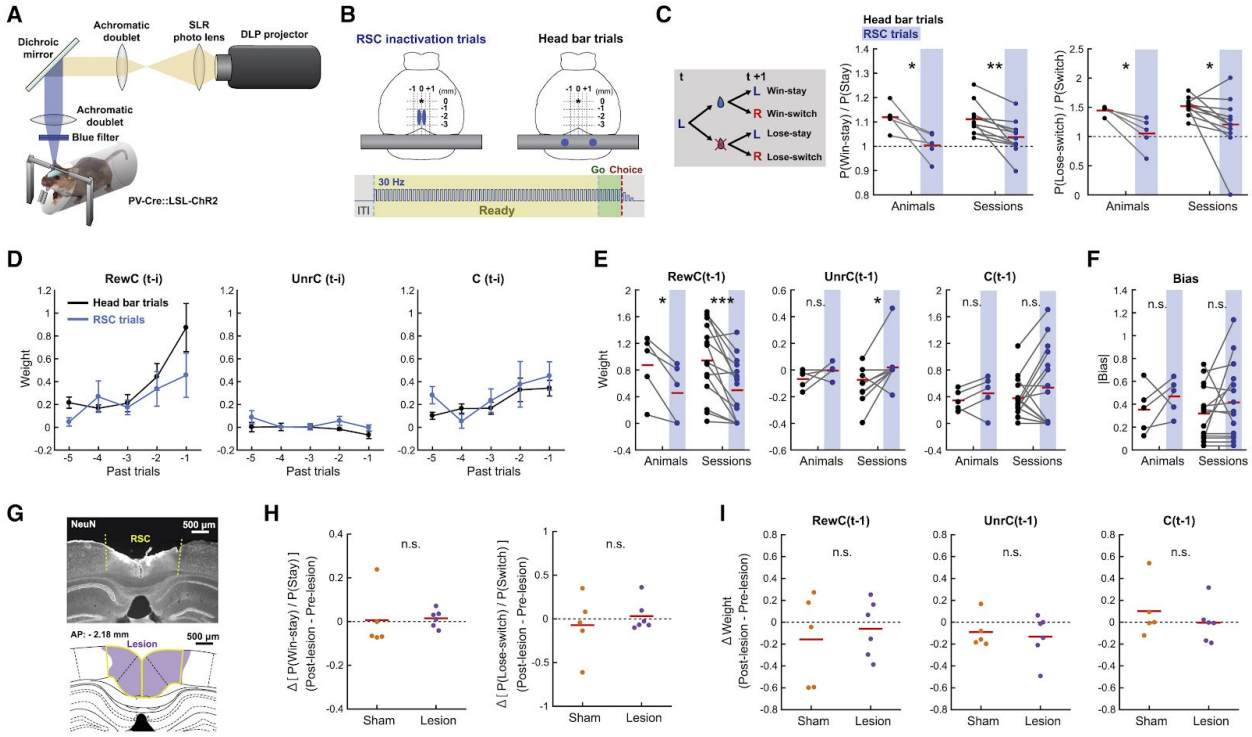

Figure 1: Acute Inactivation of RSC, but Not Its Chronic Lesion, Impairs Reward History-Based Strategy. (A) Schematic of the projector-based optical stimulation system. Patterned light is resized and focused on cortex to optogenetically activate parvalbumin-positive inhibitory neurons. (B) RSC was bilaterally inactivated in a small subset of trials within a session (5% or 15% of trials). In all other trials, the head bar was illuminated with the same light intensity and area. Elliptic illumination patterns were used for RSC inactivation trials to cover rostro-caudally elongated RSC. The illumination was applied from the onset of ready period until the choice at 30 Hz with a linear attenuation in the intensity after choice. (C) Effects of RSC inactivation on the win-stay and lose-switch probabilities (left: n = 5 mice; right: n = 12 sessions). Red line indicates the mean of each condition. Only successive choice trials were used to derive the probabilities. P(Win-stay) was normalized by the overall stay probability (the average of P(Win-stay) and P(Lose-stay)). Similarly, P(Lose-switch) was normalized by the overall switch probability (the average of P(Win-switch) and P(Lose-switch)). For the n = animals plots (left), all sessions from each mouse were pooled to calculate the probabilities. For the n = sessions plots (right), only pairs from the 15% inactivation sessions were included. RSC inactivation made the stay and switch probabilities less dependent on the reward outcomes from the 1 trials. Paired t test. (D) Behavioral dependency on rewarded choice (RewC(t-i)), unrewarded choice (UnrC(t-i)), and outcome-independent choice (C(t-i)) history in head bar trials and RSC inactivation trials. (E) Effects of RSC inactivation on behavioral dependency on the 3 types of history from 1 trial (left: n = 5 mice; right: n = 15 sessions). Pairs of head bar trials (black) and RSC inactivation trials (blue) are shown. Red lines indicate the means. RSC inactivation reduced behavioral dependency on choice-reward history, especially for the rewarded choice history. Wilcoxon signed-rank test was used for non-normally distributed UnrC(t-1) weights of n = sessions, and paired t test was used for the other comparisons. (F)

