Brain network mechanisms in learning behavior

Raphael Thomas Gerraty

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2018

© 2018
Raphael Gerraty
All rights reserved

ABSTRACT

Brain network mechanisms in learning behavior

Raphael Thomas Gerraty

The study of learning has been a central focus of psychology and neuroscience since their inception. Cognitive neuroscience's traditional approach to understanding learning has been to decompose it into discrete cognitive processes with separable and localized underlying neural systems. While this focus on modular cognitive functions for individual brain areas has led to considerable progress, there is increasing evidence that much of learning behavior relies on overlapping cognitive and neural systems, which may be harder to disentangle than previously envisioned. This is not surprising, as the processes underlying learning must involve widespread integration of information from sensory, affective, and motor sources. The standard tools of cognitive neuroscience limit our ability to describe processes that rely on widespread coordination of brain activity. To understand learning, it will be necessary to characterize dynamic co-activation at the circuit level.

In this dissertation, I present three studies that seek to describe the roles of distributed brain networks in learning. I begin by giving an overview of our current understanding of multiple forms of learning, describing the neural and computational mechanisms thought to underlie incremental feedback-based learning and flexible episodic memory. I will focus in particular on the difficulties in separating these processes at the cognitive level and in localizing them to individual regions at the neural level. I will then describe recent findings that have begun to characterize the brain's large-scale network

structure, emphasizing the potential roles that distributed networks could play in understanding learning and cognition more generally. I will end the introduction by reviewing current attempts to characterize the dynamics of large-scale brain networks, which will be essential for providing a mechanistic link to learning behavior.

Chapter 2 is a study demonstrating that intrinsic connectivity between the hippocampus and the ventromedial prefrontal cortex, as well as between these regions and distributed brain networks, is related to individual differences in the transfer of learning on a sensory preconditioning task. The hippocampus and ventromedial prefrontal cortex have both been shown to be involved in this type of learning, and this study represents an early attempt to link connectivity between individual regions and broader networks to learning processes.

Chapter 3 is a study that takes advantage of recent developments in mathematical modeling of temporal networks to demonstrate a relationship between large-scale network dynamics and reinforcement learning within individuals. This study shows that the flexibility of network connectivity in the striatum is related to learning performance over time, as well as to individual differences in parameters estimated from computational models of reinforcement learning. Notably, connectivity between the striatum and visual as well as orbitofrontal regions increased over the course of the task, which is consistent with an integrative role for the region in learning value-based associations. Network flexibility in a distinct set of regions is associated with episodic memory for object images presented during the learning task.

Chapter 4 examines the role of dopamine, a neurotransmitter strongly linked to value updating in reinforcement learning, in the dynamic network changes occurring dur-

ing learning. Patients with Parkinson's disease, who experience a loss of dopaminergic neurons in the substantia nigra, performed a reversal-learning task while undergoing functional magnetic resonance imaging. Patients were scanned on and off of a dopamine precursor medication (levodopa) in a within-subject design in order to examine the impact of dopamine on brain network dynamics during learning. The reversal provided an experimental manipulation of dynamic connectivity, and patients on medication showed greater modulation of striatal-cortical connectivity. Similar results were found in a number of regions receiving midbrain projections including the prefrontal cortex and medial temporal lobe. This study indicates that dopamine inputs from the midbrain modulate large-scale network dynamics during learning, providing a direct link between reinforcement learning theories of value updating and network neuroscience accounts of dynamic connectivity.

Together, these results indicate that large-scale networks play a critical role in multiple forms of learning behavior. Each highlights the potential importance of understanding dynamic routing and integration of information across large-scale circuits for our conception of learning and other cognitive processes. Understanding the when, where, and how of this information flow in the brain may provide an alternative or compliment to traditional theories of distinct learning systems. These studies also illustrate challenges in integrating this perspective with established theories in cognitive neuroscience. Chapter 5 will situate the studies in a broader discussion of how brain activity relates to cognition in general, while pointing out current roadblocks and potential ways forward for a cognitive network neuroscience of learning.

Table of Contents

List of Figures and Tables iii

Acknowledgments v

Chapter 1- Introduction

I. Rationale 2

II. Multiple Forms of Learning 4

III. Large-scale Functional Networks and Cognition 10

IV. A Dynamic Network Neuroscience 14

V. Overview of Present Research 18

Chapter 2- Transfer of learning relates to intrinsic connectivity between hippocampus, ventromedial prefrontal cortex, and large-scale networks 21

Introduction 23

Materials and Methods 25

Results 34

Discussion 38

Chapter 3- Dynamic flexibility in striatal-cortical circuits supports reinforcement learning 43

Introduction 45

Materials and Methods 47

Results 64

Discussion 73

Chapter 4- Dopamine induces learning-related modulation of dynamic striatal-cortical connectivity: evidence from Parkinson's disease 85

Introduction 87

Materials and Methods 91

Results 103
Discussion 109

Chapter 5- Network neuroscience: phenomena in search of a concept? 114
Introduction: the promise of a network neuroscience 116
Theoretical contributions of network neuroscience 120
Challenges 128
An example: the network neuroscience of learning 133
The future 137

References 141

List of Figures and Tables

Chapter 2

Figure 1.1 Task design 25

Table 1.1 Regions of each component included in connectivity analysis 31

Figure 1.2 Connectivity between the hippocampus and the vmPFC negatively correlates with transfer 33

Table 1.2 Correlations between behavioral transfer and functional connectivity between ROIs 35

Figure 2.3 Intrinsic Default Mode Network connectivity with vmPFC and hippocampus is associated with transfer 36

Figure 2.4 vmPFC connectivity with the frontoparietal network is negatively associated with transfer 37

Chapter 3

Figure 3.1 Task design and learning performance 48

Figure 3.2 Spatial distribution of flexibility 65

Figure 3.3 Flexibility in the striatum relates to learning performance within subjects 66

Figure 3.4 Flexibility in the striatum relates to reinforcement learning parameters across subjects 67

Figure 3.5 Allegiance between the striatum and visual cortex increases over the course of learning 69

Figure 3.6 Flexibility in cortical regions is related to learning performance 71

Figure 3.7 Network flexibility in medial prefrontal and parahippocampal cortex relates to episodic memory 72

Supplementary Figures

Supp. Figure 3.1 Grid search for optimizing hyperparameters 77

Supp. Figure 3.2 Subject-level data and fits for Bayesian hierarchical model of the effect of striatal flexibility on learning 78

Supp. Figure 3.3 The relationship between striatal flexibility and learning is robust across different lengths of connectivity window 79
Supp. Figure 3.4 Joint distribution of correlation between flexibility and reinforcement learning parameters for striatal regions 80
Supp. Figure 3.5 Module allegiance by striatal region 81
Supp. Figure 3.6 Visual and value regions change allegiance with the striatum over the course of the task 82
Supp. Figure 3.7 Exploratory uncorrected results for whole-brain effect of network flexibility on learning performance 83
Supp. Table 3.1 Regions showing effect of network flexibility on learning performance 84

Chapter 4

Figure 4.1 Reversal learning task 93
Figure 4.2 Reversal learning in Parkinson's patients ON and OFF of dopaminergic medication 103
Figure 4.3 Learning-related changes in striatal connectivity are modulated by dopamine 104
Figure 4.4 Dopamine increases dynamic connectivity across widespread brain areas 106
Figure 4.5 Dopamine modulated dynamic connectivity between the striatum and task-relevant visual regions 108

Acknowledgments

A list of the people who have made an impact on me during my time at Columbia would take another dissertation, but consider this the abridged version:

First and foremost Daphna Shohamy, who teaches me how to be a scientist in every one of our interactions. Her mentorship and support have been essential over the past six years. I would not know how to design an experiment, give a scientific talk, or write any of the papers that make up this dissertation without her. I am in awe of her abilities as a scientist and as a mentor, and I hope some part of the extent to which I have benefited from her guidance comes through in this work.

Danielle Bassett has been an inspiration as much as an advisor, and I am unbelievably lucky to have collaborated with her. Her approach to studying the brain was the missing piece of the puzzle that was my early years in graduate school. It should be glaringly obvious that the work in this dissertation would not have happened without her.

To the rest of my committee: Kevin Ochsner, James Curley, and Mariam Aly. I am lucky to have had all of them as role models and to have their feedback on this dissertation.

I have learned a great deal from (and had a great time with) everyone I worked with in the Learning Lab and at Columbia in general. Juliet Davidow, Jenna Reinen, Katherine Duncan, and Rahia Mashoodh were friends and mentors when I first got here, and I miss working with them all dearly. Madeleine Sharp and I designed the study in Chapter 4 together, and could not have carried it out without the help of two amazing research assistants: Amanda Buch and Christina Galese. Kendall Braun and I got through graduate school together, and were the lab's two "sixth-years". Akram Bakkour and Dave Barack could always be counted on for a great scientific chat over a beer, and Matti Vuorre and Bruce Dore were always game to talk to me about multilevel models. Thank you to Melina Tsitsiklis for checking in to make sure I was alive while I wrote this.

My friends have kept me sane for many years. I could not have made it through this without them. Particular thanks to Ben Castanon, Jordan Conan, and Eli Scherer for help formatting a table in Chapter 2. And thanks to eMalick Njie and everyone else in Neurostorm for reminding me that science is fun. To Alexa, thanks for everything, boo.

Thanks of course to my family. My mother Carmela raised me in the face of extreme obstacles, and has supported me unconditionally on the journey here. To my brother Myles for always being there when I needed him, and for being a little brother that I look up to.

This is for my father, who I think would have liked it.

Chapter 1

Introduction

Introduction

I. Rationale

The study of learning has been an essential theme throughout the history of psychology and neuroscience (Thorndike, 1898; Skinner, 1948; Lashley, 1950; Shepard, 1967; Schultz et al., 1997). Changing our behavior in response to experience— which can range from incrementally updating actions following repeated outcomes to generalizing from memories of unique episodes—is crucial to survival in a changing world. In modern cognitive neuroscience, a mechanistic understanding of learning remains a central goal.

Decades of research have led to substantial progress in understanding the neural underpinnings of learning in terms of distinct cognitive processes implemented locally in individual brain regions (Schacter and Tulving, 1994; Nelson, 1995; White and McDonald, 2002; Poldrack and Packard, 2003; Squire and Wixted, 2011). The striatum is known to be crucially involved in the updating of stimulus and action value via midbrain dopaminergic inputs (Fibiger et al., 1974; White, 1989b; Robbins and Brown, 1990; Blackburn et al., 1992; Knowlton et al., 1996; Shohamy et al., 2004b; Yin et al., 2006; Rutledge et al., 2010; Foerde and Shohamy, 2011b), while the hippocampus and surrounding medial temporal lobe are essential for forming more abstract spatial, relational, or statistical associations (O'Keefe and Dostrovsky, 1971; Dusek and Eichenbaum, 1997; Eichenbaum, 2000; Nakazawa et al., 2002; Mayes et al., 2007; Shohamy and Wagner, 2008; Schapiro et al., 2012; Eichenbaum, 2017). But there is increasing evidence that the cognitive processes underlying learning are overlapping, relevant to other functions, and rely on distributed rather than local neural circuits (Tulving and Markowitsch, 1997; Wesierska et al., 2005; Graham et al., 2010; Foerde and Shohamy, 2011a; Wimmer and

Shohamy, 2012; Shohamy and Turk-Browne, 2013). Learning—even in tasks as deceptively simple as classical or instrumental conditioning—is a complex process that relies on the integration of multidimensional information over time, and cognitive neuroscience has thus far had little to say concerning the underlying brain mechanisms of this widespread integration.

The approach taken in this dissertation is that, in order to fully describe learning (or any complex cognitive process), we must begin to understand the way the brain functions on the scale of distributed and dynamic circuits. The goal is to characterize learning systems by their network properties and to better understand the function of brain networks using distinct learning behaviors to constrain those functions. Focusing on how different regions integrate information from multiple sources, and under what circumstances this take place, may help to clarify the complex relationship between local brain activity, cognitive processes, and behavior.

The idea that large-scale brain networks—widely distributed groups of regions with correlated activity—provide a mechanism for integration is not new (Gallos et al., 2012; Park and Friston, 2013; Sporns, 2013). Only recently, however, has this become a tractable avenue for cognitive neuroscience research. Beginning with early studies demonstrating the existence of stable patterns of correlations across distributed groups of brain areas (Biswal et al., 1995) and the eventual discovery of the Default Mode Network (Raichle et al., 2001), neuroscientists have begun to characterize a host of resting-state networks, so called because of their presence in the absence of external stimuli. These networks exhibit similarity to patterns of co-activation across experimental tasks, suggesting an active role in cognitive processes (Mišić and Sporns, 2016). Some early at-

tempts to link brain networks to cognition (including one presented in Chapter 2) have associated stable properties of these networks at rest to learning and memory performance (McIntosh and Gonzalez-Lima, 1998; Tambini et al., 2010; Gerraty et al., 2014).

However, a major roadblock to providing a *mechanistic* explanation for the relationship between distributed patterns of brain connectivity and cognitive processes such as learning has been the lack of tools to characterize the dynamics of these networks and their direct links to specific aspects of behavior. Learning over time represents a dynamic process, one for which stable properties of brain function can provide only a limited description. Recent mathematical and statistical advances have begun, however, to enable time-resolved descriptions of large-scale network coordination, leading to an emerging field of dynamic network neuroscience (Mucha et al., 2010; Allen et al., 2012; Barttfeld et al., 2015). These early developments have paved the way for a fuller characterization of learning as a process of large-scale network coordination, in which local regions such as the striatum and hippocampus communicate dynamically with distributed networks to integrate information in the service of future actions. The studies in this dissertation explore the dynamics of large-scale networks during learning. But first, it is worth reviewing our current state of knowledge about the cognitive, neural, and computational processes underlying learning behavior.

II. Multiple Forms of Learning

Research in animals as well as in healthy and clinical human populations has supported the division between multiple learning processes that rely on distinct cognitive functions implemented in independent brain systems. Incremental, feedback-based learn-

ing depends on the striatum and its midbrain dopaminergic inputs (Kao and Powell, 1986; White, 1989b; Knowlton et al., 1996; Shohamy et al., 2004b; Yin et al., 2006; Steinberg et al., 2013), while a distinct system relies on the hippocampus and surrounding medial temporal lobe (MTL) for encoding associative or relational memories (Cohen et al., 1999; Eichenbaum, 2000; Mayes et al., 2007; Dupret et al., 2008; Shohamy and Wagner, 2008; Giovanello et al., 2009; Dupret et al., 2013). While they were originally conceived of as independent systems, there is an increasing understanding that at the very least much of learning has both incremental and associative components, requiring these systems to interact (Poldrack et al., 2001; Shohamy and Adcock, 2010; Foerde and Shohamy, 2011a).

There is a long history in psychological research of examining the effect of rewarding stimuli on incremental changes in behavior (Thorndike, 1898; Pavlov, 1927; Dews and Morse, 1958). Early lesion-based investigations into the neural processes underlying this form of learning demonstrated that the striatum and its sub-regions are critical for incorporating feedback into future actions (Kirkby and Kimble, 1968; Buerger et al., 1974; Fibiger et al., 1974). Further studies showed that the neurotransmitter dopamine plays an essential role in reward learning (Mason and Iversen, 1974; Beninger, 1983; Robbins and Everitt, 1996).

More recent investigations have begun to uncover the neurobiological and computational mechanisms underlying incremental feedback-based learning. Much of this work has focused on midbrain dopaminergic neurons and their projections to the striatum. This circuit has been shown to report a *reward prediction error*- the divergence between the expectation and receipt of a reinforcing stimulus (Schultz et al., 1997)- which is used to update the value of a cue or action. These responses are similar in nature to traditional

learning rules in psychological conditioning literature (Rescorla and Wagner, 1972), but they have been shown to reflect a more general form of updating that takes into account expected future rewards at individual time points. Notably, these dopaminergic signals resemble *temporal difference* prediction errors, which are central to a number of well-known artificial reinforcement learning algorithms, with striking precision. These algorithms are referred to as "model-free," due to the fact that they do not learn a model of statistical associations between stimuli (Sutton and Barto, 1998), reinforcing the idea that midbrain-striatal circuits are selectively involved in learning reward-based associations.

There are a number of weaknesses with the model-free learning account of midbrain and striatum function, both in terms of fully explaining the responses of neurons in both brain areas (Daw, 2011; McDannald et al., 2011; Hamid et al., 2016; Kishida et al., 2016; Gershman and Schoenbaum, 2017; Sharpe et al., 2017), and in describing learning behavior in humans and non-human animals (Tolman, 1948; Daw, 2011; Doll et al., 2012; Bornstein et al., 2017; Gershman and Daw, 2017). Model-based learning, an alternative computational framework in which an agent or organism uses statistical knowledge of the environment or task structure to predict sequences leading to rewards or to anticipate reinforcement without direct experience, accounts for some of these weaknesses, and will be discussed later in terms of its relationship to more episodic forms of learning. However, the model-free account has been largely successful and has helped elucidate the central role played by dopaminergic midbrain-striatal circuitry in incremental learning, in the form of updating predictions based on feedback.

One of the strengths of the above account is its incremental nature. Midbrain error signals are thought to update predictions or beliefs about the world through an incremen-

tal and iterative process that can converge on an adaptive response in a static environment, while changing appropriately in response to dynamic outcome contingencies. It is clear, however, that some learning happens more rapidly, with organisms able to encode episodes based on single experiences. This type of episodic learning has historically been the domain of memory research, but with increasing acknowledgment of episodic contributions to reinforcement learning and vice versa, episodic processes have become a topic of central interest to learning and memory researchers.

Research into the neural basis of memory has focused primarily on the hippocampus. Damage to this region has been known since patient H.M. to produce profound amnesia (Scoville and Milner, 1957; Squire, 2009). Following damage to the hippocampus and surrounding medial temporal lobe (MTL), patients are seemingly unable to encode any memories for new experiences. In fact, patients such as H.M. have an intact ability to learn motor sequences and to form the type of habitual associations described by the model-free learning framework described above. This finding, which led to the division of learning and memory into "declarative" and "non-declarative" forms relying on the brain systems described in this section, has had further support from studies comparing MTL-damaged amnesiac patients to individuals with Parkinson's disease, which involves damage to dopaminergic neurons in the substantia nigra resulting in decreased striatal dopamine and disrupted striatal function. Parkinson's patients have been shown across a number of studies to exhibit deficits in incremental feedback-based learning, but exhibit no impairment in episodic memory, while MTL lesion patients exhibit the opposite pattern (Knowlton et al., 1996).

Early functional magnetic resonance imaging (fMRI) studies of normal functioning hippocampi focused primarily on measuring hippocampal activation while participants encoded images of objects, then contrasting responses to objects later remembered during a subsequent memory test to forgotten objects (Paller and Wagner, 2002). While such studies are consistent with lesion studies in pointing to a role for the hippocampus in processes that fall under the rubric of "declarative" memory processes, there is reason to believe that this categorization represents an oversimplification (Shohamy and Turk-Browne, 2013). For one, a line of research in animals running parallel to the study of MTL damage in humans has clearly demonstrated that the hippocampus is essential for spatial navigation (O'KeefeIII, 1982), and that cells in this region and surrounding cortex code specifically for spatial location and time (O'Keefe and Dostrovsky, 1971; MacDonald et al., 2011). Damage to the hippocampus also impairs the ability to perform higher-order conditioning procedures such as sensory preconditioning, in which unrewarded cues are reinforced as a result of being associated with rewarded cues (see Chapter 2; (Port et al., 1987)). Other neuroimaging and lesion work in humans has shown that the MTL supports encoding statistical properties of the environment, such as whether sequentially presented pairs of objects represent common or rare transitions (Chun and Phelps, 1999; Hannula et al., 2007; Schapiro et al., 2012).

Ongoing research on the role of the hippocampus during decision-making has also begun to complicate the absolute distinction between habits learned in the striatum and memories stored in the hippocampus. The hippocampus has, for example, been shown to respond to rewarding stimuli when there is a substantial delay before the receipt of feedback (Foerde and Shohamy, 2011a), a finding in line with earlier studies of trace condi-

tioning (Solomon et al., 1986). In addition, neuroimaging work in humans has also shown that communication between the hippocampus and the striatum is related to the transfer of learned value during sensory preconditioning (Wimmer et al., 2012). Thus the hippocampus seems to play some part in incremental learning as well, at least when statistical associations between non-reinforced stimuli are relevant to the task at hand. In fact, it has been hypothesized that the model-based algorithm for learning described above is carried out in part by the hippocampus (Shohamy and Daw, 2015), which seems to be essential for forming relatively abstract relationships between stimuli (Eichenbaum et al., 1992).

While it should be clear from the above discussion that the exact functions being carried by striatal and MTL systems remain to some extent unknown, it is also apparent that both forms of learning rely on the integration of information across multiple sensory and motor domains. In learning incrementally via feedback, the striatal system must combine multidimensional sensory and motor information, while the MTL system integrates information about environmental stimuli in the form of spatiotemporal binding of flexible associations. Probing the neural processes underlying this integration requires an understanding of the brain's large-scale network properties. The central question of this dissertation is: how do specialized learning systems, particularly the striatum and the hippocampus, coordinate with distributed networks to integrate, route, and update information to contribute to different forms of learning?

III. Large-Scale Functional Networks and Cognition

In the domains of learning and memory, one compelling reason to believe that distributed networks will be an important piece of the conceptual puzzle is the seeming lack of spatial specificity, in terms of brain areas, to the processes described in the previous section. While much of the field has focused on midbrain-striatal circuitry and the hippocampus, respectively, to describe learning and memory processes, there is increasing evidence of much broader recruitment of brain regions. It should be noted that this phenomenon is not unique to learning and memory- information seems to be processed in a distributed fashion throughout the brain even in cases as simple as motion or color perception (Siegel et al., 2015)- but learning and memory will be the focus of this dissertation.

In the case of incremental reward learning, it has been shown that machine learning classifiers can decode whether a reinforcing or non-reinforcing stimulus has been shown to subjects based on fMRI activation throughout the cortex, rather than being specific to any set of "reward" related regions (Vickery et al., 2011). Subjective value and prediction error signals posited by reinforcement learning models also are also associated with changes in the BOLD signal distributed across striatal, medial prefrontal, posterior cingulate, and parietal regions, rather than being localized to the striatum alone (Kable and Glimcher, 2007; Bartra et al., 2013). Even reward prediction errors, which are localized to dopaminergic neurons in the midbrain and their ventral striatal targets, are theorized to effect updates to stimulus or action value by strengthening or weakening connections in distributed striatal cortical circuits (Glimcher 2011). Chapter 4 discusses a poten-

tial link between this theory and measures of dynamic network connectivity based in the striatum.

Episodic processes also seem to correspond to a diffuse rather than localized spatial map (Shohamy and Turk-Browne, 2013). When considering evidence for memory localization from lesion studies, it is important to note that damage to the retrosplenial cortex is known to produce retrograde and anterograde amnesia in humans (Valenstein et al., 1987) and spatial navigation impairments in rodents (Vann and Aggleton, 2002), paralleling findings following MTL damage. Lesions localized to the mediodorsal thalamic nucleus also produce memory impairments, with severe amnesia following damage to the mammillothalamic tract (Danet et al., 2015). Frontal damage also produces memory impairments, especially in cases of free or cued recall, or where ordering of spatial or temporal information is relevant to retrieval (Janowsky et al., 1989). Finally, functional imaging studies have tended to focus on the hippocampus, but meta-analyses indicate that a distributed network of regions is activated during successful memory encoding (Shohamy and Turk-Browne, 2013).

There thus seems to be a lack of regional *and* functional specificity (at least in terms of previously theorized cognitive processes) to the systems implicated in learning behavior. In parallel to the increasing evidence that cognitive processes engage widespread rather than local brain mechanisms, research has begun to characterize so-called *intrinsic* functional networks, distributed sets of regions exhibiting coupled activity even in the absence of external stimulation. Rather than testing evoked responses of individual regions to external stimuli, much of this research has focused on large-scale connectivity- the relationships between the intrinsic activity of regions across the brain, often at rest.

This has led to the characterization of consistent large-scale networks, widespread sets of brain regions whose activities are correlated (Damoiseaux et al., 2006).

This intrinsic connectivity was first demonstrated in a set of motor related regions (Biswal et al., 1995) and has more recently been expanded to a number of large-scale functional networks including default mode, frontoparietal, cingulo-opercular, and dorsal attention networks, as well as networks corresponding to known sensory and motor circuits (Power et al., 2011). Measured with standard fMRI acquisition, these networks exhibit coupling in the low frequency (.008-.12 Hz) range; however, recent fMRI studies with higher temporal resolution (Chen and Glover, 2015), as well as magnetoencephalography studies (Hipp et al., 2012), have indicated that they involve distributed communication at higher frequencies.

The importance of the brain's large-scale networks to cognition has become an area of increasing interest (Bressler and Menon, 2010). Similarity between these networks and the spatial patterns of evoked responses in task-based investigations implies a functional role for them in numerous cognitive processes (Krienen et al., 2014). In fact, it has been demonstrated using multiple techniques (e.g. univariate correlation, multivariate decomposition, and graph theoretic tools) that these patterns correspond very strongly to the large-scale functional networks described in resting subjects (Laird et al., 2011; Smith et al., 2012; Crossley et al., 2013; Cole et al., 2014). Given this correspondence, one particularly promising aspect of using large-scale networks to study cognition is the level of constraint this approach places on interpreting task-evoked regional activations, in terms of linking these activations to distinct cognitive and neural processes.

Given the arbitrary thresholding inherent in mass univariate analyses of functional brain images, it is unavoidable that statistical maps from different experiments will differ in their spatial patterns, even if the tasks engage overlapping (or even identical) neural processes. Thus researchers in different fields are prone to interpret neuroimaging results in terms of overly specific cognitive functions. Interpreting task-based activation in light of known connectivity networks provides a means to group together similar brain processes observed across tasks that may have been originally designed to test putatively distinct cognitive domains. Large-scale functional networks therefore provide potential constraints on the constituent functional systems (at least at the macroscopic level afforded by current human brain imaging techniques) on which different cognitive or behavioral processes rely. One of the biggest challenges for the network neuroscience approach moving forward will be unraveling what fundamental processes these networks are engaged in and how complex behavior emerges from these network processes.

To this end, there have been a number of recent investigations linking large-scale networks to specific cognitive domains. Relevant to learning and memory, some research has focused on changes in connectivity as a result of task-evoked activity (Harmelech et al., 2013; Gabard-Durnam et al., 2016), or of memory encoding (Vincent et al., 2006; Tambini et al., 2010). Another recent meta-analysis has linked large-scale connectivity of different sub-regions of the striatum with a set of cognitive components derived from a diverse set of tasks, with distinct striatal sub-regions involved in different functions via a connection to a number of large-scale connectivity networks (Pauli et al., 2016). It is also worth noting that changes in these networks have been observed in numerous mental disorders associated with incremental or episodic memory deficits (Greicius, 2008), includ-

ing schizophrenia (Jafri et al., 2008), depression (Sheline et al., 2010), and post-traumatic stress disorders (Sripada et al., 2012).

Yet, while there is overwhelming evidence that the brain's network properties matter for cognition, a precise description of *how* they matter has not yet been provided. Much of the reason for this has been the lack of available methods for characterizing the dynamics of networks. Most attempts to study the relationship between learning and the brain's network properties have focused on stable descriptions of the latter, often measured at rest. Graph theory, the representation of networks as mathematical structures made up of pairwise associations, has been used with great success in characterizing complex aspects of brain connectivity such as clustering, modularity, and efficiency (Bullmore and Sporns, 2009), but until relatively recently there have been no metrics for describing the evolution of graphs over time. A mechanistic understanding of the roles of large-scale brain networks in cognition, however, requires the characterization of their temporal dynamics.

IV. A Dynamic Network Neuroscience

While most studies of the brain's network properties have thus far focused on stable aspects of connectivity, measuring correlations across large windows of time often in the absence of external stimulation, there is increasing recognition of the importance of understanding the *dynamics* of brain networks in order to explain their relationship to cognitive processes such as learning (Kopell et al., 2014; Medaglia et al., 2015). To this end, recent studies have begun to investigate non-stationary aspects of large-scale brain networks (see e.g. (Fornito et al., 2012; Hutchison et al., 2013; Bassett et al., 2015)).

There are a number of reasons for the increasing push to elucidate the dynamics of large-scale networks. Most intuitively, if brain networks are important for cognitive processes such as learning and memory – and they seem to be – then our theories about the roles of brain networks in cognition must capture their dynamics. Without describing changes in brain networks over time, there is no possibility of modeling their relationship to behavior; correlations between static properties of networks and behavioral, demographic, or clinical variables can provide at best hints or snapshots of underlying mechanisms. This would hold true even if it turns out that functional brain networks do not change in terms of their *topology*, and instead only become active or inactive (or perhaps rely more on this or that individual node) during different experimental conditions or the execution of different cognitive tasks. However, while there is much more work to be done to characterize *when* and *how* large-scale brain networks change over time, there is increasing evidence that they do, and more importantly that these changes seem to involve reconfigurations in topographical structure.

Indeed, there is reason to believe that the functional connectivity patterns that make up the brain's large-scale networks change across tasks and are time-dependent even within a particular experiment. Electrophysiological studies indicate that neuronal coupling between regions can change on the order of 100 ms and that increases in large-scale coherence between the hippocampus and medial prefrontal cortex is associated with memory consolidation (Jones and Wilson, 2005). Similarly, replay of patterns of place-dependent firing of cells in the hippocampus is associated with similar replay events in sensory cortex (Ji and Wilson, 2007) and has been shown to co-occur with increased coherence between the hippocampus and the medial prefrontal cortex (Benchenane et al.,

2010). Two studies combining fMRI with electrophysiological recordings in primates have shown that sharp wave ripples in the hippocampus, which are known to relate to memory and spatial navigation, are accompanied by transient patterns of co-activation throughout the cortex, especially in regions of the default mode network (Logothetis et al., 2012; Kaplan et al., 2016).

Studies using fMRI in humans have also shown that, even when measured at rest, the brain's large-scale connectivity structure seems to exhibit dynamic changes. A number of studies measuring pairwise correlations between regions over sliding temporal windows have found robust transient changes (Liu and Duyn, 2013), with some studies using clustering algorithms to derive large-scale network "states" that the brain travels between (Allen et al., 2012). These latter findings are in line with theories that static networks measured by standard approaches to intrinsic connectivity represent "low energy" attractor states for large-scale network structure (Deco et al., 2011). Together, these findings are among a diverse program of research characterizing the "chronnectome" (Calhoun et al., 2014) or "dynome" (Kopell et al., 2014) to understanding the large-scale functional structure of the brain. At the very least, it is evident that brain networks do undergo dynamic changes. However, these changes have been difficult to quantify rigorously, which has served as a roadblock to mechanistic explanations of their relationship to cognition or behavior.

One major reason why it has been so difficult to fully characterize the apparent changes in large-scale network properties and relate them to behavior has been the lack of analytic and statistical tools to provide detailed summaries of these dynamics. In some ways this is a general problem faced by many areas of neuroscience as technical advances

allow for the collection of larger and more complex data sets. In the case of dynamic connectivity, as mentioned in the previous section, one of the most successful tools for modeling networks in general (including networks in the brain) has been graph theory, which allows for high dimensional pairs of connections to be partitioned into communities and provides metrics for summarizing the properties of these communities and their constituent nodes, which can be related to behavior (Sporns, 2003; Bassett and Bullmore, 2006; Bullmore and Sporns, 2009). Graph theory has traditionally been applied to static networks. The problem incorporating dynamics into this approach has been the inability to link nodal and community identities across time. Recent progress in network modeling has begun to change this, however.

One recent and promising theoretical advance in incorporating network dynamics into the graph theory approach has been the development of multi-slice community detection algorithms (Mucha et al., 2010). This method, explained in mathematical detail in Chapter 3, involves linking graphs across time in order to describe the evolution of their constituent sub-networks, allowing for graph theoretical descriptions of the dynamics of brain networks for the first time. One example of this has been the novel metric of *network flexibility*, a graph-based measure which indexes the extent to which brain regions couple dynamically with distinct networks over time (Bassett et al., 2011).

These dynamic graph theory methods have been successfully applied to motor learning, demonstrating that network flexibility of a number of regions predicts the extent to which individuals learn a simple visuo-motor sequence (Bassett et al., 2011). While these results represent a step forward in understanding brain network dynamics and their relationship to simple forms of learning, the details of this relationship remain unclear.

Chapters 3 and 4 in this dissertation attempt to relate network properties such as flexibility to established computational and cognitive theories of learning.

V. Overview of present research

The thrust of the above discussion is that we have come to understand a great deal about the functions of individual brain regions in multiple forms of learning, but there remains significant uncertainty about both the nature of the learning processes themselves and the specificity of the brain regions engaged in these processes. There is increasing evidence that part of this confusion is the result of a focus on mapping discrete brain regions to "known" cognitive functions. The position taken in this dissertation is that a better characterization of learning processes will entail an understanding of the brain's ability to integrate information from sensory, motor, and affective domains over time. This in turn will require us to consider the brain's large-scale circuit organization and the dynamics of activity over those circuits.

How is information relevant to a particular decision, action, or environmental context encoded and retrieved such that future behavior is modified? In addition to thinking about broad networks, answering this question will entail the characterization of the dynamic reconfiguration of these distributed brain circuits over time. Until recently, most task-evoked studies have focused on the behavior of individual regions, while most network studies have focused on static properties of brain connectivity. But, if our understanding of brain networks is to help explain how patterns of behavior are generated, we must consider brain network dynamics.

