Hippocampal Representations on the Spatial Delay Discounting Task

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Andrew E. Papale

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy

A. David Redish, adviser

June, 2015

Acknowledgements

I am most grateful to Dr. David Redish, who consistently succeeds at creating the positive atmosphere conducive to spark scientific discovery. His enthusiasm and insight were instrumental for completion of this work. I would like to thank the technicians Kelsey Seeland and Christopher Boldt without whom this research would not exist. I would like to thank Dr. Adam Johnson, whose research formed the foundations of my own. I would also like to thank Dr. Jadin Jackson and Dr. Matthijs van der Meer who both served as mentors through the early years of graduate school. Finally, I would like to thank all the members of my class: you made the early years of graduate school memorable.

Dedication

To the professors at Juniata College who inspired me to pursue science and allowed me to reach my full potential: Dr. Norman Siems, Dr. James Borgardt, Dr. Jamie White, Professor Mary Atchley, and Dr. Mark Pearson. And to my family, who's purest convictions sustained me through difficult times: Raymond Papale, Christine Papale, Keppy Arnoldsen, Ron Arnoldsen, Yvonne Koths, Steve Koths, Babcia, and Leeland Koehler.

Abstract

The hippocampus facilitates planning. While navigating an environment, hippocampal representations can be forward of the current position. During pausing, hippocampal representations can be predictive of future choices, and can represent configurable components of an environment. Here I investigate these mechanisms, both of which have been linked with decision-making involving the hippocampus. I find that they are related processes: the representations during pausing and navigation are similar on a lap-by-lap basis.

Forward representation during navigation has been linked with vicarious trial and error behavior at a choice point. To study the relationship between this behavior and the hippocampal representations that may accompany it, I describe a novel neuroeconomic task, the spatial delay discounting task. The occurrence of vicarious trial and error behavior on this task is consistent with a deliberative decision-making process. The hippocampal representations during vicarious trial and error support this interpretation. I investigate theories of vicarious trial and error behavior and its relation to hippocampal function.

Contents

List of Figures

Chapter 1

Introduction

The hippocampal contribution to decision-making is facilitated by the interaction of different network states at different times. Here I describe previously unknown similarities in the information content of two types of network activity in the hippocampus: sharp wave replay (SWR) sequences and theta sequences. Each process has a unique time course associated with specific behavior and network states, though they interact in a complimentary fashion on a neuroeconomic task where rats employ multiple strategies.

In this Thesis, I describe the behavior of rats on a spatial delay discounting task. Their behavior is consistent with a neuroeconomic model of decision-making that incorporates a representation of the value of each choice. The spatial delay discounting task is an ideal platform for testing theories of Vicarious Trial and Error (VTE) behavior. I discuss several such theories and suggest that VTE is best described in relation to deliberation. The electrophysiology of hippocampal ensembles on the spatial delay discounting task is characterized. The complex pattern of behavior on the task lends itself to discovery of previously unknown interactions between hippocampal decision-making processes. Two such phenomena have been proposed as substrates for decision-making in the hippocampus: SWR sequences and theta sequences. Here, I show that their information content is related, but their occurrence is anti-correlated. This duality allows flexible decision-making to occur on multiple timescales or during different phases of learning.

- Chapter 1 outlines the theory of deliberation, a decision-making process that includes evaluation of the outcome of choices.
- Chapter 2 describes Vicarious Trial and Error (VTE) behavior, and six theories about its function. Each theory makes unique predictions about the timing of the behavior on the spatial delay discounting task. I suggest that VTE behavior is related to the process of deliberation.
- Chapter 3 reviews the current state of the field of hippocampal electrophysiology in the awake behaving rat. Sharp wave ripple (SWR) sequences and theta sequences are characterized.
- Chapter 4 describes the role of the hippocampus in deliberation, highlighting gaps in the knowledge about the mechanism.
- In Chapter 5, I describe the development of a novel neuroeconomic task for electrophysiological recordings, the spatial delay discounting task.
- In Chapter 6, I describe the discounting behavior and timing of VTE on the spatial delay discounting task. This behavior is consistent with evaluation of goal-based outcomes of choices during deliberative search. I find that VTE occurs during periods requiring increased behavioral flexibility, suggesting it is associated with the process of deliberation.
- In Chapter 7, SWR sequence and theta sequence representations are described on the spatial delay discounting task. I find that they contain related information, but occur at different times. Their interaction provides a framework for hippocampal processing that could not be achieved with each component alone.
- Chapter 9 describes the effects of a cannabinoid agonist on the spatial delay discounting task. I find that cannabinoids increase VTE, but diminish forward representation during VTE.
- Chapter 10 provides an integrated perspective on the role of SWR and theta sequences and their relative contribution to the process of deliberation. I consider how these discoveries may relate to VTE and deliberation. Directions for future research are suggested.

Chapter 2

Theory of Deliberation

> "No single event can awaken within us a stranger totally unknown to us. To live is to be slowly born. It would be a bit too easy if we could go about borrowing ready-made souls."

Antoine de Saint-Exupery [1]

2.1 What is Deliberation?

When faced with a difficult decision, one for which there are no cached action routines or physiologically embedded responses, animals can either explore their environment in search of new information or conduct a cognitive search through pre-existing knowledge structures to find a useful experience on which to base the current decision. The quote from Antoine de Saint-Exupery captures the constructive nature of deliberation, using components of our existing neural repertoire to respond to novel situations.

Deliberation is a neural mechanism for cognitive search. It explains decision-making based on predictions about the outcome of decisions. Many decisions can not be explained without taking this processing step into account. For example, the decision about what to wear to work today may be influenced by imagining the reactions of co-workers. (If I wear that flashy tie, I might get more attention than if I wear that gray tie). Deliberation has five components: schemas, a process to search through

[1] (de Saint-Exupéry, 1969)

the schemas, a process to evaluate outcomes, an action-selection process, and working memory (Redish, 2013). This thesis will focus on schemas and the process to search through the schemas, both have been linked with the hippocampus.

Schemas describe a general property of memory: the underlying neural framework for the incorporation of new knowledge into existing places in memory (Reed, 2012). While the concept of a schema does not specify a neural mechanism, the existence of a schema is revealed by facilitation of learning different chunks of information that have some relational similarity. An example is latent learning, where pre-exposure to an environment facilitates subsequent learning of goal-directed behavior in that environment (Thistlethwaite, 1951). If a satiated rat is allowed to explore a track and locate a food source, it will find the food more readily than a naive rat if it returns to the track when hungry. In this case, the configuration of the environment (the spatial map) is the 'chunk' and learning is facilitated because there is a cognitive map upon which novel inferences can be drawn (if hungry, then go to food). The concept of a schema goes beyond chunking of spatial information. Expertise in a particular area may also be explained by the formation of schemas. A chessmaster viewing a chess board can glean more information than a novice in a shorter amount of time. After viewing a chess board, the expert player is able to chunk the pieces based on properties of mutual defense, proximity, color, and type (Chase and Simon, 1973; Linhares and Freitas, 2010).

Cognitive search through schemas may occur by evolution of cell assemblies. The cell assembly hypothesis is based on Hebb's postulate that the connection strength between neurons increases with synchronized activity of the presynaptic and postsynaptic cell (Hebb, 1949). The cell assembly is then formed from mutual excitation among cells in a distributed network across structures. Cell assemblies are chained together by temporal evolution through a multidimensional space. Harris (2005) identifies key properties of cell assemblies including 1) structured temporal patterns in spike trains that are not associated with external stimuli, and 2) coordinated variability in firing patterns among cells within an assembly, reflecting an underlying cognitive state. Another key property of cell assemblies is pattern completion, the ability to reconstruct a complete output from an incomplete or noisy input.

Hypothalamic and midbrain structures involved in attention and arousal likely

govern the relative contribution of the deliberative system in terms of load and integration time. If I look at my watch and see that I'm going to be late for work, I might short-circuit the deliberative process and grab whichever tie is in front of me. If the rat is punished for waiting at a choice point too long by delivery of an electric shock, it may begin to make faster decisions without evaluating them completely.

The 'subroutine' of deliberation may be called repeatedly, each iteration producing a new variation of potential outcomes until a choice is made. Deliberation is computationally slower than other decision-making processes owing to the increased memory demand and processing time of simulating future outcomes. Working memory is required because multiple outcomes, or the properties of those outcomes, need to be kept in mind. The simulation itself takes time to use these pieces, configuring them in different (serially ordered) ways until a decision is reached.

Structures involved in episodic future thinking are engaged during deliberation, including the hippocampus, but also potentially a constellation of other structures (Buckner and Carroll 2007; Hassabis and Maguire 2007, discussed in **Section** 5.2). The inclusion of hippocampus in the circuit may depend specifically on using a locale navigation strategy (O'Keefe and Nadel, 1978), more generally with the incorporation of any episodic information in the computation (Eichenbaum and Cohen, 2014), or most generally with any imagination of future outcomes (Schacter and Addis, 2007). Hippocampus is likely also crucial for recognition of a context that triggers a deliberative process. For example, recognizing you are at a choice point and not knowing which way to go, or recognizing the importance of which tie to wear to work.

Deliberation may include determining which dimensions to base the construction of value, an operation that would require integration of different permutations of multimodal information. For example, when confronted with the choice of which tie to wear to work, I might consider the color of my jacket, or the messiness of my closet, along with the reaction of my co-workers. (The flashy tie might clash with my black jacket. My closet might be too messy to search for the flashy tie). [2]

Evaluation of simulations from episodic future thinking is hypothesized to be carried out by the ventral striatum (vStr, or nucleus accumbens). vStr may be the gateway

[2]This would be consistent with investigation of parameters in a known state space, operations defined in Chapter 3.1 **Vicarious Trial and Error**.

to the basal ganglia, receiving excitatory projections from subiculum, CA1, medial prefrontal cortex, and amygdala. It also has reciprocal connections with ventral tegmental area and substantia nigra (Salgado and Kaplitt, 2015), structures important for a variety of decision making systems.[3] vStr medium spiny neurons respond to reward (Lavoie and Mizumori, 1994), and reward responses are reactivated during deliberation (van der Meer and Redish, 2009). vStr is a necessary component for delaying gratification in a delay discounting task (Cardinal et al., 2001). These factors make it a likely source of evaluation of episodic simulations during deliberation.

[3]Ventral tegmental area is the primary source of dopaminergic input to hippocampus and prefrontal cortex. It also receives indirect excitatory input from CA3 via the septal nucleus (Luo et al., 2011).

Chapter 3

Vicarious Trial and Error

Vicarious Trial and Error (VTE) behavior occurs at a choice point when a rat pauses and looks in one direction, then the other, before making a decision (Figure 3.1A). The term 'vicarious trial and error' carries an interpretation of the rat imagining the possible outcomes of its choice. When the term was coined, the behavior suggested this was the case (Muenzinger and Gentry, 1931), but only recently has neurophysiological evidence supported this interpretation (Johnson and Redish, 2007). This evidence has led to two related theories of deliberative decision-making that explain the occurrence of VTE as a behavioral manifestation of an underlying neural search process.

However, there is far from a consensus on this conclusion and additional theories need to be considered. Since its discovery, VTE has been referred to in the literature as turn around behavior (van der Meer and Redish, 2009), error correcting behavior, detours (Ferbinteanu and Shapiro, 2003), look-ahead scans (Gupta et al., 2013) or 'oops' behavior (Wang et al., 2014). Here, I investigate the potential role of VTE in different decision-making processes.

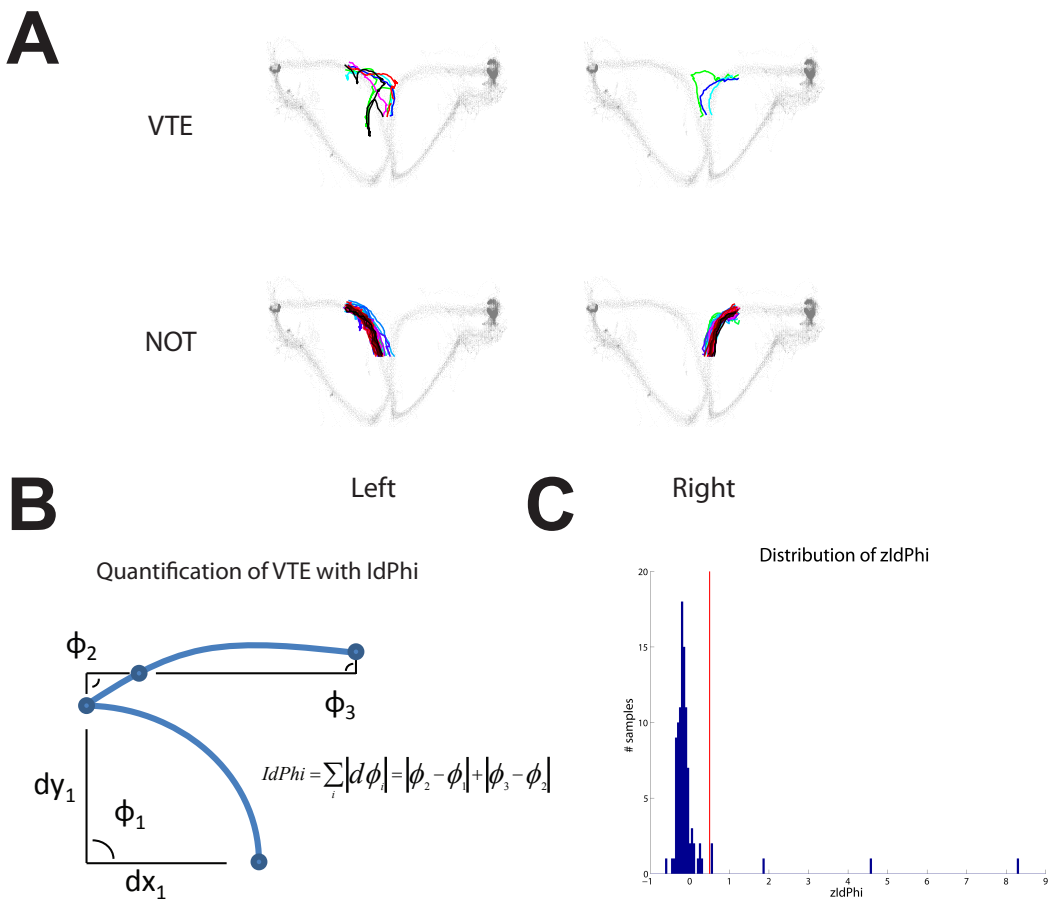

Figure 3.1: **Vicarious Trial and Error.** Vicarious Trial and Error (VTE) behavior at a choice point occurs when a rat looks in one direction and then the other before making a choice. A) Position tracking data is plotted for one example session (gray dots). VTE appears as squiggly lines on choice point passes. Trials with VTE (top row) are plotted separately for choices to the left and right, with each trial a different color. Trials that do not have VTE are ballistic trajectories through the choice point (bottom row). Here, VTE trials were chosen based on the presence of a manually-identified turn around time. B) The variability in orientation of motion through the choice point can be quantified by integrating the change in angular velocity ϕ across a choice point pass (IdPhi). VTE trajectories with high variability in orientation produce high IdPhi measurements, and non-VTE trajectories with minimal change in orientation produce low IdPhi measurements. C) The distribution of z-scored log_{10} IdPhi values for the example session. The red line indicates a typical threshold cutoff of 0.5, with five VTE events identified above this threshold.

3.1 Is VTE Conditioned Orienting?

VTE occurs during sensory discrimination tasks, and has been hypothesized to reflect conditioned orienting toward sensory cues of different associative strengths (Figure; 3.2A; Bower 1959; Spence 1960). This would place VTE as a consequence of Pavlovian action-selection processes (Redish, 2013), and its timing should be explained by simple stimulus response models. Bower (1959) hypothesized that both the choice behavior of rats and the occurrence of VTE could be described by a first-order Markov model. Orienting to each option was given a state, and choice of each option was given a state. In Bower's model, probabilities of transition between states were assigned in two different ways, giving a four or five parameter model. These specific predictions were tested by Still (1976) who found that the terms in the first-order Markov models of VTE were non-stationary and a poor description of the behavior. The model of Spence (1960), developed from a Hullian account, can be shown to be mathematically equivalent to one of the Markov models proposed by Bower (1959). It is difficult to reconcile the findings of Still (1976) with a conditioned orienting theory.

Hippocampus has been implicated in learning Pavlovian trace conditioning, when an unconditioned stimulus (US) is separated from a conditioned stimulus (CS) by a delay. VTE has not been reported during this paradigm, which does not typically use spatial cues. A binary choice trace conditioning experiment is analogous to a choice between two fixed-delay rewards on a T-maze runway task. It is unknown whether the Markov models of Bower (1959) explain VTE for the case of binary choice trace conditioning.

Initially, VTE was found to co-occur with an improvement in performance on a sensory discrimination learning task. It was therefore originally hypothesized to be a learning-associated behavior (Muenzinger and Gentry, 1931; Tolman, 1938). Visual discrimination learning entails recognition of differences in the color, pattern, or shape of visual stimuli. One paradigm Tolman employed was the Lashley jumping platform. In this experiment, the rat would have to jump from a starting platform to an escape platform, labeled by either a black or a white cue card. After initially choosing at chance levels, performance increased in parallel with an increase in VTE. Once animals reached asymptotic performance, VTE diminished. If a noticeable change in

either of the stimuli was introduced, VTE would again briefly be observed. Hippocampus is not required for simple visual discrimination tasks (Mishkin, 1954; Correll and Scoville, 1965; Samuels, 1972; Marston et al., 1993). However, Hu and Amsel (1995) found that electrolytic hippocampal lesions impaired performance on a sensory discrimination task, also decreasing the number of VTE during learning. Bett et al. (2012) tested the effect of neurotoxic lesions of hippocampus on VTE behavior during a visual discrimination task. No differences in performance or VTE were observed between groups, failing to replicate the results of Hu and Amsel (1995).[1]

Visual discrimination among multiple landmarks can be used to guide navigation. In the rodent navigation literature, using visual landmarks to guide decision-making has been called taxon navigation (O'Keefe and Nadel, 1978). One interpretation of VTE is that it is being used to perform distance estimation of distal cues using motion parallax (Ellard et al., 1984). In this study, gerbils were tested on a Lashley jumping platform and were observed to make vertical head movements before jumping, as if gauging the distance to jump. Distance estimation by visual parallax may be a confound in the interpretation of VTE behavior on the Lashley jumping platform. However, variability in VTE behavior at a choice point where the distances between landmarks is held constant cannot be explained by this theory. Neither can VTE during simple visual discrimination.

VTE behavior during deliberation may be a vestigial remnant of stimulus localization. A computational study, supported by behavioral observations in fruit bats, suggests that optimal localization of a stimulus occurs when pointing off-axis (Yovel et al., 2010). The theory that the head motion of VTE is a vestigial remnant of stimulus localization may explain reports of VTE in sharks (Graeber et al., 1978), cockroaches (Barraco et al., 1981), crayfish (Capretta and Rea, 1967), and lizards (Vance et al., 1965). I would not expect to observe more than rudimentary forms of deliberation in these animals, but would not be surprised if they could simulate trajectories to avoid moving obstacles. On the other hand, none of these reports have been replicated and further analysis is required to verify that these are not illusory movements.

[1]This effect in Hu and Amsel (1995) may have been a result of damage to inferotemporal cortex (Mishkin, 1954; Gross, 1973).

3.2 Is VTE Value Integration?

Another possibility is that VTE is a value-integration process in a race-to threshold type decision between fixed-value options (Figure 3.2B). In a race-to-threshold decision model, different neuronal subsets (in some brain structure) are associated with each choice, and they actively compete to reach a net activity threshold. As a rat looks in the direction of one option, activity of one subset of neurons increases; and when the rat looks in the opposite direction, activity of a separate subset of neurons increases. The subset that reaches a threshold of excitability first 'wins' and the rat selects that option. According to this hypothesis, more difficult decisions between similarly valued items would require longer integration times. Most studies of value-integration models have been conducted with eye tracking studies in monkeys and humans.

Saccade-fixate-saccade (SFS) visual scans in monkeys and humans may serve the same fundamental purpose as VTE behavior in rats. In a study of monkey eye movements during a discrimination learning task, it was found that the number of SFS scans increased with performance and diminished following asymptotic performance (Schrier and Povar, 1979). In a follow-up study, the authors 'punished' SFS by replacing an informative pair of stimuli with an uninformative pair if the monkey made more than one SFS. They found that discrimination learning and reversal learning were highly inefficient under the punishment condition (Schrier and Povar, 1982). Voss et al. (2011) studied a variant of SFS in healthy and amnesic subjects with hippocampal damage, noting similarities to VTE. Patients were given the option of viewing pictures that appeared in an ordered grid by moving a window using a joystick. Only one picture at a time was visible in the window. A 'spontaneous revisitation' was classified as moving the window back to a previously viewed picture. The authors found that 35% of trials had spontaneous revisitations in healthy subjects, but only 17% of trials for amnesic subjects. Healthy subjects were then asked to reposition the objects on the grid from memory. Objects that had been part of a spontaneous revisitation trial were more likely to be placed correctly. These results open up the possibility of a primate homologue to VTE.

SFS have been linked to evaluation. Krajbich et al. (2010) analyzed eye movements while human subjects made sequential binary choices about their preferred snack food.

Subjects had pre-rated the foods on a scale of -10 to 10, providing an independent measure of subjective value that was well-described by a psychometric choice curve. The authors found that the number of SFS decreased as the value difference between pairs of items increased. This result would predict the occurrence of VTE when the value between the two options on a T-maze is similar. Additionally, the authors found an increased probability of choosing an item as a function of fixation time on that object, first or last item fixated, and intrinsic gaze biases (e.g. increased likelihood of looking left). These results are consistent with the expected pattern given a race-to threshold model.

In the study of Krajbich et al. (2010), decisions were modeled as two-barrier race-to-threshold evidence accumulators, where evidence (in this case, subjective value of a snack food item) was integrated over time while the person looked at each food item. As the person looked in the direction of item A, evidence would be accumulated in one direction, and as they looked at item B, it would reverse direction. Integration would continue until the barrier for option A or the barrier for option B was reached (Krajbich et al., 2010). Presumably, the evidence accumulation would be reflected in the activity of two ensembles of neurons in some brain structure that was actively comparing value (or constructing value) online. For example, it has been suggested that ventromedial prefrontal cortex instantiates relative value signals compatible with this model (Lim et al., 2011). Stimulation of the identified brain structure may prevent either population from reaching the target threshold, impairing value-guided choice behavior. If the value-integration signal were also driving SFS, we might also expect to increase its occurrence.

3.3 Is VTE Exploration?

Exploratory learning behaviors can be divided into at least three separate cases: 1) exploration of a novel environment (estimation uncertainty or ambiguity); 2) learning about changing parameters from a known state space (expected uncertainty); and 3) spontaneous exploration due to recognition of unexpected changes (unexpected uncertainty; Winstanley et al., 2012, Redish, 2013). Exploration of novel environments is a behavioral procedure for decreasing ambiguity of an environment, reducing the

uncertainty in its size and structure. In the jumping platform experiment, this period of exploration would happen the very first time the rat was placed on the apparatus. Some questions that would be answered by exploration of a novel environment are "Where are the boundaries of the environment?", "Where are reward locations?", and "What are the routes between those locations?" The second type of learning is associated with a change in expected uncertainty. This occurs if the basic structure of the environment is known, but there are unknown parameters in it that need to be filled in. For example, if food locations in a familiar environment change each day, or if there is uncertainty about cues that would indicate which way to go. In the jumping platform example, this would occur once the rat knew that it needed to jump to one of the escape locations, but had not yet learned to use the cues to guide its behavior. The third type of exploratory learning occurs when something unpredictable changes in the environment. For example, in reversal learning tasks, reward contingencies are unexpectedly switched. To continue receiving reward after the reversal, the behavior must be updated to reflect the change. In the jumping platform task, this occurred when one of the cues was switched unexpectedly and the rat had to re-identify the correct cue. I will examine the potential role of VTE in each type of exploration.

3.3.1 Is VTE Exploration of a Novel Environment?

One possible behavioral state associated with VTE is the exploration of novel environments, those that have never before been experienced. This interpretation would imply that VTE at a 'high-cost' choice point is a residual of the variability in the paths taken by an animal during initial exploration of an unfamiliar place. Rats exhibit a somewhat stereotypical pattern when exploring a novel environment under low-stress conditions. First, a 'home base' is established in an environment. Then, repeated journeys begin from this location with each going farther from the origination point (Eilam and Golani, 1989; Russell et al., 2010). Rats will explore a novel environment spontaneously (Aggleton, 1985) and there is evidence that exploration can occur vicariously through observation (Dragoi and Tonegawa, 2011).

If VTE is simply a result of residual ambiguity at a choice point, any lateral head movements during initial exploratory behavior may constitute VTE. However, other

behaviors are present during exploration of a novel environment. Rearing and scanning behaviors are observed during exploration of a novel environment, as well as whisking and sniffing (Silva and Brandao, 2000; Monaco et al., 2014). Through scanning, whisking, and other exploratory movements, the rat on the jumping platform will begin to form a mental representation of its environment. These behaviors do not co-occur with VTE at a choice point, and this VTE is therefore unlikely to be a residual of initial exploration.

Hippocampal lesioned rats persistently explore an environment, showing a similar pattern of behavior even after repeated days of experience (Clark et al., 2005). In this study, control rats were found to have reduced measures of exploration after four days of experience in an open circular arena. These measures included total distance traveled, number of journeys from the home base, and number of head turns toward an orienting cue. In contrast, hippocampal lesioned rats had similar measures even after four days of experience. Similarly, amnesic patients have less efficient exploration of a visual scene after repeated days of experience (Yee et al., 2014) and fewer spontaneous revisitations when viewing sequences of pictures (Voss et al., 2011). Clark et al. (2005) describe head movements of rats toward and away from the orienting cue, interpreting them as VTE for the purpose of gauging distance to landmarks. This theory is discussed in the above section **Is VTE Conditioned Orienting?**. Rats with hippocampal lesions had high VTE throughout the four days of testing, linking hippocampal lesions with an increase in VTE during exploration of a novel environment.

3.3.2 Is VTE Investigation of Parameters in a Known Environment?

VTE may be a marker for exploration of alternatives in a known environment (Figure 3.2C). The basic structure of an environment can be known, but inherent variability can remain in the spatial location, temporal contiguity, or identity of environmental components. This is the what/when/where information that is thought to compose schemas of episodic memories in the hippocampus (Nyberg et al. 1996; Clayton et al. 2003; Morris 2006, discussed in **Section** 5.1). On the jumping platform example, the rat would already have a mental representation, or schema, of the experimental environment that includes the starting platform and the potential escape platforms. It will know there is a gap that it needs to jump across (or risk falling into). The rat may

also have observed and distinguished the two cue cards. However, in this state, it has not yet learned which cue card leads to an escape platform and which one leads to a dead end. Only through exploring one of the options can the rat learn this information and incorporate it into its schema of the environment. This process of investigating parameters in a known environment may be facilitated by VTE.

In a study by Bett et al. (2012), VTE timing was altered by hippocampal lesions during investigation of a familiar environment on a serial spatial reversal task. The serial spatial reversal task was a double-Y maze with two choice points and four potential goal locations. After finding the first reward, rats were given nine subsequent trials with the reward at the same location before the reward location was switched. On this task, hippocampal lesioned rats showed a persistent impairment in performance, as well as a reduced ability to remember the rewarded location in a block. On the first session, lesioned animals showed less VTE overall. Lesioned rats had reduced VTE before finding the reward in a block of trials, but had similar levels of VTE afterward (Bett et al., 2012). This result suggests that VTE in control animals was related to expected uncertainty in reward location. Because of the increase in VTE during switches in reward position, it is also related to unpredictable environmental change.

3.3.3 Is VTE Updating Expectations about Unpredictable Environmental Changes?

VTE behavior may also be associated with updating unpredictable changes of what/when/where information. The clearest example is during a reversal task, when the spatial location of the outcomes is switched. In order to continue to choose correctly on such a task, a rat has to adjust its behavior to the unprecedented change. In the jumping platform, this occurs when the experimenter swaps one of the cue cards and the rat has to update which one to choose. Reversal learning on a spatial reversal task has been shown to co-occur with increases in VTE (Blumenthal et al., 2011).

Locale navigation involves the use of a cognitive map with an internally stored representation of space (O'Keefe and Nadel, 1978; Redish, 1999). This representation is thought to be embedded within episodic memories in the hippocampus. An example of locale navigation is learning to always go to the "East" arm of a plus-maze task. Constant distal or proximal cues are used to determine "East" from any given

starting location. Lesions of the hippocampus disrupt the ability to learn this strategy (Packard and McGaugh, 1996). On a plus-maze task, Schmidt et al. (2013a) trained rats to perform concurrent locale (place) and a route (response) strategies.[2] On individual trials, an overhead light indicated if the next trial would require a place strategy: go "East." No light indicated a response strategy: turn "Right." More VTE was observed on place trials, especially when the rat was required to learn a reversal place strategy: go "West," not "East." (Schmidt et al., 2013a). However, no increase in VTE was observed when rats were required to learn a reversal response strategy: go "Left," not "Right."

Gardner et al. (2013) studied the effect of sequential locale (win-shift) or route (win-stay) choices in a two-choice task. The 'opposing Ts' task consisted of three arms joined together at the middle by an orthogonal arm. The geometry resembled the bottom of a 'W' joined to the top of an 'M.' VTE was found to occur more frequently during place strategies, in agreement with Schmidt et al. (2013a). Interestingly, the authors found that if a route choice was followed by a locale choice, the normally-observed transition from locale to route on the first choice was impaired. VTE occurred frequently during this period. This raises the possibility that habit system learning (thought to govern route learning) can be impaired by activation of the deliberative system (thought to govern locale learning, discussed in **Chapter** 2). From the results of Schmidt et al. (2013a) and Gardner et al. (2013), if VTE is involved in learning about unexpected environmental changes, its occurrence is more closely tied to those that involve locale navigation. This discussion raises an additional question.

3.4 Is VTE Conflict between Decision-Making Systems?

VTE may be a a result of transient conflict between navigation strategies, potentially mediated by different decision-making systems. The hypothesis that VTE represents a conflict between decision-making systems is, in a sense, rephrasing the question of **Section** 3.2 **Is VTE Value Integration?** where the 'value' integration step is now taking place at the level of motor output instead of revealed preference.

[2]Route navigation is defined as the sequencing of taxon (orienting) and praxic (complex response) strategies (O'Keefe and Nadel, 1978). If a rat is planning a trajectory through space, it can do so by chaining together these components.

While Schmidt et al. (2013a) found an overall increase in VTE during place strategies, an increase in VTE during error trials of all types was also found. Furthermore, an increase in VTE was observed when the potentially correct route strategy differed from the potentially correct locale strategy, putting the independent decision-making systems governing each strategy in conflict. Finally, errors potentiated increased probability of VTE on subsequent trials. This may indicate increased uncertainty about the learned strategy following an error trial, causing a 'mental check' on the subsequent trial (engaging an error-correction learning rule). This may also indicate a window of transient flexibility in responding, which might facilitate reversal learning or extinction of responding.

Differences in neurochemical signaling may bias the use of different decision making strategies. For example, blocking cholinergic receptors systemically biases rats to use praxic or taxon based navigation as opposed to locale navigation (Whishaw, 1985) and causes reduced (VTE-like) search behavior in a Morris water maze (Whishaw and Tomie, 1987). If VTE is a conflict between decision-making systems, pharmacological manipulations that impair specific decision-making strategies would be predicted to reduce conflicting signals generating VTE.

An additional possibility is that VTE represents conflict within a non-deliberative decision-making system. A rat who has run many hundreds of laps on a T-maze has cached-action responses for both left and right turns at the choice point. VTE could be a process of resolving those strategies without engaging a search process. On the other hand, the potentiation of VTE on error trials is consistent with deliberative search through known outcomes, but with a failure to find the correct outcome. The correlation between VTE and error trials could be a result of an inherent bias in the deliberative system toward novel or unfamiliar outcomes.

3.5 Is VTE Deliberative Search?

Johnson et al. (2012) propose that VTE occurs during times of maximum potential information gain. According to this theory, once an environment has become familiar, exploration can be directed to particular locations according to the expected information

to be gained at each location. This theory provides a mathematical framework to describe 1) investigating parameters in a known state space and 2) updating expectations about unpredictable environmental changes. The first type of uncertainty is dealt with by updating the expected information gained by investigating a different option. After repeated investigation of an option, the amount of information to be gained decreases, leading to either investigation of a different option or to a cessation of investigatory behavior. This, presumably leads to instantiation of a different decision-making system to exploit the information gained during exploration. The second type of uncertainty is dealt with by updating the prior on a Bayesian representation of the state space of an environment.

Spatial exploration can be contrasted with deliberation, a 'mental exploration' of imagined outcomes. By deliberating, an animal can imagine the outcomes of its decisions vicariously without actually experiencing them, a procedure providing great survival benefit. Deliberation is thought to involve the hippocampus, which uses episodic information to simulate potential future outcomes. VTE corresponds with representation of future locations in the hippocampus (Johnson and Redish, 2007). Thus, in some cases, VTE actually does represent 'vicarious trial and error.'

An important result of the study by Bett et al. (2012) is that VTE is driven by different structures depending on learning conditions. No changes in VTE were observed in hippocampal lesioned animals during a visual discrimination task, but changes related to performance were observed in a spatial serial reversal task.[3] While the behavior during pausing at a choice point and visual discrimination is indistinguishable, one process may be linked with a neural search mechanism while the other is not. This interpretation would favor restricting the definition of VTE to tasks on which evidence for neural search processes have been discovered. The 'VTE' during visual discrimination could be re-classified as a different process. This classification would lend irony to the original discovery of VTE by Muenzinger and Gentry (1931) during stimulus discrimination of auditory signals. While VTE turns out to be a mental simulation of potential outcomes during *other* behaviors, its discovery occurred during an experiment involving head movements based on optimal localization, which only superficially appeared

[3]The experiments were conducted within-subject, so these differences can not be explained by variability in the lesions.

to involve a mental search process.

Additionally, because VTE did not increase following hippocampal lesion, this argues against an interpretation of VTE being linked with a transient stage of learning before consolidation. If this were true, more VTE would be expected for a prolonged period of time. In human patients with hippocampal damage, more VTE would also be expected, but less VTE is seen (Voss et al., 2011). It should be noted that rats in the Bett et al. (2012) study were pre-trained on the serial spatial reversal task. In contrast, Clark et al. (2005) found that VTE remained persistently high in hippocampal lesioned animals during exploration of a novel environment.

In so far as VTE reveals an underlying process of neural deliberation, it is reasonable to ask why humans (who deliberate more than rats) are not moving their heads back and forth all the time. This riddle may be explained by the evolution of grasping hands and stereoscopic vision, and the decrease in information processing by olfactory and whisking sensory systems. When given a clear choice between observable stimuli, SFS sequences have similar properties to VTE. The link between evaluation and motor output, which in the rat is related to the motion of the head through space, may in primates have evolved to cortical structures involved with eye movements. However, when thinking about an abstract problem, people generally do not move their eyes back and forth either. Instead, different people recourse to different gestures, common examples being 'stroking your chin,' 'scratching your head,' or 'twirling your hair.' I hypothesize that these gestures are homologous to VTE behavior in rats.

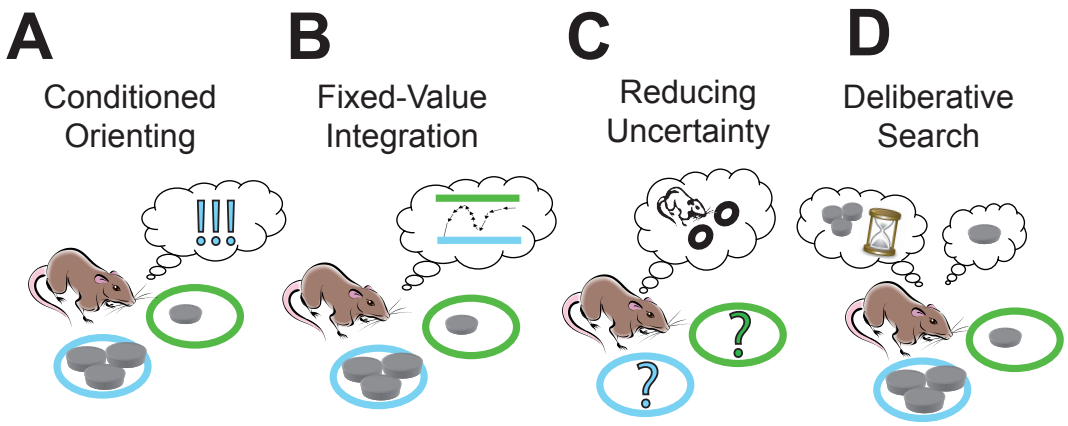

Figure 3.2: **Theories of Vicarious Trial and Error.** A) Is VTE Pavlovian Conditioned Orienting? The associative strength of predictive cues may cause orient-reorient behaviors at a choice point. If true, the probability of orienting in a particular direction can be described by a stationary variable, but this does not appear to be the case. In the diagram, the rat chooses to orient toward the larger reward because of its higher salience (represented by 3 exclamation marks). B) Is VTE Fixed-Value Integration? VTE may indicate that a neural race-to-threshold process is occurring. According to this model, looking in one direction causes value-coding activity associated with that choice (black arrows) is hypothesized to increase toward a threshold (green). When the rat looks in the other direction, value-coding activity associated with the opposite choice increases to its threshold (blue), competing with the first representation. The ensemble activity that reaches threshold first 'wins' and the rat selects that option. In this example, activity reaches the blue threshold first and the rat chooses the larger reward. If this model of VTE is true, then it would occur most frequently when the values of the two options are similar. C) Is VTE Reduction of Uncertainty? VTE may be an investigation process for learning the structure of an environment. In this cartoon, the rat does not know the value of either choice. Exploration of the two options would allow it to learn these values. If this account is true, VTE would be strongly associated with uncertainty, which can be manipulated experimentally. D) Is VTE Deliberative Search? VTE may be the behavioral marker for a deliberative search mechanism. Consideration of the potential future outcomes based on known parameters in the environment may occur during VTE. This hypothesis predicts that VTE occurs during difficult decisions.