Effect of RSC inactivation to choice bias (left: n = 5mice; right: n = 15 sessions). The absolute value of bias is shown for pairs of head bar trials (black) and RSC inactivation trials (blue). Red line indicates the mean of each condition. Paired t test. The sign of the bias was also generally unaffected by inactivation (not shown). (G) (Top) Example coronal section from a mouse with lesioned RSC. The section is stained with NeuN to visualize the presence of neurons. RSC largely lacks NeuN-positive neurons. Dashed lines indicate the borders of RSC. (Bottom) A corresponding brain atlas is shown. Yellow lines outline RSC. Purple shading indicates a lesioned area. (H) Effects of RSC lesion on win-stay and lose-switch probabilities (sham: n = 5 mice; lesion: n = 6 mice). Difference between the mean of 7 sessions before sham or lesion and the mean of 7 sessions after sham or lesion is shown. Red lines indicate the means. Two-sided t test. (I) Effects of RSC lesion on behavioral dependency on the 3 types of history from 1 trial. Two-sided t test. *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$. (Hattori et al., 2019).

Figure 2: Acute Inactivation of RSC and ACC Projections to RSC Impair Reward History-Based Strategy, while inactivation of V1 projections to RSC show no effect.

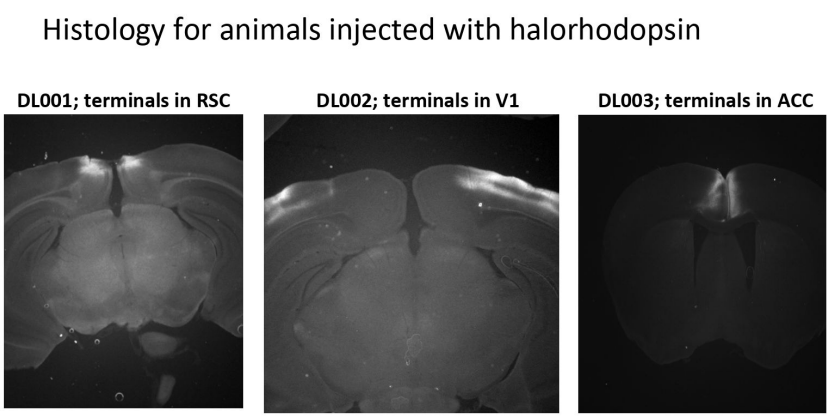

(A) Immunostaining demonstrating the injections sites.