Together, the studies presented here are intended to move towards these goals. In terms of providing a general link between networks and learning behavior, Chapter 2 examines the role of intrinsic connectivity in the transfer of value, tested using a "sensory preconditioning" paradigm. This study shows that resting-state connectivity between the hippocampus and ventromedial prefrontal cortex, as well as between these regions and default mode and fronto-parietal networks, is related to the transfer of learned value based on previously encoded stimulus-stimulus associations (Gerraty et al., 2014).

While Chapter 2 focuses on static measures of brain connectivity across individuals, Chapter 3 examines the dynamics of large-scale networks and the relevance of these dynamics during learning (Gerraty et al., 2018). This study demonstrates that *network flexibility* is associated with learning performance on a reinforcement learning task, a relationship notably present in areas of the striatum, as well as orbitofrontal, parietal, and motor cortices. Specifically, the striatum increases its connectivity with visual cortex and ventromedial prefrontal cortex during a task in which choices are differentiated by visual properties. It also provides some evidence that network flexibility in medial prefrontal and hippocampal areas is related to episodic memory.

Chapter 4 explores the role of dopamine in dynamic network changes that occur during learning, by testing patients with Parkinson's disease on and off levodopa medication, a dopamine precursor used to treat the symptoms of this disease. This study provides a direct manipulation of brain network dynamics using reversal learning and a pharmacological intervention, and demonstrates that learning-dependent cortico-striatal dynamics are modulated by dopamine. In addition, by using stimuli known to target particular areas of visual cortex (scenes and objects), this study demonstrates that dopamine modulates

dynamic changes in connectivity between the striatum and task-relevant sensory regions in a functionally specific manner.

These studies each seek to characterize a distinct aspect of the relationship between distributed brain networks and learning processes. They illustrate the promises and challenges of applying network techniques to questions in cognitive neuroscience. Importantly, they represent attempts to integrate ideas about distributed brain networks with established theoretic frameworks usually focused on local brain regions. Taken together, they may begin to provide insight into the types of functions brain networks are engaged in as we learn about the world around us.

Chapter 2

Transfer of learning relates to intrinsic connectivity between hippocampus, ventromedial prefrontal cortex, and large-scale networks

Chapter 2 was published in the Journal of Neuroscience: Gerraty, R. T., Davidow, J. Y., Wimmer, G. E., Kahn, I., & Shohamy, D. (2014). Transfer of learning relates to intrinsic connectivity between hippocampus, ventromedial prefrontal cortex, and large-scale networks. Journal of Neuroscience, 34(34), 11297-11303.

Abstract

An important aspect of adaptive learning is the ability to flexibly use past experiences to guide new decisions. When facing a new decision, some people automatically leverage previously learned associations, while others do not. This variability in transfer of learning across individuals has been demonstrated repeatedly and has important implications for understanding adaptive behavior, yet the source of these individual differences remains poorly understood. In particular, it is unknown why such variability in transfer emerges even among homogeneous groups of young healthy participants who do not vary on other learning-related measures. Here we hypothesized that individual differences in transfer of learning could be related to relatively stable differences in intrinsic brain connectivity which could constrain how individuals learn. To test this, we obtained a behavioral measure of memory-based transfer outside of the scanner, and on a separate day acquired resting-state functional MRI in 42 participants. We then analyzed connectivity across ICA-derived brain networks during rest, and tested whether intrinsic connectivity in learning-related networks was associated with transfer. We found that individual differences in transfer were related to intrinsic connectivity between the hippocampus and the ventromedial prefrontal cortex, and between these regions and large-scale functional brain networks. Together, the findings demonstrate a novel role for intrinsic brain dynamics in flexible learning-guided behavior, both within a set a functionally-specific regions known to be important for learning, as well as between these regions and the default and frontoparietal networks, which are thought to serve more general cognitive functions.

Introduction

Different people learn in different ways; some individuals readily integrate old and new information to support novel inferences, while others fail to transfer what they learn to new situations (Shohamy and Wagner, 2008; van Kesteren et al., 2010; Daw et al., 2011; Wimmer and Shohamy, 2012; Zeithamova et al., 2012). Transfer of learning is highly variable even among healthy populations (Shohamy and Wagner, 2008; Wimmer and Shohamy, 2012). Moreover, transfer is dissociable from other complementary forms of learning, such as learning based on trial-by-trial reinforcement (Shohamy and Wagner, 2008; Daw et al., 2011; Wimmer and Shohamy, 2012; Doll et al., 2015). Prior studies have shown that variability in transfer is related to differences in task-evoked brain activation during learning (Shohamy and Wagner, 2008; Wimmer and Shohamy, 2012). However, a critical open question is why this variability initially emerges. It is particularly unclear whether differences in activation merely represent transient changes related to task demands, or whether they also reflect intrinsic differences in brain organization.

Flexible learning can be measured using paradigms such as "sensory preconditioning," where participants first learn associations and then make choices about novel options (Brogden, 1939; Port et al., 1987; Wimmer and Shohamy, 2011; Jones et al., 2012; Wimmer and Shohamy, 2012). Responses to novel options provide an opportunity for spontaneous transfer of initial associations without requiring or rewarding it, capturing individual variability in this tendency. Previous work suggests that transfer may emerge from associative encoding in the hippocampus (Port et al., 1987; Wimmer and Shohamy, 2012), or from inference-based processes in the ventromedial prefrontal cortex (vmPFC) (Jones et al., 2012). However, beyond a functional role for each region during

learning and decision-making, these studies do not address the central question of whether variability in transfer behavior relates to stable variation in functional brain organization.

To answer this question, we used resting-state fMRI (rs-fMRI), which provides a stable measure of distributed brain connectivity. We measured transfer with sensory preconditioning and collected rs-fMRI on a separate day to test whether differences in behavior related to intrinsic connectivity between regions implicated in transfer, including the hippocampus and the vmPFC. RS-fMRI has also been used for characterizing large-scale networks, including the default (DMN), frontoparietal (FPN), and cingulo-opercular (CON) (Power et al., 2011; Yeo et al., 2011). Interestingly, there is overlap between learning-activated regions and a subset of DMN areas – including the hippocampus and the vmPFC, which may compose a distinct sub-network (Andrews-Hanna et al., 2010; Roy et al., 2012).

Given these findings, we hypothesized that variability in transfer would relate to differences in intrinsic connectivity between the hippocampus and the vmPFC. We additionally hypothesized that these regions' participation in large-scale functional networks, particularly the DMN, may be related to transfer. Crucially, our task was designed to test whether connectivity effects are selective to transfer versus incremental reinforcement-driven learning measured in the same task.

Methods and Materials

Experimental Design

Forty-nine healthy right-handed participants (age 21.4 ± 3.3, 30 females) were recruited from Columbia University and the surrounding community. Each participant underwent two testing sessions: a behavioral session and a separate rs-fMRI scan. Informed consent was obtained according to procedures approved by the Columbia University Institutional Review Board. Seven participants were excluded from imaging analysis for excessive motion (see fMRI preprocessing in Methods below), leaving a sample of 42. Participants were paid $20 for behavioral testing, a percentage of their earnings from the task, and $20 for the MRI scan.

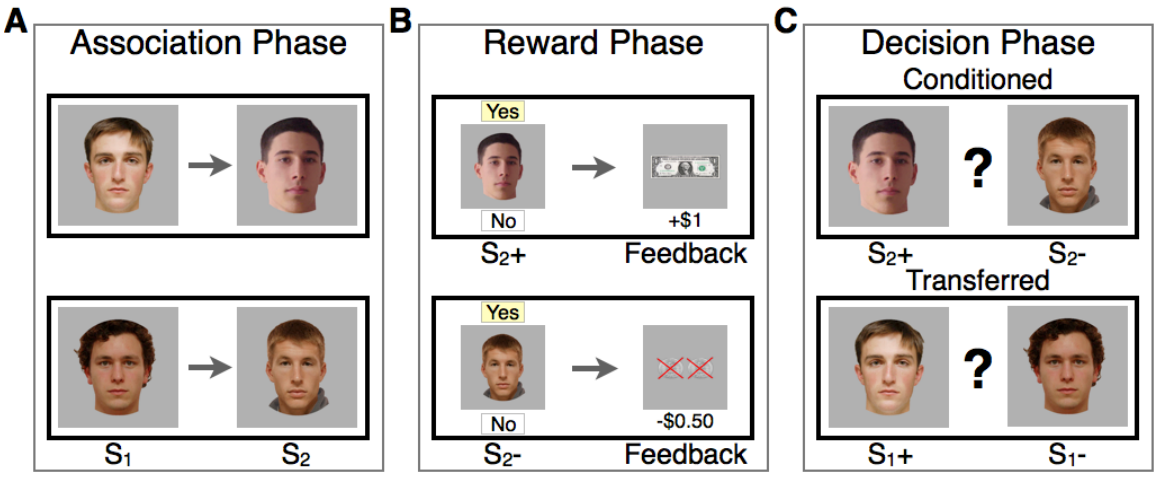

Figure 1.1 Task Design. The learning and transfer task consisted of three phases: Association phase, Reward phase, and Decision phase. (A) In the Association phase, participants viewed a series of four pairs of S_1 and S_2 face stimuli while performing a cover task. (B) In the Reward phase, participants learned that S_2 stimuli led to either monetary gain (S_{2+}), monetary loss (S_{2-}) or no outcome ($S_{2=}$), through instrumental conditioning. (C) In the Decision phase, participants chose between pairs of S_1 or pairs of S_2 stimuli. The learning of directly reinforced associations ('conditioning') was measured as participants' tendency to choose S_{2+} and avoid S_{2-}. Transfer was measured as participants' tendency to choose S_{1+} and avoid S_{1-}, reflecting transfer of learned value from S_2 to S_1 stimuli.

Task and Behavioral Analysis

Participants underwent a behavioral testing session outside the scanner, several days prior to the brain imaging session (mean time between sessions: 55.0 ± 6.7 hours). Testing was conducted first to ensure that all scanned participants had behavioral data and to screen for MRI contraindications. The transfer paradigm (**Figure 1.1**) was an adapted version of a previously-described sensory preconditioning task (Wimmer and Shohamy, 2011; Wimmer et al., 2012).

The task consists of three phases. In the *Association phase*, participants made button responses as they were exposed to sequential presentations of neutral stimuli pairs (faces, denoted S_1 and S_2). A particular S_1 stimulus always preceded its paired S_2 stimulus. Participants performed a cover task responding to "target" upside-down images, aimed at making the encoding incidental. Each of the 4 pairs was presented 10 times in pseudo-random order, intermixed with 10 target trials for 2 s, with a 2 s inter-stimulus-interval and 4 s inter-trial interval. Participants were not informed of the trial structure.

Next, during the *Reward learning phase*, participants learned to associate S_2 stimuli with a monetary gain (S_2+), a monetary loss (S_2-), or a neutral outcome (no win or loss, S_2=). On each trial, a S_2 face was presented and participants pressed a button to choose to bet or not. The S_2+ stimulus led to a reward of $1 (80% of trials 16/20) or nothing, while the S_2- stimulus led to a loss of $0.50 (80% of trials 16/20) or nothing. The S_2= stimulus led to neither gain nor loss for all 20 trials. The S_2 stimulus from the fourth pair was not presented during this phase (S_2~), providing a control stimulus with no value association for the subsequent Decision phase. Participants were instructed that they would receive a percentage of their total earnings.

In the critical *Decision phase*, participants were presented with two face stimuli and instructed to choose the face they thought would be more likely to win. Participants were instructed that they would not receive trial-by-trial feedback but that they would be rewarded for a percentage of their choices. Each choice could lead to a win of $1 or no reward. Stimuli for choice trials were either two S_1 or two S_2 faces. Choices between S_1 stimuli consisted of five trial types: positive vs. neutral (S_1+ - S_1=), positive vs. held-out (S_1+ - S_1~), positive vs. negative (S_1+ - S_1-), negative vs. neutral (S_1- - S_1=), and negative vs. held-out (S_1- - S_1~). Analogous choices were presented for S_2 stimuli. Each choice was presented three times, intermixed in pseudo-random order for 30 total trials.

We computed a transfer score for each participant by averaging the number of S_1+ stimuli chosen or S_1- stimuli avoided across the five trial types described above, providing a single estimate of individual differences in transfer, while limiting the number of comparisons with the fMRI data. Similar results were obtained using the first principal component across S_1 pairs, which explained 48.9% of the total variance as measured by the ratio of eigenvalues and was highly correlated with the mean. Population-level transfer was compared to chance performance. We used the same analysis strategy for the S_2 choices during the Decision phase, allowing a comparison of "conditioned" vs. "transfer" responses to test the selectivity of any connectivity associations.

To ensure that participants did not have pre-existing preferences for the stimuli, images were selected based on pre-task ratings. Finally, a recognition memory test for S_1-S_2 pairs was administered after the experiment to test whether subjects' transfer responses were related to explicit awareness of task structure.

MRI acquisition

Participants returned for a rs-fMRI scan that enabled us to assess intrinsic connectivity. We acquired two series of 104 interleaved T2*-weighted single-shot gradient-echo echo-planar functional images on a 3T Phillips MRI with an 8 channel head coil. Sequence parameters were as follows: TR= 3000 ms, TE=30 ms, flip angle=84 degrees, array=80x80, 41 slices, effective voxel resolution of 3 mm^3, SENSE factor=8. Participants were told they would be in the scanner for under 30 minutes, and to remain still with their eyes open looking at a fixation cross, which was back-projected and visible through a mirror mounted on the head coil. A high-resolution T1-weighted MPRAGE image was acquired for registration.

fMRI preprocessing and nuisance regression

Preprocessing of fMRI data was carried out using FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). BET was used to skull-strip anatomical and functional images (Smith, 2002). The first three volumes were removed to account for saturation effects. Functional images were slice-time corrected, motion-corrected to the median image using tri-linear interpolation with 6 degrees of freedom (Jenkinson et al., 2002), spatially smoothed with a 5 mm FWHM Gaussian kernel, grand-mean scaled, and high-pass filtered with f>0.01 hz. Functional scans were co-registered with anatomical images and linearly transformed to a standard template (T1 Montreal Neurological Institute template, voxel dimensions of 4 mm^3) using FLIRT (Jenkinson et al., 2002).

Subject-level functional data were further preprocessed with a nuisance regression using FILM with local autocorrelation implemented in FEAT (Woolrich et al., 2001).

Given the well-established effects of movement on functional connectivity (Power et al., 2012; Satterthwaite et al., 2012a), we utilized a recently-proposed extended regression technique shown to successfully control for motion artifact in connectivity measures (Satterthwaite et al., 2012b). Thirty-six nuisance variables (time-series for CSF, white matter, whole-brain, and six motion parameters, the first temporal derivative of each time-series, and the square of each time series and its derivative) were regressed against the spatially smoothed and high-pass-filtered data from initial pre-preprocessing. Each model also included single spike regressors for any time point where relative motion was greater than 2 S.D. from the global mean across subjects. Any subject with more than 20% of time-points in any scanning session "scrubbed" in this manner was excluded from further analysis, resulting in a final sample of 42 out of 49 participants. The residual time-series from each regression were then low-pass filtered at $f<0.08$ hz. For each subject, the filtered time-series from both sessions were concatenated, resulting in 202 time-points for each region and network of interest.

Connectivity Analyses

To define regions and networks of interest, we performed a spatial Independent Component Analysis (sICA), implemented in MELODIC. MELODIC's implementation of probabilistic sICA seeks to maximize the independence of latent spatial components (http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/MELODIC). For this analysis, we temporally concatenated all sessions and subjects into a 2D matrix, as this does not assume consistency of temporal patterns across subjects or sessions.

Next, regions from within one of the resulting networks, a putative subcomponent of the default mode network labeled here the Ventral Medial Network (VMN), were chosen for preliminary within-network analysis based on their association with transfer in previous reports (Jones et al., 2012; Wimmer et al., 2012; Zeithamova et al., 2012). Default mode (DMN), frontoparietal (FPN), and cingulo-opercular (CON) networks (**Table 2.1**) were included for between-network analyses examining the relationship between transfer and their connectivity with the hippocampus and vmPFC, based on previous association with the performance of cognitive tasks (Dosenbach et al., 2008; Fornito et al., 2012).

All ROIs and networks were functionally defined based on group-level ICA effects. For connectivity analysis within the ventral medial component, we isolated discrete functional ROIs using anatomical boundaries from the Harvard-Oxford atlas where available. Thus the hippocampus and striatum ROIs were defined based on overlap between group level VMN effects and anatomical masks of these regions derived from the Harvard-Oxford atlas, with a threshold probability of 25%. For the vmPFC, given its lack of clear anatomical demarcation, and taking advantage of its robust position in the VMN, the ROI was defined by thresholding the VMN component at $Z>7$, providing a region consistent with prior reports (Kumaran et al., 2009; van Kesteren et al., 2010). DMN, FPN, and CON were defined by thresholding their respective component Z statistics until there was no overlap between the networks. Mean time-series were then extracted from all regions and networks for each subject.

Table 2.1 Regions of each component included in connectivity analyses

Region	Size (voxels)	Max-Z	x	y	z
VMN					
vmpfc/hippocampus/VStr	2003	17.5	-6	42	-12
l frontal pole	17	3.86	-34	10	-36
posterior cingulate	10	3.46	-6	-54	12
r cerebellum	10	3.63	42	-42	-28
DMN					
l angular gyrus	1214	9.04	-54	-62	20
r angular gyrus	1103	8.04	58	-58	20
dorsal medial prefrontal cortex	769	9.84	2	54	24
posterior cingulate cortex	150	6.66	2	-54	32
l cerebellum	71	6.1	-26	-78	-36
r cerebellum	63	5.28	26	-78	-36
l middle frontal gyrus	39	5.04	-42	10	48
r middle frontal gyrus	17	3.48	13	35	26
frontal pole	15	4.17	2	62	-12
L FPN					
l middle/inferior frontal gyrus	1963	14.8	-50	14	32
l superior lateral occipital cortex	899	13	-34	-70	44
l middle temporal gyrus	434	9.74	-58	-46	-12
r cerebellum	359	8.21	34	-70	-44
r superior lateral occipital cortex	81	5.37	34	-66	44
r middle/inferior frontal gyrus	60	4.63	46	30	16
r middle temporal gyrus	38	5.05	8	21	14
R FPN					
r middle/inferior frontal gyrus	2303	11.5	46	22	36
r superior lateral occipital cortex	1096	13.3	46	-54	48
l cerebellum	461	10	-38	-70	-40
r middle temporal gyrus	285	7.52	62	-46	-12
l superior lateral occipital cortex	140	4.71	33	18	31
r thalamus	19	4.07	19	31	21
CON					
dACC/frontal pole/l anterior insula	3003	13.1	-26	50	20
r anterior insula	126	5.58	50	14	-4
r cerebellum	121	6.41	38	-58	-32
l cerebellum	80	6.17	-38	-58	-32

l angular gyrus	35	4.26	-58	-50	32
r caudate	35	4.37	14	22	0
r angular gyrus	26	4.47	62	-42	28

VMN, ventromedial network; DMN, default mode network; CON, cingulo-opercular network; FPN, frontoparietal network; L, left hemisphere; R, right hemisphere; dACC, dorsal anterior cingulate cortex; vmpfc, ventromedial prefrontal cortex; Vstr, ventral striatum. Components were thresholded at Z>3.1, extent > 10 for inclusion in this table. Coordinates reported in MNI space.

Within the VMN, we computed pairwise correlations between five regions: left and right hippocampus, left and right ventral striatum, and vmPFC. The correlation coefficients for each pair were Fisher *z*-transformed. Average pairwise connectivity was used to characterize the network structure. We tested the correlation between pairwise connectivity and transfer score for each set of regions (yielding 10 region-region pairs). The resulting correlations were corrected for multiple comparisons using a parametric bootstrap approach, in which we selected randomly from a 10-dimensional normal distribution with the same means and covariance as the VMN connectivity matrices across subjects. These 42×10 matrices were generated 10,000 times, with each column then correlated with composite transfer score, to calculate the family-wise error (FWE) rate for a given correlation. In addition, to probe the selectivity of any effects to transfer, we correlated pairwise connectivity with first-order conditioning performance, using both a parametric test (Meng et al., 1992) and a non-parametric bootstrap to examine the difference between these overlapping correlations.

Because of their significance for learning-based transfer (see Results), we computed connectivity between both the left hippocampus and vmPFC and three intrinsic connectivity networks: DMN, FPN, and CON. These correlations were Fisher z-transformed and included in separate linear models for each intrinsic network predicting transfer score. In order to isolate the individual effects of each region's network participa-

tion, regressors for each of these three models included connectivity between the hippocampus and the network of interest and connectivity between the vmPFC and the network of interest, as well as a nuisance regressor of hippocampus-vmPFC connectivity.

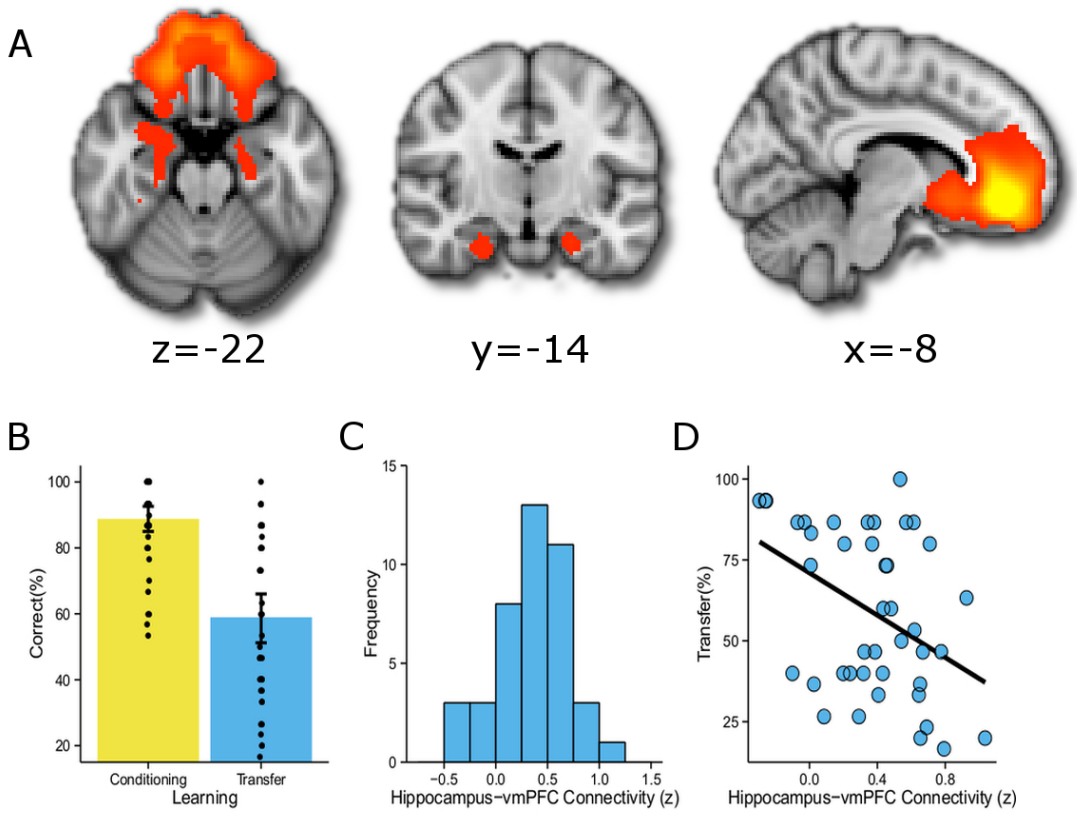

Figure 2.2 Connectivity between the hippocampus and the vmPFC negatively correlates with transfer. (A) Ventral Medial Network. Spatial ICA revealed an intrinsic functional connectivity network consisting of the ventromedial prefrontal cortex, bilateral striatum, and bilateral hippocampus. Inter-individual variability in pairwise connectivity between regions of this ventral medial network (VMN) allowed us to test the relationship between within-network connectivity and value transfer. Images thresholded at 3.1<Z and resampled to 2 mm^3 for display purposes, and are shown in radiological convention. (B) Behavioral performance during the decision phase. Average scores for first-order conditioning and transfer were both significantly above chance. Subjects exhibited a large amount of variability in transfer behavior. Bars represent mean and 95% Confidence Intervals, and points show individual subject scores. (C) The z-transformed distribution of correlation between the left hippocampus and vmPFC across participants. The y-axis represents number of participants. (D) Connectivity between the left hippocampus and vmPFC is significantly negatively correlated with transfer. A trend in the same direction was found for the right hippocampus (see text and **Table 2.2**).

Results

The sICA revealed a number of components reflecting well-characterized intrinsic connectivity networks, consistent with previous reports (Beckmann et al., 2005; Zuo et al., 2010; Fornito et al., 2012). Supporting the notion that the hippocampus, vmPFC, and ventral striatum compose a distinct sub-component of the default mode network, our sICA analysis estimated this ventral medial network (VMN) as a separate component (**Figure 2.2A**). Correlation analyses revealed significant connectivity between the vmPFC and all other regions extracted from the VMN. We also found significant inter-hemispheric connectivity for both the hippocampus and ventral striatum. These results allowed us to test whether connectivity within this network related to transfer tested on a separate day.

On the separate behavioral session, participants showed robust learning of direct S_2-outcome associations trained during the Reward phase of sensory preconditioning, evidenced by their average tendency to choose the S_{2+} and avoid the S_{2-} stimuli during the Decision phase (88.8% ± 2.0% (mean ± SE), $t(41)=19.04$, $p<0.0001$, **Figure 2.2B**). Average transfer scores were significantly greater than chance (58.9% ± 3.9%, $t(41)=2.3$, $p=0.01$ one-tailed, **Figure 2.2B**). A recognition memory test revealed no significant explicit memory for the S_1-S_2 associations (mean accuracy, 22.6% ± 3.5%, not different than chance, 25%), and the correlation between explicit memory for these pairs and transfer did not differ from chance ($r(40)=0.11$, $p=0.49$). These findings replicated prior results (Wimmer and Shohamy, 2011; Wimmer et al., 2012) indicating that any transfer in the current study was implicit and not the result of strategic reasoning. Additionally, transfer was orthogonal to first-order conditioning ($r(40)=0.02$, $p=0.93$). As anticipated,

transfer was highly variable across participants, ranging from 16.7% - 100% (**Figure 2.2B)**; thus, some participants generalized all learned associations while others showed little to no transfer.

	vmPFC	Hippocampus (R)	Hippocampus (L)	Caudate (R)	Caudate (L)
vmPFC	-	-0.29	-0.42*	0.02	0.09
Hippocampus (R)	-0.29	-	-0.18	-0.06	-0.18
Hippocampus (L)	-0.42*	-0.18	-	-0.07	-0.21
Caudate (R)	0.02	-0.06	-0.07	-	0.29
Caudate (L)	0.09	-0.18	-0.21	0.29	-

Table 2.2 Correlations between behavioral transfer and intrinsic functional connectivity between ROIs. Values in each cell represent the correlation of transfer with intrinsic functional connectivity between the ROIs listed in the corresponding row and column. Only connectivity between the vmPFC and the left hippocampus ROIs is significantly related to transfer behavior, r(40)=-0.42, p-FWE < 0.05.

We next tested the central question of whether transfer behavior relates to variability in intrinsic network connectivity (**Figure 2.2C**). Of the ten correlations between VMN regions (**Table 2.2**), only vmPFC connectivity with the left hippocampus showed a significant transfer correlation ($r(40)=-0.42$, p=0.02 FWE corrected). This correlation was negative, with lower connectivity related to higher transfer (**Figure 2.2D**). A trend in the same direction was found for vmPFC connectivity with the right hippocampus ($r(40)=-0.29$, p=0.06, uncorrected), indicating that this effect is not lateralized.

We next tested whether connectivity between the vmPFC, hippocampus, and large-scale networks of interest (**Table 2.1; Figures 2.3 and 2.4)** was related to transfer. Connectivity between the core DMN and both the hippocampus and the vmPFC was significantly related to transfer (**Figure 2.3**; DMN connectivity with the hippocampus: β=-

0.48, p=0.0002; with vmPFC: β=0.26, p=0.02; adjusted r^2=0.32, p=0.0002). Connectivity between the hippocampus and the DMN was negatively associated with transfer, while

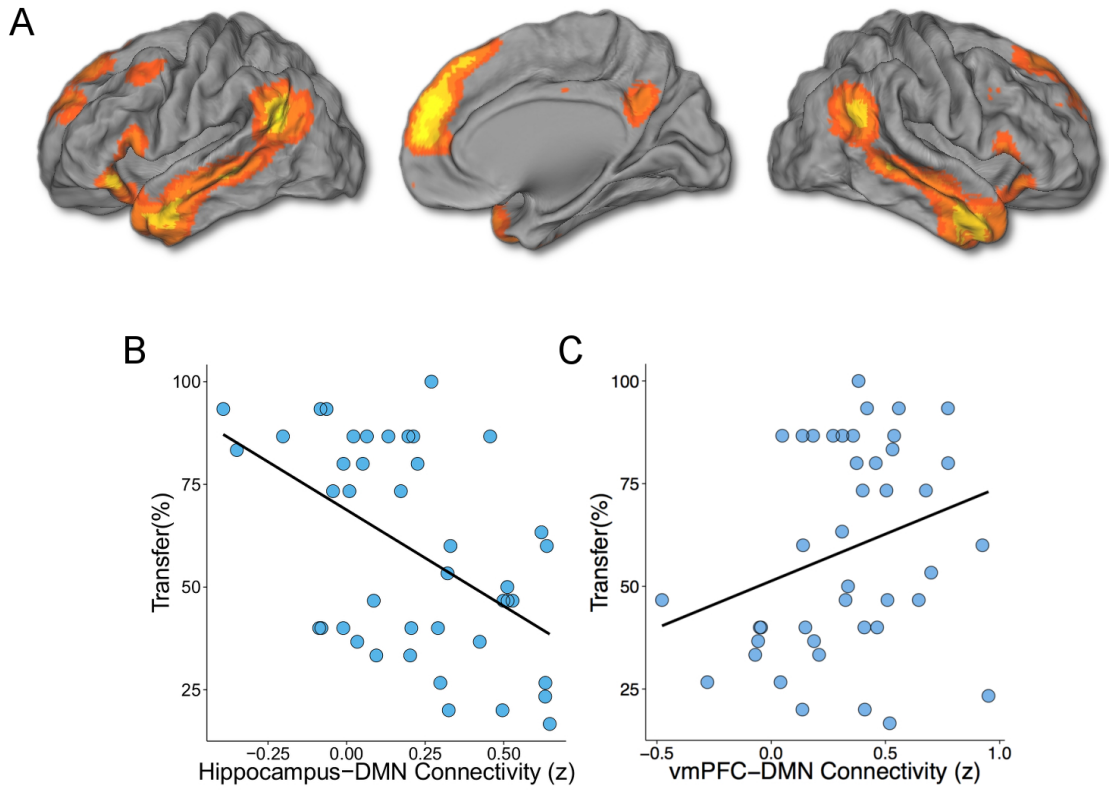

Figure 2.3 Intrinsic Default-Mode Network (DMN) connectivity with the vmPFC and the hippocampus is associated with transfer. (A) Default mode network as revealed by spatial ICA. Intrinsic functional connectivity between the DMN and (B) the hippocampus and (C) the ventromedial prefrontal cortex (vmPFC) was significantly related to transfer of learned value. Regression lines are from linear models containing DMN-hippocampus and DMN-vmPFC connectivity as predictors. Both effects remained significant when including vmPFC-hippocampus connectivity as a covariate and both effects are selective to transfer.

connectivity between the vmPFC and the DMN was positively associated with transfer. Both effects remained significant when including hippocampus-vmPFC connectivity in the model as a nuisance covariate. In keeping with the vmPFC's position as a network hub (Buckner et al., 2008; Roy et al., 2012; Sreenivas et al., 2012), its connectivity with the FPN was also significantly negatively associated with transfer (**Figure 2.4;** β=-0.38, p=0.02, adjusted r^2=0.20). This effect remained significant when including hippocampus-

vmPFC connectivity as a nuisance covariate. There were no significant effects of connectivity between the cingulo-opercular network and either the hippocampus or vmPFC.

