Report

Neuron

Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making

Highlights

- Choice point deliberation (VTE) was inversely related to sharp-wave-ripple rates
- VTE predicted a subsequent decrease in the rate of SWRs at reward sites
- SWR rate at reward sites predicted a subsequent decrease in VTE at a choice point
- Disrupting SWRs increased expression of VTE at the choice point

Authors

Andrew E. Papale, Mark C. Zielinski, Loren M. Frank, Shantanu P. Jadhav, A. David Redish

Correspondence

redish@umn.edu

In Brief

Papale et al. examine two hippocampal phenomena associated with executive processing, sharp-wave-ripple complexes occurring at reward sites and vicarious-trial-and-error behaviors occurring at decision points, and find inverse relationships at multiple timescales.

Papale et al., 2016, Neuron *92*, 975–982
December 7, 2016
http://dx.doi.org/10.1016/j.neuron.2016.10.028

CellPress

Neuron
Report

Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making

Andrew E. Papale,[1,6] Mark C. Zielinski,[2] Loren M. Frank,[3] Shantanu P. Jadhav,[4] and A. David Redish[5,7,*]
[1]Graduate Program in Neuroscience, University of Minnesota, Minneapolis, MN 55455, USA
[2]Graduate Program in Neuroscience, Brandeis University, Waltham, MA 02453, USA
[3]HHMI, Kavli Institute for Fundamental Neuroscience, Department of Physiology and Center for Integrative Neuroscience, UCSF, San Francisco, CA 94158, USA
[4]Neuroscience Program, Department of Psychology and Volen Center for Complex Systems, Brandeis University, Waltham, MA 02453, USA
[5]Department of Neuroscience, University of Minnesota, Minneapolis, MN 55455, USA
[6]Present Address: Department of Neurobiology, University of Pittsburgh, Pittsburgh, PA 15213, USA
[7]Lead Contact
*Correspondence: redish@umn.edu
http://dx.doi.org/10.1016/j.neuron.2016.10.028

SUMMARY

Current theories posit that memories encoded during experiences are subsequently consolidated into longer-term storage. Hippocampal sharp-wave-ripple (SWR) events have been linked to this consolidation process during sleep, but SWRs also occur during awake immobility, where their role remains unclear. We report that awake SWR rates at the reward site are inversely related to the prevalence of vicarious trial and error (VTE) behaviors, thought to be involved in deliberation processes. SWR rates were diminished immediately after VTE behaviors and an increase in the rate of SWR events at the reward site predicted a decrease in subsequent VTE behaviors at the choice point. Furthermore, SWR disruptions increased VTE behaviors. These results suggest an inverse relationship between SWRs and VTE behaviors and suggest that awake SWRs and associated planning and memory consolidation mechanisms are engaged specifically in the context of higher levels of behavioral certainty.

INTRODUCTION

The hippocampus has been implicated in many aspects of cognition including navigation (O'Keefe and Nadel, 1978; Redish, 1999), imagination (Buckner and Carroll, 2007; Hassabis and Maguire, 2011), and the consolidation of memories (Buzsáki, 1989; Carr et al., 2011). The working hypothesis of the hippocampal field is that hippocampal sequences underlie these cognitive processes (Skaggs et al., 1996; Foster and Wilson, 2007; Wikenheiser and Redish, 2015a). These sequences occur during two largely distinct hippocampal states (Vanderwolf, 1971; O'Keefe and Nadel, 1978): *theta* (a 6–10 Hz continuous oscillation, occurring most prominently during movement and attentive states) and *sharp-wave-ripple* complexes (SWR, a 100–200 ms transient 200 Hz oscillation, occurring during slow-wave sleep and awake stillness). When rats pause at decision points, theta sequences alternate between options (Johnson and Redish, 2007). This behavior (*vicarious trial and error* [VTE]) occurs during high uncertainty (Gardner et al., 2013; Schmidt et al., 2013; Redish, 2016), which would necessitate the exploration of future options. VTE has thus been suggested to provide a measure of uncertainty, and theta sequences during VTE may provide a neural substrate for the internal exploration of those possible future options (Johnson and Redish, 2007; Amemiya and Redish, 2016; Redish, 2016).

In contrast, SWR sequences occur both during sleep and waking. SWRs during sleep have been linked to consolidation processes (Ego-Stengel and Wilson, 2010; Girardeau and Zugaro, 2011; de Lavilléon et al., 2015), wherein repeated reactivation of sequences is thought to engrain representations in distributed hippocampal-cortical networks (Alvarez and Squire, 1994; Sutherland and McNaughton, 2000; Nadel et al., 2012; Silva et al., 2015). SWRs during waking are important for learning (Jadhav et al., 2012) and occur most often during periods of rest at rewarded locations (Karlsson and Frank, 2009; Redish, 1999). The specific role that awake SWRs play in memory processes remains unclear. There is evidence that, like theta sequences, SWR sequences contribute to internal exploration of possible future options (Pfeiffer and Foster, 2013; Singer et al., 2013), but SWR sequences also encode paths unrelated to immediate options (Davidson et al., 2009; Gupta et al., 2010) or even to the immediate task at hand (Jackson et al., 2006; Karlsson and Frank, 2009), and the prevalence of SWRs at reward locations suggests a potential role in consolidation processes that could link a previous experience to its outcome (Foster and Wilson, 2006; Singer and Frank, 2009; Buzsáki, 2015).

We report here an inverse relationship between SWRs at reward sites and VTE at a choice point. VTE became less frequent and SWRs more frequent as animals learned to exploit a rule. VTE at a choice point was associated with a decreased

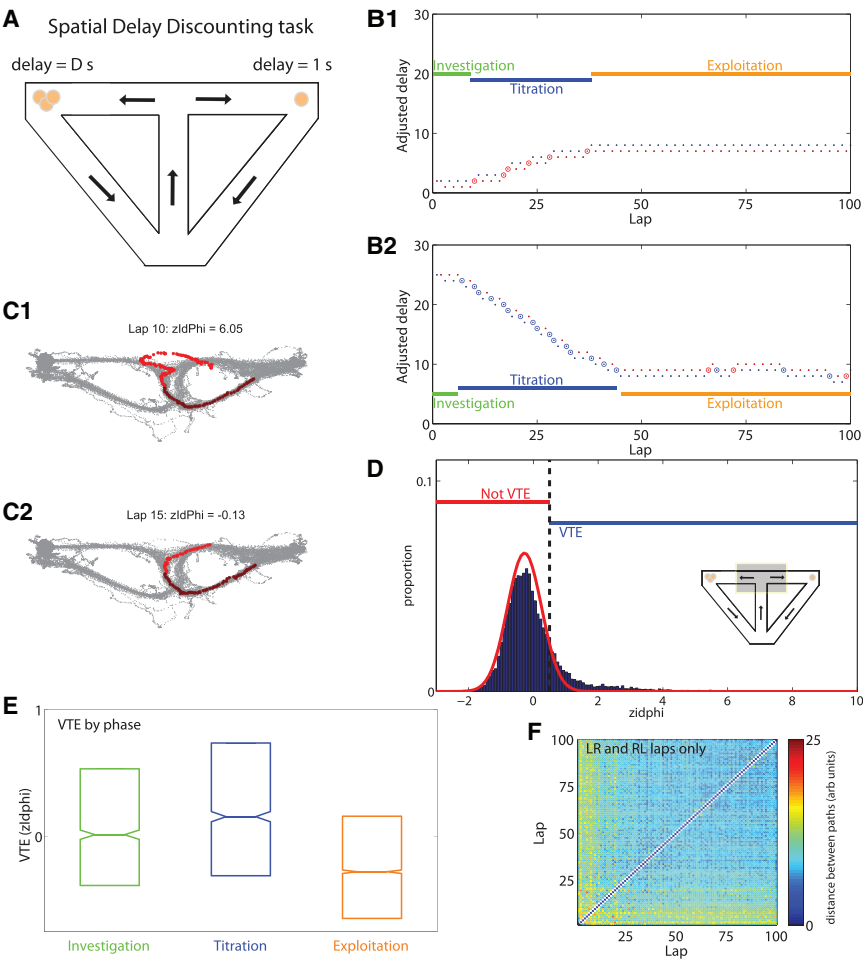

Figure 1. Behavior on the Spatial Delay Discounting Task

(A) Task layout. One side provides a larger reward (3×) after a delay D; the other side provides a smaller reward (1×) after 1 s. The delay is adjusted as a function of the animal's decisions.

(B1 and B2) Delay on the adjusted delay side by lap. Red indicates laps to the delayed side, increasing D by 1 s; blue indicates laps to the non-delayed side, decreasing D by 1 s. Small dots show LR and RL laps; circles show LL or RR laps. Behavior reveals three phases: investigation, titration, and exploitation. (B1) Upward titration; (B2) downward titration.

(C1 and C2) Gray dots show all sampled positions in a given session, with a single lap in red. VTE can be measured quantitatively with zldphi (see Supplemental Experimental Procedures). (C1) VTE pass (zldphi = 6.05); (C2) non-VTE pass (zldphi = −0.13).

(D) zldphi distributed in a skewed manner but can be separated into VTE events and non-VTE events. The threshold between VTE and not was set at zldphi = 0.5, the point where the observed distribution diverged from the expected normal distribution (see Supplemental Experimental Procedures).

(E) VTE decreased in the exploitation phase. Bars show interquartile range, line shows median, notch shows standard error of the median.

(F) Calculated distance between paths (see Supplemental Experimental Procedures). Paths became more stereotyped with time. Early laps are distant from each other as well as from the later laps. Later laps are more stereotyped, marked by a lack of distance between the paths.

rate of SWR events at the subsequent reward site, and a prevalence of SWR events at a reward site was associated with decreased VTE on the subsequent lap. Furthermore, on a second task, we found that selective interruption of SWRs during learning led to an increased prevalence of VTE. These dynamics imply a complex interaction between VTE, theta sequences, and SWR sequences, suggesting an inverse relationship between SWRs and VTE.

RESULTS

In rats, monkeys, and humans, learning often entails a transition from attentive to more automated processes (O'Keefe and Nadel, 1978; Squire, 1987; Hikosaka et al., 1999; Redish, 2013). One intriguing possibility is that the two phenomena (VTE and SWRs) occur at different times in this transition. To directly examine the relationship between VTE, SWRs, navigational planning, and behavioral flexibility, we examined the interplay between VTE and SWRs on two decision tasks, one that included a within-session transition from flexible to more automated behaviors (Papale et al., 2012), and the other that combined across-session development of environmental familiarity with the learning of a complex decision rule (Karlsson and Frank, 2009; Jadhav et al., 2012).

The first task was the spatial adjusting delay discounting (DD) task (Papale et al., 2012), a neuroeconomic task with a single decision point at a T-intersection (Figure 1A). In this task, rats are faced with a choice between a small reward delivered quickly (1× 45 mg unflavored pellet delivered after 1 s, *smaller-sooner*) or a larger reward delivered after a delay (3× food pellets delivered after Ds, *larger-later*). Larger-later and smaller-sooner sides were counterbalanced between sessions, as was the starting delay (initial range 1–30 s). The delay D was adjusted based on the rat's decisions: choosing larger-later increased D by 1 s, while choosing smaller-sooner decreased D on the larger-later side by 1 s. Thus, alternating between the two sides leaves the delay unchanged.

Behavior on this task typically proceeds through three separable phases in each daily session (Papale et al., 2012): (1) investigation—rats alternated sides to identify the delayed side and the initial starting delay D; (2) titration—rats adjusted the delay by preferentially selecting one side over the other; and (3) exploitation—rats alternated sides, holding the delay at a preferred indifference point (Figure 1B) (Papale et al., 2012; Bett et al., 2015; Breton et al., 2015; Mazur, 1997).

On a proportion of passes through the T choice, rats paused and re-oriented toward each option (vicarious trial and error [VTE]; Figure 1C1) and on a proportion of passes rats ran ballistically through (non-VTE; Figure 1C2). We quantified VTE with a *Z* scored measure of the integrated angular velocity (zIdphi, see Supplemental Experimental Procedures). Consistent with previous experiments, zIdphi was highly skewed (skewness of the histogram of all laps = 2.15, median skewness per session was 2.0 ± 0.07 SEM). VTE laps were defined to be those with zIdphi > 0.5 (see Supplemental Experimental Procedures and Figure 1D).

