

RESEARCH ARTICLE

Spatial encoding in dorsomedial prefrontal cortex and hippocampus is related during deliberation

Brendan M. Hasz¹ | A. David Redish² 

¹Graduate Program in Neuroscience,
University of Minnesota Twin Cities,
Minneapolis, Minnesota

²Department of Neuroscience, University of
Minnesota Twin Cities, Minneapolis,
Minnesota

Correspondence

A. David Redish, 6-145 Jackson Hall,
321 Church St SE, Minneapolis, MN 55455.
Email: redish@umn.edu

Funding information

National Institute of Mental Health, Grant/
Award Numbers: R01-MH080318,
R01-MH112688; National Science
Foundation, Grant/Award Number: DGE-
1069104

Abstract

Deliberation is thought to involve the internal simulation of the outcomes of candidate actions, the valuation of those outcomes, and the selection of the actions with the highest expected value. While it is known that deliberation involves prefrontal cortical areas, specifically the dorsomedial prefrontal cortex (dmPFC), as well as the hippocampus (HPC) and other brain regions, how these areas process prospective information and select actions is not well understood. We recorded simultaneously from ensembles in dmPFC and CA1 of dorsal HPC in rats during performance of a spatial contingency switching task, and examined the relationships between spatial and reward encoding in these two areas during deliberation at the choice point. We found that CA1 and dmPFC represented either goal locations or the current position simultaneously, but that when goal locations were encoded, HPC and dmPFC did not always represent the same goal location. Ensemble activity in dmPFC predicted when HPC would represent goal locations, but on a broad timescale on the order of seconds. Also, reward encoding in dmPFC increased during hippocampal theta cycles where CA1 ensembles represented the goal location. These results suggest that dmPFC and HPC share prospective information during deliberation, that dmPFC may influence whether HPC represents prospective information, and that information recalled about goal locations by HPC may be integrated into dmPFC reward representations on fast timescales.

KEYWORDS

deliberation, dorsomedial prefrontal cortex, hippocampus, spatial encoding, value

1 | INTRODUCTION

Deliberation is an approach to decision-making where the values of potential actions are estimated and compared to select the expected best action. Value-based decision making is thought to involve a deliberative process where outcomes of potential actions are estimated, the value of those hypothetical outcomes are evaluated and compared, and then the action with the highest expected value is taken (Gilbert & Wilson, 2007; Padoa-Schioppa, 2011; Rangel, Camerer, & Montague, 2008; Redish, 2013). However, the neural processes by which the brain performs deliberation are complex and poorly understood.

Deliberation is known to depend on the prefrontal cortex (PFC), hippocampus (HPC), ventral and dorsomedial striatum, among other areas (Doll, Simon, & Daw, 2012; van der Meer, Kurth-Nelson, & Redish, 2012). In rodents, the representation of prospective information by ensembles in HPC has been observed during deliberation at choice points (Johnson & Redish, 2007; Kay et al., 2020), which is thought to correspond to the simulation of the outcomes of candidate actions using an internal model stored in part by HPC (Redish, 2016). Furthermore, neural activity thought to correspond to the estimation of the value of these simulated outcomes has been observed in ventral striatum and in orbitofrontal cortex (Rich & Wallis, 2016; Steiner & Redish, 2012, 2014; Stott & Redish, 2014; van der Meer &

Redish, 2010; Wallis, 2018). However, it is unclear how these valuations are used to select between candidate actions, where and how these internal simulations are initiated, as well as where the information corresponding to the predicted outcomes are integrated.

Candidate brain regions for performing some or all of these roles include the various subregions of the prefrontal cortex (PFC). The prefrontal cortex has long been thought to support executive function (Dalley, Cardinal, & Robbins, 2004; Fuster, 2015; Kesner & Churchwell, 2011; Miller & Cohen, 2001). Areas of the dorsomedial prefrontal cortex (dmPFC) such as anterior cingulate cortex, have been implicated in conflict detection (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Dwyer, Dunn, Rhodes, & Killcross, 2010; Haddon & Killcross, 2005, 2006; Marquis, Killcross, & Haddon, 2007; Shenhav, Botvinick, & Cohen, 2013), suggesting it may be responsible for detecting the need for deliberative control. Also, dmPFC has been found to play an active role in the storage and recall of working memories (Cowen & McNaughton, 2007; Delatour & Gisquet-Verrier, 1999; Euston, Gruber, & McNaughton, 2012; Fuster & Alexander, 1971; Horst & Laubach, 2009; Preston & Eichenbaum, 2013; Ragozzino & Kesner, 1998; Tronel & Sara, 2003; Urban, Layfield, & Griffin, 2014; Yoon, Okada, Jung, & Kim, 2008), which may also translate to the storage of internally simulated outcome valuations. Some have theorized that dmPFC may initiate the internal construction of hypothetical situations, and use valuations of hypothetical outcomes to perform action selection (Hassabis & Maguire, 2009; van der Meer et al., 2012; Wang, Cohen, & Voss, 2015).

HPC and dmPFC, along with other structures, are thought to form an information-processing loop where top-down signals from dmPFC influence encoding in HPC, and information retrieved by HPC informs prefrontal representations (Eichenbaum, 2017; Jai & Frank, 2015; Redish, 2016; Shin & Jadhav, 2016; Wang et al., 2015). Lesion and inactivation studies have found that interrupting communication between PFC and HPC disrupts animals' spatial working memory (Barker et al., 2017; Floresco, Seamans, & Phillips, 1997; Maharjan, Dai, Glantz, & Jadhav, 2018; Wang & Cai, 2006). Also, theta oscillations in HPC and PFC are coherent during deliberation and appear to be important for spatial working memory (Benchenane et al., 2010; Hyman, Zilli, Paley, & Hasselmo, 2005; Jones & Wilson, 2005; O'Neill, Gordon, & Sigurdsson, 2013; Siapas, Lubenov, & Wilson, 2005). Furthermore, inactivating the medial PFC changes information representation by hippocampal place cells (Guise & Shapiro, 2017; Hok, Chah, Save, & Poucet, 2013; Navawongse & Eichenbaum, 2013; Schmidt, Duin, & Redish, 2019). There are both monosynaptic and polysynaptic projections from CA1 to dmPFC (Delatour & Witter, 2002; Ferino, Thierry, & Glowinski, 1987; Floresco & Grace, 2003; Hoover & Vertes, 2007; Jay & Witter, 1991; Swanson, 1981; Verwer, Meijer, Van Uum, & Witter, 1997), as well as a bisynaptic and bidirectional connections between CA1 and dmPFC via the nucleus reuniens and other thalamic nuclei (Cassel et al., 2013; Di Prisco & Vertes, 2006; Dolleman-van der Weel et al., 2019; Dolleman-Van der Weel, Da Silva, & Witter, 2017; Hoover & Vertes, 2012; McKenna & Vertes, 2004; Vertes, 2002, 2004; Vertes, Hoover, Do Valle,

Sherman, & Rodriguez, 2006; Vertes, Hoover, Szigeti-Buck, & Leranthe, 2007). Therefore, it is very likely that HPC and dmPFC share information relevant to deliberative decision-making on fast timescales.

However, it is unclear what specific information HPC and dmPFC share, and how quickly any information transfer occurs. Theories of the deliberative system suggest that prefrontal areas detect a need for deliberative control, instigate internal simulations of the outcomes of candidate actions (perhaps via HPC), keep track of the valuations of those outcomes (valuations which are perhaps generated by ventral striatum or orbitofrontal cortex), and use that information to decide which candidate action to enact. Therefore, likely candidates for information being passed between dmPFC and HPC include information about candidate actions, location, and reward or value. Previous work has discovered that hippocampal ensembles represent non-local spatial information which appears to correspond to internal simulations of candidate actions (Johnson & Redish, 2007; Kay et al., 2020), and other areas such as ventral striatum and orbitofrontal cortex represent value in ways suggesting they may be estimating the value of these internally simulated outcomes (Rich & Wallis, 2016; Steiner & Redish, 2012, 2014; Stott & Redish, 2014; van der Meer & Redish, 2010; Wallis, 2018). But what instigates these internal simulations? If dmPFC plays this role, then it should be possible to predict from activity in dmPFC whether non-local information is about to be represented in HPC. Also, if dmPFC is keeping track of the predicted values of candidate actions, then when the outcomes of these candidate actions are represented in HPC, downstream estimates of the value of these outcomes should affect reward encoding in dmPFC.

Here, we recorded simultaneously from ensembles in dmPFC and in CA1 of dorsal HPC in rats during performance of a spatial decision-making task. This task was a contingency switching task, where rats were required to choose actions at the choice point of the maze which was consistent with the current rule. The rule switched occasionally, preventing rats from simply automating their behavior on the task. We investigated the relationships between information encoded about spatial location and reward in dmPFC and CA1 as rats deliberated at the choice point of the maze. Spatial encoding by ensembles in dmPFC and CA1 was correlated during deliberation. While both dmPFC and CA1 encoded either prospective information or local information simultaneously, the prospective information encoded was not always identical. We found that activity in dmPFC predicted whether ensembles in HPC represented goal locations or the current location, but that the time scale of this relationship was broad, on the order of seconds. We also found that reward representation in dmPFC increased as CA1 ensembles represented the goal location, but that this increase in reward encoding was restricted to the hippocampal theta cycle during which prospective information was encoded in HPC. These results suggest that any influence of dmPFC on hippocampal prospective spatial representations may be temporally diffuse, while information about goal locations recalled by hippocampus may be integrated into prefrontal reward representations more quickly, on the order of single theta cycles.

2 | METHODS

2.1 | Subjects

Six FBNF-1 rats aged 8–14 months at the beginning of behavior were subjects for the experiment (four male, two female), bred from Fischer and Brown Norway rats (Envigo). Rats were housed on a 12 h light–dark cycle, and experimental sessions were conducted at the same time each day during the light phase. All experimental and animal care procedures complied with US National Institutes of Health guidelines for animal care and were approved by the Institutional Animal Care and Use Committee and the University of Minnesota.

