Investigating a role for hippocampus during cost based decision making

Julianne Davis

A thesis submitted in partial fulfillment

of the requirements for the degree of

Master of Science

University of Washington

2017

Committee:

Sheri J.Y. Mizumori

Jeansok J. Kim

Program Authorized to Offer Degree:

Department of Psychology

University of Washington

Abstract

Investigating a role for hippocampus during cost based decision making

Julianne Davis

Chair of the supervisory committee:

Professor Sheri J. Y. Mizumori

Department of Psychology

Hippocampus (HPC) has been implicated in decision-making as context-based features of reward have been shown to bias hippocampal place field properties. Its direct connectivity with brain structures known to be involved in determining the value of response outcomes (e.g. ventral striatum) and impulsivity related to reward-motivated behavior (e.g. orbitofrontal cortex) further suggests a role for the HPC during reward-motivated decision making. Indeed, it has been shown that the expected probability of receiving a reward biases HPC place field remapping: it seems to scale directly with the probability of receiving a reward during a probability discounting task (Tryon et al., 2017). In addition, HPC lesions increase variability in an animal's preference for large delayed rewards over small immediate rewards in a delay discounting task suggesting that the HPC is involved in cost-benefit decision making (McHugh et al., 2008, Abele and Chudasama, 2013). In this study, we sought to determine whether HPC is necessary for accurately representing and utilizing delay information by temporality inactivating HPC during a

delayed discounting T-maze task. In addition, we examined hippocampal CA1 place cell activity during performance of the same task in an effort to determine if they encode delay associated costs information.

A delay-based decision making maze task required rats to evaluate the cost of a decision in terms of the amount of time they were willing to wait in order to obtain the large reward (Jo, 2014). The HPC was infused with saline (SAL) or muscimol (MUS), a $GABA_A$ antagonist, prior to experimental sessions for the inactivation study. For the rats in the recording study, neural activity was correlated with maze region and delay context. Consistent with operant studies, the behavior analysis from 10 rats revealed the strongest preferences for the large reward when it was associated with the 10s delay, but this preference declined as the delay increased. Inactivation results suggest that HPC is necessary for normal delay discounting on this particular task. Current preliminary analysis of hippocampal place fields does not show evidence for encoding the cost of a reward in that population level place field remapping was not observed across trials associated with different delays. This suggests that the intrinsically generated value constructs, such as the expected cost of a reward, may be facilitated by HPC but it is not reflected in broad population level changes in place cell activity.

INTRODUCTION

Hippocampal place cells exhibit location-selective activity patterns which seem to form the neural basis for our sense of self location (Colgin, 2008). This has led to the long standing view of hippocampus (HPC) as a cognitive map. However, a growing body of research suggests that place cells do much more than encode current external sensory information (Anderson and Jeffery, 2003). Indeed, it has been shown that they are active in the absence of external cues, responsive to reward stimuli, necessary for accurate memory recall, and that they reflect some aspects of current motivational state (Anderson and Jeffery, 2003). It has also been shown that place cells are critical for context discrimination (Smith and Mizumori, 2006) and it is now thought that context information mediated by the HPC may be applied to facilitate the accuracy of decision making neural circuitry (Colgin, 2008; Mizumori and Tryon, 2015).

The HPC's direct connectivity with brain structures known to be involved in determining the value of response outcomes (e.g. ventral striatum) and impulsivity related to reward-motivated behavior (e.g. orbitofrontal cortex) further implicates a role for the HPC during reward-motivated decision making (Bett et al, 2015). Context-based features of reward have been shown to bias HPC place field properties as well suggests that this structure is involved in processing this kind of information (Colgin, 2008; Mizumori and Tryon, 2015). Indeed, HPC lesions increase variability in an animal's preference for large delayed rewards over small immediate rewards in a delay discounting task which suggests that the HPC is involved specifically in cost-benefit decision making (McHugh et al., 2008, Abela and Chudasama, 2013). In addition, it has been shown that the expected probability of receiving a reward biases HPC place field mapping, and this scales directly with the probability of receiving a reward during a probability discounting task (Tryon et al., 2017). HPC place cells have also been shown to

change their activity patterns in as a function of distance traveled, time on a particular path, and temporal aspects of memory in general suggesting that the HPC may be incorporating temporal components of a context when generating internal representations of its external environment (Kraus et al., 2013, Mankin et al., 2015). Thus, we hypothesized that HPC plays an important role in representing and integrating cost-related information during reward-motivated decision making and that HPC place cell activity will reflect variations in cost related information as contexts change.

To test these ideas, we first inactivated dorsal HPC to determine if this brain region is necessary for normal decision making behavior during a reward-motivated delay discounting task on an elevated T-maze. Given the findings reported in similar studies we expected that HPC inactivation would result in a decreased tolerance for delays preceding receipt of reward (i.e. increased impulsive choices) (McHugh et al., 2008, Abela and Chudasama, 2013). We then recorded place cell activity in the CA1 region of HPC in rats as they performed the same task in an effort to determine if place cells would exhibit differential activity patterns across delay conditions.

MATERIALS and METHODS

Subjects

Ten adult male Long-Evans rats (257-453 g; Simonsen Labs, Gilroy, CA) were housed individually in Plexiglas cages (Jo, 2014). All rats were kept on a 12 hour light/dark scale (lights on at 7:00 am) and were food restricted to 80-85% of their baseline body weight prior to, and over the duration of, their time as participants in the study (Jo, 2014). Animal care and experiments were conducted during the light phase, in accordance with the University of Washington's Institutional Care and Use Guidelines (Jo, 2014).

Behavioral Apparatus

An elevated T-maze, with a start arm and two goal arms (58 X 10 cm each), was located at the center of a circular curtained area (Jo, 2014). Four visual cues were hung on the curtains. Each goal arm contained a metal food cup (0.7 cm in diameter × 0.6 cm deep) at the distal end, in front of which was a painted wooden barrier (10 × 4 × 15 cm) which was used to control access to reward (Jo, 2014). A second painted wooden barrier (10.5 × 3.5 × 12.5 cm) was used to prevent the rats from moving away from the designated delay regions on the goal arms as they waited for a reward. The goal arms were hinged such that its proximal end closest to the maze center could be lowered by remote control as needed (Jo, 2014).

Presurgical training

Each rat was placed on the maze and allowed to forage for sucrose pellets (45 mg) dropped throughout the maze. Afterward, the rat was shaped to retrieve reward (4 pellets) delayed by 3 s only from two goal arms. Each animal was trained to move down the start arm and freely choose

either goal arm. Upon arrival in front of the barrier, the rat waited for 3 seconds (Jo, 2014). As the barrier was removed, and the experimenter measured the elapsed time (using a digital stopwatch) from trial onset to the time that the rat approached and consumed the reward (Jo, 2014). The experimenter gently guided the rat to the start arm for the next trial, then re-baited the food cup and put the barrier back in place (Jo, 2014). Once the rat could finish 17 trials within 20 min without being guided back to the start arm, it underwent stereotaxic implantation of cannula or hyperdrives (Jo, 2014).

