Published in final edited form as:

Hippocampus. 2017 July ; 27(7): 743–758. doi:10.1002/hipo.22720.

Hippocampal neural activity reflects the economy of choices during goal-directed navigation

Valerie L. Tryon¹, Marsha R. Penner¹, Shawn W. Heide¹, Hunter O. King¹, Joshua Larkin¹, and Sheri J. Y. Mizumori^{1,2}

¹Psychology Department, University of Washington, Seattle WA USA

²Neuroscience Program, University of Washington, Seattle WA USA

SUMMARY

Distinguishing spatial contexts is likely essential for the well-known role of the hippocampus in episodic memory. We studied whether types of hippocampal neural organization thought to underlie context discrimination is impacted by learned economic considerations of choice behavior. Hippocampal place cells and theta activity were recorded as rats performed a maze-based probability discounting task that involved choosing between a small certain reward or a large probabilistic reward. Different spatial distributions of place fields were observed in response to changes in probability, the outcome of the rats' choice, and whether or not rats were free to make that choice. The degree to which the reward location was represented by place cells scaled with the expected probability of rewards. Theta power increased around the goal location also in proportion to the expected probability of signaled rewards. Further, theta power dynamically varied as specific econometric information was obtained 'on the fly' during task performance. Such an economic perspective of memory processing by hippocampal place cells expands our view of the nature of context memories retrieved by hippocampus during adaptive navigation.

Keywords

probability discounting; place fields; theta; decision; goal-directed navigation

INTRODUCTION

The critical role of the hippocampus in episodic memory (Tulving, 2002) likely reflects pattern separation processes (Yassa and Stark, 2011) which include computations that distinguish contexts or situations by comparing their predicted features relative to those actually experienced (e.g., Mizumori et al., 1999; Smith and Mizumori, 2006; Mizumori, 2008; Duncan et al., 2012; Penner and Mizumori, 2012; Mizumori and Tryon, 2015). In support of this view, location-selective firing (i.e. *place fields*) of many hippocampal *place cells* are altered when contextual features of a familiar situation change, features such as external multisensory information, behavioral requirements, memory demands, or reward location (Smith and Mizumori, 2006; Mizumori, 2008; O'Keefe, 1976; Muller and Kubie,

1987; Fyhn et al., 2002; Leutgeb et al., 2005a and b). Further, when rats are explicitly trained on a mismatch detection task, it has been shown that hippocampal neurons fire differently on match and mismatch trials (Manns et al., 2007).

Presumably, hippocampal neural activity during goal-directed navigation reflects decisions based on an analysis of context information. These decisions should optimize the conditions and likelihood that a goal will be achieved. The types of goal information that could guide such choices and decisions is not well understood. Knowledge of the location of a goal has been shown to bias the distribution of place fields around the goal location (Hollup et al., 2001; Lee et al., 2006). When explicitly trained to do so, rats behaviorally and neurally distinguish contexts based on their different goal locations even when other sensory, behavioral, and motivational aspects of a task are held constant (Markus et al., 1995; Smith and Mizumori, 2006). Such goal location-defined neural firing patterns continue to be observed when rats are removed from the maze during the intertrial interval (Gill et al., 2011). That knowledge of the location of a current goal can exert a powerful guiding influence on the organization of neural activity in hippocampus has also been demonstrated by findings that the place fields, and their experience-dependent sequential activation, depending on recent or future goal-directed trajectories (Wood et al., 2000; Louie and Wilson, 2001; Lee and Wilson, 2002; Foster and Wilson, 2006; Diba and Buzsaki, 2007; Johnson and Redish, 2007; Karlsson and Frank, 2009; Lisman and Redish, 2009; Dragoi and Tonegawa, 2011; Ferbinteanu et al., 2011; Pfeiffer and Foster, 2013; Cei et al., 2014; Wikenheiser and Redish, 2015).

Recordings from midbrain dopamine neurons of the ventral tegmental area (VTA) show that neurons signal other aspects of a goal in addition to its location. For instance, dopamine cells signal the novel presentation of rewards, as well as the cues that predict future rewards (Schultz et al., 1997; Fiorillo et al., 2003). These reward-related responses are modulated by the learned economic conditions of reward acquisition such as the probability that rewards will be delivered at particular reward sites. This study sought to determine whether similar characteristics of a goal (other than its location) drive the characteristics of spatial representation in hippocampus. Also of interest was whether theta in the local field potentials (LFPs) respond to different types of goal information since theta appears important for spatial decision making (Belchior et al., 2014) and spatial memory (Winston, 1978; Mizumori et al., 1989).

To assess whether hippocampal neural processing reflects goal information other than its location, we developed a probability discounting maze-based task in which hippocampal activity was recorded as rats made choices that were based on their understanding of the probabilistic nature of a goal that is found in a constant location. Notably, there was no right or wrong choice; rather we assessed biases in choices that varied with changes in econometric factors of interest such as the probability of an outcome, agency (i.e. whether rats made free choices or were forced to make a specific choice), and magnitude of expected rewards. These features systematically varied within a single recording session so that the same neural (place field and LFP) responses could be related to changes in decision biases.

MATERIALS AND METHODS

Subjects

Three (for hippocampal recordings) male Long–Evans rats (350–420 g; Simonsen Laboratories) were housed individually in Plexiglas cages. The rats were maintained on a 12 h light/dark cycle (lights on at 7:00 A.M.) and all behavioral experiments were performed during the light phase. Each rat was allowed access to water *ad libitum* and food-deprived to 80% of its *ad libitum* feeding weight. All animal care and use were conducted in accordance with University of Washington's Institutional Animal Care and Use Committee guidelines.

Probability discounting maze task

Rats were trained to run on an elevated loop maze (130cm long x 86cm wide x 80cm high; Fig. 1A; developed with the assistance of Cowen, S. (<http://cowen.faculty.arizona.edu/>). As part of pretraining, rats were first habituated to the maze and maze environment. To habituate rats to the maze, rats were allowed to freely forage for sugar pellets (45mg, TestDiet) scattered around the maze. Once they consistently consumed the sugar pellets and traversed the length of the maze, pretraining on the probability task began. The beginning of pretraining consisted of a simplified version of the task during which one door was open at a time. Passage through each door was associated with a 50% probability of receiving one sugar pellet. Once rats were reliably running the maze, which took an average of 10.5 pretraining sessions, their side preference was assessed, and the large risky option was assigned to either their preferred side or their non-preferred side in a counterbalanced manner. At this time, training on the probability discounting maze task began.

The fully automated probability discounting task was modeled after one used with operant chambers (St. Onge and Floresco, 2009). Each session started when the rat was placed on the maze. After a two second delay, the start door opened and the rat turned right to arrive at the decision platform. One door was associated with the large reward/risky option (defined below), the other the small reward/certain option (also defined below), and this was held constant for each rat throughout training but counterbalanced across rats. A sensor was located on each of the doors to track the door choice of the rat. After the rat passed through one of the “decision” doors (either the large reward/risky choice or the small reward/certain choice), it triggered another sensor that closed the decision doors preventing the rat from turning back. As the rat continued down the maze, passing another sensor triggered the delivery of the food reward so that a pellet(s) would be present in the food cup once the rat arrived there. An audible click signaled the delivery of each food pellet (80209 pellet dispenser, Campden Instruments Ltd., A Lafayette Instrument Company). Choice of the small reward/certain option always delivered one pellet (one click) with 100% probability; choice of the large reward/risky option delivered four pellets (four clicks) but with varying probability. A sensor was located at the food cup to record precise arrival at the reward location. After consuming the reward, the rat had to return to the start sensor to initiate the next trial.