Behavioral analyses indicated that exploitation was marked by a decrease in VTE relative to the other two phases (ANOVA, overall effect, df = 2, n = 12,058, F = 256, $p < 10^{-51}$, post hoc Tukey test exploitation relative to investigation and titration phases, $p < 0.0001$ each, Cohen's D = 0.2) (Figure 1E). 33% of the laps in the titration phases showed VTE, and 25% of the laps in the investigation phases showed VTE, while only 14% of the laps in the exploitation phase showed VTE. Exploitation was also marked by an increase in stereotyped alternation laps as evidenced by an increase in the consistency of the rats' paths of travel (Figure 1F). Whereas non-VTE laps were almost all alternation laps (93% of non-VTE laps were alternation), VTE laps were evenly divided between alternation and adjustment laps (46% of VTE laps were alternation). These findings suggest that exploitation is a period in which the decision about where to go was made earlier, likely at the exit from the reward site, rather than at the T choice. In contrast, on VTE laps, the decision of where to go seems to be made at the T choice itself.

To determine what information was represented while the rat was at the T (Figure 2A), we divided the maze into three regions (reward sites, choice point, and the rest of the maze; Figure 2B) and decoded the animal's location from the spiking activity (see Supplemental Experimental Procedures). We measured the mean of the posterior of the decoded representation. On average, the posterior decoded probability was primarily local, even when the rat was at the choice point (Figures 2C and 2I), and the small amount of posterior decoded probability that distributed non-locally to the feeders while the rat was at the choice point was primarily toward the choice the rat would subsequently choose (Figure 2D, Wilcoxon rank-sum test, z = 46, n = 95,766, $p < 10^{-100}$, Cohen's D = 0.2).

This analysis revealed a systematic shift in representation during VTE laps (Figures 2H and 2J, ANOVA [n = 143,663], effect of region [df = 2, F = 83040, $p < 10^{-100}$, $\eta^2 = 0.3$], effect of VTE [df = 1, F = 22, $p < 10^{-100}$, $\eta^2 = 10^{-4}$], interaction [df = 2, F = 112, $p < 10^{-100}$, $\eta^2 = 10^{-3}$]). VTE laps were marked by increased decoding of the reward location (Figure 2E, Wilcoxon rank-sum test [z = 9, $p < 10^{-20}$, Cohen's D = 0.1]), and a concomitant decrease in decoding of the animal's actual location at the T (Figure 2F, Wilcoxon rank-sum test [z = −8, $p < 10^{-15}$, Cohen's D = 0.1]). There was no significant change in decoding on the rest of the maze (Figure 2G, Wilcoxon rank-sum, [z = 1.8, p = 0.07]). Consistent with previous experiments (Johnson and Redish, 2007; Amemiya and Redish, 2016), there was decoding to both sides during VTE events, but preferentially to the chosen side on non-VTE laps (Figure 2H, ANOVA [n = 95,775], effect of region [chosen/unchosen] [df = 1, F = 322, $p < 10^{-72}$, $\eta^2 = 10^{-3}$], effect of VTE [df = 1, F = 72, $p < 10^{-17}$, $\eta^2 = 10^{-3}$], interaction [df = 1,F = 142, $p < 10^{-32}$,$\eta^2 = 10^{-3}$]). This suggests that on laps in which the rats ran ballistically through the choice point (not showing VTE), they knew their target goal before arriving at the choice point, and the hippocampal sequences preferentially reflected only the chosen goal (Wikenheiser and Redish, 2015b; Amemiya and Redish, 2016). However, consistent with previous work (Johnson and Redish, 2007; Amemiya and Redish, 2016) on VTE laps, the hippocampal sequences were more equally divided between the two options. These neurophysiological findings suggested that the rat was still deciding where to go at the T on VTE laps but that on non-VTE laps, it already knew where it was going to go before it arrived at the T.

We also examined the decoded representation on the timescale of the 6–10 Hz theta rhythm. Consistent with previous work (Johnson and Redish, 2007; Gupta et al., 2012; Amemiya and Redish, 2016), theta cycles represented each side serially, even during VTE events. During any given theta cycle, the sequences were preferentially toward one side or the other and tended not to distribute simultaneously across both sides. This was true for both VTE and non-VTE laps (Figure 2K).

There are three potential explanations for the low (but significant) increased decoding to reward sites during VTE events. (1) The entire assembly could transiently jump to the reward site, with the low probability due to our inability to decode at fast enough timescales due to ensemble size. (2) A sequence could run from the current location to the goal, with the low probability due to our misalignment of decoding time bins with the timing of the sequence. (3) The representation could stretch to include both the current location and the reward site. We cannot differentiate these possibilities from the data here; however, other work looking directly at sequences (Wikenheiser and Redish, 2015b; Wang et al., 2015; Gupta et al., 2012; Feng et al., 2015) suggest that the second alternative is the most likely.

Hippocampal functional connectivity, information processing, and neural activity patterns differ between θ and the hippocampal state in which SWRs occur (LIA, marked by a more broadspectral local field potential with power in 2–4 Hz [δ]). To determine whether these two events (passes through the choice point, VTE or not, and departure from a reward site) occur during similar or different hippocampal network states, we measured the local field potentials from the hippocampal pyramidal layer and analyzed their spectral components. Over the entire session, there was increased power in both these frequency ranges (Figure S1B). Because SWRs are transient events, they vanish in the averaging process of the PSD, but they can be revealed by measuring auto-coherence plots, which measure correlations across frequencies (Masimore et al., 2004). These plots revealed transient events in the 180–200 Hz range, with cross-spectral power in lower frequencies (Figure S1A). While these SWR transients were present in the 3 s prior to departing the reward zone (Figure S1C), they were not present during passes through the T (Figure S1D), even during VTE events (Figure S1E). Supporting this distinction, δ showed stronger power than θ in the 3 s prior to departure from the reward site (Figure S1F), but θ showed greater power than δ at the choice point (Figure S1G), even during VTE (Figure S1H). The difference can be seen by subtracting

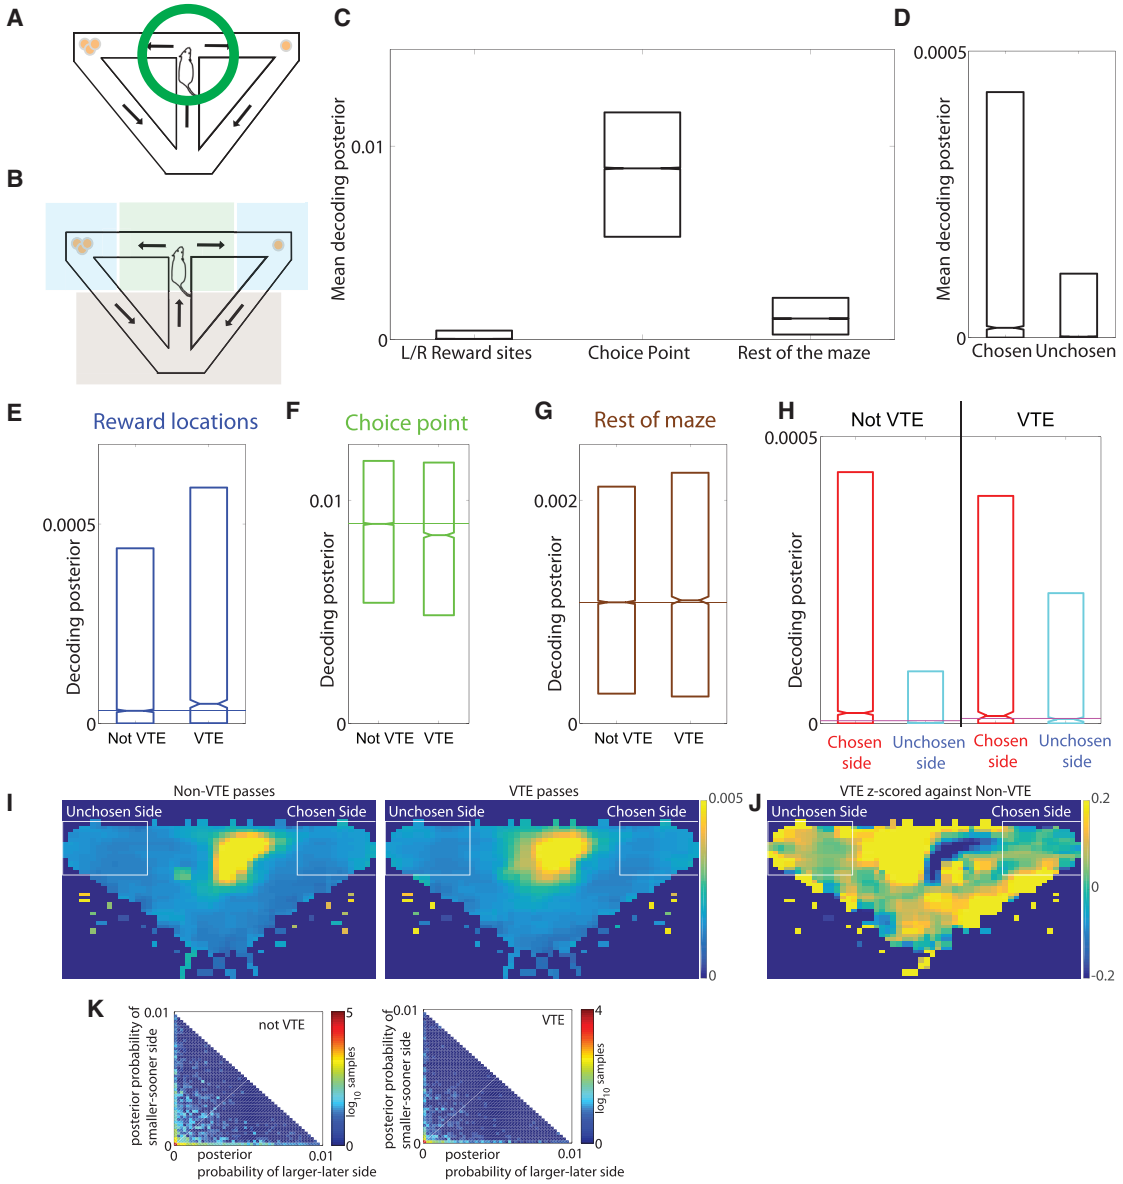

Figure 2. Hippocampal Representations during VTE

(A and B) While rats were at the choice point (A), we measured the mean Bayesian posterior probability over the choice point (green), the reward sites (blue), and the rest of the maze (brown) (B). The data in this figure come from all three phases.

(C) While the rat was at the choice point, most of the posterior remained local.

(D) The small portion of the posterior that distributed to the reward sites was significantly more distributed to the side that would subsequently be chosen.

(E and F) Decoding to the reward sites (E) increased during VTE, matched by a decrease in the decoding locally at the choice point (F).

(G) Decoding on the rest of the maze remained unchanged.

(H) On non-VTE laps, the non-local decoding to the reward sites was preferentially distributed toward the chosen side, whereas it was more evenly distributed on VTE laps.

(I) Average decoding probabilities. White boxes show the regions used to calculate the decoded posterior probability for the reward sites. Laps in which the rat went left have been flipped around the midline for display purposes. On both VTE and non-VTE passes, the majority of the decoded posterior remained at the choice point; however, the small amount of decoded posterior at the reward sites was different under VTE and non-VTE conditions.

(J) The average decoded posteriors of all VTE laps *Z* scored against the mean and standard deviations found in the non-VTE laps.

(K) Theta cycle representations encoded one side or the other but not both. For each decoded sample, we measured the proportion of the posterior assigned to each goal side (larger-later or smaller-sooner). On both VTE and non-VTE laps, when there was increased posterior to one side or the other, there was no increased posterior to the other side. This implies that the theta sequences were alternating between options serially, not simply spreading out ahead of the animal. Boxplots (C–H) show IQR (box), median (line), and standard error of the median (notch).