Surgery, and cannula and electrode preparation

Under anesthesia with isoflurane (4% mix with oxygen at a flow rate of 1L/min), rats were mounted on a stereotaxic instrument (David Kopf Instruments, Tujunga, CA) (Jo, 2014). Once the rat was unconscious the isoflurane concentration was reduced to 1-3% (Jo, 2014). The skull was exposed and adjusted to position bregma and lambda parallel to the surgery table (Jo, 2014). Five rats received bilateral implantation of guide cannulae (26 gauge) aimed at the HPC (3.8 mm posterior, 2.0 mm lateral, and 2 mm ventral to bregma). A 33 gauge dummy cannula was inserted into each guide cannula to prevent blockage (Jo, 2014). Five additional rats received unilateral implantation of recording electrodes aimed at dorsal CA1 of HPC. Tetrodes were made from 20μm lacquer-coated tungsten wires (California Fine Wire, Grover Beach, CA) and final impedance of each wire was adjusted to 0.1-0.4 M Ω (tested at 1 kHz) (Jo, 2014). Six or twelve individually drivable tetrodes were loaded onto a custom built hyperdrive, which was in turn chronically implanted in the right hemisphere dorsal to HPC (3.8 mm posterior, 2.0 mm lateral, and 2 mm ventral to bregma).

Delay discounting task

After a week of recovery from surgery, the two groups of rats performed decision making task on an elevated T-maze (Fig. 1B). All rats were first trained to perform a delay-based decision making task in which they were required to choose between a sooner small (SS) reward and a later large (LL) reward (Jo, 2014). The delay to SS reward (1 sucrose pellet) was held constant at 3 s throughout the experiments, whereas three different delays (10, 20, and 40 s) before LL reward (4 sucrose pellets) were used to test possible changes in choice preference as a function of delay to LL reward (Fig. 1C) (Jo, 2014). A daily testing session was comprised of three blocks of trials to which the three delays were assigned in a pseudo-random order and only one delay was presented in a given block (Jo, 2014). To ensure that rats were informed of the delay associated with the LL reward in a given block, each block began with 10 forced-choice trials followed by 10 free-choice trials (Jo, 2014). During the forced-choice trials, the rats performed 5 SS and 5 LL consecutive reward trials where only one goal arm and reward was made available in each trial by lowering the other goal arm (Jo, 2014). Both goal arms were made available during the free-choice trials in which animals' behavior was recorded (Jo, 2014). As the rats selected and entered the goal arms, an additional barrier was placed at its entrance to prevent them from moving out of the delay region during the longer delays (Jo, 2014). The three blocks were separated by an inter-block interval of 5 minutes during which the animals were placed on a rest area adjacent to the maze (Jo, 2014). The location of SS and LL rewards on the maze remained constant within each rat but was counterbalanced across rats (Jo, 2014).

The five rats showed discounting behavior for a minimum of three days prior to undergoing a mock infusion followed by additional training session to control for the stress of a

novel procedure. The following day, saline (SAL) or muscimol (MUS) was infused into CA1 15-20 minutes prior to the daily behavioral session. Rats received the SAL or MUS injections in an ABBA fashion counterbalanced between the rats. For four of the rats an equal delay control session was also run after completion of the ABBA portion of the experiment. For this task, following an infusion of MUS, the rats were again asked to complete the three block delay discounting task. However, this time the delay to both the small and the large reward were equal at 10, 20, or 40s. A total of 4-6 drug testing sessions were given per rat.

The tetrodes implanted in the rats included in the recording study were slowly lowered into dorsal CA1 during training over the course of 11-21 days. Once the rats discounted for a minimum of three days and putative place cells were observed, daily recording sessions began. Electrodes were advanced at the end of each completed recording session to increase the chance of stable recordings of new cells the next day. Recording sessions continued until all electrodes passed through the CA cell layer.

Intracranial microinjection

A 33 gauge injection cannula extending 1 mm below the tip of the guide cannula was connected to a 10 μl syringe (Hamilton, Reno, NV) via polyethylene tubing (PE 20) (Jo, 2014). Either 0.5 μl of muscimol (MUS; 0.5 μg/μl dissolved in saline, SAL) or the SAL vehicle was bilaterally injected at a rate of 0.5 μl/h using a microinfusion pump (KD Scientific, Holliston, MA). The injection cannula were left in place for an additional 1 minute to ensure proper diffusion (Jo, 2015). Then rats were returned to their home cage and the behavioral recording began 15-20 min later.

HPC single-unit recording

Neural activity was monitored on each recording day using a Cheetah data acquisition system (Neuralynx, Bozeman, MT) (Jo, 2014). Unit signals were amplified, filtered at 600-6000 Hz, and digitized at 16 kHz (Jo, 2014). The animal's head position was also recorded at 30 Hz by tracking two light-emitting diodes mounted on the head stage (Jo, 2014). If no isolated units were detected, tetrodes were advanced in 40 μm increments, up to 160 μm per day (Jo, 2014). Once well isolated units were found and the LFP traces displayed activity characteristic of HPC, behavioral and neural recording were conducted. All tetrodes were advanced at the end of each recording session to facilitate stable recordings of new cells the following day. During the recording session, three salient events flags were inserted into the data stream. These events were delay onset (DO), delay termination (DT), and receipt of reward (Jo, 2014). Timestamps for DO and DT were generated when an experimenter pressed a button connected to the Cheetah data acquisition system (Jo, 2014). The time of reward encounters was automatically detected by 'lick-detectors' (custom designed by Neuralynx) when the animals licked the food cups as they retrieved the sugar pellet reward (Jo, 2014).

Histology

After completion of all experiments, small marker lesions were generated by passing a 10 μA current for 10 s through two wires of each tetrode in each of the animals implanted with hyperdrives (Jo, 2014). All rats were transcardially perfused with saline and formalin. Their brains were stored in a 10% formalin/30% sucrose solution and cut in coronal sections (45 μm) on a freezing microtome (Jo, 2014). The sections were stained with cresyl violet and imaged

under light microscopy to reconstruct tetrode tracks and cannula placement in the CA1 of HPC (Jo, 2014). Data recorded only from the HPC were analyzed.

Unit classification

Single-units (> 2:1 signal-to-noise ratio) were isolated by clustering various spike waveform parameters using Offline sorter (Plexon, Dallas, TX) (Jo, 2014). For units that were recorded more than one day, the session in which the units were most clearly isolated from other units and background noise was selected for analysis (Jo, 2014). Putative hippocampal pyramidal neurons were initially identified according to mean firing rate (> 0.1 Hz, < 10 Hz) (Tryon, 2017). Additionally the place fields had to have reliability and specificity scores > 0.3 (Tryon, 2017).

RESULTS

Effect of HPC inactivation on delay-discounting behavior

HPC is known to be necessary for the encoding of many contextual cues likely used for effective decision making. For example, place cells have been shown to cluster around reward locations further supporting the hypothesis that HPC is involved in processing information related to reward, and possibly, reward-motivated decision making (Dupret et al., 2010). Indeed, it has been demonstrated that lesioning HPC resulted in increased variability in rats' behavior while performing a maze-based delay discounting decision making task (McHugh et al., 2008, Abela and Chudasama, 2013). Hippocampal place cells have also been implicated in tracking time and path duration (Mankin et al., 2015, Kraus et al., 2013). We hypothesize that HPC plays a role in integrating cost, reward, and temporal information and that it is necessary for effective decision making. We predicted that temporary inactivation of dorsal HPC with MUS would alter

choice preference for LL due to an inability to accurately represent and effectively use temporal components of a context during a reward motivated task. To address this hypothesis 5 rats with bilateral cannulae aimed at HPC were tested in the delay discounting task (Fig. 1A). SAL and MUS were used for HPC manipulations and each drug was microinfused over four-five consecutive days to test choice behavior across the three pseudo-randomly ordered delays prior to the LL.