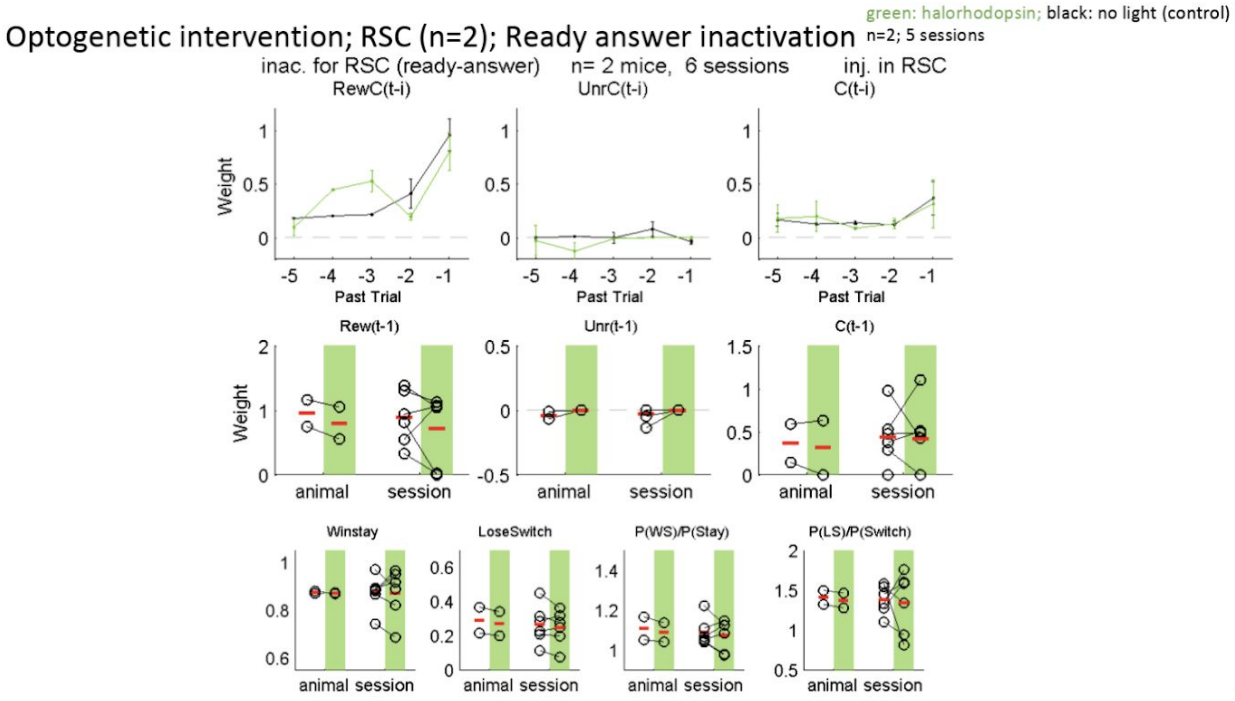

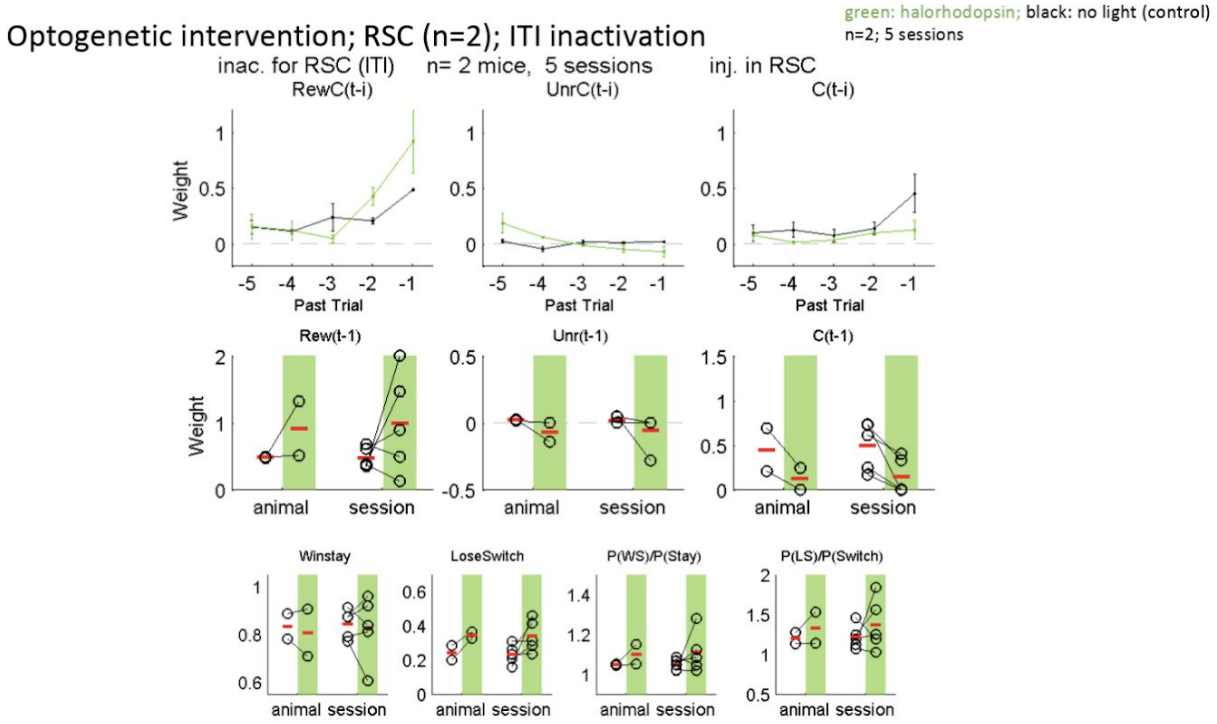

(B) Results from RSC inactivation with RSC injection sites are consistent with results from Experiment I.