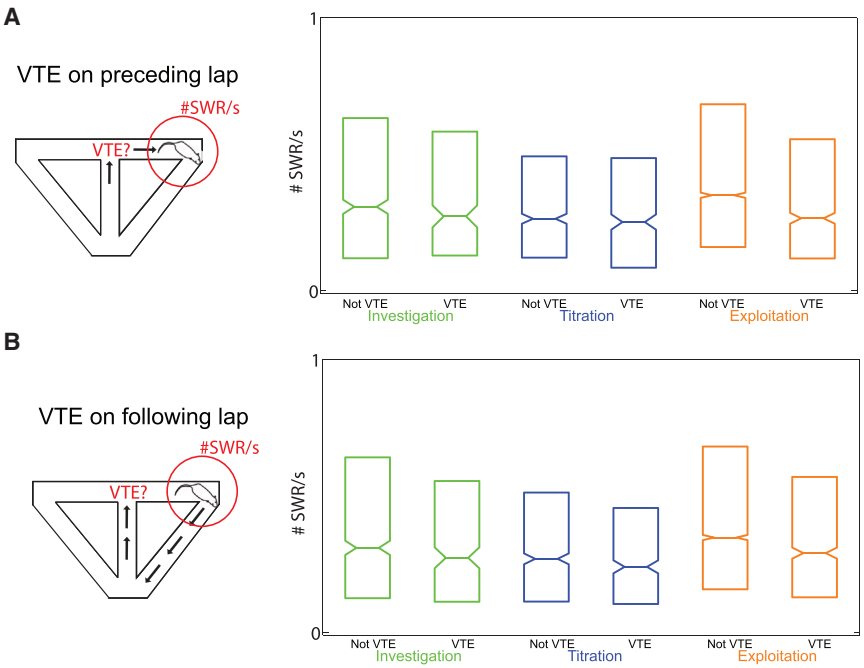

Figure 3. Sharp Waves Were More Common Before and After Non-VTE Laps

(A) We divided reward site experiences based on whether a VTE (zIdphi > 0.5) event occurred as the animal passed the choice point approaching the reward site experience. There was a lower rate of SWR events at the reward site after a VTE lap than after a non-VTE lap.

(B) If we divided reward site experiences based on whether a VTE (zIdphi > 0.5) event occurred on the subsequent lap (following the reward site experience), there was a lower rate of SWR events before a VTE lap than a non-VTE lap. Boxplots show IQR (box), median (line), and standard error of the median (notch).

the auto-coherence plots (Figure S1I) and the PSDs (Figure S1J). These two situations (leaving the feeder site and VTE events) were marked by different hippocampal network states, as can be seen in the different peri-event aligned spectrograms (Figures S1K–S1M).

To examine the interaction between SWRs and VTE events, we measured the rate of SWRs during the pause time at the reward site, after getting food and prior to leaving the reward site. We compared SWR rates across task phase (investigation/titration/exploitation) and with regards to the presence or absence of VTE during the lap approaching the reward site (Figure 3A) or after the reward site visit (Figure 3B). SWR rates increased during exploitation relative to the investigation and titration phases. The SWR rate was significantly lower when VTE occurred on the lap preceding the feeder visit in question (ANOVA [n = 11,260], effect of Rat [df = 5, F = 124, $p < 10^{-100}$, $\eta^2 = 0.04$]; effect of Session [continuous, F = 2911, $p < 10^{-100}$, $\eta^2 = 0.2$]; effect of Phase [df = 2, F = 8.0, p = 0.0003, $\eta^2 = 0.001$]; effect of VTE [df = 1, F = 5.25, p = 0.02, $\eta^2 = 10^{-4}$]). There were minor differences between individual rats (Figure S4) and an increase in SWR rate across sessions (Figure S2A), so we included rat and session in the ANOVA model as separate factors. Significances did not change if we included lap instead of phase. Even with these components, VTE on the previous lap remained a significant explanatory variable in the ANOVA model. Adding in the speed of the lap did not change the significance of preceding VTE as an explanatory variable.

SWR rates in the pause time at the reward site also predicted the prevalence of VTE on the following lap, suggesting bi-directional interactions. There were significantly fewer SWR events preceding VTE events on the subsequent lap (ANOVA [n = 11,268], effect of Rat [df = 5, F = 124, $p < 10^{-100}$, $\eta^2 = 0.04$]; effect of Session [continuous, F = 2,914, $p < 10^{-100}$, $\eta^2 = 0.2$]; effect of Phase [df = 2, F = 7.9, p = 0.0004, $\eta^2 = 0.001$]; effect of VTE [df = 1, F = 5.22, p = 0.02, $\eta^2 = 10^{-4}$]). Subsequent VTE remained a significant explanatory variable in the ANOVA model, even with the inclusion of rat, session, and phase in the model. Using lap instead of phase and adding in the speed of the lap did not change the results. Similar effects could be seen if we measured zIdphi as a function of a median split on SWR rate.

SWR events were more likely to occur when rats took a ballistic path through the T, and higher SWR rates increased the likelihood of taking a ballistic path through the next pass. These results are consistent with previous experiments showing a negative correlation between SWR rate and behavioral variability (Jackson et al., 2006) on familiar environments, although other experiments have found increased SWR rates during novel experiences (O'Neill et al., 2006; Cheng and Frank, 2008).

All animals had extensive experience on the DD task before recording, including at least 1 month of training before implantation of the recording electrodes, as well as at least 2 weeks of training after implantation; nevertheless, we observed large changes in VTE and the rate of SWRs, as well as differences in behavioral regularity across the 30-day experiment. The zIdphi measure of VTE decreased significantly over sessions (Figure S2B), driven primarily by a decrease in VTE during the exploitation phase (ANOVA [n = 11,268], effect of rat [df = 5, F = 12, $p < 10^{-100}$, $\eta^2 = 0.005$], effect of phase [df = 2, F = 18, $p < 10^{-100}$, $\eta^2 = 0.003$], effect of session [df = continuous, F = 118, $p < 10^{-100}$, $\eta^2 = 0.01$], interaction between phase and session [df = 2, F = 3.8, p = 0.02, $\eta^2 = 10^{-3}$]). In parallel, the rate of SWR events increased over sessions (Figure S2A), driven primarily by an increase in SWR rates during the exploitation phase (ANOVA [n = 11,268], effect of rat [df = 5, F = 125, $p < 10^{-100}$, $\eta^2 = 0.05$], effect of phase [df = 2, F = 0.73, p = 0.48], effect of session [df = continuous, F = 1,656, $p < 10^{-100}$, $\eta^2 = 0.1$], interaction between phase and session [df = 2, F = 7.0, p = 0.0009, $\eta^2 = 10^{-3}$]). These changes occurred along with an increase in the efficiency on the task (Figure S2C) (ANOVA [n = 11,268], effect of rat [df = 5, F = 1.8, p = 0.11,

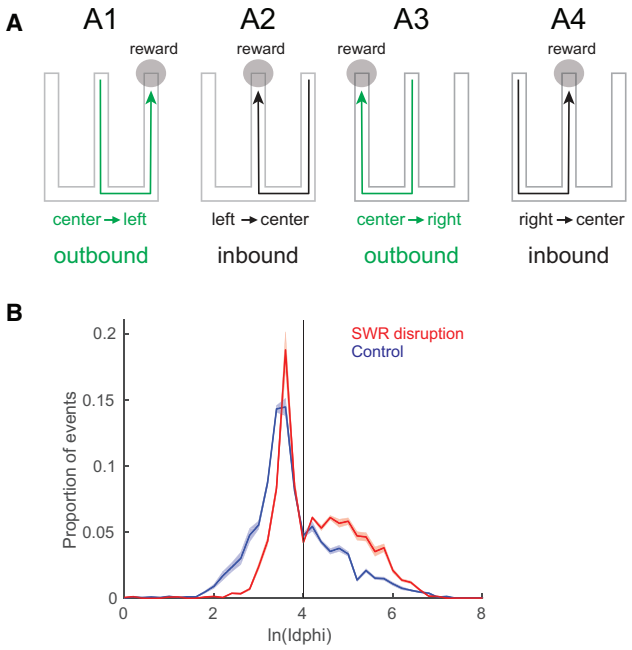

Figure 4. SWR Disruption Increases VTE Behavior in a Spatial Alternation Task

Data re-analyzed from Jadhav et al. (2012).

(A) Rats ran out from the center to a side arm (A1, outbound) and then returned to the center (A2, inbound). The following trial required visiting the alternate side arm (A3) before again returning to center (A4).

(B) VTE behavior during outbound trajectories was quantified by lnIdphi (see Supplemental Experimental Procedures). There were more high lnIdphi choice-point passes (VTE) in the SWR disruption animals as compared to the control group. Line and shaded area shows mean and SEM.

$\eta^2 = 0.06$], effect of session [df = continuous, F = 5.2, p = 0.025, $\eta^2 = 0.03$]). To measure the changes themselves, we measured the slopes of each (see Table S1).

We also found that the cumulative number of SWR events emitted within a session up to a given lap was negatively correlated with the presence or absence of a VTE event on that lap, even after controlling for the fact that VTE tended to occur on earlier laps and earlier phases of the task (ANOVA [n = 11,629], effect of phase, [df = 2, F = 45, $p < 10^{-100}$, $\eta^2 = 0.01$]; effect of lap, [df = continuous, F = 130, $p < 10^{-100}$, $\eta^2 = 0.01$]; effect of cumulative number of SWR events, [df = continuous, F = 5, $p < 0.02$, $\eta^2 = 10^{-4}$]; three-way interaction, [df = 3, F = 6, p = 0.005, $\eta^2 = 0.001$], even after including effects of rat, [df = 5, F = 18, $p < 10^{-100}$, $\eta^2 = 0.01$] and session [df = continuous, F = 378, $p < 10^{-100}$, $\eta^2 = 0.03$]). This observation is consistent with a role for awake SWRs in consolidation and the firming up of a map within a session (Buzsáki, 1989; O'Neill et al., 2006; Carr et al., 2011; Gupta et al., 2010). Moreover, it predicts that SWR disruption might influence VTE behavior.

To determine whether SWRs causally influence VTE behavior, we reanalyzed data from a spatial memory task in which SWRs were disrupted by stimulating the ventral hippocampal commissure when they were detected (Jadhav et al., 2012). In this task, rats were trained to alternate on a W-shaped maze (Figure 4A). Jadhav et al. found that disrupting SWRs led to an increase in working memory-dependent (outbound) errors, but not simple (inbound) errors. Outbound errors were defined as returning to the previous arm on the outbound journey instead of alternating. Inbound errors were defined as the animal not returning to the center arm on the return journey, whether incorrectly proceeding to the other outer arm or turning around and repeating a visit to the outer arm the animal was already in.

We hypothesized that the SWR disruptions would increase VTE events. We tested this hypothesis by comparing lnIdphi scores between SWR-disrupted and control rats. lnIdphi does not normalize within rat, allowing better comparison across groups of rats (see Supplemental Experimental Procedures).

We found that disrupting SWRs produced a dramatic increase in VTE, as evidenced by an increase in choice-point passes with high lnIdphi scores (Figure 4B, rank-sum test between distributions, [n = 1,461 control, 1,719 disruption] $z = -17$, $p < 10^{-67}$, Cohen's D = 0.6, z-proportion test for fraction of VTE events, $z = 12$, $p < 10^{-100}$). The increase in VTE due to SWR disruption could not be explained solely by the increase in error trials in the disruption animals (Figures S3A and S3B; ANOVA [n = 3,180]: control versus disruption, [df = 1, F = 138, $p < 10^{-100}$, $\eta^2 = 0.04$], correct versus error trials, [df = 1, F = 67, $p < 10^{-100}$, $\eta^2 = 0.02$], interaction, [df = 1, F = 11, $p < 0.001$, $\eta^2 = 0.003$]) (see Figure S3C). We also found that this difference in VTE behavior between disruption and control animals was seen both during initial learning and during performance in later days (Figure S3D). The SWR disruption occurred throughout learning and disrupted the rate of learning on the outbound decisions of the W-task.