Chapter 4

Hippocampal Neurophysiology

The hippocampus is a three-layer archicortex structure that is formed by two interlocking U-shaped sections, the *dentate gyrus* (DG) and the *cornu ammonis* (CA). It is located in the medial temporal lobe in humans, having migrated out phylogenetically from its medial position above the thalamus in rodents. The DG is a rare site of adult neurogenesis, with new granule cells forming functional axonal connections with DG interneurons and CA3 pyramidal cells (Toni et al., 2008). It is thought to mediate pattern separation via orthogonalization of CA3 schemas to subtle changes in context (Marr, 1971; McNaughton and Morris, 1987; Leutgeb et al., 2007; McHugh et al., 2007; Bakker et al., 2008).[1] The *cornu ammonis* is divided into three distinct parts: CA1, CA2, and CA3. The DG and CA3 receive a large fiber bundle input from entorhinal cortex (EC) layer II pyramidal cells. DG granule cells project to the dendrites of CA3 pyramidal cells. The CA3 pyramidal cells, in turn, project to the dendrites of CA1 pyramidal cells, which also receive input from EC layer III.[2] CA1 sends projections to EC layer V (Naber et al., 2001), medial prefrontal cortex (Jay et al., 1989; Jay and Witter, 1991), inhibitory projections to retrosplenial cortex (Jinno et al., 2007), projections to the septal nucleus (Risold and Swanson, 1997), and nucleus accumbens (Brog et al., 1993). Differences in projections exist along the septo-temporal (or dorso-ventral) axis. For example, ventral CA1 has reciprocal connections with amygdala (Pitkänen et al., 2000; Paxinos, 2014).

[1] The mossy fiber pathway from DG to CA3 is the only output of the DG (Paxinos, 2014).

[2] The projection from CA3 to CA1, the *Schaffer collateral*, contains axons from the majority of CA3 cells. CA3 also has a strong projection to the septal nucleus via the fornix, a back-projection to DG and, in the rat, commissural projections to contralateral *cornu ammonis* regions (Paxinos, 2014).

The major output of CA1, however, is the subiculum (Finch and Babb, 1981).

The subiculum, at the border between CA1 and EC, has similar electrophysiology and connectivity to CA1 (O'Mara et al., 2001). The subiculum receives input from EC layer III at its dendrites, dendritic and somatic excitation from CA1 pyramidal cells (Witter et al., 2000), and input from perirhinal and postrhinal cortices (Paxinos, 2014). In the rat, perirhinal cortex receives strong olfactory and auditory inputs and weak visual input from the direct pathway (Brown and Aggleton, 2001). It responds to the familiarity and recency of individual stimuli (Brown and Aggleton, 2001). Postrhinal cortex receives strong indirect and direct visual input from occipital and temporal cortices (Witter et al., 1989; Brown and Aggleton, 2001). It responds to object-place relationships, and has been implicated in attentional processing of the current context (Furtak et al., 2012).[3]

The fornix is a large fiber bundle connecting subcortical areas to the hippocampus including septal nucleus, mammillary nucleus, locus coeruleus, and supramammillary area of the hypothalamus. Additional cholinergic projections come from the nucleus of the diagonal band of Broca, serotonergic input arrives from the raphe nucleus, and dopaminergic input arrives from ventral tegmental area and substantia nigra. Inputs from these structures are distributed differentially to the different subfields of the hippocampus (Paxinos, 2014).

Owing to the unique pattern of inputs from EC, head direction cells, and their intrinsic connectivity, the CA1-CA3 regions have distinctive coding properties that intercombine to form schemas. CA1 presents a largely stable spatial map of an environment, while CA3 globally encodes subtle contextual changes. CA2 may either provide a consistent time-varying component on the scale of hours, or slow continuously-changing input to CA1, allowing dynamic regulation of the firing patterns in CA1 to different contexts (Mankin et al., 2015).

[3]The distinction between direct and indirect pathways of visual information is clearer in monkeys and humans.

4.1 Stability and dynamics of CA1 place cells

The components of hippocampal schemas are place cells. Pyramidal cells in the CA1 region of the hippocampus fire action potentials at specific places in an environment, called their place fields. Place cells have been observed in the hippocampus of rodents (O'Keefe and Dostrovsky, 1971), monkeys (Ono et al., 1993), bats (Ulanovsky and Moss, 2007) and humans (Ekstrom et al., 2003; Miller et al., 2013). Across species, 30-50% of pyramidal cells in CA1 are active in a given environment, with a large subset of those being place cells (Wilson and McNaughton, 1993; Ono et al., 1993; McHugh et al., 1996; Ulanovsky and Moss, 2007; Las and Ulanovsky, 2014).[4] In larger environments, a higher percentage of pyramidal cells form place fields, and those place fields are larger (Muller and Kubie, 1987; Fenton et al., 2008). A place cell typically has only one field, but can also have multiple fields, especially if an environment contains repeated elements (Singer et al., 2010; Cowen and Nitz, 2014).

While place cells fire action potentials in a novel environment on the first pass through their place field, a 'stable' spatial representation is achieved after approximately five minutes of exploration (Hill, 1978; Wilson and McNaughton, 1993; Frank et al., 2004). Upon entering a novel environment, interneuron activity initially decreases, allowing schemas to form rapidly (Wilson and McNaughton, 1993). While a new representation can form in the absence of NMDA activity, it requires global NMDA activity in order to remain stable beyond at least 1.5 hours (Kentros et al., 1998). In this study, Kentros et al. (1998) studied the effects of systemic injection of the NMDA receptor antagonist CPP on the formation and stability of place field schemas. They found that CPP did not affect previously-formed schemas and did not prevent the formation of schemas in novel environments. However, upon reintroduction to a novel environment, CPP caused a complete remapping. The original schema had not stabilized (Kentros et al., 1998). mGlu5 receptors (Zhang and Manahan-Vaughan, 2014) and protein synthesis (Agnihotri et al., 2004) are necessary for long-term stabilization of CA1 schemas.

Stability and maintenance of schemas may be related to ongoing spike activity in the dendritic arbor. Sheffield and Dombeck (2015) measured calcium transients in

[4]The remainder of place cell properties I describe apply to rats, unless otherwise specified, as this is where the breadth of research on *in vivo* hippocampal electrophysiology is greatest.

mouse CA1 using two-photon calcium imaging during navigation in a virtual environment. They found that about 50% of place field action potentials were accompanied by calcium spikes in the dendrites. The prevalence of dendritic calcium spikes predicted whether or not a place cell observed on the first day of recording would be present on the second day (Sheffield and Dombeck, 2015). This cellular-level effect may underlie the cell-assembly evolution that results in stable schemas in familiar environments (Karlsson and Frank, 2008). Environments that are similar have similar sets of active place cells (Leutgeb et al., 2005a). Remapping between different environments is a coordinated network phenomenon, and can not be explained by responses of independent cells (Wills et al., 2005). Place cell schemas of unchanging environments are stable over long periods of time (Thompson and Best, 1990), presumably over a lifetime.

In contrast, place cell schemas on dynamically changing environments are surprisingly flexible. New place cells in a familiar environment form during periods of head scanning (Monaco et al., 2014), behaviors associated with changes in attention or vigilance. Overdispersion (Fenton and Muller, 1998), or 'rate modulation' (Leutgeb et al., 2005b), accounts for the non-stochastic variability of firing rate during different passes through a place field.[5] Overdispersion may indicate attention to multiple reference frames in an environment (Fenton et al., 2010), switching between multiple active cell assemblies on a second-by-second basis (Jackson and Redish, 2007), or superimposition of variable non-spatial information on a framework of stable spatial representation (Colgin et al., 2008; Allen et al., 2012).[6] Partial remapping occurs when task contingency changes (Quirk et al., 1990; Markus et al., 1995; Tanila, 1999) and may be an intermediate state before the formation of a stable schema (Colgin et al., 2008). Stable CA1 schemas exhibit trajectory-dependent or context-dependent firing (Wood et al., 2000; Frank et al., 2000; Ferbinteanu and Shapiro, 2003; Pastalkova et al., 2008; Komorowski et al., 2009). During repeated traversals along a common path through an environment with changing goals, place cells will be active on one trajectory but not the other. For example, in a delayed alternation task, a place cell active on left laps may not be active during right laps (Wood et al., 2000).[7] During repeated traversals along a

[5]The term used by (Leutgeb et al., 2005b) is rate remapping, but I use the term rate modulation to emphasize the functional difference between remapping and overdispersion.

[6]It is unknown to what extent time cells (discussed in **Section** 5.1) exhibit overdispersion.

[7]This type of coding is discussed further in **Section** 5.1.

trajectory with constant goals, the center of mass of a place field shifts backward, away from the direction of travel (Markus et al., 1995; Mehta et al., 1997). Overdispersion, partial remapping, potentiation of new place fields, trajectory-dependent firing, and center of mass shifting are indications that the hippocampus is 1) actively engaged in spatial navigation and 2) dynamically updating its representation of an environment with experience.

4.2 Stability and dynamics of CA3 place cells

Place cells also exist in the CA3 region of the hippocampus. Place fields form more slowly in CA3, only after 20-30 minutes of exploration (Leutgeb et al., 2004) and tend to be more stable in non-changing environments once formed (Mizuseki et al., 2012). However, when variability is introduced in the environment, CA3 remapping occurs more readily than in CA1 (Leutgeb et al., 2004). These findings have led to the hypothesis that changing contexts challenge CA3 with a competition between pattern completion and pattern separation. CA3 NMDA receptors are crucial for both the pattern completion process (Nakazawa et al., 2002), as well as one-trial learning (Nakazawa et al., 2003). CA3 place fields may be larger or smaller than CA1 place fields, depending on the location along the transverse axis of CA3 (Lee et al., 2004). Relative to CA1, CA3 place fields have increased firing rate in the beginning of a place field (negative skweness), and represent the current location of a rat about one-half theta cycle earlier (Dragoi and Buzsaki, 2006). Like CA1, CA3 place cells exhibit overdispersion (Muller et al., 1987; Fenton and Muller, 1998), partial remapping, and trajectory-dependent firing (Ferbinteanu and Shapiro, 2003; Allen et al., 2012). During a rapid change in context, CA3 ensembles are transiently bi-stable, flickering from the old to the new context on alternate theta cycles (Jezek et al., 2011). This is consistent with the cell assembly switching account of overdispersion in CA1 (Jackson and Redish, 2007).

4.3 CA1 place cells are shaped by rhythmic inputs

CA1 place cells are shaped by the low frequency inhibitory and excitatory inputs to CA1. For example, as a mouse passes through a place field, the intracellular membrane potential ramps up asymmetrically to approximately 6mV, reaching this peak shifted at a distance a little more than halfway through the place field (Harvey et al., 2009). Local Field Potential (LFP) activity reflects the summed electrical potential at the recording site of an electrode. LFP in the hippocampus has been classified into two distinct states, theta and large-irregular activity (LIA) (Vanderwolf, 1969; O'Keefe and Nadel, 1978).[8] Theta activity accompanies movement, attentive pausing and rapid eye movement sleep (Vanderwolf, 1969; Vertes, 2011), while LIA occurs during inattentive pausing, grooming, and slow-wave sleep (Vanderwolf, 1969)

During locomotion and other attentive behavior, hippocampus has a prominent LFP network oscillation at a frequency of 6-12Hz called the theta oscillation (Vanderwolf, 1969). Theta is synchronized with exploratory movements (Hoffman et al., 2013), its instantaneous frequency is correlated with running speed (Whishaw and Vanderwolf, 1973; McFarland et al., 1975), and its power increases during periods of decision-making (Montgomery et al., 2009; Schmidt et al., 2013b; Belchior et al., 2014). GABAergic input from the medial septal nucleus onto CA1 interneurons is one source of theta (Buzsáki, 2002); however, the hippocampus is capable of spontaneous theta oscillation (Goutagny et al., 2009), and theta propagates from dorsal to ventral hippocampus as a traveling wave (Lubenov and Siapas, 2009). These observations suggest the existence of multiple theta oscillators at different stages of hippocampal processing (Patel et al., 2012). In a place field, intracellular theta power doubles, phase precessing relative to LFP theta (Harvey et al., 2009, phase precession is discussed in **Section** 4.5). When theta is disrupted through septal lesion (Winson, 1978), or pharmacological inactivation of septum with muscimol (Wang et al., 2014), deficits in spatial memory are produced. Theta is a global signal for organizing hippocampal activity across multiple timescales (Lisman and Redish, 2009; Colgin, 2013; Buzsáki and Moser, 2013).

Input from the medial septal nucleus may be necessary for stability and flexible

[8] In the early literature, theta was called *rhythmical slow activity*.

planning using CA1 ensembles. Inactivation of medial septum causes a dramatic reduction in theta power in the hippocampus (Mizumori et al., 1989) and EC (Koenig et al., 2011), and causes transient (20-60min) impairment in spatial working memory tasks (Mizumori et al., 1989; Wang et al., 2014). Septal inactivation causes about a 50% reduction in the firing rate of CA3 pyramidal cells, but has little effect on CA1 place cells in a familiar environment (Mizumori et al., 1989; Brandon et al., 2014; Wang et al., 2014). The ability to form new schemas under septal inactivation may be unimpaired (Brandon et al., 2014) or it may depend on the usefulness of landmarks (Wang et al., 2014), though schemas lack long-term stability under septal inactivation (Brandon et al., 2014; Wang et al., 2014).[9] These results suggest that pattern completion in CA3 is sufficient to maintain schemas in CA1 despite its reduced firing, but flexible planning and stability of CA1 schemas depend on global synchronizing input from septal theta.

CA1 receives asynchronous excitatory input from CA3 (low gamma) and medial EC (high gamma) on different phases of the theta cycle (Colgin et al., 2009). CA3 low gamma arrives via the Schaffer collaterals at the *Stratum radiatum* primarily on the elongated descending phase of theta. Medial EC high gamma arrives via the perforant path at the *Stratum lacunosum moleculare* during the ascending phase of theta (Schomburg et al., 2014). CA1 output is primarily entrained to the CA1-generated HFO, and is more strongly modulated by low rather than high frequency gamma. However, all frequencies strongly influence the firing of CA1 interneurons (Schomburg et al., 2014), which have a marked shaping influence on CA1 place fields (Royer et al., 2012), and theta sequences (Geisler et al., 2007; discussed in **Section** 4.5) suggesting that both gamma components are instrumental for normal CA1 function.

Input from medial EC and area CA3 are necessary for integrity of CA1 place cells, but each carries a different function. Medial EC lesions cause remapping upon reexposure to familiar environments (Van Cauter et al., 2008), suggesting a stabilizing function for schemas. Field size is reduced following medial EC lesion (Brun et al., 2008; Van Cauter et al., 2008, but see Ormond and McNaughton, 2015), but not a CA3 lesion (Brun et al., 2002). The study by Brun et al. (2002) revealed two distinct processing states for CA1, one of which was independent of CA3. Place fields in CA1 were remarkably similar after a CA3 lesion (see also (Mizumori et al., 1989)), and spatial

[9]Time cell assemblies (See **Section** 5.1) are also disrupted by septal inactivation (Wang et al., 2014).

recognition memory was preserved.[10] However, the ability to flexibly plan trajectories was impaired, as evidenced by increased escape latency during probe trials of an open Morris water maze. The behavioral result is supported by a recent study that tested the effect of CA3 versus medial EC lesion on a place (locale) or cue (praxic) strategy on a plus-maze (O'Reilly et al., 2014). Both lesions reduced correct choice using the locale strategy, with a larger and more persistent effect of CA3 lesion.[11]

The precise timing of synchronous 50-100 Hz high gamma in medial EC and CA1 is required for performance on a spatial working memory task (Yamamoto et al., 2014). Siegle and Wilson (2014) tested the effect of optogenetic inhibition of mouse CA1 at different theta phases on an H-shaped task that required retention of an informative cue across the 10cm central stem of the H. During navigation in the outer arms of the H, distal cues indicated the correct choice, permitting the mouse to use a taxon navigation strategy. Performance was improved when low gamma from CA3 was disrupted in the outer arms. In contrast, silencing high gamma from medial EC improved performance when locale navigation/working memory was guiding behavior on the central stem of the H (Siegle and Wilson, 2014). This implies that competition between EC (linked more strongly to taxon navigaton) and CA3 (linked more strongly to locale navigation) may interfere with decision-making.

4.4 Sharp Wave Ripple Sequences

Sharp wave ripples (SWRs) are generated when about 10% of CA3 pyramidal cells fire action potentials within a 100ms window (Ylinen et al., 1995; Csicsvari et al., 2000), creating a large excitatory post-synaptic potential (EPSP) in CA1 (the sharp wave) (Buzsáki et al., 1983; Suzuki and Smith, 1987). This generates 150-250Hz oscillations (the ripple) in the CA1 region, lasting 50-250ms and discharging up to 30% of its pyramidal cells (Csicsvari et al., 2000). Each activated CA1 neuron typically fires only one spike per SWR (Chrobak and Buzsáki, 1996), and a subthreshold EPSPs are generated

[10] Spatial recognition memory, a component of locale navigation, the ability to learn a location relative to a set of cues

[11] Surprisingly, O'Reilly et al. (2014) also found impairment in the praxic strategy from both lesions. The authors suggest this could be explained by the concurrent training regimen where place and cue responses were learned competitively.

from multiple CA1 pyramidal cells even at distant locations from the SWR locus (Maier et al., 2011). Ripples occur bilaterally in the hippocampus on 30-50% of SWR.[12] Although SWR are typically ($> 95\%$) a global phenomenon, they can occasionally be undetectable in adjacent longitudinal regions of the hippocampus (Chrobak and Buzsáki, 1996; Sullivan et al., 2011). Across the septal-temporal axis, SWR travel from their locus at a rate of 350cm/s (in either, or rarely, in both directions) with a spatial spread proportional to ripple magnitude (Patel et al., 2013).[13] SWR are increased 2-4 times after bilateral EC lesion (Bragin et al., 1995), suggesting that EC input acts to arrest SWR. Nakashiba et al. (2009) recorded CA1 ensembles in mice with a conditional knockout of CA3 glutamatergic output. CA3 knockout caused reduced firing of CA1 cells during SWR. The authors also reported that CA3 knockout caused a reduction in SWR frequency. However, an alternate interpretation is that it caused an increase in (100-150 Hz) HFOs, which have recently been reported in CA1 and may be dissociable from SWRs (Tort et al., 2013; Scheffer-Teixeira et al., 2013, but see Sullivan et al., 2011). Stark et al. (2014) found that local activation of a small number of pyramidal cells in CA1, in combination with fast $GABA_A$ inhibition, is necessary for SWR generation. Thus, although physiological ripples are preceded by a burst of CA3 action potentials, the mechanism for their generation exists within the CA1 network itself. During SWRs the hippocampal network is synchronized, ripples in CA1 and subiculum are coherent, and ripples occur in EC 5-30ms after a CA1 ripple (Chrobak and Buzsáki, 1996).

SWRs, which occur in monkeys (Skaggs et al., 2007) and humans (Bragin et al., 1999) as well as rodents, have been implicated in memory consolidation (Buzsáki, 1989; McGaugh, 2000). Buzsáki (1989) hypothesized that the SWRs serve as a training signal for long-term potentiation in a Hebbian network. This SWR consolidation hypothesis contrasts to the initially held assumption that the training signal is subserved by relatively non-plastic synapses reliably activating a target neuron. For example, in the model of McNaughton and Morris (1987), a subset of perforant path inputs from EC are hypothesized to act as a training signal to CA1. McNaughton et al. (1986) had

[12] While the SWR envelope is bilaterally coherent, individual cycles of the LFP ripple are not (Chrobak and Buzsáki, 1996).

[13] Patel et al. (2013) found ripples in the ventral hippocampus to be largely independent from those in dorsal to intermediate hippocampus. Additionally, ripple frequency, amplitude, and emission frequency decreased along the septal-temporal axis.

previously found that high-frequency stimulation of EC inputs to CA1 disrupted recently acquired spatial memories and impaired acquisition of new spatial memories. Subsequent research showed that medial EC input was necessary for stability of CA1 schemas (Van Cauter et al., 2008), supporting theories that medial EC acts as a sensory storage buffer for working memory (McClelland et al., 1995; Hasselmo et al., 2000; Lisman and Otmakhova, 2001) instead of a training signal for consolidation. In contrast, awake SWR disruption had little effect on schema formation and stability (Jadhav et al., 2012), but caused impaired learning with those schemas. The SWR consolidation hypothesis predicts that recently experienced events cause selective reactivation of CA1 place cells associated with those events.

The SWR consolidation hypothesis has received considerable support. Pavlides and Winson (1989) found that a single place cell in a recently-experienced location has increased firing during post-run SWR compared to one in a location that has not been visited recently. In fact, entire ensembles of place cells that were recently active in an environment before sleeping have increased likelihood of reactivation during post-run SWR (Wilson and McNaughton, 1994). In the 100-200ms SWR window, recently-activated place cells replay in ordered sequences, called *SWR sequences* (Skaggs and McNaughton, 1996). Experiences from multiple recently-experienced environments are interleaved in successive SWRs (Kudrimoti et al., 1999). Optogenetic stimulation of dopaminergic inputs to CA1 from ventral tegmental area increases SWR reactivation of previous firing patterns (McNamara et al., 2014), as would be expected for a dopaminergic goal or value encoding signal. Firing patterns during SWR are strongly linked to recently experienced events.

During slow wave sleep, SWR occur frequently, and hippocampal activity is correlated with cortical activity during these SWRs (Siapas and Wilson, 1998; Sirota et al., 2003), consistent with the hypothesis of information transfer between structures during consolidation.. Disruption of SWR during sleep slows learning on spatial tasks (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010), suggesting that consolidation and efficient locale navigation are linked. Girardeau et al. (2009) used a radial-arm maze memory task requiring the rat to remember which three of the eight arms were baited after three blocks of three minutes of exploration, with each block interspersed with three minutes of rest. In one group of this study, SWR were disrupted during the

rest phase by stimulation of the ventral hippocampal commissure, causing learning impairment. Ego-Stengel and Wilson (2010) used a task that could be solved with a praxic navigation strategy (a sequence of left and right turns) on a wheel-shaped maze. They used a two-day experimental protocol examining the effect of SWR disruption following one of two possible maze configurations. Learning was slowed only on the task followed by SWR disruption, suggesting immediate post-run rest was important for memory consolidation.

SWR sequences also occur during brief pauses in running (Diba and Buzsáki, 2007; Foster and Wilson, 2006; Carr et al., 2011). These nonlocal events over-represent goals and predict future trajectories toward goals (Pfeiffer and Foster, 2013), suggesting a role as a planning mechanism. However, they also represent infrequent or never-experienced trajectories (Gupta et al., 2010), can begin at remote locations (Davidson et al., 2009; Karlsson and Frank, 2009; Gupta et al., 2010, and others) and can reflect the time-reversed trajectory (Foster and Wilson, 2006). Awake SWR have more variability in their firing patterns than SWR during post-run rest, which replay forward trajectories from recent experience (Wikenheiser and Redish, 2013). To reconcile these results, SWR sequences during awake behavior may be viewed as groups of interchangeable components reflecting connected parts of an environment (Wu and Foster, 2014). Disruption of awake SWRs slows learning on spatial tasks (Jadhav et al., 2012), similar to the effects of disrupting SWR during sleep. Awake SWRs may be utilized for online planning during deliberation, or, consolidation may indirectly influence decision making.

4.5 Theta Sequences

Hippocampal schemas also have sequential firing during locomotion, organized by theta into subsecond patterns. While running along a maze, Muller and Kubie (1989) found that the best representation of a rat's current position was shifted forward by 120ms in CA1, approximately one theta cycle. They hypothesized that hippocampus may be involved in planning 'what might be' by calculating multiple possible trajectories across an environment dependent on variations in the shift of ensembles.

As a rat runs through a place field, the firing rate of that place cell precesses relative

to the LFP theta oscillation (O'Keefe and Recce, 1993). For a fixed extracellular theta frequency, action potentials at the beginning of a place field fire during late phases of theta (CA1 theta trough). Action potentials at the end of a place field fire during early phases of theta (CA1 theta peak). Phase precession is bimodal in CA1 (Yamaguchi et al., 2001), possibly reflecting the separation of inputs from medial EC and CA3 (see **Section** 4.3). On the other hand, CA1 and CA3 phase precession are tightly linked with a phase offset of 10° and spike timing separated by 5ms (Schmidt, 2010),[14] whereas EC layer III does not show phase precession, and other EC layers have phase precession independent of hippocampus (Hafting et al., 2008; Mizuseki et al., 2009). Therefore, CA1 phase precession may be inherited from CA3 (Schmidt, 2010). The intracellular theta phase of a place cell precesses relative to LFP theta, and spikes occur during intracellular theta peaks (Harvey et al., 2009). Phase precession occurs on the first pass through a place field in a novel environment (Schmidt et al., 2009; Feng et al., 2015) and the phase-position correlation increases within-session (Feng et al., 2015) and across-session (Cheng and Frank, 2008).

The representation of positions ahead of the rat observed by Muller and Kubie (1989) can be explained by the order-preserved sequential activation of place cells. Phase precession is coordinated across cell-assemblies in CA1, an observation called *theta sequences*. As the rat leaves one place field and enters another, spikes from the first field occur before spikes from the second field, all within one cycle of theta. In general, the order of place cell firing is preserved across theta cycles. This leads to order-preserved theta sequences (Skaggs et al., 1996; Dragoi and Buzsaki, 2006). Though the ordering is not strictly dependent on theta phase precession (Foster and Wilson, 2007), inactivation of septal theta causes the deterioration of sequential firing (Wang et al., 2014). Theta sequences begin on the second pass through a place field and continue through subsequent passes (Feng et al., 2015). There is an increase in phase-position correlation within-session from approximately the first five laps compared to subsequent laps (Feng et al., 2015) and in the first experience with a novel environment compared to subsequent visits to the same environment (Cheng and Frank, 2008).[15] The development of theta sequences can not be explained by development

[14] CA3 cells tended to fire before CA1 cells even if the CA1 place field was ahead of the CA3 field (Schmidt, 2010)

[15] The *sequence compression index*, a measure used by Cheng and Frank (2008), is the distance between

of stable schemas or lap-by-lap firing rate changes (Wang et al., 2014), suggesting a separate mechanism. Indeed, cannabinoid administration disrupts hippocampal theta sequences without changing schemas (Robbe et al., 2006; Robbe and Buzsaki, 2009). Sequences that extend farther ahead of the rat are accompanied by longer theta cycles and acceleration away from proximal landmarks (Gupta et al., 2012), and by goals that are farther ahead of the rat (Wikenheiser and Redish, 2015).[16] Cowen and Nitz (2014), while investigating a completely different question, speculate that the corners of tracks may serve as landmarks to 'reset' path-integration to an allocentric reference frame. Cei et al. 2014 and Maurer et al. 2014 tested the effects of reverse motion on theta sequences and found that they occurred in the direction of motion, despite the fact that the rat was oriented in the opposite direction (Cei et al., 2014) or walking backwards (Maurer et al., 2014).

Evidence is accumulating that theta sequences are necessary for episodic memory to impact decision-making. Theta sequences are abolished by administration of a systemic cannabinoid agonist, while schemas and theta oscillations are mostly unchanged (Robbe et al. 2006; Robbe and Buzsaki 2009, See **Chapter** 9).[17] Infusion of muscimol into the medial septal nucleus also disrupts theta sequences, though also decreasing theta power (Wang et al., 2014). Both pharmacological manipulations cause impaired choice on delayed spatial alternation tasks which place heavy demand on hippocampal function. Zugaro et al. (2004) reset the phase of theta and silenced CA1 pyramidal cells by stimulation of the ventral hippocampal commissure. The authors conclude that there was no effect on theta sequences, which appeared to resume as if nothing had happened following the disruption. However, theta sequences were assessed by visual inspection of the sequence compression index and additional evidence is required to reach this conclusion. This data argues for a role of hippocampus beyond the representation of current position by place cells.

Compelling evidence that theta sequences are part of a deliberative process would

peak firing rate of two overlapping place fields correlated with the timing of spikes in a theta cycle (Geisler et al., 2007).

[16]There is a complex relationship between look-ahead distance of theta sequences and running speed (Gupta et al., 2010).

[17]Cannabinoids also cause a reduction in gamma power and SWR power.

be to demonstrate that disruption of theta sequences, by stimulation of the hippocampal commissure (or other methods) during VTE, impairs serial forward representation. Schmidt (2010) emphasizes caution in studies of theta sequences, and suggests using correlated variability in phase precession across a cell assembly as an index of the 'noise' inherent to the representation. The utility of theta sequences for decision-making may be strongly linked to such measures of ensemble coordination. The fidelity of ensemble coordination may also be linked to ongoing consolidation.

Chapter 5

The Role of Hippocampus in Deliberation

5.1 Episodic Information

Episodic information is the what/when/where information that allows the construction of biographical memories (Tulving, 2002). For example, I can recall walking my dog (what) around the block (where) this morning (when). The hippocampus is necessary for initial formation of episodic information (Corkin, 1984; Squire and Alvarez, 1995), if it is removed, new episodic information can not be consolidated into long-term memory (Milner et al., 1998). The famous patient H.M., who had a nearly complete hippocampal lesion, if distracted, would immediately forget what he had previously been doing (Sidman et al., 1968).

To construct an episodic memory, the hippocampus encodes properties beyond the spatial domain. Three types of temporal encoding schemes have been described in hippocampal ensembles: 1) time or episode cells that fire during a specific time in an interval, 2) cells that phase precess throughout a delay, and 3) incremental timing cells that change their firing rate during a delay. Pastalkova et al. (2008) recorded ensembles of CA1 cells on a delayed alternation T-maze task where a rat ran on a running wheel for a specific duration between laps. When the authors looked at the firing of hippocampal pyramidal neurons across the 10s or 20s delay, they found that individual neurons fired at specific times during the delay. Similar to a place cell that fires

at a particular place in the environment, these 'episode cells' fired at a particular time in the delay. On a separate non-choice task, rats had to run in order to obtain a water reward, but did not have to hold any information in memory about which choice to make. During this task, correlations in temporal firing were observed during the first 2-3s, but did not persist throughout the 10s or 20s delay. However, some of these cells tended to phase precess throughout the delay. Subsequent research showed that episode or time cells develop with learning (Gill et al., 2011) and encode temporal gaps in nonspatial sequences between different sets of cues (MacDonald et al., 2011). The third type of temporal encoding scheme was observed in primate hippocampus. Naya and Suzuki (2011) and Sakon et al. (2014) found cells that increased their firing rate during context-specific intervals. Similar to a 'ramp cell' (van der Meer and Redish, 2011), these 'incremental timing cells' were observed more frequently during learning and correlated with correct performance.[1] These observations indicate that temporal representations are present in hippocampal ensembles and contribute to the formation of episodic memories.[2]

The role of hippocampus in trace conditioning may be explained by its temporal coding. It is well known that the magnitude of the effect of Pavlovian conditioning depends on the timing of the CS-US interval (Domjan, 1998), and hippocampus is involved in the learning of this precise temporal relationship (Olton, 1986). Pyramidal cells in CA1 respond to both the US and CS (Berger et al., 1983), thus encoding the duration of a trace conditioning interval (McEchron et al., 2003). Furthermore, there is learning-dependent rate modulation of the responses, and the firing rate of these cells decreases with learning (McEchron and Disterhoft, 1997). Modi et al. (2014) studied ensemble dynamics during trace conditioning, finding time cell assemblies that bridge the gap between CS and US, akin to the running wheel study of Pastalkova et al. (2008).

There is some evidence for encoding of the 'what' of episodic experience in the hippocampus. CA1 pyramidal cells will fire to the same object presented in a different location (Wood et al., 1999), during trajectory-specific choices (Wood et al., 2000) and to

[1] This type of pattern is similar to that in CA2 (Mankin et al., 2015), but over a timescale bounded by task-related events.

[2] It is not known whether time cells and place cells are integrated together in theta sequences or SWR sequences.

repeating elements in different places (Singer et al., 2010; Cei et al., 2014). Hippocampal lesions cause impaired recognition of item by position and item by context tasks (Mumby et al., 2002; Eacott and Norman, 2004). A study by McKenzie et al. (2014) comprehensively illustrates the multidimensional coding of hippocampal ensembles. McKenzie et al. (2014) used a task design where context was necessary to determine which of two objects was rewarded at a particular location. Two sets of novel objects were presented in intermixed trials where the position of the rewarded object was randomized. The authors used a Pearsons correlation coefficient analysis to ask whick dimensions explained the firing of hippocampal ensembles. Tuning curves of CA1 pyramidal cells were constructed to item, valence (rewarded or not), and set. They considered within-position, across-position, and across-context comparisons giving 16 item by position combinations. CA1 differentiated context most strongly, followed by the dimensions of position, valence, and item. For example, say object A is rewarded in context X at either position P or Q. Object A is always paired with object B, which is rewarded in context Y at either position R or S. CA1 ensembles would strongly differentiate X from Y (context) and PQRS (place). CA1 also differentiates rewarded from unrewarded object (valence by context) and A from B (item). Because item and valence were coded across-position, firing rate could not be interpreted as only encoding non-spatial properties (Leutgeb et al., 2005b). However, while conjunctive properties of objects (object by place, object by context) are coded in hippocampal ensembles, reconstruction of an episodic memory, especially the 'what' component, likely requires the contribution of other brain structures.

The clearest example of 'what' information comes from trajectory-dependent firing of place cells. Trajectory-dependent firing, depending on the future choice of the animal, may relate to the predictive role of hippocampus in deliberation. Ferbinteanu and Shapiro (2003) looked at retrospective, prospective, and trajectory-dependent coding properties in the hippocampus during repeated reversals on a plus maze task without return rails. In between trials, rats sat on a platform for 10-15s. Retrospective encoding was defined as spikes from place fields on the start arm occurring at the goal arm. Prospective encoding was defined as spikes from place fields on the (correct) goal arm firing at the start arm. The authors found a higher percentage of trajectory-dependent coding during correct trials, while retrospective coding was maintained even during

error trials. Although the choice point was excluded from analysis, the authors suggested that hippocampus was 'queried' prior to the choice point in order to facilitate decision-making, a conclusion consistent with a hippocampal role in deliberation.

5.2 Episodic Future Thinking

Patients with hippocampal damage also have difficulty imagining future events (Tulving, 1985; Klein et al., 2002; Hassabis et al., 2007; Race et al., 2011). For example, when asked, "What are your plans for the future?" a patient might respond with either absolute uncertainty or utter confabulation. The ability they seem to have lost, *episodic future thinking*, is the construction of potential future outcomes using episodic information (Atance and O'Neill, 2001).[3] There are examples of episodic future thinking in the non-human animal literature. For example, scrub jays [4] cache food that spoils at different rates and selectively return to caches that they determine to be still edible (Clayton et al., 2001). In humans, medial temporal lobe activation increases when imagining future events (Okuda et al., 2003), and hippocampal activation increases when navigating toward a goal (Maguire et al., 1998). There are plausible electrophysiological mechanisms for episodic future thinking to occur in the hippocampus, including SWR sequences (Pfeiffer and Foster, 2013; Foster and Wilson, 2006) and theta sequences (Johnson and Redish, 2007; Wikenheiser and Redish, 2015), discussed in **Section** 4.5.

Episodic future thinking may be a key component of goal-directed behavior, including the ability to delay gratification for a larger-later reward (Boyer, 2008). Cuing episodic future thinking decreases discounting (Peters and Büchel, 2010), causing subjects to choose a delayed option more frequently. Cuing episodic future thinking also results in increased correlation of BOLD activity between hippocampus and mPFC (Peters and Büchel, 2010). Increased functional coupling between medial prefrontal cortex and hippocampus has been shown to be important for correct performance on spatial working memory tasks (Jones and Wilson, 2005). Hippocampal lesions in rats increase

[3] The concept has also been called prospection, memory for the future (Ingvar, 1984), pre-experiencing (DArgembeau and Van der Linden, 2004), and mental time travel (Wheeler et al., 1997). See (Buckner and Carroll, 2007) for a review.

[4] A scrub jay is a corvid, a type of bird. This family includes jays, crows, and ravens (Emery and Clayton, 2004).

discounting when there is an inter-trial interval (Cheung and Cardinal, 2005; McHugh et al., 2008; Mariano et al., 2009; Abela and Chudasama, 2013), but do not consistently affect discounting under a free-operant procedure (Bett et al., 2014).

A fascinating distinction exists in humans between processing of linguistic and non-linguistic information. Medial temporal lobe damage may increase discounting (Palombo et al., 2014), causing subjects to choose an immediate option more frequently. However (Kwan et al., 2012) found evidence that intact semantic memory can facilitate delayed gratification even in the presence of medial temporal lobe damage and impaired episodic memory. This observation parallels patient H.M.'s ability to perform a delayed-match-to-sample task using semantic processing of trigrams (groups of three consonants, e.g. FQN), but his inability to perform the same task using visual stimuli (ellipses) (Sidman et al., 1968). Palombo et al. (2015) speculate that intertemporal choices in humans may be engaging semantic or episodic processing depending on task demands, but intertemporal choices in rats may only engage episodic processing. While common sense may argue that semantic processing in rats is nonexistent, the framework for linguistic processing may be present in non-human primates (Bornkessel-Schlesewsky et al., 2015), and proto-semantic processing may exist in the form of rats' social interactions (Jamain et al., 2008).

5.3 Deliberative Search in Hippocampus

The hippocampus contains cognitive maps (or schemas) of known environments, and machinery to generate new schemas for novel environments (O'Keefe and Nadel, 1978; Tse et al., 2007). Temporal difference reinforcement learning (TDRL) provides a mathematical framework for formalizing this concept in terms of a state space that satisfies the Markov property (Sutton and Barto, 1998).[5] A reinforcement learning algorithm can be used to search through potential future outcomes, a so-called 'tree search' (Daw et al., 2005). If a model-based TDRL framework is utilized as a search through or within hippocampal schemas, the hippocampus may be carrying out cognitive search. This is a key component of the theory of deliberation (van der Meer et al., 2012).