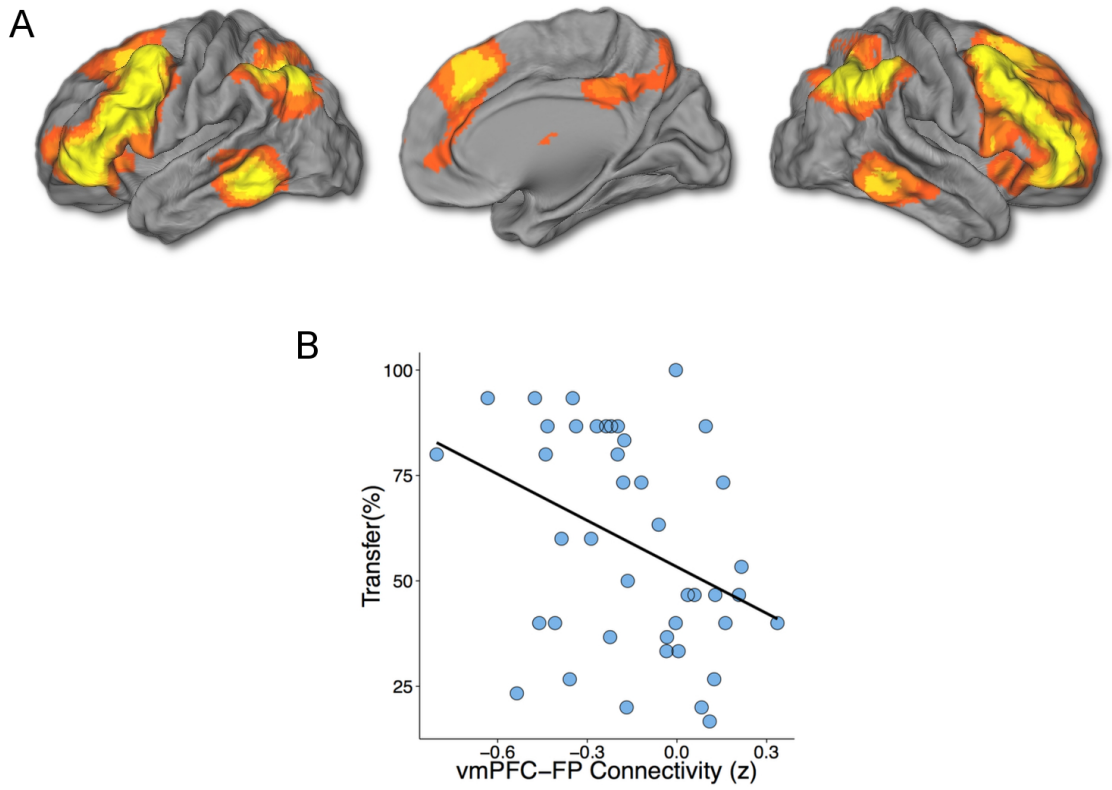

Figure 2.4 VMPFC connectivity with the frontoparietal network is negatively associated with transfer. (A) Frontoparietal network. (B) Plot showing negative relationship between value transfer and vmPFC-FPN connectivity. The regression line was computed from a linear model containing FPN-hippocampus and FPN-vmPFC connectivity as predictors. This effect remained significant when including vmPFC-hippocampus connectivity as a covariate and is specific to transfer.

Finally, we tested the selectivity of the above effects to transfer. To control for the possibility that this pattern of results broadly reflect attentional or learning tendencies, we repeated the analyses using each participant's conditioning score (i.e. the proportion of choosing S_{2+} and avoiding S_{2-}), which reflects direct learning of stimulus-outcome associations from the Reward learning phase. Importantly, the correlation between hippocampus-vmPFC connectivity and learning was selective to transfer and was not found for conditioning ($r(40)$=-0.02, p=0.90, uncorrected). Moreover, a direct comparison of the

two overlapping correlations (Meng et al., 1992) was significant (Z=-1.85, p=0.03, one-tailed); this difference was also significant when tested using a non-parametric bootstrap to account for non-normality of behavioral performance. These results provide evidence for the specific contributions of this network to transfer. Associations with between-network connectivity were also specific: conditioning did not relate to DMN connectivity with either the vmPFC or the hippocampus (hippocampus: β=-0.02, p=0.85; vmPFC: β=-0.01, p=0.85; adjusted r^2=-0.05, p=0.96). The same was true for the effect of vmPFC-FPN connectivity on learning (β=-0.06, p=0.52, adjusted r^2=-0.06).

Discussion

Individuals vary in the extent to which they leverage past learning when faced with novel choices. The sensory preconditioning paradigm provides a unique opportunity to explore the stable neural basis of variability in the tendency to transfer spontaneously, even in contexts where it provides no explicit benefit. Here we find that individual variability in the transfer of learning is related to intrinsic connectivity between the hippocampus and the vmPFC, measured during rest and on a separate day. Further, we find that individual differences in this behavior are also related to overall connectivity between these regions and the DMN and FPN. While multiple studies have demonstrated a relationship between task-evoked activation and differences in transfer (Shohamy and Wagner, 2008; Wimmer et al., 2012; Wimmer and Shohamy, 2012; Zeithamova et al., 2012), these studies do not address the potential role of intrinsic functional brain organization. Based on these studies, variability in transfer could have related solely to task-evoked patterns of brain activity, and not to *intrinsic* connectivity. The finding of a link

between these two measures advances our understanding of the role of network connectivity in behavioral variability and suggests that stable and intrinsic aspects of functional brain organization can explain individual differences in transfer of learning.

These results thus demonstrate a novel link between intrinsic connectivity and complex measures of learning. Further, their specificity to transfer of value as opposed to first-order conditioning indicates a relationship between resting-connectivity and behavior that is particular to regions involved in a given domain. While we used sICA to functionally isolate intrinsic networks, the selection of VMN regions was based on the *apriori* hypothesis that connectivity between regions known to be involved in transfer would be related to variability in this behavior. While much rs-fMRI research has focused on commonalities in normal populations, here we exploit the fact that individual connectivity is itself highly variable in healthy adults (Kahn and Shohamy, 2013) and show that this variability has implications for learning behavior.

Previous studies have separately highlighted roles for the hippocampus and the vmPFC in a wide range of behaviors that involve memory-guided learning (Port et al., 1987; Dusek and Eichenbaum, 1997; Myers et al., 2003; Greene et al., 2006; Kumaran et al., 2009; van Kesteren et al., 2010; Schoenbaum et al., 2011; Wimmer and Shohamy, 2012; Zeithamova et al., 2012). Together these regions are thought to contribute to the formation and use of flexible representations that guide transfer behavior. In the sensory preconditioning paradigm, we previously found in humans that hippocampus activation during learning is related to subsequent transfer (Wimmer and Shohamy, 2012). Jones and colleagues separately demonstrated the necessary role of a homologue to human medial orbitofrontal cortex in transfer behavior using the same paradigm in rodents (Jones et

al., 2012). The current results suggest that the strength of connectivity between the hippocampus and vmPFC at rest explains a significant portion of the variability in this kind of transfer, and links across these pieces of evidence in different species performing the same task.

The current results indicate that the correlation between hippocampal-vmPFC connectivity and transfer is negative, such that participants transferring the most are those with the weakest connectivity. The direction of this effect is consistent with a recent proposal about how connections between the hippocampus and the vmPFC relate to the incorporation of new memories into existing abstract frameworks, called "schemas". Schemas also capture an element of flexible memory-guided behavior, and studies show that the extent to which an individual will generalize a schema to new events is negatively correlated with hippocampus-vmPFC connectivity (van Kesteren et al., 2010). Such findings are consistent with neurobiological evidence of dynamic changes in hippocampal-prefrontal circuits during learning (Doyère et al., 1993; Takita et al., 1999) and have led to the proposal that the vmPFC may serve as a gateway to hippocampal learning (van Kesteren et al., 2012). Given the convergence between our results and previous findings, it is possible that the negative relationship we show reflects an overlapping learning mechanism to that proposed by van Kesteren and colleagues (van Kesteren et al., 2012). If prefrontal gating supports multiple forms of transfer, intrinsic connectivity between the hippocampus and vmPFC could be a marker of the tendency to transfer learning across a range of tasks. If so, then it would be expected that groups who show decreased transfer, such as patients with schizophrenia (Ivleva et al., 2012), may actually display increased vmPFC-hippocampus connectivity at rest.

Sensory preconditioning provides a measure of implicit transfer of value across learned associations. The findings we report are broadly relevant for other classes of paradigms that test transfer or generalization of knowledge (Kumaran and McClelland, 2012), many of which involve generalization of stimulus-stimulus associations outside the domain of reward learning. Yet other related processes have been examined within the context of reward-based decision-making, in studies of the neural and cognitive mechanisms underlying learning about structured regularities in the environment, referred to as "model-based" reinforcement learning (Daw et al., 2005; Hampton et al., 2006; Daw et al., 2011). This type of learning relies on prospective models of environmental contingencies that can flexibly guide decision-making, as opposed to the rigid and feedback-dependent updating of action or stimulus value, referred to as "model-free" learning (Doll et al., 2012). Current evidence suggests that there is significant individual variability in the respective weights placed on these types of learning (Daw et al., 2011). Although memory-based generalization and model-based frameworks vary in their origins, there is evidence for shared behavioral mechanisms (Doll et al., 2015). Furthermore, emerging findings across these domains reveal an important role for the hippocampus and vmPFC (McDannald et al., 2011; Zeithamova et al., 2012; Bornstein and Daw, 2013). Together with the current results, this suggests that variability in intrinsic functional connectivity between the hippocampus and the vmPFC may relate to a broader range of stable learning behaviors that involve a common mechanism.

In the present study we purposefully measure intrinsic connectivity in the absence of a task in order to address questions regarding variability in transfer behavior regardless

of task-evoked activity. In general, the link between activation evoked by task demands and spontaneous fluctuations remains an open question. Emerging evidence suggests that resting and evoked activity may negatively interact (He, 2013). Such a tradeoff between baseline and event-related connectivity presents a challenge to drawing inferences between these two conditions. These questions are important for interpreting studies of intrinsic connectivity, and should be tested in future studies that directly compare evoked and baseline conditions.

Learning and memory research has historically focused on distinguishing the roles of discrete and functionally specific nodes. The present findings indicate an important but unexplored role for large-scale functional networks in complex learning behavior. Our results highlight the relationship between learning and intrinsic connectivity within a ventral medial component of the default network centered on the hippocampus and the vmPFC – regions known to be involved in a host of flexible learning behaviors. In addition, we demonstrate that broader connectivity between these regions and the DMN and FPN, whose functions are currently much less clear, also relates to transfer. Taken together, these findings raise the possibility that communication between localized brain regions involved in specific cognitive functions and distributed domain-general networks may play an important role in learning. These results also suggest that this approach has potential implications for the use of rs-fMRI to further our understanding of individual differences in a range of basic learning behaviors. Because variation in these processes is likely to underlie complex behavioral traits, these results hold additional promise for answering questions about stable differences in complex behavior that are central to research in education as well as studies of clinical populations.

Chapter 3

Dynamic flexibility in striatal-cortical circuits supports reinforcement learning

A version of Chapter 3 was published in the Journal of Neuroscience: Gerraty, R. T., Davidow, J. Y., Foerde, K., Galvan, A., Bassett, D. S., & Shohamy, D. (2018). *Dynamic flexibility in striatal-cortical circuits supports reinforcement learning. Journal of Neuroscience, 38(10), 2442-2453.*

Abstract

Complex learned behaviors must involve the integrated action of distributed brain circuits. While the contributions of individual regions to learning have been extensively investigated, much less is known about how distributed brain networks orchestrate their activity over the course of learning. To address this gap, we used fMRI combined with tools from dynamic network neuroscience to obtain time-resolved descriptions of network coordination during reinforcement learning. We found that learning to associate visual cues with reward involves dynamic changes in network coupling between the striatum and distributed brain regions, including visual, orbitofrontal, and ventromedial prefrontal cortex (n=22, 13 females). Moreover, we found that this flexibility in striatal network coupling correlates with participants' learning rate and inverse temperature, two parameters derived from reinforcement learning models. Finally, we found that episodic learning, measured separately in the same participants at the same time, was related to dynamic connectivity in distinct brain networks. These results suggest that dynamic changes in striatal-centered networks provide a mechanism for information integration during reinforcement learning.

Introduction

Learning from reinforcement is central to adaptive behavior and requires integration of sensory, motor, and reward information over time. Major progress has been made in understanding how individual brain regions support reinforcement learning. However, little is known about how these brain regions interact, how their interactions change over time, and how these dynamic network-level changes relate to successful learning.

In a reinforcement learning task, participants use feedback over many trials to associate choices with probable outcomes (O'Doherty et al., 2003; Frank et al., 2004; Shohamy et al., 2004b). Computationally, this is captured by "model-free" learning algorithms, which provide a mechanistic framework for describing behavior (Sutton and Barto, 1998; Daw et al., 2005; Daw, 2011). These models account for neuronal signals underlying learning (Schultz et al., 1997; O'Doherty et al., 2003; Daw et al., 2006), demonstrating a role for the striatum and its dopaminergic inputs in updating reward predictions. However, to support reinforcement learning, the striatum must also integrate visual, motor, and reward information over time. Such a process is likely to involve coordination across a number of different circuits interconnected with the striatum.

The idea that the striatum serves an integrative role is not new (Kemp and Powell, 1971; Bogacz and Gurney, 2007; Hikosaka et al., 2014; Ding, 2015). The striatum is anatomically well positioned for integration: it receives input from many cortical areas and projects back to motor cortex (Alexander et al., 1986; Haber, 2003; Haber et al., 2006; Haber and Knutson, 2010). However, while the idea that the striatum serves such a role is theoretically appealing, it has been difficult to test empirically. Thus, it remains unknown

how the striatum interacts with cortical regions representing sensory, motor and value signals and how these network-level interactions reconfigure over the course of learning.

Here we aimed to address this gap, utilizing the emerging field of dynamic network neuroscience (Kopell et al., 2014; Medaglia et al., 2015). This area of research has been spurred by the development of tools like multi-slice community detection (Mucha et al., 2010) that can infer activated circuits and their reconfiguration from neuroimaging data. These tools have been leveraged to understand the role of dynamic connectivity in motor learning (Bassett et al., 2011; Bassett et al., 2013b; Bassett et al., 2015). A key measure is an index of a brain region's tendency to communicate with different networks over time, known as "flexibility" (Bassett et al., 2011). Prior work has shown that flexibility across a number of brain regions predicts individual differences in the acquisition speed on a simple motor task (Bassett et al., 2011), and it has been related to working memory performance (Braun et al., 2015). But its role in updating choice behavior based on reinforcement is not known.

We hypothesized that network dynamics, indexed by flexibility, support key processes underlying reinforcement learning; specifically, that reinforcement learning is associated with increases in dynamic coupling between the striatum and cortical regions processing visual and value information. We predicted that (1) reinforcement learning would involve increased flexible network coupling between the striatum and distributed brain circuits; and (2) that these circuit changes would be related to measurable changes in behavior, specifically *learning performance* (accuracy, within subjects) as well as *learning rate* and *inverse temperature*, individual difference measures derived from reinforcement learning models.

Our final prediction concerned the relationship between flexibility and episodic memory for individual events. The rationale for testing episodic memory was two-fold. First, it provided a comparison for time-on-task effects. Second, it was a question of interest given that little is known about how episodic memory is supported by brain networks. Given the extensive literature indicating that separate brain regions support episodic memory vs. reinforcement learning (Knowlton et al., 1996; Myers et al., 2003; Foerde et al., 2012; Doll et al., 2015), we predicted that (3) distinct medial temporal and prefrontal regions would exhibit a relationship between network flexibility and episodic memory.

Materials and Methods

To test these hypotheses, we used functional magnetic resonance imaging (fMRI) to measure changes in brain network structure while participants engaged in a reinforcement learning task (**Figure 3.1A**). On each trial, participants were presented with a visual cue, made a choice indicated by a key press, and then received feedback. Participants had up to 4 seconds to respond to each trial, and feedback was shown for 2 seconds. We used a task for which behavior has been well described by reinforcement learning models (Foerde and Shohamy, 2011a), and which is known from fMRI to involve the striatum (Foerde and Shohamy, 2011a) and from patient studies to depend on it (Foerde et al., 2012). The task also included trial-unique images presented during feedback, allowing us to test the role of network dynamics in episodic memory. Presentation of these images coincided with reinforcement, but they were incidental to the learning task (**Figure 3.1B**).

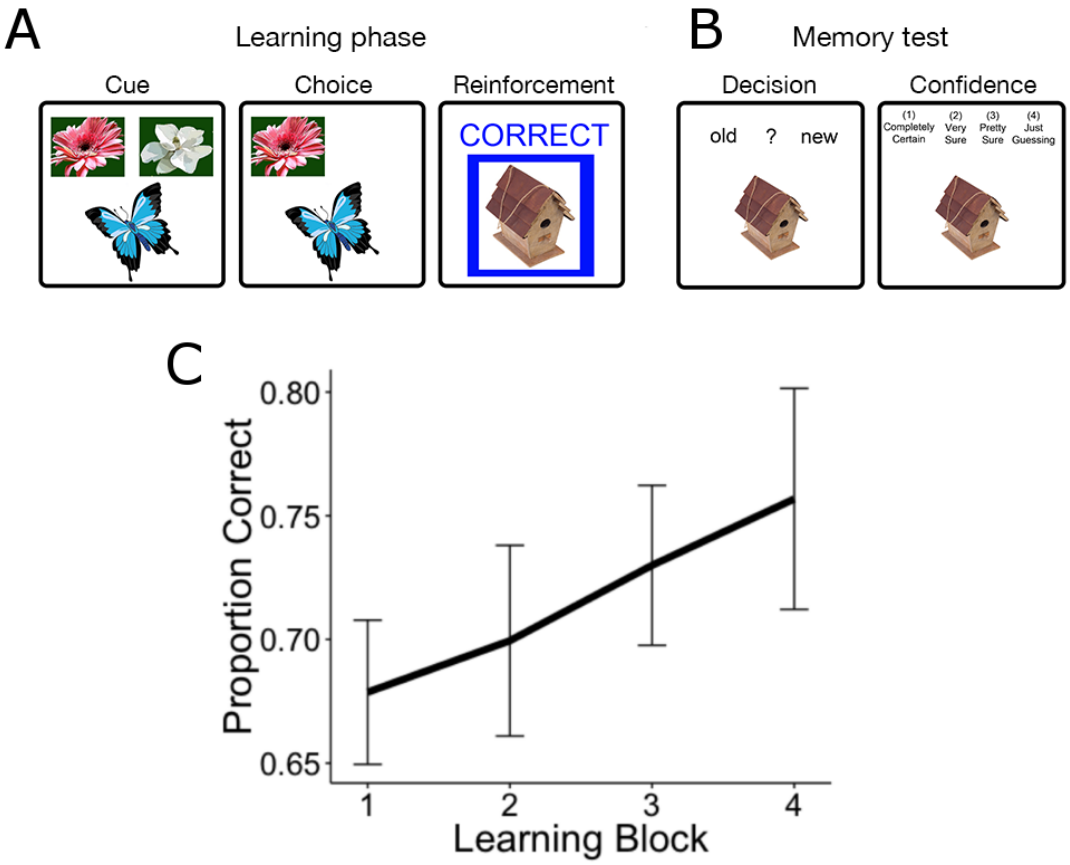

Figure 3.1 Task design and learning performance. Participants performed a modified reinforcement learning task while undergoing fMRI (Foerde and Shohamy, 2011a). **A.** Participants were instructed to associate each of 4 cues (butterflies) with one of two outcomes (flowers). Feedback was probabilistic, with positive feedback following the choice on 80% of correct trials and on 20% of incorrect trials. **B.** Each feedback event was presented with a unique image. Thirty minutes following the MRI scan, participants were given a surprise episodic memory test, testing recognition and confidence for images seen during the scan, intermixed with novel images. **C.** Average performance on the learning task improved linearly, suggesting continuous learning across all trials.

Experimental Design. Twenty-five healthy right-handed adults (age 24-30 years, mean of 27.7, standard deviation of 2.0, 13 females) were recruited from the University of California Los Angeles and the surrounding community as the adult comparison sample in a developmental study of learning (Davidow et al., 2016). All participants provided informed consent in writing to participate in accordance with the UCLA Institutional Review Board, which approved all procedures. Individuals were paid for their participation.

Participants reported no history of psychiatric or neurological disorders, or contraindications for MRI scanning. Three subjects were excluded from this analysis (two for technical issues in behavioral data collection and one for an incidental neurological finding), leading to a final sample size of 22.

Task and Behavioral Analysis. The probabilistic learning task administered to subjects undergoing an fMRI session has been previously described (Foerde and Shohamy, 2011a; Foerde et al., 2012; Foerde et al., 2013; Davidow et al., 2016). Before scanning, participants completed a practice round of 8 trials to become familiar with the task. On each trial, participants were presented with an image of one of four butterflies along with two flowers, and asked to indicate which flower the butterfly was likely to feed from, using a left or right button press. The four learning blocks were followed by a test phase, in which subjects performed the same butterfly task without feedback for 32 trials. They were given feedback consisting of the words 'Correct' or 'Incorrect'. Presentation of feedback also included an image of an object unique to each trial, shown in random order for the purpose of subsequent memory testing. For each butterfly image, one flower represented the 'optimal' choice, with a 0.8 probability of being correct, while the alternative flower had a 0.2 probability of being followed by correct feedback. Subjects performed four blocks of this probabilistic learning phase, each consisting of 30 trials. Feedback was presented for 2 seconds, and was followed by a randomly jittered inter-trial interval.

For each trial in the learning phase, we recorded the feedback the participant actually received as well as whether the optimal choice was made, and we computed the percent correct for each block based on the percent of trials on which subjects made the op-

timal choice (regardless of actual feedback). These variables enable a characterization of learning as the proportion of optimal choices in each block, as well as that in the test phase. Using this information, we fit reinforcement learning models to subjects' decisions (Sutton and Barto, 1998; Daw, 2011), utilizing a hierarchical Bayesian approach to pool uncertainty across subjects and aid in model identifiability.

Briefly, the expected value for a given choice at time *t*, Q_t, is updated based on the reinforcement outcome r_t via a prediction error δ_t:

$$Q_{t+1} = Q_t + \alpha\delta_t$$

$$\delta_t = r_t - Q_t.$$

The reinforcement learning models included two free parameters, α and β. The learning rate α is a parameter between 0 and 1 that measures the extent to which value is updated by feedback from a single trial. Higher α indicates more rapid updating based on few trials and lower α indicates slower updating based on more trials. Another parameter fit to each subject is the inverse temperature parameter β, which determines the probability of making a particular choice using a softmax function (Ishii et al., 2002; Daw, 2011), so that the probability of choosing choice 1 on trial *t* would be:

$$p(c_t = 1|\alpha, \beta) = \frac{e^{\beta Q_{1t}}}{e^{\beta Q_{1t}} + e^{\beta Q_{2t}}},$$

where $p(c_t = 1)$ refers to the probability of choice one and Q_{1t} is the value for this choice on trial *t*.

Reinforcement learning models of this form have known issues with identifiability (Gershman, 2016). To constrain the parameter space and reduce noise, we fit a hierarchical Bayesian model, which regularizes this estimation with empirical prior distributions on α and β (Daw, 2011):

$$\beta \sim Gamma(b1, b2)$$

$$\alpha \sim Beta(a1, a2),$$

where b1, a1, and a2 are shape parameters, and b2 is a scale parameter. By fitting prior parameters as part of the model, individual-level likelihood parameters are constrained by group average distributions. These group parameters were themselves regularized by weakly informative hyperprior distributions ($Cauchy^{+}(0,5)$ in all cases). Models were fit using Hamiltonian Markov Chain Monte Carlo in Stan (Carpenter et al., 2015). In addition to the benefit of constraining the parameter space, this approach produces a posterior distribution of all parameters, which incorporates uncertainty at both group and individual levels in parameter estimation, and also allows for the consideration of all plausible values of RL parameters in subsequent analyses, rather than relying on point estimates or Gaussian assumptions.

Following the fMRI session (approximately 30 minutes), subjects were given a surprise memory test for the trial-unique object images presented during feedback in the learning phase. Subjects were presented with all 120 objects shown during the conditioning phase, along with an equal number of novel objects, and asked to judge the images as "old" or "new". They were also asked to rate their confidence for each decision on a scale

of 1-4 (one being most confident; four indicating "guessing"). All responses rated 4 were excluded from our analyses (Foerde and Shohamy, 2011a).

MRI Acquisition and Preprocessing. MRI images were acquired on a 3 T Siemens Tim Trio scanner using a 12-channel head coil. For each block of the learning phase of the conditioning task, we acquired 200 interleaved T2*-weighted echo-planar (EPI) volumes with the following sequence parameters: TR = 2000 ms; TE = 30 ms; flip angle (FA) = 90°; array =64 x 64; 34 slices; effective voxel resolution = 3x3x4 mm; FOV = 192 mm). A high resolution T1-weighted MPRAGE image was acquired for registration purposes (TR = 2170 ms, TE = 4.33 ms, FA = 7º, array = 256 x 256, 160 slices, voxel resolution = 1 mm^3, FOV = 256). In addition, as part of the original study for which these data were collected, two resting state scans were acquired - one before and one after the learning task. Resting state scans were acquired with identical sequence parameters to the EPI scans described above, except that each scan consisted of 154 images (308 s). We utilized these scans to test the specificity of our results to learning, and to demonstrate the reliability of our dynamic connectivity metric across multiple scans (see *Dynamic Network Statistics- Flexibility and Allegiance*, below).

Functional images were preprocessed using FSL's FMRI Expert Analysis Tool (FEAT, (Smith et al., 2004). Images from each learning block were high-pass filtered at $f > 0.008$ Hz, spatially smoothed with a 5mm FWHM Gaussian kernel, grand-mean scaled, and motion corrected to their median image using an affine transformation with tri-linear interpolation. The first three images were removed to account for saturation effects. Functional and anatomical images were skull-stripped using FSL's Brain Extraction

Tool. Functional images from each block were co-registered to subject's anatomical images and non-linearly transformed to a standard template (T1 Montreal Neurological Institute template, voxel dimensions 2 mm^3) using FNIRT (Andersson et al., 2008). Following image registration, time courses were extracted for each block from 110 cortical and subcortical regions of interest (ROIs) segmented from FSL's Harvard-Oxford Atlas. Due to known effects of motion on measures of functional connectivity (Power et al., 2012; Satterthwaite et al., 2012a), time courses were further preprocessed via a nuisance regression. This regression included the six translation and rotation parameters from the motion correction transformation, average CSF, white matter, and whole brain time courses, as well as the first derivatives, squares, and squared derivatives of each of these confound predictors (Satterthwaite et al., 2012b). Images acquired during learning and at rest underwent identical preprocessing.

Dynamic Connectivity Analysis. To assess dynamic connectivity between the ROIs, time courses were further subdivided into sub-blocks of 25 TRs each. The selection of 25 TRs represents a compromise between the precision and range of frequencies sampled for each network, the number of networks per estimate of flexibility, and the number of flexibility measurements to compare to within-subject behavior. Smaller window sizes are more sensitive to short term-dynamics (including task-evoked changes) and also to individual differences in dynamic connectivity. In contrast, larger window sizes offer more precise estimates of connectivity and are more sensitive to inter regional variation (Telesford et al., 2016). Because we were interested in variation across time, regions, and subjects, we selected windows 25 TR (50s) in duration, representing a middle ground be-

tween time windows that exhibit high levels of temporal and individual variability (~25s) and those that show high levels of inter-regional variability (~75s).

For each 25 TR sub-block, connectivity was quantified as the magnitude-squared coherence between each pair of ROIs at f = 0.06-0.12 Hz in order to later assess modularity over short time windows in a manner consistent with previous reports (Bassett et al., 2011; Bassett et al., 2013b):

$$C_{xy}(f) = \frac{|G_{xy}(f)|^2}{G_{xx}(f)G_{yy}(f)} ,$$

where $G_{xy}(f)$ is the cross-spectral density between regions x and y, and $G_{xx}(f)$ and $G_{yy}(f)$ are the autospectral densities of signals x and y, respectively. We thus created subject-specific 110 x 110 x 32 connectivity matrices for 110 regions and 8 time windows for each of the 4 learning blocks, containing coherence values ranging between 0 and 1. The frequency range of 0.06-0.12 Hz was chosen to approximate the frequency envelope of the hemodynamic response, allowing us to detect changes as slow as 3 cycles per window with a 2 second TR. We selected this frequency band based on previous work showing that high frequency associations may not map on to task-evoked changes in connectivity (Sun et al., 2004). This observation follows if one assumes that the canonical hemodynamic response function serves as a low pass filter of any high frequency coupling. There is increasing reason to be suspicious of this assumption (see Chen and Glover, 2015; Zhang et al., 2016) for examples examining static and dynamic connectivity, respectively), but the processes leading to interactions measured by high-frequency BOLD signals are not well understood. For this reason, we decided to follow previous

studies of task-based flexibility (e.g. Bassett et al. 2011) by focusing on coupling in the frequency range associated with the canonical hemodynamic response.

Each connectivity matrix was treated as an unthresholded graph or network, in which each brain region is represented as a network node, and each functional connection between two brain regions is represented as a network edge (Bullmore and Sporns, 2009; Bullmore and Bassett, 2011). In the context of dynamic functional connectivity matrices, the network representation is a *temporal network*, which is an ensemble of graphs that are ordered in time (Holme and Saramäki, 2011). If the temporal network contains the same nodes in each graph, then the network is said to be a multilayer network where each layer represents a different time window (Kivelä et al., 2014). The study of topological structure in multilayer networks has been the topic of considerable study in recent years, and many graph metrics and statistics have been extended from the single-network representation to the multilayer network representation. Perhaps one of the single most powerful features of these extensions has been the definition of so-called *identity links*, a new type of edge that links one node in one time slice to itself in the next time slice. These identity links hard code node identity throughout time, and facilitate mathematical extensions and statistical inference in cases that had previously remained challenging.

Uncovering Evolving Circuits Using Multi-slice Community Detection. To extract modules or communities from a single-network representation, one typically applies a community detection technique such as modularity maximization (Newman, 2004). However, these single-network algorithms do not allow for the linking of communities across timepoints, thus hampering statistically robust inference regarding the reconfiguration of

communities as the system evolves (Mucha et al., 2010). In contrast, the multilayer approaches allow for the characterization of multi-layer network modularity, with layers representing time windows. In this framework, each network node in the multi-layer network is connected to itself in the preceding and following time windows in order to link networks in time. This enables us to solve the community-matching problem explicitly within the model (Mucha et al., 2010), and also facilitates the examination of module reconfiguration across multiple temporal resolutions of system dynamics (Bassett et al., 2013a). We thus constructed multilayer networks for each subject, allowing for the partitioning of each network into communities or modules whose identity is robustly tracked across time windows.

While many statistics are available to the researcher to characterize network organization in temporal and multilayer networks, it is not entirely clear that all of these statistics are equally valuable in inferring neurophysiologically relevant processes and phenomena (Medaglia et al., 2015). Indeed, many of these statistics are difficult to interpret in the context of neuroimaging data, leading to confusion in the wider literature. A striking contrast to these difficulties lies in the graph-based notion of modularity or community structure (Newman, 2004), which describes the clustering of nodes into densely interconnected groups that are referred to as *modules* or *communities* (Porter et al., 2009; Fortunato, 2010). Recent and convergent evidence demonstrates that these modules can be extracted from rest and task-based fMRI data (Meunier et al., 2010; Cole et al., 2014), demonstrate strong correspondence to known cognitive systems (including default mode, fronto-parietal, cingulo-opercular, salience, visual, auditory, motor, dorsal attention, ventral attention, and subcortical systems (Power et al., 2011; Yeo et al.,

2011)), and display non-trivial re-arrangements during motor skill acquisition (Bassett et al., 2011; Bassett et al., 2015) and memory processing (Braun et al., 2015). These studies support the utility of module-based analyses in the examination of higher order cognitive processes in functional neuroimaging data.

The partitioning of these multilayer networks into temporally linked communities was carried out using a Louvain-like locally greedy algorithm for multilayer modularity optimization (Mucha et al., 2010). The multilayer modularity quality function is given by

$$Q_{ml} = \frac{1}{2\mu}\sum_{ijlr}\left\{\left(A_{ijl} - \gamma_l \frac{k_{il}k_{jl}}{2m_l}\right)\delta_{lr} + \delta_{ij}C_{jlr}\right\}\delta(g_{il}, g_{jr}),$$

where Q_{ml} is the multilayer modularity index. The adjacency matrix for each layer *l* consists of components A_{ijl}. The variable γ_l represents the resolution parameter for layer *l*, while C_{jlr} gives the coupling strength between node *j* at layers *l* and *r* (see below for details of fitting these two parameters). The variables g_{il} and g_{jr} correspond to the community labels for node *i* at layer *l* and node *j* at layer *r*, respectively; k_{il} is the connection strength (in this case, coherence) of node *i* in layer *l*; $2\mu = \sum_{jr}\kappa_{jr}$; the multilayer node strength $\kappa_{jl} = k_{jl} + c_{jl}$; and $c_{jl} = \sum_r C_{jlr}$. Finally, the function $\delta(g_{il}, g_{jr})$ refers to the Kronecker delta function, which equals 1 if g_{il}=g_{jr}, and 0 otherwise.

Resolution and coupling parameters (γ_l and C_{jlr}, respectively) were selected using a grid search formulated explicitly to optimize Q_{ml} relative to a temporal null model (Bassett et al., 2013a). The temporal null model we employed is one in which the order of time windows in the multilayer network was permuted uniformly at random. Thus, we

performed a grid search to identify the values of γ_l and C_{jlr} that maximized $Q_{ml} - Q_{null}$, following (Bassett et al., 2013a). We selected this objective function to maximize the distance between multi-slice modularity when temporal information is removed from the network. Grid searches are visualized in **Supplementary Figure 3.1.** To ensure statistical robustness, we repeated this grid search 10 times. To maximize the stability of resolution and coupling, each subject's parameters were treated as random effects, with the best estimate of resolution and coupling generated by averaging across-individual subject estimates. This is a similar approach to that taken in computational modeling of reinforcement learning, in which learning rate and temperature parameters are averaged in order to generate prediction error estimates (Daw, 2011). With this approach, we estimated the optimal resolution parameter γ to be 1.18 (standard deviation of 0.61) and the coupling parameter C to be 1 (this was the optimal parameter for all subjects for all iterations). These values are quite similar to those chosen *a priori* (usually setting both parameters to unity) in previous reports (Bassett et al., 2011).