Prior to infusion of SAL or MUS all rats discounted for a minimum of three days and demonstrated a desire to choose the LL reward significantly more often in 10 second delay blocks as compared to the 40 second delay blocks (repeated measures ANOVA, F (2,24) = 7.3, p = 0.0027, Fig. 1D). Rats infused with SAL were more likely to choose the LL reward in the 10 second delay blocks as compared to the 40 second delay block (repeated measures ANOVA, $F_{(2,24)} = 6.7$, p = 0.00030, Fig. 1F). Infusion of MUS resulted in the abolishment of delay discounting behavior and an increased the likelihood to choose the LL reward at chance levels (Fig. 1G). No significant difference between percentage of LL and SS reward choice was found on days where the animals were infused with MUS (repeated measures ANOVA, $F_{(2,24)} = 0.37$, p = 0.69, Fig. 1E). An ANOVA with repeated measures followed by a Bonferroni *post hoc* test demonstrated a significant interaction effect between block and drug ($F_{(2,24)} = 6.8$, $p = 0.023$) with no significant effect of block ($F_{(2,24)} = 1.5$, $p = 0.29$) or drug ($F_{(1,4)} = 3.3$, $p = 0.11$).

HPC inactivation did not affect other aspects of the animals' behavior on the maze. Average block duration for the 10, 20 and 40 second delay conditions did not differ between SAL and MUS conditions (independent *t*-test, t (55) = 0.4477, $p > 0.05$; Fig. 2A). This suggests that mobility was not affected by MUS infusion. Both SAL and MUS-injected rats consumed all

rewards available on food cups, indicating no effect of HPC inactivation on consummatory behavior.

In order to rule out the possibility that inactivation of the HPC resulted in an inability to discriminate between reward sizes or in the rats forgetting the location or amount of the rewards an equal delays control was included. A subset of the rats (n = 4) were tested on this control task on the day following the completion of the delay discounting experiment. In this task animals were required to discriminate between small and large rewards after infusion of MUS into to HPC. The delays to both the small and the large rewards were both set to 10, 20, or 40 seconds depending on the block (Fig. 2B). As with the delay discounting task, each block started with 10 forced-choice trials, followed by 10 free-choice trials in which choice performance was measured. A wooden barrier was placed in front of the food cup of each goal arm to prevent rats from having access to reward during delays. While monitoring the elapsed time using a digital stopwatch, an experimenter removed the barrier at the time of delay termination. The location of SS and LL rewards in the goal arms remained stable for each rat but was counterbalanced between rats. On average, the rats increased the percentage of time they chose the large reward across all blocks (Fig. 2C). These results suggest that MUS infusion did not affect the rats' ability to discriminate between the large and small reward or prevent them from recalling this behavioral task.

Neural correlates in the HPC during the delay discounting task

In an effort to determine how HPC place cell activity might vary across different conditions on this task we recorded place cell activity from 5 rats implanted with hyperdrives.

Out of a total of 140 neurons recorded in the HPC (Fig. 3A), 55 (39.3%) neurons met the criterion for place cells.

While performing the delay discounting task, rats experienced a series of salient events, delay period conditions, and maze regions. The maze regions were defined as the start, decision, delay, and reward regions of the maze (Fig. 1B). Only data from recording trials in which the rats displayed delay discounting behavior were included (repeated measures ANOVA, $F(2, 37) = 13.65$, $p = 0.0001$, Fig.3C, D).

Initial qualitative observation of individual place field heat maps suggested remapping occurred as a function of agency, reward chosen, and delay context (Fig. 4). The first quantitative step taken to characterize the population level response was to determine if place fields were more likely to form in response to a particular cue, such as the reward, or region of the maze on our task. In order to do this the number of place cells with primary place fields in each of the four maze regions was compared. Initial observation of place field distribution on the maze suggested that a greater proportion of place fields formed at or near the reward locations in the Delay and Reward regions of the maze (Fig.5 A-B). This phenomenon has been described in previous studies and could be taken as confirmation that HPC processed reward information in similar manner during this task (Dupret et al., 2010).

Place cell remapping at a population level has been previously described as “global remapping”. This has been shown to occur in response to large scale environmental changes such as changing the room in which the experiment is conducted or changing the shape of the arena in which an open field experiment is conducted (Leutgeb et al., 2004). Given that delay context is a significant enough change to alter an animal’s behavior, we hypothesized that population level remapping of place cells may occur between delay contexts and that this would be reflected in

changing proportions of place fields in a given region of the maze across delay contexts. A Chi-square test was used to determine if the proportion of primary place fields in a given maze region differed from chance across delay conditions. The results of this test suggested that there was no difference in the proportion of primary fields in the Start, Decision, Delay, and Reward regions across the Short, Intermediate, and Long delay contexts (χ^2 (6, N = 12) = 0.00, p = 1.0, Fig. 5D). Though this does not conclusively prove that remapping is not occurring across delays it does suggest that HPC is not representing the maze in a drastically different manner across delay conditions.

Changes in the firing rate patterns of place cells within place fields whose locations do not change across conditions are indicative of rate remapping (Leutgeb et al, 2005). This well characterized phenomenon has been shown to result from changing specific variables in an animal's environment, such as the color of the walls in an arena (Leutgeb et al, 2005). We examined the average infield firing rates of the place cells recorded in this study and found that no significant difference between the firing rates of cells with primary place fields in the Short, Intermediate, or Long delay conditions (one-way ANOVA, F (2, 164) = 0.16, p = 0.85, Fig. 6D). This remained true for trials in which the rats chose the LL reward (one-way ANOVA, F (2, 78) = 2.6, p = 0.081, F (2, 86) = 0.91, p = 0.40, Fig. 6F) and for trials in which the rats chose the SS rewards (one-way ANOVA, Fig. 6E). Cells were further divided by maze region and a two-way ANOVA determined that no significant change in infield firing rate occurred between maze regions across delay conditions (F (11, 78) = 1.2, p = 0.33, Fig. 6G). Given that HPC place cells are known to remap in response to variation in reward location, the presence or absence of a reward, and the probability of receiving a reward we examined whether or not the size of the reward chosen affected the average firing rates across the different maze regions and delay

conditions. Average firing rates remained consistent across maze regions (one-way ANOVA, $F(3, 158) = 2.1$, $p = .10$, Fig. 6A). When we compared only the trials in which the rats chose the SS reward we found no difference of firing rate across the maze regions (one-way ANOVA, $F(3, 86) = 1.9$, $p = 0.13$, Fig. 6B). The same is true for trials in which the rats chose the LL reward (one-way ANOVA, $F(3, 72) = 1.1$, $p = 0.37$, Fig. 6C). A two-way ANOVA confirmed no effect of reward size or region on average infield fringing rates at a population level in this task ($F(6, 167) = 0.86$, $p = 0.54$). This suggests that drastic rate changes at a population level do not occur in response to changes in reward and delay conditions. Therefore, HPC likely represents each of these contexts, in which there is only one variable component, in a similar manner.

The information content (IC), which reflects the accuracy with which a place cell's firing represents that animals location, of all place cells within their primary place field was also compared across delay conditions and maze regions (Fig. 7). IC was found to be the same across maze regions (one-way ANOVA, $F(3, 162) = 0.39$, $p = 0.76$, Fig. 7A) delay conditions (one-way ANOVA, $F(2, 78) = 0.44$, $p = 0.65$, Fig. 7D). IC also remained unchanged across maze regions and delay condition (two-way ANOVA, $F(3, 78) = 0.84$, $p = 0.66$, Fig. 7G). The effect of reward size on the average IC values across the different maze regions and delay conditions were examined as well (Fig. 7B-C, E-F). Average IC values remained consistent across maze regions within trials when the rats chose the SS reward (one-way ANOVA, $F(3, 86) = 1.3$, $p = 0.28$, Fig.7B). The same is true for trials in which the rats chose the large reward (one-way ANOVA, $F(3, 76) = 0.61$, $p = 0.61$, Fig. 7C). IC values did not significantly change across delay conditions in which rats chose the SS (one-way ANOVA, $F(3, 86) = 1.2$, $p = 0.32$, Fig. 7E) or LL (one-way ANOVA, $F(3, 76) = 0.55$, $p = 0.58$, Fig. 7F) rewards. However, when a two-way ANOVA was used to compare IC across both delay condition and reward size chosen it was

revealed that IC was higher on average in the trials in which the rats chose the SS reward over the LL reward ($F(6, 167) = 29$, $p < 0.0070$).