[5]The Markov property is satisfied if past states are independent of future states.

The role of hippocampus in deliberation is the generation of potential future outcomes. Outcomes of choices are simulated serially (one branch at a time), with the current position serving as the root of the tree, and the nodes as the places of potential outcomes, or goal locations. In the one-dimensional case, the nodes would be goal positions at varying distances in front of the rat. In a T-maze task, the nodes would split into a binary search tree, requiring search along both branches toward each goal.

Hippocampal search on T-maze tasks has been linked with VTE behavior. VTE cooccurs with the theta rhythm, and with serial forward representations of position during learning (Johnson and Redish, 2007). These observations led to the prediction that forward theta sequences during VTE are responsible for the forward representation (Gupta et al., 2010), though the specific hypothesis has not been examined. The claim is supported by evidence that theta sequences chunk the environment into behaviorally relevant sections (Gupta et al., 2010) and predict the upcoming choice (Wikenheiser and Redish, 2015). Ferbinteanu and Shapiro (2003) found diminished prospective encoding during error trials on a multiple-reversal plus maze task.[6] In humans navigating virtual tracks, CA1 activity is related to correct performance and reaction time at a high cost choice point (Brown et al., 2014). These are properties expected from a planning mechanism that is representing potential future outcomes as nodes on a search tree.

[6]The authors did not examine VTE on the task.

Chapter 6

The Spatial Delay Discounting Task

To investigate the role of hippocampus during deliberation, and the relationship of deliberation to VTE, a task with complex choices was developed. In this chapter, I describe the rationale for development of this novel neuroeconomic task, the *spatial delay discounting task* and the behavioral procedures used to gather and analyze data using it.

6.1 Delay Discounting

Delay discounting describes the observation that animals prefer an immediately available reward over one that is delayed in time (Madden and Johnson 2010, Figure 6.1). An extensive delay discounting literature exists spanning the fields of economics, behavioral psychology, and decision neuroscience. Delay discounting appears to be the rule across species and has been observed in pigeons (Green et al., 2010), rats (Green et al., 2004), monkeys (Freeman et al., 2012), and human subjects. Successful theories of the behavior have led to its description using simple equations with time-stable parameters. The preference for immediate reward has clear evolutionary advantages, especially if resources are scarce (Stevens and Stephens, 2010).

While impulsive choice is the rule, it is also advantageous to flexibly adapt to situations where delaying gratification is beneficial. Mischel and Mischel (1983) studied delayed gratification in a group of 94 three to eleven year old children. Children were shown three marshmallows on a table in front of them. The experimenter explained

that they had to leave the room for an indefinite amount of time, but when they returned the child could have two marshmallows. However, the child was also given a bell to ring that could bring the experimenter back anytime. The experimenter explained that if the child used the bell to bring him back, then they would only get to eat one marshmallow. Children could choose to cover the marshmallows or leave them uncovered, creating a tempting environment. Older children were more successful in choosing to cover the rewards and more frequently received the two marshmallows. This ability to delay gratification has been linked to better SAT scores and an increased ability to cope with frustration and stress (Mischel et al., 1989). Its absence is a feature of substance abuse (Yi et al., 2010) and pathological gambling (Petry and Madden, 2010).

Like the marshmallow experiment, delay-discounting experiments measure choices made between taking a smaller reward sooner or waiting for a larger reward later. Several methods can be used for evaluation of discounting behavior. In humans, subjects are typically presented with a list of questions of the form, "Would you like $10 now or $100 in one week?" Nonhuman animals can be trained to use a variety of procedures to reveal their choice preferences, including the adjusting delay procedure.

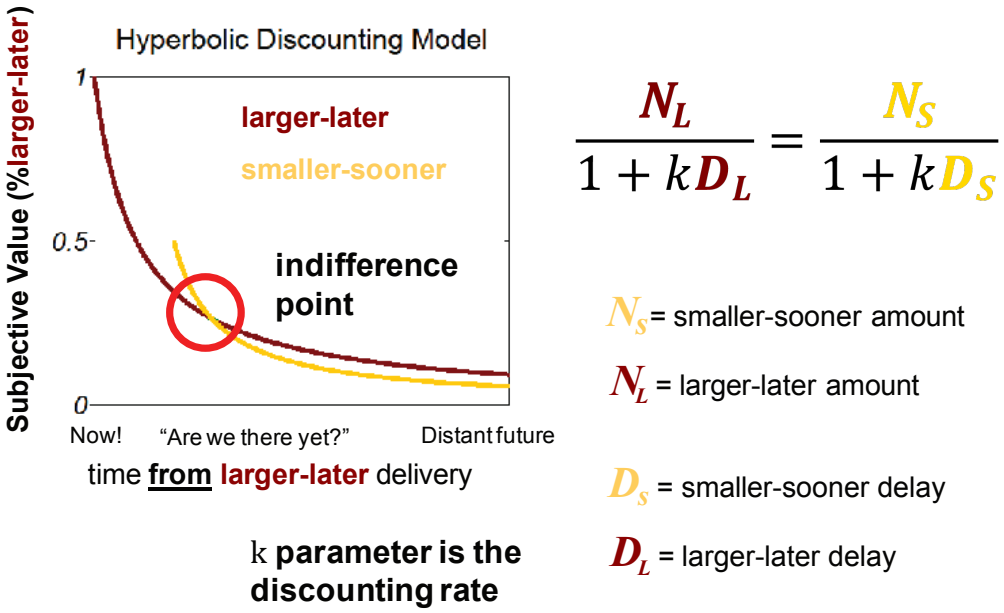

Figure 6.1: **Delay Discounting.** Delay discounting is the decrease in value of rewards that are time-delayed. It is related to both cognitive and socioeconomic variables in humans and is also observed in most non-human animals. Delay discounting experiments with human subjects typically consist of sequences of binary choices between two rewards, "Would you like $10 now or $100 in the distant future?". In this example, '$10 now' is the 'smaller- sooner' reward (gold curve) and $100 in the distant future is the 'larger- later' reward (maroon curve). Plotting the choice preference versus time from the larger-later reward delivery produces a psychometric curve for each option. Above is plotted the value of each reward as a percentage of the larger-later reward, to give a representation of the relative preference of the two options. If the larger-later reward is delivered now, it is strongly preferred over the smaller-sooner reward. If however, the larger-later reward is delivered in the distant future, the smaller-sooner reward is preferred. The point on the curve at which rewards are preferred equally is called the indifference point.

6.1.1 Adjusting Delay Procedure

The adjusting delay procedure is a common method used to study delay discounting in non-human animals (Mazur, 1997, 2001; Madden and Johnson, 2010). Typically, animals in an operant chamber are given one choice per trial, with the two options indicated by separate levers or nosepoke holes. Selecting one option gives a smaller reward after a shorter delay (smaller-sooner) and selection of the other option gives a

larger reward after a longer delay (larger-later). The larger-later delay is not constant, but is adjusted trial-by-trial (or between trial blocks) based on the choices of the animal. The choice of the larger-later option causes an increase in the larger-later delay on subsequent choice trials, while the choice of the smaller-sooner option causes a decrease in the larger-later delay on subsequent trials. Adjusting delay procedures are often presented in compound blocks (Mazur, 1997; Cardinal et al., 2002; Simon et al., 2010), with initial forced-choice trials to each option followed by free-choice trials where the delay is adjusted. Theoretically, animals should continue to adjust the delay during free-choice trials until the values of the two options are matched. Once matched, they should choose equally between the two options. However, on the lever-based versions of the task rats do not do this consistently, but show large variability in responding (Cardinal et al., 2002; Valencia Torres et al., 2011; Peterson et al., 2014). This behavior may be incompatible with evaluation of the two options.

Rats spontaneously alternate on spatial binary choice tasks (Wingfield and Dennis, 1934; Dember and Fowler, 1958). We hypothesized that the spontaneous alternation behavior of rats would allow more stable asymptotic performance on the adjusting delay procedure.

6.2 Behavioral Predictions

Given that rats tend to spontaneously alternate between sides on a spatial binary choice task, we hypothesized that a T-maze version of the adjusting delay procedure, the *spatial delay discounting task*, would produce discounting behavior with stable choice preferences.

A hyperbolic discounting model predicts the existence of an *indifference point*, a stable-state where the value of the larger-later option and the value of the smaller-sooner option are equivalent (Figure 6.1). According to the model, a reward N is discounted in time as $N/(1+kD)$ where D is time and k is a discounting parameter. Considering the discounted value of two rewards N_1 and N_2, that only exist after delays D_1 and D_2, respectively, the indifference point occurs when:

$$\frac{N_1}{1+kD_1} = \frac{N_2}{1+kD_2} \tag{6.1}$$

Conceptually, an animal making decisions based on discounted value should begin this task by investigating the initial value of each option (Which option is delayed? How long are the initial delays? How large are the rewards?). This *Investigation* phase would consist of a few alternation trials, repeated until the values are committed to memory. Then, the animal would adjust the delay until the discounted values of the two options are matched at the indifference point. This *Titration* phase would place high demand on the use of flexible decision-making or working memory. Finally, when the indifference point is reached, the animal would alternate repeatedly between the two options. In this *Exploitation* phase, the animal would take advantage of the favorable trade-off between the larger-later and smaller-sooner options and switch to an alternation strategy that would not require flexibility. The Exploitation phase is so named to draw a parallel to the exploration-exploitation trade-off problem in reinforcement learning theory (Cohen et al. 2007; Figure 6.2A,B).

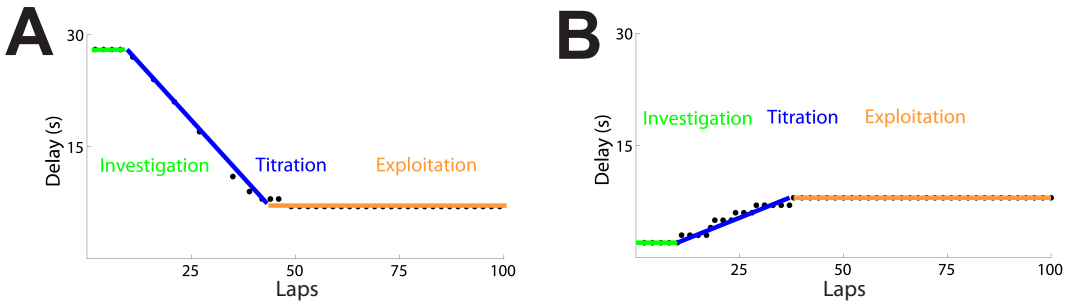

Figure 6.2: **Behavioral Predictions.** If rats make binary choices according to the discounted value of the two options, then their behavior should follow a predictable pattern. First, the rat would investigate the unknown parameters of the day to determine its strategy (Investigation, green line). Then, the rat would make repeated win-stay trials to one side, changing the adjusting delay to its preferred waiting time (Titration, blue line). Once this indifference point is reached, rats will alternate (Exploitation, orange line).

6.3 Experimental Design

6.3.1 Task Design

A trial consisted of a complete journey from the bottom of the T-maze stem through the choice point and to the feeder location without turning around and running backwards

or leaving the designated feeder sites before food was delivered (Figure 6.3A, black arrows). 100 trials (or 60min) were allotted per session and 30 sessions per rat were run with one session per day at the same time each day $\pm$2hr.

If the rat chose the larger-later side, a delay was experienced before food was delivered (Figure 6.3A, white circle). The delay was adjustable based on the choices of the rat. If the rat chose the larger-later side, the delay would be increased by 1s on the next lap (Figure 6.3B, magenta arrows). For example, if the delay was 10s, choosing the larger-later side on the next lap would result in a delay of 11s. If the rat chose the smaller-sooner side, the delay on the larger-later side would be decreased by 1s on the next lap. The delay on the smaller-sooner side was fixed at 1s. Alternation resulted in a constant delay (Figure 6.3B, gray arrows). For each rat, a pseudorandom uniform distribution of starting delays from 1s to 30s was used, with the larger-later side counterbalanced across days. Even (odd) starting delays always had the larger-later option on the left (right) side. Runs of similar delays for one side of the T-maze on subsequent days were avoided such that the moving average remained at roughly 15s.

Rewards were 45mg plain food pellets (TestDiet, Richmond, IN) delivered from automated food dispensers (Med-Associates, St. Albans, VT) through Nalgene tubes and were allowed to fall onto the track. Food rewards of one 45mg pellet for the smaller-sooner side or three 45mg pellets for the larger-later side were given on each trial at feeder sites located at the end of each T-arm.

A countdown sequence of pure tones of descending frequency accompanied the delay as an auditory cue. If the rat left the feeder zone (intentionally or not), the tone sequence stopped, indicating to the rat that he would not receive food on that lap. 1s to food and 0s to food on both sides of the maze were always accompanied by 1175 Hz and 1000 Hz tones, of duration 100ms. Higher delays were accompanied by tones of 175 Hz increase in frequency per one-second increase in delay, all of duration 100ms.

Position was tracked by an overhead camera positioned in the center of the recording environment at a resolution of 0.17cm/pixel and a speed of 60 Hz. For each timestamp, the center of mass of pixels with luminance that crossed a user-defined threshold was recorded (Cheetah, Tuscon, AZ). After 30 sessions of behavioral training, rats were implanted with multi-electrode neural recording drives and then repeated the 30 session experiment. Light emitting diodes (LEDs) on the neural recording drives were

positioned over the head of the rat, allowing precise tracking of head position (during behavioral sessions, an LED on a backpack strapped across the rat's midsection was used for tracking). Real-time position data was used to track the position of the animal for automation of the task (Matlab, The Mathworks, Natick, MA).

6.3.2 Apparatus

The track was 137.16cm $\times$ 76.2cm with an elevation of 17.78cm above the ground. The experiment was conducted in a room with consistent distal cues. The track was placed in the center of the room and was aligned each day with tape marks on the floor. The room was completely encased by a copper mesh Faraday cage. Before and after each session, the rat waited for 10-30min on an elevated flower pot having full view of the room.

During recording sessions, the headstage of the rat was attached via shielded cable to a commutator (Neuralynx, Tucson, AZ; Dragonfly, Ridgeley, WV; AirFlyte, Bayonne, NJ). Elastic tied to the commutator arm and various locations along the cable length was used to reduce the weight carried by the rat.

6.3.3 Subjects

$N = 11$ adult male Fisher 344 Brown Norway rats (Harlan, Indianapolis, IN) were used in this experiment, aged 8-12 months at the start of behavioral training. Animals were food-restricted to no less than 80% of their free-feeding body weight and water was available *ad libitum* throughout the experiment. Animals were individually housed, and were entrained to a 12 hour light/dark schedule. All procedures were conducted in full compliance with National Institute of Health guidelines for animal care and approved by the Institutional Animal Care and Use Committee at the University of Minnesota.

6.3.4 Training

Prior to beginning the first 30 days of free-choice sessions on the spatial delay discounting task, rats had approximately two weeks of training where they learned how to get food on the track by navigating the maze from the start of the T-maze stem to the feeder

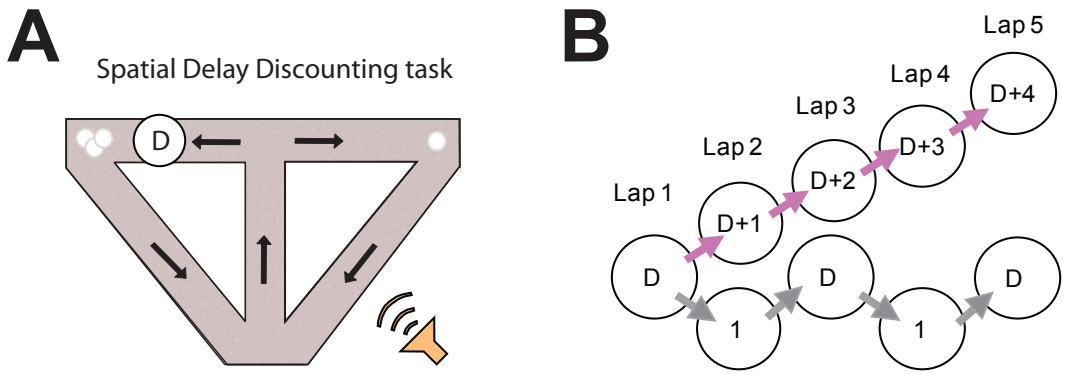

Figure 6.3: **Spatial Delay Discounting Task Diagram.** A) The spatial delay discounting task is a binary choice T-maze with return rails. A trial (black arrows) consists of a journey from the start of maze (bottom of track) up the central stem of the T to the choice point, a choice of either the left or right feeder, and then a return to the start of maze. In this example, a choice to the right feeder gives one pellet after a 1s delay and a choice to the left feeder gives three pellets after a delay of D seconds. The delay is accompanied by a tone countdown with each delay matched to a unique pitch. B) The delay D is adjusted based on the choices of the rat. Each circle indicates the delay experienced by the rat that lap. In this example, the rat has chosen the delayed side on lap one. Repeated choices to delayed side (*Titration*, magenta arrows), cause an increase in the adjusting delay by 1s for laps two through five. Choice of the non-delayed side on lap two results in a delay of 1s, followed by choice of the delayed side on lap three (*Alternation*,gray arrows) keeps the delay D at a constant interval.

arms. Each day of behavioral training, one side of the track was blocked (alternating left or right side each day), and the rats had to run to the non-blocked side to receive one 45mg pellet per trial. Attempts to run backwards were physically blocked by the experimenter. Rats quickly learned how to obtain food and completed training when they ran 100 trials per session for two days in a row.

6.3.5 Behavioral Analyses

LEDs on the headstage of implanted rats allowed precise tracking of head position over the 30 day recording sequence. During these sessions, VTE was quantified at the choice point. Position samples starting halfway up the central stem of the T-maze and ending before entry into a reward zone defined the choice point. Tone cues only occurred after a rat had left the choice point and entered one of the feeder zones, and if a rat turned around after entering a feeder zone that trial was not included for analysis. The x and y components of velocity were computed using an adaptive filtering of best-fit velocity vectors (Janabi-Sharifi et al., 2000). Then, the orientation of motion *Phi* was calculated by taking the arctangent, and unwrapped to prevent circular transitions. The change in orientation *dPhi* was then calculated using the same adaptive algorithm (Janabi-Sharifi et al., 2000). The absolute value of the change in orientation was integrated for each pass through the choice point to calculate *IdPhi* (Figure 3.1B,C). This measure was then z-scored within-session or log-transformed and z-scored within-rat to calculate *zIdPhi*. Turn around times for VTE events were identified manually.

Individual laps were classified as either Adjustment or Alternation. A win-stay strategy was classified as an *Adjustment* lap, since a repeated lap to the same side caused a change in the adjusting delay. For example, if lap $L-1$ was to the left, and lap L was also to the left, then lap L was classified as an Adjustment lap. A win-shift strategy is equivalent to an *Alternation* lap, where for example, lap $L-1$ is to the left and lap L is to the right.

A sliding-window was used to 1) classify laps according to their behavioral phase, and 2) determine the occurrence of VTE for different relative proportions of Adjustment laps. To classify each lap as Investigation, Titration, or Alternation, a five-lap sliding window was used. Titration laps were groups of five laps with at least two Adjustment laps. Investigation laps were groups of five laps that had fewer than two

adjustment laps and occurred before lap 30. Exploitation laps were groups of five laps that had fewer than two adjustment laps and occurred after lap 30. A variable sliding window from 1-10 laps was used to compare zIdPhi values for Alternation and Adjustment laps occurring at different rates.

To quantify across-session learning on the spatial delay discounting task, the number of alternation laps within $\pm$3s of the mean adjusting delay over the number of total alternation laps was computed. This measure of *alternation efficiency* quantified the precision of titration to a consistent indifference point within a session.

6.3.6 Electrophysiology

$N = 6$ rats were implanted with a neural recording apparatus (Kopf Neuro-Hyperdrive, David Kopf Instruments, Tujunga, CA) having 14 independently movable electrodes. 12/14 electrodes were tetrodes used for data collection, one electrode was placed in the hippocampal fissure and used as a theta reference, and the remaining electrode was placed in *corpus callosum* or superficial cortex and used as a reference. Rats were implanted over right dorsal hippocampus, with two at coordinates (AP= $-$3.8mm, ML= 3.5mm) and four at (AP= $-$3.8mm, ML= 3.0mm; Paxinos and Watson 1998). Over 10-14 days post-surgery, the tetrodes were slowly lowered toward CA1, at a rate inversely proportional to the distance from CA1. The distance to CA1 was monitored by observing the deflection of sharp wave ripples across tetrodes. As sharp waves (2-4Hz) flattened and ripple amplitude ($150 - 250$Hz) increased, tetrodes were gauged to be approaching the CA1 pyramidal layer (Ylinen et al., 1995).

After surgery, rats were allowed 1-3 days of recovery time before repeating the training with blocks procedure. After rats were running reliably, their recording apparatus (20-30g) was plugged into the headstage (10g), which was attached to a commutator arm by shielded cables. After rats were running reliably with this additional weight, typically 10-15 days post-surgery, the 30 days of the delay discounting experiment were repeated.

Spike waveforms were captured at a rate of 32 kHz in a 1ms window when voltage crossed a user-defined threshold on any of the four tetrode channels. Spikes were bandpass filtered (600-9,000Hz) and digitized at 32 bits (Neuralynx Cheetah, Tuscon, AZ). Local Field Potential (LFP) voltage from one channel per tetrode was sampled at

a rate of 2kHz and bandpass filtered (1-475 Hz). Off-line spike sorting was conducted using an automatic k-means clustering algorithm (KlustaKwik, Kadir et al. 2014) followed by manual spike sorting using various features of the data (MClust 3.5, Redish, et al., 2015). Individual units were manually isolated based on differences in amplitude, energy, etc. among the four tetrodes, and verified by measures of cluster quality (Schmitzer-Torbert et al., 2005) and stability across time. 1833 cells with an average firing rate of less than 10Hz were classified as CA1 pyramidal cells. An average of 12.5 cells/day were recorded over 146 days. Cell totals for each rat were (17, 209, 58, 408, 553, 558), though it is likely that some cells were recorded across multiple days.

6.3.7 Neural Analyses

Sharp Wave Ripple Detection To compute SWRs, a Hilbert transform was performed on the tetrode with the most cells for each session after bandpass filtering from $150 - 250$Hz. The amplitude obtained from the Hilbert transform was z-scored and global extrema were computed from the timeseries. Extrema that were greater than 2.5 standard deviations above the mean and greater than 250ms apart were considered SWR events. A fixed $\pm$ 125ms window around each SWR event was used for all analyses.

Bayesian decoding To analyze the content of SWR sequences and compare their information content to that in theta sequences at the choice point, a Bayesian decoding algorithm was used (Zhang et al., 1998). Spikes from pyramidal cells in CA1 were first binned into 50ms windows. Tuning curve models were derived from the occupancy-normalized activity of cells in a 56×56 spatial grid. The decoded representation was computed for each 50ms time bin, creating a probability distribution across the 56×56 space for each time bin. Analysis zones were defined by four 10×11 bin rectangles centered on the choice point, each feeder location, and the start of maze.

Bayesian decoding during sharp wave ripples To decode during SWRs, two mutually exclusive sets of spikes were used: the training set consisted of all spikes that occurred during theta epochs and **not** during a $\pm$125ms window around detected SWRs. The testing set consisted of those spikes **not** in theta epochs and all spikes in the $\pm$125ms window around detected SWRs. The maximum probability in each analysis

zone was found for each sample in a $\pm$125ms window around SWR events. For SWR events, these values were averaged across the five samples in each SWR. The corrected log-odds ratio of decoding during SWRs was computed for the delay side over the non-delay side.

Bayesian Decoding during theta sequences To decode during choice point passes, **all** spikes were used for the testing set, but only spikes in theta epochs and **not** during a $\pm$125ms window around detected SWRs were used for the training set. The maximum probability in each analysis zone was found for each position sample during a choice point pass. The corrected log-odds ratio of decoding during SWRs was computed for the delay side over the non-delay side.

Correlation of Decoding during SWR and theta sequences For comparison of decoding during SWR to decoding through a choice point pass, averaging was carried out for all SWRs within-lap and over all position samples through a choice point pass on the same lap.

Local Field Potential Analysis Local field potentials (LFPs) were analyzed using multi-taper spectrograms with a window size of 0.5s and no overlap between windows. For **Section** 8.6.1 and 8.6.2 the spectral power of all tetrodes in CA1 within each frequency band was computed according to this procedure: 1) power was averaged across all tetrodes in CA1, 2) the power was summed within each frequency band and then z-scored, 3) the average z-scored power was computed for each bin in a 56×56 grid. Filters for analysis of LFP components were set at $30 - 55$Hz for low gamma, $56 - 100$Hz for high gamma, $100 - 150$Hz for HFOs, and $150 - 250$Hz for SWRs. Positions on the grid were divided into four analysis zones at the choice point, delay side feeder, non-delay side feeder and start of maze. These zones were identical to those used for decoding analysis. For **Section** 8.6.3: 1) LFP power was z-scored within-frequency across all tetrodes in CA1 for each session, 2) z-scored power was averaged across-session, 3) Averaged z-scored power was again z-scored within-frequency. For **Section** 8.6.4: The trial-averaged spectrogram of LFP power was computed for each session across all tetrodes in CA1, 2) z-scored within-frequency across all tetrodes in

CA1 for each session, 3) binned by speed and averaged within-session, 4) averaged across-session.

6.4 Experiment Two: Cannabinoids

$N = 2$ adult male Fisher 344 Brown Norway rats (Harlan, Indianapolis, IN) were used for this experiment. All procedures were conducted in full compliance with National Institute of Health guidelines for animal care and approved by the Institutional Animal Care and Use Committee at the University of Minnesota.

The apparatus was a scaled-up version of the spatial delay discounting task and was run in the same room as the primary experiment described in **Section** 6.3.1. One rat had 14 days of training with blocks on the DUPLO task, described in (Papale et al., 2012) and (Breton et al., 2015). Then, he had 28 days of behavior on an alternate version of the spatial delay discounting task.[1] His behavioral data is included in Breton et al. (2015). After this experiment, both rats had behavioral training identical to that described in **Section** 6.3.4.

Following an implant into right dorsal CA1 and a brief three day recovery period, training resumed on the scaled-up version of the spatial delay discounting task. When rats were running reliably, a 10 day recording procedure was initiated, consisting of eight sessions with three pellets on the larger-later side, and two sessions with one pellet delivered on both sides. Initial delays were either 1s or 20s. The sessions with one pellet on both sides had starting delays of 20s. Delays were counterbalanced by side and drug/vehicle condition.

Intraperitoneal injection of the synthetic nonselective cannabinoid agonist (-)-CP55940 (Sigma-Aldrich, St. Louis, MO) at a dose of 0.05mg / 1kg were given 20-60 minutes prior to the experiment. The drug was prepared in a stock solution of 10mg drug dissolved into 3.4mL of 200 proof ethanol, and then aliquated into separate Eppendorf tube samples containing 25μL and stored in a -20° freezer. Individual samples were allowed to cool to room temperature before adding 25μL Cremophor EL to the Eppendorf tube. Saline was then added to the Eppendorf tube and mixed with the solution

[1] This rat experienced a *closed economy* version of the spatial delay discounting task. The amount of food he received each day was fixed at 9g, and the number of laps varied from day to day.

using a syringe (volume saline = weight of rat $\times$ (0.1mL vehicle / 100g rat weight) - 0.25μL solution - 0.25μL Cremophor EL + 0.2mL). A half dose from each sample was delivered by an experimenter blind to condition.

Chapter 7

Behavioral Results

7.1 Titration to consistent indifference point

Behavior on the spatial delay discounting task was consistent with predictions of the discounting model outlined in **Section** 6.2, suggesting that rats were basing their decisions on the value of outcomes. Rats titrated the adjusting-delay to an indifference point in the range of 3-11s, which remained stable 1) throughout the course of the Exploitation phase within-session, and 2) across-session (Figure 7.1A). For simplicity, the mean adjusting delay over the final 20 laps was used as an operational correlate of the indifference point, where the subjective value of the two options is equivalent. Across all sessions, the delay tended to converge to a stable value in the final 20 laps.

7.2 Adjustment and Alternation laps

Rats tended to adjust the delay during the first third of a session. The percentage of Adjustment laps peaked at 33% between laps 10-15 and decreased thereafter. The percentage of Alternation laps hovered between 60-90%, increasing steadily from its minimum at laps 10-15 until asymptoting at around 90% between laps 70-75 (Figure 7.1B). The distribution of Adjustment laps suggested division of the session into three behavioral phases, consistent with the predictions of the discounting model.

7.3 Behavioral Phase

When grouping laps by phase, the distribution of Titration laps that closely resembled the distribution of Adjustment laps, peaking around 33% between laps 10-20. Investigation laps occurred early and decreased with lap. Exploitation laps increased with lap and asymptoted at around 90% by lap 70 (Figure 7.1C).

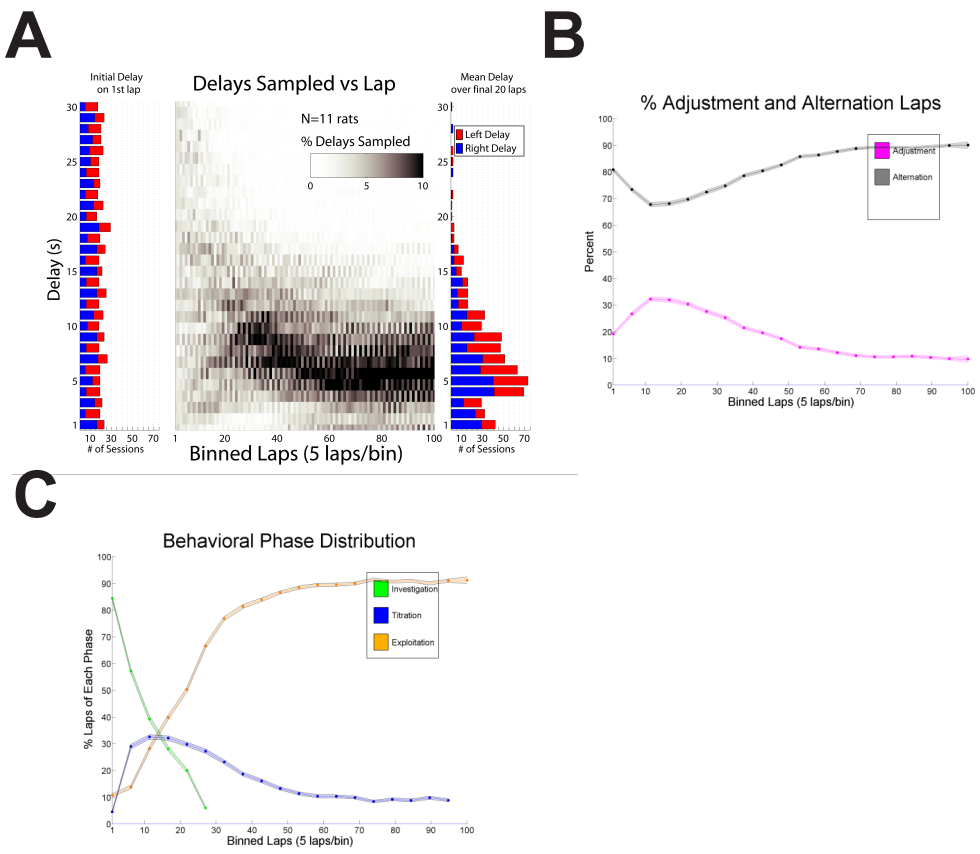

Figure 7.1: **Titration to indifference point.** Rats titrated the adjusting delay from a uniform distribution of initial delays to a final delay that was consistently around 7s. A) Initial delays on the spatial delay discounting task were uniformly distributed from 1s to 30s, counterbalanced left-to-right so the average delay on each side was approximately 15s (left panel). From this uniform distribution, rats titrated the adjusting-delay either upward or downward until it converged on the indifference point (center panel). Shading shows the percentage of delays sampled during each lap bin, normalized within-lap bin. The mean adjusting delay over the final 20 laps is consistent across-session, having a distribution with a mean ($\pm$ standard deviation) of 7.5 ± 4.8s. B) The percentage of Alternation and Adjustment laps at each lap bin shows that Adjustment laps tend to occur during the first third of the session, peaking after the first 10 laps. This indicates that there are behavioral periods dominated by Alternation laps both before and after the peak. C) Laps were categorized into behavioral phase by computing the number of Adjustment laps in a five-lap sliding window. Windows with at least two Adjustment laps were categorized as Titration phase.

7.4 Task Validity

7.4.1 model comparison

I considered whether the choices on the spatial delay discounting task were actually being made in a manner consistent with discounted value theory by using a simulation. Choices were modeled as randomly generated probabilities from a one-parameter Beta distribution, for x=0-1 where B is the beta function and $\alpha = 1 - \beta$.

$$((1-x)^{\beta-1}x^{\alpha-1}/B(\alpha,\beta))^{1/3} \tag{7.1}$$

Decimals that rounded down to zero indicated a 'right' choice and decimals that rounded up to one indicated a 'left' choice. Use of the beta function allowed modeling of the alternation preference by biasing the equation to either zero or one depending on the value of the α parameter. Different sets of conditions were used for each of the five models tested. I tested exponential as well as hyperbolic discounting models, a moving-average model, and a number of 'simple' strategies. Four simple strategies tested were random decisions, rate-maximization, maximizing total food, and minimizing the delay then 'caching out' at the end of a session. $N = 11$ simulated agents ran 30 simulated sessions of the delay discounting experimental design. Initial delays were randomly generated and counterbalanced.

What would titration look like if rats chose randomly? Random decisions were modeled with $\alpha = 0.5$, producing a uniform distribution of mean adjusting delays (16 ± 10s) with no structure to the titration (Figure 7.2A). Adjustment and Alternation laps occurred at a uniform rate of 50% throughout the session (Figure 7.2B) and no transition occurred from Titration to Exploitation (Figure 7.2C). Varying the probability α from 0-1 resulted in a 'side bias' where choice proceeded to the left or right with a constant probability. This led to increasingly divergent mean adjusting delays on the left versus right sides (Figure 7.2D).

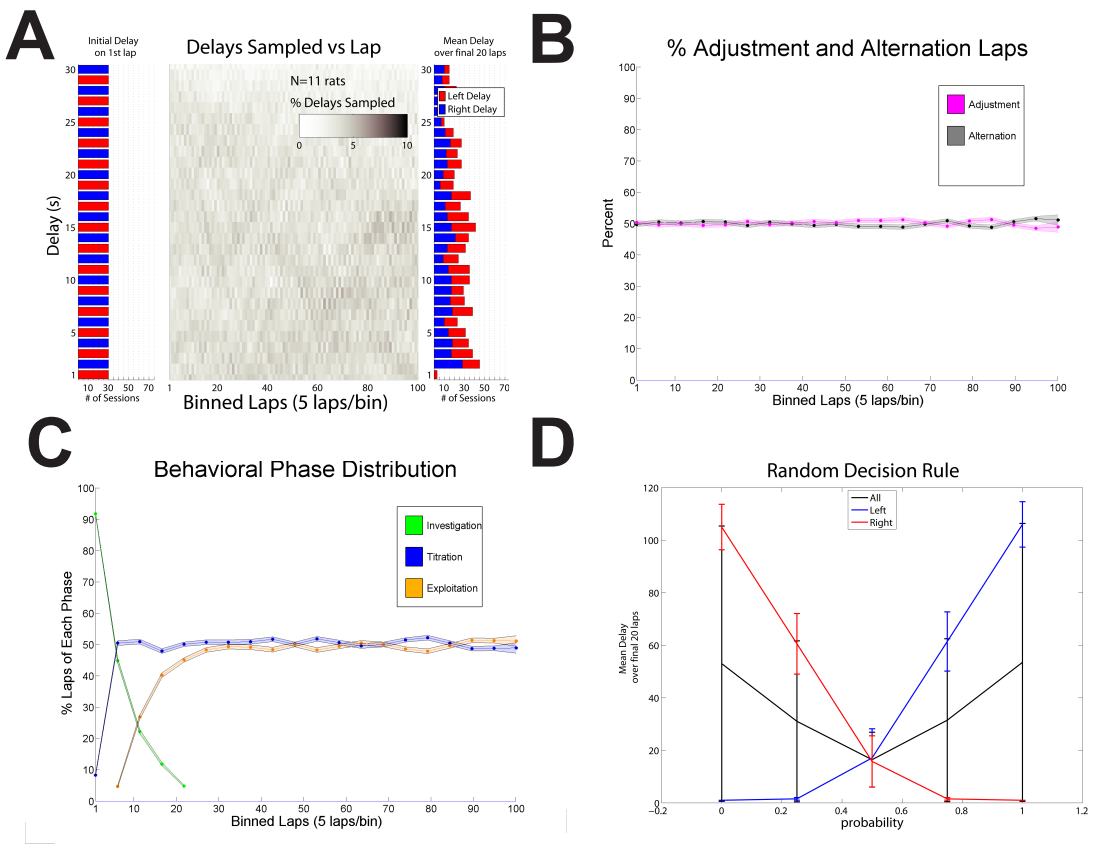

Figure 7.2: **Random Model.** A random decision model had a uniform distribution of mean adjusting delays had no structure to titration. A) This figure shows the histogram of starting delays (left panel), percentage of chosen delays in each five lap bin (center panel) and histogram of mean adjusting delays (right panel). The larger-later side is either on the right (red) or on the left (blue). B) Adjustment (magenta) and alternation (gray) occur randomly across a session. C) The distribution of Investigation (green), Titration (blue) and Exploitation (orange) laps. D) Varying the probability α term from 0-1 produced side bias. Mean adjusting delays are averaged over both sides (black), and for right side (red) and left side (blue) separately. The data in A-C are at probability of 0.5.

What would titration look like if rats maximized rate of reward? The maximization of rate of reward is a concept found in foraging theory (Stephens and Krebs, 1987; Dall et al., 2005), and is a plausible model for choice behavior. Here, rate maximization was modeled as a ratio of reward magnitude N to time per lap, including the delay D and a constant 'lap time' parameter k.

$$\frac{N_1/(D_1+k)}{(N_1/D_1+k+N_2/D_2+k)} \tag{7.2}$$

Rate maximization with lap time $k = 10$s produced a normal distribution of mean adjusting delays (Lilliefors test, $p = 0.09$, kstat=0.03) with a gradual transition from Titration to Exploitation (Figure 7.3A). The rate of Adjustment was constant at 25% throughout a session (Figure 7.3B). A lap time of 10s was chosen because it is similar to the lap time of rats in the experiment.