DISCUSSION

The working hypothesis of the hippocampal field is that information processing underlying cognitive processes depends on sequential representations expressed during theta cycles (including during VTE events) and SWR events. We examined the interplay between VTE and SWR events and found that SWR rates were diminished following VTE events and that VTE events were diminished following increased SWR rates. We also found a negative relationship between the number of SWR events emitted within a single session and the number of VTE events in that session and that a disruption of SWR events led to an increase in VTE events.

While older theories of SWR function suggested a primarily offline role in consolidation of recent memories (Buzsáki, 1989; Sutherland and McNaughton, 2000), newer observations have found SWR sequences related to immediately available future options (Diba and Buzsáki, 2007; Pfeiffer and Foster, 2013; Singer et al., 2013), as well as backward paths (Foster and Wilson, 2006; Davidson et al., 2009; Gupta et al., 2010; Wikenheiser and Redish, 2013), novel paths (Gupta et al., 2010; Dragoi and Tonegawa, 2011; Ólafsdóttir et al., 2015), and other environments (Jackson et al., 2006; Karlsson and Frank, 2009; Silva et al., 2015).

An intriguing possibility is that the non-theta hippocampal states in which SWRs occur are similar to introspective representations that may parallel the default network in humans (Buckner et al., 2008) and that theta states may parallel executive

function in humans (Garavan et al., 2002). One theory is that SWRs reflect processes exploring the cognitive space of the task to find connections (Samsonovich and Ascoli, 2005). This hypothesis would suggest that SWRs may concentrate on areas of particular interest and importance, which would be consistent with the small increase in representation of future plans (Diba and Buzsáki, 2007; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015), recent experiences (Wilson and McNaughton, 1994; Jackson et al., 2006; Singer and Frank, 2009), and novel paths in an environment (Gupta et al., 2010; Dragoi and Tonegawa, 2011). It would also be consistent with the observation that on tasks where it is important to maintain non-recently experienced portions of the maze, it is those non-recently experienced portions that are more represented (Gupta et al., 2010). Thus, a syncretic hypothesis would be that SWRs play a role in establishing, sustaining, and exploring the cognitive map, which is a form of schema development, and very much in line with a generalization of consolidation theories.

This syncretic hypothesis is supported by our data. It suggests that a disruption of SWR events should lead to increased confusion and deliberation and that increased SWR events would decrease confusion in both the short and long term. Moreover, our data that the presence of VTE diminishes the subsequent rate of SWR events suggest that in familiar environments, SWRs primarily play this role after stable behaviors are established.

Both theta and SWR sequences play roles in various hippocampal functions. We found that these two processes interacted, shifting from variable behaviors that included a preponderance of VTE events to less variable behaviors marked by an increase in SWR events. The occurrence of VTE at a decision reduced the number of subsequent SWRs at a reward site, and a preponderance of SWRs on a given lap diminished the likelihood of subsequent VTE events. A selective interruption of SWRs during learning led to an increased prevalence of VTE. These dynamics imply a complex interaction between theta and SWR sequences, suggesting that while VTE sequences may reflect an immediate decision-making process, SWR sequences may reflect ongoing consolidation and planning processes that depend on and predict uninterrupted behavior.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, four figures, and one table and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2016.10.028.

AUTHOR CONTRIBUTIONS

The DD experiment was designed by A.E.P. and A.D.R. A.E.P. collected the DD data and analyzed it with A.D.R. S.P.J. and L.M.F. designed the W-maze experiment and collected the sharp-wave disruption data, which was analyzed by M.C.Z. and S.P.J. All five authors co-wrote the paper.

ACKNOWLEDGMENTS

This work was supported by NIH grants MH080318 (A.D.R.), MH090188 (L.M.F.), and MH100284 (S.P.J.), as well as an Alfred P. Sloan Fellowship (FG-2015-65675, S.P.J.) and a NARSAD Young Investigator Grant (23566, S.P.J.).

Received: March 4, 2016
Revised: July 14, 2016
Accepted: September 22, 2016
Published: November 17, 2016

REFERENCES

Alvarez, P., and Squire, L.R. (1994). Memory consolidation and the medial temporal lobe: a simple network model. Proc. Natl. Acad. Sci. USA *91*, 7041–7045.

Amemiya, S., and Redish, A.D. (2016). Manipulating decisiveness in decision making - effects of clonidine on hippocampal search strategies. J. Neurosci. *36*, 814–827.

Bett, D., Murdoch, L.H., Wood, E.R., and Dudchenko, P.A. (2015). Hippocampus, delay discounting, and vicarious trial-and-error. Hippocampus *25*, 643–654.

Breton, Y.A., Seeland, K.D., and Redish, A.D. (2015). Aging impairs deliberation and behavioral flexibility in inter-temporal choice. Front. Aging Neurosci. *7*, 41.

Buckner, R.L., and Carroll, D.C. (2007). Self-projection and the brain. Trends Cogn. Sci. *11*, 49–57.

Buckner, R.L., Andrews-Hanna, J.R., and Schacter, D.L. (2008). The brain's default network: anatomy, function, and relevance to disease. Ann. N Y Acad. Sci. *1124*, 1–38.

Buzsáki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. Neuroscience *31*, 551–570.

Buzsáki, G. (2015). Hippocampal sharp wave-ripples: A cognitive biomarker for episodic memory and planning. Hippocampus *25*, 1073–1188.

Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. Nat. Neurosci. *14*, 147–153.

Cheng, S., and Frank, L.M. (2008). New experiences enhance coordinated neural activity in the hippocampus. Neuron *57*, 303–313.

Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. Neuron *63*, 497–507.

de Lavilléon, G., Lacroix, M.M., Rondi-Reig, L., and Benchenane, K. (2015). Explicit memory creation during sleep demonstrates a causal role of place cells in navigation. Nat. Neurosci. *18*, 493–495.

Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat. Neurosci. *10*, 1241–1242.

Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. Nature *469*, 397–401.

Ego-Stengel, V., and Wilson, M.A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. Hippocampus *20*, 1–10.

Feng, T., Silva, D., and Foster, D.J. (2015). Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. J. Neurosci. *35*, 4890–4902.

Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature *440*, 680–683.

Foster, D.J., and Wilson, M.A. (2007). Hippocampal theta sequences. Hippocampus *17*, 1093–1099.

Garavan, H., Ross, T.J., Murphy, K., Roche, R.A.P., and Stein, E.A. (2002). Dissociable executive functions in the dynamic control of behavior: inhibition, error detection, and correction. Neuroimage *17*, 1820–1829.

Gardner, R.S., Uttaro, M.R., Fleming, S.E., Suarez, D.F., Ascoli, G.A., and Dumas, T.C. (2013). A secondary working memory challenge preserves primary place strategies despite overtraining. Learn. Mem. *20*, 648–656.

Girardeau, G., and Zugaro, M. (2011). Hippocampal ripples and memory consolidation. Curr. Opin. Neurobiol. *21*, 452–459.

Gupta, A.S., van der Meer, M.A.A., Touretzky, D.S., and Redish, A.D. (2010). Hippocampal replay is not a simple function of experience. Neuron *65*, 695–705.

Gupta, A.S., van der Meer, M.A.A., Touretzky, D.S., and Redish, A.D. (2012). Segmentation of spatial experience by hippocampal θ sequences. Nat. Neurosci. *15*, 1032–1039.

Hassabis, D., and Maguire, E.A. (2011). The construction system in the brain. In Predictions in the Brain: Using our Past to Generate a Future, M. Bar, ed. (Oxford University Press), pp. 70–82.

Hikosaka, O., Nakahara, H., Rand, M.K., Sakai, K., Lu, X., Nakamura, K., Miyachi, S., and Doya, K. (1999). Parallel neural networks for learning sequential procedures. Trends Neurosci. *22*, 464–471.

Jackson, J.C., Johnson, A., and Redish, A.D. (2006). Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. J. Neurosci. *26*, 12415–12426.

Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. Science *336*, 1454–1458.

Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. J. Neurosci. *27*, 12176–12189.

Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. Nat. Neurosci. *12*, 913–918.

Masimore, B., Kakalios, J., and Redish, A.D. (2004). Measuring fundamental frequencies in local field potentials. J. Neurosci. Methods *138*, 97–105.

Mazur, J. (1997). Choice, delay, probability and conditioned reinforcement. Anim. Learn. Behav. *25*, 131–147.

Nadel, L., Hupbach, A., Gomez, R., and Newman-Smith, K. (2012). Memory formation, consolidation and transformation. Neurosci. Biobehav. Rev. *36*, 1640–1645.

O'Keefe, J., and Nadel, L. (1978). The Hippocampus as a Cognitive Map (Oxford University Press).

O'Neill, J., Senior, T., and Csicsvari, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. Neuron *49*, 143–155.

Ólafsdóttir, H.F., Barry, C., Saleem, A.B., Hassabis, D., and Spiers, H.J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. eLife *4*, e06063.

Papale, A.E., Stott, J.J., Powell, N.J., Regier, P.S., and Redish, A.D. (2012). Interactions between deliberation and delay-discounting in rats. Cogn. Affect. Behav. Neurosci. *12*, 513–526.

Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. Nature *497*, 74–79.

Redish, A.D. (1999). Beyond the Cognitive Map: From Place Cells to Episodic Memory (MIT Press).

Redish, A.D. (2013). The Mind within the Brain: How We Make Decisions and How Those Decisions Go Wrong (Oxford University Press).

Redish, A.D. (2016). Vicarious trial and error. Nat. Rev. Neurosci. *17*, 147–159.

Samsonovich, A.V., and Ascoli, G.A. (2005). A simple neural network model of the hippocampus suggesting its pathfinding role in episodic memory retrieval. Learn. Mem. *12*, 193–208.

Schmidt, B., Papale, A., Redish, A.D., and Markus, E.J. (2013). Conflict between place and response navigation strategies: effects on vicarious trial and error (VTE) behaviors. Learn. Mem. *20*, 130–138.

Silva, D., Feng, T., and Foster, D.J. (2015). Trajectory events across hippocampal place cells require previous experience. Nat. Neurosci. *18*, 1772–1779.

Singer, A.C., and Frank, L.M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. Neuron *64*, 910–921.

Singer, A.C., Carr, M.F., Karlsson, M.P., and Frank, L.M. (2013). Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. Neuron *77*, 1163–1173.

Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus *6*, 149–172.

Squire, L.R. (1987). Memory and Brain (Oxford University Press).

Sutherland, G.R., and McNaughton, B. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. Curr. Opin. Neurobiol. *10*, 180–186.

Vanderwolf, C.H. (1971). Limbic-diencephalic mechanisms of voluntary movement. Psychol. Rev. *78*, 83–113.

Wang, Y., Romani, S., Lustig, B., Leonardo, A., and Pastalkova, E. (2015). Theta sequences are essential for internally generated hippocampal firing fields. Nat. Neurosci. *18*, 282–288.

Wikenheiser, A.M., and Redish, A.D. (2013). The balance of forward and backward hippocampal sequences shifts across behavioral states. Hippocampus *23*, 22–29.

Wikenheiser, A.M., and Redish, A.D. (2015a). Decoding the cognitive map: ensemble hippocampal sequences and decision making. Curr. Opin. Neurobiol. *32*, 8–15.

Wikenheiser, A.M., and Redish, A.D. (2015b). Hippocampal theta sequences reflect current goals. Nat. Neurosci. *18*, 289–294.

Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science *265*, 676–679.

Neuron, Volume *92*

Supplemental Information

Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making

Andrew E. Papale, Mark C. Zielinski, Loren M. Frank, Shantanu P. Jadhav, and A. David Redish

Supplemental Figure 1

Associated with Figures 2 and 3.