DISCUSSION

The current study confirmed HPC's role in effective valuation of a reward with variable cost and characterized how populations of HPC place cells responded to changes in delay and reward conditions on a maze based delay discounting task. Consistent with HPC lesion studies in which the subjects completed similar tasks (McHugh et al., 2008; Abela and Chudasama, 2013), HPC inactivation resulted in discounting behavior being reduced to chance. Motor, memory, and reward magnitude discrimination functions were unaltered suggesting that the behavioral effect of inactivation was a result in changed valuation of reward or recognition of delay. This study also characterized how HPC place cells responded in a delay discounting task. Qualitative analysis of place field heat maps suggested that remapping occurred as a function of delay and reward amount chosen in a given trial. In an effort to confirm this effect place cell activity was analyzed in a variety of ways at the population level. Distribution of primary place field location across four regions of the maze was found to be unchanged across delay, region of the maze in which they were located, and reward size conditions. Infield firing rate also remained the same across delay, maze region, and reward size. IC of place cells was found to be higher for cells that were active in SS trials as compared to cells active in LL trials. IC did not change significantly across maze region and delay condition. Thus, the present study suggests that HPC is necessary for normal decision making behavior during delay discounting though this is not reflected in terms of global changes at the population level.

HPC is important for delay discounting

The HPC has been shown to encode salient contextual information (Smith and Mizumori, 2006; Colgin, 2008). More recently it has been implicated in decision making and shown to play a role in reward valuation and tracking probability of reward receipt (McHugh et al., 2008; Abela and Chudasama, 201; Tryon et al., 2017). Although HPC's involvement in processing decision making, temporal, and reward information remains to be fully characterized, the present study describes a role for HPC in integrating reward cost information in a familiar decision making task. The reduction of LL choice to chance suggests that HPC is either necessary for recalling and utilizing temporal cost information or increasing impulsivity as suggested in previous studies (McHugh et al., 2008). Given that inactivation resulted in abolishment of delay discounting behavior rather than shifting the discounting curve down we believe that these results support the former hypothesis. These results provide compelling evidence for the necessity of HPC in order to optimize cost-based decision making.

Responses of HPC place cells in the delay discounting task

Previous work indicates that HPC conveys reward (Tryon et al., 2017) and temporal information (Kraus et al., 2013). In support of these findings, HPC inactivation disrupted baseline behavior in a reward-motivated decision making task (Fig. 1E). Visual examination of place filed heat maps also suggest that place cells encode temporal and reward aspects of our task (Fig. 4). However, quantitative analysis of the recorded place cell population did not yield results suggesting significant change as a result of delay context maze region, or reward amount (Fig. 5-7).These results do not conclusively prove that HPC place cells are not

responding to these conditions. Rather they suggest that HPC is representing these conditions in a sufficiently similar manner such that any potential increase or decrease in activity is likely undetectable at the population level. Alternatively, it could be that a different neural network represents the different delay conditions. In this case, if some cells increased firing, while others decreased firing, the population analysis would show a null effect. Future analyses will include population vector analysis in order to determine the degree of correlation between place field properties between different conditions (Leutgeb, 2004). Peristimulus time histogram analysis will be used in order to determine if place cells fire at successive moments during the delay periods on this task. If variation in activity occurs between delay conditions it would support the argument that place cells provide an additional temporal dimension that is integrated with spatial mapping (Eichenbaum, 2014). Further evidence for robust representation of both time and space in the HPC during this task would describe a new role for this structure in reward motivated decision making mechanism.

References

Abela AR, Chudasama Y (2013) Dissociable contributions of the ventral hippocampus and orbitofrontal cortex to decision-making with a delayed or uncertain outcome. Eur J Neurosci 37:640-7

Anderson MI, Jeffery KJ (2003) Heterogeneous modulation of place cell firing by changes in context. J Neurosci. 23:8827-35

Bett D, Murdoch LH, Wood ER, Dudchenko PA. (2014) Hippocampus, delay discounting, and vicarious trial-and-error. Hippocampus 25:643-54.

Colgin LL, Moser EI, Moser MB (2008) Understanding memory through hippocampal remapping. Trends Neurosci 31:469-77

Dupret D, O'Neill J, Pleydell-Bouverie B, Csicsvari J (2010) The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat Neurosci 10:995-1004

Eichenbaum H (2014) Time cells in the hippocampus: a new dimension for mapping memories. Nat Rev Neurosci 15:732-44

Jo YS (2014) Prefrontal Modulation of Midbrain Dopamine Systems during Navigation-based Decision Tasks (Doctoral thesis). Retrieved from the University of Washington ResearchWorks Archive (http://hdl.handle.net/1773/26127)

Jo YS, Mizumori SJY (2016) Prefrontal Regulation of Neuronal Activity in the Ventral Tegmental Area. Cereb Cortex 26:4057-4068. By permission of Oxford University Press.

Kraus BJ, Robinson RJ 2nd, White JA, Eichenbaum H, Hasselmo ME. (2013) Hippocampal "time cells": time versus path integration. Neuron 78:1090-101

Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005) Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles. Science 309:619-623

Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI (2004) Distinct ensemble codes in hippocampal areas CA3 and CA1. Science 305:1295–1298

Mankin EA, Diehl GW, Sparks FT, Leutgeb S, Leutgeb JK (2015) Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts. Neuron 85:190-201

McHugh SB, Campbell TG, Taylor AM, Rawlins JN, Bannerman DM (2008) A role for dorsal and ventral hippocampus in inter-temporal choice cost-benefit decision making. Behav Neurosci 122:1-8

Mizumori, SJY, & Tryon, VL (2015) Integrative hippocampal and decision-making neurocircuitry during goal-relevant predictions and encoding. Progress in Brain Research, 219: 217-242.

Tryon VL, Penner MR, Heide SW, King HO, Larkin J, Mizumori SJ (2017) Hippocampal neural activity reflects the economy of choices during goal-directed navigation. Hippocampus. In press

Smith DM & Mizumori SJY (2006) Learning-related development of context-specific neuronal responses to places and events: The hippocampal role in context processing. Journal of Neuroscience, 26:3154-3163.