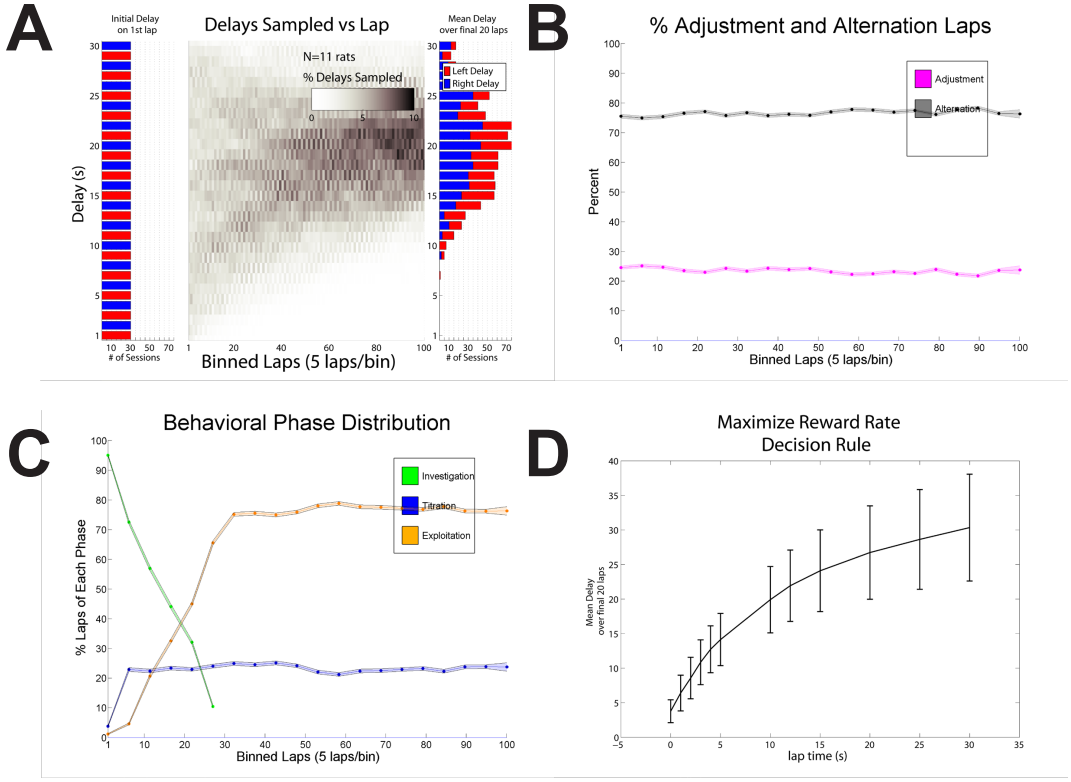

Figure 7.3: **Rate Model.** The rate model produced titration to a higher indifference point than seen in the experimental data as well as a uniform rate of Adjustment and Alternation. A) This figure shows the histogram of starting delays (left panel), percentage of chosen delays in each five lap bin (center panel) and histogram of mean adjusting delays (right panel). The larger-later side is either on the right (red) or on the left (blue). B) Adjustment (magenta) and alternation (gray) occur at a uniform rate throughout the session. C) The distribution of Investigation (green), Titration (blue) and Exploitation (orange) laps. D) Varying the lap time k term from 0-30s produced increasing mean adjusting delays. The data in A-C are at a lap time of 10s.

What would titration look like if rats maximized total food intake? To maximize the total food intake, rats would adjust the delay upward to a high value and then alternate. The Titration phase would always be upward to a delay depending on the running speed of the rat, and the Exploitation phase would be at a high delay above 30s. Although there is no experimental ceiling on the length of the delay, in practice, the maximum delay that can be sustained for 100 laps in 45 minutes is around 70s. To

model this behavior, a calculation was performed that computed the maximum possible delay without exceeding a 45 minute session time. A constant lap time was assumed. The agent titrated to that delay and then alternated. (Figure 7.4A). 'Decisions' for this model were hard-coded so no beta distribution was used for action selection. Increasing the lap time parameter to 30s caused the mean adjusting delay to decrease linearly to 55s (Figure 7.4D).

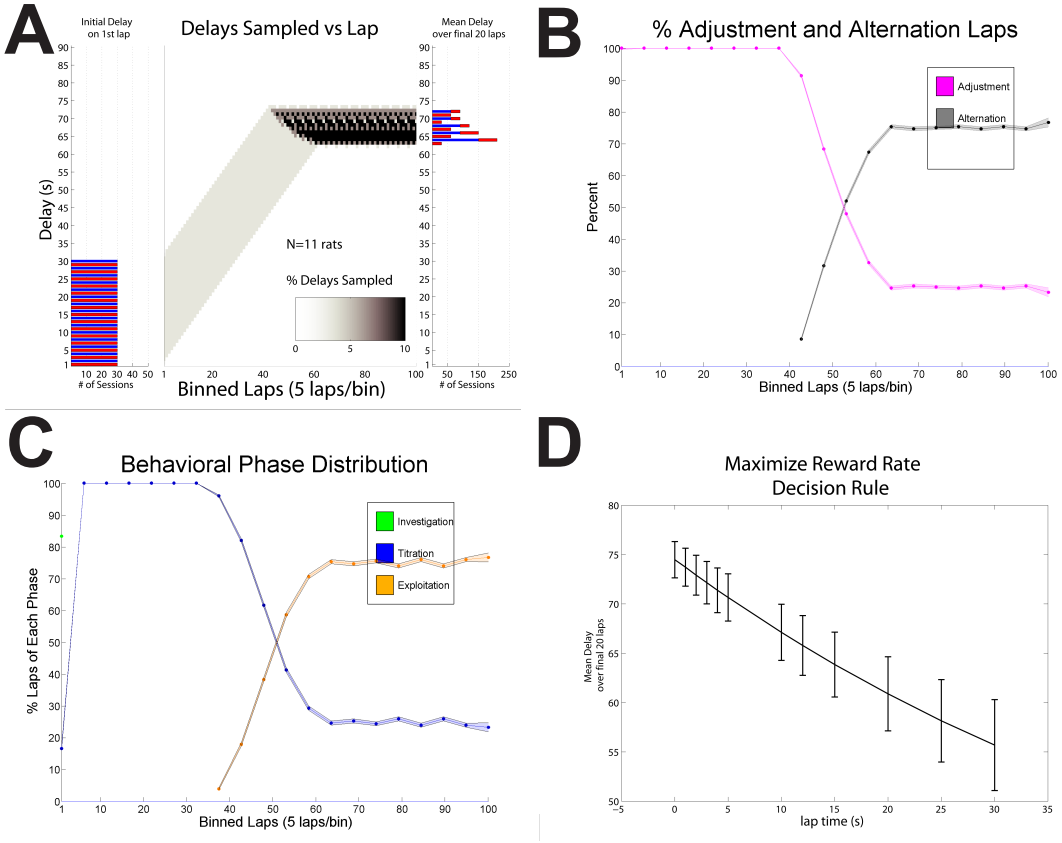

Figure 7.4: **Max Food Model.** A) This figure shows the histogram of starting delays (left panel), percentage of chosen delays in each five lap bin (center panel) and histogram of mean adjusting delays (right panel). The larger-later side is either on the right (red) or on the left (blue). B) Adjustment (magenta) and alternation (gray) have a sharp inflection point due to the hard-coding of the decision rule. C) Similarly, the distribution of Titration (blue) and Exploitation (orange) laps have a sharp inflection. Investigation (green) does not occur as agents adjust the delay from the first lap using this decision rule. D) Varying the lap time from 0s to 30s caused a decrease in the mean adjusting delay from 70s to 55s.

What would titration look like if rats avoided the larger-later side and then 'cached out' at the end of a session? During some sessions, rats would adjust the delay to 1s and alternate, with a second upward Titration phase in the second half of the session. I tested the possibility that the second Titration phase could be a result of aversion to the larger-later side coupled with a 'caching out' at the end of a session to obtain enough food per day. The delay aversion was set by a single parameter in the beta equation α, and caching out began on a randomly chosen lap from 70 to 90. The mean adjusting delay for $\alpha = 0.78$ over the last 20 laps had a higher standard deviation than the experimental data (Figure 7.5A). Adjustment occurred at a uniform rate of 33% until the caching out inflection point (Figure 7.5B). The Exploitation phase occurred in a 'U-shaped' pattern (Figure 7.5C). Increasing the delay aversion parameter caused the mean adjusting delay to asymptote around 5s. Note that delay 'aversion' for $\alpha < 0.5$ actually resulted in preference for the larger-later side, causing upward titration. (Figure 7.5D).

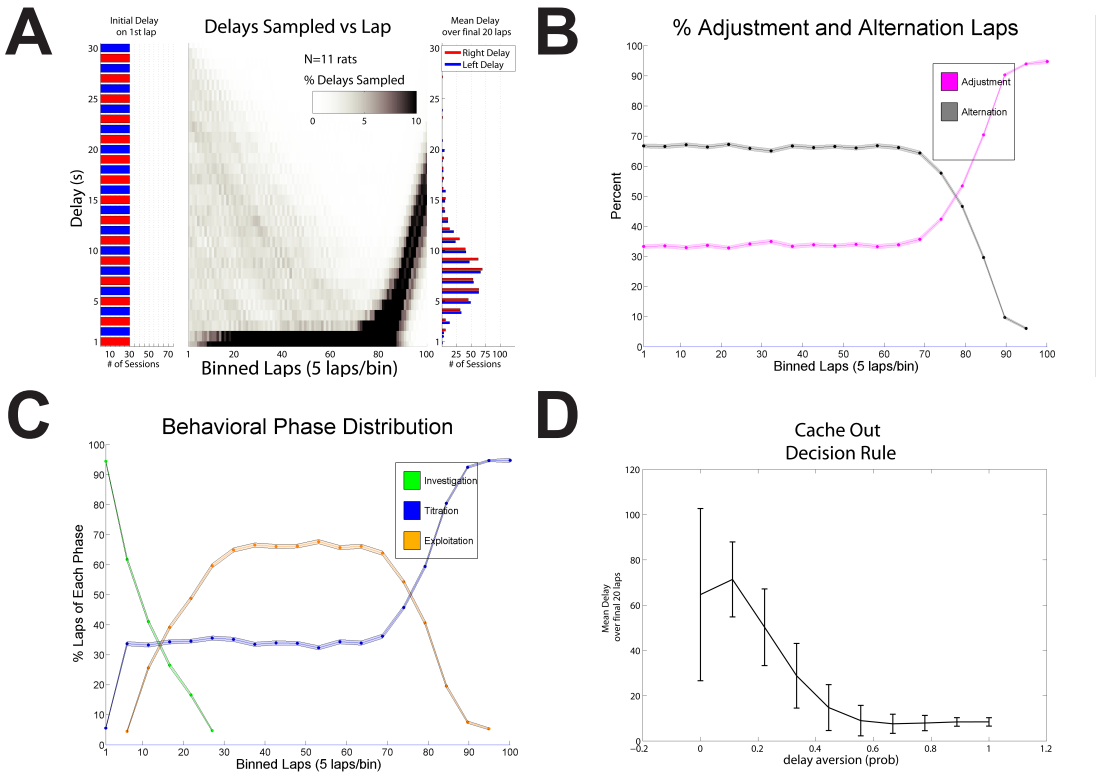

Figure 7.5: **Cache Out Model.** A) This figure shows the histogram of starting delays (left panel), percentage of chosen delays in each five lap bin (center panel) and histogram of mean adjusting delays (right panel). The larger-later side is either on the right (red) or on the left (blue). B) Adjustment (magenta) and alternation (gray) have a sharp inflection point due to the hard-coding of the cache-out decision rule. C) The distribution of Investigation (green), Titration (blue) and Exploitation (orange) laps. D) Varying the delay aversion parameter α from 0 to 1 caused a decrease in the average and standard deviation of the mean adjusting delay. Note that $\alpha < 0.5$ produced delay preference, causing upward titration.

What would titration look like if rats used the recent history of the delay to base their decisions? To answer this question, I used a moving-average model of the delay over a window of previous laps' delays. The moving average model computed value according to a two-parameter equation where $f(L)$ ($f(R)$) indicate a choice to the left (right) and nL (nR) indicate the number of choices to the left (right). The symbols for number of pellets N and delay D remain the same. The window parameter nW counts backwards from the current lap iL. The lap time parameter k is added to the chosen delays on each side in the denominator of the equation.

$$\frac{\frac{(nL*N_1(L)}{(\sum_{iW=iL-nW}^{iL} D_1(L)+nL*k)}}{\frac{((nL*N_1(L)}{(\sum_{iW=iL-nW}^{iL} D_1(L)+nL*k)} + \frac{(nR*N_2(R)}{(\sum_{iW=iL-nW}^{iR} D_2(R)+nR*k))}} \tag{7.3}$$

For a lap time of $k = 10$s and a window of $nW = 3$ laps, the moving average model produced titration to 20s (Figure 7.6A). The proportion of Adjustment laps increased briefly in the first five laps because the decision rule produced random decisions when the number of laps was less than nW. Otherwise, titration was constant at 25% (Figure 7.6B). Varying the lap time parameter k from 0-25s and the lap window parameter nW from 0-7 laps produced this series of curves for the mean adjusting delay. The mean adjusting delay was only weakly related to the lap window, while both the average and standard deviation of the mean adjusting delay increased with lap time. For a lap window of zero laps, all lap time simulations gave similar results (Figure 7.6D).

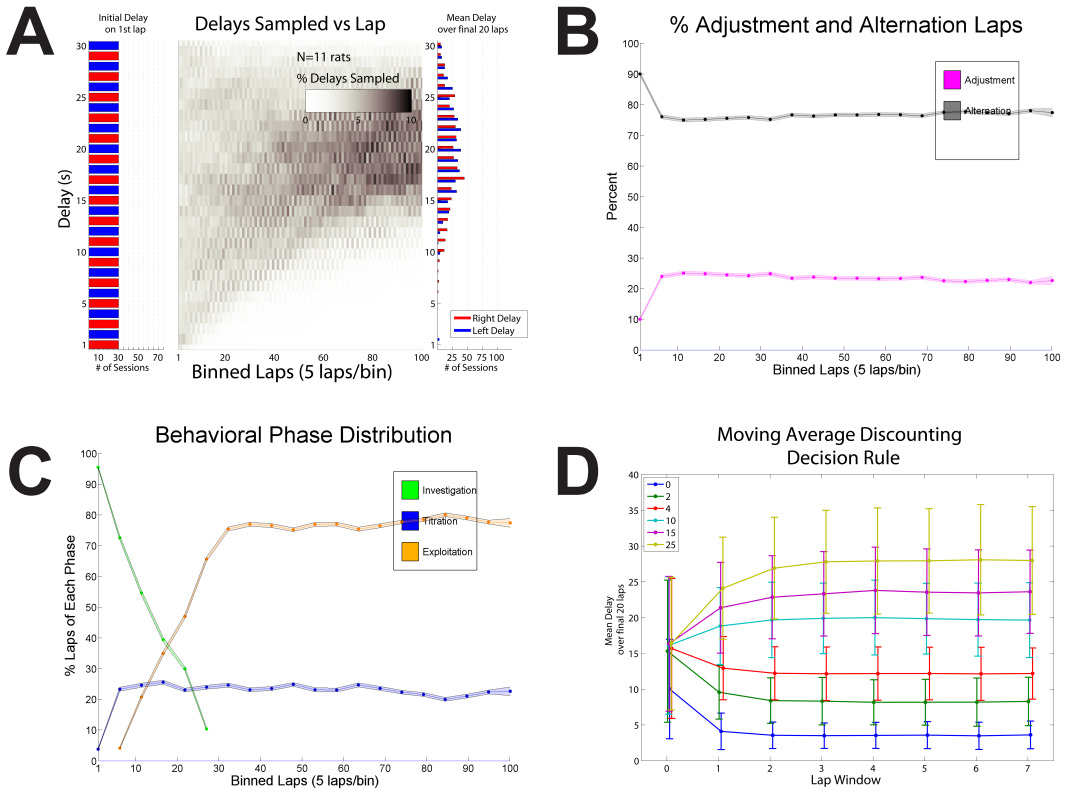

Figure 7.6: **Moving Average Model.** A) This figure shows the histogram of starting delays (left panel), percentage of chosen delays in each five lap bin (center panel) and histogram of mean adjusting delays (right panel). The larger-later side is either on the right (red) or on the left (blue). Data is from the moving average model with a lap time of $k = 10$s and a window of $nW = 3$ laps. B) Adjustment (magenta) and Alternation (gray) occur at a constant rate throughout the session. C) The distribution of Investigation (green), Titration (blue) and Exploitation (orange) laps. D) The mean adjusting delay as a function of the lap time parameter k and the window parameter nW. Mean adjusting delay increases with lap time, and is only weakly changed by window length. Data from A-C are at a lap time of 10s and a lap window of three.

What would titration look like if rats evaluated their options based on an exponential function of time? The exponential model implemented a value-based decision discounted according to a one-parameter equation where $N_2 = 1$, $D_2 = 1$, $N_1 = 3$, D_1 is the adjusting delay, and k is the discounting parameter.

$$\frac{N_1 * e^{-k*D_1}}{(N_1 * e^{-k*D_1} + N_2 * e^{-k*D_2})} \tag{7.4}$$

The exponential discounting model for discounting parameter $k = 0.25$ produced steep titration to consistent mean adjusting delays (Figure 7.7A). Adjustment peaked above 50% on early laps and decreased to about 20% by lap 30, remaining at that level for the remainder of the session (Figure 7.7B). The distribution of Investigation (green), Titration (blue) and Exploitation (orange) laps shows that Titration peaks early in the session then declines to a lower level. Varying the discounting parameter k from 0 to 1.25 caused a steep decrease in the average and standard deviation of the mean adjusting delay. (Figure 7.7D).

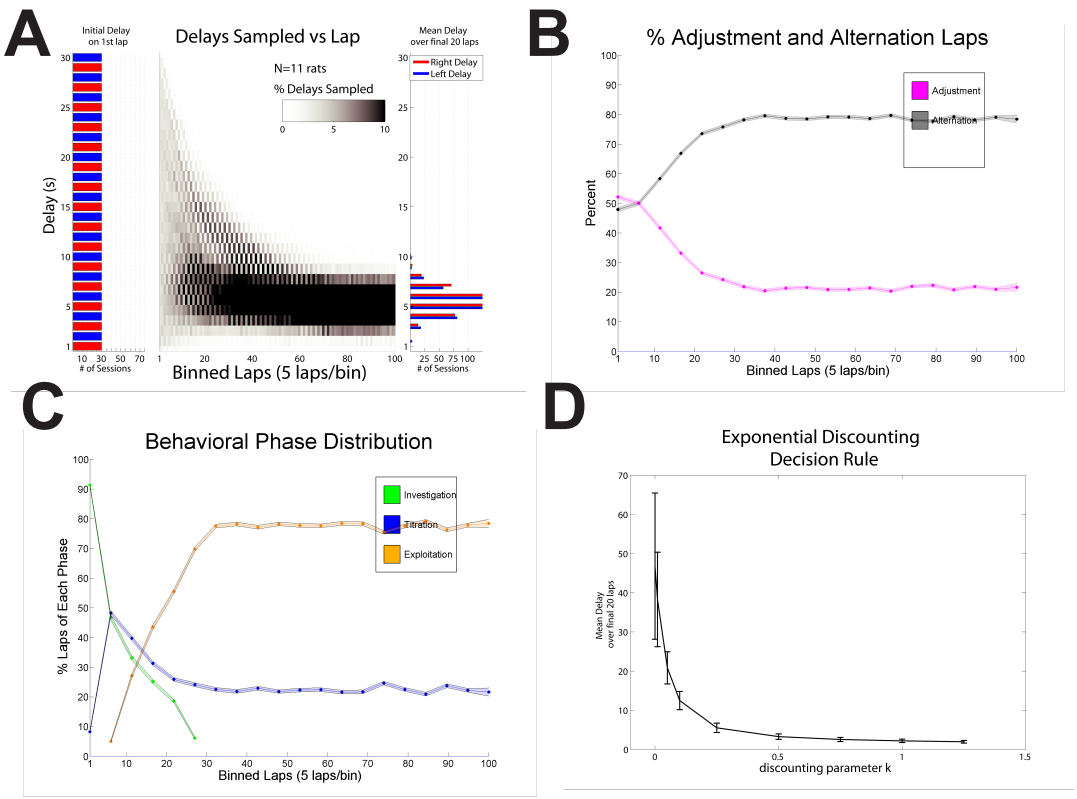

Figure 7.7: **Exponential Model.** A) This figure shows the histogram of starting delays (left panel), percentage of chosen delays in each five lap bin (center panel) and histogram of mean adjusting delays (right panel). The larger-later side is either on the right (red) or on the left (blue). Data is from the exponential model with a discounting parameter of $k = 0.25$. B) Adjustment (magenta) and Alternation (gray) occur at roughly equal proportions early in the session, but asymptote to steady-state levels by lap 30. C) The distribution of Investigation (green), Titration (blue) and Exploitation (orange) laps shows that the Titration phase peaks around 50% in laps 5-10, then decreases to a steady-state level around 25% by lap 30. D) Varying the discounting parameter k from 0-1.25 caused the average and standard deviation of the mean adjusting delay to decrease.

What would titration look like if rats were evaluating options based on a hyperbolic function? The hyperbolic model implemented a value-based decision discounted according to a one-parameter equation where $N_2 = 1$, $D_2 = 1$, $N_1 = 3$, D_1 is the adjusting delay, and k is the discounting parameter.

$$\frac{N_1/(1+e^{k*D_1})}{(N_1/(1+e^{k*D_1}) + N_2/(1+e^{k*D_2}))} \tag{7.5}$$

The hyperbolic discounting model produced titration that was shallower than the exponential model but not as shallow as the moving average model (Figure 7.8A). Adjustment gradually decreased during the first half of the session, from 33% to about 20% (Figure 7.8B). The distribution of Titration laps peaked at lap 5-10 around 33% and decreased steadily thereafter (Figure 7.8C). Varying the k parameter from 0-1.5 produced a decrease in the mean and standard deviation of the mean adjusting delay (Figure 7.8D).

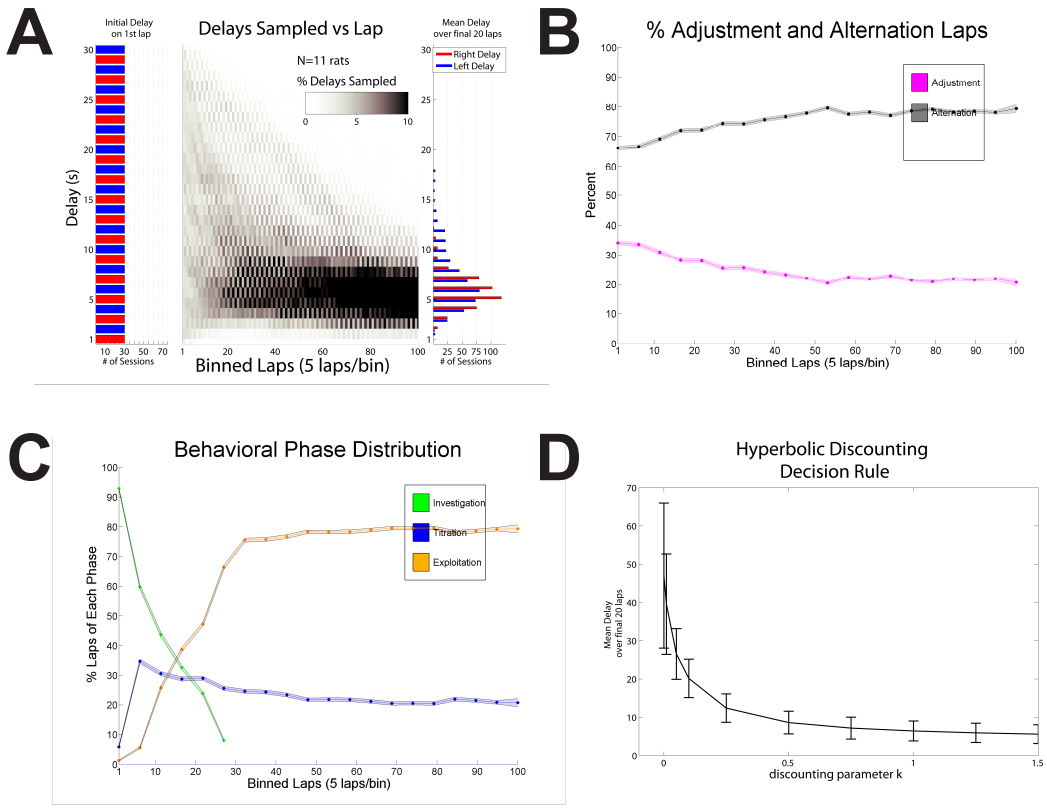

Figure 7.8: **Hyperbolic Model.** A) This figure shows the histogram of starting delays (left panel), percentage of chosen delays in each five lap bin (center panel) and histogram of mean adjusting delays (right panel). The larger-later side is either on the right (red) or on the left (blue). Data is from the hyperbolic model with a discounting parameter of $k = 1$. B) Adjustment (magenta) and alternation (gray) asymptote to steady-state levels by lap 60. C) The distribution of Investigation (green), Titration (blue) and Exploitation (orange) laps shows that the Titration phase peaks around 35% in laps 5-10, then decreases to a steady-state level around 20% by lap 30. D) Varying the discounting parameter k from 0-1.5 produces a decrease in the average and standard deviation of the mean adjusting delay.

None of the models captured all of the behavioral dynamics of rats on the spatial delay discounting task. Hyperbolic and exponential discounting models produced the closest matches to the Adjustment/Alternation and behavioral phase distribution plots. However, titration for the exponential model was too abrupt and the transition from the Titration to Exploitation phase in the hyperbolic model was too gradual. The mean adjusting delay distribution for the behavioral data is non-normal, and resembles the broader distribution of the rate model or cache out model. Examining the titration of rats (Figure 7.1A), the delays sampled by lap had a slight curvature to it which was not adequately captured by any of the models. Future research may wish to modify the beta decision rule to better fit the data, as the cubed root implementation was somewhat arbitrary. Alternately, behavior might best be modeled by a transition from a goal-directed to a habitual decision-making system involving different types of valuation (Daw et al., 2005; Keramati et al., 2011).

7.4.2 Sensitivity to Value

While titration to a consistent indifference point suggested rats were using representations of value to base their decisions (Papale et al., 2012), to further support this claim we tested additional predictions about the relationship between the indifference point and the magnitude of the larger-later reward.[1] If rats were basing decisions on a representation of value, then the indifference point should increase with increased larger-later reward magnitude N_2, reflecting its increased value relative to a fixed smaller-sooner option N_1. Specifically, the hyperbolic discounting model predicts a linear relationship between the indifference point D_2 and the magnitude of the larger-later reward (Bradshaw and Szabadi 1992; Equation 7.6). The model also predicts that discounting can be described by a single parameter k, which appears in both the slope m (Equation 7.7) and intercept b (Equation 7.8) of the linear equation. In this experiment we varied the larger-later reward from one pellet to five pellets, while keeping the smaller-sooner reward constant at one pellet.

$$D_2 = \frac{1/k + D_1}{N_1} N_2 - \frac{1}{k} \tag{7.6}$$

[1] For details on the methods of this experiment, see (Papale et al., 2012).

$$k = \frac{1}{(m * N_1) - D_1} \tag{7.7}$$

$$k = -\frac{1}{b} \tag{7.8}$$

The larger the magnitude of the larger-later reward, the higher the mean- adjusting delay over the final 20 laps (Kruskal-Wallis, $p < 10^{-10}$, $\chi^2(114) = 55.03$). This effect was well-characterized by a linear relationship ($R^2 = 0.42$, $\beta = 2$, tstat = 9, $p < 10^{-10}$). These observations indicate that behavior on the variable-reward discounting task was consistent with hyperbolic discounting.

The slope and intercept of a least-squares fit linear regression for each rat was used to calculate the discounting parameter k for each rat. For certain possible scenarios, k parameters calculated from actual data can be negative. This occurs when either $m < D_1/N_1$, or $b > -1$. For our task, $D_1 = 1$ and $N_1 = 1$, so $k < 0$ if $b < 1$. If $m = 1$ or $b = 0$, the discounting parameter is undefined. All slopes were within these parameters. I found discounting parameters of $4.11, 0.66, 0.88, 1.16$ derived from the slope and $0.53, 0.22, 1.10, 0.78$ derived from the intercept (Figure 7.9B).

On individual sessions, k can be estimated as $(1 - N_2)/(N_2 - D_2)$. If $D_2 < N_2$, then $k < 0$, and if $D_2 = N_2$, then k is undefined. I found that $79/115$ (69%) of individual sessions were within these parameters. Taking the mean and standard deviation of individual sessions gave estimates of k to be (1.54 ± 1.05, 0.87 ± 0.31, 0.33 ± 0.11, 0.75 ± 0.19; Figure 7.9C).

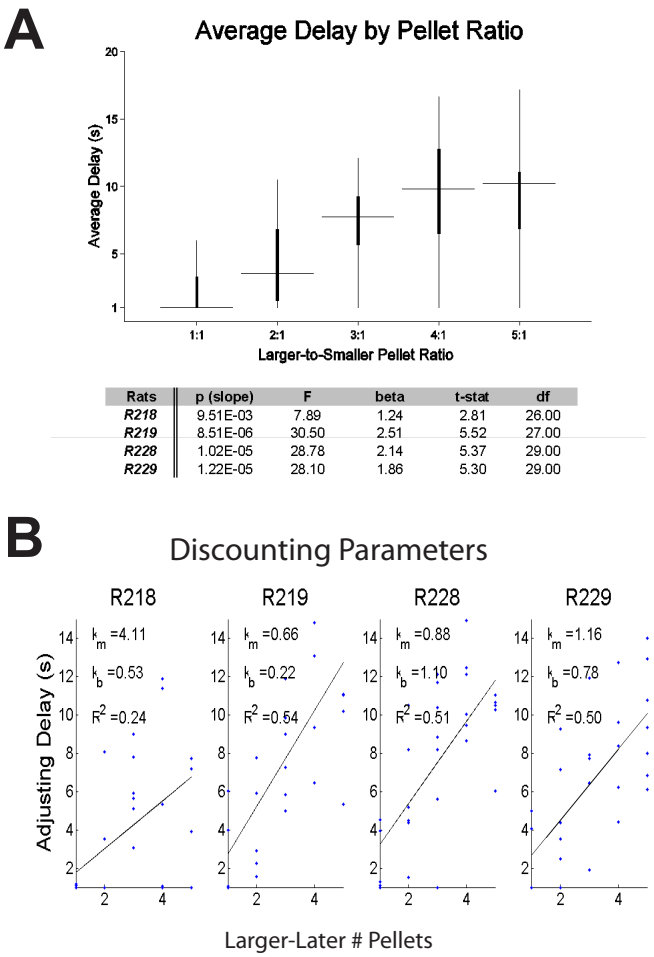

Rats	p (slope)	F	beta	t-stat	df
R218	9.51E-03	7.89	1.24	2.81	26.00
R219	8.51E-06	30.50	2.51	5.52	27.00
R228	1.02E-05	28.78	2.14	5.37	29.00
R229	1.22E-05	28.10	1.86	5.30	29.00

Figure 7.9: **Sensitivity to Value.** A separate experiment was conducted where the magnitude of the larger-later reward was varied from 1-5 pellets. The mean adjusting delay over the final 20 laps was linearly related to the reward magnitude. When both rewards were 1 pellet, rats titrated the delay to 1s and alternated for the remainder of the session. When the larger-later reward was three pellets, rats titrated the delay to the 7s range, consistent with the first experiment. The adjusting delay for the four and five pellet rewards was similar, possibly due to satiation at large reward values. The table shows the regression statistics for each rat.

7.5 VTE occurred on Adjustment Laps

On the spatial delay discounting task, VTE occurred most frequently on Adjustment laps when animals used a win-stay strategy. The z-scored integrated change in orientation (zIdPhi) was used to quantify VTE behavior at the choice point (See **Section** 6.3.5). zIdPhi was higher on Adjustment laps (Wilcoxon rank sum, $p = 2.8 \times 10^{-132}$, zval = 24.47). Different within-session explanatory variables were considered: Adjustment/Alternation, adjusting delay, lap number, behavioral phase, and delay zone. A stepwise linear regression was performed to determine the effect size of each variable on zIdPhi. Overall, the order of effects was Adjustment/Alternation ($\beta = 0.54$, $p = 7.3 \times 10^{-90}$, tstat = 20.33) and behavioral phase ($\beta = -0.061$, $p = 6.6 \times 10^{-5}$, tstat = -3.99). While these terms suggested plausible explanations for VTE, several factors warranted additional analysis. The β term on behavioral phase was small, and both terms are strongly correlated with lap number and adjusting delay. Therefore, we looked at each variable independently.

A significant decrease in VTE across lap was observed (Kruskal-wallis, $p < 10^{-10}$, $\chi^2(99) = 583.6$). suggesting that VTE was more likely to occur earlier during the session (Figure 7.10E). However, when analyzing Adjustment and Alternation laps separately, the effect was limited to Alternation laps, while VTE occurred at a similar level on Adjustment laps throughout a session. To examine the relative contribution of lap number versus Adjustment/Alternation, I performed a two-way ANOVA. The main significant effect was Adjustment/Alternation ($p < 10^{-10}$, df = 1, F = 1, 578); although there was a significant effect of lap number ($p < 10^{-10}$, df = 98, F = 2.28), the effect size was small.

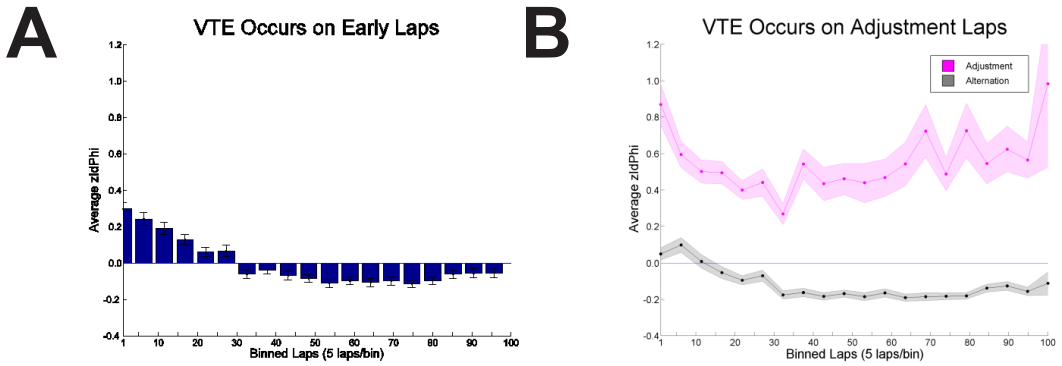

Figure 7.10: **VTE occurs on Early Laps laps.** Within-session, the occurrence of VTE was related most strongly with the choice of an Adjustment lap win-stay strategy. A) VTE was measured for laps grouped by their behavioral phase. Overall, VTE was higher during Titration than during Exploitation phase. During early laps, VTE was higher regardless of behavioral phase. B) When phase classification was divided into the component Adjustment and Alternation laps, within-group, VTE was greater on Adjustment laps.

A significant increase in VTE with adjusting delay was observed (Kruskal-wallis, $p < 10^{-10}$, $\chi^2(29) = 266$), suggesting that VTE occurred during high delays (Figure 7.11A). However, when analyzing Adjustment and Alternation laps separately, the effect was limited to Alternation laps, while VTE occurred at a uniform rate for Adjustment laps (Figure 7.11B). A two-way ANOVA was used to assess the relative effect size of adjusting delay and Adjustment/Alternation on VTE. The main effect was Adjustment/Alternation ($p < 10^{-10}$, df = 1, F = 1,501); although there was a significant effect of adjusting delay ($p < 10^{-10}$, df = 29, F = 4.57), the effect size was small.

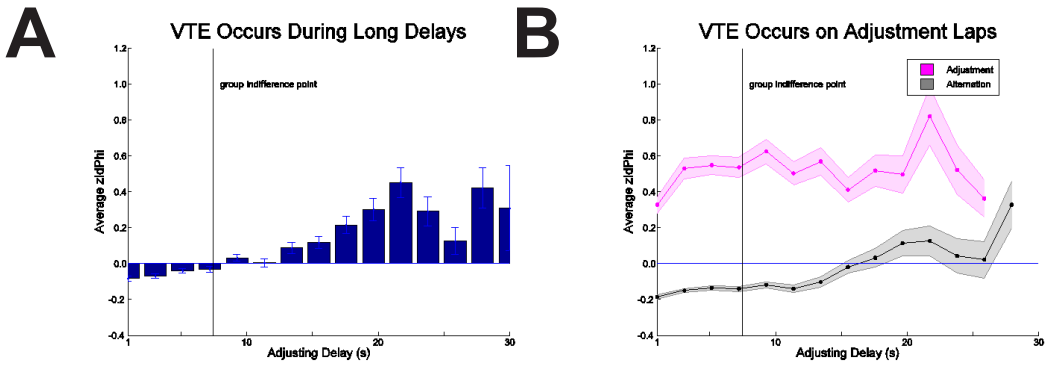

Figure 7.11: **VTE occurs on High Delays.** Within-session, the occurrence of VTE was related most strongly with the choice of an Adjustment lap win-stay strategy. A) VTE was more likely to occur at high delays. VTE was compared with the adjusting-delay on each lap (regardless of whether or not the rat chose the larger-later side). The indifference point across all rats is $7.5 \pm 4.8s$. B) When each group of delays was broken up into component Adjustment and Alternation laps, across all delays, VTE was more likely to occur on Adjustment laps. VTE also increased with delay on Alternation laps.