Finally, we note that maximization of the modularity quality function is NP-hard, and the Louvain-like locally greedy algorithm we employ is a computational heuristic with non-deterministic solutions. Due to the well known near-degeneracy of Q_{ml} (Good et al., 2010; Mucha et al., 2010; Bassett et al., 2013a), we repeated the multi-slice community detection algorithm 500 times using the resolution and coupling parameters estimated from the grid search procedure outlined above. This approach ensured an adequate sampling of the null distribution (Bassett et al., 2013a). Each repetition produced a hard partition of nodes into communities as a function of time window: that is, a community or

module allegiance identity for each of the 110 brain regions in the multilayer network. We used these community labels to compute flexibility and module allegiance statistics.

Dynamic Network Statistics—Flexibility and Allegiance. To characterize the dynamics of these temporal networks and their relation to learning, we computed the *flexibility* of each node, which measures the extent to which a region changes its community allegiance over time (Bassett et al., 2011). Intuitively, flexibility can be thought of as a measure of a region's tendency to communicate with different networks during learning. Flexibility is defined as the number of times a node displays a change in community assignment over time, divided by the number of possible changes (equal here to the number of time windows in a learning block minus 1). This was computed for each region in each block. In addition, average measures of flexibility were computed across the brain and across all blocks. We also computed the *module allegiance* of each ROI with respect to regions of the striatum during each learning block. Module allegiance is the proportion of time windows in which a pair of regions is assigned the same community label, and thus tracks which regions are most strongly coupled with each other at a given point in time. To obtain stable estimates, we averaged both flexibility and allegiance scores for each ROI over the 500 iterations of the multilayer community detection algorithm. To test the reliability of flexibility as a measure of dynamic connectivity, we calculated the correlation between flexibility averaged across the striatum (i) during the first resting state scan (prior to learning) and (ii) during the second resting state scan (after learning).

We hypothesized that flexibility would be positively related to learning as measured by performance across blocks of the task. To examine the effect of flexibility on

learning from feedback, we estimated a generalized mixed-effects model predicting optimally correct choices with flexibility estimates for each block with a logistic link function, using the Maximum Likelihood (ML) approximation implemented in the *lme4* package (Bates et al., 2015). Thus, each participant's average flexibility in an *a priori* striatum ROI was calculated for each learning block and was used to predict the proportion of optimal choices in each block. The ROI included bilateral caudate, putamen, and nucleus accumbens regions from the Harvard-Oxford atlas. We included a random effect of subject, allowing for different effects of flexibility on learning for each subject, while constraining these effects with the group average. Average flexibility across sessions was included as a fixed effect in the model to ensure that our estimates represented within-subject learning effects. Additionally, to examine distinct effects in different striatal sub-regions, we included striatal ROI as a varying effect. The maximum likelihood estimate of the variance by region was 0, indicating that there is little variability across regions and not enough data to distinguish these small effects. To further explore differences in this effect across sub-regions of the striatum, we fit separate models for each sub-region ROI, essentially assuming that this inter-regional variance is infinite. Even with this assumption, there were no significant differences between estimates for any pair of striatal regions. We also estimated this relationship between performance and whole-brain flexibility, which has been related to several cognitive functions in previous reports (Bassett et al., 2011; Braun et al., 2015). To rule out motion influences, we also included the block-level averages of the root mean squared relative displacement, a common output measure from motion correction preprocessing pipelines.

To provide appropriate posterior inference about the plausible parameter values indicated by our data, and to account for uncertainty about all parameters, we also fit a fully Bayesian extension of the ML approximation described above for the effect of striatal flexibility on learning performance. We used the 'brms' package for fitting flexibility-performance models in the Stan language (Carpenter et al., 2015). These were similar to the likelihood approximation models, but included a covariance parameter for subject-level slopes and intercepts (which could not be fit by the above approximation), and weakly informative prior distributions to regularize parameter estimation:

$$\beta \sim N(0, 10^2)$$

$$\tau \sim Cauchy^{+}(0, 5),$$

where β represents the "fixed effects" parameters (slope and intercept), τ represents the "random effects" variance for subject-level estimates sampled from β, and $Cauchy^{+}$ is a positive half-t distribution with one degree of freedom (Gelman, 2006b). Similarly, we used an lkj prior with $\eta = 2$ for correlations between subject-level intercept and slope estimates (Lewandowski et al., 2009). This approach also allowed us to fit and visualize subject-level estimates of the relationship between flexibility and performance.

To examine the relationship between flexibility and parameters estimated from reinforcement learning models, we tested whether striatal flexibility was correlated with the learning rate α and inverse temperature β for each subject, using Spearman's correlation coefficient due to the non-Gaussian distribution of these parameters. We hypothesized that inverse temperature, which is tightly linked to overall optimal choice performance, would be positively correlated with flexibility in the striatum. Learning rate has a more

complex relationship with performance on reinforcement learning tasks; different learning environments will afford distinct optimal learning rates. However, lower learning rates reflect wider temporal averaging, and are preferable once the correct choice is learned. In that sense, lower learning rates reflect more stable processing of sensory information indicating an optimal choice. Given the hypothesis that dynamic connectivity in the striatum underlies such processing, it is reasonable to suspect that lower learning rates would be associated with higher flexibility in this region. To account for joint uncertainty in these parameters at the group and subject level, correlations were computed over the full posterior distributions from the reinforcement-learning model. We computed correlations with flexibility in each sub-region of the striatum, as well as with flexibility averaged across the striatum. To ensure that these associations were specific to network dynamics evoked during learning, rather than reflecting intrinsic network characteristics unrelated to task performance, we also calculated the correlation coefficient between both reinforcement learning parameters and the average flexibility in the striatum during the resting state scan acquired before the learning task.

Our hypothesis that dynamic connectivity in the striatum allows for the integration of sensory and value information during learning led us to predict that, as decisions are learned, the striatum should increase its tendency to couple with relevant sensory (in this case visual) areas and regions processing value, such as the vmPFC. To determine which regions changed coupling with the striatum during the course of the task, we fit mixed effects models using learning block to predict log-transformed module allegiance. This analysis was computed first using the average of each ROI's allegiance across striatal sub-regions, and then separately for each sub-region of the striatum. In both cases this

analysis was carried out for each region of the brain, treating block as a factor so as to avoid assumptions about the linearity or direction of changes in allegiance. We controlled the false discovery rate across all ROI-striatum pairs.

To explore other regions exhibiting effects of dynamic connectivity on learning performance, we separately modeled the effect of flexibility in each brain region on reinforcement learning using the ML approximation implemented in the *lme4* package. We applied a false discovery rate correction for multiple comparisons across regions (Benjamini and Hochberg, 1995). While regions passing this threshold are reported, we also visualize the results using an exploratory uncorrected threshold of $p<0.05$.

To explore the relationship between network dynamics and other forms of learning, we also regressed flexibility statistics from each ROI against subsequent memory scores for the trial-unique objects presented during feedback. If the effects of striatal flexibility were relatively selective to incremental learning, we expected to find no significant association even at an uncorrected threshold with memory in the regions comprising our striatal ROI. In addition, this provided an exploratory analysis to examine the regions in which network flexibility plays a potential role in episodic memory. Given a host of previous studies on multiple learning systems, we reasoned it might be possible to detect an effect of dynamic network coupling on episodic memory in regions traditionally associated with this form of learning.

Results

Reinforcement learning performance. Participants learned the correct response for each cue. The percentage of optimal responses increased continuously from 68% in the first block to 76% in the final block, on average. Using a mixed-effects logistic model, we observed a significant effect of block on learning performance, as measured by the proportion of optimal responses during each block (**Figure 3.1C**; $\beta = 0.28$, Standard Error (S.E.) = 0.11, p = 0.01 (Wald approximation, (Bates et al., 2015). We also fit reinforcement learning models (Sutton and Barto, 1998; Daw, 2011) to participants' trial-by-trial choice behavior, utilizing hierarchical Bayesian models to aid estimation and pool information across subjects (Gershman, 2016).

We fit standard reinforcement learning parameters to subjects' choice data using a hierarchical Bayesian model. Fit was evaluated using the Widely Applicable Information Criterion, a measure of expected out-of-sample deviance (Watanabe, 2013; Vehtari et al., 2017). The full model fit better than a null model with no learning rate (WAIC difference = 894.2, S.E.= 47.4), as well as a model where a single learning rate (α) was estimated as fixed across all subjects (WAIC difference = 21.6, S.E. = 8.2). The average learning rate (α) was 0.41 with a standard deviation of 0.14; the average inverse temperature (β) was 3.84, with a standard deviation of 4.31. These α and β parameters provide a mechanistic probe of individual differences in learning, which allowed us to characterize the relationship between network dynamics and distinct sources of learning variability.

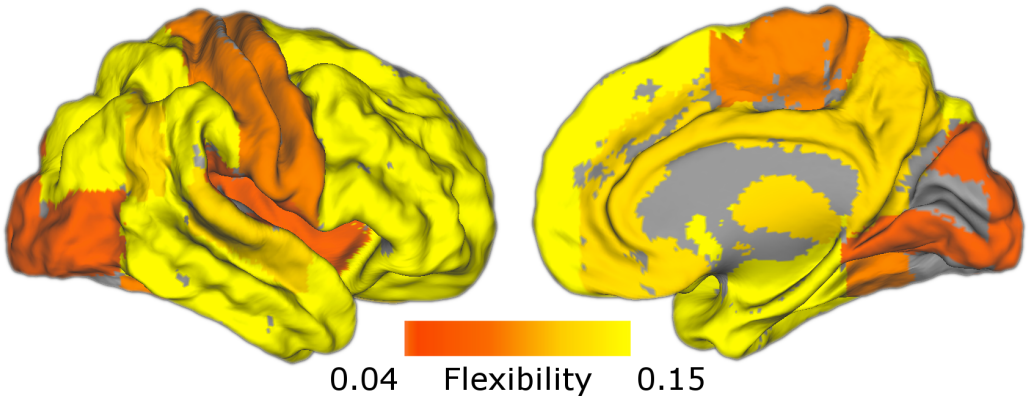

Figure 3.2. Spatial distribution of flexibility. Network flexibility was computed for each ROI in the Harvard-Oxford atlas. Here flexibility is averaged across learning blocks to visualize the spatial distribution. Flexibility is highest in subcortical regions and association cortex, and lowest in primary sensory areas. See **Supplementary Figure 3.1** for grid search for selecting resolution and coupling parameters of the multislice community detection algorithm used when computing flexibility.

Flexibility in the striatum relates to reinforcement learning. To characterize spatial and temporal properties of dynamic brain networks during the task, we constructed dynamic functional connectivity networks for each subject in 50 s windows, and used a recently developed multi-slice community detection algorithm (Mucha et al., 2010) to partition each network into dynamic communities: groups of densely connected brain regions that evolve in time. Our analyses included 110 cortical and subcortical ROIs from the Harvard-Oxford atlas, including bilateral nucleus accumbens, caudate, and putamen subregions of the striatum. We computed a flexibility statistic for each learning block, which measures the proportion of changes in each region's allegiance to large-scale communities over time (Bassett et al., 2011). Overall, flexibility was highest in association cortex and subcortical areas, and lowest in primary sensory and motor regions **(Figure 3.2).**

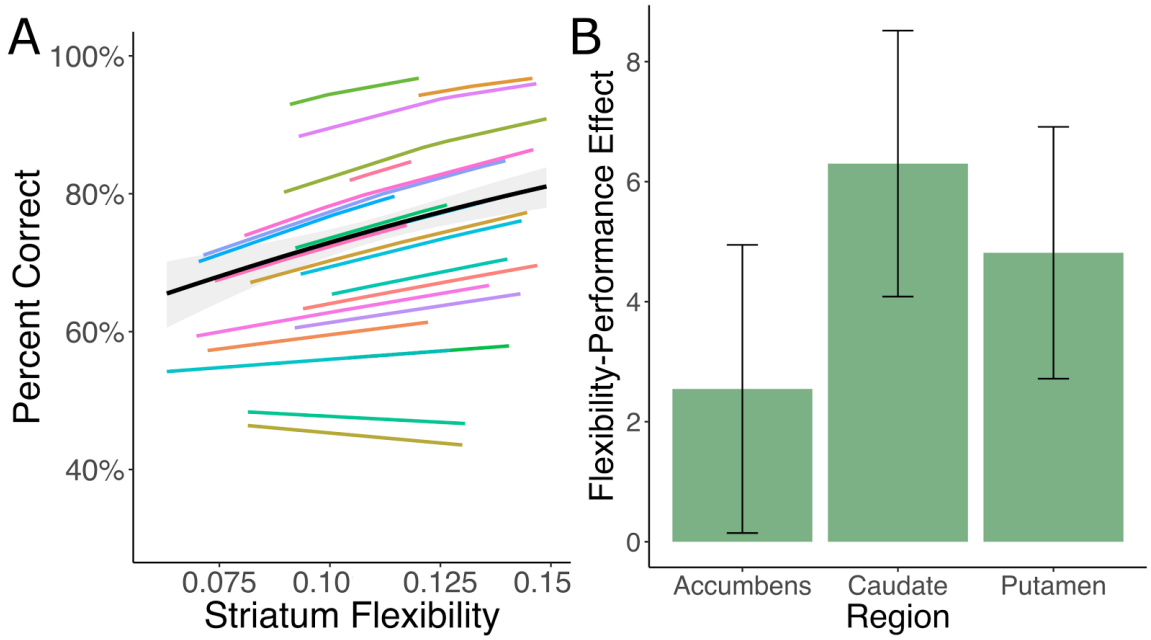

Figure 3.3 Flexibility in the striatum relates to learning performance within subjects. A. Mixed-effects model fit for the association between network flexibility in an *a priori* striatum ROI and learning performance. The black line represents the fixed effect estimate and the gray band represents the 95% confidence interval for this estimate. Color lines are predictions based on subject-level random effects estimates of the flexibility-performance relationship. See **Supplementary Figure 3.2** for a Bayesian extension of this model and full model fits with uncertainty for individual subjects. **B.** This effect was not distinguishable across regions of the striatum (bar plots and error bars represent estimates and standard errors from separate mixed-effects models for each ROI). **Supplementary Figure 3.3** shows the relationship between flexibility and performance with flexibility computed across different temporal window lengths.

To test whether flexibility in the striatum's network coupling is related to learning performance, we fit a mixed effects logistic regression (Bates et al., 2015) using average flexibility across the Harvard Oxford striatum ROIs during individual learning blocks to predict performance. Striatal flexibility computed for each block was significantly associated with the proportion of optimal responses (**Figure 3.3A**; $\beta = 9.45$, S.E. = 2.75, $p < 0.001$ (Wald approximation (Bates et al., 2015))). This effect could not be distinguished statistically across sub-regions of the striatum (**Figure 3.3B**). For appropriate posterior inference, we fit a Bayesian extension of this model (Carpenter et al., 2015) to generate a posterior 95% credible interval of [3.53, 14.99] (**Supplementary Figure 3.2**). To ensure that this approach reflected a within-subjects relationship between flexibility and learn-

ing, we included each subject's average flexibility across blocks in the model. This model produced similar results (β = 9.79, S.E. = 2.82, p = 0.0005), indicating that increases in dynamic striatal connectivity are associated with increased reinforcement-learning performance. Including average motion for each learning block in the model did not alter these results (β = 8.47, S.E. = 3.03, p = 0.005).

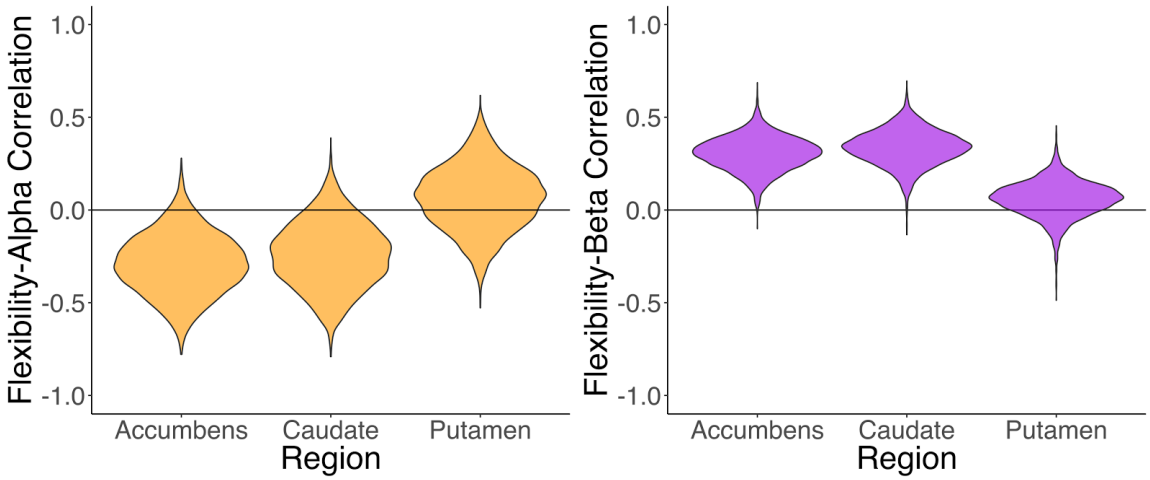

Figure 3.4. Flexibility in the striatum relates to reinforcement learning parameters across subjects. Violin plots showing posterior distributions of the correlation between parameters from RL models and flexibility in striatal regions. **A**. Learning rate, which indexes reliance on single trials for updating value, is negatively correlated with flexibility in the nucleus accumbens and caudate. **B.** Inverse temperature, which measures overall use of learned value, is positively correlated with flexibility in the same regions. Plotting the joint distribution and utilizing partial correlations indicate that these effects are separable (**Supplementary Figure 3.4**).

Individual differences in reinforcement learning parameters correlate with striatal flexibility. We next explored the relationship between flexibility and reinforcement learning model parameters, which account for individual differences in learning behavior. We were most interested in the learning rate α, which quantifies the extent to which individuals weigh feedback from single trials when updating the value of a choice (Sutton and Barto, 1998; Daw, 2011). Learning rate was negatively correlated with network flexibility in the nucleus accumbens (Spearman's correlation coefficient ρ= -0.29, p(ρ>0) = 0.04, **Figure 3.4A**) and to a lesser extent the caudate (ρ= -0.24, p(ρ>0) = 0.09 **Figure**

3.4A); that is, participants with a lower learning rate (indicating more integration of value across multiple trials) had more flexibility in these regions. Inverse temperature was positively correlated with flexibility in the same regions (accumbens ρ= 0.30, p(ρ<0) < 0.001; caudate ρ= 0.33, p(ρ<0) = 0.002), indicating that subjects relying more on learned value overall showed more dynamic striatal connectivity (**Figure 3.4B**). Averaging flexibility across all striatal regions, the correlation between learning rate and flexibility was ρ= -0.20 (p(ρ>0) = 0.12) and the correlation between inverse temperature and flexibility was ρ= 0.27 (p(ρ<0) = 0.007)**.** The effects of these two parameters were separable, as indicated by partial correlations and the joint posterior distribution (**Supplementary Figure 3.4**)**.**

Finally, we sought to address the specificity and reliability of these results. We took advantage of resting state data obtained in the same individuals. We found that striatal flexibility as measured during the resting state scan, prior to learning, was not related to individual differences in learning parameters (learning rate ρ = -0.01, p(ρ>0) = 0.54; inverse temperature ρ = 0.07, p(ρ<0) = 0.21), indicating that these effects are specific to dynamic connectivity evoked by learning. Additionally, we assessed the reliability of this measure by comparing striatal flexibility measured during a resting-state scan *before* the learning task with that measured in a resting-state scan *after* learning. We found that flexibility was positively correlated across scans (r=0.59, p=0.004) indicating significant stability for this measure of striatal connectivity dynamics.

Together, these results demonstrate that reinforcement learning involves dynamic changes in network structure centered on the striatum. They also suggest that distinct sources of individual differences in learning – reliance on individual trial feedback and

overall use of learned value – are related to differences in this dynamic striatal coupling. We next sought to examine *which* regions the striatum connects with during the task, and how such connections change over the course of learning.

Striatal allegiance with visual and value regions increases during learning. While an increase in flexible striatal network coupling is associated with learning within and between individuals, this leaves open the critical question of which regions are involved in this process. To address this question we used a dynamic graph theory metric known as *module allegiance*, which measures the extent to which each pair of regions shares a common network during a given time window (Bassett et al., 2015). Using the community labels described above, we first estimated the allegiance between the striatal sub-regions and every other ROI in the brain for each time window, and then probed their relationship to learning.

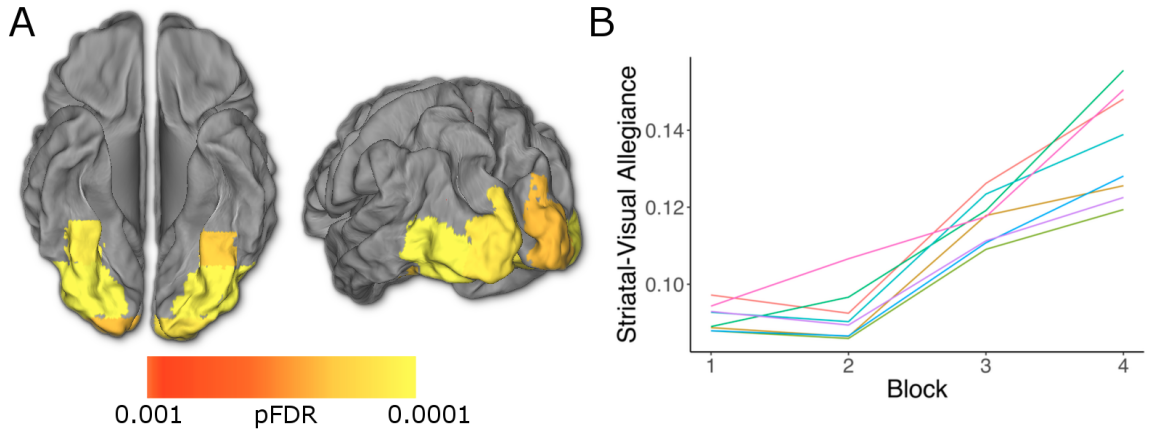

Figure 3.5. Allegiance between the striatum and visual cortex increases over the course of learning. **A**. Module allegiance between the striatum and a number of visual cortex ROIs changes over time (whole-brain corrected, pFDR < 0.05). **B.** Average striatal allegiance increases in each of these visual ROIs (color lines represent the mean for each ROI passing FDR threshold across subjects and striatal regions). Allegiance is averaged across striatal sub-regions in both panels **A.** and **B**. See **Supplementary Figure 3.5** for allegiance averaged over time and **Supplementary Figure 3.6** for results presented by sub-region.

We found that overall, the nucleus accumbens and the caudate showed strongest connectivity with midline prefrontal, temporal, and retrosplenial structures, while the putamen exhibited its highest allegiance with motor cortices and insula (**Supplementary Figure 3.5**). To address the key question of which regions changed coupling with subregions of the striatum during learning, we ran whole-brain searches of separate mixed-effects ANOVAs for each region, predicting striatal allegiance with learning block. This analysis does not assume any shape or direction to these temporal changes. A number of regions of visual cortex showed an increase in striatal allegiance over the course of learning (all FDR $p<0.001$, **Figure 3.5**). Examining sub-regions of the striatum (correcting for multiple comparisons across allegiance of all ROIs with all striatal regions) revealed increases in visual coupling in the nucleus accumbens and putamen, as well as between the putamen and orbitofrontal and ventromedial prefrontal cortices, regions known for their role in value processing (Padoa-Schioppa and Assad, 2006) (**Supplementary Figure 3.6**).

Flexibility relates to learning in a distributed set of brain regions. In a number of reports on dynamic networks, averaged whole brain flexibility has been used as a marker of global processes and associated with cognition (Bassett et al., 2011; Braun et al., 2016). Indeed, we found that whole-brain flexibility related to learning performance within subjects ($\beta = 11.84$,S.E. = 3.91, $p<0.005$) and, to some extent, learning rate across subjects (ρ= -0.26, $p(\rho>0) = 0.07$). We were thus interested in regions outside of the striatum that exhibit dynamic connectivity related to learning.

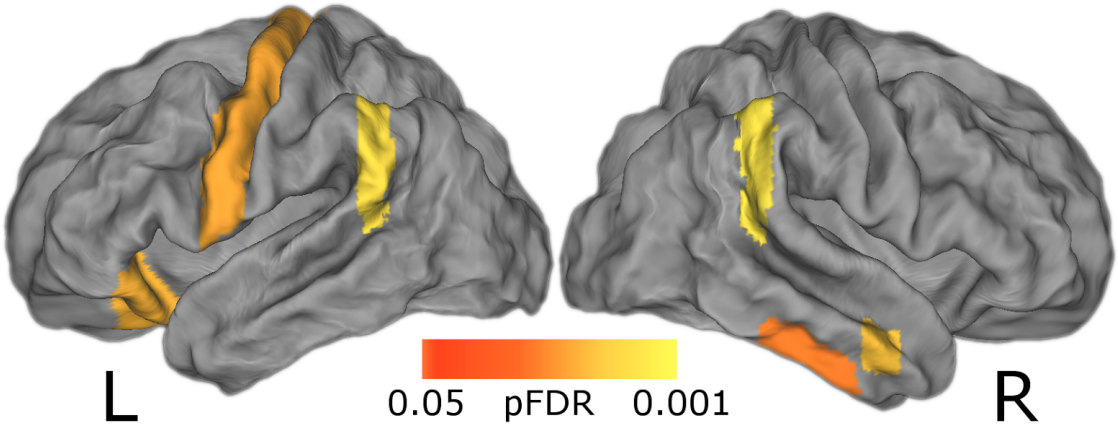

Figure 3.6 Flexibility in cortical regions is related to learning performance. Regions passing FDR correction following a univariate whole-brain analysis using the same mixed-effects model as the *a priori* striatum ROI. Regions passing this threshold include left motor cortex, bilateral parietal cortex, and right orbitofrontal cortex. See **Supplementary Table 3.1 and Supplementary Figure 3.7** for a full list of regions and an exploratory uncorrected map.

We conducted the learning performance analysis for each of the 110 ROIs, in addition the analysis in the striatum reported above (**Figure 3.3**). This analysis again revealed a significant effect of flexibility in striatal sub-regions (the right putamen and left caudate) surviving FDR correction. In addition, the whole-brain corrected results, presented in **Figure 3.6**, indicate that network flexibility in regions of the motor cortex, parietal lobe, and orbital frontal cortex (among others, see **Supplementary Table 3.1 and Supplementary Figure 3.7** for full list and uncorrected map), are associated with reinforcement learning.

Flexibility in medial cortical regions is associated with episodic memory. Finally, our task also included trial-unique objects presented simultaneously with reinforcement, allowing us to measure subjects' episodic memory, a process thought to rely on distinct cognitive and neural mechanisms to feedback-based incremental learning. We tested whether network flexibility was associated with episodic memory for these trial-unique

images, as assessed in a later surprise memory test (**Figure 3.1**). Having a measure of episodic memory for the same trials in the same participants allowed us to determine whether striatal network dynamics are correlated with any form of learning, or whether these two forms of learning, occurring at the same time, are related to distinct network dynamics.

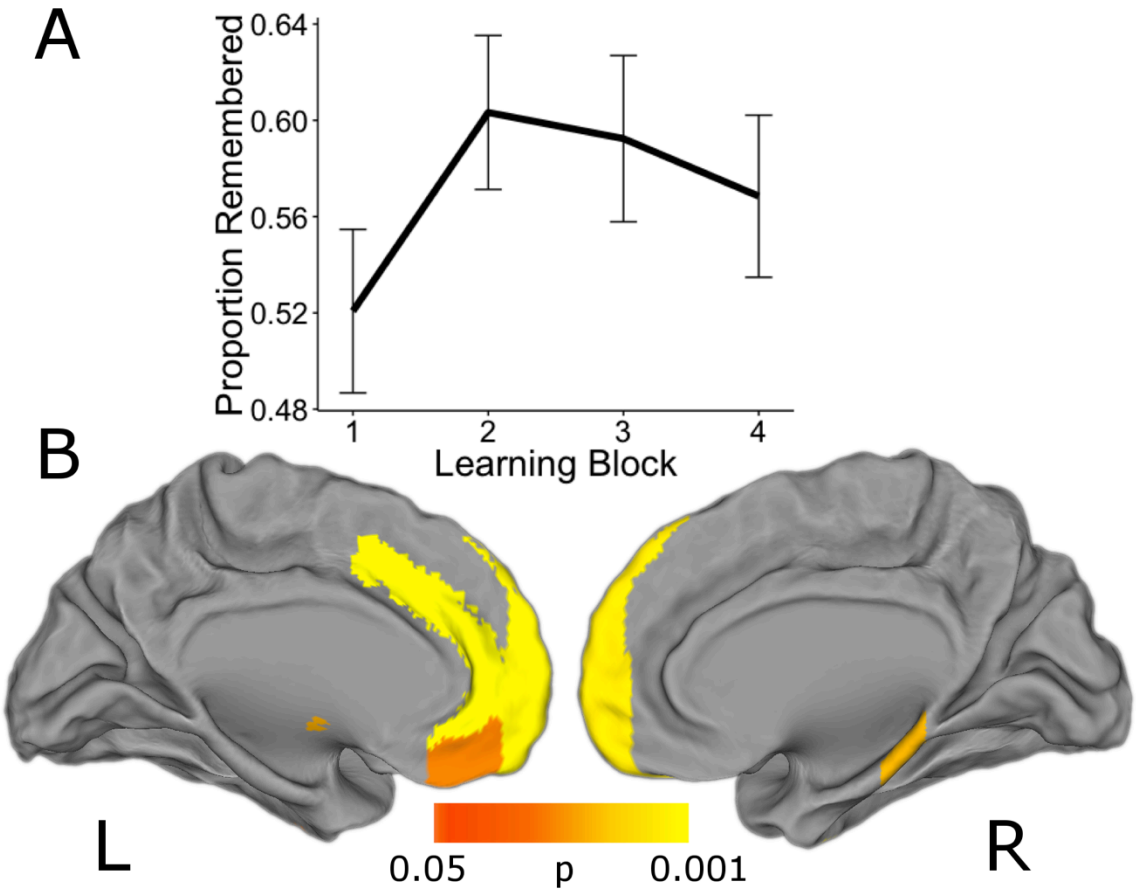

Figure 3.7 Network flexibility in medial prefrontal and parahippocampal cortex relates to episodic memory. An exploratory analysis showed effects of network flexibility on episodic memory performance in medial prefrontal and temporal lobes. **A**. Average memory (proportion remembered) across blocks. Participants' recollection accuracy varied across blocks. Line represents group average and bars represent standard errors. **B**. A number of medial prefrontal regions as well as the right parahippocampal gyrus passed an exploratory uncorrected threshold of $p<0.05$ for the effect of flexibility on subsequent episodic memory. The effect in the left paracingulate gyrus survived FDR correction**.**

Behaviorally, participants' memory was better than chance (d-prime = 0.93, $t_{21} = 7.27$, $p < 0.0001$). Memory performance ("hits") varied across learning blocks, allowing

us to assess within-subject associations between network flexibility and behavior (**Figure 3.7A**). Memory performance was not correlated with incremental learning performance (mixed effects logistic regression,$\beta = 0.41$, Standard Error (S.E.) = 0.51, p = 0.42 (Wald approximation)). We tested the effect of flexibility on memory performance (proportion correct) in each of the 110 ROIs. A whole-brain FDR-corrected analysis revealed one region where flexibility was associated with episodic memory, the left paracingulate gyrus. An exploratory uncorrected analysis revealed regions in the medial prefrontal and medial temporal (parahippocampal) cortices where flexibility was associated with episodic memory ($p<0.05$ uncorrected, **Figure 3.7B**). None of the sub-regions from our *a priori* striatum ROI passed even this low threshold for an effect of flexibility on episodic memory.

Discussion

The current study reveals that reinforcement learning involves dynamic coordination of distributed brain regions, particularly interactions between the striatum and visual and value regions in the cortex. Increased dynamic connectivity between the striatum and large-scale circuits was associated with learning performance as well as with parameters from reinforcement learning models. Together, these findings suggest that network coordination centered on the striatum underlies the brain's ability to learn to associate values with sensory cues.

Our results indicate that during learning the striatum increases the extent to which it couples with diverse brain networks, specifically with regions processing value and relevant sensory information. This may represent the formation of efficient circuits for inte-

grating and routing decision variables. Our reinforcement learning model findings are consistent with this idea. Striatal flexibility is negatively related to learning rate, suggesting that increased dynamic coupling with relevant cortical areas may lead to less trial-level weighting of prediction errors during learning. Flexibility in striatal circuits is also positively related to inverse temperature, indicating that this increased dynamic coupling is associated with stronger reliance on learned value during decision-making. Thus, our findings support a framework wherein network flexibility underlies information integration during learning.