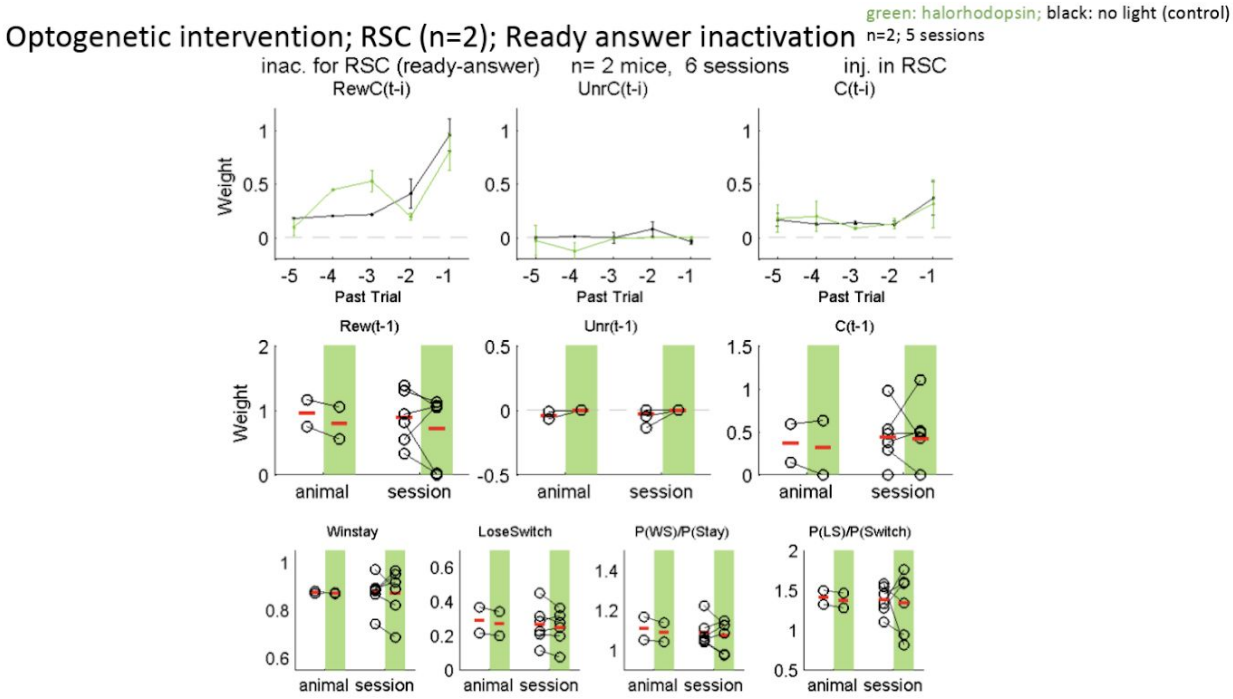

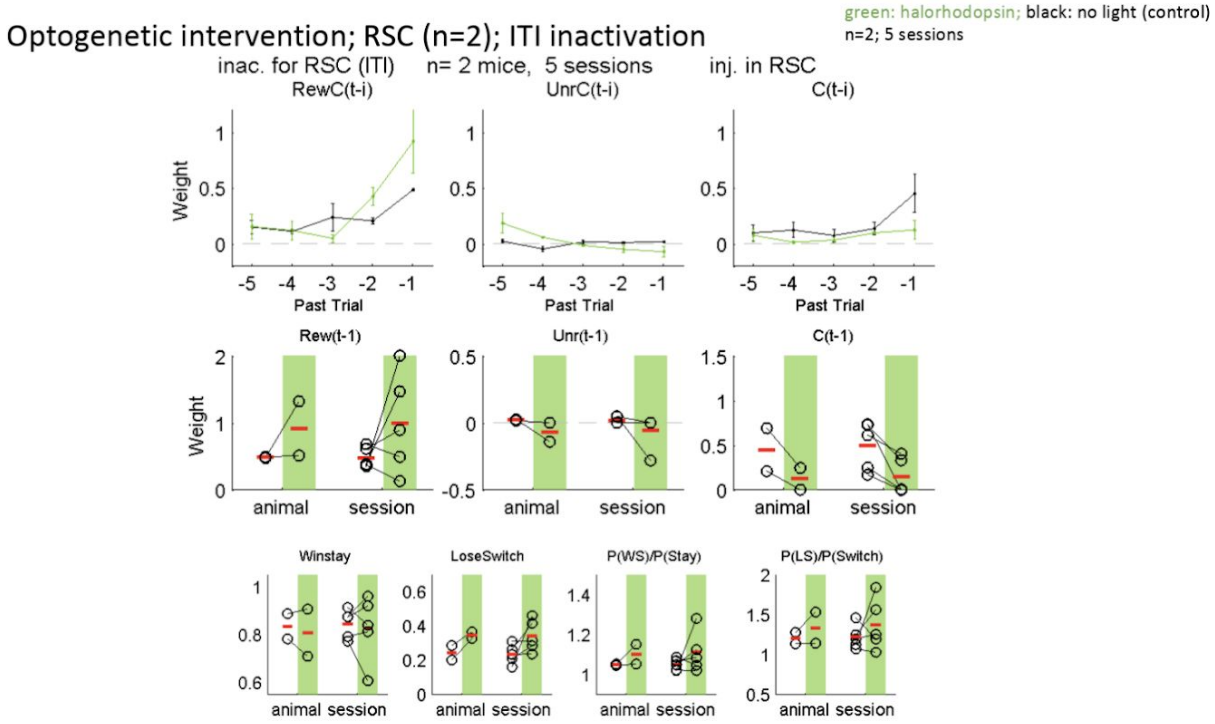

(C) RSC inactivation with ACC injection sites show ITI effects similar to ITI effects of RSC injection sites, but different RA effects compared to RS effects of RSC injection sites.