VTE was higher during the Titration and Investigation phases, indicating a general increase during the period of flexible responding when the change in adjusting delay is greatest (Kruskal-wallis, $p < 10^{-10}$, $\chi^2(2) = 348.36$; Figure 7.12A). To determine the relative effect size of behavioral phase versus Adjustment/Alternation, a two-way ANOVA was performed. The main significant effect was Adjustment/Alternation ($p < 10^{-10}$, df = 1, F = 1,626); although there was a significant effect of phase ($p < 10^{-10}$, df = 2, F = 29.11), the effect size was small.

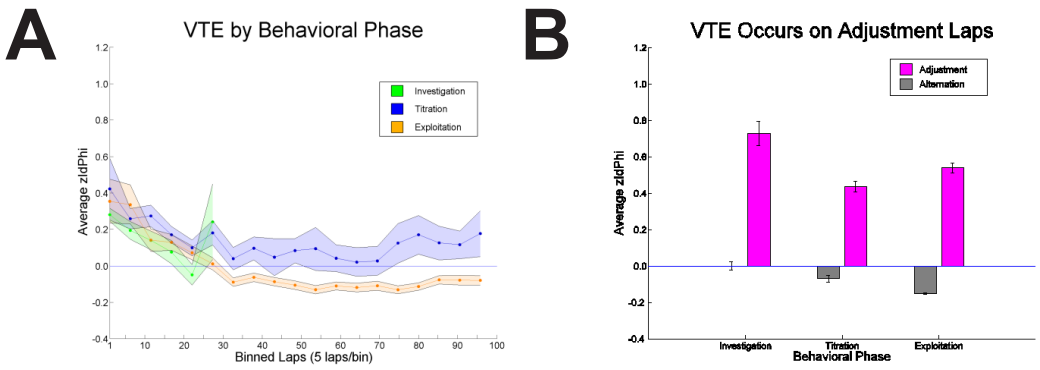

Figure 7.12: **VTE occurs on Titration laps.** Within-session, the occurrence of VTE was related most strongly with the choice of an Adjustment lap win-stay strategy. A) VTE was more likely to occur on early laps. VTE was computed for each lap in five lap bins. VTE decreased with lap for the first half of the session, asymptoted at three-fourths of the way through the session, and increased towards the end of the session. B) When each lap was categorized as an Alternation or Adjustment, VTE was found to be higher on Adjustment laps. VTE decreased with lap on Alternation laps.

Interestingly, a given VTE event had a roughly equal probability of occurring on either an Adjustment or an Alternation lap, in agreement with previous results (Stott and Redish, 2014). This is not inconsistent with an increased occurrence of VTE on Adjustment laps, because the there are roughly four times more Alternation laps per session. It is interesting to speculate whether more VTE would be observed on Alternation laps if a task was designed such that the proportion of win-stay trials was greater than win-shift trials.

7.6 VTE occurred during periods of flexible responding within-session

Adjustment laps occurred during the Titration phase when flexible responding between win-stay and win-shift strategies was greatest. To test whether VTE was more likely to occur during groups of Adjustment laps or single, isolated Adjustment laps, a sliding window was used to compute VTE for different ratios of Alternation laps to total laps. For a ten-lap window, VTE increased on Adjustment laps with the proportion

of Alternation laps in the window. VTE on Alternation laps was constant across the window. (Figure 7.13A). Varying the sliding window from one-to-ten laps produced similar results (Figure 7.13B). VTE was most likely to occur during isolated Adjustment laps and less likely to occur during groups of Adjustment laps.

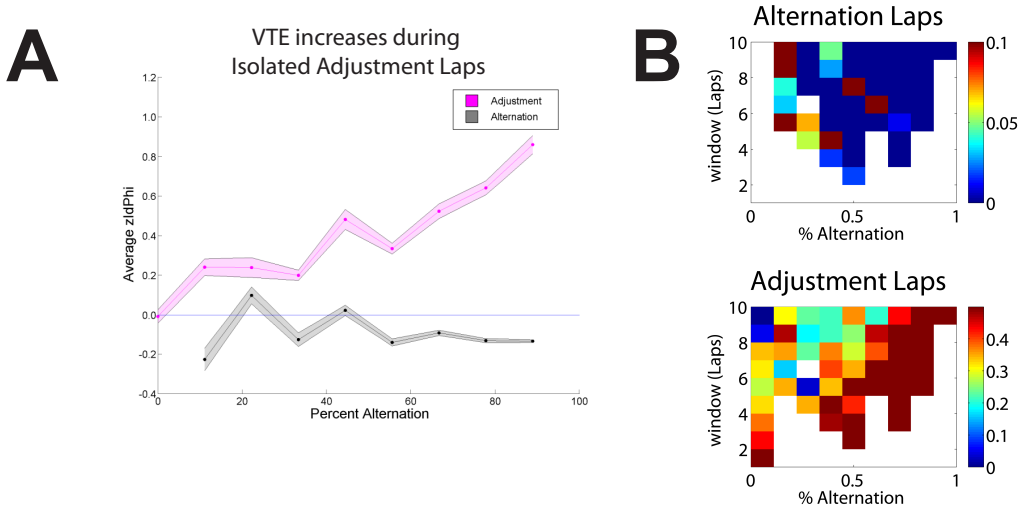

Figure 7.13: **VTE occurred during isolated Adjustment laps.** VTE was higher on isolated Adjustment laps that occurred amid a group of Alternation laps. A) The ratio of Alternation laps to total laps was computed for each lap in a ten-lap sliding window and zIdPhi was averaged for each bin. As the percentage of Alternation laps increased, VTE was more likely to occur on an Adjustment lap. VTE on Alternation laps did not change as a function of the number of Alternation laps in the window. VTE was more likely to occur during periods of variable behavior when an Adjustment lap occurred amid a group of Alternation laps. B) The sliding window was varied from one-to-ten laps and the analysis repeated, with similar results. When varying window size, Alternation laps (left panel) showed a lower average VTE and little variability. Color indicates average zIdPhi from 0 to 0.1. Adjustment laps (right panel) show more VTE, with increasing values for less-frequent Adjustment laps. Color indicates average zIdPhi from 0 to 0.5. VTE occurred more frequently during variable behavior.

7.7 VTE occurred during optimization of strategy across-session

To quantify VTE across-session, I log-transformed IdPhi and z-scored within-rat. VTE was higher on initial sessions, and decreased across-session (ANOVA, $p = 1.4 \times 10^{-204}$, df = 8, F = 128.8; Figure 7.14A). Papale et al. (2012) found that VTE occurred during the Titration phase (also see Figure 7.12). When considering Titration and Exploitation

separately, there was an interaction between phase and session (two-way ANOVA, $p = 7 \times 10^{-5}$, df = 16, F = 2.94). VTE decreased with session more strongly during the Exploitation phase (ANOVA, $p = 6.6 \times 10^{-134}$, df = 8, F = 84.69; Figure 7.14B). VTE decreased somewhat during the Titration phase (ANOVA, $p = 2.6 \times 10^{-31}$, df = 8, F = 21.7).

This suggested that there was an across-session learning effect on the spatial delay discounting task. Alternation efficiency, a measure of the number of Alternation laps at the indifference point over the total number of Alternation laps, increased with session (ANOVA, $p = 0.032$, df = 8, F = 2.15). These observations confirm that there is an optimization of Titration across-session.

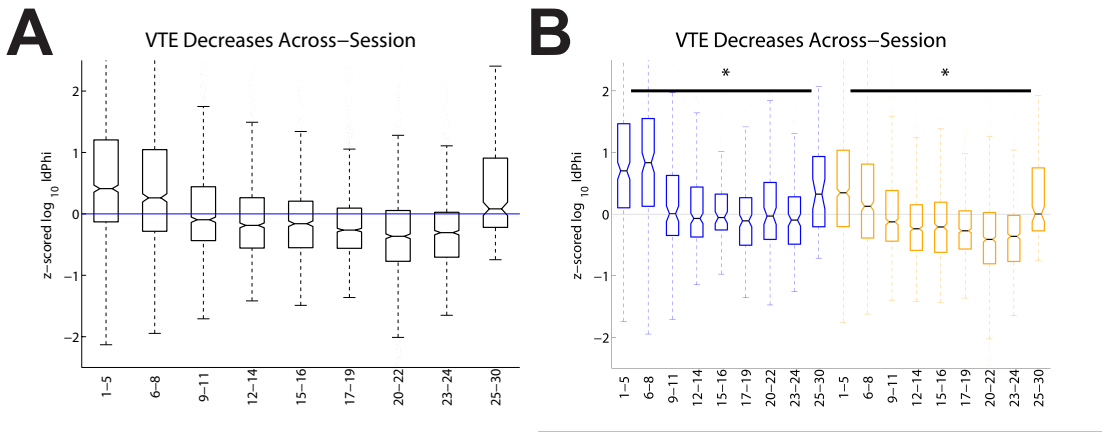

Figure 7.14: **VTE occurred during optimization of strategy across-session.** A) VTE decreased across-session. During initial sessions, VTE is, on average, around 0.33. Average zIdPhi decreased to about -0.33 towards later sessions. Note that VTE for the last five sessions did not follow this trend, appearing to increase. B) VTE decreased across-session more strongly during the Exploitation phase. Across-session, VTE decreased during Titration (blue) and Exploitation (orange). The effect was more pronounced during Exploitation.

7.8 Conclusions

On the spatial delay discounting task, rats titrated the adjusting delay to a consistent indifference point. In doing so, they passed through three distinct behavioral phases: Investigation, Titration, and Exploitation. VTE on the spatial delay discounting task was observed during Adjustment laps, which occurred primarily during the Titration

phase. VTE was also observed during early laps, at high delays, and during isolated Adjustment laps. VTE decreased across-session, while alternation efficiency increased.

Chapter 8

SWR Sequences and Theta Sequences

8.1 Sharp Wave Ripples (SWRs)

8.1.1 Detection of SWRs

SWRs occurred throughout the maze, but were concentrated at feeder locations (Figure 8.1A). $47,500/123,464$ (38%) of SWRs occurred after feeder fire, prior to entry into the next lap. The remainder were uniformly distributed in space across the track. $10,871/47,500$ SWRs at feeder locations contained at least one spike, $47,352$ had a z-score of the global extrema less than 30, $47,097$ occurred during laps with no technical errors or backwards trajectories, and $6,132$ had non-zero probability of representation at the choice point or start of maze analysis zones. Subsequent analyses were done on the $6,132$ SWRs that occurred at feeder locations and met these criteria (Figure 8.1B). Results were qualitatively unchanged if these criteria were relaxed.

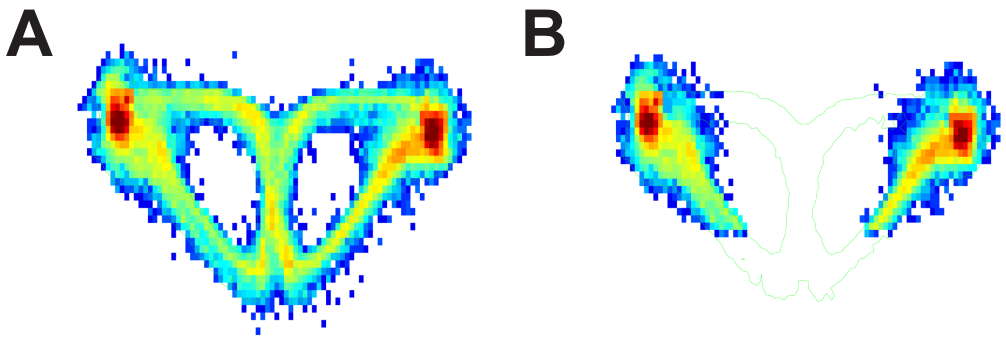

Figure 8.1: **Location of Sharp Wave Ripples.** SWRs are concentrated at feeder locations, with 38% occurring after feeder fire before beginning the next lap. Of these, 13% are used for analysis of SWR content. SWRs are detected by computing a Hilbert transform on a tetrode bandpass filtered at $150 - 250$ Hz. A) The spatial histogram of SWR occurrence on the spatial delay discounting task is shown. The task is divided into 56×56 spatial bins. The outline of the task is shown (green line). The color axis shows the log_{10} number of SWRs from zero (blue) to three (red). Note that the majority of SWRs are located at the feeders (the upper-left and upper-right corners). B) The spatial histogram of analyzed SWRs on the spatial delay discounting task is shown. Analyzed SWR are restricted to feeder locations. SWR on each lap that occurred after feeder fire and before entry into the start of maze zone are analyzed. The color axis shows the log_{10} number of SWR from zero (blue) to three (red). The outline of the task is shown (green line).

8.1.2 More SWRs occurred during non-VTE Laps and the Exploitation phase

Prior research suggested that SWR rate increased with experience (Jackson et al., 2006). On the spatial delay discounting task, rats transitioned from an Adjustment strategy (Titration phase) to an Alternation strategy (Exploitation phase) with experience on each session (Figure 7.1C). More VTE was observed during the Titration phase (Papale et al., 2012; Figure 7.12). I therefore hypothesized that the rate of SWR emission would increase during the Exploitation phase and non-VTE laps. This prediction was confirmed, as there were more SWR/s during the Exploitation phase (two-way ANOVA; $p = 2 \times 10^{-5}$, df = 1, F = 18.21), and during non-VTE laps (two-way ANOVA; $p = 2.7 \times 10^{-26}$, df = 1, F = 113.15). An interaction effect between VTE/non-VTE and Titration/Exploitation was also observed ($p = 2.3 \times 10^{-3}$, df = 1, F = 9.3; Figure 8.2).

Examining the frequency of SWRs for VTE or non-VTE on Adjustment or Alternation laps, more SWRs occurred on non-VTE (two-way ANOVA; $p = 2.4 \times 10^{-13}$, df

= 1, F = 50.78) and Alternation ($p = 5.8 \times 10^{-17}$, df = 1, F = 70.25), with an interaction effect ($p = 1.3 \times 10^{-6}$, k = 1, F = 23.42). The effect of Adjustment/Alternation (three-way ANOVA; $p = 9.4 \times 10^{-9}$, df = 1, F = 24.07) was greater than the effect of Titration/Exploitation ($p = 0.02$, df = 1, F = 5.43).

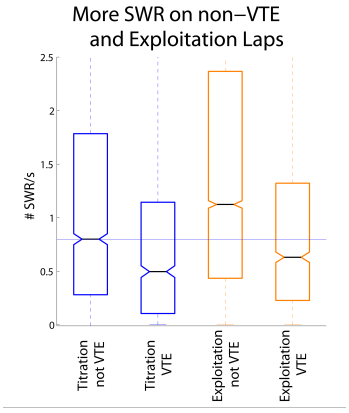

Figure 8.2: **SWRs occur during non-VTE laps and the Exploitation phase.** A boxplot of SWR/s for Titration (blue) and Exploitation (orange) during non-VTE and VTE laps. Two main effects explain the within-session timing of SWRs. SWRs occur more frequently on Exploitation laps and on non-VTE laps. SWRs are less likely to occur on laps in the Titration phase and if the upcoming pass through the choice point contains a VTE behavior.

8.1.3 SWR frequency increased and pairwise coactivity during SWRs decreased with session

Awake SWR emission has been observed to increase with learning across-session on a spatial working memory task (Jackson et al., 2006). From an initial rate of 0.5 SWR/s during the earliest sessions, SWR occurrence increased to greater than 1.5 SWR/s during the latest sessions. (Figure 8.3A).

Coordinated activity of firing decreased across-session. Pairwise co-activity among pyramidal cells was measured during SWRs and then z-scored to the average co-activity of 500 shuffled ISIs. Using this measure of coordinated activity during SWRs, I found that the number of co-active cells decreased steadily across-session. The number of co-active cells for the earliest sessions was greater than one standard deviation above chance levels. During the latest sessions, the pairwise co-activity among cells decreased to chance levels (Figure 8.3B).

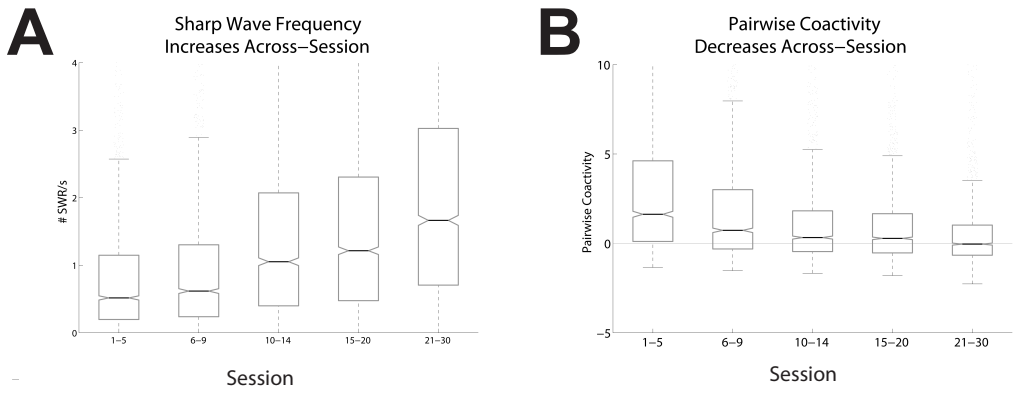

Figure 8.3: **SWR frequency increases and pairwise coactivity decreases across-session.** A) On early sessions, SWRs occur at a rate of about one every other lap. This frequency triples during later sessions. B) On early sessions, SWR pairwise coactivity is high, indicating that many different trajectories are being represented during SWRs. On later sessions, SWR pairwise coactivity is reduced to chance levels, indicating that the content of SWRs had become more focused.

8.2 Bayesian decoding during SWRs

Having observed task-relevant changes in SWRs, I next examined the spatial content of SWRs using a Bayesian decoding algorithm. An example SWR is shown with its average decoded representation on the spatial delay discounting task. The outline of the task is in green and the position of the rat is marked by a blue circle. The average decoded position is in black-to-white shading and shows a strong representation of both the leftward trajectory and left feeder (Figure 8.4A).

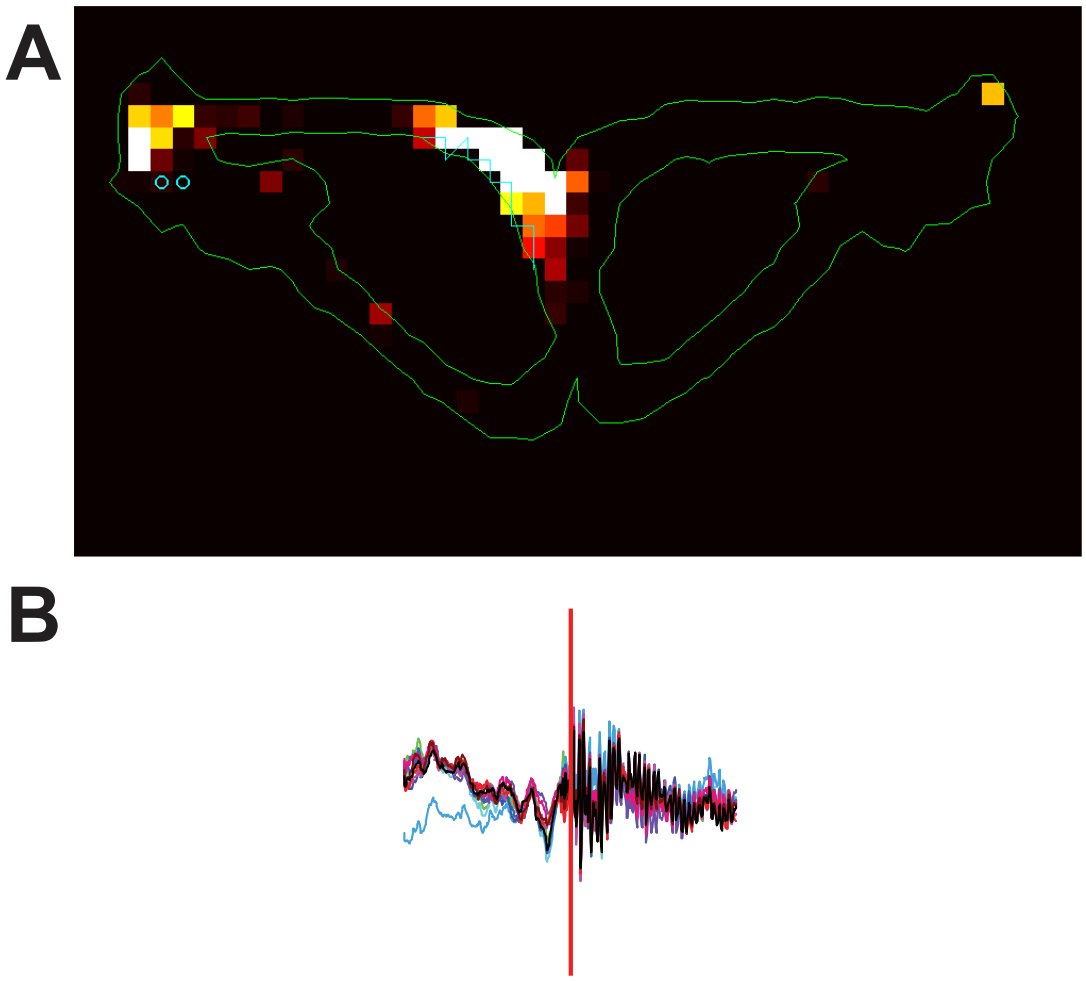

Figure 8.4: **Bayesian decoding during SWRs.** A Bayesian decoding algorithm is used to determine the position of the rat represented during SWRs at feeder locations. Nonlocal representations are observed during some SWRs. A) The spatial histogram of average decoded representation is displayed across one SWR. This SWR has a nonlocal representation of the choice point and local representation of the left feeder. Shading indicates average representation from 0% (black) to 1% (white). The outline of the maze is in green. The location of the rat during the SWR is shown by two cyan circles at the left feeder. The trajectory of the rat on the subsequent lap is shown by the cyan line. B) SWRs are detected reliably. This SWR (red line) with the LFP trace $\pm$125ms for all 12 tetrodes shows a synchronized network ripple oscillation in CA1. The majority of detected SWRs share this signature.

8.2.1 SWR decoding reflected the recent past

The content of SWRs after feeder fire was largely explained by the most recent choice of the rat. Decoding following a choice of the larger-later feeder predominantly represented the larger-later feeder (z-test; $p < 1 \times 10^{-100}$, z = -34.9) and decoding following a choice of the smaller-sooner feeder predominantly represented the smaller-sooner side feeder (z-test, $p < 1 \times 10^{-100}$, z = 47.2). Because awake SWRs have also been shown to represent future trajectories (Diba and Buzsáki, 2007; Pfeiffer and Foster, 2013), the content of SWRs was examined as a function of time after feeder fire. For this analysis, time after feeder fire was normalized from zero (time of feeder fire) to one (time of start of next lap). The log-odds ratio for SWRs at the larger-later (gold line) was above zero, indicating a higher representation of the delay-side feeder. The log-odds ratio for SWRs at the smaller-sooner feeder (maroon line) was below zero, indicating a higher representation of the non-delayed side. The representation of recent past increased significantly over this interval for the larger-later side (ANOVA; $p = 1 \times 10^{-48}$, df = 10, F = 26.28) and for the smaller-sooner side (ANOVA; $p = 3 \times 10^{-40}$, df = 10, F = 22.11). The representation varied from about 20-60%, indicating the possibility of unexplained variability in the decoded representation (Figure 8.5A).

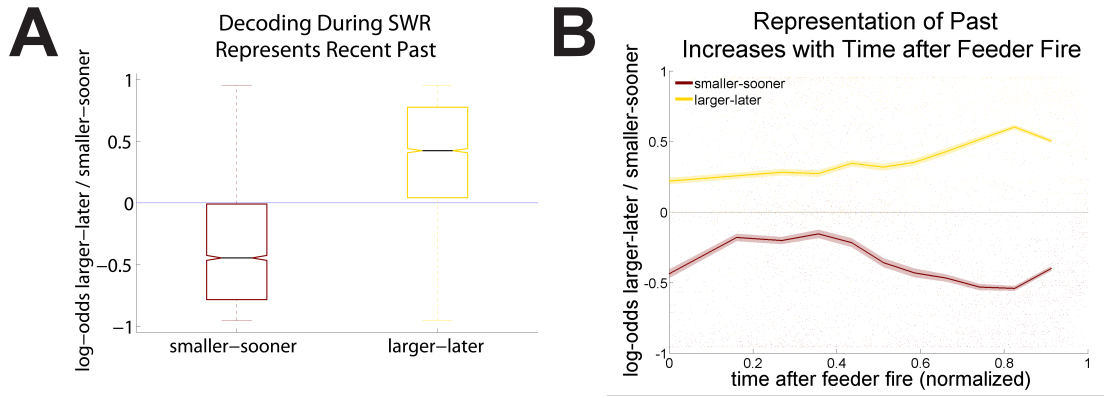

Figure 8.5: **SWR decoding reflects the recent past.** SWRs frequently represent the most recent choice of the rat. A) The boxplot shows the median log-odds representation of the larger-later versus smaller-sooner sides. Values above zero indicate a greater representation of the larger-later side, while values below zero indicate a greater representation of the smaller-sooner side. On smaller-sooner choices, SWR predominantly reflect the smaller-sooner side (maroon). On larger-later choices, SWR predominantly reflect the larger-later side (gold). B) The representations vary as a function of time after feeder fire. On the larger-later side (gold), the representation of recent past increase nearly linearly as a function of time. On the smaller-sooner side (maroon), the representation of recent past initially decreases, then increases with time from feeder fire. Note that as the log-odds approaches -1, the representation for the smaller-sooner side increases relative to the larger-later side.

8.2.2 Following non-VTE laps, SWR decoding to the opposite side of future choice increased

The previous analysis described a main effect of recent past on the representation of position during SWRs. Awake SWR representations also represent future trajectories (Diba and Buzsáki, 2007; Pfeiffer and Foster, 2013). Therefore, I considered the effect of future choice and current choice together. An effect of future choice was discovered (two-way ANOVA; $p = 0.0001$, df = 1, F = 15.58). Following from the anti-relation of SWR frequency with VTE frequency (Figure 8.2), it is possible that the content of SWR is different for VTE and non-VTE trials. Specifically, SWR may only be predictive of future trajectories following non-VTE laps, when theta sequences are not guiding decision-making. To test this relationship a four-way ANOVA was performed with the variables of current side, future choice, VTE on prior lap, and time after feeder fire. Significant main effects were observed for current side ($p = 1 \times 10^{-9}$, df = 1, F = 36.67),

and future choice ($p = 0.0013$, df = 1, F = 10.33). Significant interaction effects were observed for current side x time after feeder fire ($p = 5 \times 10^{-24}$, df = 1, F = 102.9), current side x VTE on previous lap ($p = 0.0116$, df = 1, F = 6.37), and future choice x time after feeder fire ($p = 0.0198$, df = 1, F = 5.43).

To investigate the interactions, I plotted the log-odds ratio of feeder decoding as a function of time after feeder fire separately for SWR that occurred following VTE laps and non-VTE laps. Representation before a choice of the larger-later side (gold line) decreased as a function of time after feeder fire (ANOVA; $p = 7 \times 10^{-20}$, df = 10, F = 11.8), indicating an increased representation of the *smaller-sooner side*. Representation before choosing the smaller-sooner side (maroon line) increased as a function of time after feeder fire (ANOVA; $p = 6.8 \times 10^{-24}$, df = 10, F = 13.68), indicating an increased representation of the *larger-later side*. SWR decoding after a non-VTE lap therefore represented the *opposite side* of the future choice. This is consistent with previous observations on a binary choice task (Gupta et al., 2010). Furthermore, the probability of representing the opposite side increased with time after feeder fire. In contrast, SWR decoding after a VTE lap did not show this effect. Representation of both sides was equal and did not deviate with time after feeder fire. (Figure 8.6B)

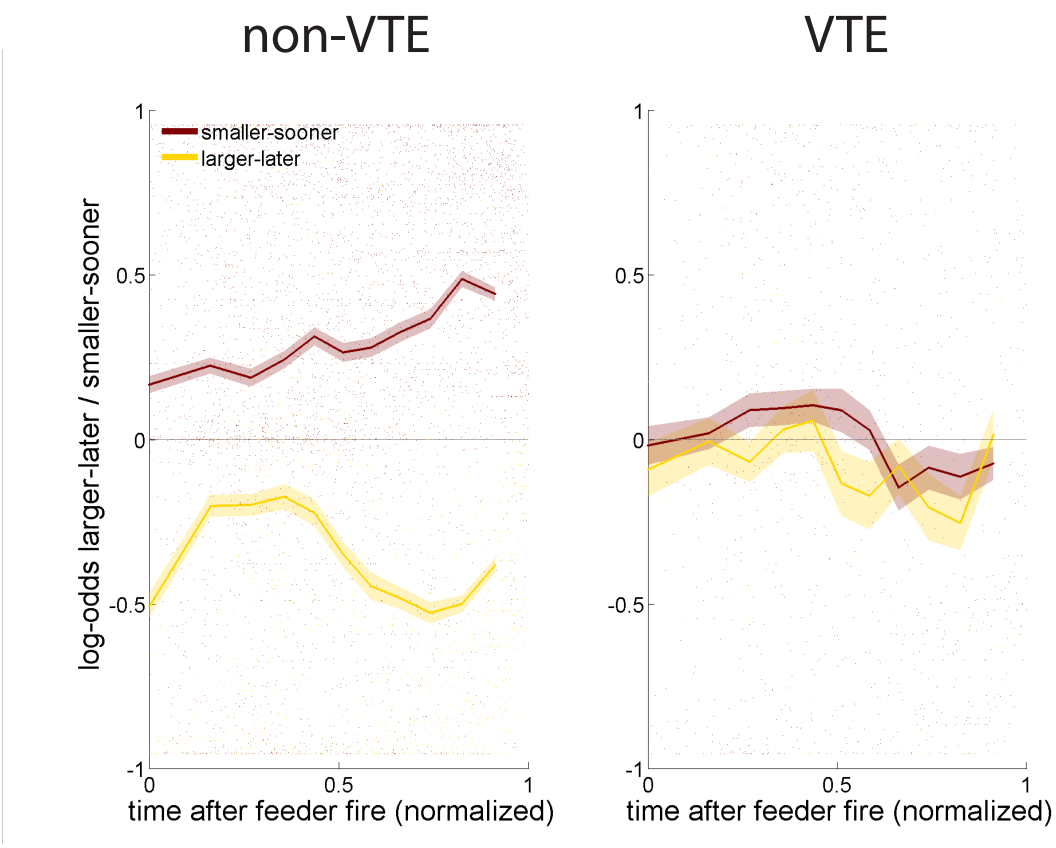

Figure 8.6: **Following non-VTE laps, SWR decoding to the opposite side of future choice increases.** SWRs following non-VTE laps predict the opposite of the future choice. A) The log-odds representation of the larger-later to smaller-sooner side is shown. Values above zero indicate a greater representation of the larger-later side, while values below zero indicate a greater representation of the smaller-sooner side. Values are plotted separately for non-VTE laps (left panel) and VTE laps (right panel). On laps with a future choice to the smaller-sooner side (maroon), the representation of the *larger-later side* increases with time after feeder fire following non-VTE laps (left panel), but not following VTE laps (right panel). In compliment, on laps previously having no VTE and a future choice to the larger-later side (gold), the representation of the *smaller-sooner side* side increases with time after feeder fire (values become more negative, indicating greater representation of the smaller-sooner side). The representations are biased toward the *opposite* of the future choice.

8.3 Bayesian Decoding during Theta Sequences

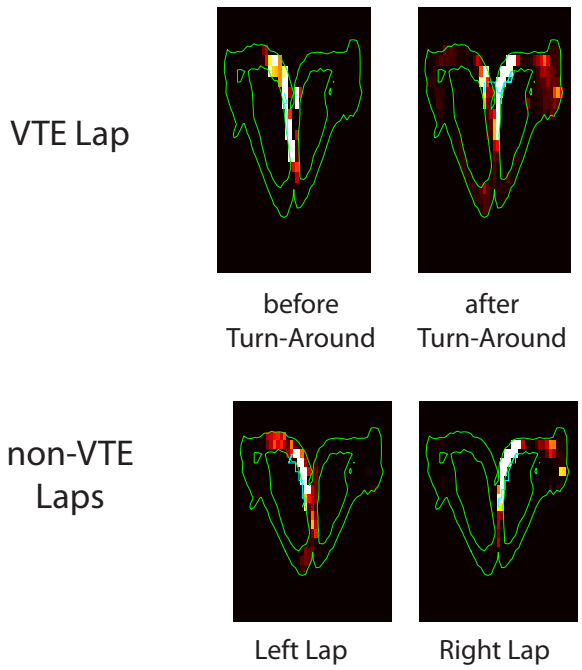

Figure 8.7: **Bayesian decoding during Theta Sequences.** A Bayesian decoding algorithm is used to determine the position of the rat represented during theta sequences at the choice point. Nonlocal representations are observed during some choice point passes. A) The spatial histogram of average representation is displayed across one choice point pass. A VTE lap (top row) contains a turn-around event at the choice point where the rat looks left, then pauses and looks right, before continuing its journey. Decoding before the turn-around (top left) shows representation of the left trajectory, including retrospective coding behind the rat and forward decoding in front of the rat. Decoding after the turn-around (top right) shows decoding to both potential feeder locations. In comparison, decoding for a non-VTE lap to the left (bottom left) and a non-VTE lap to the right (bottom right) are shown. Note that the representation lies entirely along one trajectory during non-VTE laps. Shading indicates average representation from 0% (black) to 1% (white). The outline of the maze is in green. The location of the rat during the SWR is shown by two cyan circles at the left feeder.

8.3.1 Decoding During Theta Sequences represented future choice

Previous research suggested that theta sequences predict the choice of a rat (Wikenheiser and Redish, 2015) and I tested this possibility on the spatial delay discounting

task. Decoding of the side chosen by the rat was significantly increased during non-VTE (Wilcoxon signed rank; $p < 1 \times 10^{-99}$, zval = -58.82), and VTE (Wilcoxon signed rank; $p = 1 \times 10^{-32}$, zval = -11.89) suggesting that theta sequences predict future choices. The effect was greater for non-VTE laps (Wilcoxon rank sum; $p = 5.1 \times 10^{-34}$, zval = -12.16; Figure 8.8A).

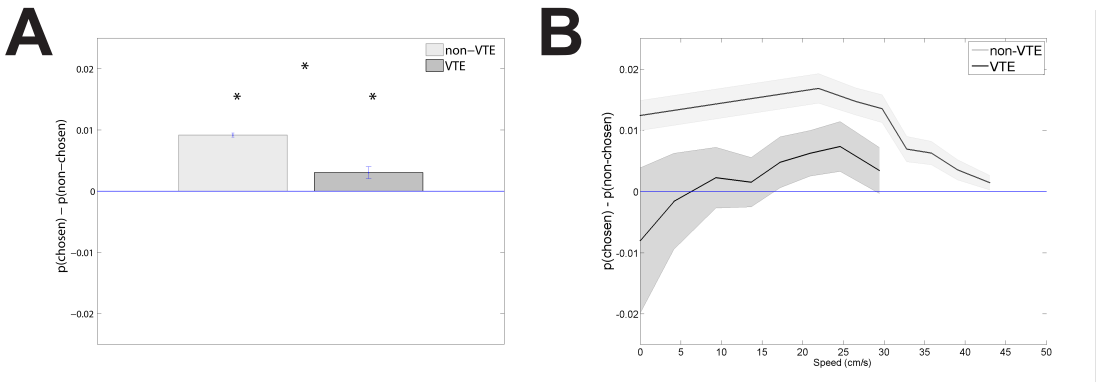

Figure 8.8: **Decoding During Theta Sequences represents future choice.** The difference in representation between chosen and non-chosen sides is compared for VTE and non-VTE laps. For this analysis, larger-later and smaller-sooner choices are combined into one group. A) The average representation of the difference favors the chosen side for both non-VTE (light gray) and VTE (dark gray). The representation during non-VTE is also significantly greater than for VTE. B) The difference in representation between chosen and non-chosen sides is speed-modulated. For VTE (dark gray), representation of the chosen side increases with speed from 0 cm/s to 30 cm/s, while for non-VTE (light gray) it remains constant over this range, decreasing at speeds exceeding those seen during VTE laps.

8.3.2 Forward representation during theta sequences was higher during VTE

The previous analysis examined average log-odds ratios, an approach that would mask independent representations of each side. While forward representation was not predictive of future choice during VTE, deliberation theory hypothesizes that both potential outcomes are represented at a choice point (van der Meer et al., 2012). I therefore compared the decoded probability of the four analysis zones separately, regardless of future choice, for VTE and non-VTE laps.

Decoded representation of the larger-later side and smaller-sooner side was higher during VTE (Wilcoxon rank sum; larger-later side, $p = 1.3 \times 10^{-30}$, zval = -11.5;

smaller-sooner side, $p = 4.6 \times 10^{-18}$, zval = 8.66). In contrast, decoded representation at the choice point was reduced during VTE (Wilcoxon rank sum; $p = 0.049$, zval = -1.97) and representation of the start of maze was unchanged (Wilcoxon rank sum; $p = 0.48$, zval = -0.71; Figure 8.9A). Importantly, representation of the remainder of the maze (the spatial bins **not** in any of the four analysis zones) decreased during VTE (Wilcoxon rank sum; $p = 0.49$, zval = -0.69). This suggested that decoding was focused toward goal locations during VTE and not simply distributed over the entire maze. Forward representation of future goal locations increased during VTE behavior.