Each recording session assessed choice behavior and neural activity across 64 trials that were separated into 4 probability blocks of 16 trials each (Fig. 1B). The probability blocks

of trials delivered a large reward with 100%, 50%, 25% or 12.5% probability (presented in descending order) when rats chose the large reward/risky option. During each probability block, the first 8 trials were forced choice trials (in which only one option was available) while the next 8 trials were free choice trials (in which both doors were open). The forced choice trials informed the rats as to the current probability of obtaining a large reward when choosing the large reward/risky option. Selection of small reward/certain option always resulted in a small reward 100% of the time. Thus, blocks 1 (forced choice) and 2 (free choice) were associated with 100% probability of receiving the large reward, 3 (forced) and 4 (free) with 50% probability, 5 (forced) and 6 (free) with 25% probability, and 7 (forced) and 8 (free) with 12.5% probability if the large reward/risky option was chosen. After rats were consistently completing all trials in a session across multiple days and displayed discounting behavior, they underwent a surgical procedure for the implantation of recording electrodes. Each rat ran an average of 13.5 probability discounting sessions before surgery.

Surgery

Each rat was placed in an induction chamber and deeply anesthetized under isoflurane (4% mix with oxygen at a flow rate of 1 L/min). While unconscious, the animal was placed in a stereotaxic instrument (David Kopf Instruments) and anesthesia was maintained throughout surgery by isoflurane (1 ~ 2.5%) delivered via a nosecone. The skull was exposed and adjusted to place bregma and lambda on the same horizontal plane. A microdrive array (described below) was unilaterally implanted so that the tetrodes could be placed dorsal to the hippocampus (3.6–3.7 mm posterior from Bregma, 3.5 mm lateral to the midsagittal suture, and 2.0 mm ventral from the dural surface). The drive assembly was secured in place with anchoring screws and dental cement. After surgery, an analgesic (Metacam, 1mg/kg) and an antibiotic (Baytril, 5mg/kg) were administered. Rats were allowed to recover for 7 days, during which they were fed *ad libitum* and handled daily.

Single-unit recording and postsurgical procedures

Recording tetrodes were constructed from 20 μ m lacquer-coated tungsten wires (California Fine Wire) and mounted on an array of twelve independently adjustable microdrives (Harlan 12 Drive, Neuralynx). Tetrode tips were gold-plated to reduce impedance to 0.2–0.4 M Ω (tested at 1 kHz). After a week of recovery from surgery, rats were returned to a food-restricted diet and spontaneous neural activity in the hippocampus was monitored as follows on a daily basis: the electrode interface board of the microdrives was connected to preamplifiers, and the outputs were transferred to a Cheetah digital data acquisition system (Neuralynx). Hippocampal single unit (0.6 – 6.0 kHz) and LFP (0.3 – 300 Hz) signals were filtered then digitized at 16 kHz. Neuronal spikes were recorded for 2 ms after the voltage deflection exceeded a predetermined threshold at 500–7000 $\times$ amplification. At the end of each recording session, tetrodes were lowered in 40 μ m increments, up to 160 μ m per day, to target new units for the following day. Experimental sessions continued until tetrodes passed through the hippocampus based on the distance traveled from the brain surface. A video camera mounted on the ceiling tracked infrared LED signals attached to the preamplifier and subsequent position data were fed to the acquisition system.

Once clearly isolated and stable units were found, the rats were recorded during a probability discounting session. The session began by placing the rat on the start segment of the maze. Two sec after the start sensor was triggered, the start door opened and the session continued until all 64 trials were completed. The sensor-driven opening and closing of maze doors was controlled automatically by a custom program (Z-basic). A timestamp of each triggered sensor was fed into the Neuralynx system so that they can be incorporated into future spike and video analyses

Histology

After the completion of all recording sessions, tetrode locations were verified histologically. Rats were deeply anesthetized under 4% isoflurane. The final position of each tetrode was marked by passing a 15 μ A current through a subset of the tetrode tips for 15 s. Then, the animals were given an overdose of sodium pentobarbital and transcardially perfused with 0.9% saline and a 10% formaldehyde solution. Brains were stored in a 10% formalin–30% sucrose solution at 4°C for 72 h. The brains were frozen, and then cut in coronal sections (40 μ m) on a freezing microtome. The sections were then mounted on gelatin-coated slides, stained with Cresyl violet, and examined under light microscopy. Only cells verified to be recorded in hippocampus were included in the data analysis.

Data analysis

Hippocampal single units were isolated using an Offline Sorter (Plexon). Various waveform features, such as the relative peak, valley, width, and principle component, were compared across multiple units simultaneously recorded from four wires of a tetrode. Only units showing good recording stability across blocks were included. Further analysis of the sorted units was performed with custom Matlab software (Mathworks).

Putative hippocampal pyramidal neurons were initially identified, as in prior studies, according to mean firing rate (less than 10 Hz) and exhibition of high-frequency complex-spike bursts as revealed in autocorrelograms and interspike interval histograms. Data from cells that did not exhibit complex spike bursts were not included in subsequent analyses.

Place Field Identification and Characterization—Maps were constructed for the spike and occupancy analysis by binning pixels using square bins ~2.5 cm. The spike map and occupancy map were separately smoothed using 2D convolution with a nine bin discrete Gaussian kernel (sigma ~0.8493) centered on each bin. The rate map was generated by dividing each pixel of the smoothed spike map by each pixel in the smoothed occupancy map. Place fields were located by constructing a binary rate map from the smoothed rate map using a threshold of 20% of the peak session smooth rate. In addition, the block conditions being compared had to show place fields with a peak rate of at least 1 Hz, and the primary place field was defined as the largest field that had an infield rate greater than or equal to 20% of the peak rate. Given the location of each field, values for the number of spikes within a field were determined using the smoothed rate maps. Similar criteria have been used in previous studies (e.g. Smith et al., 2006, 2012; Gill et al., 2011; Martig et al., 2011), and 97% of the cells analyzed showed mean firing rates less than 3 Hz, with the average rate less than 0.8 Hz. Therefore, the likelihood of including interneurons in our

sample is very low. Common metrics of place field stability (information content, reliability and spatial specificity) were calculated as previously published (Smith and Mizumori, 2006; Skaggs et al., 1993).

Stability of place field organization—The stability of location selective firing of place cells was calculated only for trial blocks that had a primary field that met the following criteria: number of infield spikes ≥ 1 ; area of the field $\geq 30\%$ of the total visited pixels for that block; infield/outfield rate ratio score ≥ 0.3 (which corresponds to the infield rate being twice that of that outfield rate); reliability ≥ 0.3 ; mean session rate ≥ 0.1 Hz and < 10 Hz.

Linearized Maze—The maze was linearized by using an ellipse to fit the path of the maze, which is roughly ellipsoidal. The ellipse was then segmented into 100 equal length segments and assigned a corresponding bin number. Each position from the video was assigned to the linear bin that was closest, as defined by the Euclidean distance. In order to correct for small deviations in the rat's position, the linear bins were smoothed using a lowess differentiation filter.

LFP Analysis—Values for velocity, firing rate, and power were only included when the rat's position monotonically increased with respect to the bin numbers (i.e. they were moving forward). Values were averaged when successive frames had repeating values for a bin. Values were also averaged across blocks if needed for a particular analysis. Power was calculated using the multitaper Fourier analysis, `mtspectrogramc`, from the Chronux toolbox, using a 500 ms window with a 50 ms step. The resulting spectrogram was filtered and averaged over each frequency band and converted to dB using the relation $10 \cdot \log_{10}(\mu V^2 / \text{Hz})$. The theta frequency band used in this analysis was 6–10 Hz.