2.2 | Training and behavior

The contingency switching task was a spatial reversal task where rats were required to adjust their behavioral strategies after uncued rule changes. The maze consisted of several low-cost choice points followed by a high-cost choice point between two actions: left or right (Figure 1a). The maze was constructed using LEGO Duplo blocks on a white canvas. The configuration of the low-cost choice points at the center of the maze was determined by a single wall in the middle of the maze, which switched back and forth from the left to right side randomly each day. Each lap, if rats chose the action at the high-cost choice point which was consistent with the current contingency, they received one unflavored food pellet (45 mg each, Test Diet, Richmond, IN) at one of two reward sites on the side of the maze, and an additional food pellet at the rear of the maze. If their choice was inconsistent with the current contingency, no reward was delivered and rats had to circle around to the start of the maze to initiate a new

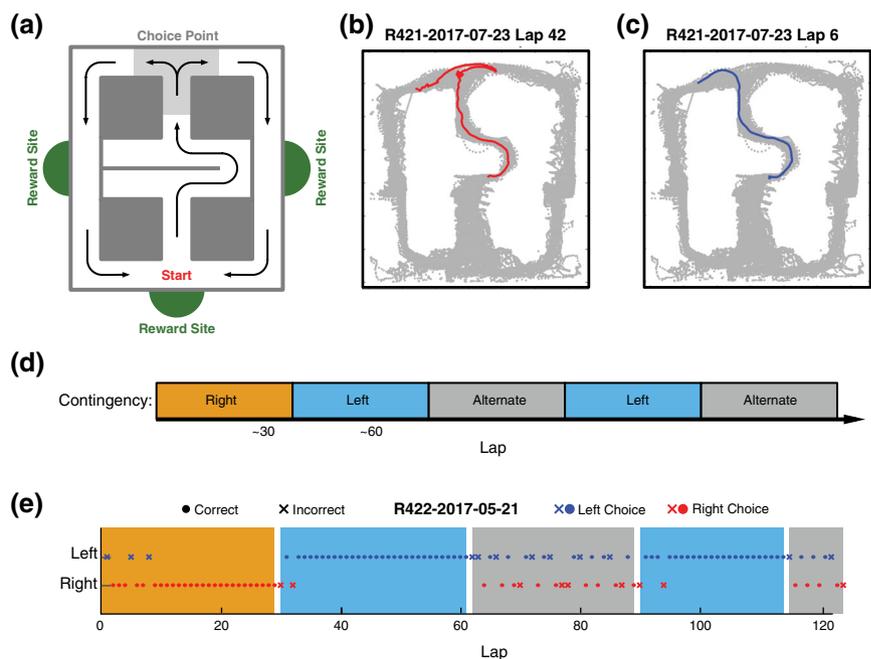
lap. Two different auditory cues also signaled to the rats whether their decision was correct or incorrect: a swept-frequency sinewave “chirp” from 1 to 3 kHz for correct, and two shorter 1 kHz square wave tones for incorrect. The contingency on any given lap was either Left (only left choices at the choice point lead to reward), Right (only right choices), or Alternate (the opposite choice from the previous lap was required for reward).

The first contingency type of a session was selected randomly, and contingencies were presented in blocks: every 30 ± 5 laps, the contingency changed randomly to one of the other two contingencies. Rats were allowed to run laps freely on the task for 1 hr each day, and their daily food allowance came only from performing the task. However, rats were fed extra food after running the task if their weight dropped below 80% of their free-feeding weight. During the experiment, none of the six rats dropped below 80% of their free-feeding weight, and so no post-feeding was required, but rats were post-fed after 7 out of 212 sessions during training (3%).

Rat positions on the maze were tracked using a video camera (Cohu, Inc., San Diego, CA) placed above the maze. Custom Matlab (The MathWorks, Natick, MA) software determined animal position from the video, and controlled the state of the task (the current contingency, food pellet release, the presentation of audio cues, etc.). The Matlab software also interfaced with an Arduino Uno Rev3 which ran custom software and triggered the release of food pellets from food pellet dispensers (MedAssociates, St. Albans, VT).

Rats were trained over the course of 4 weeks. Starting on the first week, rats were deprived of the freely available food in their home cages, but continued to have free access to water. Rats were handled and offered up to 15 g of food pellets each day for half an hour, to train them to eat the food pellets which would be available while performing the contingency switching task. For the second week, rats

FIGURE 1 Rat behavior on the contingency switching task. (a) The contingency switching task is a spatial maze consisting of a central maze segment, a choice point, and reward zones where rats receive rewards dependent on making choices consistent with the current contingency. (b) Example of a pass through the choice point where the rat displayed VTE and (c) a non-VTE pass. Bold lines show the position of the rat's head during a pass through the choice point. Gray dots show the position of the rat's head throughout the entire session. (d) Contingencies are presented in blocks of laps lasting 30 ± 5 trials. (e) Example behavioral data from a single session [Color figure can be viewed at wileyonlinelibrary.com]



performed a simplified version of the task where the contingency was either Left or Right, and the contingency stayed constant throughout each session but changed randomly from session to session. Rats were rewarded with two food pellets per feeder at all feeder sites for the second week. For the third week, again there were no within-session contingency switches, but all three contingencies were possible (including Alternate), and only one food pellet was delivered at the rear food delivery site. For the final week of training, only one food pellet was delivered at all feeder sites, but the task was otherwise the same as during Week 3.

After training, rats were given free access to food for at least 3 days, and then surgerized. After 3 days of post-surgery recovery, rats were again food deprived and re-trained for 1–2 weeks on the final training phase of the task (all three contingencies possible, but no within-session contingency switches, and one pellet per feeder). Finally, rats performed the full version of the task including within-session contingency switches and neural recordings for 2–3 weeks.

2.3 | Surgery and electrode targets

After training, rats were given free access to food for at least 3 days, and then chronically implanted with a hyperdrive containing 24 tetrodes, and a separate drive containing a 32-site silicon probe (Cambridge NeuroTech, Cambridge, UK). The hyperdrives for five rats contained two bundles of 12 tetrodes each, targeting the CA1 region of dorsal hippocampus bilaterally (3.8 mm posterior and ± 3.0 mm lateral from bregma). The hyperdrive for one rat contained a single bundle of 24 tetrodes targeting the right hippocampus. The silicon probes consisted of two 16-site shanks which were implanted 3.8 mm anterior to and 0.7 mm lateral from bregma at a 25° angle (targeting dmPFC on the right hemisphere, such that the final target was 2.3 mm A/P, -0.7 mm M/L, and 3.9 mm D/V, all coordinates relative to bregma). The hyperdrives and silicon probe drives were made in-house, and protective shrouds around the drives and amplifier boards were printed on a Form 2 3D printer (Formlabs, Somerville, MA).

Animals were anaesthetized with and maintained on isoflurane (0.5 – 2% isoflurane vaporized in O₂) for the duration of the surgery. Rats were placed in a stereotaxic apparatus (Kopf, Tujunga, CA) and were given penicillin (Combi-Pen-48) intramuscularly in each hindlimb, and carprofen (Rimadyl) subcutaneously. Rats' heads were shaved and disinfected with Betadine (Purdue Rederick, Norwalk, CT) before making an incision to reveal the skull. Three to five jewelers' screws were used to anchor the drives to the skull, one of which was used as ground for the tetrodes, and a separate screw used as ground for the probe.

Three craniotomies were opened: two for the bilateral tetrode bundles using a surgical trephine, and one for the silicon probes using a burr. The dura was removed with forceps, the probe and tetrode drives were positioned with the stereotax, and then silicone gel (Dow Corning, Midland, MI) was applied to the craniotomies. A layer of

MetaBond (Parkell, Edgewood, NY) and then dental acrylic (The Hygenic Corporation, Cuyahoga, OH) was applied to secure the drives to the skull. After surgery, the probes and tetrodes were turned down 640 μm . Rats were subcutaneously injected with carprofen on the day of surgery and for 2 days after surgery, as well as enrofloxacin (Enroflox) the day of surgery and for 5 days post-surgery.

2.4 | Data acquisition and electrophysiology

Neural data from all rats was acquired on an Intan RHD2000 recording system (Intan Technologies, Los Angeles, CA), using four RHD2132 amplifier boards (three for the tetrodes and one for the silicon probe). The digitized signals were passed through a 24-channel commutator (Moog, East Aurora, NY) to allow the rats to move freely throughout recording sessions. To synchronize behavior with the neural recordings, timestamps were sent from the Matlab (Version 2017a, The MathWorks, Inc., Natick, MA) software running the task to digital input ports on the Intan RHD2000 USB Interface Board via an Arduino Uno.

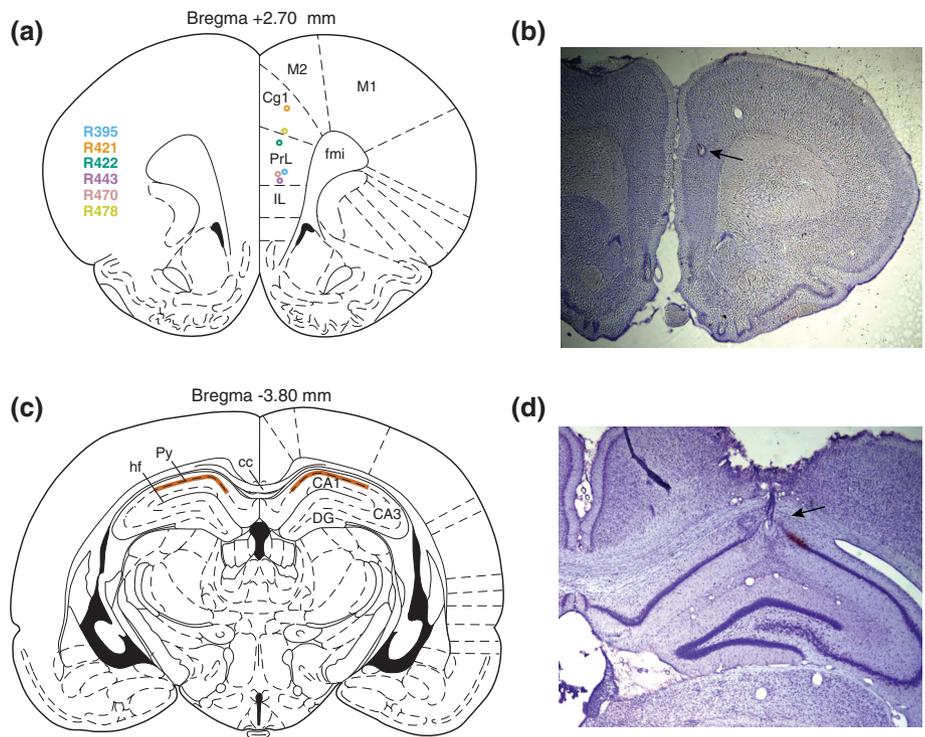
Tetrodes were slowly advanced toward the hippocampal pyramidal layer, and the probes toward dmPFC, over the course of around 2 weeks, as the rats recovered and were re-trained on the task. The pyramidal layer was identified by the size of ripples and the direction of sharp wave deflection, as well as spike bursts during these SWRs.