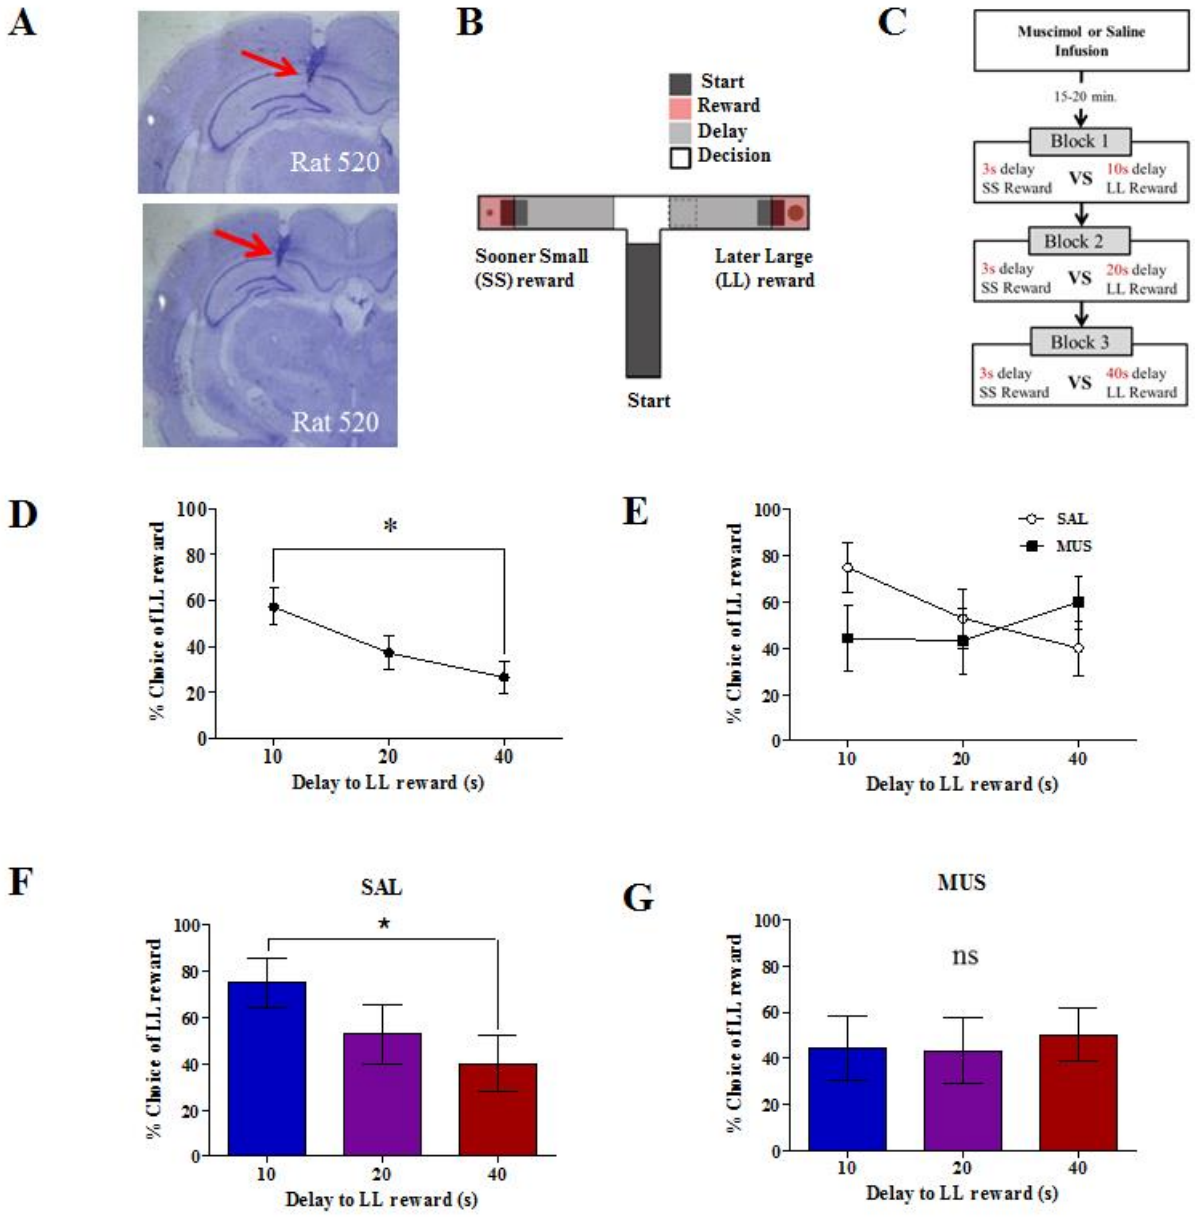

Figure 1. Choice performance on a delay-based decision making task. (**A**) Nissl-stained Section for cannula placement in CA1 of HPC. (**B**) Illustration of the T-maze with maze regions. Wooden barriers (black squares) were placed in front of food cups on the goal arms to control animals' access to rewards. When rats entered onto a goal arm, an additional barrier (dashed

rectangle) was placed behind the rat to prevent them from exiting the goal arm during the waiting period (adapted from Jo, 2016). (**C**) Daily behavioral procedures. Three different lengths of delay to LL reward were randomly ordered and tested in separate trials blocks. The delay to SS reward remained constant throughout the experiment. Each block consisted of ten forced-choice trials, followed by ten free-choice trials (adapted from Jo, 2016). (**D**) Pre-surgical delay discounting behavior. Error bars indicate SEM. (**E**) Choice preference for LL reward as a function of delay to LL reward. Error bars indicate SEM.