Phase Precession—Theta phase precession refers to successive spikes of a place cell bursts that occur at progressively earlier phases of a theta cycle (O'Keefe and Recce, 1993). We examined whether the degree of phase precession varied as a function of the econometric and behavioral factors described above for place fields. Continuous recordings of local field potential (extracellular potential) were filtered for components in the theta frequency range (6–10 Hz) using a Chebyshev Type 2 bandpass filter. The Hilbert transform was applied to get instantaneous phase over time. Distance through the field was measured by calculating the length of the rat's path until the time when a spike occurs. Distance through the field for each pass was normalized for the total length of that pass through the field. Each pass had to have a minimum of 5 spikes in at least 4 theta cycles, a mean velocity of at least 3 cm/s, and a maximum interspike interval of 1 second (Schlesiger et al., 2015). Additionally the place fields had to have reliability and specificity scores greater than 0.3, an area less than a third of all visited pixels (in the block), and firing rates between 0.1 Hz and 10 Hz. Phases for each spike were calculated by matching the closest timestamp of the spike to the LFP timestamps and using the phase as calculated from the Hilbert transformation. Data from multiple passes in each block were pooled into one data set then analyzed together; single pass data were not analyzed. Theta phase precession was calculated as the slope of spike phase versus normalized distance into the field. A circular-linear regression method was applied (Kempster et al., 2012).

Statistical analysis—A multiple linear regression was performed on every cell to determine if agency, probability, or outcome could significantly predict changes in primary place field location and/or rate. To do this, the location and mean rate of the primary place field for each cell was determined for every trial in a session. If the primary place field did not meet the above criteria, the trial was assigned as having no field. Then, the econometric variables (i.e. agency, probability, and outcome) were assigned as the predictor variables in the regression to predict either location or rate (regressions were ran separately for location and rate). All regressions were determined using SPSS version 22, IBM. α was set at $p = 0.05$.

To assess whether probability of reward, choice outcome and/or agency similarly contributed to the changes observed in theta power throughout the probability discounting task, the maze was divided into eight discrete regions (defined in the results section) that were associated with different task phases. The power for each tetrode during each trial was determined and a multiple linear regression analysis was run for each region with agency, outcome, probability, and velocity as predictors (SPSS version 22, IBM). α was set at $p = 0.05$.

To determine whether or not place fields were distributed in a biased fashion across the maze, the frequencies of place fields in the same eight regions used for the LFP analysis were assessed using the Chi square test. The null hypothesis was that the place fields were evenly distributed across the different regions of the maze, adjusted for the size of each of the eight regions. Place fields distributions were considered unequal with α set at 0.05. Additionally, to specifically test if there were more place fields in the feeder region of the maze, a binomial test was performed comparing the observed and expected number of place fields in the feeder region and the rest of the maze. Expected numbers of place fields were based on the size of the region (i.e. the feeder region was 22% of the maze area so 22% of place fields would be expected there). $p = 0.05$ indicates there were significantly more or less fields in the feeder region than expected.

RESULTS

Probability discounting behavior on a maze

A total of 146 sessions were recorded (36.5 sessions per rat on average, ranging from 20 to 54 sessions) from 3 rats. In 89 of 146 recording sessions, rats exhibited probability discounting behavior, which was defined as the choice for the large risky reward during the highest probability (100%) always being at least 10% greater than the choice for the lowest probability (12.5%) Indeed, there was a significant main effect of probability on choice of the large reward (repeated measures ANOVA, $F(3, 88) = 111.5$, $p < 0.001$), and choice for the large reward during 100% probability was greater than the three other probabilities (Bonferroni post-hoc test for pairwise comparisons, $p < 0.05$). Figure 1C shows that during these sessions rats preferred the risky (large reward) option when the probability of receiving a large reward was high. As the probability of receiving the reward declined, so did the rats' preference for choosing the risky option.

Place fields remain spatially selective despite changes in the probability of reward or agency

Out of 909 recorded hippocampal neurons, 573 were identified as putative pyramidal cells according to standard criteria (Smith and Mizumori, 2006). Of these, 424 units were recorded from CA1 and 149 from CA3. These groups were combined for analyses since no differences were observed on any measure (Supplementary Fig. 1). Place cells showed strong location-specific firing regardless of changes in expected reward probability (i.e., 100%, 50%, 25%, or 12.5%), agency (i.e., forced choice or free choice), or choice outcome (i.e., win or lose), (Supplementary Fig. 2). The whole-session mean firing rate and the maximum firing rate within place fields did not significantly vary as a function of expected reward probability, agency, or choice outcome (all p 's > .05). Generally, information content also did not change as a function of reward probability or agency with the exception of win trials, which had significantly lower information content compared to the other outcomes (1-way ANOVA, $F(2,748) = 15.58$, $p < 0.0001$). This pattern may reflect the tendency of place fields to be located around the goal location after the feeder signal confirms to the rat that food will be delivered on a win trial (see below).

The locations of place fields were biased to occur close to the goal location when rewards are expected, but the degree of bias depended on the expected probability of receiving rewards

When considering the overall recording session, place fields were observed across the entire length of the (linearized) maze. A differential distribution of place fields became apparent, however, when trials were sorted according to reward probability, agency and trial outcome (Fig. 2; Supplementary Table 1). Significantly more place fields were found around the reward location for the 12.5% win trials in both forced and free choice conditions (binomial test, $p < 0.001$; Supplementary Table 2). While not significant, place fields also tended to accumulate around the reward site for free-choice certain trials while place fields were not preferentially found near the reward location on lose trials in any condition. This finding is consistent with prior reports that place fields tend to be found near significant locations such as reward and goal locations (e.g., Hollup et al., 2001). Specifically, these data show that the learned probability of receiving rewards and agency are important factors that determine goal-related place cell activity. This preference for place field locations emerged prior to arrival at the feeder, indicating that the auditory cue that signaled impending reward was at least in part responsible for the subsequent place field reorganization. The presence of the cue per se, however, could not have been the only factor initiating this response as the bias for place fields to be located near the reward location was most striking during low reward probability (win) trials. In sum, the scaling of the redistribution of place fields relative to the probability conditions of trials was observed only for win trials (under both forced and free choice conditions). In contrast, goal location preferences of place fields were observed only when rats freely chose the safe, certain option regardless of the probability condition.

While analyzed at the population level, the ability of the econometric variables to influence spatially selective firing was observed at the single-field level Figure 3 shows the rats' trajectory in gray and all of the recorded spikes for a cell over the entire session are depicted in red. Figure 3A and B show a hippocampal neuron's clear place field during free choice

trials (Choice) relative to forced choice trials (Forced) of the same probability block. This place field did not vary significantly across probability blocks, but rather by agency. Other place cells preferentially fired during trials of a particular outcome rather than agency or probability (Fig. 3C&D). Also shown is a place field around the location of the reward but only on lose trials. For another hippocampal neuron, it can be seen that different spatial firing patterns of place fields were observed when the probability of reward varied from 100% to 12.5% even though the sensory environment, behavioral responses, and motivation were the same (Fig. 3E, F). It can be seen that little firing was observed for the 100% and 25% blocks, and the highest firing rates were seen for the 50% and 12.5% blocks of trials. Forced and free choice trials were collapsed in this example as no differences were observed.