Local Field Potentials differed between departing the reward site and passing through the choice point. A Spectral auto-coherence (SAC) showing increased correlation at sharp wave frequencies (between 150 and 250 Hz). Spectral auto-coherence is calculated as the correlation of the spectrogram, collapsed across time. These auto-coherence plots can be used to identify transient events of consistent frequency. The sharp lines at 60 and 180 Hz are line noise that is uncorrelated to neural signals. **B**: Power spectral density (PSD) showing that there is power in the delta (2-4 Hz) and theta (6-10 Hz) bands. Panels **A** and **B** are taken over all data from all sessions on the DD task. **C**,**F**: SAC and PSD plots for the 3s preceding departure from the feeder-zone. Note the strong SWR events in the auto-coherence plot and the lack of strong theta power ($\theta < \delta$). **D**,**G**: SAC and PSD plots for the 3s preceding leaving the choice point. Note that there is little to no SWR signals in the choice point passes, but there is strong theta power ($\theta > \delta$). **E**,**H**: SAC and PSD plots for the 3s surrounding VTE events. Note that there is little to no SWR signals in the choice point passes, but there is strong theta power ($\theta > \delta$). The difference can be seen clearly in the subtraction plots (**I**,**J**). **K**,**L**,**M**: Frequency-normalized average spectrograms aligned to the departure from the reward site, VTE events, and delivery of reward. Note the gamma frequencies at the VTE event show increased high gamma and high-frequency oscillations as the animals re-orient at the turn around point.

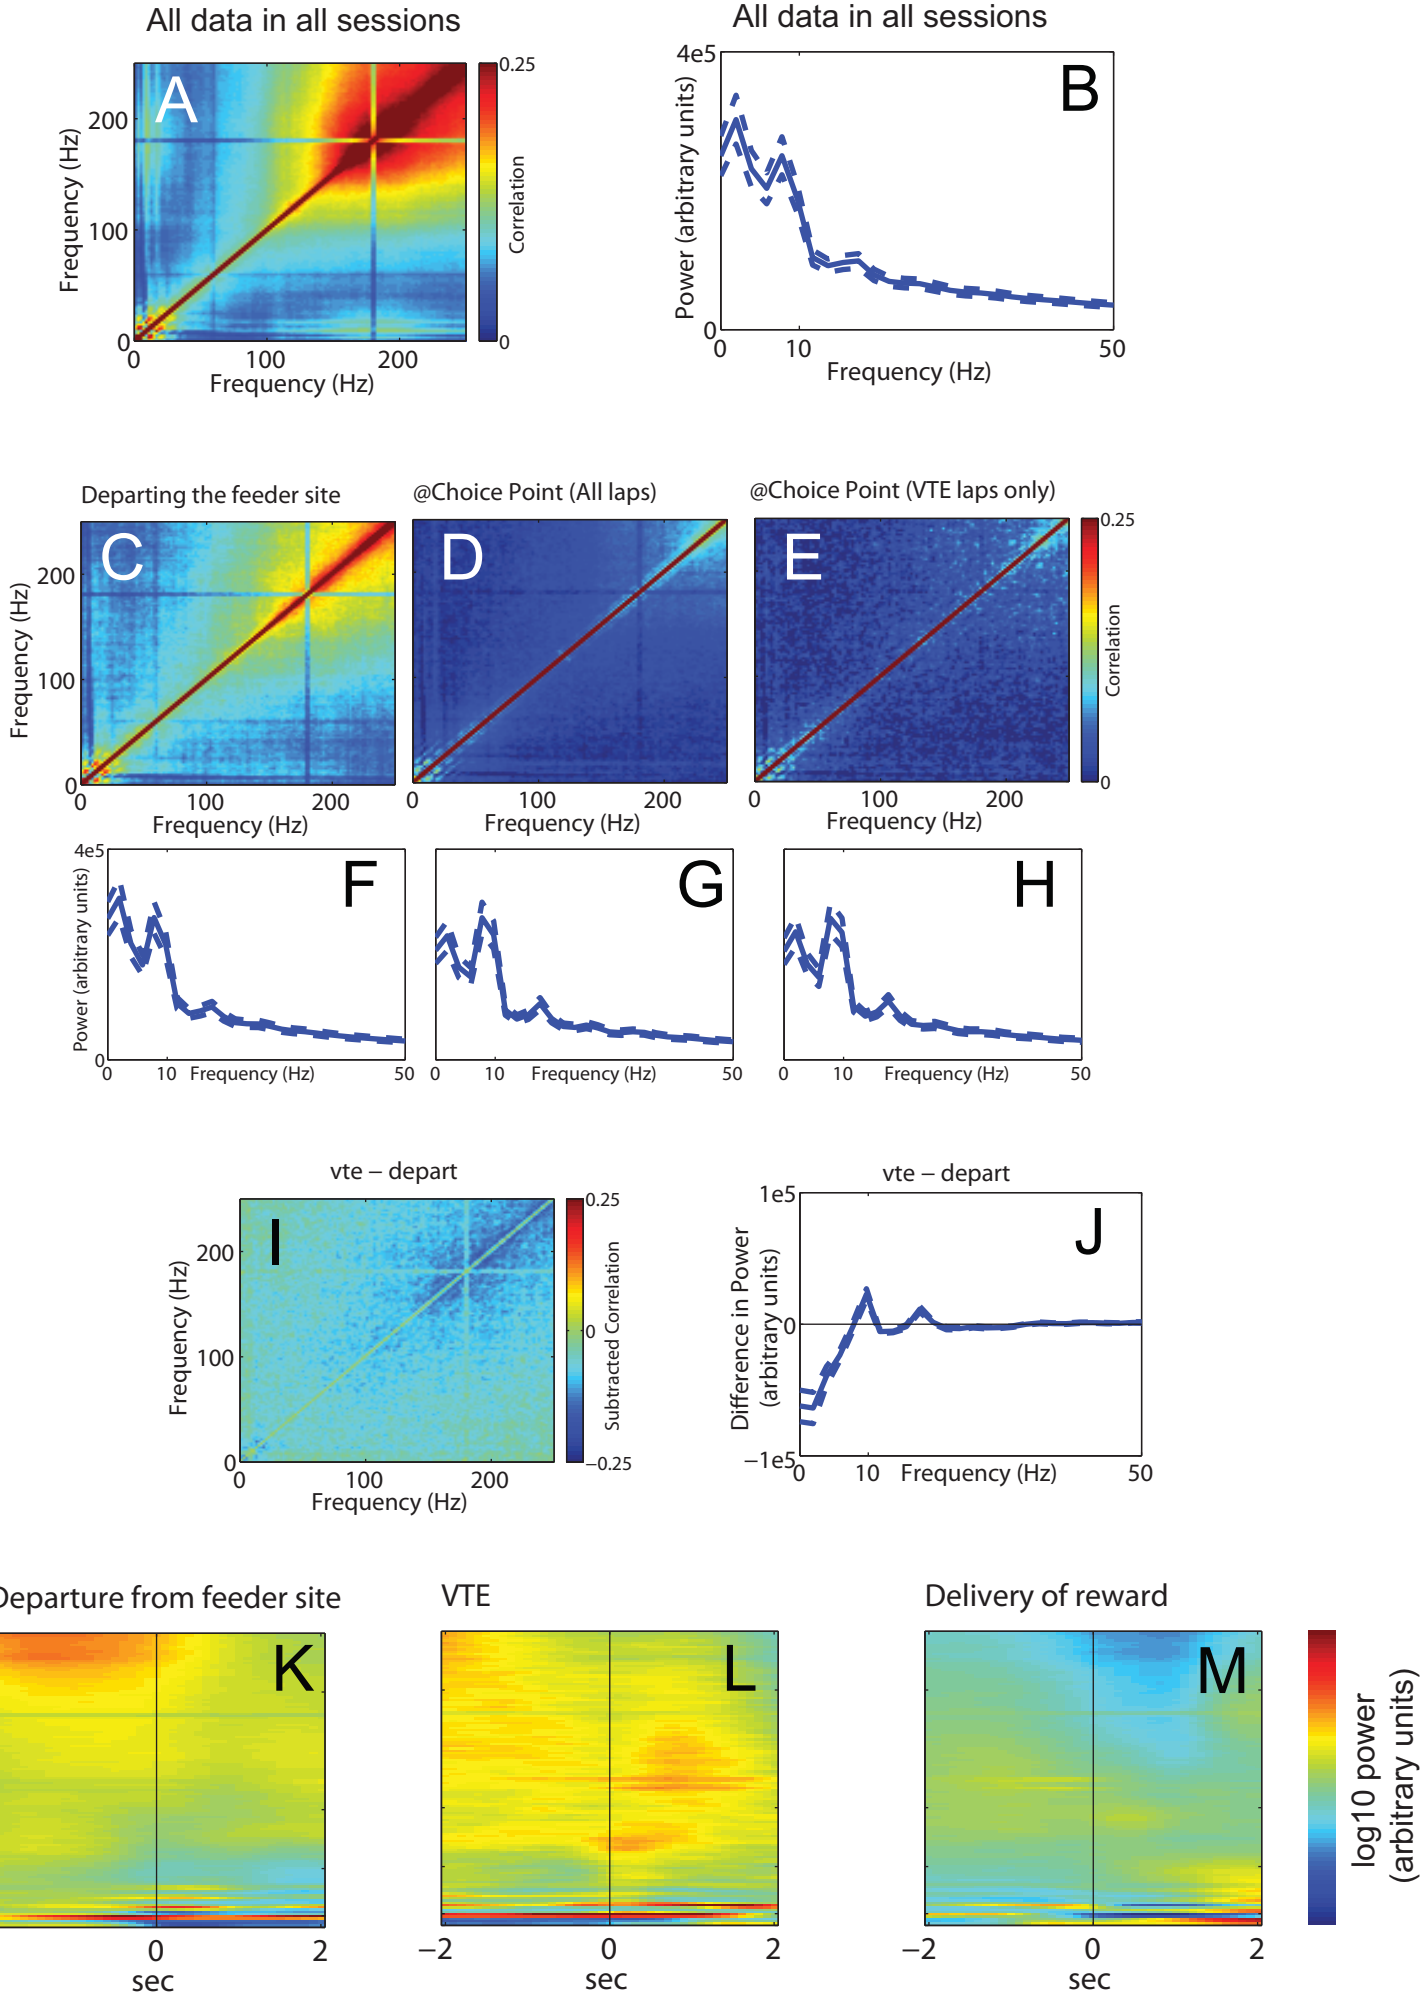
All data in all sessions
A
Frequency (Hz)
Correlation
All data in all sessions
B
Power (arbitrary units)
Frequency (Hz)
Departing the feeder site
@Choice Point (All laps)
@Choice Point (VTE laps only)
C
D
E
F
G
H
vte – depart
I
Subtracted Correlation
vte – depart
J
Difference in Power (arbitrary units)
Departure from feeder site
VTE
Delivery of reward
K
L
M
Hz
sec
log10 power (arbitrary units)

Supplemental Figure 2

Associated with Figure 3.

A: The rate of SWR events increased over sessions. This increase occured during all three phases, but was most consistent during the exploitation phase. The three inset panels show the same axes as the main panel, but only for laps restricted to each phase (investigation, titration, exploitation). **B:** The amount of VTE (measured by zIdphi) decreased over sessions. This decrease was driven entirely by changes during the exploitation phase. The three inset panels show the same axes as the main panel, but only for laps restricted to each phase (investigation, titration, exploitation). **C:** The alternation efficiency (see **supplemental methods**) increased over sessions. Boxplots show IQR (box), median (line), and standard error of the median (notch).