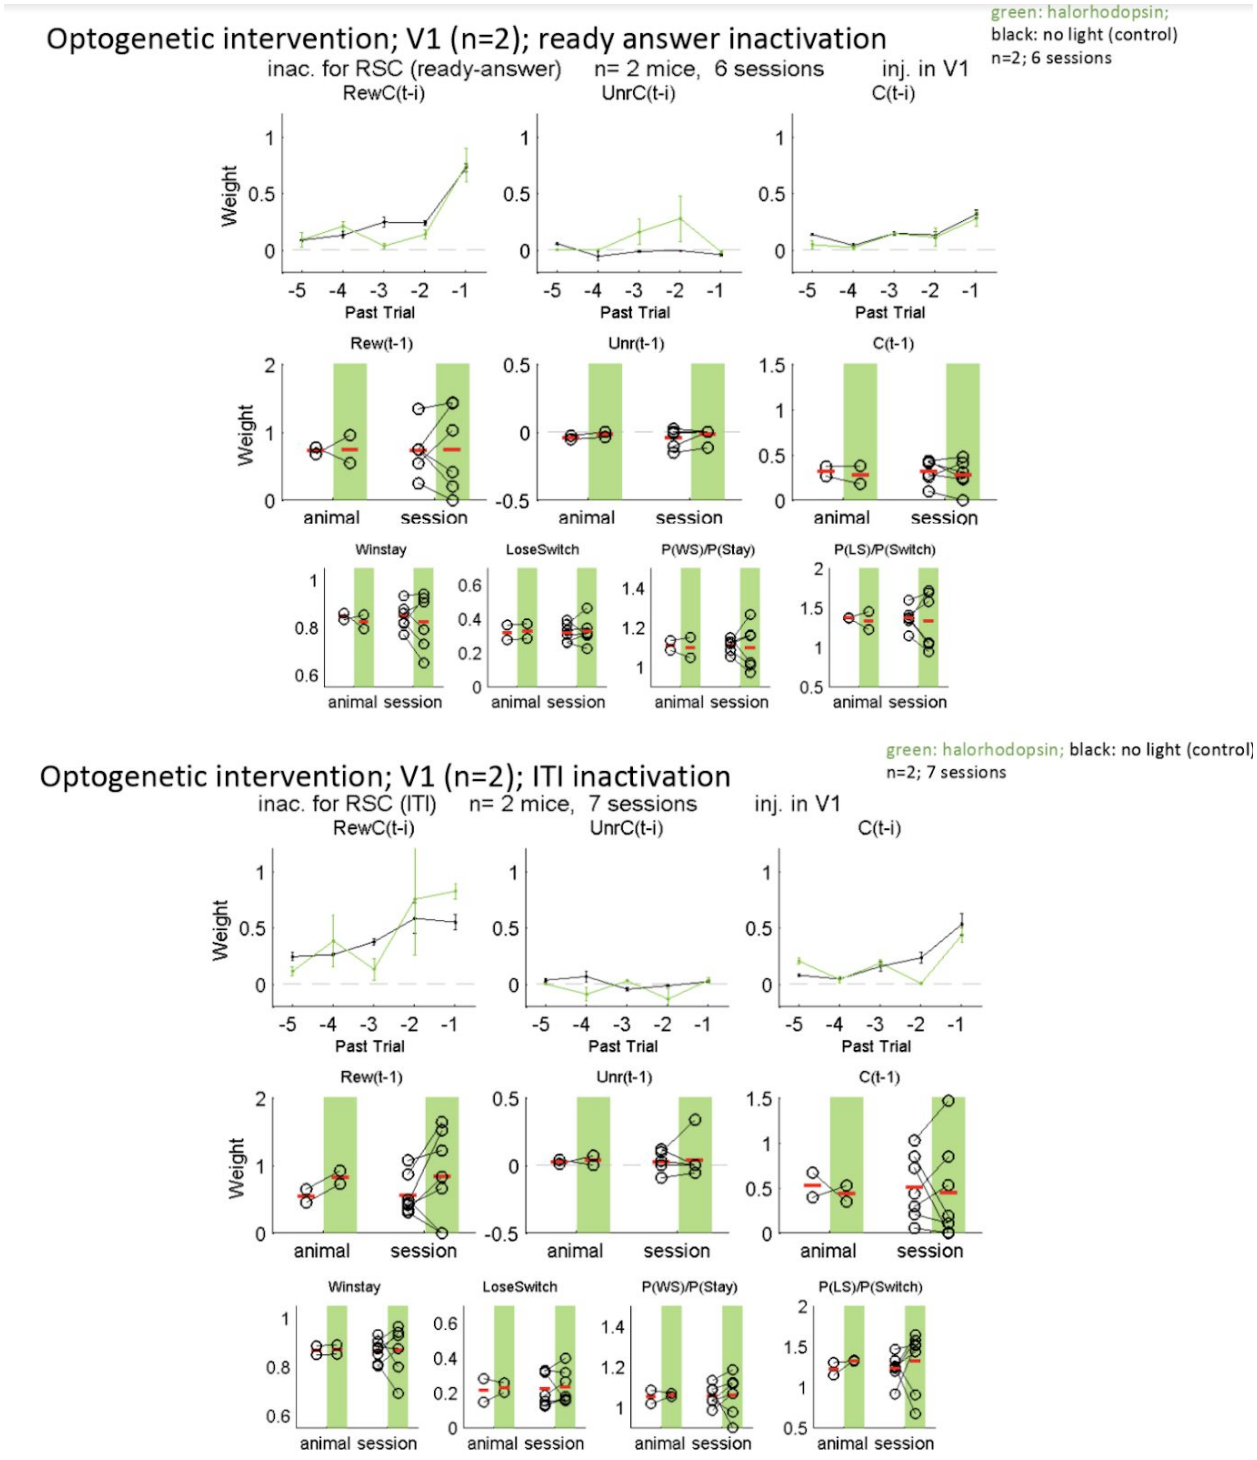

(D) RSC inactivation with V1 injection sites shows no effect.

Experiment III

Immunostaining of the two mice with CA1 and two mice with subiculum target injections of pAAV-CaMKlla-eNpHR 3.0-EYFP (halorhodopsin virus) showed the expression in its respective injection site areas - although, it appears that the virus spilled into the surrounding areas as well, potentially affective RSC as well (Figure 3A).