Place fields change primary field location and mean firing rate as a function of expected probability of reward, agency and choice outcome

A multiple regression was run to predict the location of the primary place field from agency, probability, and outcome (Fig. 4). First, neurons were not included in further analyses if the mean rate over the entire recording session was not at least 0.1 Hz or greater than 10 Hz. A total of 364 neurons met this initial rate criterion. The three econometric variables significantly predicted the location of the primary place field for 166 neurons, ($p < .05$, mean $R^2 = .24$). Only one cell had all three variables as significant predictor of location whereas 40 cells had 2 significant predictors of location. Of the 40 cells with two significant predictors of location, 10 were agency and outcome, 10 were agency and probability, and 20 were probability and outcome. 125 cells had only one significant predictor; $p < .05$. These variables did not significantly predict location of the primary place field for 197 neurons (54%).

Another multiple regression was run on the same 364 neurons to predict the rate of the primary place field as a function of agency, probability, and outcome. These variables statistically significantly predicted the mean firing rate of the primary place field for 168 (or 46% of) neurons, ($p < .05$, mean $R^2 = .23$). Only two cells had all three variables as significant predictor of rate whereas 44 cells had two significant predictors of rate. Of the 44 cells with two significant predictors of rate, 8 were agency and outcome, 10 were agency and probability, and 26 were probability and outcome. 124 cells had only one significant predictor; $p < .05$. These variables did not significantly predict the rate of the primary place field for 195 neurons (54%).

If one of the variables was a significant predictor of rate or location in the multiple regression, it is represented in Figure 4 where the mean beta (unstandardized) coefficient value represents the degree to which significant predictors influence either rate or location. At a population level, outcome had the strongest impact on place fields location followed by probability (which matches the place field distribution data described below). Of those place fields whose locations were affected by econometric factors, the mean unstandardized coefficient shows that agency and outcome had the biggest impact, causing the largest shift in location. When looking at rate, probability and outcome had similar impacts. Mean

unstandardized coefficient weight data show that of the impacted cells, agency caused the largest change in firing rate, regardless of location.

Hippocampal theta oscillations encode econometric information

LFP data were continuously recorded from the same tetrodes that recorded single unit data. Described below are theta band (6–10 Hz) power changes in response to different economic conditions of choice.

Theta power correlates with economic aspects of choice options—Figure 5 illustrates changes in theta power as rats traversed (a linearized representation of) the maze. Also shown (black dashed line) is the average velocity with which rats moved across the maze. Figure 5A separates theta power as a function of choice outcome (win, lose, certain).

To illustrate the impact of probability and agency conditions on the differential theta response to choice outcome, Figure 5B presents theta power as a function of outcome (win trials-top row, certain trials-middle row, and lose trials-bottom row) for forced choice trials (left column) and free choice trials (right column).

The changes in theta power observed (Fig. 5) suggest that theta is related to behavior in different ways during task performance. The loop maze was designed so that different segments of the maze were associated with the presentation/processing of specific types of information and opportunities (e.g., choices) to the rat. In this way, it was possible to compare the relative contribution of not only probability, agency, choice outcome, and velocity to behavioral and neural responses, but also to demonstrate the contributions of meaningful and specific information gleaned by rats as they traversed the maze. Figure 7A illustrates and defines these different informational regions of the maze. For each maze segment, theta power for each tetrode during each trial was determined and a multiple linear regression analysis was run for each region with agency, outcome, probability, and velocity as predictors. As predicted from numerous past studies (e.g., Whishaw and Vanderwolf, 1973; Buzsaki 2002), velocity was a significant predictor of theta power fluctuation in all segments (Fig 7B). Therefore, only the econometric variables that influenced theta power will be discussed below.

While the rat resided in the pre-decision area of the maze (*region 1*), probability and agency were significant predictors of theta power while outcome was not (model $F(4, 38919) = 665.18, p < 0.001$, unstandardized coefficient for agency: 0.07, $p = .02$, probability: 0.11, $p < 0.001$, outcome: 0.02, $p = .68$, and velocity: 0.16, $p < 0.001$). This might have been expected as agency and probability information (based on the previous trial) were known to the rat in this region. For the decision platform, (*region 2*) agency information should be salient to the animal because they are either forced to make a particular decision or are able to freely choose between the two options. Consistent with this view, we found that only agency could explain a significant amount of the variability observed in theta power in region 2 (model $F(4, 38243) = 2161.08, p < 0.001$, unstandardized coefficient for agency: 0.11, $p = 0.01$, probability: 0.01, $p = 0.71$, outcome: 0.05, $p = .06$, and velocity: 0.08, $p < 0.001$).

The immediate post decision region of the maze was demarcated by the sensor that was triggered when the animal passed through the decision doors but ended before the rat knew the outcome of its choice. While rats were on this maze segment (*region 3*), they may have retained information about the nature of the choices made, resulting in continued modulation of theta power by agency. In addition, probability information again significantly modulated theta power, while outcome did not, presumably because the outcome was not yet known for risky choices (model $F(4, 38253) = 291.50, p < 0.001$, unstandardized coefficient for agency: 0.23, $p < 0.001$, probability: 0.28, $p < 0.001$, outcome: 0.06, $p = 0.14$, and velocity: 0.11, $p < 0.001$). The next segment (*region 4*) was defined as the region after exiting the decision platform, a time when the rat heads toward the feeder trigger. All three econometric variables significantly influenced theta power in this region even though the rat did not yet know the outcome of risky choices (model $F(4, 38243) = 469.65, p < 0.001$, unstandardized coefficient for agency: 0.12, $p = 0.03$, probability: 0.25, $p < 0.001$, outcome: 0.36, $p < 0.001$, and velocity: 0.24, $p < 0.001$). It is possible that expected or predicted outcome of a choice influenced the rats' evaluation of the decision.

New information was provided in *region 5* as it contained the feeder trigger sensor which provided the rat with information regarding the outcome of its most recent choice, as well as information about the magnitude of the expected reward (i.e one click was audible per pellet dispensed). Therefore, it was hypothesized that if theta is related to specific learned goal information, outcome should now account for a disproportionately greater amount of the variance in theta power. Indeed, we observed a large increase in the variability in theta power that could be explained by outcome. Agency was still related to theta power in this region while probability was not (model $F(4, 38243) = 104.67, p < 0.001$, unstandardized coefficient for agency: 0.19, $p < 0.001$, probability: 0.01, $p = 0.64$, outcome: 1.24, $p < 0.001$, and velocity: 0.08, $p < 0.001$).

The post-feeder trigger region (*region 6*) was the maze area between the time when animals received the outcome information for that trial, and the time when the rat arrived at the food cup. In this region, outcome continued to have a significant effect on theta power, as did probability. Agency, however, had little or no impact on theta power at this point of maze performance (model $F(4, 38243) = 61.39, p < 0.001$, unstandardized coefficient for agency: 0.04, $p = 0.55$, probability: 0.13, $p < 0.001$, outcome: 0.48, $p < 0.001$, and velocity: 0.01, $p = 0.01$). This could be due, at least in part, to the fact that the ability to make a choice is irrelevant at this point in the task.