Signals were filtered and spikes and LFP signals were extracted using in-house software written in Matlab and C. Spikes recorded on tetrodes in the hippocampus were manually clustered using the MClust 4.4 software package (Redish, 2017). Only well-separated clusters were kept and used for analysis. The median isolation distance was 21.2, and the median L-ratio was 0.0899 (Schmitzer-Torbert, Jackson, Henze, Harris, & Redish, 2005). Spikes recorded on silicon probes were sorted offline using Kilosort (Pachitariu, Steinmetz, Kadir, Carandini, & Harris, 2016) into putative clusters, and then manually refined using Phy (Rossant, Kadir, Goodman, Hunter, & Harris, 2016).

2.5 | Histology

After rats were finished running the experiment, both tetrode and silicon probe recording locations were marked with electrolytic lesions. Ten μA was passed through a channel on each tetrode, and every fourth channel on the silicon probes, for 10 s. At least 2 days after the lesions were made, the rats were anesthetized with a pentobarbital sodium solution (150 mg/kg, Fatal-Plus) and then perfused transcardially with saline followed by 10% formalin. Brains were stored in formalin, and then in a 30% sucrose formalin solution until slicing. Coronal slices were made through the hippocampus and prefrontal cortex (sagittal slices were made instead in PFC for four rats) using a cryostat, and the slices were stained with cresyl violet and imaged to determine tetrode and silicon probe recording locations (Figure 2).

FIGURE 2 Electrode recording locations and histology. (a) Silicon probe recording locations in dmPFC for each rat. (b) Example photo of a cresyl violet stained coronal slice through PFC, showing the electrolytic lesion created to mark the recording location. (c) Tetrode recording region (highlighted in orange) was in the pyramidal layer of CA1 of dorsal HPC. (d) Example photo of a cresyl violet stained coronal slice through HPC, showing an electrolytic lesion and electrode track. Anatomy diagrams in (a) and (c) are from Paxinos and Watson (2006). *IL*, infralimbic cortex; *PrL*, prelimbic cortex; *Cg1*, cingulate cortex area 1. *M2*, secondary motor cortex; *M1*, primary motor cortex; *fmi*, forceps minor of the corpus callosum; *hf*, hippocampal fissure; *Py*, pyramidal layer. *cc*, corpus callosum; *DG*, dentate gyrus [Color figure can be viewed at wileyonlinelibrary.com]



2.6 | Statistics and general data analysis

All analyses were performed using Matlab (Version 2017a, The MathWorks, Inc., Natick, MA), except the binomial tests which were performed using the SciPy package for Python (Virtanen et al., 2020). For all neural analyses, we included only sessions where ≥ 10 cells were recorded in dmPFC and ≥ 10 cells were recorded in CA1 simultaneously.

2.7 | Vicarious trial and error (VTE)

To measure vicarious trial and error behaviors, we used *zldPhi*, a metric which captures both how long the rat hesitates at the choice point, and how quickly it moves its head back and forth (Papale, Stott, Powell, Regier, & Redish, 2012). When x and y are the x and y positions of the rat's head,

$$ldPhi = \int_{\text{zone entry}}^{\text{zone exit}} \left| \frac{\delta}{\delta t} \text{atan2} \left(\frac{\delta y}{\delta t}, \frac{\delta x}{\delta t} \right) \right| \delta t$$

and *zldPhi* is the *ldPhi*, z -scored within session.

To distinguish VTE events from non-VTE events, we fit a half-Gaussian distribution to values less than the mode of the *zldPhi* distribution. We then assumed that *zldPhi* values under a full Gaussian distribution with the same mean and SD as the fit half-Gaussian corresponded to non-VTE events, and passes through the choice point with greater *zldPhi* values corresponded to VTE events.

2.8 | Bayesian decoding

We used cross-validated Bayesian decoding (Zhang, Ginzburg, McNaughton, & Sejnowski, 1998) to decode several variables from ensemble activity in dmPFC or CA1. In Figures 3, 4, and 6, we decoded spatial position from ensemble firing rates in CA1 or dmPFC, using a 32×32 grid of spatial bins. In Figure 4, we trained a binary Bayesian decoder on ensemble activity in dmPFC to predict whether CA1 represented the choice point or the reward zone. In Figure 5a, we trained a binary decoder on ensemble activity in dmPFC to predict whether reward was received on the current trial or not. In we also trained a binary decoder on ensemble activity in dmPFC to predict whether reward was received on the current trial or not. However, the training period used for this decoder was from 1 to 3 s after reward zone entry, and the test period was during rats' deliberation at the choice point.

The time bins used for decoding corresponded to peak-to-peak times of the local field potential recorded in the hippocampal fissure, band-passed between 6 and 11 Hz (theta oscillations). We used these hippocampal theta time bins even for the decoders trained on ensemble activity in dmPFC. In Figure 5a, we instead used 200 ms time bins aligned to the time of rats' entry into the goal zone.

To compute the decoded posterior probability of different zones, we took the average posterior probability within the different zones shown in Figure 3d. The zones used were the choice point (green dotted box in Figure 3d), the left goal site (red dotted box), the right goal site (blue dotted box), and either goal site (the area within both the red and blue dotted boxes).

3 | RESULTS

3.1 | Rat behavior on the contingency switching task

We recorded from dmPFC and CA1 simultaneously in six rats as they performed the contingency switching task, in which rats ran through a spatial maze with one main choice point between two actions: left or right (Figure 1a). Rats received a food pellet reward if they chose the action at the choice point which was correct under the current contingency. If their choice was inconsistent with the current contingency, no reward was delivered and rats had to circle around to the start of the maze to initiate a new lap. The contingency on any given lap was either Left (only left choices led to reward), Right (only right choices), or Alternate (the opposite choice from the previous lap was required for reward). Contingencies were presented in blocks of laps: every 30 ± 5 laps, the contingency switched to one of the other two contingencies (Figure 1d).

Rats ran 137.7 ± 31.7 laps per session (mean \pm SD), and there were 4.1 ± 1.2 contingency switches per session. Rats made choices which were consistent with the current contingency on 78.0 ± 3.0 percent of laps. Rats sometimes displayed VTE behaviors as they passed through the choice point, where rats moved their heads back and forth as if considering which path to take (Figure 1b). VTE is considered a behavioral marker of deliberation in rodents (Redish, 2016). For more details of rat behavior during the contingency switching task see Hasz and Redish (2020).

3.2 | Non-local spatial representations were correlated between CA1 and dmPFC

To determine how spatial information represented in dmPFC and CA1 was related during deliberation, we performed cross-validated Bayesian decoding of spatial position from ensemble spiking activity in dmPFC, and also separately for simultaneously recorded ensembles in CA1. Using the decoded spatial posterior distributions, we computed how far ahead of or behind rats' actual positions CA1 and dmPFC represented, and how the spatial representations in each area were related (Figure 3a).

Previous work has found that HPC represents locations further ahead of the rat's actual location during vicarious trial and error (Amemiya & Redish, 2016; Amemiya & Redish, 2018; Johnson & Redish, 2007; Kay et al., 2020; Papale, Zielinski, Frank, Jadhav, & Redish, 2016). Replicating this previous work, we found that during passes through the choice point on the contingency switch task, indeed CA1 encoded positions further ahead of the rat during VTE than during non-VTE choice point passes ($p = .0156$, one-sided Wilcoxon signed rank test, $N = 6$ rats, Figure 3c). Similarly, the dmPFC also encoded positions further ahead of the rat during VTE at the choice point ($p = .0469$, one-sided Wilcoxon signed rank test, $N = 6$ rats), though the size of this effect was less prominent than that

observed in CA1 (Figure 3c). This suggests that both CA1 and dmPFC represented prospective spatial information during deliberation.

However, simply because both areas represented prospective information during deliberation does not necessarily mean they represented identical spatial information at the same time. To determine whether spatial representations in dmPFC and CA1 were tightly locked on a fast timescale, we correlated the distance ahead or behind rats being represented by ensembles over the course of single theta cycles. We found that there was a small correlation between the relative position represented in CA1 and the relative position represented in dmPFC (a mean Spearman's correlation coefficient of .0458, which was significantly greater than zero, $p = .0469$, one-sided Wilcoxon signed rank test, $N = 6$ rats, Figure 3b).

To determine whether these correlations were due to both areas representing goal locations simultaneously, as opposed to simply being a result of minute positional differences in the spatial representations across theta cycles, we analyzed the decoded spatial locations by which zone was being represented. Because the Bayesian decoding resulted in a decoding posterior across the entire maze, we were able to define three zones of interest: the choice point, the reward zone on the left side of the maze, and the reward zone on the right side of the maze (Figure 3d). To perform the following correlations, we computed the sum of the decoding posterior within each zone (normalized by area) per theta cycle. Encoding of the choice point versus either reward zone was correlated between dmPFC and CA1 (a median Spearman's correlation coefficient of .129, which was significantly greater than zero, $p = .0313$, Wilcoxon signed rank test, $N = 6$ rats, Figure 3e). However the identity of the reward site being encoded by dmPFC and CA1 was not significantly correlated ($p = .156$, Wilcoxon signed rank test, $N = 6$ rats, Figure 3f). The correlation coefficients for choice point and either reward zone representation were greater than the correlation coefficients for the identity of the encoded reward site (paired Wilcoxon signed rank test, $p = .0313$, $N = 6$ rats, Figure 3e vs Figure 3f).

These results suggest that while both the dmPFC and HPC may have been in deliberative modes simultaneously, and both seemed to represent either local or prospective information within the same theta cycles, that they were not necessarily representing identical information at the same time.

3.3 | dmPFC predicted non-local spatial representation in CA1

If dmPFC instigates internal simulations and evaluations of candidate actions, carried out by the hippocampus and other structures, then neural signatures corresponding to that initiation should be present in dmPFC ensemble activity. Therefore, it should be possible to predict from ensemble activity in PFC whether hippocampus will represent prospective information.