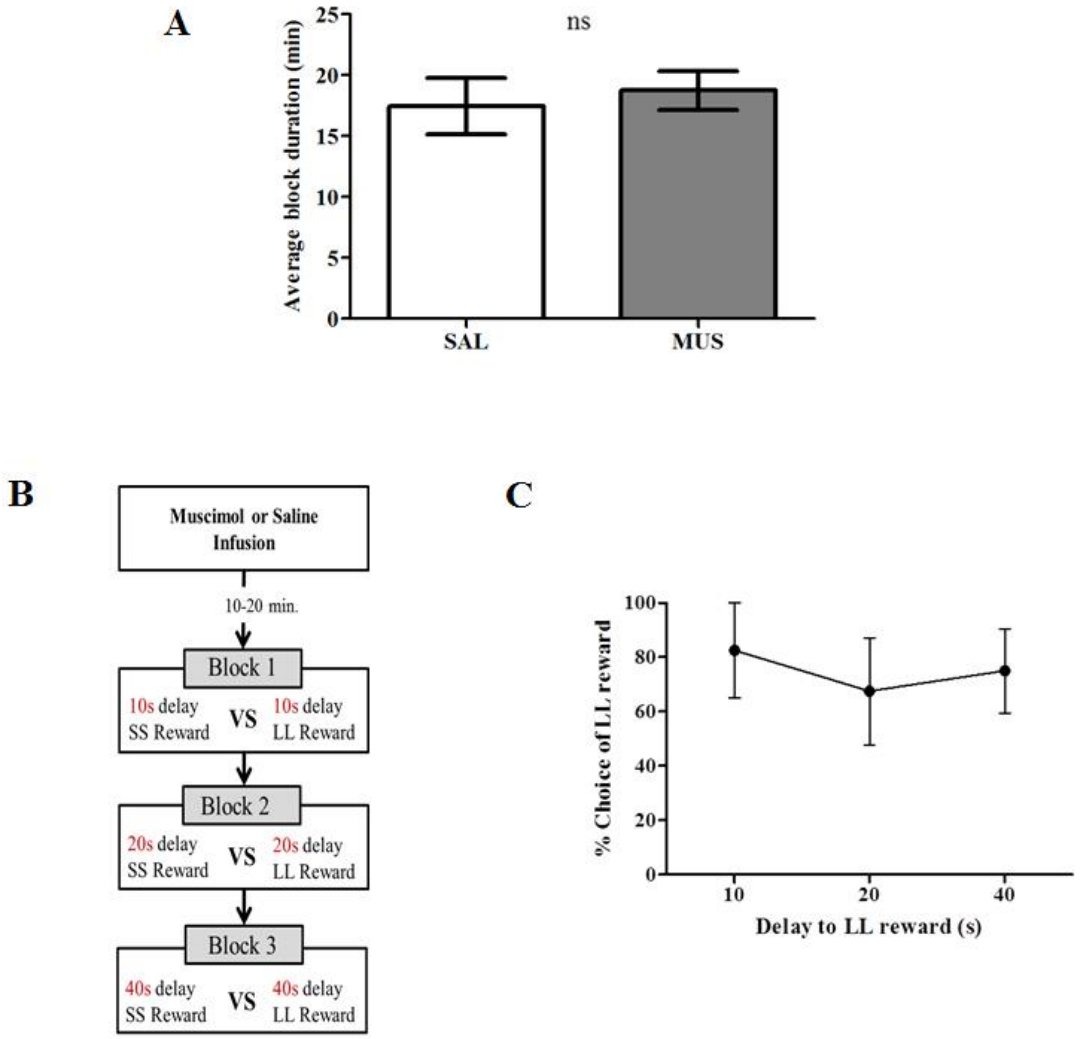

Figure 2. Control delay-based decision making task. (**A**) Average duration of one block of behavior. MUS injection into the HPC did not significantly increased the duration, compared to SAL injection. All graphs show mean ± SEM. (**B**) Equal delays control behavioral procedures. Three different lengths of delay to LL and SS reward were randomly ordered and tested in separate trial blocks. Each block consisted of ten forced-choice trials, followed by ten free-choice trials (adapted from Jo, 2016). (**C**) Average choice preference for LL reward in the control portion of the experiment. Error bars indicate SEM.

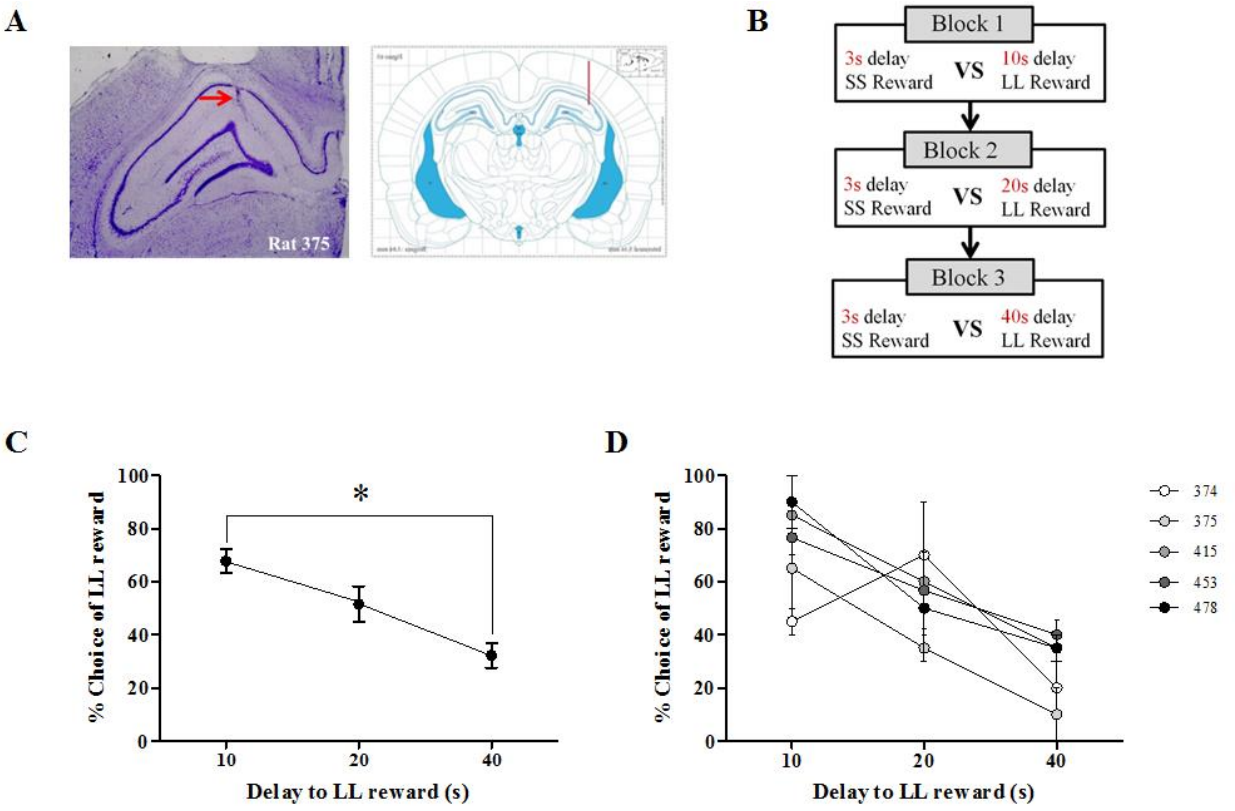

Figure 3. Choice performance on a delay-based decision making task. (**A**) Nissl-stained section showing the final location of the tetrode tips in the HPC (indicated by the arrows). (**B**) Daily behavioral procedures. Three different lengths of delay to LL reward were randomly ordered and tested in separate trial blocks. The delay to SS reward remained constant throughout the experiment. Each block consisted of ten forced-choice trials, followed by ten free-choice trials (adapted from Jo, 2016). (**C**) Choice preference for LL reward as a function of delay to LL reward for the trials included in the recording study. Error bars indicate SEM. (**D**) Choice preference for LL reward as a function of delay to LL reward for the individual rats included in the recording study. Error bars indicate SEM.