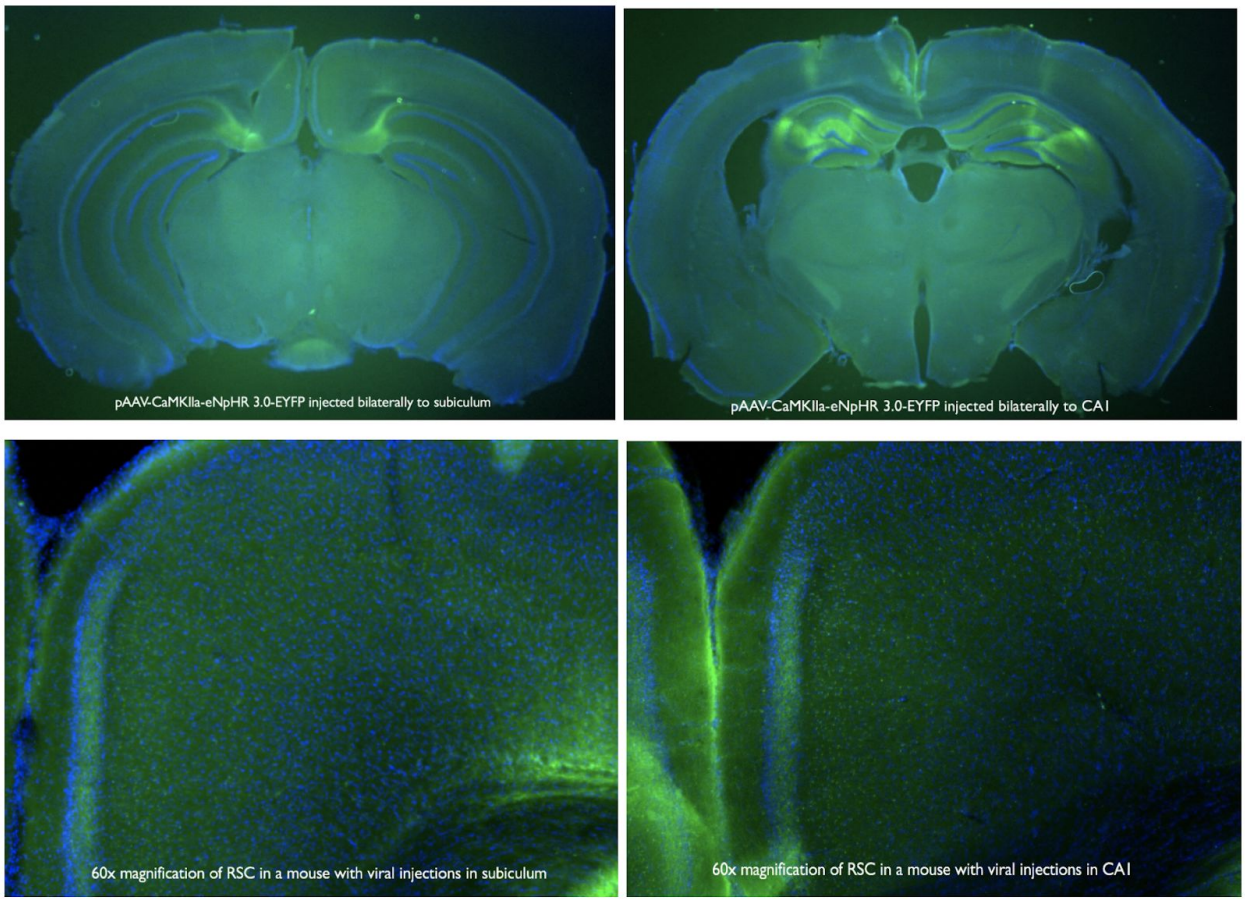

Figure 3: Immunostaining of pAAV-CaMKlla-eNpHR 3.0-EYFP Expression in Hippocampal Areas. (A) (Top, left) Virus was expressed in the subiculum and partially in its surrounding areas. (Top, right) Virus was expressed in CA1, and it appears it leaked to its surrounding areas and corex, as well - potentially affecting RSC too. (B) (Bottom) There is a sparse fluorescence in RSC magnified images in both a mouse with viral injection in subiculum (left) and a mouse with viral injection in CA1 (right).

Attempting to identify axonal expression and assure that projections indeed reached RSC, 60x magnified images of CA1 and subiculum were taken, and although these do show some green fluorescence of what could be real axons, it could also be autofluorescence or virus leak in the cortex (Figure 3B).

DISCUSSION

By adjusting the dynamic foraging task to head-fixed mice, a paradigm was established in which mice learn to make a decision based on history-dependent value of given choices. The study done by Hattori and colleagues (2019) revealed that RSC improves encoding of history information when used for decision making, and keeps value-related information persistently.

A notable finding in the current research is that RSC uniquely maintains persistent value-related population activity that tracked value updates on a trial-by-trial basis. Value-related information has to be accessible to action selection circuitry even if actions are not on a fixed temporal schedule (e.g., variable inter-trial intervals as in the task at hand). The persistent value coding in RSC might allow other brain areas to retrieve this information at any moment. However, it is noted that persistent population activity might not be the only mechanism by which value information is maintained in the