To determine whether activity in dmPFC carried information about whether CA1 was representing prospective information, we

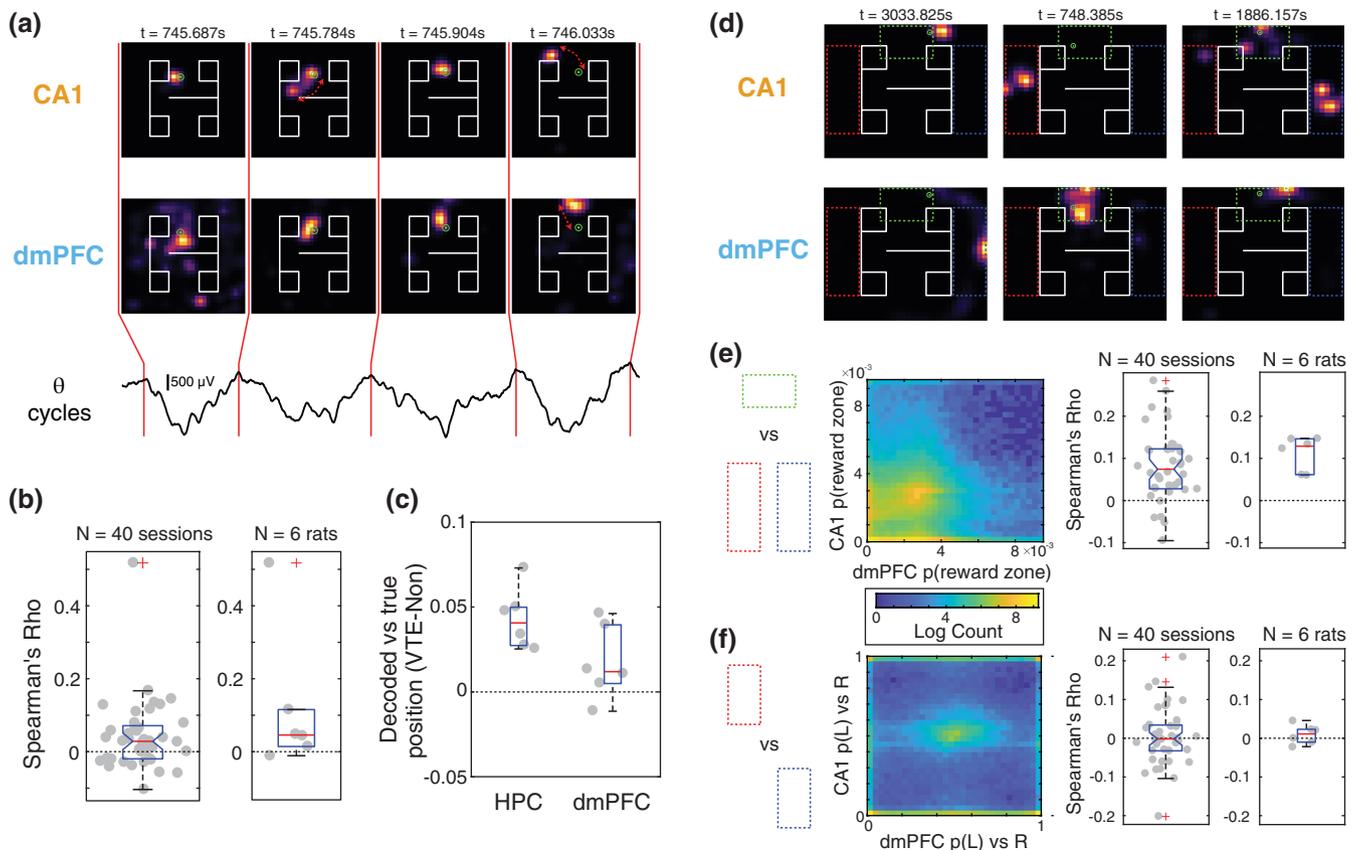


FIGURE 3 Correlations between location encoding in dmPFC and CA1. (a) Examples of Bayesian decoding of spatial location was performed on ensemble activity in dmPFC and CA1 using time bins corresponding to hippocampal theta cycles. Shown are example decoded posterior distributions from four consecutive theta cycles. Green dots indicate the actual position of the animal's head. Red arrows indicate discrepancies between the animal's true and the decoded locations. (b) Correlation coefficients of the distance relative to the rats true position decoded from ensemble activity between dmPFC and CA1, for each session and for each rat. (c) Difference in the average decoded position relative to the actual position of the rat at the choice point between VTE and non-VTE passes. Higher values indicate positions further ahead of the rat were represented during VTE passes. Units of the y-axis are laps (full revolutions around the maze). (d) Examples of Bayesian decoding and how it was used to compute the posterior probability of three different zones: the choice point (green dotted box), the left reward zone (red dotted box), and the right reward zone (blue dotted box). The top row shows decoding from CA1, the bottom row shows decoding from dmPFC, and each column shows decoding from the same theta cycle. (e) Correlation between dmPFC and CA1 encoding of the choice point versus either reward site. The probability that CA1 represented either reward zone as opposed to the choice point was computed from CA1 2D spatial posteriors for each theta cycle (left panel, y-axis of the 2D histogram). The probability that dmPFC represented either reward zone as opposed to the choice point was also computed from dmPFC 2D spatial posteriors for each theta cycle (left panel, x-axis of the 2D histogram). These values, where each datapoint corresponded to a theta cycle during a pass through the choice point, were correlated to produce the correlation coefficients grouped by session (middle panel) or rat (rightmost panel). (f) Correlation between dmPFC and CA1 encoding of the reward site identity. Same as (e), except the probability that either area represented the right reward zone as opposed to the left reward zone (instead of either reward zone vs. the choice point). Figures S2–S5 show these analyses for the first 10 and last 10 laps of each block [Color figure can be viewed at wileyonlinelibrary.com]

first used cross-validated Bayesian decoding to decode the location represented by CA1 on a per-theta-cycle basis. We categorized hippocampal theta cycles while rats passed through the choice point as either “local” (the top 10% of theta cycles with the highest posterior density in the choice point) or “prospective” (the top 10% of theta cycles with the highest posterior density in the reward zones – either the left or right reward zone, as in Figure 3e). Then, we trained a classifier (also a cross-validated Bayesian decoder) to predict from ensemble activity in dmPFC whether simultaneous hippocampal activity was representing local or prospective information. We measured the

performance of the classifier using the area under the receiver operating characteristic curve (AUROC) metric. An AUROC value of 0.5 indicates a classifier is performing at chance, while an AUROC value of 1.0 would indicate the decoder is correctly classifying every single theta cycle.

We found that this classifier trained on dmPFC activity performed well above chance, with AUROC values generally in the 0.6–0.8 range (Figure 4a). To ensure that the classifier was performing above chance, we repeatedly trained models on shuffled dmPFC firing rates (we used 1,000 shuffles per session). The performance of the classifier

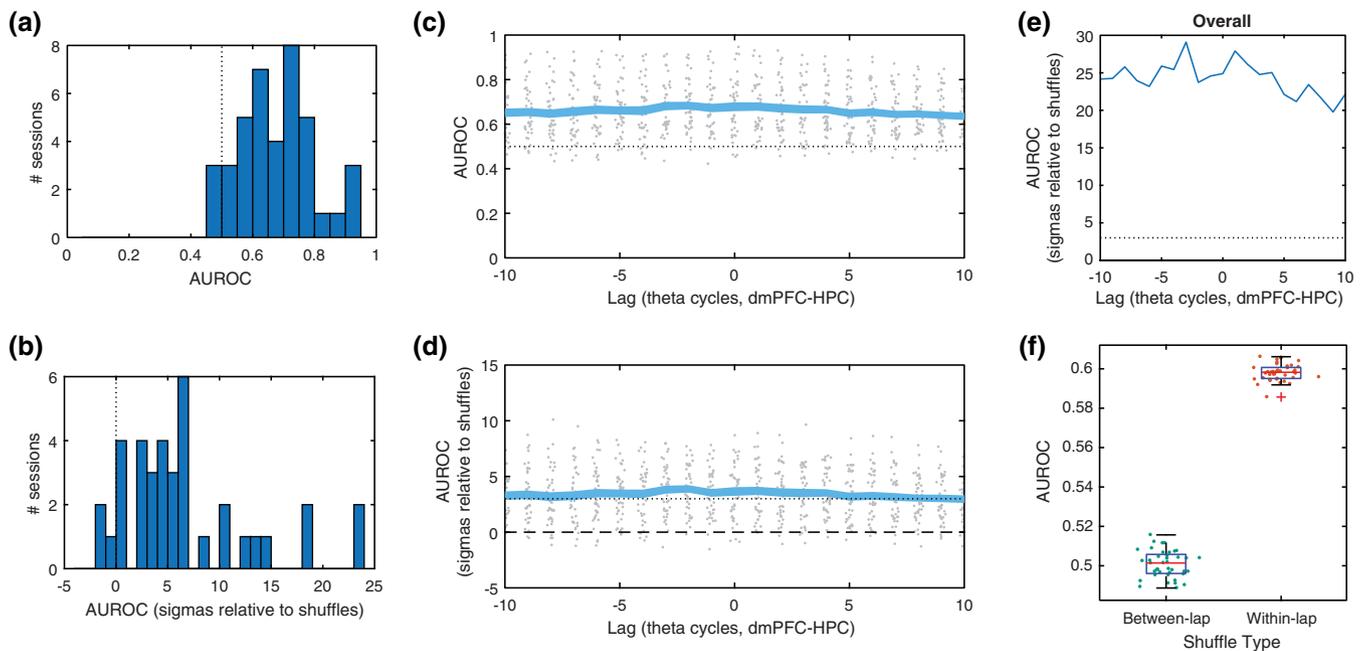


FIGURE 4 Predictions of prospective representation in CA1 from activity in dmPFC. (a) Area under the receiver operating characteristic curve per session for the classifier predicting whether CA1 is representing local or prospective spatial information, trained on dmPFC firing rates. (b) The AUROC relative to the AUROC of models fit to shuffled dmPFC firing rates. The x-axis is the number of SDs above the shuffle distribution. (c) AUROC of the classifier as a function of when dmPFC firing rates were used to predict representation in CA1. Negative values on the horizontal axis correspond to when dmPFC firing rates were used to predict upcoming representation in CA1, and positive values on the horizontal axis correspond to when dmPFC firing rates were used to predict past representation in CA1. Line and shaded area shows the mean \pm standard error, $N = 40$ sessions, while the gray dots show the AUROC for each individual session. Horizontal dotted line shows 0.5, corresponding to the AUROC expected from noise. (d) Same as in panel (c), but showing the SDs of the AUROC above the corresponding shuffle distributions as in panel (b). The horizontal dotted lines show the median of the shuffle distribution and 3σ . (e) The AUROC relative to shuffle distributions for all sessions combined. (f) The AUROC of Bayesian decoders trained to predict whether HPC ensembles were representing nonlocal information from dmPFC ensemble firing rates, when dmPFC firing rates were shuffled within or between laps. Figure S6 shows these analyses for the first 10 and last 10 laps of each block [Color figure can be viewed at wileyonlinelibrary.com]

trained on actual dmPFC firing rates was better than shuffles for all but three sessions, and was significantly above the shuffle distributions for the vast majority of sessions (Figure 4b). The AUROC value for the decoder across all sessions was over 20 SDs above the distribution resulting from applying the decoder to shuffled firing rates (Figure 4e, lag of 0).