Central to this interpretation is the finding that regions of the striatum increase connectivity with visual and orbitofrontal areas over the course of this task. If the striatum's role is in part to integrate motivational and sensory information in support of decisions, then such an increase is expected given that the relevant information on this task is the value of stimuli differentiated by their visual properties. The fact that bilateral primary auditory cortices are the only regions showing a significant decrease in connectivity with the striatum is also consistent with this explanation. It has been suggested that the striatum may serve to gate information coming in and out of the neocortex (Frank and Badre, 2011), and that this gating function may be related to a broad role for the striatum serving as a hub for controlling decisions (Shadlen and Shohamy, 2016). Reinforcement learning may be characterized in part by the dynamic formation of circuits linking areas processing sensory and value information to the striatum, and by the updating of these circuits via prediction error inputs from the midbrain, to ultimately control actions that reflect a decision. While we focused primarily on the striatum, given its established importance in reinforcement learning, our whole-brain results suggest that the relationship

between dynamic connectivity and learning is not localized to this region: it is plausible that a number of regions integrate information during learning via changes in dynamic connectivity patterns.

This framework offers clear and testable predictions for future studies. It suggests that flexibility will play a larger role in learning the more that learning depends on widespread information integration, and also that this process is specific to regions known to support the particular demands of learning in a given situation. For example, instrumental conditioning involving complex audio-visual stimuli (Kehoe and Gormezano, 1980) would be expected to associate more strongly with striatal flexibility than the task presented here and would be expected to involve increases in striatal interactions with auditory as well as visual cortex. Learning that relies on other forms of integration across time or space (Eichenbaum, 2000; Staresina and Davachi, 2009; Wimmer et al., 2012; Wimmer and Shohamy, 2012), is predicted to be associated with network flexibility in medial temporal and prefrontal regions.

There are a number of limitations to this report. First, given the static feedback probabilities, the relationship between network flexibility and learning performance could be affected by the time on task for each subject. This seems unlikely to fully explain the relationship because flexibility was related to multiple aspects of learning behavior and episodic memory, which was associated with network flexibility in a distinct set of brain regions, did not increase over time. Nonetheless, future studies incorporating reversal periods to dissociate time from performance will be important for addressing this issue. Another limitation is the hard-partitioning approach for network assignments provided by multi-slice community detection, which necessarily underemphasizes uncertainty about

community labels. There have been very recent attempts to formalize probabilistic models of dynamic community structure (Durante et al., 2016; Palla et al., 2016), but most work examining dynamic networks in the brain have used deterministic community assignment (Bassett et al., 2011; Bassett et al., 2015; Shine et al., 2016). More work is needed to develop and validate these probabilistic models and apply them to neuroscience data. Finally, while the spatial resolution of fMRI makes it an appealing method to characterize dynamic networks, studies using modalities with higher temporal resolution such as ECoG (Khambhati et al., 2016) and MEG (Siebenhühner et al., 2013) will be important for providing more fine-grained temporal information.

To summarize, we report a link between reinforcement learning and dynamic changes in networks centered on the striatum. While most descriptions of reinforcement learning have focused on the role of individual regions, recent advances in network theory are beginning to make the role of dynamic communication between individual regions and broader networks in this process a tractable area of research (Bassett et al., 2011; Kopell et al., 2014; Bassett et al., 2015; Braun et al., 2015). Here we show that incremental learning based on reinforcement is associated with dynamic changes in network structure, across time and across individuals. Our results suggest that the striatum's ability to dynamically alter connectivity with sensory and value-processing regions provides a mechanism for information integration during decision-making and that learning may be characterized by the formation of these dynamic circuits.

Supplementary Figures

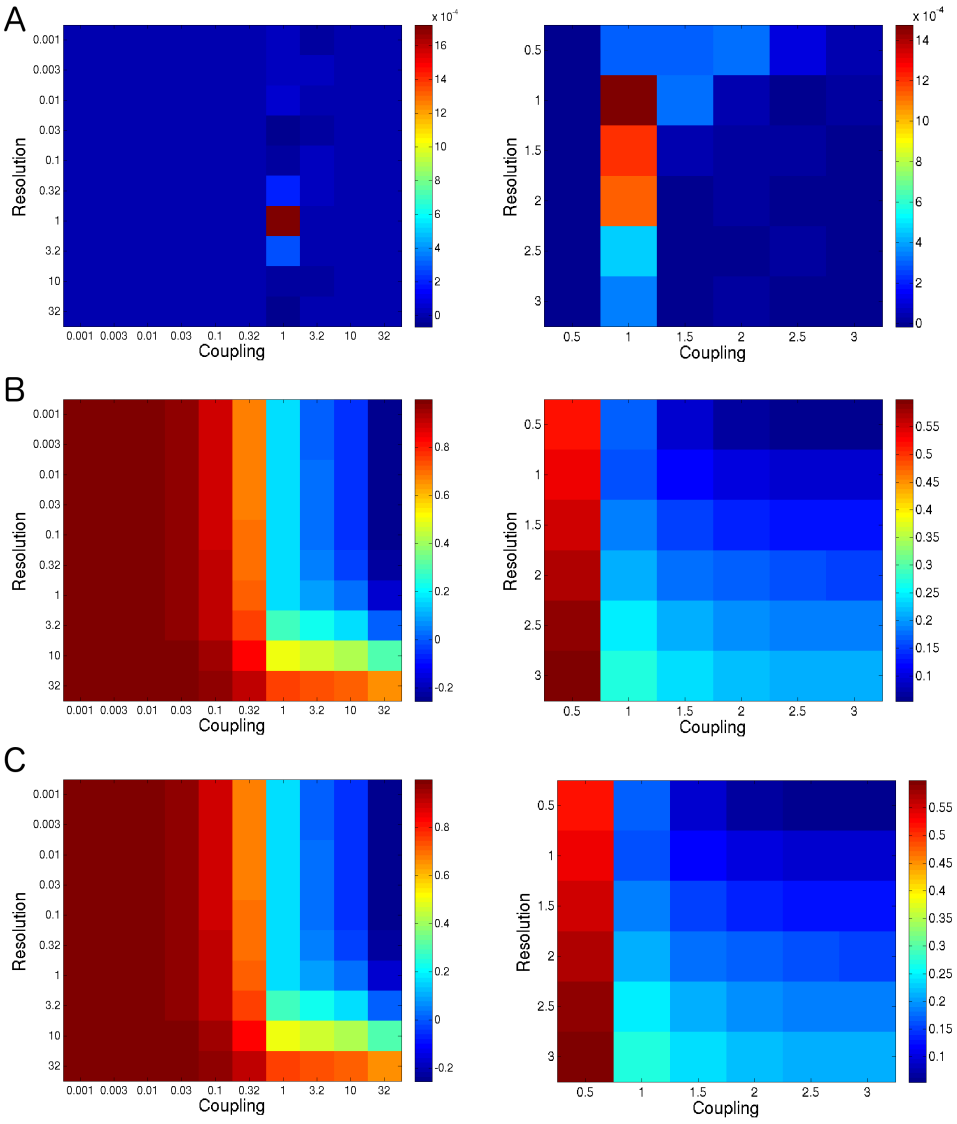

Supplementary Figure 3.1. Two grid searches at distinct scales illustrating that our resolution and coupling parameters (1.18 and 1, respectively) fall at the peak of our optimization function, which is the difference between multi-slice modularity from our data and a null temporal model ($Q_{ml} - Q_{null}$, shown in **A**). We first used a wide grid to cover a larger range of potential parameters (left). The average peak of this search across subjects and iterations was used to select resolution and coupling terms for multi-slice community detection. To ensure that the large grid steps did not affect this selection, we repeated the search on a smaller scale (right). Our parameters are clearly within the optimum range on both grids. Panels **B** and C show the same grids for Q_{ml} (**B**) and Q_{null} (**C**). Qml and Qnull show similar changes with resolution and coupling due to their sensitivity to shared features such as network size and edge strength distribution (Good et al., 2010).

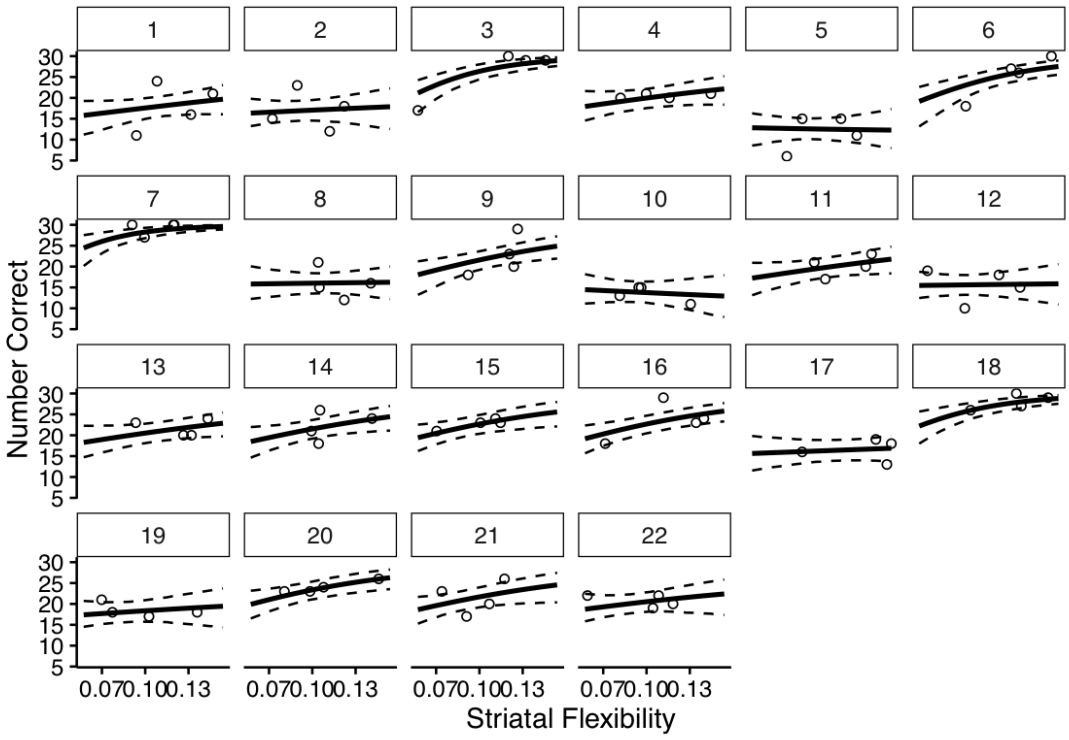

Supplementary Figure 3.2 Subject-level data and fits for Bayesian hierarchical model of the effect of striatal flexibility on learning. For posterior inference on the effect of striatal flexibility on learning performance, we fit a Bayesian hierarchical model. Each subplot displays data (open circles) from a single subject. Solid lines represent model estimates for the effect of flexibility on learning, while dotted lines represent 95% credible intervals.

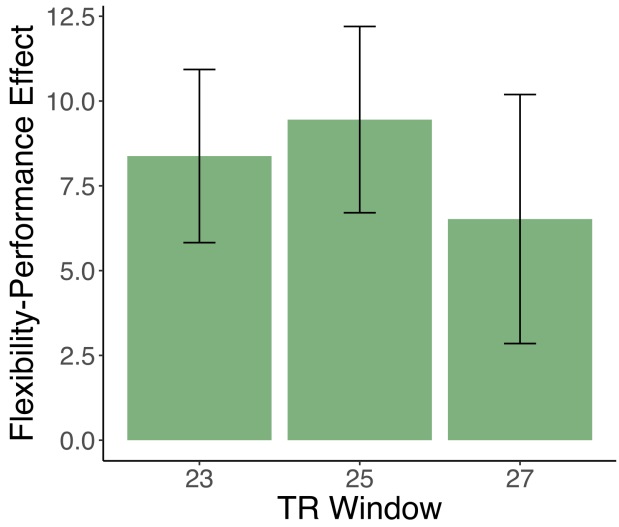

Supplementary Figure 3.3 The relationship between striatal flexibility and learning is robust across different lengths of connectivity windows. We recomputed coherence over 23- and 27-TR windows, and ran the multi-slice community detection algorithms to extract flexibility. The relationship between average flexibility and performance was similar across time windows. Regression coefficients and standard errors (SE) from mixed-effects models are plotted here (23 TR: effect=8.38, SE=2.55; 27 TR: effect=6.52, SE=3.67).

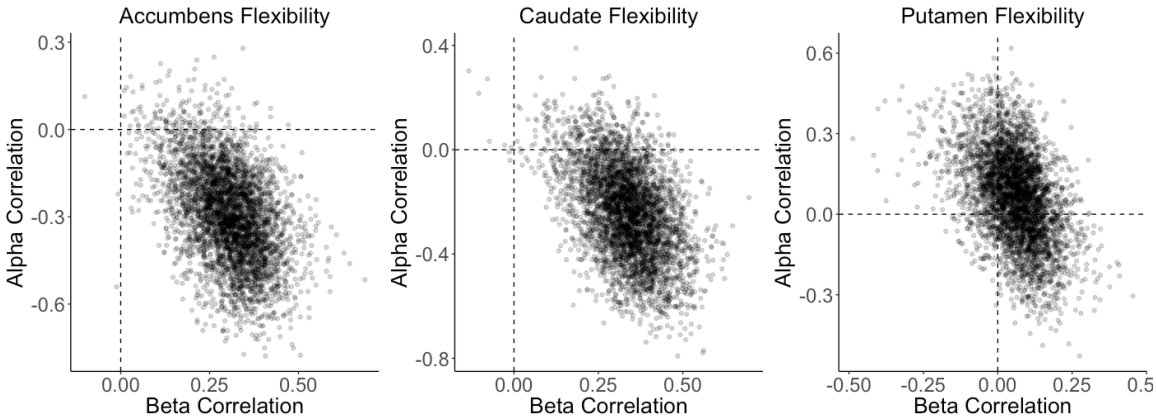

Supplementary Figure 3.4 Joint distributions of correlation between flexibility and reinforcement learning model parameters for striatal regions. Because we computed Spearman correlations over the posterior distributions of learning rate and inverse temperature from hierarchical Bayesian models, our inferences can be most fully expressed with the joint distributions of correlations between flexibility and each parameter. While there is some covariance between these correlations, the two effects are clearly separable. This is further supported by partial correlations, which did not substantially alter inference (accumbens-alpha ρ = -0.23, caudate-alpha ρ = -0.16, accumbens-beta ρ = 0.23, caudate-beta ρ = 0.28).

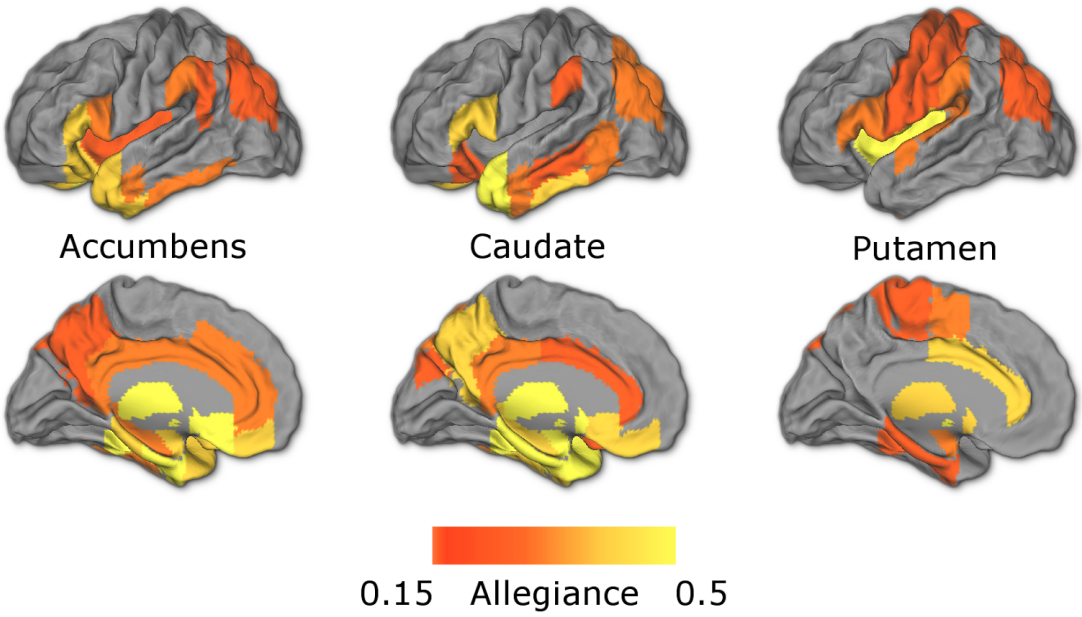

Supplementary Figure 3.5 Module allegiance broken down by striatal region. Maps show regions in the 50th percentile of allegiance for each striatal ROI, averaged over all learning blocks. Consistent with anatomical and functional connectivity, the nucleus accumbens and caudate show stronger allegiance with midline frontal, temporal, and retrosplenial regions, while the putamen shows relatively stronger allegiance with motor regions.

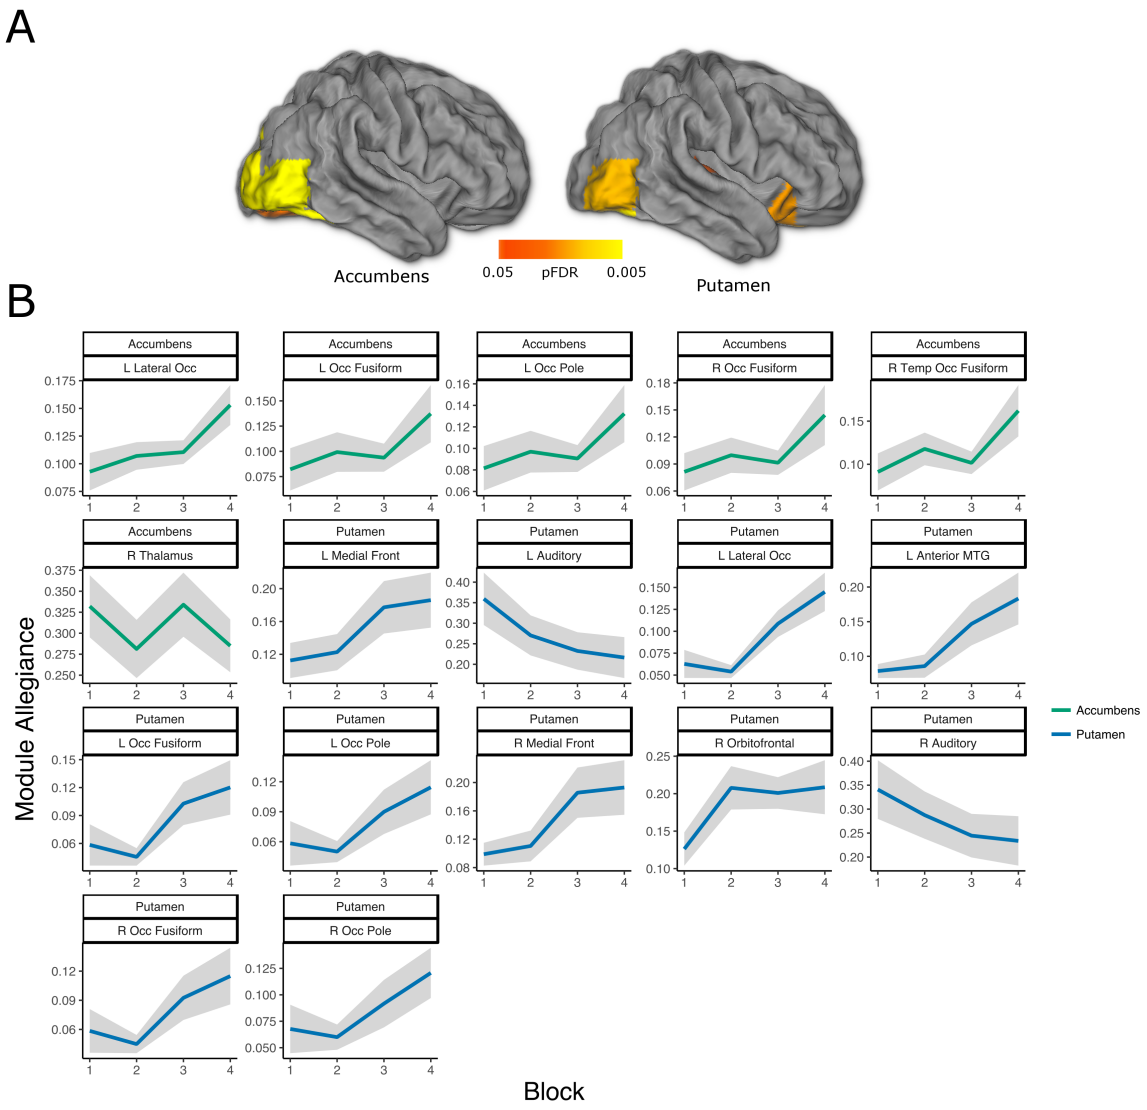

Supplementary Figure 3.6 Visual and value regions change allegiance with the striatum over the course of the task. A. Regions where time-dependent changes to striatal allegiance exceed a threshold of pFDR<0.05, corrected for all ROIs' allegiance with all three striatal regions (109 x 3 comparisons). These results were generated using mixed-effects ANOVAs and contain no assumptions about the shape or direction of changes. **B.** Panel plot showing the change in allegiance over time for every pair passing the above threshold. Lines and bands represent and bands standard errors. As with average striatal allegiance, the nucleus accumbens and putamen increase coupling with visual regions during the task. In addition, the putamen exhibits an increase in coupling with the right orbitofrontal and ventromedial prefrontal cortex and a decrease in coupling with primary auditory cortex. No regions' allegiance with the caudate survived correction for multiple comparisons. Abbreviations: Occ=Occipital, Temp=Temporal, Front=Frontal, and MTG=Middle Temporal Gyrus.

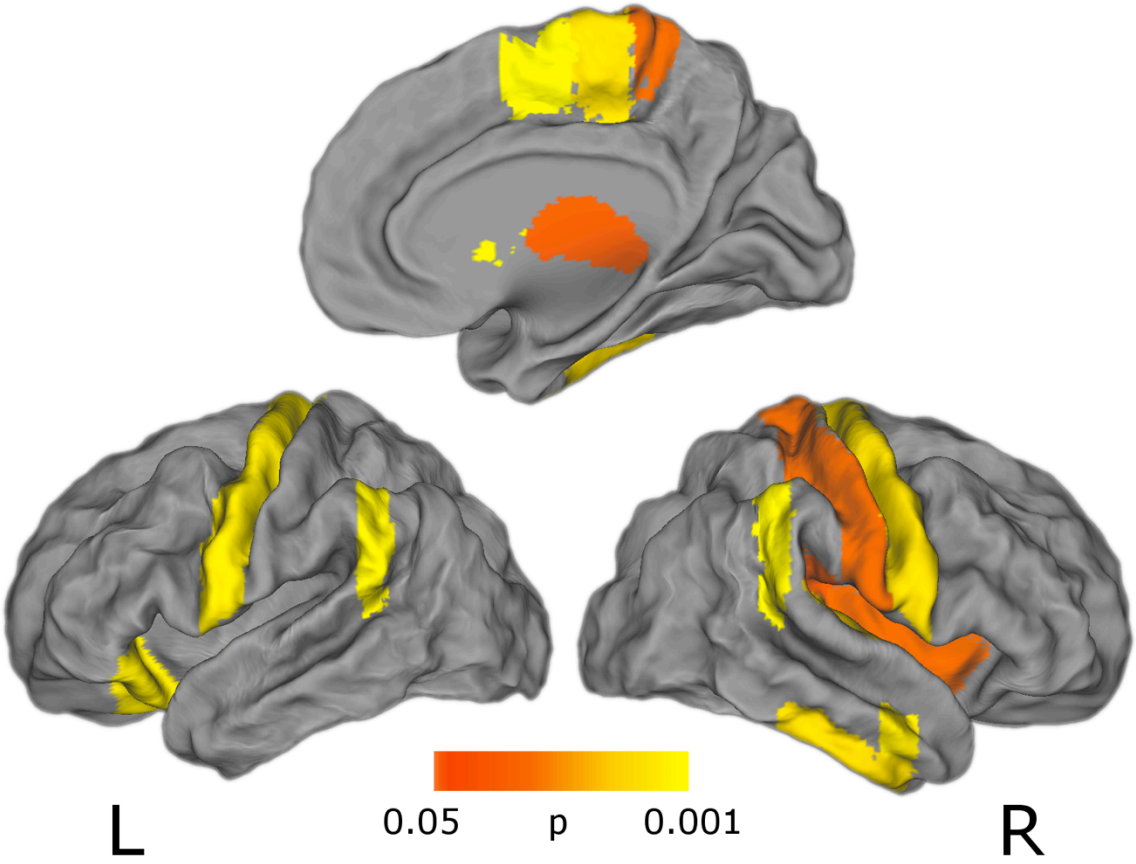

Supplementary Figure 3.7 Exploratory uncorrected (p<0.05) results for whole-brain effect of network flexibility on reinforcement learning performance.

Supplementary Table 3.1

Harvard Oxford Region	Regression Coefficient	p
Left Caudate	5.93	0.003
Left Orbitofrontal Cortex	4.88	0.006
Left Planum Polare	10.97	0.000
Left Precentral Gyrus	4.79	0.006
Left Supramarginal Gyrus	5.38	0.003
Right Inferior Temporal Gyrus, anterior	4.75	0.01
Right Inferior Temporal Gyrus, posterior	4.66	0.01
Right Supplemental Motor Cortex	8.54	0.0009
Right Middle Temporal Gyrus	3.51	0.006
Right Planum Temporal	8.19	0.01
Right Putamen	6.10	0.003
Right Supramarginal Gyrus	6.48	0.003

Chapter 4

Dopamine induces learning-related modulation of dynamic striatal-cortical connectivity: evidence from Parkinson's disease

Chapter 4 is being prepared for submission with Madeleine Sharp, Amanda Buch, Danielle Bassett, Daphna Shohamy

Abstract

Dopaminergic inputs from the midbrain to the striatum signal a reward prediction error that is thought to update the value of actions by modifying connections in widespread cortico-striatal circuits. Decades of research have described the activity of individual regions in reinforcement learning; however, a broader role for dopamine in modulating network-level processes has been difficult to capture empirically. Here, we address this gap by characterizing the effects of dopamine on learning-related dynamic connectivity in patients with Parkinson's disease undergoing fMRI. Patients with Parkinson's disease have severe dopamine depletion in the striatum and are treated with dopamine enhancing drugs, providing an opportunity to compare learning and brain activity when patients are in a low dopamine state (off drugs) vs. a high dopamine state (on drugs). Patients performed a probabilistic reversal learning task with domain-specific visual choice options, in order to measure the relationship between dopamine and dynamic connectivity centered on the striatum. We found that reversal learning altered dynamic network flexibility in the striatum and that this effect was dependent on dopaminergic state. We also found that dopamine modulated changes in connectivity between the striatum and specific task-relevant visual areas of inferior temporal cortex, providing empirical support to theories that value is updated through changes in cortico-striatal circuits. These results suggest that the effect of dopamine on learning-related brain activity is not localized only to trial-specific updating in the ventral striatum, but has effects on widespread circuits in the brain.

Introduction

Updating actions based on feedback is critical for survival in a changing environment. Accordingly, reinforcement has been central to our understanding of learning in psychology and neuroscience (Thorndike, 1898; Skinner, 1948; Schultz et al., 1997). Learning from reinforcement relies on the formation of associations between sensory cues, actions, and the value of outcomes, which must involve coordinated routing and processing of information across widespread brain areas. However, much of the research examining the neural and cognitive mechanisms of reinforcement learning has focused on describing the roles of individual brain regions. And while circuit-level changes play a major role in theories of how value is updated in reinforcement learning (White, 1989a; Glimcher, 2011), such changes have been the focus of surprisingly little empirical work.

In studies of reinforcement learning, particular attention has been paid to the striatum and its dopaminergic inputs from the midbrain. The striatum has long been theorized to play an integrative role in brain function due to its widespread projections to cortical areas and its output to the motor system (Kemp and Powell, 1971; Bogacz and Gurney, 2007; Hikosaka et al., 2014; Ding, 2015). Through a diverse set of inputs and the ability to gate and amplify outputs, the region is thought to be important for behavioral selection (Yin and Knowlton, 2006). Unsurprisingly given these properties, it has been demonstrated repeatedly in human and animal studies that the striatum is necessary for the process of learning from feedback (White, 1989b; Robbins and Brown, 1990; Yin et al., 2006; Vo et al., 2014). Furthermore, striatal function has been shown to depend on dopaminergic inputs from neurons in the ventral tegmental area and substantia nigra, two

nuclei in the midbrain (Kao and Powell, 1986; Westerink and Kwint, 1996; Steinberg et al., 2013).

Indeed, the precise signal being reported by midbrain dopaminergic neurons has been relatively well characterized. The firing of these neurons in response to sensory cues and reward outcomes provides a striking match to a specific reward prediction error computation found in a number of artificial reinforcement learning algorithms (Schultz et al., 1997; Sutton and Barto, 1998; Pagnoni et al., 2002; Bayer and Glimcher, 2005; Daw et al., 2006). Thus dopaminergic projections from the midbrain to the striatum and cortex are thought to update learned associations based on feedback via prediction error signaling. The exact mechanism of this update has not been demonstrated, but it is theorized to take place via modulation of cortico-striatal circuits. Neurons in the striatum and frontal cortex that are active at the time of an outcome as a result of their involvement in the preceding actions are thought to be "transiently bathed" in the dopaminergic prediction error (Glimcher, 2011). Because the actual strengthening or weakening of value-based associations is putatively mediated by changes in striatal-cortical connectivity, a mechanistic description of the precise nature of these circuit-level interactions is essential for an understanding of the neural substrates for learning. Any account of reinforcement learning that does not describe striatal cortical connectivity changes will necessarily be incomplete.

While reinforcement learning theories in neuroscience posit this widespread circuit-level mechanism, until recently there have been few investigations into the time-varying interactions between distributed regions that facilitate reinforced behavior. This dearth of study has been due in part to difficulties in testing the existence and nature of these interregional interactions empirically; while measures of the activation of individual

regions provide a dynamic portrait of brain processes, most measures of interactions between regions have been based on pairwise structure in static correlations of regional activity time series. Recent studies in network neuroscience have begun to address these difficulties using emerging tools from graph theory to characterize the evolution of dynamic connectivity patterns over the same time scales as behavior change (Bassett et al., 2011; Bassett et al., 2015; Shine et al., 2016). With particular relevance to our study, some reports have explicitly linked dynamic and large-scale connectivity centered on the striatum to value- and feedback-based learning (Mattar et al., 2016; Gerraty et al., 2018).

Part of the support for the role of striatal dopamine in learning comes from studies of patients with Parkinson's disease, who experience loss of dopaminergic midbrain neurons, and who exhibit subtle deficits in reinforcement learning in addition to their more pronounced motor disturbances (Knowlton et al., 1996; Frank et al., 2004; Foerde et al., 2012). The main pharmaceutical treatment for Parkinson's disease is levodopa, a dopamine precursor that leads to increased levels of the neurotransmitter in the striatum. Thus, medication manipulations in patients with the disorder have been used to model dopamine's effects on learning, by comparing behavioral performance and brain activity on versus off medication. In this study, we used this approach to test dopamine's effect on dynamic network connectivity in the striatum during reinforcement learning in patients with Parkinson's disease.

We have recently suggested that changes in connectivity between the striatum and areas of cortex that process sensory and value information provide a mechanism for dynamic information integration during reinforcement learning, allowing for the routing of evidence that is relevant for a given choice (Gerraty et al., 2018). In this framework, as

the optimal choice is learned, the striatum increases its coupling with sensory areas conveying decision-relevant information, and with regions processing value or motivational salience. This idea is consistent with reinforcement learning theories which hold that learned associations are updated via dopamine-induced changes in striatal-cortical circuitry. However, the relationship between dynamic connectivity and dopamine has not been characterized, diminishing the ability to link these recent findings to prominent theories of value updating. Moreover, research on dynamic connectivity during learning has previously relied on correlational data, and it is therefore not yet known whether experimental manipulations of learning can modulate dynamic connectivity.

Given prior findings demonstrating a relationship between dynamic cortico-striatal connectivity and value learning, and the success of reinforcement learning theory in describing dopaminergic inputs to the striatum, we hypothesized that dopamine would modulate learning-related changes in dynamic connectivity between the striatum and task-relevant cortical areas. Specifically, we predicted that (i) reversal learning would modulate flexible network connectivity in the striatum, (ii) this modulation would depend on levels of the neurotransmitter dopamine, and (iii) dopamine would specifically alter dynamic connectivity between the striatum and cortical regions involved in processing task-relevant visual category information.

We assessed the validity of these predictions in patients with Parkinson's disease undergoing fMRI while engaged in a reversal learning task with visual-category-specific choice options. We compared behavior and dynamic connectivity metrics across sessions in which patients were tested on and off of dopaminergic medication using a within-subject design. Our findings indicate that learning-induced changes in dynamic striatal-

cortical connectivity are modulated by dopamine, and that these changes are most pronounced in connections with task-specific sensory regions, providing a link between reinforcement learning and network mechanisms underlying the acquisition of learned behavior.

Methods

Patients

Participants were thirty patients with idiopathic Parkinson's disease (7 females, mean ± SD age: 61.86 ± 6.25 years). Patients were recruited from either the Center for Parkinson's Disease and other Movement Disorders at Columbia University Medical Center or from the Michael J Fox Foundation Trial Finder website. All patients provided informed consent and were compensated $100 per day for taking part in the study. All aspects of the study were approved by Columbia University's Institutional Review Board.