The rat arrived at the reward location during the next segment of the maze (*region 7*). The end of this region is demarcated by the sensor at the feeder. Theta power was dramatically elevated just prior to reward encounters, on average regardless of whether the reward was large (win condition) or small (certain condition). In contrast, a smaller increase in theta was observed at the reward location on lose trials. Such a distinguished response was also evident in single tetrode records (Fig. 6). This is the first identification of a specific type of economic information (and not just ongoing behavior) encoded by theta. In alignment with this qualitative observation that reward and goal information significantly shapes theta, outcome again predicted a large amount of variability in theta power in this region (model $F(4, 38243) = 331.32, p < 0.001$, unstandardized coefficient for agency: 0.06, $p = 0.24$,

probability: 0.22, $p < 0.001$, outcome: 1.42, $p < 0.001$, and velocity: 0.11, $p < 0.001$). During rewarded trials, it appears that the increase in theta power prior to reward encounters scaled inversely with the probability condition of the trial. That is, theta power is greatest during the lowest probability condition. The elevated theta could not be due to the presence of reward per se since the rats received the same type and magnitude of rewards for all win trials (regardless of probability condition), or due to the specific behaviors of the rats (since in all win trials, rats expected rewards at the feeder and stopped to consume the food). Rather theta appeared to respond to the known probability condition of a block. Probability scaling of theta was observed not only in the population summary (Fig. 5) but also for individual tetrode recording sites (Fig. 6B). Thus, not surprisingly, probability was also significant predictor of theta power in this region, which could explain the interaction between probability of receiving reward and the magnitude of reward being reflected in the scaling of theta power at the reward location (Fig 5B.). Finally, the last region of the maze (*region 8*), includes the maze area just after the food encounter. Here, we found that only agency and outcome had a modest influence on the variability observed in theta power (model $F(4, 38010) = 7.96$, $p < 0.001$, unstandardized coefficient for agency: 0.13, $p = 0.03$, probability: 0.05, $p = 0.13$, outcome: 0.09, $p = 0.03$, and velocity: 0.02, $p < 0.001$).

Since theta power increases around the goal appeared similar to that observed for place fields, we were interested to see if dynamic changes in theta power were related to place field characteristics, such as changes in in-field firing rate and location of primary place fields. To assess this, a Pearson's correlation analysis was performed to compare mean theta power over the session for each linearized bin of the maze with the in-field firing rate of a primary place field that occurred in the same linear bin. We also tested whether mean theta power over the session for each linearized bin correlated with the frequency of place fields that occurred in that corresponding bin. Theta power was not correlated with in-field firing rate of primary place fields ($r^2 = 0.02$, $p = 0.15$) but was significantly correlated with the frequency at which place fields occurred in a particular location ($r^2 = 0.12$, $p = 0.05$) (Supplementary Fig. 3A & B). From these data, it is apparent that instantaneous theta power has a significant relationship with the number of primary place fields that occurred at particular locations rather than the mean rate of those fields (Fig. S3).

Overall, econometric variables clearly and differentially influenced the strength of theta oscillations depending on the area of maze that the rat was experiencing, and thus what information was available and useful to the rat during task performance. This was evident from the finding that the outcome variable predicted the greatest amount of fluctuation in theta power after the outcome was known, supporting the hypothesis that theta is significantly influenced by goal and reward information.

Theta phase precession does not reflect economic information—No differences in theta phase precession were observed as a function of the expected probability of reward, agency, choice outcome. However, the percent of place fields showing significant phase precession did vary by segment of the maze (2-way ANOVA, main effect of region, $F(2,7) = 3.497$, $p = 0.02$. No effect of outcome $F(2,7) = 0.225$, $p = 0.80$) (Supplementary Fig. S4). It is noteworthy that the proportion of place fields showing phase precession steadily increased until rats were given information about the outcome of their last choice. This suggests that

the association between place fields and theta precession is related to information processing that remains engaged until the outcomes of choices are known.

DISCUSSION

Many past studies clearly show that knowledge about the location of goals impacts decisions and hippocampal neural processing during adaptive navigation. The present study tested directly whether economic considerations of a desired goal also bias behavioral decisions and neural processing by recording place cells and LFPs as rats ran a maze-based probability discounting task. Place fields aggregated near the reward location after the rat received a signal that reward was imminent (i.e. on win trials), and the strength of this biased distribution was most pronounced on low reward probability (win) trials. Since all win trials resulted in the same amount of reward, the probability effect cannot be due to the mere expectation or consumption of a reward. Place fields also distinguished forced and free choice trials in cases when the probability, magnitude, and location of rewards were the same. Theta power (commonly considered to exert a movement-related modulatory effect on place cell excitability, or to be involved in temporally sequencing place fields) was shown to dynamically and differentially vary as a function of expected reward probability, agency, and choice outcome depending on the information that was available to the rat as it traversed different maze regions. Specifically, the theta response scaled to the trial-specific *learned probability* of reward rather than movement or reward acquisition per se. Thus, in addition to goal location, economic factors such as the likelihood of finding reward at a particular location regulates patterned neural activity in hippocampus.

Place fields show selective encoding of learned probability, agency and choice outcome information

Many reports of hippocampal place field responses to changes in task and behavioral requirements or spatial context support the view that hippocampus is essential for learning new context-relevant information and for signaling changes in familiar contexts. The critical nature of this learned context information during navigation has been shown to include features of the external and internal environments (i.e. external sensory and intrinsic motivational information), and knowledge of the behavioral requirements of the task. This information is organized spatially and temporally, perhaps to facilitate efficient interactions with connected brain areas. Our results show that the economic condition or state likely comprises another type of information that contributes to the definition of recalled contextual memories that bias place field organization. Although the largest change in place field distributions occurred during the trials with the lowest probability of reward delivery, and these trials occurred at the end of recording sessions, it is not likely that unintentional electrode movement produced the observed effect. If that were the case, we should have observed significant place field reorganization for all trial types during the last block of recording. This was not the case: the pattern of place field re-distribution differed across win, lose and certain trial types, as we across free and forced choice condition during certain trials. Also, changing the economic conditions of the task did not alter the ability of place cells to exhibit clear and spatially-selective place fields. Therefore, it seems that altering economic conditions that biased choices resulted in dramatic redistribution of place fields

relative to key task events (such as obtaining rewards) and the presentation of task-relevant information.

The strong impact of learned econometric features on hippocampal neural activity suggests that such information is an integral component of the memory that defines place field characteristics in familiar contexts. The economic state surrounding choice options may organize the matrix of sensory, motivational, and behavioral information that place fields are known to represent, and in this way assist in the organization of adaptive place field sequences that may enable hippocampal-dependent memories (e.g., Lee and Wilson 2002; Foster and Wilson, 2006; Diba and Buzsaki, 2007; Johnson and Redish., 2007; Singer and Frank, 2007; Gupta et al., 2012). It is worth noting that responses of neurons of the midbrain-striatal-prefrontal decision making system are also modulated by the expected probability and magnitude of reward (Penner and Mizumori, 2012; Schultz et al., 1997; Lansink et al., 2009; Lisman and Redish, 2009; van der Meer and Redish, 2009; Schultz, 2007). Further, dopamine reward signals are context-dependent in a manner similar to that observed for hippocampal place fields (e.g., Puryear et al., 2010), and loss of dopamine function impairs hippocampal long term potentiation (Frey and Schroeder, 1990), and diminishes reward-place associations (McNamara, 2014). Thus the dopamine system may confer stability of reward processing, memory, and hippocampal physiology under stable context and task conditions. Indeed, reward related dopaminergic activity has been related to hippocampal memory processing (e.g., Wittman et al., 2005; McNamara et al., 2014) and VTA inactivation results in less stable hippocampal place fields (e.g., Martig et al., 2009). Analogous to the interpretation given for the enhanced dopamine response to rewards when they are the least expected (Fiorillo et al., 2003), one could interpret our observation of an overrepresentation of place fields at the goal during low probability win trials (Fig. 5) as signaling outcomes of greater surprise or salience. Such place field redistributions may have been signaled by dopamine responses to the reward predicting auditory cues. This place field salience signal may then update existing memories and attention mechanisms accordingly.