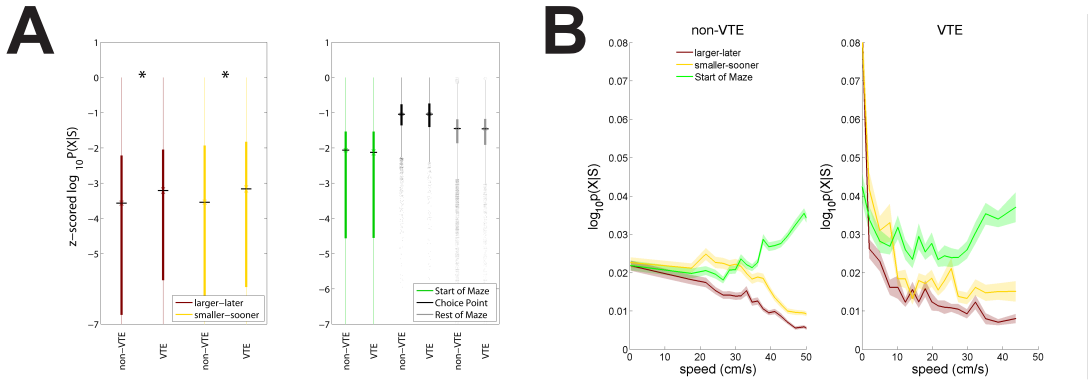

Figure 8.9: **Forward representation during Theta Sequences is higher during VTE.** The forward representation is compared separately for larger-later and smaller-sooner sides for both VTE and non-VTE. This analysis reveals that forward decoding is higher on VTE. A) Boxplots show the z-scored log-transformed forward (left panel) and other (right panel) decoding. Forward decoding to the larger-later side (maroon) and smaller-sooner side (gold) is higher during VTE. Decoding to other regions of the maze is not changed significantly during VTE. B) Representation during theta sequences is speed-modulated. Log-transformed average decoding for the larger-later (maroon), smaller-sooner (gold), and start of maze (green) zones is shown for non-VTE (left) and VTE (right). For non-VTE, decoding to the start of maze increases with speed and forward decoding decreases with speed. For VTE, forward decoding is an exponentially decreasing function of speed. In comparison, decoding to the start of maze is fairly constant.

8.3.3 Forward decoding to non-chosen side increased during VTE

Deliberation theory predicts that both potential choices are represented at a choice point during VTE (van der Meer et al., 2012). During VTE, decoded representation to the non-chosen side increased (Wilcoxon rank sum; larger-later side, $p = 3.5 \times 10^{-60}$,

zval = 16.36; smaller-sooner side, $p = 6.1 \times 10^{-27}$, zval = 10.75), confirming this prediction. Decoded representation to the chosen side was higher overall, but was increased for non-VTE laps relative to VTE laps, complimenting the above log-odds analysis (Figure 8.8A). On a choice to the smaller-sooner side, decoding at the choice point increased to the larger-later side. On a choice to the larger-later side, decoding at the choice point increased to the smaller-sooner side. This suggested that CA1 was representing both potential future paths (Figure 8.10A).

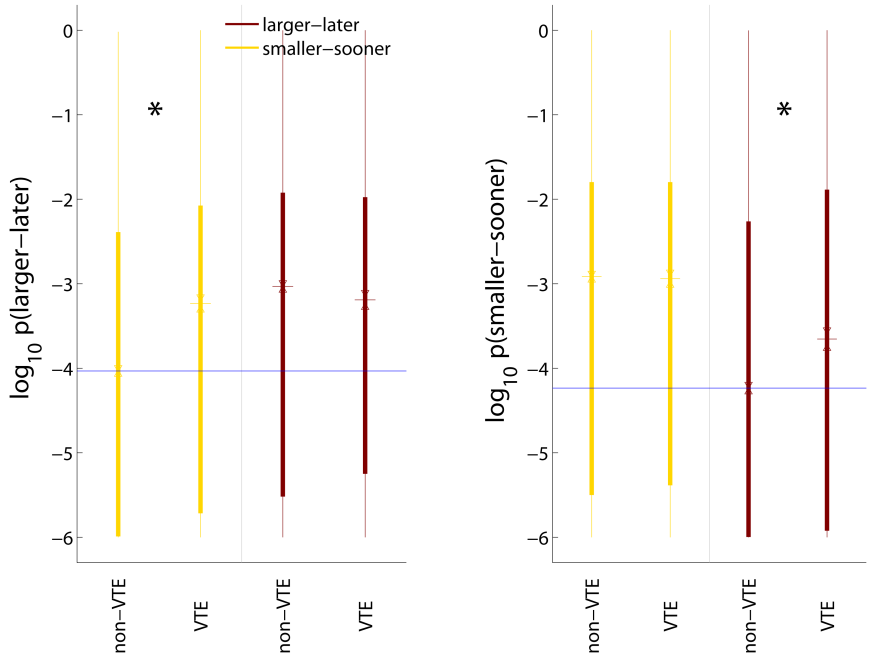

Figure 8.10: **Forward decoding to non-chosen side increases during VTE.** Decoding to the non-chosen side increases during VTE. Log-transformed decoding is compared for the larger-later and smaller-sooner sides during VTE and non-VTE laps. A) Boxplots shows larger-later (maroon) and smaller-sooner (gold) side decoding for choices to the larger-later side (left) and smaller-sooner side (right). For choice to the larger-later side, decoding to the smaller-sooner side increases during VTE. In complimentary fashion, for choice to the smaller-sooner side, decoding to the larger-later side increases during VTE. Note that decoding to the chosen side (maroon on the left, gold on the right) is higher overall for both choices, consistent with previous findings.

8.3.4 Forward representation was serial

The previous analyses reveal that representation of both feeders increased during VTE. Deliberation theory predicts that the representation to each side occurs in a serial fashion, so that the representation to the larger-later side occurs at a different time than the representation to the smaller-sooner side (van der Meer et al., 2012). To test this prediction, I plotted the two-dimensional histogram of the feeder representation and computed the correlation coefficient between delay and non-delay sides. If the representations of the larger-later and smaller-sooner sides were orthogonal, then they were being represented in a serial fashion. I found that this was the case, with representations clustering along the axes. The correlation coefficient analysis yielded a significant correlation for non-VTE (R=0.02, $p = 1.9 \times 10^{-5}$) and no correlation for VTE (R=0.018, p=0.09). During VTE, theta sequence representations were serial. (Figure 8.11).

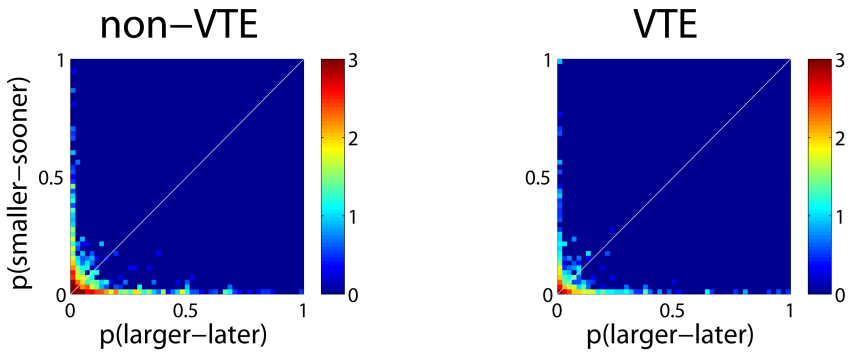

Figure 8.11: **Forward representation is serial.** Representations of the larger-later side is mutually exclusive to representation of the smaller-sooner side. A two-dimensional histogram is shown of representation on the smaller-sooner versus larger-later side. Representations along the axes indicate that decoding is serial, or mutually exclusive. Representations along the diagonal indicate parallel or joint representation of both sides. A) For both non-VTE and VTE laps, representation tends to be in one direction or the other.

8.4 Forward representation was speed modulated

Network states in the hippocampus are speed-modulated. For example, theta (Whishaw and Vanderwolf, 1973) and gamma (Ahmed and Mehta, 2012) frequency increase with running speed, while SWR occur during pausing.[1] VTE is inversely correlated with average speed, and there is a wider range of speed on VTE laps as compared to non-VTE laps. Therefore, I investigated the relationship between forward decoding and speed.

I found that representation of feeder locations was speed-modulated. On VTE laps, representation of the feeders was greatest at low velocities (< 2 cm/s) but decreased at high velocities (ANOVA; larger-later side, maroon line, $p = 5.0 \times 10^{-78}$, df = 10, F = 40.04; smaller-sooner side, gold line, $p = 1.7 \times 10^{-58}$, df = 10, F = 30.47). On non-VTE laps, representation of feeders diverged at high velocities (ANOVA; larger-later side, maroon line, $p = 6.3 \times 10^{-88}$, df = 10, F = 44.07; smaller-sooner side, gold line, $p = 5.5 \times 10^{-70}$, df = 10, F = 35.55; Figure 8.9B). The representation of the future choice was speed-modulated on non-VTE laps. Representation of future choice decreased with speed (ANOVA; light gray line, $p = 3.7 \times 10^{-60}$, df = 10, F = 30.88). Representation did not vary with speed on VTE laps (ANOVA; dark gray line, $p = 0.2$, df = 10, F = 1.34).

8.5 SWR and theta sequences contained similar information on a lap-by-lap basis

Two mechanisms to explain hippocampal involvement in planning have been proposed, SWR sequences and theta sequences, but their interactions have never been examined. To investigate the relationship between SWR sequences and theta sequences on a lap-by-lap basis, I plotted the two-dimensional histogram of feeder representations during theta and SWR. On average, SWR sequence representations were an order of magnitude stronger than theta sequence representations. SWR and theta sequence representations of the larger-later and smaller-sooner side were similar. For SWR sequences where the decoding to the larger-later side was high, theta sequences had high decoding to the larger-later side (Figure 8.12A, top left). For SWR sequences where the

[1] Since position is tracked from the head, and SWR are observed during grooming and pauses in running which still involve head movements, there may not be a strict relationship between head speed and SWR.

decoding to the smaller-sooner side was high, theta sequences also had high decoding to the smaller-sooner side (Figure 8.12A, top right). Additionally, SWR and theta sequence representations to opposite sides were anti-correlated. If the SWR sequence representation of the smaller-sooner (larger-later) side was high (Figure 8.12A, bottom left), theta sequence representation of the smaller-sooner (larger-later) side was low (Figure 8.12A, bottom right).

To quantify this relationship, I binned the decoding during SWR sequences into two bins. Same side representations increased for the larger-later side (t-test; $p = 2.38 \times 10^{-9}$, df = 3033, tstat = -5.99) and for the smaller-sooner side (t-test; $p = 1.08 \times 10^{-10}$, df = 3033, tstat = -6.5). Opposite side representations decrease for the delay to smaller-sooner side (t-test; $p = 8.12 \times 10^{-15}$, df = 3033, tstat = 7.8) and smaller-sooner to larger-later side (t-test; $p = 2.61 \times 10^{-5}$, df = 3033, tstat = 4.21). (Figure 8.12B).

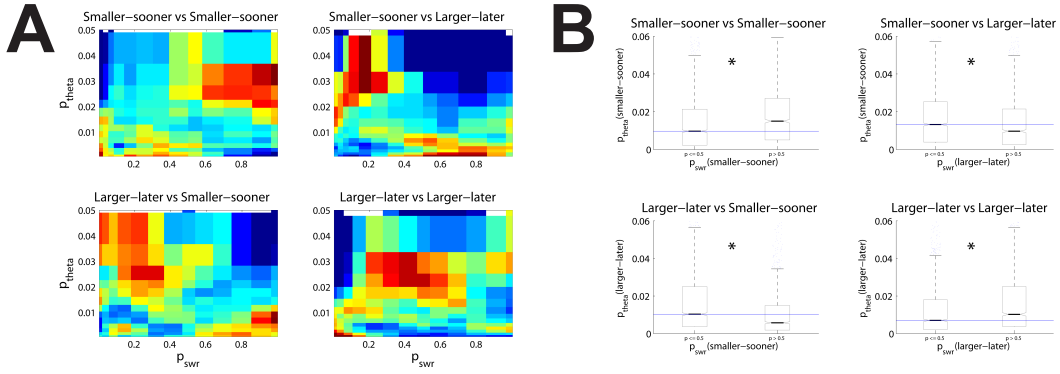

Figure 8.12: **SWR and theta sequences contain similar information on a lap-by-lap basis.** SWR and theta sequences on the same lap have related information content. A) A two-dimensional histogram is computed comparing the average decoding for all SWR sequences on a lap versus that for choice point theta sequences. On a lap-by-lap basis, representations are related. Decoding to the smaller-sooner (theta sequence) versus smaller-sooner (SWR sequence) side is positively correlated (top left). In compliment, decoding to the larger-later (theta sequence) versus larger-later (SWR sequence) side is positively correlated. Decoding in either combination of opposite sides (top right, bottom left) is anti-correlated. B) The histograms are binned into SWR representation $<= 0.5$ or > 0.5 for statistical analysis. Both same side representations significantly increase and both opposite side representations significantly decrease.

8.6 Local Field Potentials

LFP oscillations reflect the underlying network state of a structure and may reveal inter-regional communication with other structures (See **Section** 4.3). I examined the timing and spatial distribution of gamma and high frequency oscillation (HFO) LFP events, which are linked to different processing modes in CA1.

8.6.1 Gamma increases during Investigation

The power of 30-55 Hz low and 56-100 Hz high gamma was higher during the Investigation phase in all zones (ANOVA; low gamma, $p = 7.7 \times 10^{-55}$, df = 2, F = 134.54; high gamma, $p = 1 \times 10^{-100}$, df = 2, F = 265.65). At the choice point and start of maze, low gamma (two-way ANOVA; $p = 1.4 \times 10^{-12}$, df = 3, F = 19.73) and high gamma (two-way ANOVA; $p = 2.7 \times 10^{-6}$, df = 3, F = 9.61) were increased, and both decreased across behavioral phase (two-way ANOVA; low gamma, $p = 0.0052$, df = 6, F = 3.09; Figure 8.15; high gamma $p = 0.0001$, df = 6, F = 4.54; Figure 8.14). The power of 100-150 Hz HFOs was higher during Investigation (ANOVA; $p = 2.7 \times 10^{-9}$, df = 2, F = 19.98), and tended to be higher in the start of maze and choice point (two-way ANOVA; $p = 4.5 \times 10^{-5}$, df = 3, F = 7.64) with an decrease across phase (two-way ANOVA; $p = 0.01$, df = 6, F = 2.57; Figure 8.13). Overall, all LFP gamma frequencies increased during the Investigation phase, when the rat was learning the unknown parameters in a known environment.[2] Additionally, low and high gamma were higher at the choice point.

8.6.2 High and Low gamma increased at the choice point during VTE

Prior research has found that 40-120 Hz gamma increased at the central stem on a delayed alternation task. This increase was accompanied by increased CA1-CA3 gamma coherence, which could not be explained by differences in motor behavior (Montgomery and Buzsáki, 2007). 65-120 Hz gamma increased at the choice point during a delayed non-match to place task, with a delayed synchrony between CA1 and medial EC during VTE laps (Yamamoto et al., 2014). Therefore, I examined gamma power

[2]For a discussion of this type of learning in comparison with other learning processes, see **Section** 3.3.

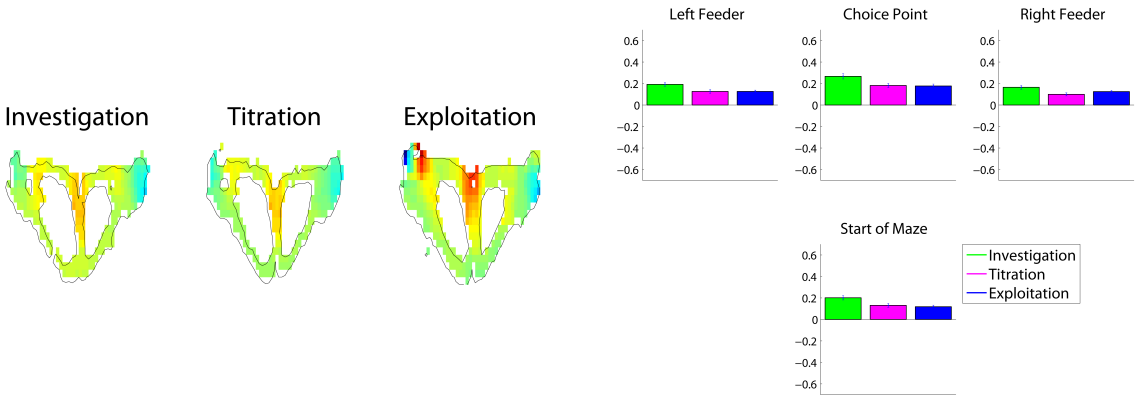

Figure 8.13: **HFO Spatial Distribution for Different Behavioral Phases.** Local field potentials (LFPs) display different spatiotemporal dynamics on the spatial delay discounting task. (left panels) Two-dimensional histograms are shown for each analyzed LFP frequency band in a 56×56 position grid averaged within-session, across-session, then across-rat. Color range is from -0.05 (blue) to 0.05 (red). (right panels) z-scored power is averaged within-session, then across-session for each analysis zone. HFOs are uniformly distributed across the maze. HFOs are higher during Investigation for each analysis zone.

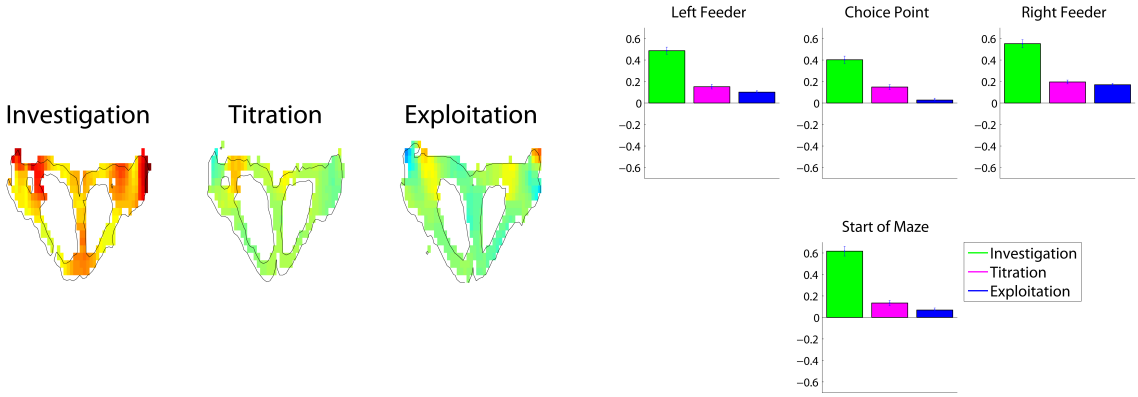

Figure 8.14: **High Gamma Spatial Distribution for Different Behavioral Phases.** Local field potentials (LFPs) display different spatiotemporal dynamics on the spatial delay discounting task. (left panels) Two-dimensional histograms are shown for each analyzed LFP frequency band in a 56×56 position grid averaged within-session, across-session, then across-rat. Color range is from -0.05 (blue) to 0.05 (red). (right panels) z-scored power is averaged within-session, then across-session for each analysis zone. Overall, high gamma power is the strongest of the four LFP bands analyzed. High gamma power is higher during Investigation for each analysis zone. In addition, it decreases from Titration to Exploitation at the choice point and start of maze.

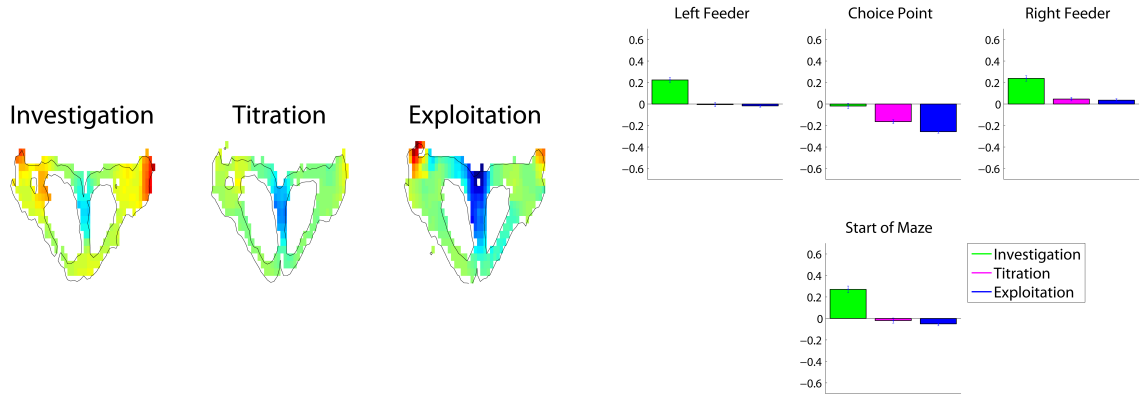

Figure 8.15: **Low Gamma Spatial Distribution for Different Behavioral Phases.** Local field potentials (LFPs) display different spatiotemporal dynamics on the spatial delay discounting task. (left panels) Two-dimensional histograms are shown for each analyzed LFP frequency band in a 56×56 position grid averaged within-session, across-session, then across-rat. Color range is from -0.05 (blue) to 0.05 (red). (right panels) z-scored power is averaged within-session, then across-session for each analysis zone. Low gamma is higher during the Investigation phase for each analysis zone. Low gamma decreases with phase in the choice point and start of maze zones. Low gamma power at the choice point is lower than elsewhere.

separately for VTE and non-VTE laps. During VTE, the power of high and low gamma increased (Wilcoxon rank sum; low gamma, $p = 0.028$, zval = 2.19; high gamma, $p = 1.8 \times 10^{-4}$, zval = 3.74). This effect was clearest at the choice point (Figure 8.17; Figure 8.18). In contrast, HFO power did not change during VTE laps (Wilcoxon rank sum; $p = 0.47$, zval = -717; Figure 8.16).

8.6.3 The timing of gamma is synchronized around VTE events

The previous analysis found that gamma power was increased during VTE laps. Prior research has found that synchronization of CA1 and medial EC gamma at the choice point of a delayed non-match to place task was necessary for correct performance. During VTE, the synchronization was delayed (Yamamoto et al., 2014).[3] Additionally, van der Meer and Redish (2009) and Stott and Redish (2014) found increased reward decoding in ventral striatum before the turn around time of a VTE event and Steiner

[3]This desynchronization provides a mechanistic explanation for behavioral impairments observed from stimulation of the perforant path connection McNaughton et al. (1986).

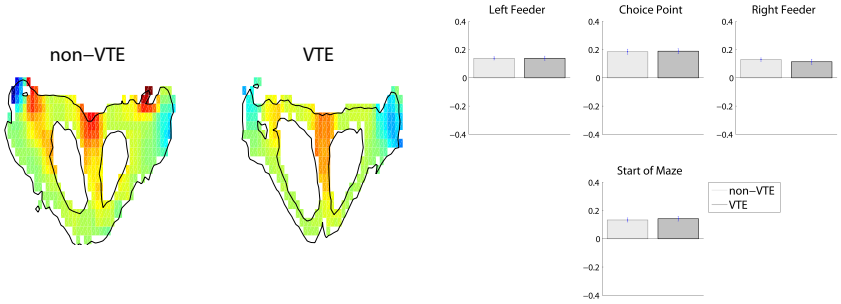

Figure 8.16: **HFO Spatial Distribution for non-VTE versus VTE.** Local field potentials (LFPs) are analyzed for non-VTE and VTE laps (left panels) Two-dimensional histograms are shown for each analyzed LFP frequency band in a 56×56 position grid averaged within-session, across-session, then across-rat. Color range is from -0.05 (blue) to 0.05 (red). (right panels) z-scored power is averaged within-session, then across-session for each analysis zone. HFOs increase at the choice point, but do not vary by VTE.

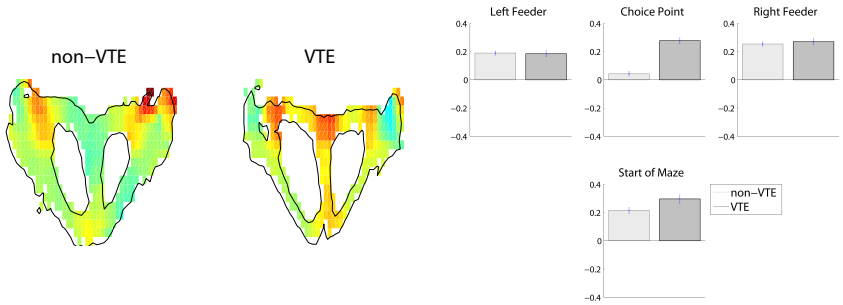

Figure 8.17: **High Gamma Spatial Distribution for non-VTE versus VTE.** Local field potentials (LFPs) are analyzed for non-VTE and VTE laps (left panels) Two-dimensional histograms are shown for each analyzed LFP frequency band in a 56×56 position grid averaged within-session, across-session, then across-rat. Color range is from -0.05 (blue) to 0.05 (red). (right panels) z-scored power is averaged within-session, then across-session for each analysis zone. High gamma increases at the choice point during VTE. High gamma does not vary for any other analysis zone.

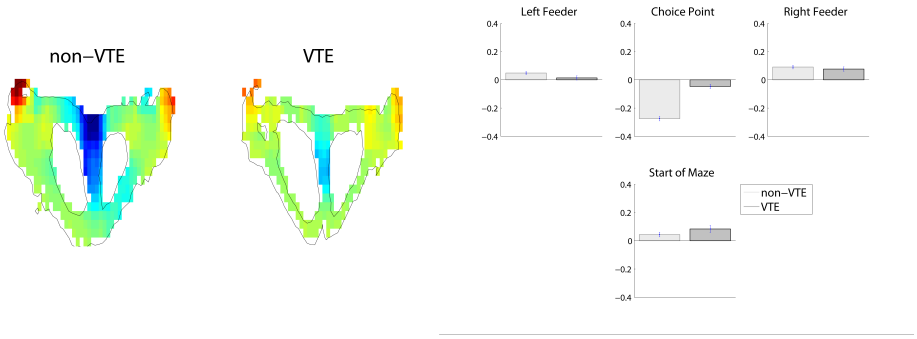

Figure 8.18: **Low Gamma Spatial Distribution for non-VTE versus VTE.** Local field potentials (LFPs) are analyzed for non-VTE and VTE laps (left panels) Two-dimensional histograms are shown for each analyzed LFP frequency band in a 56×56 position grid averaged within-session, across-session, then across-rat. Color range is from -0.05 (blue) to 0.05 (red). (right panels) z-scored power is averaged within-session, then across-session for each analysis zone. Low gamma increases at the choice point during VTE. On average, low gamma power is lowest at the choice point.

and Redish (2012) and Stott and Redish (2014) found increased reward decoding in orbitofrontal cortex after the turn around time. These structures are part of the deliberative network, discussed in **Section** 2. These results suggest that high gamma in CA1 is related to deliberation of upcoming rewards. To investigate the timing of gamma during VTE, I performed an LFP peri-event time histogram aligned to VTE turn around events. I found that 30-50 Hz low gamma (Wilcoxon rank sum; $p < 2.4 \times 10^{-18}$, zval = -8.74) and 50-100Hz high gamma (Wilcoxon rank sum; $p < 2.3 \times 10^{-21}$, zval = -9) increased after the turn-around time of a VTE event (Figure 8.19B,D). In addition, 100-150 Hz HFOs increased after the turn-around time (Wilcoxon rank sum, $p = 1.8 \times 10^{-3}$, zval = -3.11). SWR power did not change (Wilcoxon rank sum; $p = 0.07$, zval = 1.77). On a VTE lap, the turn around time was often the time of minimum velocity. Therefore as a control, I computed the LFP peri-event time histogram around the minimum velocity of non-VTE laps. No differences were seen before versus after the minimum velocity for low gamma (Wilcoxon rank sum; $p = 0.07$, zval = 1.80), high gamma (Wilcoxon rank sum; $p = 0.2$, zval = 1.22), or HFO ($p = 0.056$, zval = 1.91). However, an increase in SWR power was observed after the minimum velocity ($p = 6.7 \times 10^{-5}$, zval = 3.98; Figure 8.19A,C).

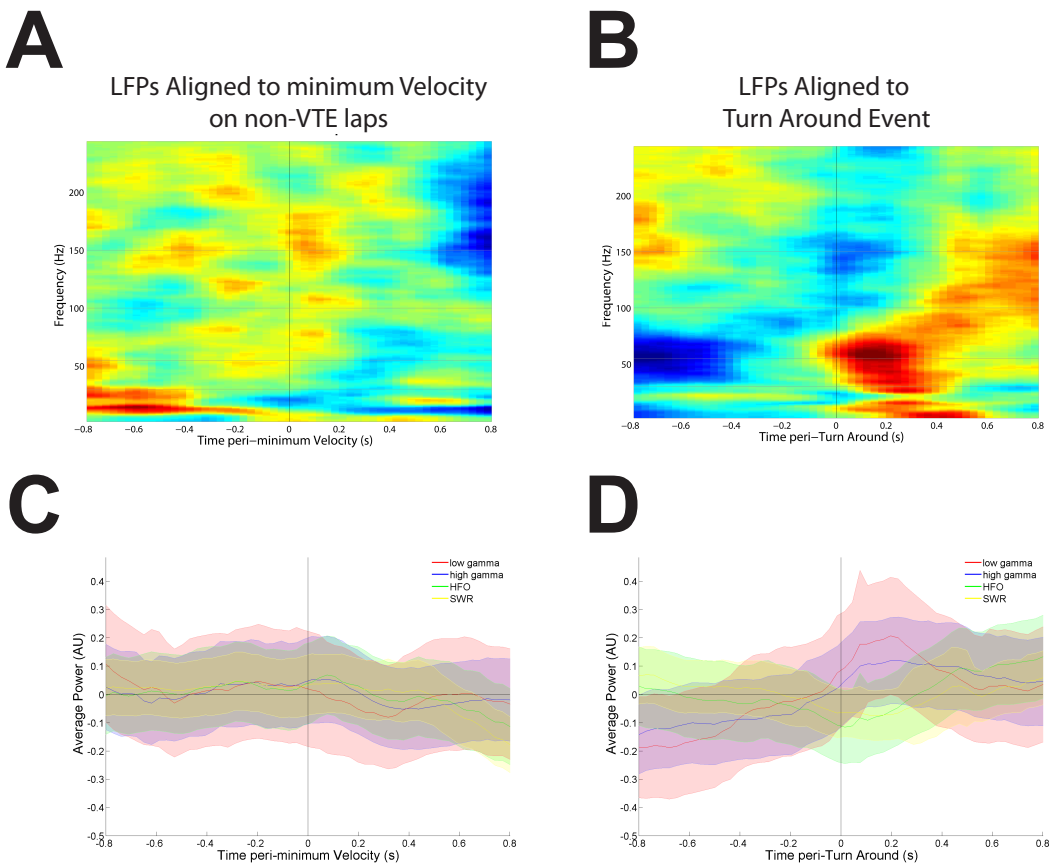

Figure 8.19: **Gamma timing is synchronized around VTE events.** LFP spectral power is aligned to the turn-around time of VTE trials or the minimum velocity of non-VTE trials. A) Aligning LFP to the minimum velocity of non-VTE trials, no clear changes in gamma or SWR power are visible. A decrease in theta power is visible during deceleration to minimum velocity. B) Aligned to the turn-around time of VTE trials, low and high gamma increase following the turn around time. As the rat accelerates away from the turn-around, an increase in gamma frequency is also observed. C,D) Spectral power is averaged within frequency bands: 30-50Hz for low gamma, 50-100 Hz for high gamma, 100-150 Hz for HFOs, and 150-250 Hz for SWR. An increase in gamma power is visible following turn-around times for VTE laps. Shading denotes standard error. Significant increases are observed for low gamma, high gamma, and high frequency oscillations after the turn around time. Except for an increase in SWR power, no significant changes are observed for non-VTE laps aligned to minimum velocity.

In contrast with the motor-independent effects seen in Montgomery and Buzsáki (2007) and Yamamoto et al. (2014), Ahmed and Mehta (2012) found a strong correlation between 30-120 Hz gamma and running speed. The authors suggested this

was a means for achieving speed-independent scaling of theta sequence representations. Kemere et al. (2013) found that a correlation between gamma power and running speed in CA1 was experience-dependent. The authors suggested that changes in functional coupling based on speed allowed hippocampus to function differently at different times. At slow speeds, CA3 input dominated CA1 and the hippocampus was geared toward consolidation, while at high speeds CA1 was geared to more accurately represent position. In contrast to both of these theories, Maurer et al. (2012) found that theta sequences scaled with running speed, so the faster a rat was running, the further the look-ahead distance. Nevertheless, as with theta, there appeared to be speed-dependent modulation of gamma frequency. To examine this effect, I binned spectral power by running speed across the entire maze (Figure 8.20A), and separately for each of the behavioral phases (Figure 8.20B). There was a correlation between running speed and gamma frequency. Future research may wish to examine this relationship more closely.

8.6.4 Gamma frequency increases with speed

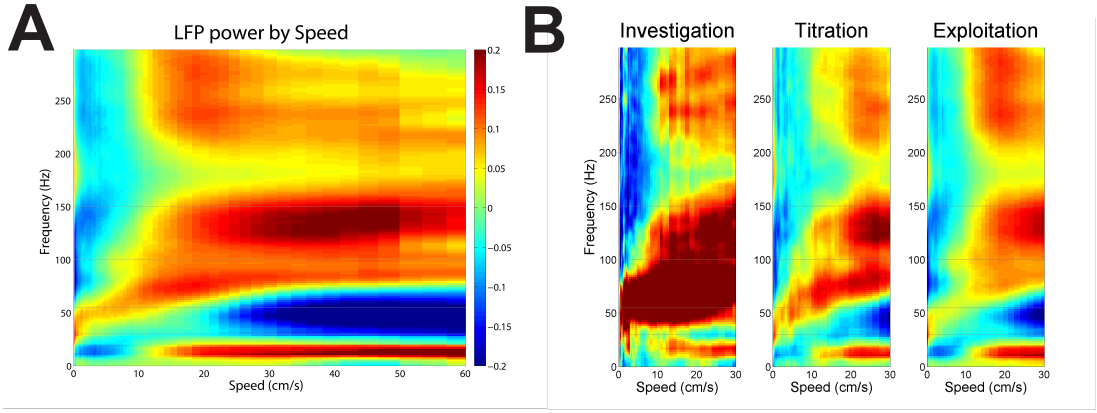

Figure 8.20: **Gamma increases with running speed.** A) Spectral power is binned by running speed throughout the maze. From 0 cm/s to about 30 cm/s average gamma frequency increases from about 25 Hz to 100 Hz. Note that at low running speed, 150-250 Hz SWR power is increased. The increase in spectral power above 200 Hz at high running speeds likely reflects spike or movement-related artifacts. Note that an increase in 100-150 Hz HFOs is also seen above 20 cm/s. Spectral power is z-scored within session and is shown from -0.2 to 0.2 in arbitrary units. B) Spectral power is binned by running speed separately for each behavioral phase. Note the increased gamma power during the Investigation phase, consistent with previous findings. The gamma frequency versus running speed relationship may vary by behavioral phase, though additional research is necessary to determine this. Average spectral power is shown from -0.2 to 0.2 in arbitrary units.

8.7 Conclusions

The results in this section provide evidence for engagement of hippocampus during decision-making on the spatial delay discounting task. Both theta sequence sweeps at a choice point and SWR sequences at feeder locations predict future choice, and their representations are correlated on a lap-by-lap basis. However, SWR and VTE tend to occur at different times. VTE occurs during Titration and SWR occurs during Exploitation. This suggests a complimentary interaction between the two phenomena, facilitating decision-making across different timescales between different decision-making systems.

The data presented here are compatible with the theory that theta sequences sweep forward to one side, then the other during VTE. Forward decoding during VTE occurs

when the rat is paused at the choice point, near the turn around time. Theta sequence sweeps during VTE are accompanied by increased gamma power aligned to the turn around time.

To gain a better understanding of how these sequences are involved in decision-making, it would be useful to perturb them without altering the fundamental hippocampal schemas upon which they are based. Robbe et al. (2006) found that the synthetic cannabinoid agonist CP55940 disrupts theta sequences while leaving place cell schemas intact. In this study, the firing rate of place cells did not change, but correlated firing between pairs of cells was diminished. Specifically, cannabinoids reduced phase precession and caused an increase in 'temporal jitter' in theta sequences (Robbe and Buzsaki, 2009).[4] Behaviorally, the drug impaired the ability of rats to perform a working memory task (delayed alternation). Cannabinoids therefore allow manipulation of theta sequences without affecting schemas. I predict that the diminished spike synchrony under cannabinoids eliminates ordered sequences during VTE, preventing forward decoding and search through potential outcomes. This effect may underlie the impaired decision-making under cannabinoids.

Cannabinoids also cause reduced gamma power (Robbe et al., 2006), and they may disrupt the synchrony between CA1 and EC that has been shown to facilitate decision-making (Yamamoto et al., 2014). On the spatial delay discounting task, I predict that cannabinoids prevent the increase in gamma power seen during VTE after the turn around time.

Cannabinoids cause reduced SWR power and frequency (Robbe et al., 2006), but their effect on SWR sequences is unknown. Given their disruption of theta sequences, I would predict that SWR sequences are also disrupted. They would also, presumably, slow consolidation and prevent transition from flexible to habitual decision-making strategies.

[4]Considering two place cells, the distance between peak firing rate in their place fields is correlated with the timing of spikes in a theta cycle (Geisler et al., 2007). This is called the *sequence compression index*. CP55940 disrupted the cross-correlogram of firing rates in overlapping place fields, reducing the sequence compression index. This index was found to be proportional to performance on the delayed alternation task (Robbe and Buzsaki, 2009).

Chapter 9

Cannabinoids

9.1 Background

The endocannabinoid system has functional roles in immune response (Pacher and Mechoulam, 2011), anti-inflammation (Mazzari et al., 1996; Richardson et al., 1998), anti-nociception (Yaksh, 1981; Martin et al., 1995; Lichtman et al., 1996; Malan et al., 2001) regulation of feeding behavior (Williams and Kirkham, 1999; Koch et al., 2015).[1], and regulation of stress and anxiety (Onaivi et al., 1990; Marco et al., 2004; Viveros et al., 2005; Rubino et al., 2008; Hill et al., 2011; Gray et al., 2015).[2] Additionally, cannabinoids modulate dopamine-releasing neurons in the ventral tegmental area (French et al., 1997) through local activation of CB_2 receptors (Zhang et al., 2014), and activation of CB_1 receptors in medial prefrontal cortex (Draycott et al., 2014) and hippocampus (Loureiro et al., 2015; Lisboa et al., 2015). This links cannabinoids to addiction and motivational processes (Serrano and Parsons, 2011).