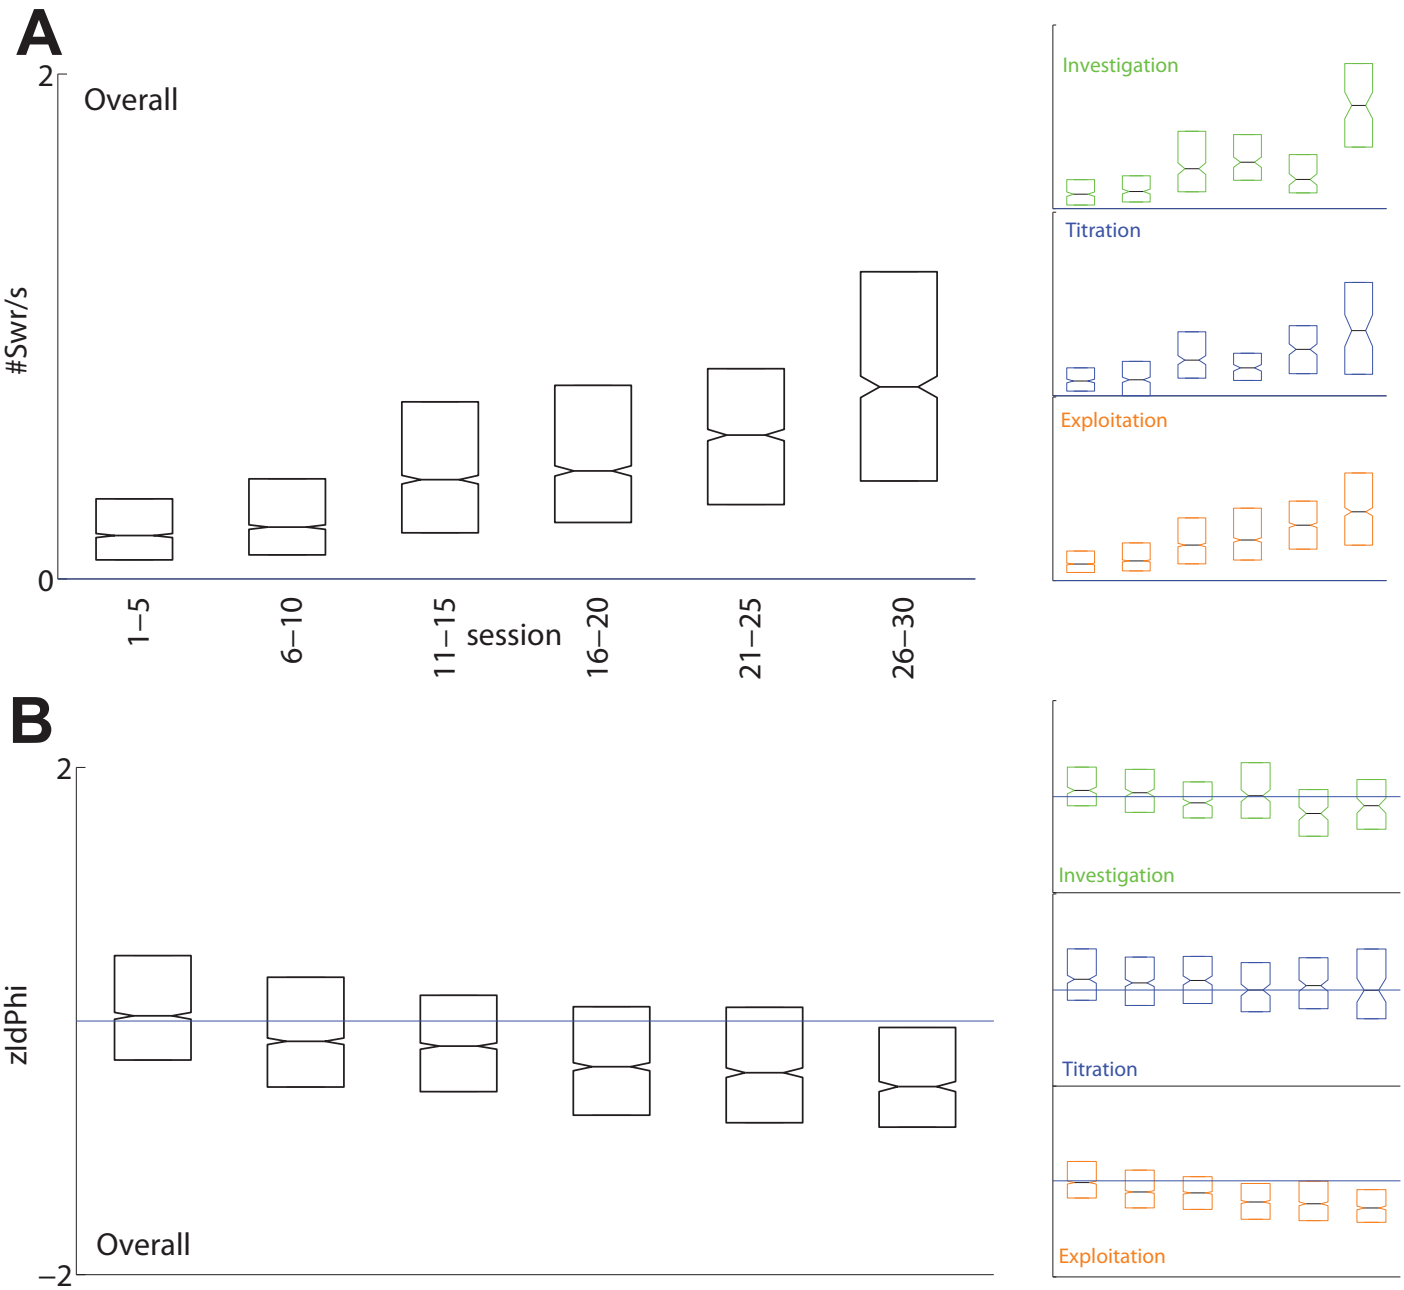
A
Overall
2
0
#Swr/s
1–5
6–10
11–15
16–20
21–25
26–30
session
Investigation
Titration
Exploitation

B
2
−2
zIdPhi
Overall
1–5
6–10
11–15
16–20
21–25
26–30
session
Investigation
Titration
Exploitation

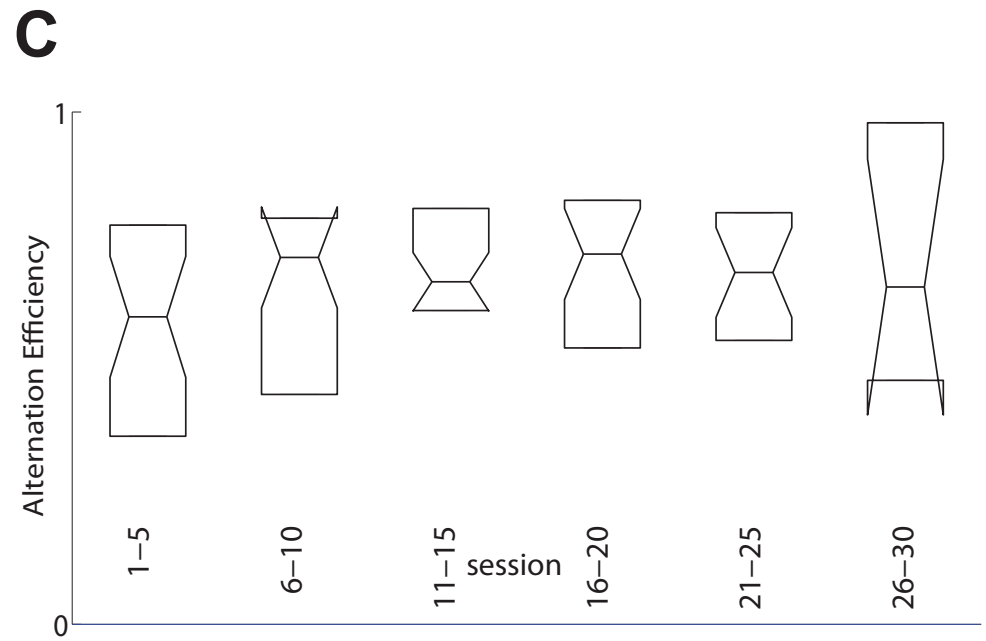
C
1
0
Alternation Efficiency
1–5
6–10
11–15
16–20
21–25
26–30
session

Supplemental Figure 3

Associated with Figure 4.

Effect of SWR disruption on VTE is not due to increase in errors, and persists across learning and performance of the alternation task. A. To determine whether the increase in VTE associated with SWR disruption was simply the result of more outbound errors, we separated our analysis of the effect of SWR disruption on VTE events into correct and incorrect trials. Distribution of lnIdphi limited to correct trials still showed an increase in high lnIdphi choice point passes for disruption animals as compared to control ($p < 10^{-6}$). **B**. We also saw an increase in high lnIdphi choice point passes for disruption animals during incorrect trials ($p < 10^{-10}$). **C**. There was a significant interaction ($p < 0.001$, see text) between the control / distruption groups and correct / error trials. Error bars show SEM. Asterisks indicate significant comparisons $p < 0.05$ as measured by a post-hoc Tukey test.

D. SWR disruption was maintained throughout the experimental period, which allowed us to ask what impact SWR disruption had on VTE behavior over learning. We saw a general decrease in VTE behavior over learning for both control and disruption animals, in agreement with delay discounting task. On the first day of task acquisition (when animals explored the novel W-track), the SWR disruption group matches control in proportion of VTE events. A difference in proportion of VTE events in SWR disruption animals was seen on all days except the first day. Differences were seen even on the final days (days 7 and 8) when animals performed mostly above chance-level (Jadhav et al., 2012).