It is worth noting that the locus coeruleus also provides dopamine input to the hippocampus, and this input plays an important role in the hippocampal response to environmental novelty (Takeuchi et al., 2016). In the same study, optogenetic stimulation of the locus coeruleus enhanced memory in novel environments, and VTA inactivation did not attenuate the locus coeruleus stimulation-induced memory enhancement effect. Therefore, it was suggested that the locus coeruleus and VTA dopamine systems may make different functional contributions to hippocampus. Rather than coding environmental novelty, VTA dopamine may instead signal reward expectancy. Thus, future work should aim to assess whether the dopamine system is indeed responsible for the dynamic encoding of economic variables observed in the hippocampus, and if so, whether it is VTA dopamine that is responsible for signaling reward expectation.

Hippocampal theta dynamically reflects different types of econometric information during active navigation

Theta power dynamically varied during task performance depending on the changing nature of salient economic information available to the rat (Fig. 6). These variations could not be

accounted for by passage of time or sensory input since the changes in theta power did not correlate with specific behaviors, locations, reward or sensory input per se. This result illustrates that theta plays rather specific roles in information processing within hippocampus, perhaps to coordinate neuronal responses according to the details of expected economic information. In this way, theta may effectively up- or down-regulate neural plasticity mechanisms as available information varies ‘on-the-fly’. Indeed it is known that theta enhances synaptic plasticity mechanisms such as long term potentiation (Larson et al., 1986). Also, the proportion of place fields that showed phase precession increased as animals approached the location where they received information about the outcome of their choice. After receiving that information, the proportion of place fields that showed phase precession declined. Interestingly, this coincided with the time that theta power dramatically increased and place field reorganization around the goal location was observed. The increased theta around the goal location may reflect the occurrence of surprising, or low probability events that could trigger the necessary plasticity to update behavioral or neural states. In summary, the type of relationship between place fields and theta may vary according to the information that is currently available to, and being processed by, the animal.

The observed pattern of regulation of theta power by economic information is consistent with the hypothesis that in hippocampus, spatial context information is organized according to an understanding of choice options. Such an information framework could facilitate communication between other brain areas. For example, the reciprocally connected prefrontal cortex (Jay et al., 1989; Rajasethupathy et al., 2015) is known to be critical for adaptive response selection based on economic and predictive information (e.g., Pratt and Mizumori, 2001; Matsumoto et al., 2003; Baker and Ragozzino, 2014), and the prefrontal cortex and hippocampus show theta comodulation that coincides with decisions during spatial task performance (Jones and Wilson, 2005; Hyman et al., 2005). In the future, it will be of interest to determine if the economic regulation of hippocampal neural activity is abolished when the prefrontal cortex is taken off line. Conversely, the predictive and behavioral planning functions of prefrontal codes may be disrupted if hippocampus is prevented from transmitting information about the surprising success of choices.

Conclusion

The presents study reveals that the spatial organization of hippocampal place fields and theta power are strongly influenced by specific economic information that biases decisions and choices during goal-directed navigation. In this way, behavioral economic information may exert control over hippocampal neural activity by providing a broad decision-based framework within which predicted spatial and temporal contextual information is interpreted and processed to form or update episodic memories.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Grant support: This research was supported by NIMH grant 58755 and University of Washington (UW) Royalty Research Fund, UW Mary Gates Scholarship, and NIA grant 5T32AG000057

References

- Baker PM, Ragozzino ME. The prelimbic cortex and subthalamic nucleus contribute to cue-guided behavioral switching. *Neurobiol Learn Mem.* 2014; 107:65–78. [PubMed: 24246555]
- Belchior H, Lopes-dos-Santos V, Tort AB, Ribeiro S. Increase in hippocampal theta oscillations during spatial decision making. *Hippocampus.* 2014; 24:693–702. [PubMed: 24520011]
- Buzsáki G. Theta oscillations in the hippocampus. *Neuron.* 2002; 33:325–340. [PubMed: 11832222]
- Cei A, Girardeau G, Drieu C, Kambi KE, Zugaro M. Reversed theta sequences of hippocampal cell assemblies during backward travel. *Nat Neurosci.* 2014; 17:719–724. [PubMed: 24667574]
- Diba K, Buzsáki G. Forward and reverse hippocampal place-cell sequences during ripples. *Nat Neurosci.* 2007; 10:1241–1242. [PubMed: 17828259]
- Dragoi G, Tonegawa S. Preplay of future place cell sequences of hippocampal cellular assemblies. *Nature.* 2011; 469:397–401. [PubMed: 21179088]
- Duncan K, Ketz N, Inati SJ, Davachi L. Evidence for area CA1 as a match/mismatch detector: A high-resolution fMRI study of the human hippocampus. *Hippocampus.* 2012; 22:389–398. [PubMed: 21484934]
- Ferbinteanu J, Shirvalkar P, Shapiro ML. Memory modulates journey-dependent coding in the rat hippocampus. *J Neurosci.* 2011; 31:9135–9146. [PubMed: 21697365]
- Fiorillo CD, Tobler PN, Schultz W. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science.* 2003; 299:1898–1902. [PubMed: 12649484]
- Foster DJ, Wilson MA. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature.* 2006; 440:680–683. [PubMed: 16474382]
- Fremeau RT, Duncan GE, Fornaretto MG, Dearry A, Gingrich JA, Breese GR, Caron MG. Localization of D1 dopamine receptor mRNA in brain supports a role in cognitive, affective, and neuroendocrine aspects of dopaminergic neurotransmission. *Proceedings of the National Academy of Sciences.* 1991; 88(9):3772–3776.
- Frey U, Schroeder H. Dopaminergic antagonists prevent long term maintenance of posttetanic LTP in the CA1 region of rat hippocampal slices. *Brain Res.* 1990; 522:69–75. [PubMed: 1977494]
- Fyhn M, Molden S, Hollup S, Moser MB, Moser E. Hippocampal neurons responding to first-time dislocation of a target object. *Neuron.* 2002; 35:555–566. [PubMed: 12165476]
- Gill PR, Mizumori SJY, Smith DM. Hippocampal episode fields develop with learning. *Hippocampus.* 2011; 21:1240–1249. [PubMed: 20665593]
- Gupta AS, van der Meer MA, Touretzky DS, Redish AD. Segmentation of spatial experience by hippocampal theta sequences. *Nat Neurosci.* 2012; 15:1032–1039. [PubMed: 22706269]
- Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J Neurosci.* 2001; 21:1635–1644. [PubMed: 11222654]
- Hyman JM, Zilli EA, Paley AM, Hasselmo ME. Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus.* 2005; 15:739–749. [PubMed: 16015622]
- Jay TM, Glowinski J, Thierry AM. Selectivity of the hippocampal projection to the prelimbic area of the prefrontal cortex in the rat. *Brain Res.* 1989; 505:337–340. [PubMed: 2598054]
- Jin X, Costa RM. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature.* 2010; 466:457–462. [PubMed: 20651684]
- Johnson A, Redish AD. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci.* 2007; 27:12176–12189. [PubMed: 17989284]
- Jones MW, Wilson MA. Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol.* 2005; 3:e402. [PubMed: 16279838]