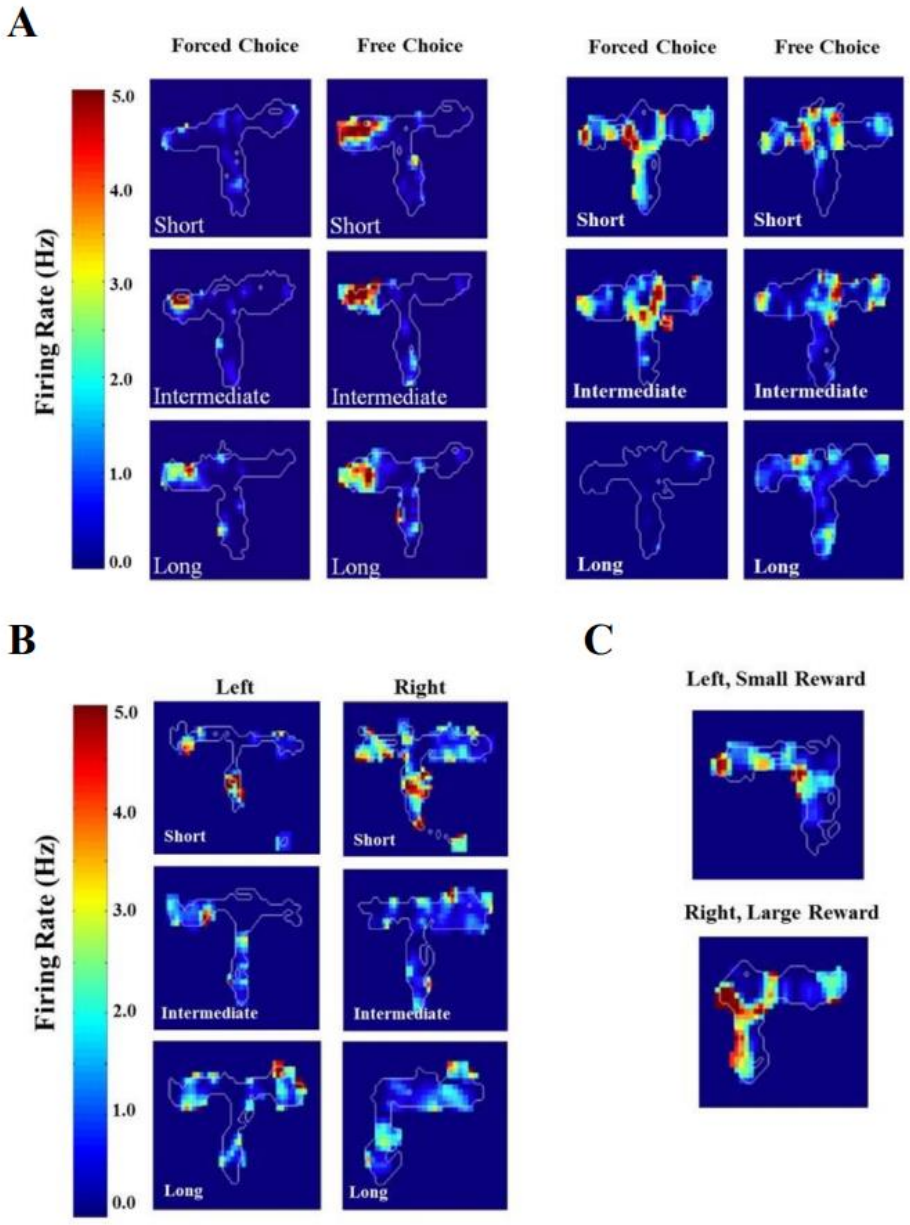

Figure 4. Example heat maps from place cells that respond to delay, agency, and reward size. (**A**) Firing rate maps of two cells who's place fields (PF's) appear to reflect changes in delay and agency. (**D**) Firing rate map a cell who's PF's appear to reflect changes in delay conditions.

(**C**) Firing rate map of a cell who's PF's appear to reflect changes in differential expectation of delayed rewards.

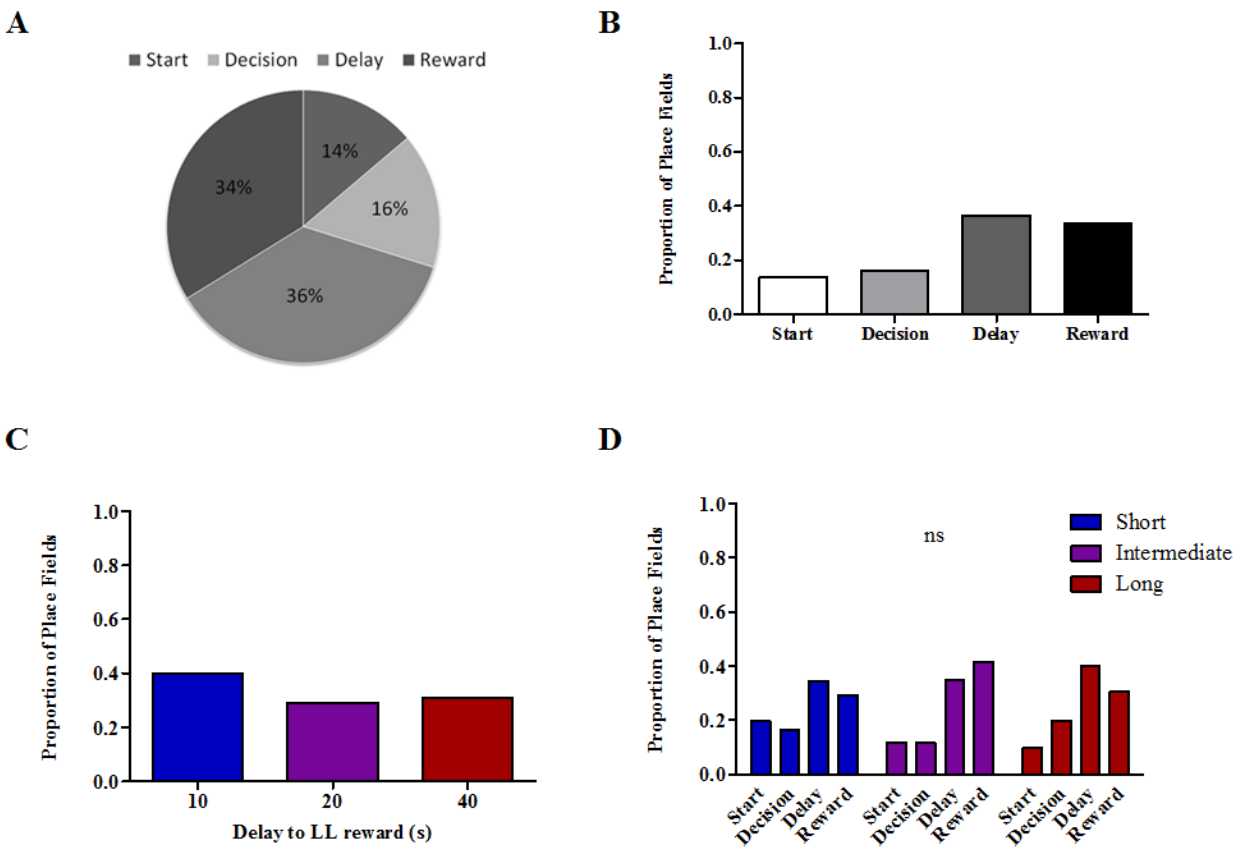

Figure 5. Place field distribution differs across maze regions and delay conditions. (**A**) Percentage of cells with primary place fields in each of the four maze regions over all trials. (**B**) Proportion of cells with primary place fields in the four regions of the maze. (**C**) Proportion of primary place fields that fall into one of the three delay conditions. (**D**) Proportion of primary place fields across delay conditions and maze region.