brain. Persistent activity seems especially fitting for short-term maintenance of value information when the brain needs to frequently retrieve value information and update it, as in the task in which mice are making hundreds of decisions in each behavioral session. It is hypothesized that this mechanism might coexist with additional mechanisms that prefer a more long-term, stable storage of value information when it might only need to be accessed in the distant future. Such long-term storage mechanisms might include steady modifications in synaptic weights.

One other important finding is that history coding in RSC is flexible. RSC increased history coding when mice learned to use history for decision making. The other regions examined in Hattori et al. 's (2019) study did not show such correlations with the ongoing behavioral approach. These results indicate that although history coding is widespread and its flexibility is area-specific.

History coding in RSC is particularly sensitive to the continuing behavioral strategy. It was also discovered that RSC is required for reward-history based decision making. Acute optogenetic inactivation of RSC during the pre-choice period selectively impaired behavioral dependence on rewarded and unrewarded choice history. These behavioral effects contrast with the previously reported effects of mouse medial prefrontal cortex (mPFC) inactivation, which specifically affected choice bias but not the dependence on reward history, indicating that different areas mediate different aspects of the behavior. This idea fits with the findings that history-related and value-related information is widespread across many areas. The partial redundancy of value coding likely ensures the robustness of value-based decision making, a fundamental and evolutionarily conserved behavior.