However, the predictive power of dmPFC activity did not seem to be highly temporally specific. We repeated the classification analysis, but using neural activity from dmPFC at different time lags relative to the HPC theta cycle for which the classifier was trying to predict. If, say, activity in dmPFC were causing the next theta cycle in HPC to represent non-local information, but had no influence over hippocampal representations during preceding or subsequent theta cycles, then we would expect to see a sharp peak in the classifier's performance at a lag of -1 theta cycle. In contrast to that hypothesis, the performance of the classifier had a very broad performance profile as a function of lag (Figure 4c–e). Furthermore, the performance of classifiers trained on dmPFC firing rates which were shuffled within-lap was better than the performance of classifiers trained on dmPFC firing rates which were shuffled between-lap (Figure 4f). These results suggest that medial prefrontal cortex may be affecting the hippocampal state

in a way which corresponds to strategy in general, such as whether to deliberate or not (Hasz & Redish, 2020; Powell & Redish, 2014), and that this influence may last over longer time periods such as entire trials.

3.4 | Non-local representation in CA1 corresponded to an increase in reward encoding in dmPFC

Our results thus far suggest that dmPFC could play some role in influencing hippocampal circuitry to enter into prospective modes, but it is also possible that hippocampal activity has an effect on prefrontal representations. Next, we investigated if and how the encoding of reward in prefrontal cortex is affected by spatial representations in hippocampus.

First, we examined whether ensembles in dmPFC encoded whether animals received reward or not. We performed cross-validated Bayesian decoding on ensemble activity in dmPFC as rats entered the reward zone, and decoded whether rats received reward on that lap or whether there was a lack of reward. The task included

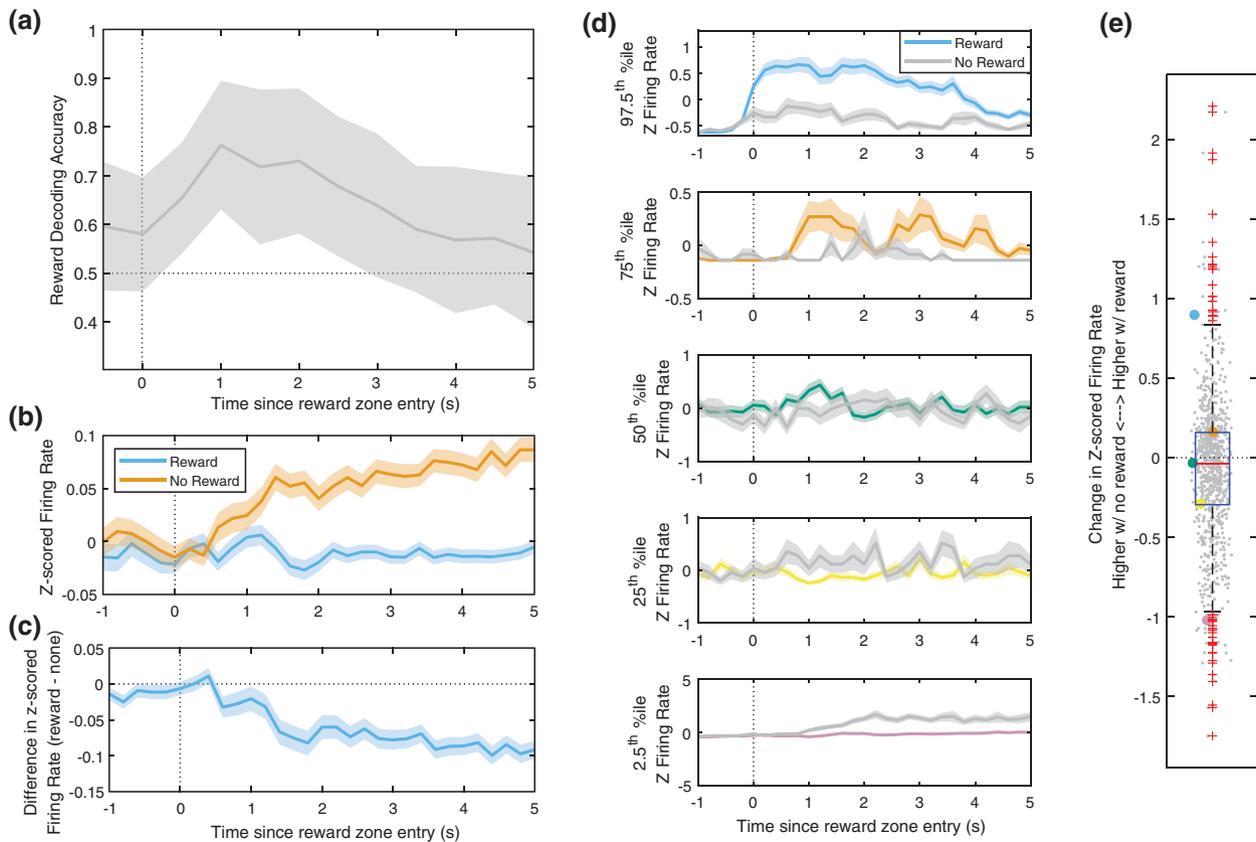


FIGURE 5 Reward representation in dmPFC. (a) Accuracy of Bayesian decoding of reward (vs. the lack thereof) from dmPFC ensembles. Line shows the mean and shaded area shows the SD across $N = 40$ sessions. Vertical dotted line shows the time of reward zone entry. (b) z-scored firing rates of individual units following reward zone entry, split by whether reward was received on that lap or not. Vertical dotted line shows the time of reward zone entry. (c) Difference in z-scored firing rates between reward and lack thereof. Panels (b) and (c) show the mean \pm standard error across $N = 824$ units. (d) z-scored firing rates upon reward zone entry for individual example cells. Color of the line in each panel corresponds to the colored dots in (e). (e) Difference in z-scored firing rates between reward and lack thereof for all units. Greater values indicate greater firing rates following reward than following a lack of reward, while lesser values indicate a lesser firing rate following reward than following a lack of reward. Note that the median is only slightly less than 0 (the averaging effect seen in panels (b) and (c)), but that this difference is small compared to the spread of the entire distribution [Color figure can be viewed at wileyonlinelibrary.com]

an audio cue as to whether a choice was correct or incorrect upon reward zone entry, and so rats had multiple sources of information as to whether their choice was correct or incorrect (both the audio cue and the presence or absence of food reward at the feeder site). The presence or absence of reward was most accurately decoded from dmPFC activity between around 1 and 3 s after reward zone entry (Figure 5a). This indicates that dmPFC does in fact carry information about reward, at least at the time of reward delivery.

One potential confound, however, is that rats' running velocity may have differed between reward receipt and the lack thereof: rats stopped to consume food when they received reward, but often continued on to start a new lap when they did not. In this case, dmPFC may have been simply representing running velocity, and not reward per se, but decoding of reward from ensemble activity could still be possible. Therefore, we performed reward decoding from ensemble activity in dmPFC, but using only times when rats were running (which we defined as times when rats were moving more quickly than 10 cm/s). We found that we were still able to decode reward receipt

from the lack thereof with similar accuracy during the same time period following reward delivery (Figure S1). This suggests that dmPFC represents information about reward independently of velocity.

How was this reward information represented? Firing rates were, on average, significantly higher after an incorrect choice than after reward was delivered (Figure 5b,c). However, to say that firing rates decreased following reward receipt would be a mischaracterization of the data, as the effect was mostly due to averaging. While cells showed a decreased firing rate upon reward on average, there was a wide distribution of responses to reward across individual units. Some cells' firing rates increased in response to reward while other cells had greater firing rates after incorrect choices (Figure 5e, and see Figure 5d for examples of cells' responses across the distribution). This indicates that the reward encoding we observed in dmPFC was not simply due to "reward cells" or cells which decreased their firing rate in the same direction upon reward receipt, but rather that the representation of reward was due to encoding across the entire ensemble.