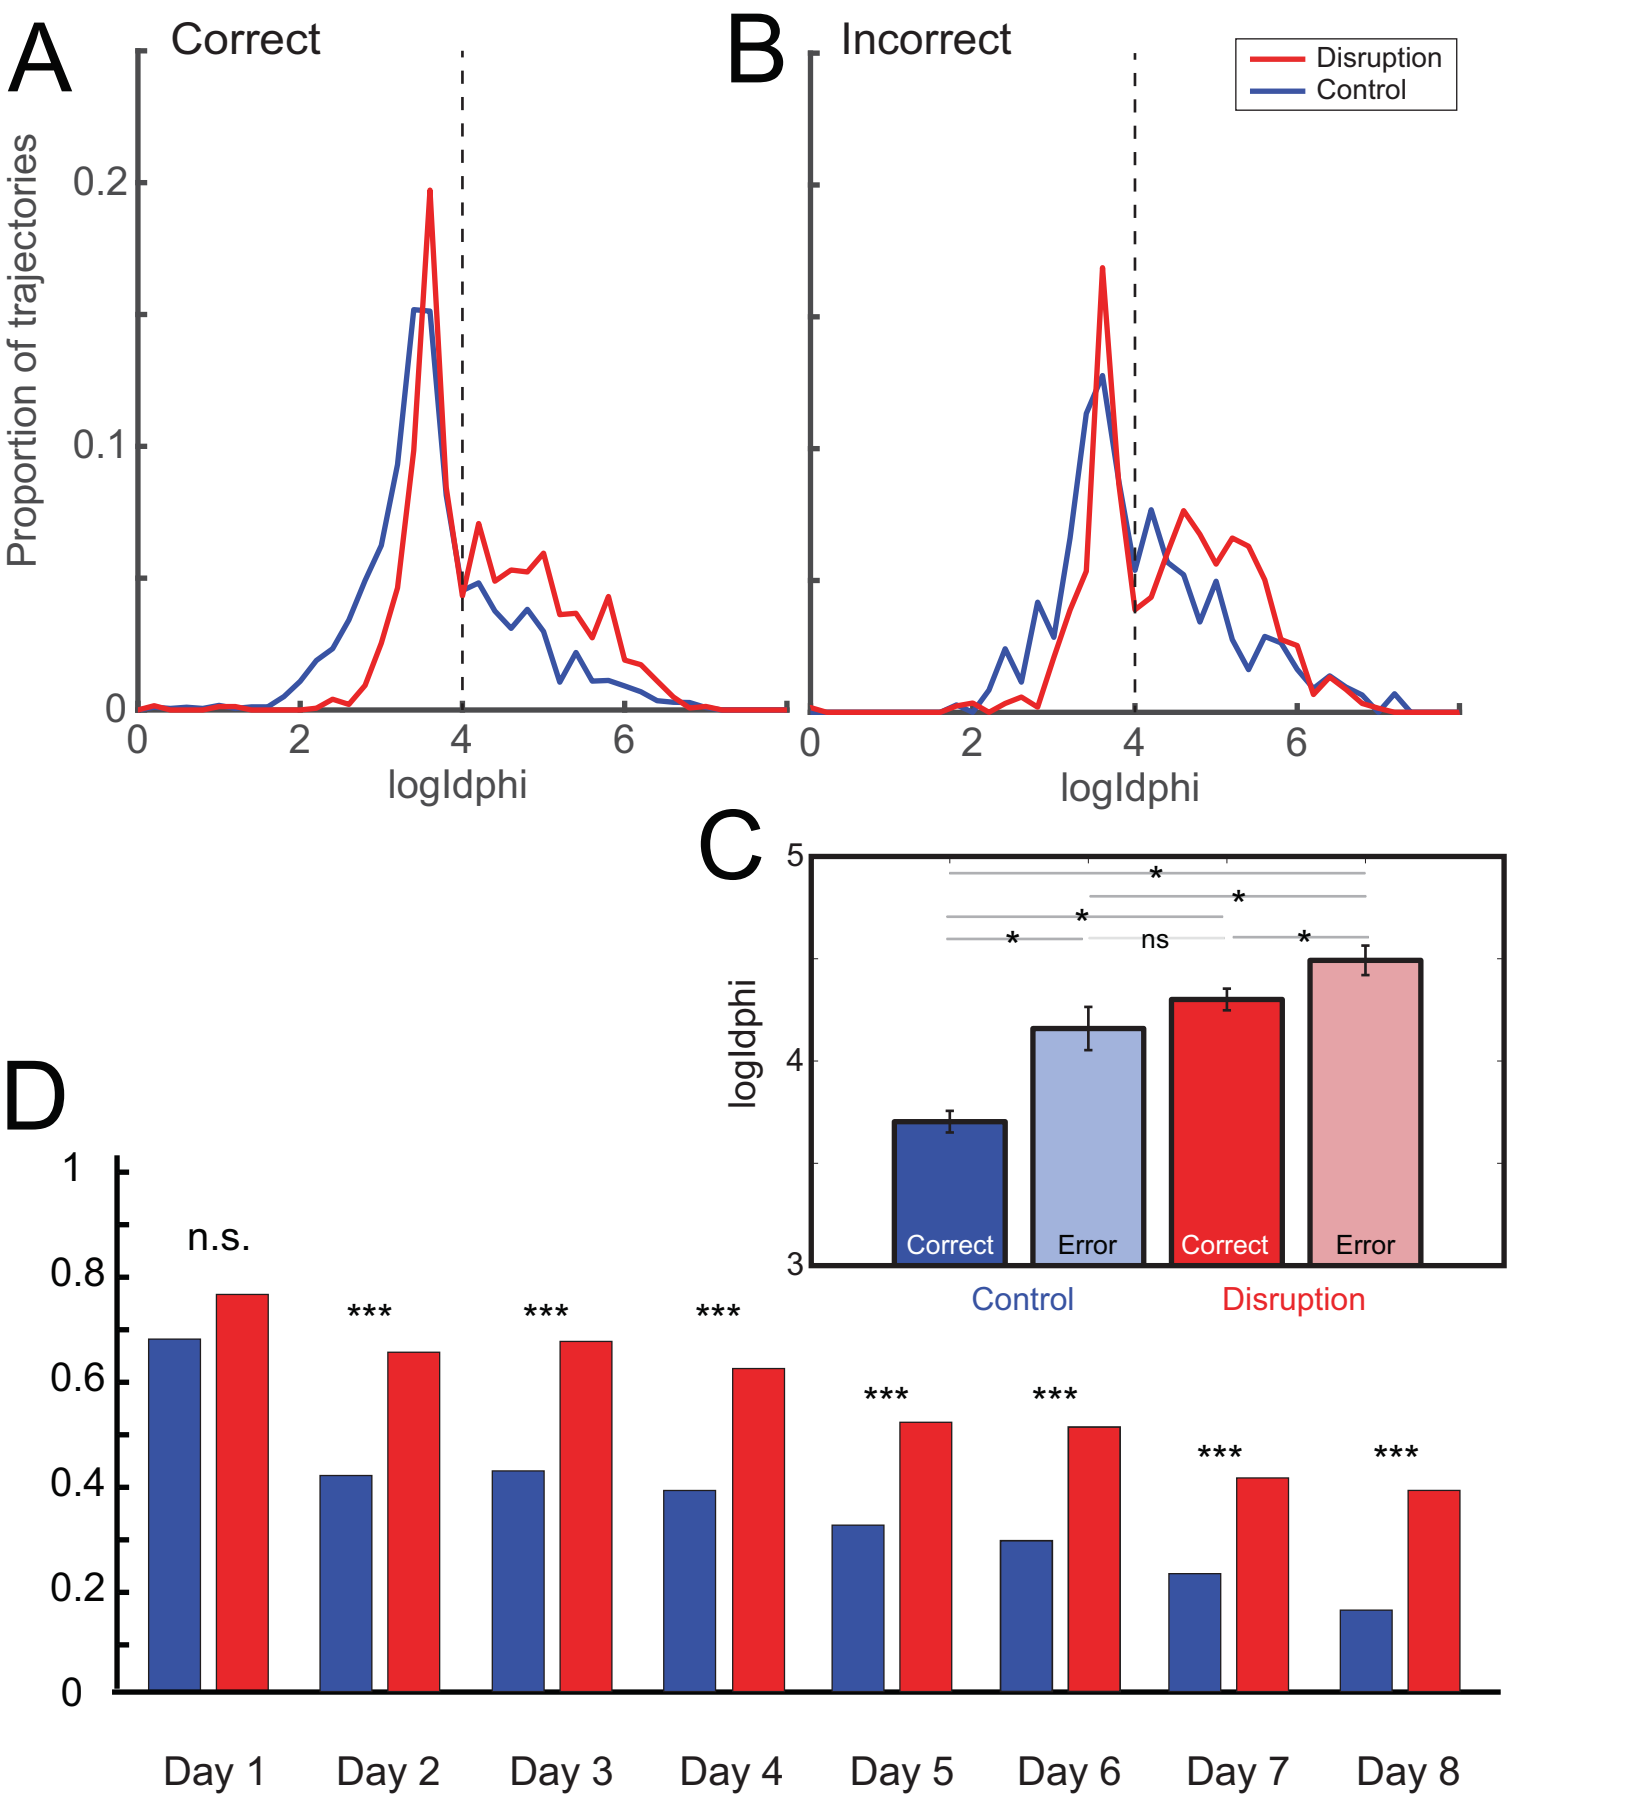
A
Correct
B
Incorrect
Disruption
Control
Proportion of trajectories
0.2
0.1
0
0
2
4
6
logIdphi
C
5
4
3
logIdphi
*
*
*
*
*
ns
Correct
Error
Correct
Error
Control
Disruption
D
1
0.8
0.6
0.4
0.2
0
n.s.
***
Day 1
Day 2
Day 3
Day 4
Day 5
Day 6
Day 7
Day 8

Supplemental Figure 4

Associated with Figures 1 and 4.

Normalization procedures applied to the two experiments. The questions associated with the DD task require comparisons between laps occuring in each animal, so we normalized the lnIdphi measures (**A**) by z-scoring each measurement within rat (**B**). Because the questions associated with the W-maze task require between-animal comparisons, we did not normalize within-rat (which can obscure differences) and used the lnIdphi measure directly. See **supplemental methods**.

Experiment 1: VTE and SWR interactions on the DD task.
(within rat comparisons require within-rat normalization)

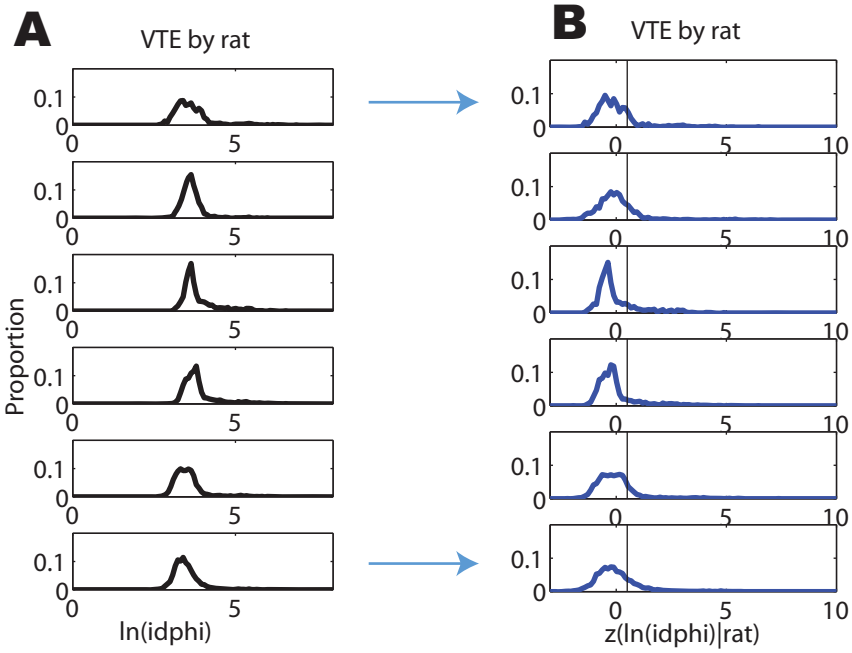

Experiment 2: Effect of SWR disruption on VTE on the W-maze task
(between rat comparisons preclude within-rat normalization)

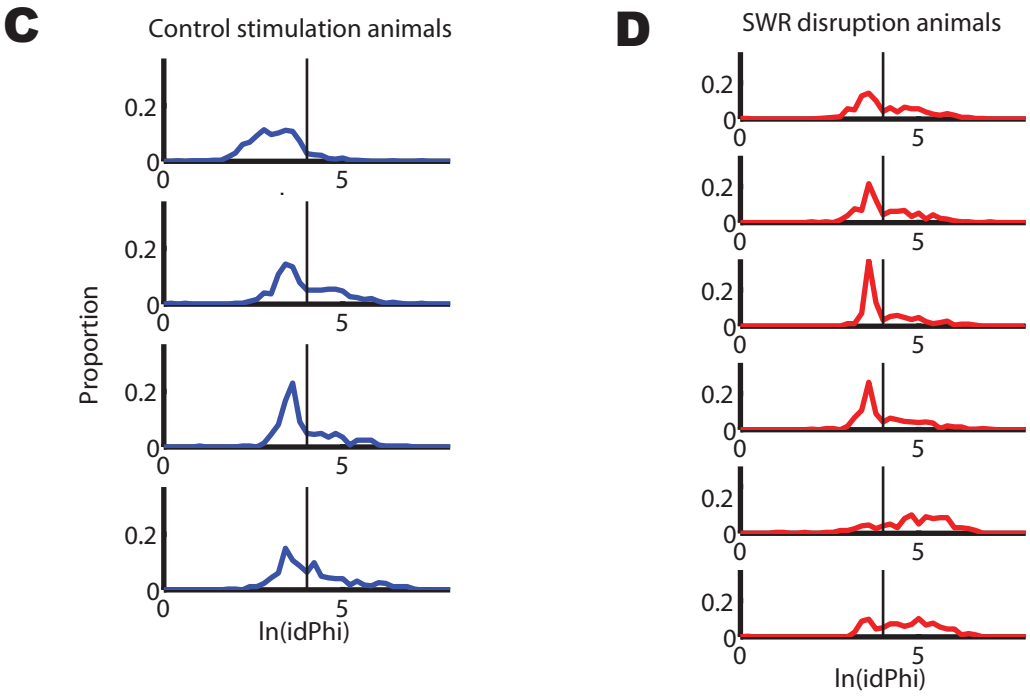

Supplemental Material (5): Table 1

Statistics for changes observed over days in the DD task. We applied ANOVAs to each analysis, taking into account phase of the task and session number, and measuring their interaction. For each phase, for each process, we measured the slope using a single-variable linear regression and measured the probability that the slope as measured was different from zero.