Because of how densely RSC is interconnected with many other areas, including the hippocampus and related cortical areas, premotor cortex, basal ganglia, and thalamus, additional experiments were performed utilizing optogenetic inactivation to assess the effects of ACC, VA1, CA1, and subiculum projections on RSC. Although the experiment involving hippocampal projections was not completed due to the pandemic, a pilot experiment that attempted to assess the effect ACC and VA1 projections to RSC indicated that while ACC may be involved in value-based decision making, inhibiting V1 projections to RSC showed no effect on behavior.

With this hub-like connectivity, RSC might be an ideal area for the computation and persistent maintenance of history-dependent value with access to choice and outcome information as well as action selection circuits. These findings pave the way for future assessment of circuit mechanisms for the plasticity and persistence of value coding with RSC as a central locus.

Parts of the discussion section is a rewrite of the material as it appears in Area-Specificity and Plasticity of History-Dependent Value Coding During Learning, 2019, Ryoma Hattori, Bethanny Danskin, Zeljana Babic, Nicole Mlynaryk, and Takaki Komiyama published in Cell 2019. The thesis author was one of the authors of this publication.

REFERENCES

Alabi, O. O., Fortunato, M. P., & Fuccillo, M. V. (2019). Behavioral Paradigms to Probe Individual Mouse Differences in Value-Based Decision Making. *Frontiers in Neuroscience*, *13*. doi: 10.3389/fnins.2019.00050

Dehaene, S., & Changeux, J.-P. (2000). Reward-dependent learning in neuronal networks for planning and decision making. *Progress in Brain Research Cognition, Emotion and Autonomic Responses: The Integrative Role of the Prefrontal Cortex and Limbic Structures*, 217–229. doi: 10.1016/s0079-6123(00)26016-0

Hattori, R., Danskin, B., Babic, Z., Mlynaryk, N., & Komiyama, T. (2019). Area-Specificity and Plasticity of History-Dependent Value Coding During Learning. *Cell*, *177*(7). doi: 10.1016/j.cell.2019.04.027

Filla, I., Bailey, M. R., Schipani, E., Winiger, V., Mezias, C., Balsam, P. D., & Simpson, E. H. Striatal dopamine D2 receptors regulate effort but not value-based decision making and alter the dopaminergic encoding of cost. *Neuropsychopharmacol* 43, 2180–2189 (2018). https://doi.org/10.1038/s41386-018-0159-9

Krasheninnikova, A., Höner, F., O'Neill, L., Penna, E., & Bayern, A. M. P. V. (2018). Economic Decision-Making in Parrots. *Scientific Reports*, *8*(1). doi: 10.1038/s41598-018-30933-5

Milczarek, M. M., Vann, S. D., & Sengpiel, F. (2018). Spatial Memory Engram in the Mouse Retrosplenial Cortex. *Current Biology*, *28*(12). doi: 10.1016/j.cub.2018.05.002

Radwanska, A., Debowska, W., Liguz-Lecznar, M., Brzezicka, A., Kossut, M., & Cybulska-Klosowicz, A. (2010). Involvement of retrosplenial cortex in classical conditioning. *Behavioural Brain Research*, *214*(2), 231–239. doi: 10.1016/j.bbr.2010.05.042

Roy, Dheeraj S., et al. "Distinct Neural Circuits for the Formation and Retrieval of Episodic Memories." *Cell*, vol. 170, no. 5, 2017, doi:10.1016/j.cell.2017.07.013.

Sousa, A. F. D., Cowansage, K. K., Zutshi, I., Cardozo, L. M., Yoo, E. J., Leutgeb, S., & Mayford, M. (2019). Optogenetic reactivation of memory ensembles in the retrosplenial cortex induces systems consolidation. *Proceedings of the National Academy of Sciences*, *116*(17), 8576–8581. doi: 10.1073/pnas.1818432116

Tanaka, K. Z., Pevzner, A., Hamidi, A. B., Nakazawa, Y., Graham, J., & Wiltgen, B. J. (2014). Cortical Representations Are Reinstated by the Hippocampus during Memory Retrieval.

Neuron, *84*(2), 347–354. doi: 10.1016/j.neuron.2014.09.037

Vedder, L. C., Miller, A. M. P., Harrison, M. B., & Smith, D. M. (2016). Retrosplenial Cortical Neurons Encode Navigational Cues, Trajectories and Reward Locations During Goal Directed Navigation. *Cerebral Cortex*. doi: 10.1093/cercor/bhw192

Vogt, B. A. (2016). Midcingulate cortex: Structure, connections, homologies, functions and diseases. *Journal of Chemical Neuroanatomy*, *74*, 28–46. doi: 10.1016/j.jchemneu.2016.01.010