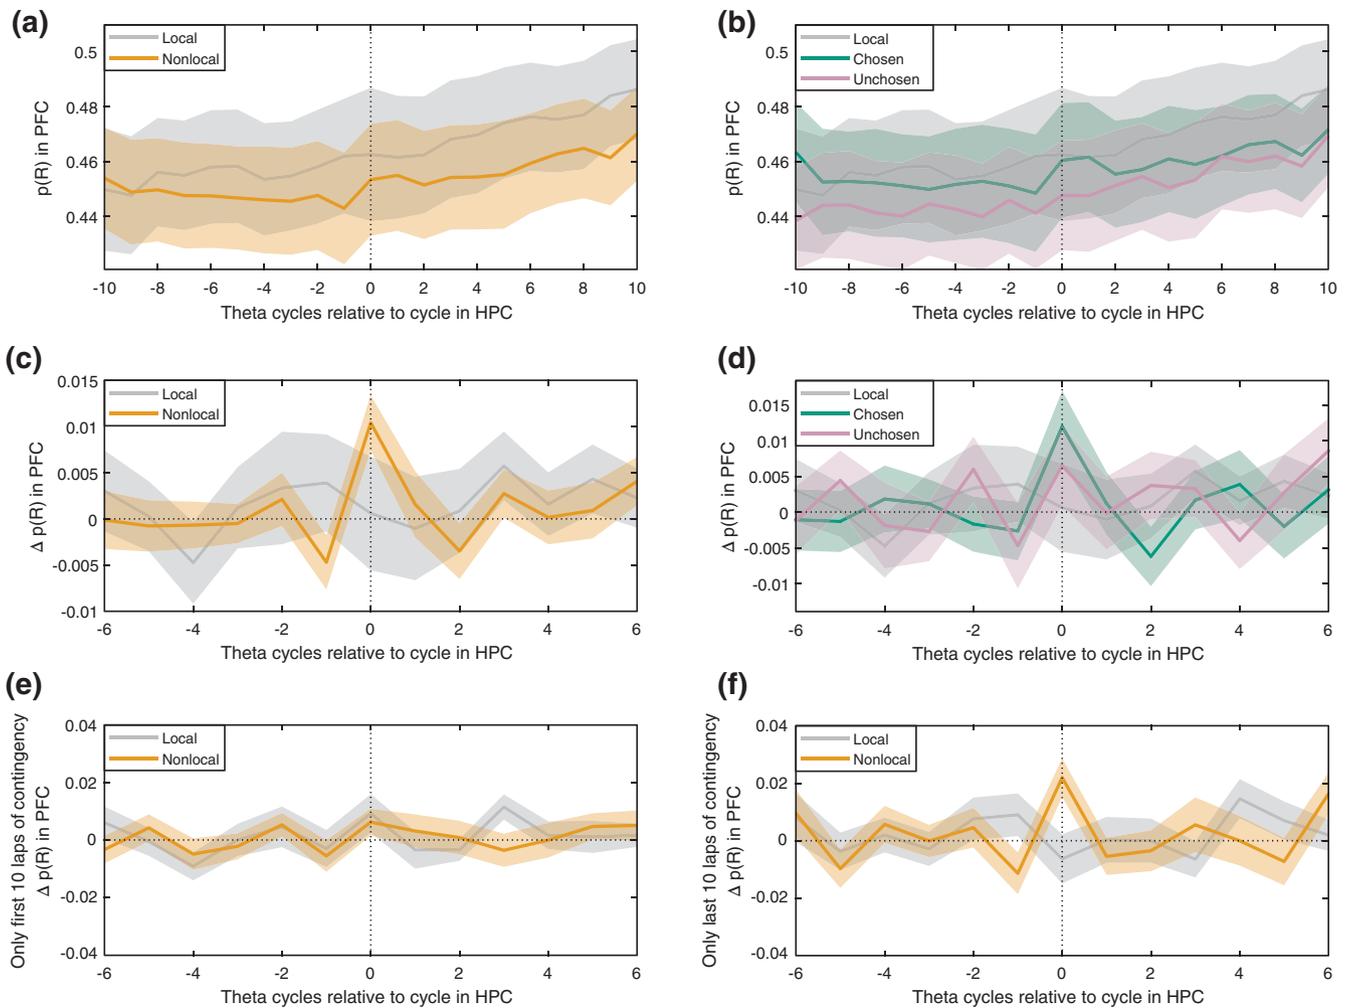


FIGURE 6 Reward encoding in dmPFC and its dependence on goal encoding in CA1. (a) Decoded reward probability from dmPFC as a function of hippocampal theta cycles since the reference theta cycle, split by whether hippocampal ensembles represented the goal location or the current location during the reference theta cycle. (b) Same as in panel (a), but theta cycles where CA1 represented non-local information have been further split by whether CA1 represented the goal location the rat ultimately chose on that lap (“Chosen”) or the opposite goal location (“Unchosen”). (c) Change in the decoded reward probability as a function of lag (the derivative of the reward signal in panel (a)). (d) Change in the decoded reward probability as a function of lag (the derivative of the reward signal shown in panel (b)). Shown in all panels is the mean \pm standard error, $N = 40$ sessions. $p(R)$ indicates the posterior probability of reward, as decoded from ensemble firing rates in dmPFC. (e) Same as in (c) but showing only data from the first 10 laps of contingency blocks. (f) Same as in (c) but showing only data from the last 10 laps of contingency blocks [Color figure can be viewed at wileyonlinelibrary.com]

This is consistent with previous work which finds that single units in prefrontal areas tend to have highly mixed selectivity (Asaad, Rainer, & Miller, 1998; Mansouri, Matsumoto, & Tanaka, 2006; Mante, Sussillo, Shenoy, & Newsome, 2013; Rigotti et al., 2013).

Did prefrontal representations of reward change in response to the encoding of non-local information in hippocampus? We performed cross-validated Bayesian decoding of reward from ensemble activity in dmPFC and also cross-validated Bayesian decoding of spatial location from simultaneously recorded ensemble activity in CA1, while rats deliberated at the choice point of the contingency switch task. The CA1 decoder was trained on location and neural activity across the entire session, with 100-fold cross validation, using time bins corresponding to hippocampal theta cycles. The decoder used for

dmPFC was trained to decode reward versus the lack thereof, using only epochs 1–3 s after reward zone entry, which was the time during which we found dmPFC to be most reliably encoding reward (Figure 5a). This decoder was then used to predict the probability of reward from ensemble activity in dmPFC while rats passed through the choice point. The time bins used for dmPFC decoding were the same as those used for hippocampal decoding (theta cycles of the hippocampal LFP). We split hippocampal theta cycles into those where hippocampus was representing “local” spatial information and those where hippocampus was representing non-local, reward zone information. “Local” theta cycles were those where the decoded spatial posterior distribution had a greater posterior probability in the choice point than anywhere else on the maze. In contrast, the “non-local” theta cycles

were those during which the reward zone had a greater posterior probability than the rest of the maze combined (so, they were not just non-local in a broad sense, but specifically those theta cycles during which hippocampus represented the reward zone most strongly).

We found that during theta cycles during which CA1 was representing the reward zone, there was not a significantly different amount of reward encoding in dmPFC than during theta cycles when CA1 was representing the choice point (Figure 6a). However, there was a significant change in the reward representation in dmPFC. Reward encoding in dmPFC significantly increased during theta cycles where CA1 represented the reward zone (Figure 6c, orange line at zero lag; $p = .0021$, two-sided Wilcoxon signed rank test comparing the change in decoded reward probabilities from the previous theta cycle, $N = 40$ sessions). The rate of this change was approximately 1.02% per theta cycle, where a 100% change would correspond to a transition from complete certainty that reward would not be received to complete certainty that reward would be received within one theta cycle. In contrast, dmPFC reward encoding did not change during theta cycles where CA1 represented only local information (Figure 6c, gray line at zero lag; $p = .85$, two-sided Wilcoxon signed rank test, $N = 40$ sessions). However, this change in dmPFC reward encoding was not significantly greater during theta cycles in which HPC represented the reward zone than during theta cycles in which HPC represented locally (Figure 6c, gray vs orange line at zero lag; $p = .12$, two-sided Wilcoxon rank sum test, $N = 40$ sessions).

To determine the temporal specificity of this effect, we performed a lag analysis similar to that used in Figure 4c,d. We examined the change in reward representation in dmPFC as a function of the number of theta cycles by which the reference theta cycle in CA1 led or lagged the decoding time bin used for reward decoding from dmPFC. Unlike the results in the previous section (Figure 4c,d), we found that the change in reward representation in dmPFC was very tightly locked to the theta cycle during which CA1 represented the reward zone (Figure 6c). The only offset during which the reward encoding in dmPFC significantly increased was when the same time bin was used for reward decoding from dmPFC and location decoding from CA1, during which CA1 represented the reward zone (Figure 6c, orange line; two-sided Wilcoxon signed rank tests, $p = .0021$ for a theta cycle lag of zero; for all other lags $p > .0039$, the Bonferroni-corrected significance level for $\alpha = .05$). Theta cycles during which CA1 represented only local information showed no systematic change in reward encoding in dmPFC, at any lag (Figure 6c, gray line; two-sided Wilcoxon signed rank tests, for all lags including zero theta cycle lag $p > .0039$, the Bonferroni-corrected significance level for $\alpha = .05$). Although on average the reward encoding in dmPFC increased by a greater amount on theta cycles during which CA1 represented specifically the reward zone on the side which the rat ended up choosing on that lap, this was not significantly different between theta cycles during which CA1 represented the chosen versus the unchosen side (Figure 6b,d, two-sided Wilcoxon signed rank test, $p = .158$).

How does this change in reward representation in dmPFC correspond to the changing contingencies? If non-local representation in HPC is driving reward expectancy changes in dmPFC, then we may

expect to see these changes occurring most strongly when the task contingencies are stable and any hippocampal prospective information reaching dmPFC is most confident about its predictions. While our previous results did not seem to be related to whether the rat was early in a contingency block or late in a contingency block (Figures S2–S6), the increase in reward encoding in dmPFC concurrent with nonlocal representation in CA1 was especially prominent late within contingency blocks (Figure 6e,f). In the first 10 laps of contingency blocks, there was no significant increase in reward encoding in dmPFC on theta cycles when HPC representations were the most nonlocal (Figure 6e at a lag of 0 theta cycles, an average reward encoding increase of 0.62% per lap, $p = .202$, two-sided Wilcoxon signed rank test), and the change in dmPFC reward encoding was not significantly greater during theta cycles when HPC representations were the most nonlocal than during theta cycles when HPC representations were the most local (Figure 6e at a lag of 0 theta cycles, gray vs. orange lines, $p = .881$, two-sided Wilcoxon rank sum test). In contrast, during the last 10 laps of contingency blocks, there was a significant increase in reward encoding in dmPFC on theta cycles when HPC representations were the most nonlocal (Figure 6f at a lag of 0 theta cycles, an average reward encoding increase of 2.2% per lap, $p = .00285$, two-sided Wilcoxon signed rank test), and this change in reward encoding by dmPFC was significantly greater during theta cycles when HPC represented the reward zone most strongly than during theta cycles when HPC representations were local (Figure 6f at a lag of 0 theta cycles, gray vs. orange lines, $p = .030$, two-sided Wilcoxon rank sum test). Taken together, these results suggest that hippocampal spatial representations may have a fast, temporally precise effect on reward encoding in dmPFC, specifically at times when the environment is stable.

4 | DISCUSSION

We examined the relationship between hippocampal and prefrontal representations during deliberation, and found that they were related in complex ways. While CA1 and dmPFC spatial representations were correlated in that they represented either local information or prospective information together, the two areas did not appear to always represent identical prospective information at the same time. Furthermore, activity in dmPFC predicted whether CA1 was representing local or prospective information, but this relationship occurred across a very broad timescale, on the order of seconds. On the other hand, representations of goal location in CA1 appeared to have very temporally specific effects on reward encoding in dmPFC, across timescales on the order of a single theta cycle.