VTE, change over sessions

ANOVA: effect of phase $p=10^{-27}$, effect of session $p=10^{-76}$, interaction $p=10^{-29}$

Phase	Slope	p-value (>0)
OVERALL	-0.19	$p<10^{-93}$
Investigation	-0.13	$p<10^{-05}$
Titration	-0.09	$p<0.001$
Exploitation	-0.21	$p<10^{-82}$

SWR/s, change over sessions

ANOVA: effect of phase $p=10^{-20}$, effect of session $p=10^{-75}$, interaction $p=10^{-100}$

Phase	Slope	p-value (>0)
OVERALL	+0.43	$p<10^{-100}$
Investigation	+0.46	$p<10^{-66}$
Titration	+0.39	$p<10^{-56}$
Exploitation	+0.42	$p<10^{-100}$

Alternation efficiency, change over sessions

ANOVA: effect of session $p<0.025$

Phase	Slope	p-value (>0)
OVERALL	+0.16	$p<0.05$

Supplemental Material: Experimental Procedures

DD task

Subjects. Six adult male Fisher 344 Brown Norway rats (Harlan, Indianapolis, IN) age 8-12 months at the beginning of the experiment were trained on the Delay Discounting (DD) task. Rats were individually housed on a 12 hour light / dark cycle and food restricted to no more than 80% of ad lib weight with water available ad lib. All experimental procedures were conducted in compliance with National Institute of Health guidelines for animal experimentation and approved by the Institutional Animal Care and Use Committee at the University of Minnesota.

Behavioral Task and Training The behavioral training, task design, and lap categorization have been described previously (Papale et al., 2012). At the same time each day (±2hr), rats were trained to run 100 laps / session on a T-maze version of an adjusting-delay discounting task. Each session, one reward site would deliver a larger-later reward (3×45mg sucrose pellets, Test Diet) and the other would deliver a smaller-sooner reward (1×45mg pellet). Rats were required to wait from 1-30s prior to delivery of the larger-later reward and 1s prior to delivery of the smaller-sooner reward. The delay preceding the larger-later reward was adjusted based on the decisions of the rat. Repeated choices to the larger-later side increased the delay by 1s on the subsequent lap, while repeated choices to the smaller-sooner side decreased it by 1s on the subsequent lap. These are called *adjustment* laps. *Alternation* between sides left the delay unchanged. Delays were accompanied by a tone-sequence countdown with each delay matched to a specific pitch. The tones began after leaving the choice point, so analysis of choice point behavior (and associated neurophysiology) did not include tone cues.

Rats ran one session per day. Each rat ran 30 sessions before being implanted with tetrodes for neural ensemble recordings, were then re-trained as tetrodes were lowered into the hippocampus. Once tetrodes were recording hippocampal neurons, rats ran an additional 30 sessions (1 session / day). The recording sessions are analyzed here.

Rats were tracked by LEDs on the recording headstage from an overhead camera with a resolution of 0.17 cm / pixel and at a speed of 60Hz.

Vicarious Trial and Error (VTE) behavior was quantified with the zIdphi measure (the z-scored, integrated absolute angular velocity ($d\phi$), see Papale et al., 2012). In short, zIdphi measures the integrated angular velocity of the orientation of motion. After calculating the Idphi measure, we took the $\ln$ of it to reduce the skewness. Because the first experiment (DD) depended on comparisons between laps occuring in each animal, we then z-scored the $\ln$ Idphi measure using means and standard deviation calculated from all laps for a given rat. These transformations are shown in **Fig. S4**. In order to identify a threshold to separate VTE laps from non-VTE laps, we found the 5 most populous bins in the histogram shown in **Fig. 1D**, took the mean zIdphi of those 5 bins, then took all the samples below that mean, reflected them around the mean, and used that pseudo-sample to define an expected Gaussian distribution of non-VTE laps. The observed distribution diverged from the expected Gaussian at zIdphi =0.5. Therefore, for the DD task, VTE laps were classified as those with zIdphi > 0.5.

The Idphi measure includes both speed through the choice point as well as variability in the path through the choice point. While Idphi does not provide a sharp decision boundary between VTE and not, previous studies have found that choice point passes with Idphi above the decision point tend to be VTE laps with pauses and a high behavioral variability, while those below tend to be fast and regular passes with a consistent speed indicative of more ballistic movements (van der Meer and Redish, 2010; Blumenthal et al., 2011; Steiner and Redish, 2012; Stott and Redish, 2014; Amemiya and Redish, 2016). This decision variable (defined a priori from previous studies) was sufficient to identify significant differences in the hippocampal ensemble representation (see **Fig. 2.**)

Behavioral analysis: Rats on the DD task typically show behaviors indicative of three phases: an investigation phase, a titration phase, and an exploitation phase (Papale et al., 2012). Importantly, these phases were not imposed on the task, but rather measured from the rat's behavior. *Investigation* was defined as laps before the first adjustment lap. *Exploitation* was defined as laps following the first time the rat reached the indifference point for that day. The **indifference point** was measured as the mean adjusted delay on the delayed side over the final 20 laps of the session. *Titration* was defined as the laps between investigation and exploitation.

Alternation efficiency was defined as the number of alternation laps when the delay was within 3s of the final measured indifference point divided by the total number of alternation laps. **Path stereotypy** was measured by first defining laps as proceeding from

departure from one feeder (measured as the time crossing a spatial threshold) and arrival at the next (measured, again, as the time crossing a spatial threshold). The 2D $\langle x, y \rangle$ path was resampled to 1000 samples. For each pair of laps (L_i, L_j), the distance between the paths was defined as the mean distance between the matched samples. Thus paths that are both spatially and temporally matched will have small distance measures. Separate analyses were done for left-left (LL), right-right (RR), left-right (LR), and right-left (RL) laps. Analyses were done for each session, and then averaged across sessions. Lap pairs from different paths (LL x RR, LR x RL, etc) were defined as `NaN` and did not enter into the analyses.

Neural Recording. Rats were implanted with 12-tetrode hyperdrives (Kopf) using standard procedures targeting right dorsal CA1 (-3.8mm AP, -3.0mm ML). Over the next 10-14 days, tetrodes were lowered at a rate of 20-320μm / day, with larger adjustments initially and finer adjustments as target depth was approached (1200-1800μm). Local field potentials were used to align tetrodes as they approached CA1, and CA1 depth was confirmed when putative pyramidal cells that spiked during flat sharp wave ripples were observed.

Action potentials were recorded at a rate of 32kHz within a 1ms window if, on any of the four channels, voltage exceeded a user-defined threshold. Spikes were then filtered (600-9000Hz) and digitized (Neuralynx, Cheetah, Tuscon, AZ). The local field potential voltage was sampled from one channel per tetrode at 2kHz and bandpass filtered (1-475Hz). Spike sorting was carried out offline using an automatic k-means clustering algorithm (KlustaKwik, Harris et al.), and then manually using waveform parameters in a multi-dimensional space (MClust 3.5, Redish et al.).

SWR Analysis. SWRs were analyzed while rats paused at the feeders before beginning the next lap. Pause time was taken from the time of reward delivery to the time that rats left the feeder zone, defined as a circle of pixels recorded from the overhead camera. SWR rate was defined as the number of SWR at the feeder site normalized by the pause time.

SWRs were detected using the tetrode with the most cells per session. The local field potential from that tetrode was bandpass filtered from 150-250Hz and then a Hilbert transform was applied. The instantaneous amplitude obtained from the Hilbert was

z-scored and global extrema were computed with those greater than 2.5 standard deviations above mean amplitude were considered SWR. A SWR that occurred less than 250ms from a prior event was discarded.

Bayesian Decoding. A Bayesian algorithm was applied to analyze the information content of CA1 ensembles during SWR and choice point passes (Zhang et al., 1998). If a cell had an average firing rate of less than 10Hz, its firing rate was binned into 50ms windows. Tuning curves were defined from the cell's occupancy-normalized firing in a 56×56 spatial grid using data from the training set. All spikes during theta epochs and those not occurring during the ± 125ms window around SWR events were used for the training set. The posterior probability was then computed for each 50ms time bin, producing a probability distribution across the 56 x 56 grid at each time step. Analysis zones were defined as four rectangles (see **Fig. 2B**). Four of the six rats had sufficiently large neural ensembles (> 10 cells / session) to do Bayesian decoding analyses. All six rats were included in the SWR and behavioral analyses.

Spectral Analyses. Spectral analyses were done using standard spectral methods (spectrogram function, Matlab) with a Hamming window sized to be able to access 2 Hz. Power spectral density plots were generated by averaging across these spectra to remove the temporal dimension. Frequency-frequency autocoherence plots (Masimore et al., 2004) were generated by correlating across these spectra to remove the temporal dimension.

W-Maze task

Animals. 14 male Long Evans rats weighing 450-600 grams were used in this study. All procedures were approved by the Institutional Animal Care and Use Committee at the University of California, San Francisco and conformed to National Institutes of Health guidelines.

Behavior. Animals were food deprived to 85-90% of their baseline weight and were pre-trained to alternate between two ends of a linear track for liquid food reward (evaporated milk) dispensed automatically. Animals were tested on the novel W track for

8 days, two 15min sessions per day. Food wells at the end of each arm automatically dispensed liquid food reward according to alternation rules previously described (Jadhav et al., 2012). In brief, center arm visits were rewarded when the previously rewarded (n-1) food well was located in one of the side arms. A side arm well visit was rewarded when the previously rewarded food well (n-1) was located in the center arm, and the previously visited side arm (n-2) was the opposite side arm, that is, the animal had successfully alternated visiting side arms. Repeated visits to the same reward well and incorrect alternations were not rewarded. At the conclusion of each run session, animals were placed in a familiar rest box for 15min.

Stimulation properties. Surgical procedures to implant electrodes were as previously described (Jadhav et al., 2012). In addition to 12-tetrodes targeting CA1 (-3.6 mm AP,+2.2 mm ML), one or two tungsten stimulation electrodes were implanted targeting vHC (-1.3 mm AP, $\pm$1 mm ML) ipsilaterally or bilaterally to CA1 recording tetrodes. Stimulation pulses were biphasic and 0.2 ms in duration, with amplitude calibrated to induce 100 ms of inhibition of CA1 multiunit activity (ranging from 40-180μA).

Real time SWR detection and disruption. For full details, see (Jadhav et al., 2012). Briefly, we chose local field potentials for 5-6 tetrodes, filtering in the ripple band (20 tap band-pass IIR filter, 100-400 Hz). To establish a threshold for disruption, we calculated smoothed means and standard deviations of the absolute value of the LFP on each tetrode. The threshold for disruption was then set to 4-6 s.d. above the mean on at least two tetrodes. For control animals, we introduced a latency of 150-200 ms between onset of stimulation and SWR detection, effectively decoupling the two and maintaining SWR content.

Position tracking and reconstruction. Illuminated IR LEDs were attached to the front and back of the animals' headstage during recording for position and speed reconstruction. Rat behavior sessions were then tracked via fixed overhead monochrome CCD camera at 30 Hz with a 0.45 cm / pixel resolution. We reconstructed position using a semi-automated thresholding analysis using custom software (MATLAB, Mathworks). Position was then smoothed using a non-linear method (Jadhav et al., 2012) and used to separate the behavior into individual trials (behavior trajectories) based on well position

and W alternation sequence order (returning to center well, inbound; proceeding to alternating side arm, outbound). Trajectories were classified as correct or incorrect based on correspondence between alternation sequence order and reward wells visited. Outbound, memory dependent trajectories were selected for further analysis.

Choice zone identification. The choice zone of each trajectory was defined as a circle with a radius of 15 cm, centered at the intersection point of the middle maze arm and the remainder of the W track. The first instance of the animal crossing into the choice zone and the first instance of the animal leaving the choice zone were considered the start and end of the choice zone trajectory. These positions were used for subsequent ldphi and VTE calculations.

lnIdphi and VTE calculation. We calculated ldphi for each choice-zone trajectory, using equivalent methods to that used in the DD task. Because the W-maze analyses depend on between-animal comparisons, we did not normalize the VTE measures within rats, but used the lnIdphi measure directly. For the W-maze task, VTE laps were classified as those with a lnIdphi measure greater than 4.

lnIdphi and VTE behavior was quantified similarly for SWR disruption (n = 6) and control stimulation (n = 4) animals.