Cannabinoids disrupt working memory in humans (Tinklenberg et al., 1970; Abel,

[1] Feeding behavior is mediated by regulation of hypothalamic pro-opiomelanocortin (POMC) neurons. To regulate feeding, CB_1 receptors may act to inhibit GABA release from Agouti-related protein neurons onto POMC neurons, also activating a molecular switch that increases production of β-endorphin in POMC neurons, causing selective activation of downstream μ-opioid receptors that promote feeding behavior (Koch et al., 2015; Patel and Cone, 2015).

[2] A mechanism for cortical feedback to terminate the stress response was proposed by Hill et al. (2011). Glucocorticoid increase in medial PFC from the hypothalamic-pituitary axis stress response causes a time-delayed increase in the production of an endocannabinoid (2-arachidonoylglycerol), which inhibits GABA transmission in prefrontal cortex, disinhibiting principal cells that may function to terminate the stress response (Hill et al., 2011).

1971; Miller et al., 1977; Ranganathan and Dsouza, 2006) and rats (Carlini et al., 1970; Heyser et al., 1993; Lichtman et al., 1995; Miyamoto et al., 1995). Genetic knockout of astrocytic CB_1 but not neuronal CB_1 receptors in mice causes impaired working memory on a delayed-match-to-place version of the Morris water maze and impaired LTD at CA3-CA1 synapses (Han et al., 2012). CB_1 receptors are also expressed on CCK containing interneurons which primarily target the soma and distal dendrites of pyramidal cells (Katona et al., 1999). Cannabinoid agonists also cause a reduction in the firing rate and synchronization of these interneurons (Katona et al., 1999; Holderith et al., 2011), which may be responsible for the decreased gamma power in CA1 (Robbe et al., 2006) and CA3 (Holderith et al., 2011). Intrahippocampal infusion of a cannabinoid agonist produces deficits in spatial working memory in the 8-arm radial maze (Lichtman et al., 1995) and decreases ripple ($150 - 250$ Hz) power in CA1 (Robbe et al., 2006). Cannabinoid disruption of working memory is likely mediated by CA1 astrocyte CB_1 activity in combination with network desynchronization caused by inhibition of fast spiking interneurons.

The effect of cannabinoid manipulations on delay discounting experiments are complex. The CB_1 antagonist and inverse agonist rimonabant has been found to increase discounting (Boomhower and Rasmussen, 2014), or have no effect on discounting (Pattij et al., 2007). CB_1 agonists have been found to have no effect on discounting in rats (Pattij et al., 2007) or humans (McDonald et al., 2003). I did not make any predictions about the effect of cannabinoids on the indifference point, but hypothesized that they would reduce titration efficiency (defined in **Section** 7.7) by interfering with working memory.

Cannabinoids modulate motor responses (Consroe et al., 1986), and running speed after administration of cannabinoid agonists is generally slower on runway tasks (Robbe et al., 2006). VTE has an inverse relationship with running speed (**Section** 8.4). The Robbe et al. (2006) study showed what appeared to be increased VTE at the choice point of the delayed alternation task, although the effect was not quantified. These observations led me to the hypothesis that cannabinoid agonists would increase VTE.

9.2 Cannabinoids increase VTE

VTE was quantified at the choice point using log IdPhi z-scored within rat. Across all sessions, VTE was significantly higher with cannabinoids compared with saline control (Mann-Whitney U; $p = 2.2 \times 10^{-7}$, z = 5.18). Papale et al. (2012) found that VTE was higher during Adjustment laps on the spatial delay discounting task (**Section** 7.5A). Therefore, I investigated if this effect was true with cannabinoids. I found that VTE was higher during Adjustment laps for saline sessions, confirming our previous findings (Mann-Whitney U; $p = 0.04$, z = 2.04). Additionally, cannabinoids increased VTE to a much greater extent on Adjustment laps than Alternation laps (Mann-Whitney U; $p = 4.6 \times 10^{-4}$, z = 3.5; Figure 9.1B,C).

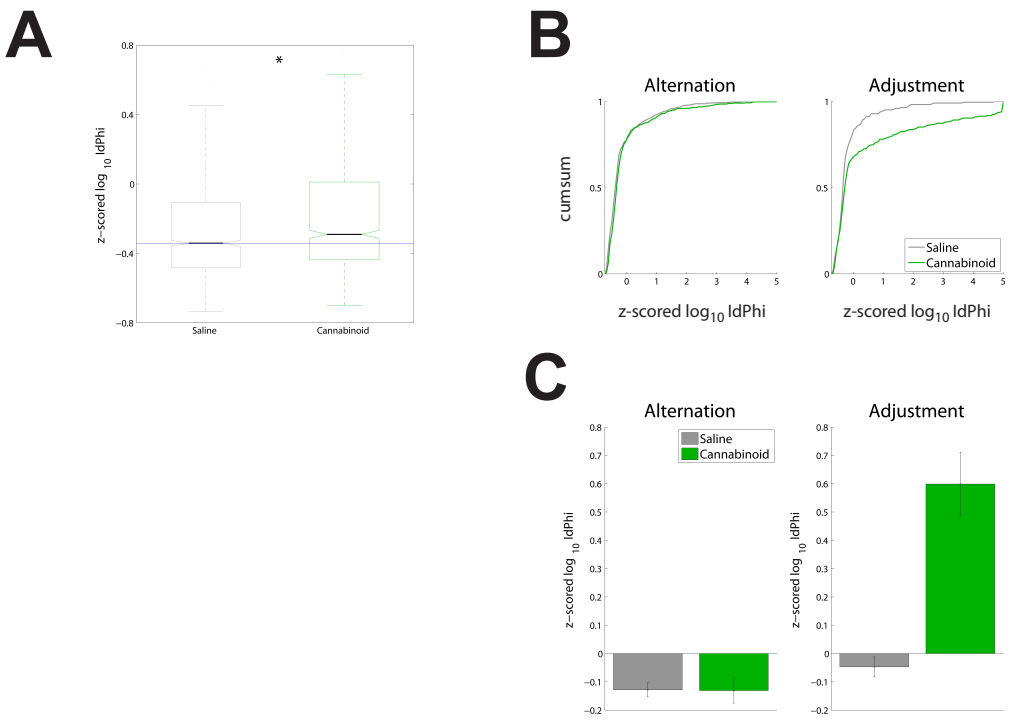

Figure 9.1: **Cannabinoids increase VTE.** VTE is measured at the choice point using the log-transformed IdPhi measure and then z-scored by rat. Over all laps, VTE is higher during cannabinoid sessions. The increase in VTE during cannabinoids is restricted to Adjustment laps. A) A boxplot shows the median VTE measure for saline (gray) and cannabinoid (green) groups. Median VTE is higher in the cannabinoid group. B) Normalized cumulative sum of the histograms for z-scored log-transformed IdPhi distributions during Alternation (left) and Adjustment (right) lap types reveals a clear increase in VTE during Adjustment laps in the cannabinoid group. C) Average VTE for Alternation (left) and Adjustment (right) laps confirms the increase in VTE during Adjustment laps for the cannabinoid group.

9.3 Conclusions

Cannabinoids provide a selective way to disrupt theta sequences without changing hippocampal schemas. The data from this section show that cannabinoids increase VTE during Adjustment laps, suggesting that cannabinoids impair deliberative decision-making. Preliminary data (not shown) also suggests that forward decoding during

VTE is diminished by cannabinoids. Future research may confirm this effect, implicating it in the spatial working memory deficit caused by the drug. The effect of cannabinoids on SWR sequences is unknown, though it decreases both SWR power and frequency. Future research may focus on investigating SWR sequence composition, as well as the effect of cannabinoids on consolidation and schema learning. Additionally, the effects of cannabinoids on the dynamic aspects of hippocampal coding (overdispersion, remapping, trajectory-dependent firing, etc.) have not been examined. Systemic administration of cannabinoid agonists may prove useful for investigation of network properties of deliberative decision-making.

On the other hand, limitations in systemic administration of the drug are its effects on a diverse array of central and peripheral systems. It can not be said for certain that behavioral effects are due to the action of the drug in the hippocampus. Optogenetic tools offer the promise of a more selective manipulation of theta sequences.

Chapter 10

Discussion and Conclusion

10.1 Deliberation

The deliberative decision-making system is characterized by a cognitive search mechanism. During some difficult decisions, the hippocampus engages in a search process through memory schemas. This generates potential future outcomes, to be evaluated downstream. The temporal evolution of cell assemblies in hippocampal ensembles may account for the dynamic 'sweeps' of SWR sequences and theta sequences. The deliberation 'subroutine' may be called once or repeatedly during decision-making. While sweeps were originally discovered in the CA3 region of hippocampus, they have also been characterized in the CA1 region, and have been hypothesized to occur in the entorhinal cortex. The relationship between sweeps in CA1 and CA3 is unknown, and future research may focus on answering this question. Deliberation has been constrained to include search through hippocampal schemas (van der Meer et al., 2012), but I speculate that it is a more dynamic process, involving different combinations of brain regions depending on task demands.

10.2 Vicarious Trial and Error

VTE is related to the deliberative system, though the relationship is complex. Alternative theories fail to explain the timing of this seemingly simple behavior. If the definition of deliberation is constrained to include search through hippocampal schemas

(van der Meer et al., 2012), then VTE behavior alone is not sufficient to indicate a deliberative decision-making strategy. Deliberation occurs on a spatial task in the absence of VTE (Wikenheiser and Redish, 2015), and VTE occurs on spatial tasks in the absence of hippocampus (Bett et al., 2014). However, its occurrence on a binary choice task does strongly correlate with deliberation (Johnson and Redish, 2007). I hypothesize that VTE on hippocampal-dependent binary choice tasks (e.g. delayed alternation) would more strongly dissociate deliberative from non-deliberative strategies. If deliberation is more broadly defined as a search through cognitive schemas, then VTE may be more strongly coupled with deliberation in its many forms. On the other hand, a more restrictive definition of VTE may be appropriate, as 'VTE' during visual discrimination may be completely unrelated to VTE during a spatial working memory task. VTE in rats is related to SFS sequences in primates. I hypothesize that stereotypical behavioral gestures in humans are also homologues to VTE in rodents.

10.3 Spatial Delay Discounting Task

The delay discounting concept has advanced the study of choice behavior. The ability to delay gratification is related to real-world issues like child development and addiction. The spatial delay discounting task is a free-operant open runway experiment used to study discounting in rats. Because rats prefer to alternate given a binary spatial choice between two equal-value options on a T-maze, the spatial delay discounting task may produce more stable indifference points than the classic adjusting delay procedure. Rats titrate the adjusting delay in three distinct phases, consistent with behavioral predictions. Future research may wish to focus on testing aspects of the hyperbolic discounting model on the spatial delay discounting task. For example, the model predicts that decisions will be very sensitive to increases in the delay of the smaller-sooner option. Increasing the delay before the smaller-sooner option should cause a multiplicative increase in the indifference point.[1] Additional experiments may

[1]Preliminary evidence suggests this is not the case, at least in overtrained rats. I performed a pilot experiment where the delay to the smaller-sooner option was increased within-session. A small increase of 5s did not appreciably alter behavior whether the change occurred during the Titration or Exploitation phase. A large increase of 10s caused only a slight increase in the indifference point (Papale, A.E. *Unpublished observation*).

wish to vary the inter-trial interval, or the variability in travel distance of the navigation circuit. Both of these manipulations would place heavier demand on hippocampal processing, perhaps increasing the occurrence of deliberation and the link between deliberation and VTE. These experiments may be able to dissociate between competing decision-making models of discounting behavior.

VTE on the spatial delay discounting task occurs during Adjustment laps, when behavioral flexibility is at a premium. VTE increases during isolated Adjustment laps as compared to runs of consecutive Adjustment laps. VTE decreases with session while alternation efficiency increases, suggesting refinement of titration strategies with experience. Hippocampal lesions on the spatial delay discounting task produce a prolonged Investigation phase (Bett et al., 2014), suggesting a role for hippocampus in increasing the efficiency to reach the indifference point.

10.4 Sharp Waves Sequences and Theta Sequences

I found that SWR frequency increases across-session and during non-VTE Exploitation laps. This is consistent with the role of SWR in consolidation of episodic memories. Hippocampus may be disengaged from decision-making during non-VTE and Exploitation, and consolidation may be more efficient when the deliberative system is offline. On the other hand, awake SWR predict future trajectories during goal-directed tasks (Pfeiffer and Foster, 2013) and impact learning of a spatial working memory task (Jadhav et al., 2012). This suggests a potential role for SWR in deliberation. How can this redundancy be explained? How are the processes of deliberation and consolidation linked together?

One possibility is that there are two different ways to plan future trajectories by querying hippocampus. During active movement, an online planning mechanism might be bound to representations of options that are either being approached or departed. In support of this, theta sequences 'chunk' an environment into pieces based on landmarks (Gupta et al., 2010) and look-ahead distance is predictive of upcoming choice (Wikenheiser and Redish, 2015). Prospective and retrospective encoding have been linked with performance of locale navigation on a plus-maze (Ferbinteanu and

Shapiro, 2003).[2] Perhaps only the active cell-assembly can be searched during theta sequence sweeps; or, map-switching (Jackson et al., 2006; Jezek et al., 2011) on alternating theta cycles may allow sweeps to 'jump' across cell assemblies linked by a common topology. In contrast, more complex situations may be considered during offline planning, including never-experienced routes (Gupta et al., 2012), and imagining trajectories not bound to the current location (Karlsson and Frank, 2009). Plasticity during SWR may allow for dynamic search across cell assemblies that is not possible during active movement.

Another possibility is that the process of consolidation may directly bias trajectory planning. The trajectory that is being consolidated at the time of journey departure may be more likely to influence the upcoming decision. Correlations between theta sequences and SWR sequences have been observed during awake SWR (Jackson et al., 2006) and during sleep (O'Neill et al., 2008) and I found that SWR sequences contain similar representations to the theta sequences on the previous pass through the choice point (**Section** 8.5). The majority of SWR represented these recently activated place cells (See **Section** 8.2.1), consistent with consolidation of recent experience. Similar results were obtained for SWR representation during exploration of novel environments (O'Neill et al., 2006). A subset of SWR also represented the *opposite* of the future choice (See **Section** 8.2.2). How can this finding be reconciled with the findings of (Pfeiffer and Foster, 2013) where SWR representations predict upcoming trajectories? It is possible that on a binary choice task predictive SWR indicate the direction *not* to go. Similar results were obtained on a two-T reversal task (Gupta et al., 2010). Since the representation of recent experience is strongest on the spatial delay discounting task, and the majority of laps were Alternation laps, it is possible that trajectory planning of the opposite side reflects the peculiar tendency of rats to spatially alternate on binary choice tasks. To test this idea, SWR sequences could be measured on a binary choice plus-maze where the third option is blocked, then opened during testing. If the binary choice task is trained first, a specific prediction would be that there are 1) behaviorally more errors in the plus-maze that are a result of alternation to the opposite side from the side of departure, and, 2) that these errors would be preceded by representation of SWR to the side of departure. The way a problem is analyzed in real time may

[2]VTE has similarly been linked with locale navigation (Schmidt et al., 2013a; Gardner et al., 2013).

fundamentally influence its encoding into long-term memory.

In seeming contrast to the above idea, Singer et al. (2013) found roughly equal chosen and non-chosen trajectories during SWR sequences on a W-shaped track. However, results could differ due to running of inbound and outbound journeys on the same topology. In support of this idea, Jackson (2006) examined correlated overdispersion among overlapping place fields on different tasks. He found that cell assemblies on the open field, but not a linear track, had correlated overdispersion. Anti-correlations between different cell assemblies on the linear track may have been overwhelming positive correlations within cell assemblies. The phenomenon of between-assembly dilution of within-assembly correlations may therefore also be present in SWR sequences. In support of this idea, (Wu and Foster, 2014) found that SWR sequences with rats running bidirectional trajectories on a Y-shaped track tended to reverse direction at the spatial choice point.

During VTE, potential future choices are represented serially. This replicates the results of Johnson and Redish (2007). Additionally, forward representations occur during reduced running speed, and potentially during synchronized gamma bursts following the turn around time of a VTE event. During non-VTE, the actual future choice is represented more strongly, consistent with Gupta et al. (2012) and Wikenheiser and Redish (2015). Future research may confirm the prediction that forward representations during VTE are actually due to alternating theta sequence sweeps. The data in the present study was not clean enough to confirm this prediction. Given this limitation, and the generally noisy quality of theta sequences, an important question to address in future research is, 'Are theta sequences coherent enough to facilitate inter-regional information transfer?' It is possible that this question will be resolved by recording of larger numbers of cells. Theoretically, given an estimate in the temporal noise inherent to theta sequences, it should be possible to compute the level of signal redundancy necessary to facilitate the communication of the trajectory to structures downstream from CA1.

Cannabinoids increase the 'noise' in the CA1 network, disrupting theta sequences in hippocampal schemas. On T-maze tasks, they increase VTE, potentially also diminishing forward decoding during VTE. This may underlie the performance deficits caused by cannabinoids. The effect of cannabinoids on SWR sequences is unknown,

but it is likely to disrupt them in a similar way to theta sequences. As such, cannabinoids are not ideal for dissociating mechanisms of deliberation from those of consolidation. However, cannabinoids are useful for testing network-level effects in the deliberative system, such as the level of ensemble noise that the system can tolerate in order to facilitate communication with a downstream structure.

10.5 Conclusion

In this thesis, I ask the question, 'What is VTE?' The conclusion is reached that VTE is often a marker for deliberation involving hippocampus - but not always. Sometimes it is linked to stimulus localization or sensory integration, as in the case of sensory discrimination. This behavior may eventually be classified as something other than VTE. I have outlined the state of the field of *in vivo* hippocampal electrophysiology and explored the role of the hippocampus in deliberation. I have described a novel neuroeconomic procedure, the spatial delay discounting task. Behavioral results on this task suggest that it is well-suited for testing theories of discounting behavior. Rats consistently titrate an adjusting delay to a stable indifference point and then alternate. VTE on the task occurs during periods of flexible decision-making, consistent with a role in deliberation.

I found that SWR sequences and theta sequences contain similar information on a lap-by-lap basis. Though both contain information about future choices, SWR and VTE occur at different times on the task. While VTE tends to occur during Titration, SWR occur more frequently during Exploitation. SWR also occur more frequently during non-VTE laps. The redundant information in these representations may be a result of multiple ways of querying the hippocampus, or biases in deliberation due to ongoing consolidation. Future research involving cannabinoid agonists to disrupt theta sequences may reveal important network-level effects for effective communication during deliberation.

References

Abel EL (1971) Marihuana and memory: acquisition or retrieval? *Science* 173:1038–1040.

Abela AR, Chudasama Y (2013) Dissociable contributions of the ventral hippocampus and orbitofrontal cortex to decision-making with a delayed or uncertain outcome. *European Journal of Neuroscience* 37:640–647.

Aggleton J (1985) One-trial object recognition by rats. *The Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology* 37:279–294.

Agnihotri NT, Hawkins RD, Kandel ER, Kentros C (2004) The long-term stability of new hippocampal place fields requires new protein synthesis. *Proceedings of the National Academy of Sciences of the United States of America* 101:3656–3661.

Ahmed OJ, Mehta MR (2012) Running speed alters the frequency of hippocampal gamma oscillations. *The Journal of Neuroscience* 32:7373–7383.

Allen K, Rawlins JNP, Bannerman DM, Csicsvari J (2012) Hippocampal place cells can encode multiple trial-dependent features through rate remapping. *The Journal of Neuroscience* 32:14752–14766.

Atance CM, O'Neill DK (2001) Episodic future thinking. *Trends in Cognitive Sciences* 5:533–539.

Bakker A, Kirwan CB, Miller M, Stark CE (2008) Pattern separation in the human hippocampal CA3 and dentate gyrus. *Science* 319:1640–1642.

Barraco D, Lovell K, Eisenstein E (1981) Effects of cycloheximide and puromycin on learning and retention in the cockroach, p. americana. *Pharmacology Biochemistry and Behavior* 15:489–494.

Belchior H, dos Santos VL, A.B.L. Tort S, Ribeiro (2014) Increase in hippocampal theta oscillations during spatial decision making. *Hippocampus* 24:693–702.

Berger TW, Rinaldi PC, Weisz DJ, Thompson RF (1983) Single-unit analysis of different hippocampal cell types during classical conditioning of rabbit nictating membrane response. *Journal of Neurophysiology* 50:1197–1219.

Bett D, Allison E, Murdoch LH, Kaefer K, Wood ER, Dudchenko PA (2012) The neural substrates of deliberative decision making: contrasting effects of hippocampus lesions on performance and vicarious trial-and-error behavior in a spatial memory task and a visual discrimination task. *Frontiers in behavioral neuroscience* 6.

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

Blumenthal A, Steiner A, Seeland KD, Redish AD (2011) Effects of pharmacological manipulations of NMDA-receptors on deliberation in the Multiple-T task. *Neurobiology of Learning and Memory* 95:376–384.

Boomhower SR, Rasmussen EB (2014) Haloperidol and rimonabant increase delay discounting in rats fed high-fat and standard-chow diets. *Behav Pharmacol* 25:705–716.

Bornkessel-Schlesewsky I, Schlesewsky M, Small SL, Rauschecker JP (2015) Neurobiological roots of language in primate audition: common computational properties. *Trends in Cognitive Sciences* 19:142–50.

Bower GH (1959) Choice-point behavior In *Studies in Mathematical Learning Theory*, chapter 6, pp. 109–124. Stanford University Press, Stanford,CA.

Boyer P (2008) Evolutionary economics of mental time travel? *Trends in Cognitive Science* 12:219–224.

Bradshaw CM, Szabadi E (1992) Choice between delayed reinforcers in a discrete-trials schedule: the effect of deprivation level. *The Quarterly Journal of Experimental Psychology* 44B:1–16.

Bragin A, Jandó G, Nádasdy Z, van Landeghem M, Buzsáki G (1995) Dentate EEG spikes and associated interneuronal population bursts in the hippocampal hilar region of the rat. *Journal of Neurophysiology* 73:1691–1705.

Bragin A, Engel J, Wilson CL, Fried I, Buzsáki G (1999) High-frequency oscillations in human brain. *Hippocampus* 9:137–142.

Brandon MP, Koenig J, Leutgeb JK, Leutgeb S (2014) New and distinct hippocampal place codes are generated in a new environment during septal inactivation. *Neuron* 82:789–796.

Breton YA, Seeland KD, Redish AD (2015) Aging impairs deliberation and behavioral flexibility in inter-temporal choice. *Frontiers in aging neuroscience* 7.

Brog JS, Salyapongse A, Deutch AY, Zahm DS (1993) The patterns of afferent innervation of the core and shell in the accumbens part of the rat ventral striatum: Immunohistochemical detection of retrogradely transported fluoro-gold. *Journal of Comparative Neurology* 338:255–278.

Brown MW, Aggleton JP (2001) Recognition memory: what are the roles of the perirhinal cortex and hippocampus? *Nature Reviews Neuroscience* 2:51–61.

Brown TI, Hasselmo ME, Stern CE (2014) A high-resolution study of hippocampal and medial temporal lobe correlates of spatial context and prospective overlapping route memory. *Hippocampus* 24:819–839.

Brun VH, Otnæss MK, Molden S, Steffenach HA, Witter MP, Moser MB, Moser EI (2002) Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science* 296:2243–2246.

Brun VH, Leutgeb S, Schwarcz HQWR, Witter MP, Moser EI, Moser MB (2008) Impaired spatial representation in ca1 after lesion of direct input from entorhinal cortex. *Neuron* 57:290–302.

Buckner RL, Carroll DC (2007) Self-projection and the brain. *Trends in Cognitive Sciences* 11:49–57.

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

Buzsáki G, Leung LW, Vanderwolf CH (1983) Cellular bases of hippocampal EEG in the behaving rat. *Brain Research* 287:139–171.

Buzsáki G (2002) Theta oscillations in the hippocampus. *Neuron* 33:325–340.

Buzsáki G, Moser EI (2013) Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature neuroscience* 16:130–138.

Capretta PJ, Rea R (1967) Discrimination reversal learning in the crayfish. *Animal behaviour* 15:6–7.

Cardinal RN, Daw N, Robbins T, Everitt BJ (2002) Local analysis of behaviour in the adjusting-delay task for choice of delayed reinforcement. *Neural Networks* 15:617–634.

Cardinal RN, Pennicott DR, Lakmali C, Robbins TW, Everitt BJ et al. (2001) Impulsive choice induced in rats by lesions of the nucleus accumbens core. *Science* 292:2499–2501.

Carlini E, Hamaoui A, Bieniek D, Korte F (1970) Effects of (–) 9-trans-tetrahydrocannabinol and a synthetic derivative on maze performance of rats. *Pharmacology* 4:359–368.

Carr MF, Jadhav SP, Frank LM (2011) Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience* 14:147–153.

Cei A, Girardeau G, Drieu C, El Kanbi K, Zugaro M (2014) Reversed theta sequences of hippocampal cell assemblies during backward travel. *Nature neuroscience* 17:719–724.

Chase WG, Simon HA (1973) Perception in chess. *Cognitive psychology* 4:55–81.

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

Cheung THC, Cardinal RN (2005) Hippocampal lesions facilitate instrumental learning with delayed reinforcement but induce impulsive choice in rats. *BMC Neuroscience* 6:36.

Chrobak JJ, Buzsáki G (1996) High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. *Journal of Neuroscience* 16:3056–3066.

Clark BJ, Hines DJ, Hamilton DA, Whishaw IQ (2005) Movements of exploration intact in rats with hippocampal lesions. *Behavioural brain research* 163:91–99.

Clayton NS, Bussey TJ, Dickinson A (2003) Can animals recall the past and plan for the future? *Nature Reviews Neuroscience* 4:685–691.

Clayton NS, Griffiths D, Emery N, Dickinson A (2001) Elements of episodic–like memory in animals. *Philosophical Transactions of the Royal Society B: Biological Sciences* 356:1483–1491.

Cohen JD, McClure SM, Angela JY (2007) Should i stay or should i go? how the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362:933–942.

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser MB, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* 462:353–357.

Colgin LL, Moser EI, Moser MB (2008) Understanding memory through hippocampal remapping. *Trends in Neurosciences* 31:469–477.

Colgin LL (2013) Mechanisms and functions of theta rhythms. *Annual review of neuroscience* 36:295–312.

Consroe P, Sandyk R, Snider SR (1986) Open label evaluation of cannabidiol in dystonic movement disorders. *International Journal of Neuroscience* 30:277–282.

Corkin S (1984) Lasting consequences of bilateral medial temporal lobectomy: Clinical course and experimental findings in hm In *Seminars in Neurology*, Vol. 4, pp. 249–259.

Correll R, Scoville W (1965) Effects of medial temporal lesions on visual discrimination performance. *Journal of comparative and physiological psychology* 60:175.

Cowen SL, Nitz DA (2014) Repeating Firing Fields of CA1 Neurons Shift Forward in Response to Increasing Angular Velocity. *The Journal of Neuroscience* 34:232–241.

Csicsvari J, Hirase H, Mamiya A, Buzsáki G (2000) Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. *Neuron* 28:585–594.

Dall SR, Giraldeau LA, Olsson O, McNamara JM, Stephens DW (2005) Information and its use by animals in evolutionary ecology. *Trends in ecology & evolution* 20:187–193.

DArgembeau A, Van der Linden M (2004) Phenomenal characteristics associated with projecting oneself back into the past and forward into the future: Influence of valence and temporal distance. *Consciousness and cognition* 13:844–858.

Davidson TJ, Kloosterman F, Wilson MA (2009) Hippocampal replay of extended experience. *Neuron* 63:497–507.

Daw ND, Niv Y, Dayan P (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience* 8:1704–1711.

de Saint-Exupéry A (1969) *Flight to Arras* Houghton Mifflin Harcourt.

Dember WN, Fowler H (1958) Spontaneous alternation behavior. *Psychological Bulletin* 55:412–428.

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

Domjan M (1998) *The principles of learning and behavior* Brooks/Cole, fourth edition.

Dragoi G, Buzsaki G (2006) Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* 50:145–157.

Dragoi G, Tonegawa S (2011) Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* 469:1–7.

Draycott B, Loureiro M, Ahmad T, Tan H, Zunder J, Laviolette SR (2014) Cannabinoid transmission in the prefrontal cortex bi-phasically controls emotional memory formation via functional interactions with the ventral tegmental area. *J Neurosci* 34:13096–13109.

Eacott MJ, Norman G (2004) Integrated memory for object, place, and context in rats: A possible model of episodic-like memory? *J. Neurosci.* 24:1948–1953.

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

Eichenbaum H, Cohen NJ (2014) Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron* 83:764–770.

Eilam D, Golani I (1989) Home base behavior of rats (*rattus norvegicus*) exploring a novel environment. *Behavioral Brian Research* 34:199–211.

Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL, Fried I (2003) Cellular networks underlying human spatial navigation. *Nature* 425.

Ellard CG, Goodale MA, Timney B (1984) Distance estimation in the mongolian gerbil: the role of dynamic depth cues. *Behavioural Brain Research* 14:29–39.

Emery NJ, Clayton NS (2004) The mentality of crows: convergent evolution of intelligence in corvids and apes. *Science* 306:1903–1907.

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

Fenton AA, Muller RU (1998) Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences, USA* 95:3182–3187.

Fenton AA, Kao HY, Neymotin SA, Olypher A, Vayntrub Y, Lytton WW, Ludvig N (2008) Unmasking the CA1 Ensemble Place Code by Exposures to Small and Large Environments: More Place Cells and Multiple, Irregularly Arranged, and Expanded Place Fields in the Larger Space. *J. Neurosci.* 28:11250–11262.

Fenton AA, Lytton WW, Barry JM, Lenck-Santini PP, Zinyuk LE, Kubik S, Bures J, Poucet B, Muller RU, Olypher AV (2010) Attention-like modulation of hippocampus place cell discharge. *J. Neurosci.* 30:4613–4625.

Ferbinteanu J, Shapiro ML (2003) Prospective and retrospective memory coding in the hippocampus. *Neuron* 40:1227–1239.

Finch DM, Babb TL (1981) Demonstration of caudally directed hippocampal efferents in the rat by intracellular injection of horseradish peroxidase. *Brain research* 214:405–410.

Foster DJ, Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440:680–683.

Foster DJ, Wilson MA (2007) Hippocampal theta sequences. *Hippocampus* 17:1093–1099.

Frank LM, Brown EN, Wilson M (2000) Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* 27:169–178.

Frank LM, Stanley GB, Brown EN (2004) Hippocampal plasticity across multiple days of exposure to novel environments. *The Journal of neuroscience* 24:7681–7689.

Freeman KB, Nonnemacher JE, Green L, Myerson J, Woolverton WL (2012) Delay discounting in rhesus monkeys: Equivalent discounting of more and less preferred sucrose concentrations. *Learning & behavior* 40:54–60.

French ED, Dillon K, Wu X (1997) Cannabinoids excite dopamine neurons in the ventral tegmentum and substantia nigra. *Neuroreport* 8:649–652.

Furtak SC, Ahmed OJ, Burwell RD (2012) Single neuron activity and theta modulation in postrhinal cortex during visual object discrimination. *Neuron* 76:976–988.

Gardner RS, Uttaro MR, Fleming SE, Suarez DF, Ascoli GA, Dumas TC (2013) A secondary working memory challenge preserves primary place strategies despite overtraining. *Learning & Memory* 20:648–656.

Geisler C, Robbe D, Zugaro M, Sirota A, Buzsáki G (2007) Hippocampal place cell assemblies are speed-controlled oscillators. *PNAS* 104:8149–8154.

Gill PR, Mizumori SJY, Smith DM (2011) Hippocampal episode fields develop with learning. *Hippocampus* 21:1240–9.

Girardeau G, Benchenane K, Wiener SI, Buzs'aki G, Zugaro MB (2009) Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience* 12:1222–1223.

Goutagny R, Jackson J, Williams S (2009) Self-generated theta oscillations in the hippocampus. *Nature neuroscience* 12:1491.

Graeber RC, Schroeder DM, Jane JA, Ebbesson SO (1978) Visual discrimination following partial telencephalic ablations in nurse sharks (ginglymostoma cirratum). *Journal of Comparative Neurology* 180:325–344.

Gray JM, Vecchiarelli HA, Morena M, Lee TTY, Hermanson DJ, Kim AB, McLaughlin RJ, Hassan KI, Kühne C, Wotjak CT, Deussing JM, Patel S, Hill MN (2015) Corticotropin-releasing hormone drives anandamide hydrolysis in the amygdala to promote anxiety. *J Neurosci* 35:3879–3892.

Green L, Myerson J, Holt DD, Slevin JR, Estle SJ (2004) Discounting of delayed food rewards in pigeons and rats: Is there a magnitude effect? *Journal of the Experimental Analysis of Behavior* 81:39–50.

Green L, Myerson J, Calvert AL (2010) Pigeons'discounting of probabilistic and delayed reinforcers. *Journal of the experimental analysis of behavior* 94:113–123.

Gross CG (1973) Visual functions of inferotemporal cortex In *Visual Centers in the Brain*, pp. 451–482. Springer.

Gupta AS, van der Meer MAA, Touretzky DS, Redish AD (2010) Hippocampal replay is not a simple function of experience. *Neuron* 65:695–705.

Gupta AS, van der Meer MA, Touretzky DS, Redish AD (2012) Segmentation of spatial experience by hippocampal theta sequences. *Nature neuroscience* 15:1032–1039.

Gupta K, Erdem U, Hasselmo M (2013) Modeling of grid cell activity demonstrates in vivo entorhinal look-aheadproperties. *Neuroscience* 247:395–411.

Hafting T, Fyhn M, Bonnevie T, Moser MB, Moser EI (2008) Hippocampus-independent phase precession in entorhinal grid cells. *Nature* 453:1248–1252.

Han J, Kesner P, Metna-Laurent M, Duan T, Xu L, Georges F, Koehl M, Abrous DN, Mendizabal-Zubiaga J, Grandes P et al. (2012) Acute cannabinoids impair working memory through astroglial CB 1 receptor modulation of hippocampal LTD. *Cell* 148:1039–1050.

Harris KD (2005) Neural signatures of cell assembly organization. *Nature Reviews Neuroscience* 6:399–407.

Harvey CD, Collman F, Dombeck DA, Tank DW (2009) Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* 461:941–946.

Hassabis D, Kumaran D, Vann SD, Maguire EA (2007) Patients with hippocampal amnesia cannot imagine new experiences. *PNAS* 104:1726–1731.

Hassabis D, Maguire EA (2007) Deconstructing episodic memory with construction. *Trends in cognitive sciences* 11:299–306.

Hasselmo ME, Fransen E, Dickson C, Alonso AA (2000) Computational Modeling of Entorhinal Cortex. *Ann NY Acad Sci* 911:418–446.

Hebb DO (1949) *The Organization of Behavior* Wiley, New York Reissued 2002 LEA.

Heyser CJ, Hampson RE, Deadwyler SA (1993) Effects of delta-9-tetrahydrocannabinol on delayed match to sample performance in rats: alterations in short-term memory associated with changes in task specific firing of hippocampal cells. *Journal of Pharmacology and Experimental Therapeutics* 264:294–307.

Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. *Experimental Neurology* 62:282–297.

Hill MN, McLaughlin RJ, Pan B, Fitzgerald ML, Roberts CJ, Lee TTY, Karatsoreos IN, Mackie K, Viau V, Pickel VM, McEwen BS, Liu Qs, Gorzalka BB, Hillard CJ (2011) Recruitment of prefrontal cortical endocannabinoid signaling by glucocorticoids contributes to termination of the stress response. *J Neurosci* 31:10506–10515.

Hoffman KL, Dragan MC, Leonard TK, Micheli C, Montefusco-Siegmund R, Valiante TA (2013) Saccades during visual exploration align hippocampal 3–8 Hz rhythms in human and non-human primates. *Frontiers in systems neuroscience* 7.

Holderith N, Németh B, Papp OI, Veres JM, Nagy GA, Hájos N (2011) Cannabinoids attenuate hippocampal gamma oscillations by suppressing excitatory synaptic input onto CA3 pyramidal neurons and fast spiking basket cells. *The Journal of physiology* 589:4921–4934.

Hu D, Amsel A (1995) A simple test of the vicarious trial-and-error hypothesis of hippocampal function. *PNAS* 92:5506–5509.

Ingvar DH (1984) " Memory of the future": an essay on the temporal organization of conscious awareness. *Human neurobiology* 4:127–136.

Jackson J, Redish AD (2007) Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus* 17:1209–1229.

Jackson JC, Johnson A, Redish AD (2006) Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. *Journal of Neuroscience* 26:12415–12426.

Jackson J (2006) Network Consistency and Hippocampal Dynamics: Using the properties of cell assemblies to probe the hippocampal representation of space Ph.D. diss., University of Minnesota.

Jadhav SP, Kemere C, German PW, Frank LM (2012) Awake hippocampal sharp-wave ripples support spatial memory. *Science* 336:1454–1458.

Jamain S, Radyushkin K, Hammerschmidt K, Granon S, Boretius S, Varoqueaux F, Ramanantsoa N, Gallego J, Ronnenberg A, Winter D et al. (2008) Reduced social interaction and ultrasonic communication in a mouse model of monogenic heritable autism. *Proceedings of the National Academy of Sciences* 105:1710–1715.

Janabi-Sharifi F, Hayward V, Chen CSJ (2000) Discrete-time adaptive windowing for velocity estimation. *IEEE Transactions on Control Systems Technology* 8:1003–1009.

Jay TM, Glowinski J, Thierry AM (1989) Selectivity of the hippocampal projection to the prelimbic area of the prefrontal cortex in the rat. *Brain research* 505:337–340.

Jay TM, Witter MP (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* 313:574–586.

Jezek K, Henriksen EJ, Treves A, Moser EI, Moser MB (2011) Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478:246–249.

Jinno S, Klausberger T, Marton LF, Dalezios Y, Roberts JDB, Fuentealba P, Bushong EA, Henze D, Buzsáki G, Somogyi P (2007) Neuronal diversity in gabaergic long-range projections from the hippocampus. *The Journal of neuroscience* 27:8790–8804.