- Karlsson MP, Frank LM. Awake replay of remote experiences in the hippocampus. *Nat Neurosci*. 2009; 12:913–918. [PubMed: 19525943]
- Kempler R, Leibold C, Buzsaki G, Diba K, Schmidt R. Quantifying circular-linear associations: hippocampal phase precession. *J Neurosci Methods*. 2012; 207:113–124. [PubMed: 22487609]
- Lansink CS, Goltstein PM, Lankelma JV, McNaughton BL, Pennartz SMA. Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biol*. 2009; 7:e:1000173. [PubMed: 19688032]
- Larson J, Wong D, Lynch G. Patterned stimulation at the theta frequency is optimal for the induction of hippocampal long-term potentiation. *Brain Res*. 1986; 368:347–50. [PubMed: 3697730]
- Lee AK, Wilson MA. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*. 2002; 36:1183–1194. [PubMed: 12495631]
- Lee I, Griffin AL, Zilli EA, Eichenbaum H, Hasselmo ME. Gradual translocation of spatial correlates of neuronal firing in the hippocampus toward prospective reward locations. *Neuron*. 2006; 51:639–650. [PubMed: 16950161]
- Leutgeb JK, Leutgeb S, Treves A, Meyer R, Barnes CA, McNaughton BL, Moser MB, Moser EI. Progressive transformation of hippocampal neuronal representations in “morphed” environments. *Neuron*. 2005a; 48:345–358. [PubMed: 16242413]
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*. 2005b; 309:619–623. [PubMed: 16040709]
- Lisman J, Redish AD. Prediction, sequences, and the hippocampus. *Philos Trans R Soc B: Biol Sci*. 2009; 364:1193–1201.
- Louie K, Wilson MA. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*. 2001; 29:145–156. [PubMed: 11182087]
- Mankin EA, Diehl GW, Sparks FT, Leutgeb S, Leutgeb JK. Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts. *Neuron*. 2015; 85:190–201. [PubMed: 25569350]
- Manns JR, Zilli EA, Ong KC, Hasselmo ME, Eichenbaum H. Hippocampal CA1 spiking during encoding and retrieval: relation to theta phase. *Neurobiol Learn Mem*. 2007; 87:9–20. [PubMed: 16839788]
- Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J Neurosci*. 1995; 15:7079–7094. [PubMed: 7472463]
- Martig AK, Mizumori SJ. Ventral tegmental area disruption selectively affects CA1/CA2 but not CA3 place fields during a differential reward working memory task. *Hippocampus*. 2011; 21(2):172–184. [PubMed: 20082295]
- Martig AK, Jones GL, Smith KE, Mizumori SJY. Context dependent effects of ventral tegmental area inactivation on spatial working memory. *Behav Brain Res*. 2009; 203:316–320. [PubMed: 19447146]
- Matsumoto K, Suzuki W, Tanaka K. Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science*. 2003; 301:229–232. [PubMed: 12855813]
- McNamara C, Tejero-Cantero A, Trouche S, Campo-Urriza N, Dupret D. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nat Neurosci*. 2014; 17:1658–1660. [PubMed: 25326690]
- Mizumori SJY, McNaughton BL, Barnes CA, Fox KB. Preserved spatial coding in hippocampal CA1 pyramidal cells during reversible suppression of CA3 output: Evidence for pattern completion in hippocampus. *J Neurosci*. 1989; 9:3915–3928. [PubMed: 2585060]
- Mizumori SJY, Ragozzino KE, Cooper BG, Leutgeb S. Hippocampal representational organization and spatial context. *Hippocampus*. 1999; 9:444–451. [PubMed: 10495025]
- Mizumori, SJY. A context for hippocampal place cells during learning. In: Mizumori, SJY., editor. *Hippocampal place fields: Relevance to learning and memory*. New York: Oxford University Press; 2008. p. 16–43.
- Mizumori SJY, Tryon VL. Integrative hippocampal and decision-making neurocircuitry during goal-relevant predictions and encoding. *Prog Brain Res*. 2015; 219:217–242. [PubMed: 26072241]

- Muller RU, Kubie JL. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci*. 1987; 7:1951–1968. [PubMed: 3612226]
- O'Keefe J. Place units in the hippocampus of the freely moving rat. *Exp Neurol*. 1976; 51:78–109. [PubMed: 1261644]
- O'Keefe J, Recce ML. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*. 1993; 3:317–330. [PubMed: 8353611]
- Otmakhova, N., Duzel, E., Deutch, AY., Lisman, J. Intrinsically motivated learning in natural and artificial systems. Springer; Berlin Heidelberg: 2013. The hippocampal-VTA loop: the role of novelty and motivation in controlling the entry of information into long-term memory; p. 235-254.
- Penner MR, Mizumori SJY. Neural systems analysis of decision making during goal-directed navigation. *Prog Neurobiol*. 2012; 96:96–135. [PubMed: 21964237]
- Pfeiffer BE, Foster DJ. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*. 2013; 497:74–79. [PubMed: 23594744]
- Pratt WE, Mizumori SJY. Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behav Brain Res*. 2001; 123:165–183. [PubMed: 11399329]
- Puryear CB, Kim MJ, Mizumori SJY. Conjunctive encoding of reward and movement by ventral tegmental area neurons: context control during adaptive navigation. *Behav Neurosci*. 2010; 124:234–247. [PubMed: 20364883]
- Rajasethupathy R, Sankaran S, Marshel JH, Kim CK, Ferenczi E, Lee SY, Berndt A, Ramakrishnan C, Jaffe A, Lo M, Liston C, Deisseroth K. Projections from neocortex mediate top-down control of memory retrieval. *Nature*. 2015; doi: 10.1038/nature15389.
- Rezayof A, Zarrindast MR, Sahraei H, Haeri-Rohani A. Involvement of dopamine receptors of the dorsal hippocampus on the acquisition and expression of morphine-induced place preference in rats. *Journal of Psychopharmacology*. 2003; 17(4):415–423. [PubMed: 14870954]
- Schlesiger MI, Cannova CC, Boubilil BL, Hales JB, Mankin EA, Brandon MP, Leutgeb JK, Liebold C, Leutgeb S. The medial entorhinal cortex is necessary for temporal organization of hippocampal neuronal activity. *Nat Neurosci*. 2015; 18:1123–1132. [PubMed: 26120964]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Schultz W. Multiple dopamine functions at different time courses. *Ann Rev Neurosci*. 2007; 30:259–288. [PubMed: 17600522]
- Singer AC, Frank LM. Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron*. 2007; 64:910–921.
- Skaggs, WE., McNaughton, BL., Gothard, KM., Markus, EJ. An information-theoretic approach to deciphering the hippocampal code. In: Hanson, SJ.Cowan, JD., Giles, CL., editors. *Advances in neural information processing systems*. San Mateo CA: Morgan Kaufmann; 1993. p. 1030-1038.
- Smith DM, Barredo J, Mizumori SJY. Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus*. 2012; 22:1121–1133. [PubMed: 21630374]
- Smith DM, Mizumori SJY. Learning-related development of context-specific neuronal responses to places and events: the hippocampal role in context processing. *J Neurosci*. 2006; 26:3154–3163. [PubMed: 16554466]
- St Onge JR, Floresco SB. Dopaminergic modulation of risk-based decision making. *Neuropsychopharmacol*. 2009; 34:681–697.
- Takahashi T, Duszkiwicz AJ, Sonneborn A, Spooner PA, Yamasaki M, Watanabe M, Smith CC, Fernandez G, Deisseroth K, Greene RW, Morris RGM. Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature*. 2016; 537:357–362. [PubMed: 27602521]
- Tulving E. Episodic memory: from mind to brain. *Ann Rev Psychol*. 2002; 53:1–25. [PubMed: 11752477]
- van der Meer MAA, Redish AD. Covert expectation-of-reward in rat ventral striatum at decision points. *Front Neurosci*. 2009; 3:1–15. [PubMed: 19753087]