Patients were in the mild-to-moderate stage of disease, as rated on the Unified Parkinson's Disease Rating Scale (UPDRS) in the OFF medication phase by a neurologist specializing in movement disorders (mean ± SD: 20.35 ± 9.22). Two patients were not rated and 5 other patients were missing ratings for one session, due to the lack of an available neurologist. The range of disease duration was 1-17 years. All patients had been receiving levodopa treatment for at least 6 months; the mean total daily levodopa or equivalent dose was 796.76 ± 404.91 mg. Fifteen patients were also taking dopamine agonists. In addition to the learning task described below, patients also completed the Montreal Cognitive Assessment (MoCA) neuropsychological battery, as well as an extended Digit Span, Starkstein Apathy Scale, and a Beck Depression Inventory. Patients did not

exhibit dementia (indicated by MoCA < 26) and had no other history of major neurological or psychiatric illness except for Parkinson's disease.

Medication state

Each participant was tested in two sessions 24 hours apart and counterbalanced for order of medication phase. For the OFF phase, patients were asked to undergo an overnight withdrawal from all medication taken for Parkinson's, lasting at least 16 hours in duration which is at least 10 half-lives for levodopa and 2 half-lives for dopamine agonists. In the ON phase, the same patients were tested 1-1.5 hours after taking their normal dose of levodopa. To isolate the effect of levodopa, patients who were additionally taking dopamine agonists were asked to take only levodopa for the ON testing day. UPDRS scores decreased between OFF and ON sessions (mean $\pm$ standard error ON versus OFF difference: -10.48 ± 0.96, $t(22) = -10.87$, $p < 0.0001$).

Task

We designed a reinforcement learning task with two broad goals (**Figure 4.1**). First, we included a reversal of the optimal choice during each session, in order to manipulate learned choice contingencies and to characterize the effect of this manipulation on dynamic connectivity. Second, choice options were images of objects and scenes, categories known to evoke specific patterns of activation in the ventral visual stream, in order to characterize the dynamics of specific striatal-cortical circuits during learning. We predicted that the use of these categories during learning would lead to changes in striatal connectivity with domain-specific sensory areas. Each patient was tested in ON and OFF

dopaminergic states to characterize the effect of dopamine on both of these aspects of the task.

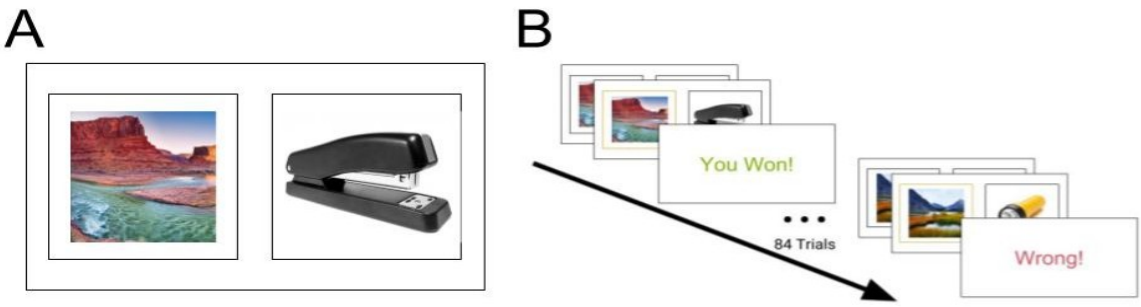

Figure 4.1. Reversal learning task. **A.** On each trial, participants were asked to choose between an image of a scene and an image of an object. Each trial contained unique images that were randomly placed on the left or right side of the screen, and they indicated their decision with a button press. **B**. One category led to positive feedback with a probability of 0.7, while the other category led to positive feedback with a probability of 0.3. The category-outcome contingencies reversed after 85 trials.

On each trial, participants were asked to choose between images of a scene and of an object, which changed on every trial. Images were randomly presented on the right and left sides of the screen and participants responded using index and middle finger responses, respectively. There were 150 trials on each day, broken up into 5 scan runs to allow short breaks for patients. The optimal category led to positive feedback ('You Win!') with a probability of 0.7, while the non-optimal category led to positive feedback with a probability of 0.3. Negative feedback ('Wrong!') was shown with a probability of 0.3 and 0.7 for 'correct' and 'incorrect' choices, respectively. Participants had 2.5 seconds to respond, followed by a 1.5 second crosshair and feedback shown for 1 second. Reversals took place on the 85th trial of each session. Participants were instructed that the correct option could change at any point during the task, and they were not told how many times a change could take place. Optimal categories at the start of the session were

counterbalanced across subject-session pairs, such that half of the subjects experienced the same optimal category at the beginning of both sessions. Trials were separated by an inter-trial interval (ITI) drawn from an exponential distribution with a mean of 3 seconds, and a minimum of 1.5 seconds. Twenty-five participants also completed a surprise subsequent memory test for chosen objects 24 hours following their second session; we note that these data were not analyzed for this report.

Behavioral analysis

On each trial, we recorded whether participants chose the option with the higher probability of correct feedback, as well as reaction time and medication state. For behavioral analysis, we divided trials into 10 learning blocks (2 blocks per scan run). We estimated a mixed effects logistic regression using *lme4* (Bates et al., 2015) with average performance and medication effect varying randomly by subject and by learning block. In *glmer* syntax for clarity:

$$Opt \sim Med + Day + (Med \mid Subject) + (Med + Day \mid Block),$$

Where *Med* indicates ON or OFF medication session (coded as 0.5 or -0.5) and *Day* indexes the first or second session (also coded as 0.5 or -0.5). For inference, we performed a parametric bootstrap using the varying effects estimated for learning block. To characterize task- and dopamine-related changes in reaction time (RT), we also fit a linear mixed effects model of *log(RT)* with identical predictors and varying effects.

Imaging Acquisition

Images were acquired on a 3T General Electric Signa MRI scanner using a 32-channel head coil. Functional images were acquired using a multiband pulse sequence

with the following parameters: TR=850 ms, TE=25 ms, flip angle = 60°, field of view (FOV) = 192 mm. A high-resolution (1 mm isotropic) T1-weighted image was also acquired using the BRAVO pulse sequence for image co-registration. Functional runs for 5 sessions (no more than 1 run per session) were lost during acquisition due to errors in multiband image reconstruction.

Preprocessing

Functional images were preprocessed using FSL FMRI Expert Analysis Toolbox (FEAT; (Smith et al., 2004). Images were corrected for baseline magnetic field inhomogeneity using FUGUE. The first five images from each block were removed to account for saturation effects. The first of these five saturated images from each scan run were averaged together across blocks to form a functional template for registration. Images were high-pass filtered at $f < 0.008$ Hz, spatially smoothed with a 5 mm Gaussian kernel, grand-mean scaled, and motion corrected to the averaged template image using an affine transformation with trilinear interpolation. Due to well-characterized motion artifacts in measures of functional connectivity (Power et al. 2012; Satterthwaite et al. 2012), we used a previously validated nuisance regression strategy (Satterthwaite et al. 2013) to further preprocess functional images. Predictors in this nuisance regression included the 6 translation and rotation parameters from the motion correction registration as well as CSF, white matter, and whole brain average time course, in addition to the square, derivative, and squared derivative of each confound. This method has been shown to outperform a number of other strategies for motion correction—including PCA- and ICA-based decomposition, global signal regression, and motion regression techniques with fewer

parameters—on measures of connectivity-motion and modularity-motion correlations and as well as on measures of network identifiability (Ciric et al. 2018).

Preprocessed residual images were registered to individuals' anatomical images using linear boundary-based registration (BBR) (Greve and Fischl, 2009), and were then registered to a standard space MNI template using a nonlinear registration implemented in FNIRT. To characterize the effect of learning and dopamine on dynamic changes in functional connectivity, we used the Harvard-Oxford atlas of 110 cortical and subcortical regions (including 6 striatal regions: bilateral caudate, putamen, and nucleus accumbens) in keeping with our previous report on dynamic flexibility during reinforcement learning (Gerraty et al. 2018). After registration to 2mm MNI space, average time courses were extracted for each region in the atlas.

Dynamic connectivity estimation

Time series were concatenated across blocks, and we used the Multiplication of Temporal Derivatives (Shine et al., 2015) to characterize dynamic connectivity between regions. This method has been shown to be capable of detecting changes in community structure with more sensitivity than standard sliding window techniques. In this technique, dynamic coupling between each pair of regions *i* and *j* at each point in time *t* is calculated as the ratio of the product of their temporal derivatives and the product of their standard deviations. This estimate is then smoothed using a running average:

$$A_{ijt} = \frac{1}{2w+1} \sum_{t-w}^{t+w} \frac{dt_{it} dt_{jt}}{\sigma_i \sigma_j} \quad .$$

A_{ijt} is the dynamic functional connectivity between regions *i* and *j* at time point *t*, dt_{it} is the temporal derivative and σ_i the standard deviation of the average timeseries for region *i,* and *w* is the number of time points on each side of time *t* used for smoothing. We used

a window length of 13 TRs for temporal smoothing, corresponding to w=6 and averaging across 5.1 seconds on each side of the connectivity estimate for each time point. Using this metric, we computed a smoothed measure of network coupling for each pair of regions, leading to an N x N x T (e.g., 110 x 110 x 1477) connectivity matrix for each session. We use this matrix to represent a temporal network in which network nodes represent brain areas in the whole-brain parcellation, and network edges represent the magnitude of functional connectivity estimated using the method of Multiplication of Temporal Derivatives.

Temporal community detection

One of the most useful tools for characterizing network structure in the brain has been a set of techniques for community (or "module") detection (Bassett 2013). These techniques allow for the partitioning of the brain into internally dense, externally sparse groups of nodes based on connectivity strength. Such communities can be extracted at rest or during task performance, and provide a striking match to known cognitive systems identified and characterized through other methods. Communities in brain networks have been shown to undergo non-trivial rearrangement during reinforcement learning (Gerraty et al., 2018), as well as other cognitive and motor tasks (Bassett et al., 2011; Braun et al., 2015). To characterize the evolution of network structure in this experiment, we used a recently developed multilayer community detection algorithm (Mucha et al., 2010), which uses identity links to connect networks in neighboring time windows, in order to solve the community-matching problem and provide time-dependent labels for community assignment (Bassett 2013).

Each connectivity matrix was treated as a unthresholded graph and because the graph extends in time, it represents a *temporal network*, or an ensemble of graphs ordered in time (Holme and Saramäki, 2011). Because each layer has the same number of slices, it can be represented as a *multilayer network* (Kivelä et al., 2014). Here, we partitioned regions in the multilayer network into temporal communities using a Louvain-like locally greedy algorithm for multilayer modularity maximization (Mucha et al., 2010; Jutla et al., 2011; Bassett et al., 2013a). The quality function maximized in this algorithm is:

$$Q = \frac{1}{2\mu}\sum_{ijlr}(A_{ijl} - \gamma_l \frac{k_{il}k_{jl}}{2m_l})\delta_{lr} + \delta_{ij}C_{jlr}\delta(g_{il}, g_{jr}) \ ,$$

where the adjacency matrix for each layer *l* consists of pairwise connectivity components A_{ijl}. In this case, *l* is equivalent to a time point *t* in the multiplication of temporal derivatives coupling measure. The variable γ_l represents the resolution parameter for layer *l*, and C_{jlr} indexes the coupling strength between node *j* at layer *l* and node *j* at layer *r* (due to the large number of layers being decomposed into temporal communities in this study, we set both the resolution and coupling to 1 rather than optimizing the pair of hyperparameters); k_{il} is the coupling strength (the sum of edge weights over all connections) of node *i* in layer *l*; and m_l is the average coupling strength (the sum of all edge weights over all connections for nodes *i* and *j* divided by 2). Finally, the variables g_{il} and g_{jr} correspond to the community labels for node *i* at layer *l* and node *j* at layer *r*, respectively, and $\delta(g_{il}, g_{jr})$ is the Kronecker delta function, which equals 1 if $g_{il} = g_{jr}$, and 0 otherwise.

Network Diagnostics

We utilized two network diagnostics to characterize dynamic connectivity changes during learning and their relationship to dopaminergic state. First, we computed the *flexibility* of each brain region, which measures the proportion of time points during which the community assignment for a node changes. This metric has been shown to relate to reinforcement learning in a previous report (Gerraty et al. 2018), as well as to motor sequence learning (Bassett et al. 2013) and executive function (Braun et al. 2015). While there are potential ambiguities in the measure related to uncertainty in node assignment, we take it to index roughly the extent to which a region is coupling with multiple networks during any given time period. We divided the community assignments into 10 blocks (2 for each task period in the scanner), and computed the flexibility for each region in each learning block.

To characterize the network connectivity of the striatum in more detail— specifically, to characterize *which* regions the striatum couples with and whether these connections are altered by dopaminergic state—we estimated the community *allegiance* for each pair of regions in each time learning block. Allegiance measures the proportion of time points in a given window in which each pair of regions is assigned to the same community. The measure has been linked to motor learning (Bassett et al., 2015), as well as reinforcement learning (Gerraty et al., 2018). To disentangle domain-specific visual areas of the inferior temporal lobe, we used the Brainnetome atlas, a finer parcellation of 246 regions (Fan et al., 2016), to characterize which areas change their allegiance with the striatum during reversal learning. This atlas includes the following 12 striatal sub-regions: bilateral ventromedial and dorsolateral putamen, ventral and dorsal caudate, globus palli-

dus, and nucleus accumbens. For analyses of striatal allegiance, we used sub-regions overlapping with the Harvard-Oxford atlas, thus excluding the globus pallidus. Time series from each of these 246 regions underwent identical connectivity and community detection analyses to those described above.

Flexibility Analysis

To characterize the time course of flexible striatal connectivity during learning, and the effect of choice reversal and dopamine on these dynamics, we estimated a linear mixed effects model analogous to that used to model behavior. This model predicted flexibility with estimates of average flexibility (intercept) and medication effect varying by subject, learning block, and striatal sub-region. In *lmer* syntax:

$$Flexibility \sim Med + (Med \mid Subject) + (Med \mid Block) + (Med \mid ROI) .$$

To characterize uncertainty about the time course of flexibility ON and OFF of medication, we performed a parametric bootstrap using the varying intercept and medication effects estimated by learning block. To uncover regions outside of the striatum showing medication effects on network flexibility, we extended this analysis by including all Harvard-Oxford ROIs, with medication estimates varying by region. We bootstrapped these estimates in the same fashion, but using varying effects estimated by region rather than by learning block, thus regularizing our estimates of the medication effect for each ROI using the distribution of effects across regions. For both striatal and whole-brain analyses we subtracted each region's flexibility in the first learning block and fit the above models to blocks 2-10 in order to estimate changes in flexibility from baseline during learning.

Allegiance Analysis

To analyze changes in striatal connectivity with specific target regions, we extracted the allegiance between each striatal sub-region and every other region, using the higher resolution Brainnetome atlas (Fan et al., 2016). To characterize temporal changes in striatal-cortical connectivity and the effect of dopaminergic state on these changes, we estimated the following linear mixed effects model, presented in *lmer* syntax:

$$log(allegiance) \sim Med + Block + Med{:}Block + (\ Med + Block + Med{:}Block \mid\ Sub) + (\ Med + Block + Med{:}Block \mid\ ROI)\ .$$

Allegiance is the community allegiance computed between each pair of striatal sub-regions and non-striatal target regions, *Med* is the participants' medication state for each session, and *Block* is a factor indexing learning blocks. In this model we divided the task into 5 blocks rather than 10, corresponding in this case to scan run, due to the large number of parameters in the medication x block interaction. Effects varied randomly by participant (*Sub*) and striatal region (*ROI*). Due to the large number of regions in this more fine-grained atlas (multiplied by the number of striatal regions), we estimated this model separately for each target region and approximated a *p*-value for each term using a Wald Chi-square test, rather than allowing effects to vary randomly by target region and bootstrapping.

To evaluate the functional specificity of any time- or medication- related changes in connectivity with the striatum, we constructed a metric for correspondence to known scene or object processing areas for each region in the atlas, using reverse inference maps from Neurosynth (Yarkoni et al., 2011). We thresholded reverse inference maps for the terms 'object' and 'place' at $Z>3.1$ and binarized them to create masks. The object-scene selective metric was simply the number of voxels in a given region overlapping with the

object mask minus the number of voxels overlapping with the place mask, divided by the total number of voxels in the region. This approach gave each region a weight for correspondence to object- or scene- related areas, and the weight was positive for objects and negative for scenes.

To test whether connectivity between the striatum and any domain-specific visual processing regions showed time or medication effects, we first fit the above model to regions in the top 5% of the absolute value of object-scene processing weights and corrected for multiple comparisons for each parameter using False Discovery Rate (FDR) with $q < 0.05$ across these selective regions (Benjamini and Hochberg 1995). To test for the specificity of this effect to relevant scene or object processing regions, we used a more liberal threshold of $p < 0.01$ and calculated the average object-scene processing weight across regions passing the threshold. We then compared this average to means bootstrapped from samples of weights from the same number of regions, taking the absolute value to capture both object and scene area correspondence. We performed both of these procedures for learning block, medication, and block x medication interaction parameters in the above model.

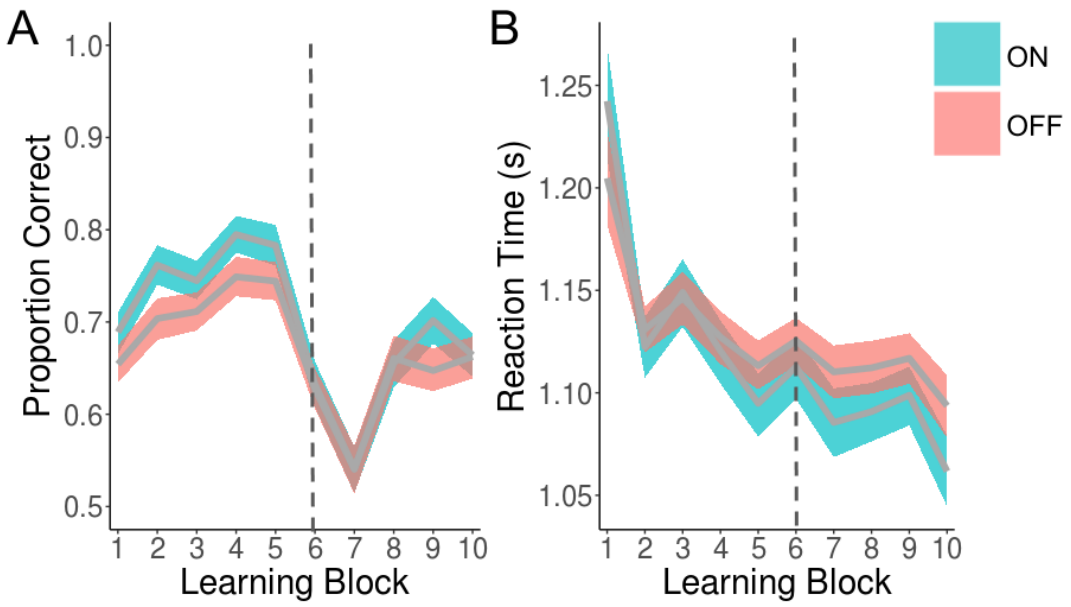

Figure 4.2 Reversal learning in Parkinson's patients ON and OFF of dopaminergic medication. A. Behavioral performance, as measured by the proportion of optimal responses, is plotted against trials binned into 10 learning blocks. The grey lines indicate the average proportion and the ribbons indicate standard errors, estimated by bootstrapping a generalized logistic mixed effects model of optimal choice. **B.** Reaction time (RT, seconds) is plotted against trials binned into the same 10 learning blocks. Grey lines indicate the geometric mean RT and the ribbons indicate standard errors, estimated by bootstrapping a linear mixed effects model of log RT. Colors show medication state (ON levodopa = blue, OFF levodopa = red). The dotted lines in both panels illustrate the reversal of outcome contingencies in the 6th leanring block.

Results

Behavioral Results

Overall, participants learned to track the correct choice over the course of the experiment. We fit a mixed effects logistic regression to characterize performance and the effects of medication. Somewhat surprisingly, we did not observe a robust overall difference in performance between patients ON and OFF of medication (parameter estimate β = 0.08, 95% confidence interval (CI) = [-0.03, 0.17], increase in proportion correct CI = [- 0.006, 0.04]; **Figure 4.2A**). There was some evidence for a small effect of dopamine in individual learning blocks, particularly before reversal (p(ON<OFF) < 0.05 in block 2, and p(ON<OFF) < 0.15 in blocks 1-5 and block 9). While Parkinson's patients have been

shown to exhibit deficits in multiple forms of feedback-based learning, and levodopa has been shown to alleviate such deficits, these effects are not always present, and some studies have found the reverse effect (Shohamy et al., 2004a; Cools et al., 2007; Schonberg et al., 2010; Grogan et al., 2017; Timmer et al., 2017). We observed no effect of medication on reaction times, which decreased over the course of the task, particularly after the first block (**Figure 4.2B**).

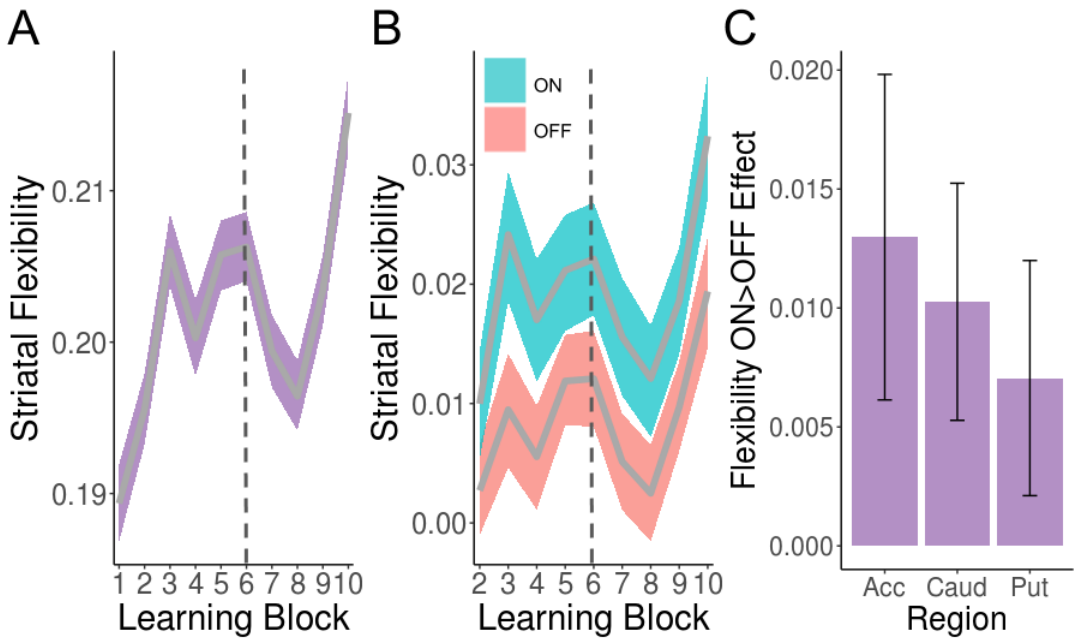

Figure 4.3 Learning-related changes in dynamic striatal connectivity are modulated by dopamine. A. Reversal learning changes network flexibility in the striatum. Striatal flexibility as measured by the proportion of network changes in a time window, is plotted against learning block. The reversal of outcome contingencies provided an experimental manipulation of flexibility, which decreased following reversal before recovering, mirroring behavioral performance (**Figure 4.2**). **B.** Dopaminergic state modulates changes in striatal flexibility. Plot shows baseline-subtracted network flexibility in the striatum over learning blocks. Grey lines indicate the mean flexibility and the ribbons indicate standard errors, estimated by bootstrapping a linear mixed effects model of flexibility varying by subject ROI, learning block, and striatal sub-region. Colors show medication state (ON levodopa = blue, OFF levodopa = red, average across states = purple). The dotted lines in panels A and B illustrate the reversal of outcome contingencies in the 6th learning block. **C.** The effect of medication was similar across sub-regions of the striatum. Bars show parameter estimates for separate models fit to each ROI; error bars indicate standard errors. Acc = Nucleus Accumbens, Caud = Caudate, and Put = Putamen.

Learning and dopamine both modulate dynamic flexibility in the striatum

Previous studies of flexibility have been limited in interpretation by the fact that they have relied on correlations with behavioral measures. We designed this task with a

reversal in order to provide an experimental manipulation of this network metric. We extracted time series from regions distributed across the brain, and computed the dynamic connectivity between all pairs of regions at each time point during the task. We summarized dynamic connectivity in the form of temporal multilayer networks and extracted the reconfiguration of functional modules using a multilayer community detection technique. After splitting the task into 10 learning blocks per session (2 per scanner run), we calculated the *network flexibility* in each block to capture the dynamic connectivity of regions during different periods of the task. We were particularly interested in flexibility in striatal connectivity, as we have previously shown this network metric to relate to reinforcement learning (Gerraty et al. 2018).

To test the effects of dopamine on dynamic striatal flexibility, we fit a linear mixed effects model with average flexibility and medication state varying by subject and by striatal sub-region. In both the ON and OFF medication conditions, network flexibility in the striatum decreases as a result of reversal before recovering, mirroring changes in learning performance (**Figure 4.3A**). Overall flexibility was higher in the OFF medication condition than in the ON medication condition (regression β = -0.004, CI = [-0.006, -0.002], p(OFF>ON) < 0.01), which could potentially be due to medication-related differences in motion or in overall differences in dynamic brain connectivity. Because we were most interested in *changes* in network flexibility during learning, we subtracted the flexibility measured in the first learning bin from the flexibility measured in all other bins. As can be seen in **Figure 4.3B**, this baseline-subtracted flexibility was higher in patients ON dopamine medication than OFF dopamine medication (regression β = 0.01, CI = [0.009, 0.02], p(ON<OFF) < 0.01), indicating greater modulation of flexible striatal-network

coupling in the presence of dopamine. This effect was similar in all sub-regions of the striatum (**Figure 4.3C**).

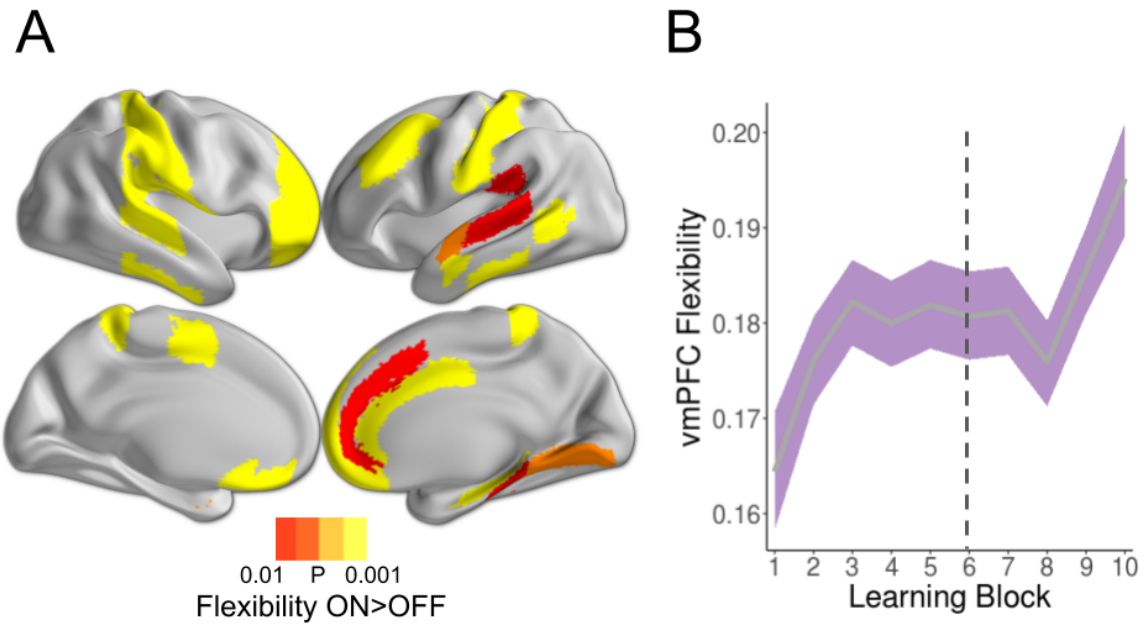

Figure 4.4 Dopamine increases dynamic connectivity across widespread brain areas.
A. Results from a linear mixed effects model of baseline-subtracted flexibility, with medication effect varying by subject, learning block, and ROI. Medication effects were parametrically bootstrapped from the varying ROI estimates, and thresholded at $p(\text{OFF}>\text{ON}) < 0.01$. Lighter color indicates lower $p(\text{OFF}>\text{ON})$. The hippocampus, vmPFC, and regions across limbic, parietal, and prefrontal cortex exhibit greater learning-related changes in dynamic connectivity ON dopamine medication relative to OFF dopamine medication. **B.** The ventromedial prefrontal cortex is also modulated by the reversal task. The grey line indicates the mean flexibility and the ribbons indicate standard errors, estimated by bootstrapping a linear mixed effects model of vmPFC flexibility.

Dopamine has widespread effects on network flexibility

Dopamine has widespread cortical and subcortical targets, and we have previously shown that network flexibility in regions of parietal and prefrontal cortex also relates to reinforcement learning. Thus, we were interested in whether other regions exhibited greater flexibility in network coupling ON versus OFF dopamine medication. To determine whether such differences existed, we estimated a linear mixed effects model with medication effects in network flexibility varying randomly by subject, learning block, and brain region. Region-level differences surpassing $p<0.01$ included the nucleus ac-

cumbens, hippocampus, ventromedial and dorsolateral prefrontal cortex, as well as supplementary motor and parietal cortex (**Figure 4.4**). All regions passing this threshold exhibited greater flexibility ON dopamine medication relative to OFF dopamine medication. Given its anatomical position as a major target of midbrain dopamine, and its established role in value-based learning, we also extracted the time course of flexibility in the ventromedial prefrontal cortex (vmPFC) region showing a medication effect. As seen in **Figure 4.4B**, the flexibility in this region was also affected by reversal, further supporting a link between dopamine, learning, and widespread changes in network connectivity.

Dopamine modulates dynamic coupling between the striatum and task-specific visual areas

Having shown that dynamic connectivity in the striatum changes over the course of learning and is modulated by dopamine, we were interested in *which* regions vary in their striatal coupling. The task was designed to capture the effect of dopamine on dynamic coupling between the striatum and areas of visual cortex. In particular, we hypothesized that the striatum would exhibit greater changes in connectivity with domain-specific areas that contribute to the processing of scene or object stimuli, and that these changes would be affected by dopaminergic state.

To capture more distinct areas of cortex, we used a more fine-grained atlas of 246 regions (Fan et al., 2016). Using Neurosynth (Yarkoni et al., 2011), we constructed a measure of correspondence between each region in this atlas and areas contributing to the processing of scene or object information (see **Methods**). We calculated the *community allegiance* between the striatum and each of these regions. To test whether category-specific visual regions exhibited dynamic changes in connectivity with the striatum dur-

ing the task, and whether any such changes were affected by dopamine, we estimated a linear mixed-effects model of changes in striatal allegiance with regions showing the strongest (top 5%) correspondence to scene- or object-specific areas. In this model, we included learning block, medication state, and the interaction between the two.

No regions showed an overall effect of medication on average striatal allegiance. A region in the right fusiform gyrus with a higher correspondence to object areas exhibited a significant effect of learning block (p<0.05 FDR corrected; **Figure 4.5A**, top), as well as a significant dopaminergic modulation of these temporal changes (p<0.05 FDR corrected; **Figure 4.5B**).

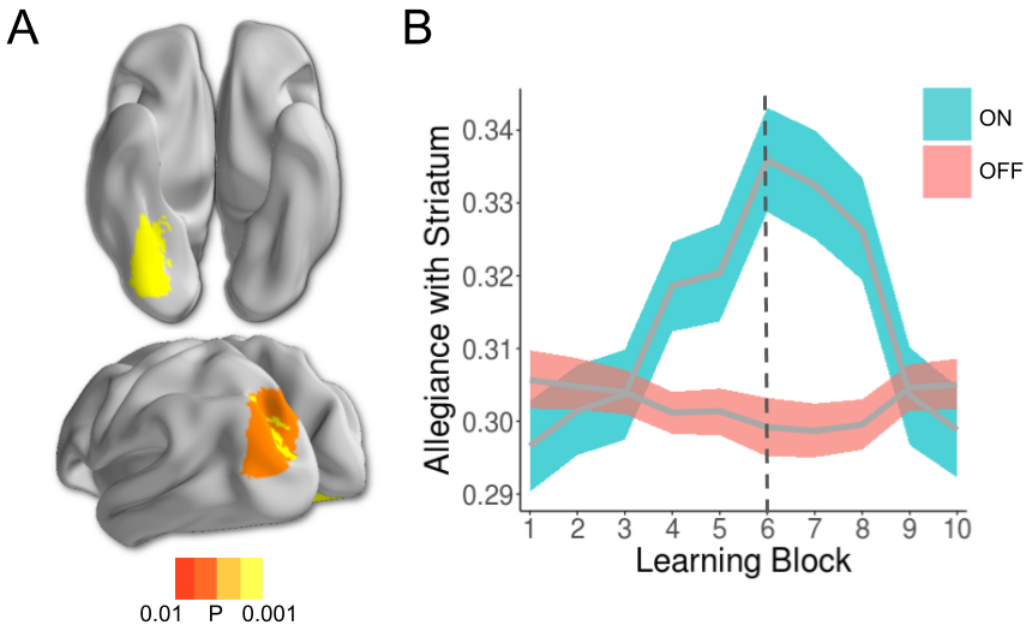

Figure 4.5 Dopamine modulates dynamic connectivity between the striatum and task-relevant visual regions. A. Regions exhibiting changes in module allegiance with the striatum over the course of the task. Shown at a threshold of $p < 0.01$, these regions include the fusiform gyrus (*top*), inferior parietal lobe (*bottom*), and lateral superior occipital gyrus (*bottom*). Based on Neurosynth maps, these regions showed a greater correspondence to object-processing areas than would be expected by chance alone. **B.** Interaction between medication state and changes in allegiance in the fusiform gyrus region shown in panel **A** (*top*). This area of inferior temporal cortex showed significant temporal changes in striatal allegiance, as well as a significant dopaminergic effect on these changes (both FDR p<0.05). Solid lines show mean allegiance with the striatum and ribbons show standard error. Colors indicate medication group (blue=ON, red=OFF).