Furthermore, we found that hippocampal goal encoding corresponded to an increase in reward encoding in dmPFC late in contingency blocks (Figure 6f), when animals were generally more confident as to the contingency identity and deliberated less, but not early in contingency blocks where deliberation was more common (Figure 6e). This is consistent with previous observations that hippocampal representations predict goal locations when animals show no signs of

deliberation, but not when animals display markers of deliberation such as vicarious trial and error (Johnson & Redish, 2007; Kay et al., 2020; Papale et al., 2016). Taken together with these findings, our result suggests that hippocampal representations of goal locations may have a more robust effect on reward representations in dmPFC when hippocampal predictions of future outcomes are more confident.

If the prefrontal cortex uses estimates of reward or value associated with candidate actions to perform action selection, then presumably it or other areas must keep track of these candidate actions and the corresponding value estimates. However, there are different mechanisms the brain may use to perform this computation, and different brain areas may play different roles in this process. One possibility is that the prefrontal cortex uses a similar mechanism for value-based decision making as sensory areas do for sensory-based decision making, by integrating evidence until some decision threshold is reached. These so-called “drift diffusion models” or “sequential sampling models” have been found to explain both behavior and neural activity in sensory areas during sensory decision-making tasks (Forstmann, Ratcliff, & Wagenmakers, 2016; Hanes & Schall, 1996; Ratcliff, 1978; Ratcliff & McKoon, 2008; Stone, 1960). However, it is unclear whether the brain uses a similar mechanism for making value-based decisions, where the properties of the options being decided between are entirely internal, as opposed to sensory decision-making where those properties are external and directly observable. Some work suggests that, at least behaviorally, reaction times during value-based decision making tasks can be explained using drift-diffusion models (Krajbich & Rangel, 2011).

However, an alternative possibility is that there is no slow integration process, but rather the brain considers options serially and discretely, and makes a decision without evidence accumulation per se. Recent work indicates that, at least in the orbitofrontal cortex, value signals switch back and forth suddenly, suggesting that options are being considered serially during value-based decisions (Rich & Wallis, 2016; Wallis, 2018). Such stepwise representational jumps during decision-making have been observed even during sensory-based decisions (Latimer, Yates, Meister, Huk, & Pillow, 2015; Sadacca et al., 2016; Yang & Shadlen, 2007). Yet a third possibility is that both of these mechanisms occur simultaneously: certain brain areas could consider and evaluate options simultaneously, while others accumulate evidence for each decision and trigger the corresponding action to be taken when some decision threshold is reached. While our results cannot distinguish these three possibilities, they do indicate that valuations arising from the representation of prospective information in HPC change reward encoding in dmPFC.

Our results are consistent with recent work by Zielinski, Shin, and Jadhav (2019), who also find that representations of location in PFC and CA1 are correlated during a spatial decision-making task. Our data further indicates that HPC and dmPFC spatial representations are coherent not only in general, but specifically that they encode prospective information simultaneously. However, our data suggests that dmPFC and CA1 may not be encoding identical prospective spatial information simultaneously during deliberation. Instead, we found that

the goal zone represented by dmPFC and HPC was uncorrelated during prospective encoding by both structures, but that reward encoding in dmPFC changed during hippocampal theta cycles where CA1 was encoding goal locations. This suggests that while dmPFC and HPC represent prospective information together, the relationship between prospective information represented in these areas may be more nuanced than a simple duplication of spatial information, and also involve decision making variables such as reward or value.

ACKNOWLEDGMENTS

We thank Christopher Boldt, Kelsey Seeland, and Ayaka Sheehan for technical support and for building the tetrode drives, Christopher Boldt for building the silicon probe drives, Ayaka Sheehan for performing the histology, as well as Onni Rauhala and Daniel Min for help training rats. This work was funded by NSF IGERT Neuroengineering grant DGE-1069104, NIH R01-MH080318, and NIH R01-MH112688.

CONFLICT OF INTERESTS

The authors declare no potential conflict of interests.

AUTHOR CONTRIBUTIONS

Brendan M. Hasz and A. David Redish conceived and designed the experiment and analyses, and wrote the manuscript. Brendan M. Hasz collected and analyzed the data.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

A. David Redish  <https://orcid.org/0000-0003-3644-9072>

REFERENCES

- Amemiya, S., & Redish, A. D. (2016). Manipulating decisiveness in decision making: Effects of clonidine on hippocampal search strategies. *Journal of Neuroscience*, *36*(3), 814–827.
- Amemiya, S., & Redish, A. D. (2018). Hippocampal theta-gamma coupling reflects state-dependent information processing in decision making. *Cell Reports*, *22*(12), 3328–3338.
- Asaad, W. F., Rainer, G., & Miller, E. K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron*, *21*(6), 1399–1407.
- Barker, G. R. I., Banks, P. J., Scott, H., Ralph, G. S., Mitrophanous, K. A., Wong, L., ... Warburton, E. C. (2017). Separate elements of episodic memory subserved by distinct hippocampal–prefrontal connections. *Nature Neuroscience*, *20*(2), 242–250.
- Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P. L., Gioanni, Y., Battaglia, F. P., & Wiener, S. I. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal–prefrontal network upon learning. *Neuron*, *66*(6), 921–936.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*(3), 624–652.
- Cassel, J., De Vasconcelos, A. P., Loureiro, M., Cholvin, T., Dalrymple-Alford, J. C., & Vertes, R. P. (2013). The reuniens and rhomboid nuclei: Neuroanatomy, electrophysiological characteristics and behavioral implications. *Progress in Neurobiology*, *111*, 34–52.

- Cowen, S. L., & McNaughton, B. L. (2007). Selective delay activity in the medial prefrontal cortex of the rat: Contribution of sensorimotor information and contingency. *Journal of Neurophysiology*, 98(1), 303–316.
- Dalley, J. W., Cardinal, R. N., & Robbins, T. W. (2004). Prefrontal executive and cognitive functions in rodents: Neural and neurochemical substrates. *Neuroscience & Biobehavioral Reviews*, 28(7), 771–784.
- Delatour, B., & Gisquet-Verrier, P. (1999). Lesions of the prelimbic–infralimbic cortices in rats do not disrupt response selection processes but induce delay-dependent deficits: Evidence for a role in working memory? *Behavioral Neuroscience*, 113(5), 941–955.
- Delatour, B., & Witter, M. P. (2002). Projections from the parahippocampal region to the prefrontal cortex in the rat: Evidence of multiple pathways. *European Journal of Neuroscience*, 15(8), 1400–1407.
- Di Prisco, G. V., & Vertes, R. P. (2006). Excitatory actions of the ventral midline thalamus (rhomboid/reuniens) on the medial prefrontal cortex in the rat. *Synapse*, 60(1), 45–55.
- Doll, B. B., Simon, D. A., & Daw, N. D. (2012). The ubiquity of model-based reinforcement learning. *Current Opinion in Neurobiology*, 22(6), 1075–1081.
- Dolleman-Van der Weel, M. J., Da Silva, F. H. L., & Witter, M. P. (2017). Interaction of nucleus reuniens and entorhinal cortex projections in hippocampal field ca1 of the rat. *Brain Structure and Function*, 222(5), 2421–2438.
- Dolleman-van der Weel, M. J., Griffin, A. L., Ito, H. T., Shapiro, M. L., Witter, M. P., Vertes, R. P., & Allen, T. A. (2019). The nucleus reuniens of the thalamus sits at the nexus of a hippocampus and medial prefrontal cortex circuit enabling memory and behavior. *Learning & Memory*, 26(7), 191–205.
- Dwyer, D. M., Dunn, M. J., Rhodes, S. E. V., & Killcross, A. S. (2010). Lesions of the prelimbic prefrontal cortex prevent response conflict produced by action–outcome associations. *Quarterly Journal of Experimental Psychology*, 63(3), 417–424.
- Eichenbaum, H. (2017). Prefrontal–hippocampal interactions in episodic memory. *Nature Reviews Neuroscience*, 18(9), 547–558.
- Euston, D. R., Gruber, A. J., & McNaughton, B. L. (2012). The role of medial prefrontal cortex in memory and decision making. *Neuron*, 76(6), 1057–1070.
- Ferino, F., Thierry, A. M., & Glowinski, J. (1987). Anatomical and electrophysiological evidence for a direct projection from ammon's horn to the medial prefrontal cortex in the rat. *Experimental Brain Research*, 65(2), 421–426.
- Floresco, S. B., & Grace, A. A. (2003). Gating of hippocampal-evoked activity in prefrontal cortical neurons by inputs from the mediodorsal thalamus and ventral tegmental area. *Journal of Neuroscience*, 23(9), 3930–3943.
- Floresco, S. B., Seamans, J. K., & Phillips, A. G. (1997). Selective roles for hippocampal, prefrontal cortical, and ventral striatal circuits in radial-arm maze tasks with or without a delay. *Journal of Neuroscience*, 17(5), 1880–1890.
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, 67, 641–666.
- Fuster, J. (2015). *The prefrontal cortex*. Cambridge, MA: Academic Press.
- Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, 173(3997), 652–654.
- Gilbert, D. T., & Wilson, T. D. (2007). Propection: Experiencing the future. *Science*, 317(5843), 1351–1354.
- Guise, K. G., & Shapiro, M. L. (2017). Medial prefrontal cortex reduces memory interference by modifying hippocampal encoding. *Neuron*, 94(1), 183–192.
- Haddon, J. E., & Killcross, A. S. (2005). Medial prefrontal cortex lesions abolish contextual control of competing responses. *Journal of the Experimental Analysis of Behavior*, 84(3), 485–504.
- Haddon, J. E., & Killcross, S. (2006). Prefrontal cortex lesions disrupt the contextual control of response conflict. *Journal of Neuroscience*, 26(11), 2933–2940.
- Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, 274(5286), 427–430.
- Hassabis, D., & Maguire, E. A. (2009). The construction system of the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1263–1271.
- Hasz, B. M., & Redish, A. D. (2020). Dorsomedial prefrontal cortex and hippocampus represent strategic context even while simultaneously changing representation throughout a task session. *Neurobiology of Learning and Memory*, 171, 107215.
- Hok, V., Chah, E., Save, E., & Poucet, B. (2013). Prefrontal cortex focally modulates hippocampal place cell firing patterns. *Journal of Neuroscience*, 33(8), 3443–3451.
- Hoover, W. B., & Vertes, R. P. (2007). Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Structure and Function*, 212(2), 149–179.
- Hoover, W. B., & Vertes, R. P. (2012). Collateral projections from nucleus reuniens of thalamus to hippocampus and medial prefrontal cortex in the rat: A single and double retrograde fluorescent labeling study. *Brain Structure and Function*, 217(2), 191–209.
- Horst, N. K., & Laubach, M. (2009). The role of rat dorsomedial prefrontal cortex in spatial working memory. *Neuroscience*, 164(2), 444–456.
- Hyman, J. M., Zilli, E. A., Paley, A. M., & Hasselmo, M. E. (2005). Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus*, 15(6), 739–749.
- Jai, Y. Y., & Frank, L. M. (2015). Hippocampal–cortical interaction in decision making. *Neurobiology of Learning and Memory*, 117, 34–41.
- Jay, T. M., & Witter, M. P. (1991). Distribution of hippocampal ca1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology*, 313(4), 574–586.
- Johnson, A., & Redish, A. D. (2007). Neural ensembles in ca3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45), 12176–12189.
- Jones, M. W., & Wilson, M. A. (2005). Theta rhythms coordinate hippocampal–prefrontal interactions in a spatial memory task. *PLoS Biology*, 3(12), e402.
- Kay, K., Chung, J. E., Sosa, M., Schor, J. S., Karlsson, M. P., Larkin, M. C., ... Frank, L. M. (2020). Constant sub-second cycling between representations of possible futures in the hippocampus. *Cell*, 180(3), 552–567.
- Kesner, R. P., & Churchwell, J. C. (2011). An analysis of rat prefrontal cortex in mediating executive function. *Neurobiology of Learning and Memory*, 96(3), 417–431.
- Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33), 13852–13857.
- Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244), 184–187.
- Maharjan, D. M., Dai, Y. Y., Glantz, E. H., & Jadhav, S. P. (2018). Disruption of dorsal hippocampal–prefrontal interactions using chemogenetic inactivation impairs spatial learning. *Neurobiology of Learning and Memory*, 155, 351–360.
- Mansouri, F. A., Matsumoto, K., & Tanaka, K. (2006). Prefrontal cell activities related to monkeys' success and failure in adapting to rule changes in a Wisconsin card sorting test analog. *Journal of Neuroscience*, 26(10), 2745–2756.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474), 78–84.