- Whishaw IQ, Vanderwolf CH. Hippocampal EEG and behavior: change in amplitude and frequency of RSA (theta rhythm) associated with spontaneous and learned movement patterns in rats and cats. *Behav Biol.* 1973; 8:461–484. [PubMed: 4350255]
- Wikenheiser AM, Redish AD. Hippocampal theta sequences reflect current goals. *Nat Neurosci.* 2015; 18:289–94. [PubMed: 25559082]
- Winson J. Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science.* 1978; 201:160–163. [PubMed: 663646]
- Wittmann BC, Schott BH, Guderian S, Frey JU, Heinze HJ, Düzel E. Reward-related FMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron.* 2005; 45:459–467. [PubMed: 15694331]
- Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron.* 2000; 27:623–633. [PubMed: 11055443]
- Yassa MA, Stark CE. Pattern separation in the hippocampus. *Trends Neurosci.* 2011; 34:515–525. [PubMed: 21788086]

Figure 1.
A. Schematic of the loop maze used to assess probability discounting behavior. Movement sensors were located at key locations on the maze (e.g., start door, door exit, feeder trigger, feeder). **B.** Each session assessed choice behavior and neural activity across 64 trials that were separated into 4 probability blocks of 16 trials each. The different probability blocks of trials denoted that a large reward would be delivered with 100%, 50%, 25% or 12.5% probability (presented in descending order) when rats chose the risky option. During each probability block, the first 8 trials were forced choice trials (in which only one door was

open) while the next 8 trials were free choice trials (in which both doors were open). The forced choice trials informed the rats as to the current probability of large reward. Selection of the safe, certain option always resulted in a small reward 100% of the time. **C.** A total of 146 sessions were recorded (36.5 sessions per rat on average, ranging from 20 to 54 sessions) from 3 rats. In 89 sessions, rats displayed clear discounting behavior (error bars represent SEM). Data shown were taken only from sessions during which rats exhibited discounting behavior. **D.** Histological reconstruction of the hippocampal recording sites show that most of the recordings occurred in dorsal CA1 with a smaller number in dorsal CA3. Since the responses of the two populations of place cells were similar, their data were combined for analysis.

Figure 2. The distribution of place fields varied depending on the expected reward probability as shown on linearized representations of the loop maze

Distributions of place fields during forced choice (left) or free choice (right) trials are shown as a function of outcome (win-top, certain-middle, and lose-bottom). Within each plot, we illustrate the proportion of place fields as a function of probability of reward. During win trials, a larger proportion of place fields were found around encounters with the feeder, and this was most pronounced on low probability trials. This pattern was also observed, albeit to a lesser extent, during free choice, certain trials, but not during forced choice, certain or lose trials.

Figure 3. Individual examples of place field responses to the expected probability of receiving a large reward, agency, and trial outcome

Plots shown in **A**, **C**, and **E** show the rats' trajectory in gray and all of the recorded spikes for a cell over the entire session is depicted in red. Heat maps in **B**, **D**, and **F** depict the spatial firing patterns during different conditions. **A, B:** A hippocampal neuron shows clear place field activity during free choice trials (Choice) relative to forced choice trials (Forced) of the same probability block. This place field did not vary significantly across probability blocks, but rather by agency. **C, D:** Other place cells preferentially fired during trials of a particular outcome rather than agency or probability. Shown is a place field that was observed around the location of the reward but only on lose trials. **E, F:** For another hippocampal neuron, it can be seen that different spatial firing patterns of place fields were observed when the probability of reward varied from 100% to 12.5% even though the sensory environment, behavioral responses, and motivation were the same. It can be seen that little firing is observed for the 100% and 25% blocks, and the highest firing rates were seen for the 50% and 12.5% blocks of trials. Forced and free choice trials were collapsed in this example as no differences were observed.

Figure 4. Place fields change primary field location and mean firing rate as a function of expected probability of reward, agency and choice outcome

A multiple linear regression was performed on 364 neurons to determine if agency, probability, or outcome could significantly predict changes in primary place field location and/or rate. **A.** Econometric variables significantly predicted the location the of the primary place field for 166 neurons, ($p < .05$, mean $R^2 = .24$). The y axis represents the proportion of the tested neurons for which a single econometric variable was a significant predictor of location of the primary place field. Agency was a significant predictor of location for 9.89% of neurons, probability for 18.68%, and outcome for 27.75% of the neurons tested. **B.** The

mean B coefficient values (y axis) are the degree to which significant predictors influence location of the primary place field. A whole integer represents the bin location on the linearized maze. **C.** Econometric variables significantly predicted the mean firing rate of the primary place field for 168 neurons, ($p = .05$, mean $R^2 = .23$). The y axis represents the proportion of the tested neurons for which a single econometric variable was a significant predictor of mean firing rate of the primary place field. Agency was a significant predictor of rate for 11.81% of neurons, probability for 24.45%, and outcome for 21.15% of the neurons tested. **D.** The mean B coefficient values (y axis) are the degree to which significant predictors influence the mean firing rate of the primary place field. α was set at $p = 0.05$. Error bars represent the $\pm$ SEM.

Figure 5. Power of the theta rhythm varied according to changing economic conditions of choice
A. Mean ($\pm$ SEM) theta power is compared across choice outcomes on a linearized representation of the loop maze. While some variability is observed, particularly striking is the peak in theta power that starts prior to reward (feeder) encounters. Most notable, the greatest theta response is observed around the feeder, particularly for win and certain trials. The mean velocity of the rat is shown by the black dashed line. It can be seen that the overall fluctuations in theta power did not track directly variations in velocity, suggesting that other factors contributed to the theta changes. **B.** The whole session data shown in A) are broken down according to agency (forced choice-left column; free choice-right column) and choice outcome (win, certain, lose) within each panel. Theta data are further broken down according to the expected probability of reward. Several patterns of results are clear. The increased theta power around the feeder is not only greatest for win trials, followed by certain trials (and very little for lose trials), but the magnitude of the increased theta power scales to the expected probability of reward. That is, the smallest increase in theta was observed during 100% win trials (after both forced and free choices), and the largest increase in theta power was observed during 12.5% win trials. A similar type of scaling was observed during the forced choice, certain trials.

Figure 6. Theta power modulation according to the expected reward outcome and probability is observed on single tetrode traces

A. Theta power was dramatically elevated just prior to reward encounters at the feeder, with the increase greatest during win trials, followed by certain trials. Lose trials showed the smallest increase. **B.** Single tetrode example of the inverse scaling of the increase in theta power as a function of the expected probability of reward during win trials. These data were collected during a single recording session. The fact that the scaled response can be seen in a single recording site illustrates that the group summary shown in Figure 5 is not to the averaging across a large number of sessions.

Figure 7. The impact of reward probability, agency and expected choice outcome on theta power varies depending upon the economic information available at different regions of the loop maze
A. The maze was divided into eight regions that were associated with varying types of information processing (see text). **B.** For each maze region, the amount of variability in theta power (for each tetraode during each trial) that could be accounted for by probability, agency and choice outcome was determined with a multiple linear regression analysis (* $p < .05$), using agency, probability, outcome and velocity as predictors. Velocity significantly modulated theta across the entire maze. In contrast, theta was differentially modulated depending on the region of the maze the rat was in, and thus the information that was available at particular task phases. Importantly, agency and probability tended to have the strongest influence until the rat received the reward predicting cue. After the cue

presentation, outcome was significantly related to theta power. Together, these data illustrate that theta power is modulated according different types of economic information that is available during goal-directed navigation.