This result demonstrates that dynamic connectivity between the striatum and a category-specific visual region is altered by dopamine, but the analysis was limited to visual areas. To test the specificity of these effects, we used a lower threshold (p<0.01, uncorrected) across all brain regions. Two more regions in occipital and inferior parietal lobe (IPL) passed this more liberal threshold for an effect of time on striatal allegiance (**Figure 4.5A**, bottom), and the same IPL region passed this threshold for a significant dopamine-by-time interaction. These regions showed a significantly higher level of object correspondence than would be expected due to chance alone (p=0.014 for the effect of learning block; p < 0.005 for the block-by-medication interaction; see **Methods** for description of bootstrapping procedure). Collectively, these results provide further evidence of a dopaminergic effect on the dynamics of striatal connectivity with task-relevant regions during learning.

Discussion

There is growing evidence that dynamic changes in cortico-striatal connectivity play an important role in value-based learning (Mattar et al., 2016; Gerraty et al., 2018). Here, we provide a learning-based experimental manipulation of dynamic cortico-striatal connectivity, providing causal support to previously demonstrated associations between network flexibility in the striatum and learning. In addition, we show that changes in striatal connectivity are modulated by the neurotransmitter dopamine, which has been theorized to facilitate learning by updating associations via changes in striatal cortical connections (White, 1989a; Glimcher, 2011). Our finding that dopamine-induced changes in

cortico-striatal coupling are most pronounced in domain-specific sensory regions provides empirical support for this theory.

Dopamine is known to play an essential role in reinforcement learning. Evidence for this role comes from experimental manipulation of dopamine in non-human animals (Beninger, 1983), studies in healthy controls (Pizzagalli et al., 2008), and studies in patients with Parkinson's disease (Knowlton et al., 1996; Shiner et al., 2012). Further evidence comes from the finding that midbrain dopaminergic neurons report an analogue to the temporal difference prediction error postulated in reinforcement learning algorithms (Schultz et al., 1997; Bayer and Glimcher, 2005). The latter finding has been particularly influential, as it provides a putative computational mechanism for dopamine's role in learning. While the prediction error account provides a parsimonious description of the firing of midbrain neurons during learning, it leaves open the question of how this signal actually affects learned associations.

The dominant theory for how prediction errors update associations between stimuli, actions, and value outcomes is via changes in synaptic weights to targets of midbrain dopamine neurons that are active following a given state and action (Glimcher, 2011). While this theory is intuitive and consistent with known anatomy and neurobiology, it has proven difficult to test empirically. This is because most measures have focused on the activity of single neurons or brain regions. Even measures that can be used to characterize brain networks have been hampered by an inability to describe dynamic changes in these circuits. Our report takes advantage of recent developments in network neuroscience that afford precisely this ability to characterize circuit dynamics (Bassett et al., 2013a; Khambhati et al., 2017; Sizemore and Bassett, 2017).

Network neuroscience has provided a wealth of experimental evidence that the characteristics of distributed brain circuits are relevant to cognition (see (Medaglia et al., 2015; Petersen and Sporns, 2015; Bassett and Sporns, 2017) for review). Despite this promise, it has been difficult to link findings from network neuroscience to established theories in cognitive neuroscience. Thus it has been hard to articulate the specific implications of our increasing study of brain networks for our understanding of cognition in general. Our findings in this study provide precisely such a bridge, and illustrate the potential utility of characterizing the brain's dynamic network structure for answering essential questions about the mind. Here we show that cortico-striatal circuits change as a result of reinforcement learning associations, and that these changes are modulated by dopamine.

More recently, the completeness of the reward prediction error account of dopamine's role in reinforcement learning has been called into question. A number of findings indicate that so-called model-free reinforcement learning or value-based decision making tasks evoke activity in systems associated with episodic or associative memory (Shadlen and Shohamy, 2016; Bornstein et al., 2017; Bakkour et al., 2018). In addition, there is increasing evidence that midbrain dopamine itself is essential for more complex so-called model-based associations, which rely on cognitive maps of environmental or conceptual contingencies (Deserno et al., 2015; Sharp et al., 2015; Sadacca et al., 2016; Sharpe et al., 2017). A causal role for dopaminergic circuits in model-based learning would indicate that this neurotransmitter updates associations that extend beyond cue- and action-outcome links, perhaps reporting a successor representation of statistical associations between states of the world (Dayan 1993; Gershman and Schoenbaum 2017). If dopaminer-

gic updating does underlie the formation of complex stimulus-stimulus associations, one would expect to find dynamic changes in striatal-cortical and cortical-cortical connectivity modulated by dopamine on more complex tasks. The same network tools utilized in this report are of potential use for providing this characterization.

This report has a number of limitations. The multilayer community detection algorithm we used provides deterministic community labels, and are unable to assess the possibility of regions coupling with multiple communities at a given time point. Addressing this limitation may be important for providing a fuller account of dynamic changes in brain networks during learning and other cognitive processes. The further development of probabilistic and generative models, which pose serious statistical and computational challenges, will be essential in this regard (Durante et al., 2016; Palla et al., 2016; Betzel and Bassett, 2017). In addition, the use of fMRI, which affords a crucial level of spatial specificity, limits the temporal resolution with which we can describe brain network dynamics. To better characterize network changes underlying learning, it will be useful to measure the effect of individual prediction errors on network structure, and network studies using ECoG or combining fMRI and EEG may prove useful in this regard. A limitation that is unique to this study is the use of a patient group with no control participants. We chose to focus on a within-subject pharmacological manipulation so as to limit the number of terms in the statistical interactions necessary to characterize dopaminergic effects, thereby limiting the uncertainty in our inferences. It is possible that these results are specific to patients with Parkinson's disease; however, this concern may be mitigated by the fact that patients learned the correct choice and updated their decisions following re-

versal. Nevertheless, further research could usefully evaluate similar hypotheses in neurologically intact humans.

While network science has produced a growing number of results linking large-scale brain circuits to cognitive processes such as learning, it is sometimes unclear what the theoretical implications of these results are. In this study we sought to characterize, based on convergence between previous studies of dynamic network changes during learning and computational theories of reinforcement, the effect of dopamine on learning-related changes in striatal-cortical coupling. In addition to providing experimental support linking dynamic connectivity to feedback-based learning and a neurotransmitter system known to underlie this process, this study illustrates the potential utility of integrating network neuroscience with more established theoretic frameworks for understanding cognition.

Chapter 5

Network neuroscience: phenomena in search of a concept?

Such notions are speculative and vague, but we seem to have no choice but to be vague or to be wrong, and I believe that a confession of ignorance is more hopeful for progress than a false assumption of knowledge.

-K.S. Lashley, 1930, Basic Neural Mechanisms in Behavior

The preceding chapters in this dissertation provide evidence that large-scale network connectivity, and in particular dynamic connectivity, plays a role in multiple forms of learning. In doing so, they highlight the promise of a network approach for linking brain activity to behavior. They also illustrate some of the challenges in gleaning theoretical insight into cognition from this network characterization, as well as the potential of surmounting these challenges.

Chapter 2 demonstrated that value transfer on a sensory precondition task is associated with intrinsic connectivity between the hippocampus and ventromedial prefrontal cortex, two regions known to be essential to this behavior (Port et al., 1987; Jones et al., 2012), as well as between these regions and more domain-general connectivity networks (Gerraty et al., 2014). Understanding the mechanisms underlying this link is impossible, however, without measuring changes in network connectivity that take place during learning. The study presented in Chapter 3 used emerging techniques from network science to describe changes in network connectivity that take place during reinforcement learning (Gerraty et al., 2018). The striatum exhibited an increase in dynamic connectivity and coupled more with visual- and value-processing regions over the course of a learn-

ing task, and this dynamic connectivity correlated with performance. In an attempt to ground interpretation of this result in light of known cognitive and neural processes, Chapter 4 tested patients with Parkinson's disease on and off of dopaminergic medication on a reversal learning task, allowing for the experimental manipulation of both learning and dopamine. This study further demonstrates that learning involves dynamic changes in cortico-striatal connectivity, and shows that these changes are modulated by dopamine, providing empirical evidence for network-level processes that have been theorized to underlie reinforcement learning.

Each study presents evidence for the relationship between network connectivity and learning behavior, and each faced serious interpretational roadblocks. Hopefully it is clear from the trajectory of these chapters, with each study building on the limitations of the last, that these challenges are not insurmountable. However, many difficulties remain for integrating network science into cognitive neuroscience. The goal of this chapter is to situate the studies presented in this dissertation in a broader description of the promise and challenges facing the network neuroscience in terms of linking empirical findings about brain circuits to an understanding of how cognition might emerge from their dynamics.

Introduction: **the promise of a network neuroscience**

Debates over whether brain functions are best described as local or distributed have a long and storied history in psychology and neuroscience. From foundational arguments between proponents of the neuron doctrine and reticular theory (Yuste, 2015), to the cranial-based phrenology of Gall and lesion-based localization of Broca (Zola-

Morgan, 1995), Karl Lashley's arguments against the cortical reflex arc (Lashley, 1930), and ever-ongoing debates about connectionist neural networks versus modular computational architectures (Fodor, 1975; Smolensky, 1990; Botvinick et al., 2017; Lake et al., 2017), the local vs. distributed question has been a recurrent backdrop to the history of studying the mind.

It was noted almost twenty years ago that, while the advent of whole-brain human neuroimaging techniques had led to recognition that most cognitive tasks evoke distributed rather than localized activation, the relationship between cognitive processes and large-scale brain activity had largely not been worked out (McIntosh, 2000). In fact, it is easy to point to successes from functional neuroimaging in which we have learned something about the computations of an individual brain region, often for perceptual processes such as face perception (Kanwisher et al., 1997); however, while large-scale neuroimaging has taught us that most brain activity involves coordinated interactions across wide swaths of the cortex (Bressler, 1995; Power et al., 2011; Raichle, 2015), today it is still difficult to articulate what we have learned about the inner workings of the mind from this fact.

More recently, graph theoretic tools have been brought to bear on this question, characterizing the topological structure apparent in interactions between brain areas (Sporns, 2003; Bassett and Bullmore, 2006; Bullmore and Sporns, 2009). Over the past decade, driven in part by advances in dynamic applications of graph theory, these approaches have coalesced into a subfield known as network neuroscience (Muldoon and Bassett, 2014; Medaglia et al., 2015; Bassett and Sporns, 2017). The tools and approach of network neuroscience have been touted as potential solutions to a wide array of prob-

lems we face in our attempts to understand the brain. It has been claimed that this emerging field can help us solve previously intractable issues in development (Fair et al., 2009), learning (Bassett and Mattar, 2017), psychopathology (Menon, 2011; Braun et al., 2016), and the general ontology of cognition (Bertolero et al., 2015), among many others. Indeed, there has been an ever-increasing diversity of empirical findings linking brain networks to many aspects of cognition and behavior, from basic perceptual processes (Moussa et al., 2011; Ekman et al., 2012) to social interactions (Barrett and Satpute, 2013) and consciousness (Lee et al., 2010).

However, it is not clear what such findings mean, at least for our theories of cognition. We have described many network phenomena in the brain, and related some of them to behavior, but in many respects we lack a conceptual framework for understanding this relationship. There is still much work to be done to integrate network ideas into established cognitive neuroscience theory. In this review, I will describe what I see as the biggest challenges to this integration—those facing network measures as they are currently used and interpreted, as well as problems with established theories in cognitive neuroscience as they relate to (or fail to relate to) the dynamics of brain circuits—and look forward towards areas of research that seem ripe for incorporating network ideas. It may be premature or unfair to evaluate the potential of a field that has just begun to cohere, but with that caveat in mind this review represents something of an internal audit of the promises and pitfalls facing the study of large-scale brain networks.

Before taking stock it is important to highlight a distinction between two ways in which network neuroscience could be useful for understanding the link between the brain and behavior. One can consider the utility of the emerging battery of network methods,

which provide new ways of measuring and interpreting brain activity and relating it to behavior. One can also consider the contributions to a theoretical understanding of cognitive or behavioral processes afforded by modeling the brain as a network. Of course, there is no way to completely separate methodological approaches from theoretical contributions. In fact, as should become clear, our theoretical vocabulary is often dictated or at least constrained by the measurement devices available to us at any point in time, and it has been argued that a focus on single-neuron (as opposed to circuit) activity as the fundamental functional unit of the brain has been largely influenced by technological constraints (Yuste, 2015). Still, it will be useful to consider each in turn when evaluating the network approach. While it is clear that network neuroscience has been extremely useful in terms of emerging methods, providing theoretical contributions to our understanding of cognition presents a much more difficult challenge.

As another note before continuing, one of the strengths of network science is that it provides a mathematical framework that generalizes across multiple domains. Network theory can be applied to data from sociology, political science, ecology, and many other areas of investigation outside of neuroscience. Even within neuroscience, it can be applied to the behavior of proteins, neurons, or of brain areas. Because this review is meant to summarize the relevance of networks to cognitive neuroscience and the challenges facing this link, the focus will be on network approaches applied to interactions across distributed brain regions. The focus of cognitive neuroscience has been on the relationship between brain activity at the macro- and meso-scopic scale, and much of the work linking networks to cognition has followed suit. But many of the points of discussion will easily apply to thinking about neuron-level circuits and their dynamics.

Theoretical contributions of network neuroscience

A methodological frontier

What exactly does network neuroscience bring to the table in terms of theoretical insight into how the brain produces cognition? Maintaining the distinction sketched out above, one obvious answer is a suite of new and interesting methods for measuring, summarizing, and partitioning relational structures among brain regions. Network methods provide descriptions of (usually pairwise, but see (Giusti et al., 2016)) interactions between brain regions, and at its most cutting edge, the evolution of these interactions over time. These methods and their application to brain data are genuinely novel, and provide descriptions of brain activity that were previously impossible.

An example of this dynamic relevant throughout this review is the development of multiplex community detection in graph theory (Mucha et al., 2010). Graph theory, one of the most widely used network tools for analyzing brain data, summarizes brain activity (or any measurement) in terms of pairwise adjacency (correlation, coherence, etc.) between nodes V on graph G, with adjacency represented by edges E. A detailed description of graph theory's application to brain data is outside of the scope of this review, but there are a number of sources where this is provided (see (Sporns, 2010) for comprehensive review). One of the most useful tools for describing graphs is community detection, a set of algorithms for decomposing a network into a set of communities or modules (Fortunato, 2010). On historic limitation of community detection tools had been their inability to link communities across time. In some ways the 2010 development of multiplex community detection (Mucha et al., 2010) marks a turning point in the application of

networks to cognitive neuroscience, allowing researchers to link network changes to dynamic behaviors. A purely mathematical tool opened up an expansive and fertile area for asking questions about the relationship between the brain, cognition, and behavior.

But what exactly are the questions being asked by network neuroscientists with these tools, and how are they relevant to understanding cognition? What specific aspects of neural function would a focus on networks uniquely elucidate? Surely, the fact that network neuroscience is emerging as its own sub-discipline, with a journal, conferences, and dozens of review papers on the topic, indicates that the approach has something more to offer than novel tools, even if those tools can potentially lead to future insights. Network neuroscience should (and seems to) offer a novel theoretical vantage point for understanding the relationship between the brain and cognition. But the exact nature of the insight offered by the field for theories of cognition has been difficult to articulate.

Theory in cognitive (neuro)science

Part of the reason for this difficulty in discerning genuine theoretical insights about cognition from network neuroscience seems both unavoidable and essential to the emergence of the field. Because cognitive neuroscience has been focused primarily in mapping cognitive processes into activity in local and distinct brain areas (Posner and DiGirolamo, 2000), there is a lack of vocabulary for describing the potential roles of interactions between brain regions in cognition.

Cognitive science is a framework describing the processes underlying mental function as series of operations over abstract symbolic representations (Pylyshyn, 1980). These cognitive processes are evaluated empirically by measuring behavioral output such

as reaction times and are agnostic with respect to physical brain mechanisms. In fact, the computational and philosophical foundations of cognitive science dictate that these processes must necessarily be invariant with respect to physical instantiation (Cummins, 1989).

Cognitive neuroscience has traditionally and in large-part been a search for brain regions underlying these cognitive operations, providing a map of the *structures* in the brain that support cognitive processes (Posner et al., 1988). The mostly implicit assumption has been that cognitive functions will be *localized* to individual brain regions. This assumption is inherent in the traditional tendency to "reverse infer" function from regions passing statistical thresholds for association with a specific dimension of an experiment (Poldrack, 2011). Associations between activity, or more recently multivariate patterns of activity, in a brain region, and a behavioral or putative computational variable are taken as evidence that this brain region carries out a specific cognitive operation. There is reason to believe, however, that a network representation of the brain will be relevant for describing the processes underlying cognition.

One finding from human neuroimaging that has been difficult to square with this standard cognitive view has been the discovery of large-scale brain networks. Efforts to decompose brain activity into independent or functional components, even in the absence of stimuli, has revealed that outside of sensory cortices this decomposition does not lead to local clusters, but rather to patterns distributed across the whole brain (Beckmann et al., 2005; Power et al., 2011; Smith et al., 2012; Raichle, 2015). These networks are extremely robust; their stability has been demonstrated across individuals and over time (Zuo et al., 2010; Wang et al., 2011).

Most methods for characterizing these networks begin with a fundamentally different way of representing brain activity, using interactions between regions as the fundamental component of representation. It should come as no surprise, then, that it has been a challenge incorporating findings from network neuroscience, presented in terms of interactions, into the traditional landscape of cognitive functions. The standard inference that a statistical association provides evidence of a local underlying computation is clearly insufficient when your dependent variable is a set of distributed associations.

The theoretical stance implicit in network neuroscience, then, is that we must consider interactions between brain regions (or units at other levels of brain activity) in order to understand function. That is, a sufficient description of cognitive processes and the production of complex behavior must include descriptions not only of the activities of individual neurons or brain regions, but of interactions between them. Part of the issue in integrating network neuroscience findings into established theories of cognition is the lack of recognition in cognitive neuroscience that dynamic activity over distributed circuits may represent an irreducible component of the production of the complex behaviors of interest to psychologists and cognitive scientists. In this way, in addition to tools for examining complex interactions in the brain and their relationship to behavior, network neuroscience provides the insight that understanding such interactions represents an essential goal for psychology and cognitive science.

And studies have demonstrated relationships between measures of network function and a host of cognitive and behavioral processes, including but not limited to perception (Moussa et al., 2011; Ekman et al., 2012), motor function (Xu et al. 2014; Bassett et al. 2015), attention and executive function (Wang et al., 2009; Bartolomeo et al., 2012),

learning (Bassett et al., 2011; Gerraty et al., 2014; Gerraty et al., 2018), memory (Nyberg et al., 2000; Wang et al., 2010), and social cognition (Barrett and Satpute, 2013). The ubiquity of this general finding that network-level measures are relevant to many cognitive and behavioral processes validates the foundational claim of network neuroscience: that we must understand circuit-level brain interactions to understand cognition. In some ways however, this apparent fact raises more questions than it answers. What is the nature of these network-level processes? What aspects of cognition and behavior rely on them? Are they domain general, with the same aspects of network processing relevant for all of cognition, or are their distinct network mechanisms underlying particular cognitive domains?

The need for a theory of brain networks

What may be missing for answering these questions is a specific description of the ways in which internal processing of particular stimuli or the production of a particular behavior *requires* dynamic activity over large-scale circuits. While fixed, anatomical connectivity between regions is consistent with and even included in many cognitive models, the potential that given circuits can turn "on" or "off" or be weighted differently depending on context is missing. I will argue here that not only is this possible, but that some behaviors depend on such dynamic changes in neural circuitry, in a way that cannot be reduced to the standard descriptions of cognitive neuroscience. It is in this area that network neuroscience seems to have the most theoretic potential.

To begin with a very simple behavior, imagine a cell which controls eye movements to a particular point in visual space. Cells in the Frontal Eye Fields do just this,

possessing receptive fields for subsequent saccades(Mohler et al., 1973; Bruce and Goldberg, 1985). Current theories about perceptual decisions indicate that such a cell would receive input from specific neurons in the Lateral Inferior Parietal (LIP) cortex, which accumulate evidence for initiating eye movements, exhibiting ramping activity reflecting the likelihood of taking a specific action, such as saccading to the cell's receptive field (Gold and Shadlen, 2007). It is clear that animals can be trained to use different sources of information for an eye movement in different contexts.

These can range from random dot motion in different directions, which would depend on the same LIP neurons interacting with distinct cells in area MT (Newsome et al., 1989; Hanks et al., 2006), or other visual stimulus dimensions (which would depend on accumulating evidence from different visual areas). An animal could learn that the optimal action depends on evidence from stimuli in different parts of the visual field, which would depend on the same LIP neurons interacting with different cells in any visual area with retinotopic mapping. The animal could also learn that the eye movement should depend on properties of an auditory stimulus (Mazzoni et al., 1996) or on a mixture of multimodal stimuli, which would depend on widespread integration of evidence. Perhaps most critically, the animal could learn to rapidly switch between evidence sources depending on context, for instance when primed with an instructive stimulus before each trial (Siegel et al., 2015).

Thus the same behavioral output can rely on a dynamic repertoire of input stimuli and internal processing via dynamic activity over distributed brain circuits. It is the network of regions active and the interactions between them—rather than the activity in any one cell or brain region—which characterize the internal processes transforming sensory

input to behavioral output, even in this extremely simple task. Add to this the fact that the animal could clearly map any of these stimulus examples to multiple motor outputs in a similar fashion and the necessity of dynamic circuit activity should become clear. These circuits are not necessarily local, but distributed across the brain, involving the widespread integration of information; and they are dynamic as well as context dependent, performing different functions depending on the set of co-activated regions with which they are interacting.

How exactly this takes place is unclear. It has been proposed that a group of interneurons in the thalamus control the routing of such information across circuits in the cortex (Kastner and Saalman, 2012; Mitchell et al., 2014; Roth et al., 2016). Ongoing technological development in systems- and circuit-level neuroscience of tools for recording from neurons based on their large-scale connectivity will be necessary to begin to tease apart the mechanisms underlying this process, but they will play an integral role in understanding how networks function in cognition, providing a detailed description down to the cellular level.

The ideas presented above are similar to a concept introduced twenty years ago by A.R. McIntosh in an earlier attempt to integrate networks into cognitive science (McIntosh, 2000). Drawing on the brain's widespread anatomical connectivity as well as the notion of *transient response plasticity*—the ability of neurons throughout cortex to modify their response properties depending on context—McIntosh introduced the concept of "neural aggregates" in contrast to the notion that cognitive processes are mediated by specific and discrete brain systems. Instead, he argued that cognitive functions are the result of aggregated properties of active regions, and are thus determined by dynamic and

context dependent interactions. Network neuroscience provides tools to examine such neural aggregation empirically.

Mapping networks to cognition

The challenge moving forward is to find a language or framework for describing cognitive processes in terms of overlapping patterns of groups of interacting brain regions. There are a number of possibilities for how this might develop, some more challenging to traditional ideas in cognitive neuroscience than others.

It is possible that linking brain networks to cognitive neuroscience will be as simple as showing which large-scale networks map onto which cognitive domain. In this scheme, brain networks will play the same role as brain regions in older theories: they will be the (distributed) locus of cognitive processes. This tendency is already somewhat prevalent with the increase in use of network methods: networks are often given (sometimes conflicting) cognitive labels such as the dorsal and ventral attention networks, the emotional salience networks, theory of mind network, and others, indicating that they map directly onto a specific cognitive task. This speaks to the extent to which one-to-one functional mapping is ingrained into the cognitive neuroscience framework.

Alternatively, it is unclear that the domains carved out by cognitive science represent fundamental or discrete computations that we should expect to find in the brain at all. It is possible that, in terms of processes in which the brain is engaged, traditional cognitive domains do not represent the right constructs. If we imagine decomposing complex behavior (say across a large number of experimental tasks) into independent constituent parts, we may find that traditional cognitive processes to not describe the re-

sulting factors, in much the same way that brain regions turn out to be poor candidates for constituent (macroscopic) components of brain activity.

To further complicate matters, it is evident that the brain's connectivity structure is more complex than canonical large-scale networks appear. Brain networks have been shown to undergo non-trivial fluctuations in topology (Allen et al., 2012; Smith et al., 2012; Liu and Duyn, 2013), indicating that as currently characterized they represent static snapshots of much more dynamic phenomena. The question for a new cognitive neuroscience, one that is able to incorporate a network perspective, is: what are the relationships between dynamic brain networks and the constituent processes underlying complex behavior? Given the uncertainty inherent in this question, it will be important to focus on specific behaviors and develop theories of network mechanisms that might underlie them.

Challenges facing network neuroscience

Theoretical contributions to cognitive neuroscience

While it is clear that network neuroscience provides novel ways of describing brain activity in terms of interactions and relating these interactions to behavior, many questions remain about the theoretical implications of treating the brain as a network. What do brain networks explain about cognitive processes that standard approaches, in which cognition is described by computations generated by complex populations of cells localized to individual brain regions? What does the discovery of different large-scale networks tell us about cognitive ontology? Do different networks serve distinct cognitive functions akin to traditional cognitive or neuropsychological domains previously attributed to brain regions? Or do cognitive domains represent constructs that each rely on the

interaction of multiple brain networks? Perhaps most importantly, how do these networks change over time, and how do network dynamics relate to cognition?

One way to frame these questions is in terms of computational modeling. If one were to develop a model of a specific cognitive process incorporating the findings of network neuroscience, what would be the form of the resulting model? How would such a model be distinct from traditional computational models? It is essential to explore these questions if network neuroscience is to be integrated with our current theories and impact our understanding of the mechanisms giving rise to cognition.

Underlying neural mechanisms

There has been some detailed work in modeling the neurophysiological underpinnings of the large-scale correlations that generate distributed brain networks at rest (Deco et al., 2011). Much less is known about the effect of task-evoked activation in these models, and even less about their relation to dynamic connectivity changes during task performance. So in addition to challenges in articulating the *roles* of dynamic connectivity changes in cognition, it has been very difficult to explain how these changes are generated by underlying neuron- or population-level processes. One example is the network measure of flexibility, which is thought to index dynamic coupling between regions and broader networks. Are fluctuations in network connectivity the result of changes at the synapse? Or would a gating mechanism, perhaps controlled by thalamic interneurons, provide a better explanation? Again, it is helpful to frame these questions in terms of explicit models. If we were to build a model of dynamic changes in distributed inter-region interactions from spiking neurons or mean field populations, what would it look like and

how would it relate to behavior? A key aspect of building theoretic insights from the empirical successes of network neuroscience will be developing generative models of the underlying neurophysiology.

Interpretation of network metrics

Despite large gaps in our understanding of the neurophysiological underpinnings of network-level processes, in network neuroscience studies topological measures and network diagnostics are often given specific interpretations in relation to neural or cognitive data. It can be unclear to what extent these interpretations are appropriate or valid.

For instance, communities into which networks are decomposed are often known as modules, and graphs in network science are often said to possess modular structure, meaning that there are more connections clustered within groups than would be expected by random connectivity. What is the relationship between modules in community detection and the modular architectures proposed in traditional cognitive theories? It is clear that graph theory modules are neither localized nor encapsulated, but it has been asserted that the modular organization of brain networks is related to the modular organization of the mind (Bertolero et al., 2015). However, a precise mapping of the general tendency towards information segregation captured by graph theoretic modularity and the encapsulation and serial processing conveyed by cognitive notions of modularity has not yet been provided.

As another example, many studies have shown that measures of efficiency from graph theory relate to a number of cognitive processes, both at rest and during task performance (Heitger et al., 2012; Langer et al., 2012; Ma et al., 2012; Gießing et al., 2013;

Shine et al., 2016). High clustering is often interpreted as evidence of efficient local information flow (Bullmore and Sporns, 2009; Sporns, 2011), while measures of path length are taken to index long-distance transmission of information or global segregation (Achard and Bullmore, 2007). While those measures have specific meaning in network science, there has been no empirical validation that they accurately index information flow in neuronal firing, or concepts of efficiency in cognitive science.

The notion of flexibility, which will be crucial to the studies described in the next section, is also somewhat ambiguous. Network flexibility has been interpreted as evidence of sequential switching of the networks each brain region communicates with (Bassett et al., 2011). And in a useful illustration of many of the challenges discussed in this review, network flexibility has been correlated with cognitive flexibility (Braun et al., 2015), however the mechanism by which processes captured by this network metric would lead to increases in flexibility as understood cognitively remains largely unclear. Given the ambiguities in terms of underlying neurophysiology described above and potential issues in estimation at the algorithmic level described in the next section, caution should be exercised in the interpretation of this and other network diagnostics. And crucially, we should lay out clearly potential explanatory mechanisms when relating these measures to behavior, and test these mechanisms empirically.

Probabilistic modeling

Many of the successes in network neuroscience have come from the incorporation of tools from signal processing and complex systems theory. For example, pairwise relationships between regions in a network are often summarized using measures like mean

squared coherence, networks are decomposed into communities, and diagnostics are used to summarize networks' topological properties. Many of the tools used are deterministic: they do not incorporate uncertainty in their estimates. Community detection algorithms, for example, minimize a loss (or maximize a quality) function to provide a single network assignment to each node (Bassett et al., 2013a; Garcia et al., 2017). While these tools have been enormously successful, when fitting network models to brain data to relate brain networks to behavior, we are faced with uncertainty at every level of analysis.

Representing this uncertainty with probability would be beneficial for a number of reasons. It would allow us to pool our uncertainty across observations, regularizing our estimates for individual subjects or brain regions (Gelman, 2006a). Using probability to represent uncertainty also provides the optimal inference for a specified model, and avoids pitfalls associated with point estimates in high dimensional or bounded parameter spaces (Chung et al., 2015). In addition, current network models often have hyperparameters, which are either set arbitrarily or maximized with respect to a desired function (Bassett et al., 2013a; Garcia et al., 2017). In a probabilistic model, uncertainty about hyperparameters could be easily incorporated into inferences at all levels of analysis. Finally, as described in more detail below, it would relieve some of the ambiguity present in measures of dynamic changes in network structure, by separately modeling fluctuations in network interactions, "simultaneous" interactions with multiple networks, and noise in our estimates of community structure. There has been a recent increase in research into probabilistic generative models of network structure (Durante et al., 2016; Palla et al., 2016; Betzel and Bassett, 2017), however these models remain to be fully validated and

are not yet widely applied. These endeavors will clarify a number of interpretational issues in network neuroscience.

An example: the network neuroscience of learning

The studies presented in this dissertation provide a useful illustration of the promises and challenges described thus far. Given the repeated findings linking large-scale network connectivity to multiple cognitive processes, the study presented in Chapter 2 aimed to examine the relationship between intrinsic connectivity and learning. Using standard static approaches to quantifying interactions between brain regions, this study showed that connectivity between the hippocampus and ventromedial prefrontal cortex was related to a specific aspect of learning performance (value transfer) on a sensory preconditioning task (Gerraty et al., 2014). The same was true of connectivity between these two regions, both of which are known to be essential for performing this task, and the default mode and frontoparietal networks, distributed sets of regions which seem to play a more general role in cognition. This finding provides further support to the thrust of network neuroscience's theoretical stance: that interactions between regions matter for cognition. They also reveal the weakness of static network approaches for filling in the details of this relationship: the inability to measure changes in network interactions during learning limits the theoretical insight offered by the finding. What is the mechanism through which differences in these large-scale patterns relate to performance? With static measures of correlation, we are left with the unsatisfying conclusion that connectivity must matter *somehow*.

Chapter 3 was motivated by the idea that given new methods to characterize the dynamics of this network, we could circumvent this issue. Recognizing the explanatory limitations of static measures of connectivity, this study examined the relationship between dynamic changes in network connectivity centered on the striatum and reinforcement learning, which is known to depend on this region (Gerraty et al., 2018). We used multiplex community detection, measuring temporal changes in network structure, and network flexibility, which as described above indexes the extent to which each region is changing the communities or networks it couples with over time.

It is worth pointing out a few ambiguities with the dynamic network metrics used in this study. Flexibility is technically the proportion of time windows during which a region "changes" its community allegiance. That is to say, it is the proportion of time windows for which the community detection algorithm provides a different community label for a given region. There are at least two aspects of this measure that are already ambiguous.

The first has to do with the method of community detection used, and the second has to do with the interpretation of a change in community. First, the methods used for community detection are deterministic: each region is labeled with a single community and no measure of uncertainty is provided. Thus it is unclear in general whether a change in community structure could be due to uncertainty about community assignment. Averaging over multiple estimations of community detection will alleviate this concern somewhat, but a probabilistic generative model of community detection would provide more precise differentiation between changes in community structure and uncertainty

about community assignment, among other benefits (see **Challenges to Network Neuroscience**, above).