- Marquis, J., Killcross, S., & Haddon, J. E. (2007). Inactivation of the prelimbic, but not infralimbic, prefrontal cortex impairs the contextual control of response conflict in rats. *European Journal of Neuroscience*, 25(2), 559–566.
- McKenna, J. T., & Vertes, R. P. (2004). Afferent projections to nucleus reuniens of the thalamus. *Journal of Comparative Neurology*, 480(2), 115–142.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1), 167–202.
- Navawongse, R., & Eichenbaum, H. (2013). Distinct pathways for rule-based retrieval and spatial mapping of memory representations in hippocampal neurons. *Journal of Neuroscience*, 33(3), 1002–1013.
- O'Neill, P., Gordon, J. A., & Sigurdsson, T. (2013). Theta oscillations in the medial prefrontal cortex are modulated by spatial working memory and synchronize with the hippocampus through its ventral subregion. *Journal of Neuroscience*, 33(35), 14211–14224.
- Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M., & Harris, K. D. (2016). Fast and accurate spike sorting of high-channel count probes with kilosort. In *Advances in neural information processing systems*, 2016, 4448–4456.
- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: A good-based model. *Annual Review of Neuroscience*, 34, 333–359.
- Papale, A. E., Stott, J. J., Powell, N. J., Regier, P. S., & Redish, A. D. (2012). Interactions between deliberation and delay-discounting in rats. *Cognitive, Affective, & Behavioral Neuroscience*, 12(3), 513–526.
- Papale, A. E., Zielinski, M. C., Frank, L. M., Jadhav, S. P., & Redish, A. D. (2016). Interplay between hippocampal sharp-wave-ripple events and vicarious trial and error behaviors in decision making. *Neuron*, 92(5), 975–982.
- Paxinos, G., & Watson, C. (2006). *The rat brain in stereotaxic coordinates: Hard cover edition*, Cambridge, MA: Elsevier.
- Powell, N. J., & Redish, A. D. (2014). Complex neural codes in rat prelimbic cortex are stable across days on a spatial decision task. *Frontiers in Behavioral Neuroscience*, 8(120), 1–16.
- Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, 23(17), R764–R773.
- Ragozzino, M. E., & Kesner, R. P. (1998). The effects of muscarinic cholinergic receptor blockade in the rat anterior cingulate and prelimbic/infralimbic cortices on spatial working memory. *Neurobiology of Learning and Memory*, 69(3), 241–257.
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7), 545–556.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922.
- Redish, A. D. (2013). *The mind within the brain: How we make decisions and how those decisions go wrong*, Oxford, UK: Oxford University Press.
- Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, 17(3), 147–159.
- Redish, A. D. (2017). Mclust. Version 4.4.07. <http://redishlab.neuroscience.umn.edu/MClust/MClust.html>.
- Rich, E. L., & Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex. *Nature Neuroscience*, 19(7), 973–980.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451), 585–590.
- Rossant, C., Kadir, S., Goodman, D., Hunter, M., & Harris, K. (2016). Phy. Version 1.0.9.
- Sadacca, B. F., Mukherjee, N., Vladusich, T., Li, J. X., Katz, D. B., & Miller, P. (2016). The behavioral relevance of cortical neural ensemble responses emerges suddenly. *Journal of Neuroscience*, 36(3), 655–669.
- Schmidt, B., Duin, A. A., & Redish, A. D. (2019). Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision-making. *Journal of Neurophysiology*, 121, 1981–2000.
- Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K., & Redish, A. D. (2005). Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience*, 131(1), 1–11.
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, 79(2), 217–240.
- Shin, J. D., & Jadhav, S. P. (2016). Multiple modes of hippocampal-prefrontal interactions in memory-guided behavior. *Current Opinion in Neurobiology*, 40, 161–169.
- Siapas, A. G., Lubenov, E. V., & Wilson, M. A. (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron*, 46(1), 141–151.
- Steiner, A. P., & Redish, A. D. (2012). The road not taken: Neural correlates of decision making in orbitofrontal cortex. *Frontiers in Neuroscience*, 6(131), 1–21.
- Steiner, A. P., & Redish, A. D. (2014). Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nature Neuroscience*, 17(7), 995–1002.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3), 251–260.
- Stott, J. J., & Redish, A. D. (2014). A functional difference in information processing between orbitofrontal cortex and ventral striatum during decision-making behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655), 20130472.
- Swanson, L. W. (1981). A direct projection from ammon's horn to prefrontal cortex in the rat. *Brain Research*, 217(1), 150–154.
- Tronel, S., & Sara, S. J. (2003). Blockade of nmda receptors in prelimbic cortex induces an enduring amnesia for odor-reward associative learning. *Journal of Neuroscience*, 23(13), 5472–5476.
- Urban, K. R., Layfield, D. M., & Griffin, A. L. (2014). Transient inactivation of the medial prefrontal cortex impairs performance on a working memory-dependent conditional discrimination task. *Behavioral Neuroscience*, 128(6), 639–643.
- van der Meer, M., Kurth-Nelson, Z., & Redish, A. D. (2012). Information processing in decisionmaking systems. *The Neuroscientist*, 18(4), 342–359.
- van der Meer, M. A. A., & Redish, A. D. (2010). Expectancies in decision making, reinforcement learning, and ventral striatum. *Frontiers in Neuroscience*, 3(6), 29–37.
- Vertes, R. P. (2002). Analysis of projections from the medial prefrontal cortex to the thalamus in the rat, with emphasis on nucleus reuniens. *Journal of Comparative Neurology*, 442(2), 163–187.
- Vertes, R. P. (2004). Differential projections of the infralimbic and prelimbic cortex in the rat. *Synapse*, 51(1), 32–58.
- Vertes, R. P., Hoover, W. B., Do Valle, A. C., Sherman, A., & Rodriguez, J. J. (2006). Efferent projections of reuniens and rhomboid nuclei of the thalamus in the rat. *Journal of Comparative Neurology*, 499(5), 768–796.
- Vertes, R. P., Hoover, W. B., Szigeti-Buck, K., & Leranath, C. (2007). Nucleus reuniens of the midline thalamus: Link between the medial prefrontal cortex and the hippocampus. *Brain Research Bulletin*, 71(6), 601–609.
- Verwer, R. W. H., Meijer, R. J., Van Uum, H. F. M., & Witter, M. P. (1997). Collateral projections from the rat hippocampal formation to the lateral and medial prefrontal cortex. *Hippocampus*, 7(4), 397–402.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., ... van der Walt, S. J. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272.

- Wallis, J. D. (2018). Decoding cognitive processes from neural ensembles. *Trends in Cognitive Sciences*, 22(12), 1091–1102.
- Wang, G., & Cai, J. (2006). Disconnection of the hippocampal–prefrontal cortical circuits impairs spatial working memory performance in rats. *Behavioural Brain Research*, 175(2), 329–336.
- Wang, J. X., Cohen, N. J., & Voss, J. L. (2015). Covert rapid action-memory simulation (crams): A hypothesis of hippocampal–prefrontal interactions for adaptive behavior. *Neurobiology of Learning and Memory*, 117, 22–33.
- Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, 447(7148), 1075–1080.
- Yoon, T., Okada, J., Jung, M. W., & Kim, J. J. (2008). Prefrontal cortex and hippocampus subserve different components of working memory in rats. *Learning & Memory*, 15(3), 97–105.
- Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, 79(2), 1017–1044.
- Zielinski, M. C., Shin, J. D., & Jadhav, S. P. (2019). Coherent coding of spatial position mediated by theta oscillations in the hippocampus and prefrontal cortex. *Journal of Neuroscience*, 39(23), 4550–4565.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Hasz BM, Redish AD. Spatial encoding in dorsomedial prefrontal cortex and hippocampus is related during deliberation. *Hippocampus*. 2020;1–15. <https://doi.org/10.1002/hipo.23250>