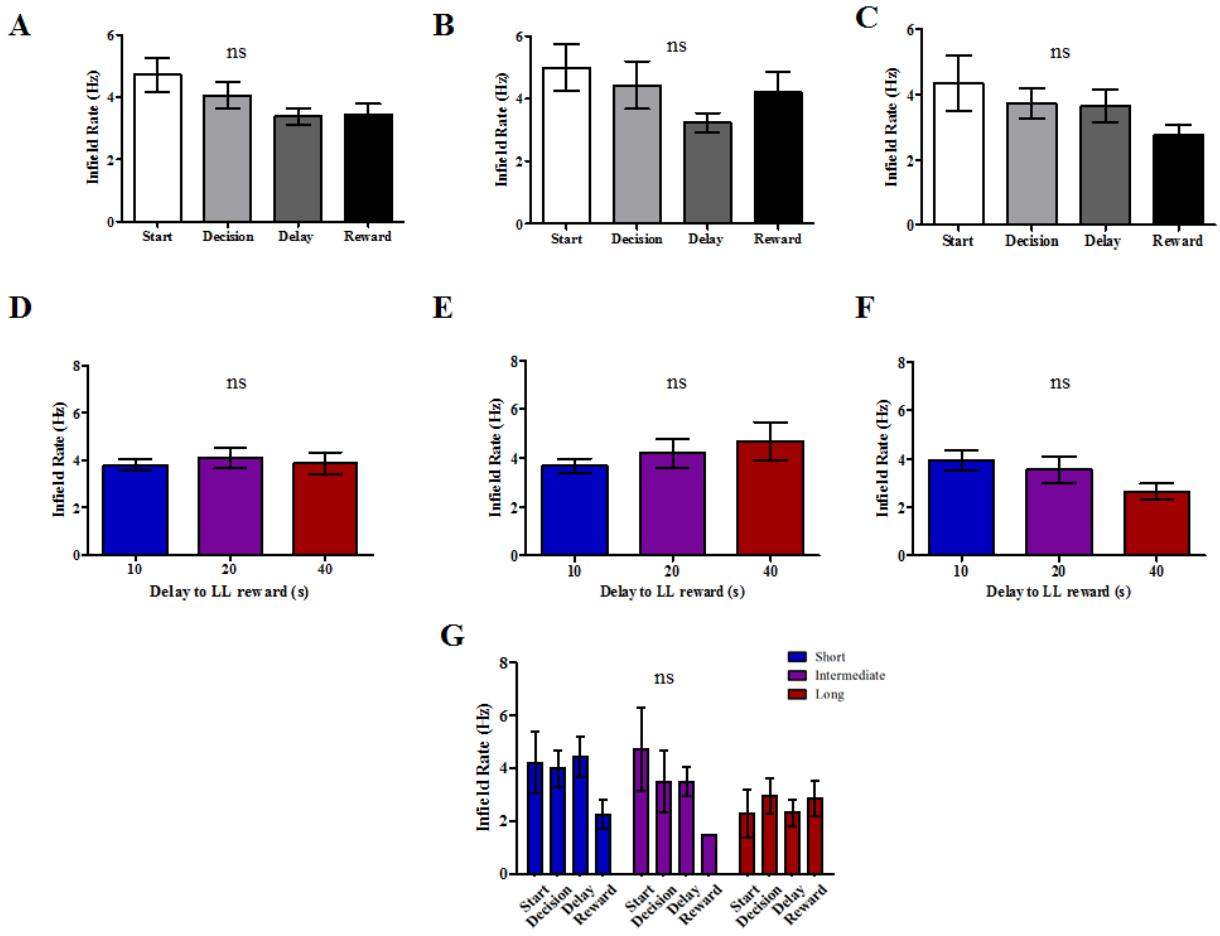

Figure 6. Infield firing rate averages across delay conditions and maze regions. (**A**) Average infield firing rates for all free choice trials for the start, decision, delay, and reward regions of the maze. Error bars indicate SEM. (**B**) Average infield firing rates for free choice trials in which rats chose the SS reward for the start, decision, delay, and reward regions of the maze. Error bars indicate SEM. (**C**) Average infield firing rates for free choice trials in which rats chose the LL reward for the start, decision, delay, and reward regions of the maze. Error bars indicate SEM. (**D**) Average infield firing rates for all free choice trials for the short, intermediate, and long delay conditions. Error bars indicate SEM. (**E**) Average infield firing rates for free choice trials in which rats chose the SS reward for the short, intermediate, and long delay conditions. Error

bars indicate SEM. (**F**) Average infield firing rates for free choice trials in which rats chose the LL reward for the short, intermediate, and long delay conditions. Error bars indicate SEM. (**G**) Average infield firing rate for all free choice trials for the short, intermediate, and long delay conditions across the four maze regions. Error bars indicate SEM.

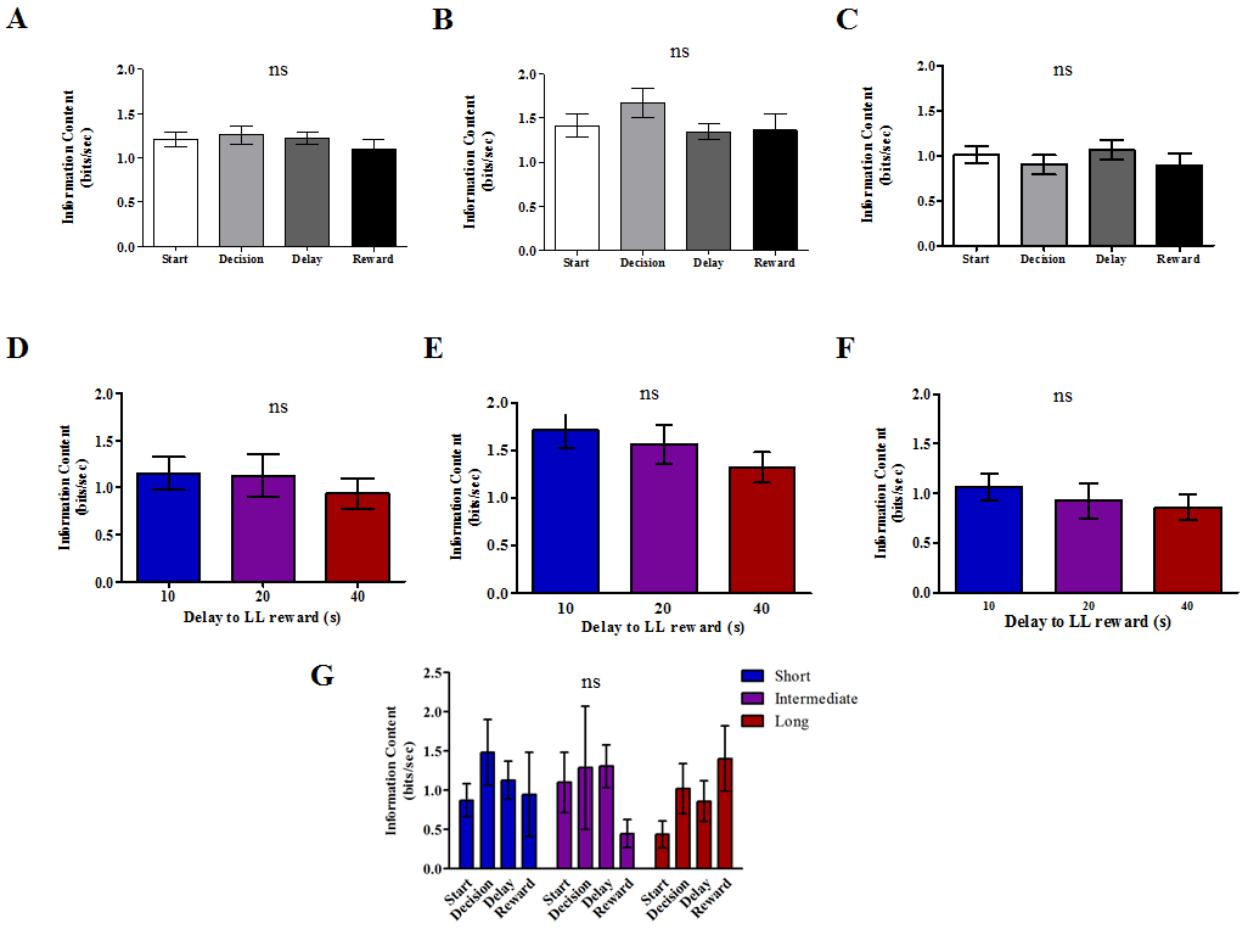

Figure 7. Information content (IC) averages across delay conditions and maze regions. (**A**) Average IC for all free choice trials for the start, decision, delay, and reward regions of the maze. Error bars indicate SEM. (**B**) Average IC for free choice trials in which rats chose the SS reward for the start, decision, delay, and reward regions of the maze. Error bars indicate SEM. (**C**) Average IC for free choice trials in which rats chose the LL reward for the start, decision, delay, and reward regions of the maze. Error bars indicate SEM. (**D**) Average IC for all free choice trials for the short, intermediate, and long delay conditions. Error bars indicate SEM. (**E**) Average IC for free choice trials in which rats chose the SS reward for the short, intermediate, and long delay conditions. Error bars indicate SEM. (**F**) Average IC for free choice trials in

which rats chose the LL reward for the short, intermediate, and long delay conditions. Error bars indicate SEM. (**G**) Average IC for all free choice trials for the short, intermediate, and long delay conditions across the four maze regions. Error bars indicate SEM.