Johnson A, Redish AD (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience* 27:12176–12189.

Johnson A, Varberg Z, Benhardus J, Maahs A, Schrater P (2012) The hippocampus and exploration: dynamically evolving behavior and neural representations. *Fronteirs in Human Neuroscience* 6.

Jones MW, Wilson MA (2005) Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLOS Biology* 3:e402.

Kadir SN, Goodman DF, Harris KD (2014) High-dimensional cluster analysis with the masked em algorithm .

Karlsson MP, Frank LM (2009) Awake replay of remote experiences in the hippocampus. *Nature Neuroscience* 12:913–918.

Karlsson MP, Frank LM (2008) Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *The Journal of neuroscience* 28:14271–14281.

Katona I, Sperlágh B, Sık A, Käfalvi A, Vizi E, Mackie K, Freund T (1999) Presynaptically located CB1 cannabinoid receptors regulate GABA release from axon terminals of specific hippocampal interneurons. *The Journal of neuroscience* 19:4544–4558.

Kemere C, Carr MF, Karlsson MP, Frank LM (2013) Rapid and continuous modulation of hippocampal network state during exploration of new places. *PloS one* 8:e73114.

Kentros C, E H, Hawkins RD, Kandel ER, Shapiro M, Muller RV (1998) Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280:2121–2126.

Keramati M, Dezfouli A, Piray P (2011) Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Computational Biology* 7:e1002055.

Klein SB, Loftus J, Kihlstrom JF (2002) Memory and temporal experience: the effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *Social Cognition* 20:353–379.

Koch M, Varela L, Kim JG, Kim JD, Hernández-Nuño F, Simonds SE, Castorena CM, Vianna CR, Elmquist JK, Morozov YM, Rakic P, Bechmann I, Cowley MA, Szigeti-Buck K, Dietrich MO, Gao XB, Diano S, Horvath TL (2015) Hypothalamic POMC neurons promote cannabinoid-induced feeding. *Nature* 519:45–50.

Koenig J, Linder AN, Leutgeb JK, Leutgeb S (2011) The spatial periodicity of grid cells is not sustained during reduced theta oscillations. *Science* 332:592–595.

Komorowski RW, Manns JR, Eichenbaum H (2009) Robust conjunctive item–place coding by hippocampal neurons parallels learning what happens where. *The Journal of Neuroscience* 29:9918–9929.

Krajbich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nature* 13:1292–1298.

Kudrimoti HS, Barnes CA, McNaughton BL (1999) Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *Journal of Neuroscience* 19:4090–4101.

Kwan D, Craver CF, Green L, Myerson J, Boyer P, Rosenbaum RS (2012) Future decision-making without episodic mental time travel. *Hippocampus* 22:1215–1219.

Las L, Ulanovsky N (2014) Hippocampal neurophysiology across species In *Space, Time and Memory in the Hippocampal Formation*, pp. 431–461. Springer.

Lavoie AM, Mizumori SJY (1994) Spatial-, movement- and reward-sensitive discharge by medial ventral striatum neurons in rats. *Brain Research* 638:157–168.

Lee I, Rao G, Knierim J (2004) A Double Dissociation between Hippocampal Subfields: Differential Time Course of CA3 and CA1 Place Cells for Processing Changed Environments. *Neuron* 42:803–815.

Leutgeb JK, Leutgeb S, Moser MB, Moser EI (2007) Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus. *Science* 315:961–966.

Leutgeb JK, Leutgeb S, Treves A, Meyer R, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005a) Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron* 48:345–358.

Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005b) Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309:619–623.

Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI (2004) Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. *Science* 305:1295–1298.

Lichtman AH, Cook SA, Martin BR (1996) Investigation of brain sites mediating cannabinoid-induced antinociception in rats: evidence supporting periaqueductal gray involvement. *J Pharmacol Exp Ther* 276:585–593.

Lichtman AH, Dimen KR, Martin BR (1995) Systemic or intrahippocampal cannabinoid administration impairs spatial memory in rats. *Psychopharmacology* 119:282–290.

Lim SL, O'Doherty JP, Rangel A (2011) The decision value computations in the vmpfc and striatum use a relative value code that is guided by visual attention. *The Journal of Neuroscience* 31:13214–13223.

Linhares A, Freitas AET (2010) Questioning Chase and Simon's (1973) Perception in Chess: The Experience Recognition Hypothesis. *New Ideas in Psychology* 28:64–78.

Lisboa SF, Borges AA, Nejo P, Fassini A, Guimarães FS, Resstel LB (2015) Cannabinoid cb1 receptors in the dorsal hippocampus and prelimbic medial prefrontal cortex

modulate anxiety-like behavior in rats: additional evidence. *Prog Neuropsychopharmacol Biol Psychiatry* 59:76–83.

Lisman J, Redish AD (2009) Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society B Biological Sciences* 364:1193–1201.

Lisman JE, Otmakhova NA (2001) Storage, recall, and novelty detection of sequences by the hippocampus: elaborating on the socratic model to account for normal and aberrant effects of dopamine. *Hippocampus* 11:551–568.

Loureiro M, Renard J, Zunder J, Laviolette SR (2015) Hippocampal cannabinoid transmission modulates dopamine neuron activity: impact on rewarding memory formation and social interaction. *Neuropsychopharmacology* 40:1436–1447.

Lubenov EV, Siapas AG (2009) Hippocampal theta oscillations are travelling waves. *Nature* 459:534–9.

Luo AH, Tahsili-Fahadan P, Wise RA, Lupica CR, Aston-Jones G (2011) Linking context with reward: a functional circuit from hippocampal ca3 to ventral tegmental area. *Science* 333:353–357.

MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H (2011) Hippocampal time cells bridge the gap in memory for discontiguous events. *Neuron* 71:737–749.

Madden GJ, Johnson PS (2010) A delay-discounting primer In Madden G, Bickel W, editors, *Impulsivity: The Behavioral and Neurological Science of Discounting*, pp. 11–37. APA books.

Maguire EA, Burgess N, Donnett JG, Frackowiak RSJ, Frith CD, O'Keefe J (1998) Knowing where and getting there: A human navigation network. *Science* 280:921–924.

Maier N, Tejero-Cantero Á, Dorrn AL, Winterer J, Beed PS, Morris G, Kempter R, Poulet JF, Leibold C, Schmitz D (2011) Coherent phasic excitation during hippocampal ripples. *Neuron* 72:137–152.

Malan J T, Ibrahim MM, Deng H, Liu Q, Mata HP, Vanderah T, Porreca F, Makriyannis A (2001) CB2 cannabinoid receptor-mediated peripheral antinociception. *Pain* 93:239–245.

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

Marco EM, Pérez-Alvarez L, Borcel E, Rubio M, Guaza C, Ambrosio E, File SE, Viveros MP (2004) Involvement of 5-HT1A receptors in behavioural effects of the cannabinoid receptor agonist CP 55,940 in male rats. *Behav Pharmacol* 15:21–27.

Mariano T, Bannerman D, McHugh S, Preston T, Rudebeck P, Rudebeck S, Rawlins J, Walton M, Rushworth M, Baxter M et al. (2009) Impulsive choice in hippocampal but not orbitofrontal cortex-lesioned rats on a nonspatial decision-making maze task. *European Journal of Neuroscience* 30:472–484.

Markus EJ, Qin Y, Leonard B, Skaggs WE, McNaughton BL, Barnes CA (1995) Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience* 15:7079–7094.

Marr D (1971) Simple memory: A theory of archicortex. *Philosophical Transactions of the Royal Society of London* 262:23–81.

Marston H, Everitt B, Robbins T (1993) Comparative effects of excitotoxic lesions of the hippocampus and septum/diagonal band on conditional visual discrimination and spatial learning. *Neuropsychologia* 31:1099–1118.

Martin WJ, Patrick SL, Coffin PO, Tsou K, Walker JM (1995) An examination of the central sites of action of cannabinoid-induced antinociception in the rat. *Life Sci* 56:2103–2109.

Maurer AP, Burke SN, Lipa P, Skaggs WE, Barnes CA (2012) Greater running speeds result in altered hippocampal phase sequence dynamics. *Hippocampus* 22:737–747.

Maurer AP, Lester AW, Burke SN, Ferng JJ, Barnes CA (2014) Back to the future: Preserved hippocampal network activity during reverse ambulation. *The Journal of Neuroscience* 34:15022–15031.

Mazur JE (2001) Hyperbolic value addition and general models of animal choice. *Psychological Review* 108:96–112.

Mazur J (1997) Choice, delay, probability and conditioned reinforcement. *Animal Learning and Behavior* 25:131–147.

Mazzari S, Canella R, Petrelli L, Marcolongo G, Leon A (1996) N-(2-hydroxyethyl)hexadecanamide is orally active in reducing edema formation and inflammatory hyperalgesia by down-modulating mast cell activation. *Eur J Pharmacol* 300:227–236.

McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* 102:419–457.

McDonald J, Schleifer L, Richards JB, de Wit H (2003) Effects of THC on behavioral measures of impulsivity in humans. *Neuropsychopharmacology* 28:1356–1365.

McEchron MD, Disterhoft JF (1997) Sequence of single neuron changes in CA1 hippocampus of rabbits during acquisition of trace eyeblink conditioned responses. *Journal of Neurophysiology* 78:1030–1044.

McEchron MD, Tseng W, Disterhoft JF (2003) Single neurons in CA1 hippocampus encode trace interval duration during trace heart rate (fear) conditioning in rabbit. *The Journal of neuroscience* 23:1535–1547.

McFarland WL, Teitelbaum H, Hedges EK (1975) Relationship between hippocampal theta activity and running speed in the rat. *Journal of comparative and physiological psychology* 88:324.

McGaugh JL (2000) Memory–a century of consolidation. *Science* 287:248–251.

McHugh S, Campbell T, Taylor A, Rawlins J, Bannerman D (2008) A role for dorsal and ventral hippocampus in inter-temporal choice cost-benefit decision making. *Behavioral neuroscience* 122:1.

McHugh TJ, Blum KI, Tsien JZ, Tonegawa S, Wilson MA (1996) Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell* 87:1339–1349.

McHugh TJ, Jones MW, Quinn JJ, Balthasar N, Coppari R, Elmquist JK, Lowell BB, Fanselow MS, Wilson MA, Tonegawa S (2007) Dentate gyrus NMDA receptors mediate rapid pattern separation in the hippocampal network. *Science* 317:94–99.

McKenzie S, Frank AJ, Kinsky NR, Porter B, Rivière PD, Eichenbaum H (2014) Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* 83:202–215.

McNamara CG, Tejero-Cantero Á, Trouche S, Campo-Urriza N, Dupret D (2014) Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature neuroscience* .

McNaughton BL, Barnes CA, Rao G, Baldwin J, Rasmussen M (1986) Long-term enhancement of hippocampal synaptic transmission and the acquisition of spatial information. *Journal of Neuroscience* 6:563–571.

McNaughton BL, Morris RGM (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences* 10:408–415.

Mehta MR, Barnes CA, McNaughton BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences, USA* 94:8918–8921.

Miller JF, Neufang M, Solway A, Brandt A, Trippel M, Mader I, Hefft S, Merkow M, Polyn SM, Jacobs J et al. (2013) Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science* 342:1111–1114.

Miller L, Cornett T, Brightwell D, McFarland D, Drew W, Wikler A (1977) Marijuana: Effects on storage and retrieval of prose material. *Psychopharmacology* 51:311–316.

Milner B, Squire LR, Kandel ER (1998) Cognitive neuroscience and the study of memory. *Neuron* 20:445–468.

Mischel HN, Mischel W (1983) The development of childrens knowledge of self-control strategies. *Child Development* 54:603–619.

Mischel W, Shoda Y, Rodriguez MI (1989) Delay of gratification in children. *Science* 244:933–938.

Mishkin M (1954) Visual discrimination performance following partial ablations of the temporal lobe: II. Ventral surface vs. hippocampus. *Journal of comparative and physiological psychology* 47:187.

Miyamoto A, Yamamoto T, Watanabe S (1995) Effect of repeated administration of δ 9-tetrahydrocannabinol on delayed matching-to-sample performance in rats. *Neuroscience letters* 201:139–142.

Mizumori SJY, McNaughton BL, Barnes CA, Fox KB (1989) Preserved spatial coding hippocampus CA1 pyramidal cells during reversible suppression in CA3c output: Evidence for pattern completion in hippocampus. *Journal of Neuroscience* 9:3915–3928.

Mizuseki K, Royer S, Diba K, Buzsáki G (2012) Activity dynamics and behavioral correlates of CA3 and CA1 hippocampal pyramidal neurons. *Hippocampus* 22:1659–1680.

Mizuseki K, Sirota A, Pastalkova E, Buzsáki G (2009) Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* 64:267–280.

Modi MN, Dhawale AK, Bhalla US (2014) CA1 cell activity sequences emerge after reorganization of network correlation structure during associative learning. *Elife* 3:e01982.

Monaco JD, Rao G, Roth ED, Knierim JJ (2014) Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. *Nature neuroscience* 17:725–731.

Montgomery SM, Buzsáki G (2007) Gamma oscillations dynamically couple hippocampal CA3 and CA1 regions during memory task performance. *PNAS* 104:14495–14500.

Montgomery SM, Betancur MI, Buzsáki G (2009) Behavior-dependent coordination of multiple theta dipoles in the hippocampus. *The Journal of Neuroscience* 29:1381–1394.

Morris R (2006) Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity, synaptic tagging and schemas. *European Journal of Neuroscience* 23:2829–2846.

Muenzinger KF, Gentry E (1931) Tone discrimination in white rats. *Journal of Comparative Psychology* 12:195–206.

Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience* 7:1951–1968.

Muller RU, Kubie JL (1989) The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience* 9:4101–4110.

Muller RU, Kubie JL, Ranck J JB (1987) Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *Journal of Neuroscience* 7:1935–1950.

Mumby DG, Gaskin S, Glenn MJ, Schramek TE, Lehmann H (2002) Hippocampal damage and exploratory preferences in rats: memory for objects, places, and contexts. *Learning & Memory* 9:49–57.

Naber PA, Lopes da Silva FH, Witter MP (2001) Reciprocal connections between the entorhinal cortex and hippocampal fields CA1 and the subiculum are in register with the projections from CA1 to the subiculum. *Hippocampus* 11:99–104.

Nakashiba T, Buhl DL, McHugh TJ, Tonegawa S (2009) Hippocampal ca3 output is crucial for ripple-associated reactivation and consolidation of memory. *Neuron* 62:781–787.

Nakazawa K, Quirk MC, Chitwood RA, Watanabe M, Yeckel MF, Sun LD, Kato A, Carr CA, Johnston D, Wilson MA, Tonegawa S (2002) Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science* 297:211–218.

Nakazawa K, Sun LD, Quirk MC, Rondi-Reig L, Wilson MA, Tonegawa S (2003) Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron* 38:305–315.

Naya Y, Suzuki WA (2011) Integrating what and when across the primate medial temporal lobe. *Science* 333:773–776.

Nyberg L, McIntosh AR, Cabeza R, Habib R, Houle S, Tulving E (1996) General and specific brain regions involved in encoding and retrieval of events: what, where, and when. *Proceedings of the National Academy of Sciences* 93:11280–11285.

O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Research* 34:171–175.

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

O'Keefe J, Recce M (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.

Okuda J, Fujii T, Ohtake H, Tsukiura T, Tanji K, Suzuki K, Kawashima R, Fukuda H, Itoh M, Yamadori A (2003) Thinking of the future and past: The roles of the frontal pole and the medial temporal lobes. *Neuroimage* 19:1369–1380.

Olton DS (1986) Hippocampal function and memory for temporal context In *The hippocampus*, pp. 281–298. Springer.

O'Mara SM, Commins S, Anderson M, Gigg J (2001) The subiculum: a review of form, physiology and function. *Progress in neurobiology* 64:129–155.

Onaivi ES, Green MR, Martin BR (1990) Pharmacological characterization of cannabinoids in the elevated plus maze. *J Pharmacol Exp Ther* 253:1002–1009.

O'Neill J, Senior T, Csicsvari J (2006) Place-selective firing of ca1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron* 49:143–155.

O'Neill J, TJ S, K A, JR H, Csicsvari J (2008) Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience* 11:209–215.

Ono T, Nakamura K, Nishijo H, Eifuku S (1993) Monkey hippocampal neurons related to spatial and nonspatial functions. *Journal of Neurophysiology* 70:1516–1529.

O'Reilly KC, Alarcon JM, Ferbinteanu J (2014) Relative contributions of CA3 and medial entorhinal cortex to memory in rats. *Frontiers in behavioral neuroscience* 8.

Ormond J, McNaughton BL (2015) Place field expansion after focal mec inactivations is consistent with loss of fourier components and path integrator gain reduction. *Proceedings of the National Academy of Sciences* 112:4116–4121.

Pacher P, Mechoulam R (2011) Is lipid signaling through cannabinoid 2 receptors part of a protective system? *Prog Lipid Res* 50:193–211.

Packard MG, McGaugh JL (1996) Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory* 65:65–72.

Palombo DJ, Keane MM, Verfaellie M (2014) The medial temporal lobes are critical for reward-based decision making under conditions that promote episodic future thinking. *Hippocampus* .

Palombo DJ, Keane MM, Verfaellie M (2015) How do lesion studies elucidate the role of the hippocampus in intertemporal choice? *Hippocampus* 25:407–408.

Papale AE, Stott JJ, Powell NJ, Regier PS, Redish AD (2012) Interactions between deliberation and delay-discounting in rats. *Cognitive, Affective, & Behavioral Neuroscience* 12:513–526.

Pastalkova E, Itskov V, Amarasingham A, Buzsaki G (2008) Internally generated cell assembly sequences in the rat hippocampus. *Science* 321:1322–1327.

Patel J, Fujisawa S, Berényi A, Royer S, Buzsáki G (2012) Traveling theta waves along the entire septotemporal axis of the hippocampus. *Neuron* 75:410–417.

Patel J, Schomburg EW, Berényi A, Fujisawa S, Buzsáki G (2013) Local generation and propagation of ripples along the septotemporal axis of the hippocampus. *The Journal of Neuroscience* 33:17029–17041.

Patel S, Cone RD (2015) Neuroscience: a cellular basis for the munchies. *Nature* 519:38–40.

Pattij T, Janssen MCW, Schepers I, González-Cuevas G, de Vries TJ, Schoffelmeer ANM (2007) Effects of the cannabinoid CB1 receptor antagonist rimonabant on distinct measures of impulsive behavior in rats. *Psychopharmacology (Berl)* 193:85–96.

Pavlides C, Winson J (1989) Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *Journal of Neuroscience* 9:2907–2918.

Paxinos G, Watson C (1998) *The Rat Brain in Stereotaxic Coordinates* Academic Press, New York, fourth edition.

Paxinos G (2014) *The rat nervous system* Academic Press.

Peters J, Büchel C (2010) Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron* 66:138–148.

Peterson JR, Hill CC, Kirkpatrick K (2014) Measurement of impulsive choice in rats: Same-and alternate-form test–retest reliability and temporal tracking. *Journal of the experimental analysis of behavior* .

Petry NM, Madden GJ (2010) *Impulsivity: The behavioral and neurological science of discounting*, chapter Discounting and Pathological Gambling, pp. 273–294 American Psychological Association, Washington DC.

Pfeiffer BE, Foster DJ (2013) Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497:74–79.

Pitkänen A, Pikkarainen M, Nurminen N, Ylinen A (2000) Reciprocal connections between the amygdala and the hippocampal formation, perirhinal cortex, and postrhinal cortex in rat: a review. *Annals of the New York Academy of Sciences* 911:369–391.

Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. *The Journal of Neuroscience* 10:2008–2017.

Race E, Keane MM, Verfaellie M (2011) Medial temporal lobe damage causes deficits in episodic memory and episodic future thinking not attributable to deficits in narrative construction. *The Journal of Neuroscience* 31:10262–10269.

Ranganathan M, Dsouza DC (2006) The acute effects of cannabinoids on memory in humans: a review. *Psychopharmacology* 188:425–444.

Redish AD (1999) *Beyond the Cognitive Map: From Place Cells to Episodic Memory* MIT Press, Cambridge MA.

Redish AD (2013) *The Mind within the Brain* Oxford University Press, New York, NY.

Reed S (2012) *Cognition: Theories and applications* CENGAGE learning.

Richardson JD, Kilo S, Hargreaves KM (1998) Cannabinoids reduce hyperalgesia and inflammation via interaction with peripheral CB1 receptors. *Pain* 75:111–119.

Risold P, Swanson L (1997) Connections of the rat lateral septal complex. *Brain research reviews* 24:115–195.

Robbe D, Buzsaki G (2009) Alteration of Theta Timescale Dynamics of Hippocampal Place Cells by a Cannabinoid Is Associated with Memory Impairment. *J. Neurosci.* 29:12597–12605.

Robbe D, Montgomery SM, Thome A, Rueda-Orozco PE, McNaughton BL, Buzsaki G (2006) Cannabinoids reveal importance of spike timing coordination in hippocampal function. *Nature Neuroscience* 9:1526–1533.

Royer S, Zemelman BV, Losonczy A, Kim J, Chance F, Magee JC, Buzsáki G (2012) Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nature neuroscience* 15:769–775.

Rubino T, Realini N, Castiglioni C, Guidali C, Viganó D, Marras E, Petrosino S, Perletti G, Maccarrone M, Di Marzo V, Parolaro D (2008) Role in anxiety behavior of the endocannabinoid system in the prefrontal cortex. *Cereb Cortex* 18:1292–1301.

Russell JC, McMorland AJ, MacKay JW (2010) Exploratory behaviour of colonizing rats in novel environments. *Animal Behaviour* 79:159–164.

Sakon JJ, Naya Y, Wirth S, Suzuki WA (2014) Context-dependent incremental timing cells in the primate hippocampus. *Proceedings of the National Academy of Sciences* 111:18351–18356.

Salgado S, Kaplitt M (2015) The nucleus accumbens: A comprehensive review. *Stereotactic and functional neurosurgery* 93:75–93.

Samuels I (1972) Hippocampal lesions in the rat: Effects on spatial and visual habits. *Physiology & behavior* 8:1093–1097.

Schacter DL, Addis DR (2007) The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B* 362:773–786.

Scheffer-Teixeira R, Belchior H, Leo R, Ribeiro S, Tort A (2013) On high-frequency field oscillations ($>$ 100 Hz) and the spectral leakage of spiking activity. *The Journal of Neuroscience* 33:1535–1539.

Schmidt B, Papale A, Redish A, Markus E (2013a) Conflict between place and response navigation strategies: Effects on vicarious trial and error (VTE) behaviors. *Learning & Memory* 20:130–138.

Schmidt B, Hinman JR, Jacobson TK, Szkudlarek E, Argraves M, Escabí MA, Markus EJ (2013b) Dissociation between dorsal and ventral hippocampal theta oscillations during decision-making. *The Journal of Neuroscience* 33:6212–6224.

Schmidt R (2010) Hippocampal correlation coding: phase precession and temporal patterns in CA3 and CA1 Ph.D. diss., Berlin, Humboldt-Univ., Diss., 2010.

Schmidt R, Diba K, Leibold C, Schmitz D, Buzsáki G, Kempter R (2009) Single-trial phase precession in the hippocampus. *The Journal of neuroscience* 29:13232–13241.

Schmitzer-Torbert N, Jackson J, Henze D, Harris K, Redish A (2005) Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* 131:1–11.

Schomburg EW, Fernández-Ruiz A, Mizuseki K, Berényi A, Anastassiou CA, Koch C, Buzsáki G (2014) Theta phase segregation of input-specific gamma patterns in entorhinal-hippocampal networks. *Neuron* 84:470–485.

Schrier A, Povar M (1979) Eye movements of stumptailed monkeys during discrimination learning: VTE revisited. *Animal Learning and Behavior* 7:239–245.

Schrier A, Povar M (1982) Eye movements of monkeys during discrimination learning: role of visual scanning. *Journal of Experimental Psychology: Animal Behavior Processes* 8:33–48.

Serrano A, Parsons LH (2011) Endocannabinoid influence in drug reinforcement, dependence and addiction-related behaviors. *Pharmacol Ther* 132:215–241.

Sheffield ME, Dombeck DA (2015) Calcium transient prevalence across the dendritic arbour predicts place field properties. *Nature* 517:200–204.

Siapas AG, Wilson MA (1998) Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* 21:1123–1128.

Sidman M, Stoddard L, Mohr J (1968) Some additional quantitative observations of immediate memory in a patient with bilateral hippocampal lesions. *Neuropsychologia* 6:245–254.

Siegle JH, Wilson MA (2014) Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus. *Elife* 3:e03061.

Silva R, Brandao M (2000) Acute and chronic effects of gepirone and fluoxetine in rats tested in the elevated plus maze: An ethological analysis. *Pharmacology Biochemistry and Behavior* .

Simon NW, LaSarge CL, Montgomery KS, Williams MT, Mendez IA, Setlow B, Bizon JL (2010) Good things come to those who wait: Attenuated discoutning of delayed rewards in aged Fischer-344 rats. *Neurobiology of Aging* 31:853–862.

Singer AC, Carr MF, Karlsson MP, Frank LM (2013) Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* 77:1163–1173.

Singer AC, Karlsson MP, Nathe AR, Carr MF, Frank LM (2010) Experience-dependent development of coordinated hippocampal spatial activity representing the similarity of related locations. *The Journal of Neuroscience* 30:11586–11604.

Sirota A, Csicsvari J, Buhl D, Buzsáki G (2003) Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences* 100:2065–2069.

Skaggs WE, McNaughton BL (1996) Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271:1870–1873.

Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6:149–173.

Skaggs WE, McNaughton BL, Permenter M, Archibeque M, Vogt J, Amaral DG, Barnes CA (2007) EEG sharp waves and sparse ensemble unit activity in the macaque hippocampus. *Journal of neurophysiology* 98:898–910.

Spence K (1960) Conceptional models of spatial and non-spatial selective learning In Spence K, editor, *Behavior Theory and Learning: Selected Papers*, pp. 366–392. Prentice-Hall.

Squire LR, Alvarez P (1995) Retrograde amnesia and memory consolidation: A neurobiological perspective. *Current Opinion in Neurobiology* 5:169–177.

Stark E, Roux L, Eichler R, Senzai Y, Royer S, Buzsáki G (2014) Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. *Neuron* 83:467–480.

Steiner AP, Redish AD (2012) The road not taken: neural correlates of decision making in orbitofrontal cortex. *Frontiers in neuroscience* 6.

Stephens DW, Krebs JR (1987) *Foraging Theory* Princeton.

Stevens JR, Stephens DW (2010) *Impulsivity: The Behavioral and Neurological Science of Discounting*, chapter The adaptive nature of impulsivity, pp. 361–388 American Psychological Association.

Still A (1976) An evaluation of the use of Markov models to describe the behaviour of rats at a choice point. *Animal Behaviour* 24:498–506.

Stott JJ, Redish AD (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* 369:20130472.

Sullivan D, Csicsvari J, Mizuseki K, Montgomery S, Diba K, Buzsáki G (2011) Relationships between hippocampal sharp waves, ripples, and fast gamma oscillation: influence of dentate and entorhinal cortical activity. *The Journal of Neuroscience* 31:8605–8616.

Sutton RS, Barto AG (1998) *Reinforcement Learning: An introduction* MIT Press, Cambridge MA.

Suzuki SS, Smith GK (1987) Spontaneous EEG spikes in the normal hippocampus. I. behavioral correlates, laminar profiles and bilateral synchrony. *Electroencephalography and clinical neurophysiology* 67:348–359.

Tanila H (1999) Hippocampal place cells can develop distinct representations of two visually identical environments. *Hippocampus* 9:235–246.

Thistlethwaite D (1951) A critical review of latent learning and related experiments. *Psychological Bulletin* 48:97.

Thompson LT, Best PJ (1990) Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Research* 509:299–308.

Tinklenberg JR, Melges FT, Hollister LE, Gillespie HK (1970) Marijuana and immediate memory. *Nature* 226:1171–1172.

Tolman EC (1938) The determiners of behavior at a choice point. *Psychological Review* 45:1–41.

Toni N, Laplagne DA, Zhao C, Lombardi G, Ribak CE, Gage FH, Schinder AF (2008) Neurons born in the adult dentate gyrus form functional synapses with target cells. *Nature neuroscience* 11:901–907.

Tort AB, Scheffer-Teixeira R, Souza BC, Draguhn A, Brankačk J (2013) Theta-associated high-frequency oscillations (110–160Hz) in the hippocampus and neocortex. *Progress in neurobiology* 100:1–14.

Tse D, Langston RF, Kakeyama M, Bethus I, Spooner PA, Wood ER, Witter MP, Morris RGM (2007) Schemas and memory consolidation. *Science* 316:76–82.

Tulving E (1985) Memory and consciousness. *Canadian Psychology/Psychologie Canadienne* 26:1.

Tulving E (2002) Episodic memory: From mind to brain. *Annual Review of Psychology* 53:1–25.

Ulanovsky N, Moss CF (2007) Hippocampal cellular and network activity in freely moving echolocating bats. *Nature Neuroscience* 10:224–233.

Valencia Torres L, da costa Araujo S, Sanchez CO, Body S, Bradshaw C, Szabadi E (2011) Transitional and steady-state choice behavior under and adjusting-delay schedule. *Journal of the Experimental Analysis of Behavior* 95:57–74.

Van Cauter T, Poucet B, Save E (2008) Unstable CA1 place cell representation in rats with entorhinal cortex lesions. *European journal of Neuroscience* 27:1933–1946.

van der Meer M, Kurth-Nelson Z, Redish AD (2012) Information processing in decision-making systems. *The Neuroscientist* 18:342–359.

van der Meer MA, Redish AD (2009) Covert expectation-of-reward in rat ventral striatum at decision points. *Frontiers in integrative neuroscience* 3.

van der Meer MA, Redish AD (2011) Theta phase precession in rat ventral striatum links place and reward information. *The Journal of Neuroscience* 31:2843–2854.

Vance VJ, Richardson AM, Goodrich RB (1965) Brightness discrimination in the collared lizard. *Science* 147:758–759.

Vanderwolf CH (1969) Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalography and clinical neurophysiology* 26:407–418.

Vertes RP (2011) Hippocampal theta rhythm of rem sleep. *Rapid Eye Movement Sleep: Regulation and Function* p. 151.

Viveros MP, Marco EM, File SE (2005) Endocannabinoid system and stress and anxiety responses. *Pharmacol Biochem Behav* 81:331–342.

Voss JL, Warren DE, Gonsalves BD, Federmeier KD, Tranel D, Cohen NJ (2011) Spontaneous revisitation during visual exploration as a link among strategic behavior, learning, and the hippocampus. *Proceedings of the National Academy of Sciences* 108:E402–E409.

Wang Y, Romani S, Lustig B, Leonardo A, Pastalkova E (2014) Theta sequences are essential for internally generated hippocampal firing fields. *Nature neuroscience* .

Wheeler MA, Stuss DT, Tulving E (1997) Toward a theory of episodic memory: the frontal lobes and autonoetic consciousness. *Psychological bulletin* 121:331.

Whishaw IQ (1985) Cholinergic receptor blockade in the rat impairs locale but not taxon strategies for place navigation in a swimming pool. *Behavioral neuroscience* 99:979.

Whishaw IQ, Tomie JA (1987) Cholinergic receptor blockade produces impairments in a sensorimotor subsystem for place navigation in the rat: evidence from sensory, motor, and acquisition tests in a swimming pool. *Behavioral neuroscience* 101:603.

Whishaw I, Vanderwolf CH (1973) Hippocampal EEG and behavior: change in amplitude and frequency of RSA (theta rhythm) associated with spontaneous and learned movement patterns in rats and cats. *Behavioral biology* 8:461–484.

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

Wikenheiser AM, Redish AD (2015) Hippocampal theta sequences reflect current goals. *Nature neuroscience* 18:289294.

Williams CM, Kirkham TC (1999) Anandamide induces overeating: mediation by central cannabinoid (cb1) receptors. *Psychopharmacology (Berl)* 143:315–317.

Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J (2005) Attractor dynamics in the hippocampal representation of the local environment. *Science* 308:873–876.

Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.

Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* 265:676–679.

Wingfield R, Dennis W (1934) The dependence of the rat's choice of pathways upon the length of the daily trial series. *Journal of Comparative Psychology* .

Winson J (1978) Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science* 201:160–163.

Winstanley CA, Balleine BW, Brown JW, Buchel C, Cools R, Durstewitz D, O'Doherty JP, Pennartz CM, Redish AD, Seamans JK, Robbins TW (2012) "Search, goals, and the brain" MIT Press.

Witter MP, Groenewegen HJ, Lopes da Silva FH, Lohman AHM (1989) Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region. *Progress in Neurobiology* 33:161–253.

Witter MP, Wouterlood FG, Naber PA, van Haeften T (2000) Anatomical organization of the parahippocampal-hippocampal network. *Annals of the New York Academy of Sciences* 911:1–24.

Wood ER, Dudchenko PA, Eichenbaum H (1999) The global record of memory in hippocampal neuronal activity. *Nature* 397:613–616.

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H (2000) Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* 27:623–633.

Wu X, Foster DJ (2014) Hippocampal replay captures the unique topological structure of a novel environment. *The Journal of Neuroscience* 34:6459–6469.

Yaksh TL (1981) The antinociceptive effects of intrathecally administered levonantradol and desacetyllevonantradol in the rat. *J Clin Pharmacol* 21:334S–340S.

Yamaguchi Y, Aota Y, McNaughton BL, Lipa P (2001) Bimodality of theta phase precession in hippocampal place cells in freely running rats. *Journal of Neurophysiology* 87:2629–2642.

Yamamoto J, Suh J, Takeuchi D, Tonegawa S (2014) Successful execution of working memory linked to synchronized high-frequency gamma oscillations. *Cell* 157:845–857.

Yee LT, Warren DE, Voss JL, Duff MC, Tranel D, Cohen NJ (2014) The hippocampus uses information just encountered to guide efficient ongoing behavior. *Hippocampus* 24:154–164.

Yi R, Mitchell SH, Bickel WK (2010) *Delay discounting and substance abuse-dependence.*, chapter Delay discounting and substance abuse-dependence., pp. 191–211 American Psychological Association.

Ylinen A, Bragin A, Nadasdy Z, Jando G, Szabo I, Sik A, Buzsáki G (1995) Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: Network and intracellular mechanisms. *Journal of Neuroscience* 15:30–46.

Yovel Y, Falk B, Moss CF, Ulanovsky N (2010) Optimal localization by pointing off axis. *Science* 327:701–704.

Zhang HY, Gao M, Liu QR, Bi GH, Li X, Yang HJ, Gardner EL, Wu J, Xi ZX (2014) Cannabinoid cb2 receptors modulate midbrain dopamine neuronal activity and dopamine-related behavior in mice. *Proc Natl Acad Sci U S A* 111:E5007–E5015.

Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ (1998) Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology* 79:1017–1044.

Zhang S, Manahan-Vaughan D (2014) Place field stability requires the metabotropic glutamate receptor, mGlu5. *Hippocampus* 24:1330–1340.

Zugaro MB, Monconduit L, Buzsáki G (2004) Spike phase precession persists after transient intrahippocampal perturbation. *Nature Neuroscience* 8:67–71.

Appendix A

Glossary and Acronyms

Care has been taken in this thesis to minimize the use of jargon and acronyms, but this cannot always be achieved. This appendix defines jargon terms in a glossary, and contains a table of acronyms and their meaning.

A.1 Glossary

- **Vicarious Trial and Error (VTE)** – A behavior where an animal looks toward one option, then the other, before making a choice.
- **Sharp Wave Ripple (SWR)** – A hippocampal network event, common during pauses in running, where pyramidal cells in the hippocampus oscillate at a frequency of 150-250Hz.
- **High Frequency Oscillation (HFO)** – A CA1 network oscillation, often synchronized to local pyramidal cell spiking, occurring at a frequency of 100-150Hz.
- **Adjustment** – A classification for lap-pairs on the spatial delay discounting task. Adjustment laps are repeated laps to the same side, or choices made using a win-stay strategy. Adjustment laps can cause the adjusting delay to increase by $+1$s for a repeated choice to the delay side or decrease by -1s for a repeated choice to the non-delay side.
- **Alternation** – A classification for lap-pairs on the spatial delay discounting task.

Alternation laps occur when the rat chooses the opposite side from the previous choice. The adjusting delay remains constant during alternation laps.

- **Investigation** – A classification for laps on the spatial delay discounting task. Investigation laps occur in the first 30 laps before the beginning of the Titration phase. During Investigation, the rat learns the unknown parameters of the task for the given day.
- **Titration** – A classification for laps on the spatial delay discounting task. Titration laps occur when there are two or more Adjustment laps in a five lap window (all five laps are then classified as Titration). During Titration, the rat adjusts the delay to the indifference point.
- **Exploitation** – A classification for laps on the spatial delay discounting task. Exploitation laps occur when there are fewer than two Adjustment laps in a five lap window after the first Titration laps (all five laps are then classified as Exploitation). During Exploitation, the rat takes advantage of the favorable trade-off at the indifference point.
- **Indifference Point** – In a delay discounting task, when the values of the two options are, presumably, equal such that the rat chooses both options equally.
- **SWR sequence** – A Sharp Wave Ripple sequence is the firing of CA1 pyramidal cells during ripples such that the order (or reverse order) of activation is repeated from when the rat was previously running along a trajectory.
- **Theta sequence** – A theta sequence is the ordered firing of CA1 pyramidal cells while running along a trajectory within the time window of each theta oscillation.
- **Spatial delay discounting task** – The experimental procedure used in this thesis, the spatial delay discounting task is a spatial T-maze version of the adjusting-delay discounting task, which is typically conducted in an operant chamber.

A.2 Acronyms

Table A.1: Acronyms

Acronym	Meaning
VTE	Vicarious Trial and Error
SWR	Sharp Wave Ripple
HFO	High Frequency Oscillation
dlStr	dorsolateral Striatum
EPSP	Excitatory Post-Synaptic Potential
TDRL	Temporal Difference Reinforcement Learning
SFS	Saccade Fixate Saccade
EC	Entorhinal Cortex
LED	Light Emitting Diode