Second, and relatedly, it is unclear whether, or perhaps more precisely *when*, a region's flexible connectivity represents time-dependent changes in community or simply interactions with multiple communities. In other words, in addition to algorithm-level uncertainty, flexibility could mean that a region is switching the networks in which it takes part, or it could mean that a region is communicating with multiple networks. The first point here is about potential noise, and the second is about potential signal, that deterministic community detection models are unable to detect or differentiate.

Regardless of the exact interpretation of flexibility, it was shown to relate to reinforcement learning in this study. Dynamic patterns of network coupling between the striatum and cortical areas processing visual stimuli and value information increased during learning, and flexible striatal connectivity was shown to increase with learning performance. Flexible connectivity in the striatum also correlated with individual differences in parameters from standard reinforcement learning models. This finding is consistent with previous reports relating flexibility with motor learning (Bassett et al., 2011), but also provides some anatomical specificity, showing that network flexibility centered on the striatum, which is essential for learning based on feedback, was related to reinforcement learning performance.

Given the ambiguities described above, the study in Chapter 3 also sketched out a potential theoretical framework for linking dynamic connectivity to the processes underlying reinforcement learning behavior. In this framework, dynamic connectivity metrics such as flexibility and allegiance capture the routing or gating of information that must be

integrated for a given decision. According to this idea, an essential part of learning the optimal actions in a given context is learning *which* relevant environmental variables to evaluate as evidence towards *what* motor or decision output. Thus "carving out" the relevant circuits for a decision in the right behavioral context is a necessary part of learning for which dynamic coupling between distributed brain areas plays an essential role.

This notion of dynamic connectivity in the learning process is consistent with the putative mechanism by which dopamine updates choice values in established theories of reinforcement learning. In fact, it can be argued that network neuroscience measures of striatal-cortical coupling fills a gap in these theories, providing the means to measure large-scale circuit-level changes thought to be effected by dopamine. Dopamine neurons in the midbrain are known to signal temporal difference prediction errors to striatal and prefrontal targets (Schultz et al., 1997; Pagnoni et al., 2002; Bayer and Glimcher, 2005). In reinforcement learning algorithms developed in engineering and computer science, these prediction errors are used to update the value of states and actions based on the surprise of outcomes (Sutton and Barto, 1998). The neural mechanisms underlying this update (that is the nature of changes in prediction error targets) are uncertain, but it is theorized that the levels of dopamine generated by midbrain prediction errors strengthen or weaken connections in striatal-cortical circuits that are active at the time of reward receipt (Glimcher, 2011). Thus, reinforcement learning theory also posits widespread circuit-level changes as the locus of associative learning. Measures of dynamic connectivity in network neuroscience can help add necessary detail to this theory, and allow for the manipulation and measurement of changes in striatal-cortical connectivity.

While these ideas seem to converge, there has been no empirical evidence of a relationship between dopamine and dynamic connectivity. The study in chapter 4 establishes this connection, testing patients with Parkinson's disease on and off of dopaminergic medication during a reversal-learning task. The results provide further evidence that learning alters dynamic connectivity between the striatum and relevant sensory areas, and demonstrates that these learning-related changes are modulated by dopamine. This is an encouraging finding, one that highlights the utility of network neuroscience methods (and a focus on the dynamics of brain networks, in particular) for helping to unravel a theoretical question in cognitive neuroscience research: how a midbrain dopamine prediction error comes to update value associations.

The future

One of the lessons that can be drawn from the above discussion is that, while there may be fundamental principles of brain network function that hold across behaviors and cognitive domains, it is helpful to specify particular behaviors and develop theories about how they depend on dynamics over relevant brain circuits. With this in mind, this review will conclude with a look forward to areas of research outside of the reinforcement learning theories described above for which network neuroscience may provide a particularly useful perspective.

The role of memory sampling in decision-making

A number of recent reports indicate that decisions between options which are either familiar or have been experienced repeatedly rely in part on episodic memory processes (Bornstein et al. 2017; Bakkour et al. 2018; Gerraty et al. in preparation). These findings provide an example of the difficulty of strictly mapping behaviors onto individual cognitive domains, as such decisions were previously thought to rely on incremental learning and cached valuation systems. The involvement of episodic memory indicates that a description based on these systems is insufficient. Relevant to the purposes of this review, a recent theory about how memory enters into the decision process (Shadlen and Shohamy, 2016), suggests that the hippocampus activates cortical representations associated with memories related to choice options. These representations are theorized to enter into evidence for specific actions via a thalamo-cortical "pipe" routing information from distributed regions to areas such as LIP, which accumulate evidence for motor output. Crucially, these memory associations are thought to be sampled sequentially for any given input, meaning that the process of evidence accumulation takes place via dynamic activity over changing circuits as memories are sampled. While there have been no empirical investigations examining this process, it is clear that network neuroscience would provide the perspective and tools to characterize the dynamic circuit-level activation underlying sampling from memory for value-based decisions.

Information routing in neural network models

Recent years have seen a resurgence in the use of artificial neural networks to model the neural computations underlying cognition. Driven by advances in convolutional and recurrent architectures, along with computational capacity for deeper networks,

neural network models have begun to match human performance in a number of domains (Krizhevsky et al., 2012; Mnih et al., 2015). They have also been shown to provide state of the art models of representational geometry of cortical responses (Khaligh-Razavi and Kriegeskorte, 2014). It has been suggested, however, that current models may be limited by the fact that their architectures are set by researchers a priori, and that improvements could be made by including a gating mechanism for learning which modules or layers to route information through in different contexts (Hafner et al., 2017; Sabour et al., 2017; Hayworth and Marblestone, 2018). In fact some of these proposals have been explicitly inspired by ideas of thalamic information routing described above (Hafner et al., 2017; Hayworth and Marblestone, 2018). Given the current bidirectional flow of progress between artificial intelligence and cognitive neuroscience, theoretical insight from network neuroscience into questions of cognitive function could inspire further information routing architectures in neural networks, while the potential success of these dynamic architectures could inspire theoretical and empirical investigations into how the brain flexibly routes information.

Conclusion

While the emergence of network neuroscience holds immense potential for our understanding of the brain, there are serious challenges ahead in terms of incorporating network approaches into our theories of cognition. From methodological concerns about fitting network models, to conceptual ambiguities about network diagnostics applied to brain and cognitive data, and most importantly to the need for developing network mechanisms with testable predictions, there are a number of dimensions along which network

neuroscience must continue to progress in order to change the way we think about the mind.

References

Achard S, Bullmore E (2007) Efficiency and cost of economical brain functional networks. PLoS computational biology 3:e17.

Alexander GE, DeLong MR, Strick PL (1986) Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annual review of neuroscience 9:357-381.

Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD (2012) Tracking whole-brain connectivity dynamics in the resting state. Cerebral cortex:bhs352.

Andersson J, Smith S, Jenkinson M (2008) Fnirt-fmrib's non-linear image registration tool. Human Brain Mapping 2008.

Andrews-Hanna JR, Reidler JS, Sepulcre J, Poulin R, Buckner RL (2010) Functional-anatomic fractionation of the brain's default network. Neuron 65:550-562.

Bakkour A, Zylberberg A, Shadlen MN, Shohamy D (2018) Value-based decisions involve sequential sampling from memory. bioRxiv:269290.

Barrett LF, Satpute AB (2013) Large-scale brain networks in affective and social neuroscience: towards an integrative functional architecture of the brain. Current opinion in neurobiology 23:361-372.

Bartolomeo P, Thiebaut De Schotten M, Chica AB (2012) Brain networks of visuospatial attention and their disruption in visual neglect. Frontiers in human neuroscience 6:110.

Bartra O, McGuire JT, Kable JW (2013) The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. Neuroimage 76:412-427.

Barttfeld P, Uhrig L, Sitt JD, Sigman M, Jarraya B, Dehaene S (2015) Signature of consciousness in the dynamics of resting-state brain activity. Proceedings of the National Academy of Sciences 112:887-892.

Bassett DS, Bullmore E (2006) Small-world brain networks. The neuroscientist 12:512-523.

Bassett DS, Mattar MG (2017) A network neuroscience of human learning: potential to inform quantitative theories of brain and behavior. Trends in cognitive sciences 21:250-264.

Bassett DS, Sporns O (2017) Network neuroscience. Nature neuroscience 20:353.

Bassett DS, Yang M, Wymbs NF, Grafton ST (2015) Learning-induced autonomy of sensorimotor systems. Nature Neuroscience.

Bassett DS, Porter MA, Wymbs NF, Grafton ST, Carlson JM (2013a) Robust detection of dynamic community structure in networks. Chaos 23:013142.

Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST (2011) Dynamic reconfiguration of human brain networks during learning. Proceedings of the National Academy of Sciences 108:7641-7646.

Bassett DS, Wymbs NF, Rombach MP, Porter MA, Mucha PJ, Grafton ST (2013b) Task-based core-periphery organization of human brain dynamics. PLoS computational biology 9:e1003171.

Bates D, Mächler M, Bolker B, Walker S (2015) Fitting Linear Mixed-Effects Models Usinglme4. Journal of Statistical Software 67.

Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. Neuron 47:129-141.

Beckmann CF, DeLuca M, Devlin JT, Smith SM (2005) Investigations into resting-state connectivity using independent component analysis. Philosophical Transactions of the Royal Society of London B: Biological Sciences 360:1001-1013.

Benchenane K, Peyrache A, Khamassi M, Tierney PL, Gioanni Y, Battaglia FP, Wiener SI (2010) Coherent theta oscillations and reorganization of spike timing in the hippocampal-prefrontal network upon learning. Neuron 66:921-936.

Beninger RJ (1983) The role of dopamine in locomotor activity and learning. Brain Research Reviews 6:173-196.

Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society Series B (Methodological):289-300.

Bertolero MA, Yeo BT, D'Esposito M (2015) The modular and integrative functional architecture of the human brain. Proceedings of the National Academy of Sciences 112:E6798-E6807.

Betzel RF, Bassett DS (2017) Generative models for network neuroscience: prospects and promise. Journal of The Royal Society Interface 14:20170623.

Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar mri. Magnetic resonance in medicine 34:537-541.

Blackburn JR, Pfaus JG, Phillips AG (1992) Dopamine functions in appetitive and defensive behaviours. Progress in neurobiology 39:247-279.

Bogacz R, Gurney K (2007) The basal ganglia and cortex implement optimal decision making between alternative actions. Neural computation 19:442-477.

Bornstein AM, Daw ND (2013) Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. PLoS Computational Biology 9:e1003387.

Bornstein AM, Khaw MW, Shohamy D, Daw ND (2017) Reminders of past choices bias decisions for reward in humans. Nature Communications 8:15958.

Botvinick M, Barrett DG, Battaglia P, de Freitas N, Kumaran D, Leibo JZ, Lillicrap T, Modayil J, Mohamed S, Rabinowitz NC (2017) Building machines that learn and think for themselves. Behavioral and Brain Sciences 40.

Braun U, Schäfer A, Walter H, Erk S, Romanczuk-Seiferth N, Haddad L, Schweiger JI, Grimm O, Heinz A, Tost H (2015) Dynamic reconfiguration of frontal brain networks during executive cognition in humans. Proceedings of the National Academy of Sciences 112:11678-11683.

Braun U, Schäfer A, Bassett DS, Rausch F, Schweiger JI, Bilek E, Erk S, Romanczuk-Seiferth N, Grimm O, Geiger LS (2016) Dynamic brain network reconfiguration as a potential schizophrenia genetic risk mechanism modulated by NMDA receptor function. Proceedings of the National Academy of Sciences 113:12568-12573.

Bressler SL (1995) Large-scale cortical networks and cognition. Brain Research Reviews 20:288-304.

Bressler SL, Menon V (2010) Large-scale brain networks in cognition: emerging methods and principles. Trends in cognitive sciences 14:277-290.

Brogden W (1939) Sensory pre-conditioning. Journal of Experimental Psychology 25:323.

Bruce CJ, Goldberg ME (1985) Primate frontal eye fields. I. Single neurons discharging before saccades. Journal of neurophysiology 53:603-635.

Buckner RL, Andrews-Hanna JR, Schacter DL (2008) The brain's default network. Annals of the New York Academy of Sciences 1124:1-38.

Buerger A, Gross C, Rocha-Miranda C (1974) Effects of ventral putamen lesions on discrimination learning by monkeys. Journal of comparative and physiological psychology 86:440.

Bullmore E, Sporns O (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. Nature Reviews Neuroscience 10:186-198.

Bullmore ET, Bassett DS (2011) Brain graphs: graphical models of the human brain connectome. Annual review of clinical psychology 7:113-140.

Calhoun VD, Miller R, Pearlson G, Adalı T (2014) The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. Neuron 84:262-274.

Carpenter B, Gelman A, Hoffman M, Lee D, Goodrich B, Betancourt M, Brubaker MA, Guo J, Li P, Riddell A (2015) Stan: a probabilistic programming language. Journal of Statistical Software.

Chen JE, Glover GH (2015) BOLD fractional contribution to resting-state functional connectivity above 0.1 Hz. NeuroImage 107:207-218.

Chun MM, Phelps EA (1999) Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. Nature neuroscience 2:844.

Chung Y, Gelman A, Rabe-Hesketh S, Liu J, Dorie V (2015) Weakly informative prior for point estimation of covariance matrices in hierarchical models. Journal of Educational and Behavioral Statistics 40:136-157.

Cohen NJ, Ryan J, Hunt C, Romine L, Wszalek T, Nash C (1999) Hippocampal system and declarative (relational) memory: summarizing the data from functional neuroimaging studies. Hippocampus 9:83-98.

Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE (2014) Intrinsic and task-evoked network architectures of the human brain. Neuron 83:238-251.

Cools R, Lewis SJ, Clark L, Barker RA, Robbins TW (2007) L-DOPA disrupts activity in the nucleus accumbens during reversal learning in Parkinson's disease. Neuropsychopharmacology 32:180.

Crossley NA, Mechelli A, Vértes PE, Winton-Brown TT, Patel AX, Ginestet CE, McGuire P, Bullmore ET (2013) Cognitive relevance of the community structure of the human brain functional coactivation network. Proceedings of the National Academy of Sciences 110:11583-11588.

Cummins RC (1989) Meaning and mental representation.

Damoiseaux J, Rombouts S, Barkhof F, Scheltens P, Stam C, Smith SM, Beckmann C (2006) Consistent resting-state networks across healthy subjects. Proceedings of the national academy of sciences 103:13848-13853.

Danet L, Barbeau EJ, Eustache P, Planton M, Raposo N, Sibon I, Albucher J-F, Bonneville F, Peran P, Pariente J (2015) Thalamic amnesia after infarct The role of the mammillothalamic tract and mediodorsal nucleus. Neurology 85:2107-2115.

Davidow JY, Foerde K, Galván A, Shohamy D (2016) An Upside to Reward Sensitivity: The Hippocampus Supports Enhanced Reinforcement Learning in Adolescence. Neuron 92:93-99.

Daw ND (2011) Trial-by-trial data analysis using computational models. Decision making, affect, and learning: Attention and performance XXIII 23:3-38.

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

Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ (2006) Cortical substrates for exploratory decisions in humans. Nature 441:876-879.

Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ (2011) Model-based influences on humans' choices and striatal prediction errors. Neuron 69:1204-1215.

Deco G, Jirsa VK, McIntosh AR (2011) Emerging concepts for the dynamical organization of resting-state activity in the brain. Nature Reviews Neuroscience 12:43-56.

Deserno L, Huys QJ, Boehme R, Buchert R, Heinze H-J, Grace AA, Dolan RJ, Heinz A, Schlagenhauf F (2015) Ventral striatal dopamine reflects behavioral and neural signatures of model-based control during sequential decision making. Proceedings of the National Academy of Sciences 112:1595-1600.

Dews P, Morse W (1958) SOME OBSERVATIONS ON AN OPERANT IN HUMAN SUBJECTS AND ITS MODIFICATION BY DEXTRO AMPHETAMINE1. Journal of the Experimental Analysis of Behavior 1:359-364.

Ding L (2015) Distinct dynamics of ramping activity in the frontal cortex and caudate nucleus in monkeys. Journal of neurophysiology 114:1850-1861.

Doll BB, Simon DA, Daw ND (2012) The ubiquity of model-based reinforcement learning. Current opinion in neurobiology 22:1075-1081.

Doll BB, Shohamy D, Daw ND (2015) Multiple memory systems as substrates for multiple decision systems. Neurobiology of learning and memory 117:4-13.

Dosenbach NU, Fair DA, Cohen AL, Schlaggar BL, Petersen SE (2008) A dual-networks architecture of top-down control. Trends in cognitive sciences 12:99-105.

Doyère V, Burette F, Rédini-Del Negro C, Laroche S (1993) Long-term potentiation of hippocampal afferents and efferents to prefrontal cortex: implications for associative learning. Neuropsychologia 31:1031-1053.

Dupret D, O'Neill J, Csicsvari J (2013) Dynamic reconfiguration of hippocampal interneuron circuits during spatial learning. Neuron 78:166-180.

Dupret D, Revest J-M, Koehl M, Ichas F, De Giorgi F, Costet P, Abrous DN, Piazza PV (2008) Spatial relational memory requires hippocampal adult neurogenesis. PloS one 3:e1959.

Durante D, Mukherjee N, Steorts RC (2016) Bayesian Learning of Dynamic Multilayer Networks. arXiv preprint arXiv:160802209.

Dusek JA, Eichenbaum H (1997) The hippocampus and memory for orderly stimulus relations. Proceedings of the National Academy of Sciences 94:7109-7114.

Eichenbaum H (2000) A cortical–hippocampal system for declarative memory. Nature Reviews Neuroscience 1:41-50.

Eichenbaum H (2017) The role of the hippocampus in navigation is memory. Journal of neurophysiology 117:1785-1796.

Eichenbaum H, Otto T, Cohen NJ (1992) The hippocampus—what does it do? Behavioral and neural biology 57:2-36.

Ekman M, Derrfuss J, Tittgemeyer M, Fiebach CJ (2012) Predicting errors from reconfiguration patterns in human brain networks. Proceedings of the National Academy of Sciences 109:16714-16719.

Fair DA, Cohen AL, Power JD, Dosenbach NU, Church JA, Miezin FM, Schlaggar BL, Petersen SE (2009) Functional brain networks develop from a "local to distributed" organization. PLoS computational biology 5:e1000381.

Fan L, Li H, Zhuo J, Zhang Y, Wang J, Chen L, Yang Z, Chu C, Xie S, Laird AR (2016) The human brainnetome atlas: a new brain atlas based on connectional architecture. Cerebral cortex 26:3508-3526.

Fibiger H, Phillips A, Zis A (1974) Deficits in instrumental responding after 6-hydroxydopamine lesions of the nigro-neostriatal dopaminergic projection. Pharmacology Biochemistry and Behavior 2:87-96.

Fodor JA (1975) The language of thought: Harvard University Press.

Foerde K, Shohamy D (2011a) Feedback timing modulates brain systems for learning in humans. The Journal of Neuroscience 31:13157-13167.

Foerde K, Shohamy D (2011b) The role of the basal ganglia in learning and memory: insight from Parkinson's disease. Neurobiology of Learning and Memory 96:624-636.

Foerde K, Braun EK, Shohamy D (2012) A trade-off between feedback-based learning and episodic memory for feedback events: evidence from Parkinson's disease. Neurodegenerative Diseases 11:93-101.

Foerde K, Race E, Verfaellie M, Shohamy D (2013) A role for the medial temporal lobe in feedback-driven learning: evidence from amnesia. The Journal of Neuroscience 33:5698-5704.

Fornito A, Harrison BJ, Zalesky A, Simons JS (2012) Competitive and cooperative dynamics of large-scale brain functional networks supporting recollection. Proceedings of the National Academy of Sciences 109:12788-12793.

Fortunato S (2010) Community detection in graphs. Physics reports 486:75-174.

Frank MJ, Badre D (2011) Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: computational analysis. Cerebral cortex 22:509-526.

Frank MJ, Seeberger LC, O'reilly RC (2004) By carrot or by stick: cognitive reinforcement learning in parkinsonism. Science 306:1940-1943.

Gabard-Durnam LJ, Gee DG, Goff B, Flannery J, Telzer E, Humphreys KL, Lumian DS, Fareri DS, Caldera C, Tottenham N (2016) Stimulus-elicited connectivity influences resting-state connectivity years later in human development: a prospective study. Journal of Neuroscience 36:4771-4784.

Gallos LK, Makse HA, Sigman M (2012) A small world of weak ties provides optimal global integration of self-similar modules in functional brain networks. Proceedings of the National Academy of Sciences 109:2825-2830.

Garcia JO, Ashourvan A, Muldoon S, Vettel JM, Bassett DS (2017) Applications of community detection techniques to brain graphs: Algorithmic considerations and implications for neural function. bioRxiv:209429.

Gelman A (2006a) Multilevel (hierarchical) modeling: what it can and cannot do. Technometrics 48:432-435.

Gelman A (2006b) Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). Bayesian analysis 1:515-534.

Gerraty RT, Davidow JY, Wimmer GE, Kahn I, Shohamy D (2014) Transfer of learning relates to intrinsic connectivity between hippocampus, ventromedial prefrontal cortex, and large-scale networks. The Journal of Neuroscience 34:11297-11303.

Gerraty RT, Davidow JY, Foerde K, Galvan A, Bassett DS, Shohamy D (2018) Dynamic flexibility in striatal-cortical circuits supports reinforcement learning. Journal of Neuroscience 38:2442-2453.

Gershman SJ (2016) Empirical priors for reinforcement learning models. Journal of Mathematical Psychology 71:1-6.

Gershman SJ, Daw ND (2017) Reinforcement learning and episodic memory in humans and animals: an integrative framework. Annual review of psychology 68:101-128.

Gershman SJ, Schoenbaum G (2017) Rethinking dopamine prediction errors. bioRxiv:239731.

Gießing C, Thiel CM, Alexander-Bloch AF, Patel AX, Bullmore ET (2013) Human brain functional network changes associated with enhanced and impaired attentional task performance. Journal of Neuroscience 33:5903-5914.

Giovanello KS, Schnyer D, Verfaellie M (2009) Distinct hippocampal regions make unique contributions to relational memory. Hippocampus 19:111-117.

Giusti C, Ghrist R, Bassett DS (2016) Two's company, three (or more) is a simplex. Journal of computational neuroscience 41:1-14.

Glimcher PW (2011) Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. Proceedings of the National Academy of Sciences 108:15647-15654.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annual review of neuroscience 30.

Good BH, de Montjoye Y-A, Clauset A (2010) Performance of modularity maximization in practical contexts. Physical Review E 81:046106.

Graham KS, Barense MD, Lee AC (2010) Going beyond LTM in the MTL: a synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. Neuropsychologia 48:831-853.

Greene AJ, Gross WL, Elsinger CL, Rao SM (2006) An FMRI analysis of the human hippocampus: inference, context, and task awareness. Journal of cognitive neuroscience 18:1156-1173.

Greicius M (2008) Resting-state functional connectivity in neuropsychiatric disorders. Current opinion in neurology 21:424-430.

Greve DN, Fischl B (2009) Accurate and robust brain image alignment using boundary-based registration. Neuroimage 48:63-72.

Grogan JP, Tsivos D, Smith L, Knight BE, Bogacz R, Whone A, Coulthard EJ (2017) Effects of dopamine on reinforcement learning and consolidation in Parkinson's disease. eLife 6.

Haber SN (2003) The primate basal ganglia: parallel and integrative networks. Journal of chemical neuroanatomy 26:317-330.

Haber SN, Knutson B (2010) The reward circuit: linking primate anatomy and human imaging. Neuropsychopharmacology 35:4-26.

Haber SN, Kim K-S, Mailly P, Calzavara R (2006) Reward-related cortical inputs define a large striatal region in primates that interface with associative cortical connections, providing a substrate for incentive-based learning. The Journal of Neuroscience 26:8368-8376.

Hafner D, Irpan A, Davidson J, Heess N (2017) Learning Hierarchical Information Flow with Recurrent Neural Modules. In: Advances in Neural Information Processing Systems, pp 6727-6736.

Hamid AA, Pettibone JR, Mabrouk OS, Hetrick VL, Schmidt R, Vander Weele CM, Kennedy RT, Aragona BJ, Berke JD (2016) Mesolimbic dopamine signals the value of work. Nature neuroscience 19:117-126.

Hampton AN, Bossaerts P, O'doherty JP (2006) The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. The Journal of Neuroscience 26:8360-8367.

Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nature neuroscience 9:682.

Hannula DE, Ryan JD, Tranel D, Cohen NJ (2007) Rapid onset relational memory effects are evident in eye movement behavior, but not in hippocampal amnesia. Journal of cognitive neuroscience 19:1690-1705.

Harmelech T, Preminger S, Wertman E, Malach R (2013) The day-after effect: long term, Hebbian-like restructuring of resting-state fMRI patterns induced by a single epoch of cortical activation. The Journal of Neuroscience 33:9488-9497.

Hayworth KJ, Marblestone AH (2018) How thalamic relays might orchestrate supervised deep training and symbolic computation in the brain. bioRxiv:304980.

He BJ (2013) Spontaneous and task-evoked brain activity negatively interact. The Journal of Neuroscience 33:4672-4682.

Heitger MH, Ronsse R, Dhollander T, Dupont P, Caeyenberghs K, Swinnen SP (2012) Motor learning-induced changes in functional brain connectivity as revealed by means of graph-theoretical network analysis. Neuroimage 61:633-650.

Hikosaka O, Kim HF, Yasuda M, Yamamoto S (2014) Basal ganglia circuits for reward value-guided behavior. Annual review of neuroscience 37:289.

Hipp JF, Hawellek DJ, Corbetta M, Siegel M, Engel AK (2012) Large-scale cortical correlation structure of spontaneous oscillatory activity. Nature neuroscience 15:884-890.

Holme P, Saramäki J (2011) Temporal Networks. Physics Reports 519:97-125.

Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, Della Penna S, Duyn JH, Glover GH, Gonzalez-Castillo J (2013) Dynamic functional connectivity: promise, issues, and interpretations. Neuroimage 80:360-378.

Ishii S, Yoshida W, Yoshimoto J (2002) Control of exploitation–exploration meta-parameter in reinforcement learning. Neural networks 15:665-687.

Ivleva EI, Shohamy D, Mihalakos P, Morris DW, Carmody T, Tamminga CA (2012) Memory generalization is selectively altered in the psychosis dimension. Schizophrenia research 138:74-80.

Jafri MJ, Pearlson GD, Stevens M, Calhoun VD (2008) A method for functional network connectivity among spatially independent resting-state components in schizophrenia. Neuroimage 39:1666-1681.

Janowsky JS, Shimamura AP, Kritchevsky M, Squire LR (1989) Cognitive impairment following frontal lobe damage and its relevance to human amnesia. Behavioral neuroscience 103:548.

Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage 17:825-841.

Ji D, Wilson MA (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. Nature neuroscience 10:100-107.

Jones JL, Esber GR, McDannald MA, Gruber AJ, Hernandez A, Mirenzi A, Schoenbaum G (2012) Orbitofrontal cortex supports behavior and learning using inferred but not cached values. Science 338:953-956.

Jones MW, Wilson MA (2005) Theta rhythms coordinate hippocampal–prefrontal interactions in a spatial memory task. PLoS Biol 3:e402.

Jutla IS, Jeub LG, Mucha PJ (2011) A generalized Louvain method for community detection implemented in MATLAB. URL http://netwiki amath unc edu/GenLouvain.

Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. Nature neuroscience 10:1625-1633.

Kahn I, Shohamy D (2013) Intrinsic connectivity between the hippocampus, nucleus accumbens, and ventral tegmental area in humans. Hippocampus 23:187-192.

Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. Journal of neuroscience 17:4302-4311.

Kao K, Powell D (1986) Lesions of substantia nigra retard Pavlovian somatomotor learning but do not affect autonomic conditioning. Neuroscience letters 64:1-6.

Kaplan R, Adhikari MH, Hindriks R, Mantini D, Murayama Y, Logothetis NK, Deco G (2016) Hippocampal sharp-wave ripples influence selective activation of the default mode network. Current Biology 26:686-691.

Kastner S, Saalman YB (2012) Role of the pulvinar in regulating information transmission between cortical areas. Journal of Vision 12:1374-1374.

Kehoe EJ, Gormezano I (1980) Configuration and combination laws in conditioning with compound stimuli. Psychological Bulletin 87:351.

Kemp JM, Powell T (1971) The connexions of the striatum and globus pallidus: synthesis and speculation. Philosophical Transactions of the Royal Society of London B: Biological Sciences 262:441-457.

Khaligh-Razavi S-M, Kriegeskorte N (2014) Deep supervised, but not unsupervised, models may explain IT cortical representation. PLoS computational biology 10:e1003915.

Khambhati AN, Sizemore AE, Betzel RF, Bassett DS (2017) Modeling and interpreting mesoscale network dynamics. Neuroimage.

Khambhati AN, Davis KA, Lucas TH, Litt B, Bassett DS (2016) Virtual cortical resection reveals push-pull network control preceding seizure evolution. Neuron 91:1170-1182.

Kirkby RJ, Kimble DP (1968) Avoidance and escape behavior following striatal lesions in the rat. Experimental neurology 20:215-227.

Kishida KT, Saez I, Lohrenz T, Witcher MR, Laxton AW, Tatter SB, White JP, Ellis TL, Phillips PE, Montague PR (2016) Subsecond dopamine fluctuations in human striatum encode superposed error signals about actual and counterfactual reward. Proceedings of the National Academy of Sciences 113:200-205.

Kivelä M, Arenas A, Barthelemy M, Gleeson JP, Moreno Y, Porter MA (2014) Multilayer networks. Journal of Complex Networks 2:203 - 271.

Knowlton BJ, Mangels JA, Squire LR (1996) A neostriatal habit learning system in humans. Science 273:1399-1402.

Kopell NJ, Gritton HJ, Whittington MA, Kramer MA (2014) Beyond the connectome: the dynome. Neuron 83:1319-1328.

Krienen FM, Yeo BT, Buckner RL (2014) Reconfigurable task-dependent functional coupling modes cluster around a core functional architecture. Phil Trans R Soc B 369:20130526.

Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems, pp 1097-1105.

Kumaran D, McClelland JL (2012) Generalization through the recurrent interaction of episodic memories: a model of the hippocampal system. Psychological review 119:573.

Kumaran D, Summerfield JJ, Hassabis D, Maguire EA (2009) Tracking the emergence of conceptual knowledge during human decision making. Neuron 63:889-901.

Laird AR, Fox PM, Eickhoff SB, Turner JA, Ray KL, McKay DR, Glahn DC, Beckmann CF, Smith SM, Fox PT (2011) Behavioral interpretations of intrinsic connectivity networks. Journal of cognitive neuroscience 23:4022-4037.

Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ (2017) Building machines that learn and think like people. Behavioral and Brain Sciences 40.

Langer N, Pedroni A, Gianotti LR, Hänggi J, Knoch D, Jäncke L (2012) Functional brain network efficiency predicts intelligence. Human brain mapping 33:1393-1406.

Lashley KS (1930) Basic neural mechanisms in behavior. Psychological review 37:1.

Lashley KS (1950) In search of the engram.

Lee U, Oh G, Kim S, Noh G, Choi B, Mashour GA (2010) Brain networks maintain a scale-free organization across consciousness, anesthesia, and recoveryevidence for adaptive reconfiguration. Anesthesiology: The Journal of the American Society of Anesthesiologists 113:1081-1091.

Lewandowski D, Kurowicka D, Joe H (2009) Generating random correlation matrices based on vines and extended onion method. Journal of multivariate analysis 100:1989-2001.

Liu X, Duyn JH (2013) Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proceedings of the National Academy of Sciences 110:4392-4397.

Logothetis NK, Eschenko O, Murayama Y, Augath M, Steudel T, Evrard H, Besserve M, Oeltermann A (2012) Hippocampal-cortical interaction during periods of subcortical silence. Nature 491:547-553.

Ma S, Calhoun VD, Eichele T, Du W, Adalı T (2012) Modulations of functional connectivity in the healthy and schizophrenia groups during task and rest. Neuroimage 62:1694-1704.

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

Mason S, Iversen S (1974) Learning impairment in rats after 6-hydroxydopamine-induced depletion of brain catecholamines. Nature 248:697.

Mattar MG, Thompson-Schill SL, Bassett DS (2016) The network architecture of value learning. Network Neuroscience:1-27.

Mayes A, Montaldi D, Migo E (2007) Associative memory and the medial temporal lobes. Trends in cognitive sciences 11:126-135.

Mazzoni P, Bracewell RM, Barash S, Andersen RA (1996) Spatially tuned auditory responses in area LIP of macaques performing delayed memory saccades to acoustic targets. Journal of neurophysiology 75:1233-1241.

McDannald MA, Lucantonio F, Burke KA, Niv Y, Schoenbaum G (2011) Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. The Journal of Neuroscience 31:2700-2705.

McIntosh A, Gonzalez-Lima F (1998) Large-scale functional connectivity in associative learning: interrelations of the rat auditory, visual, and limbic systems. Journal of Neurophysiology 80:3148-3162.

McIntosh AR (2000) Towards a network theory of cognition. Neural Networks 13:861-870.

Medaglia JD, Lynall M-E, Bassett DS (2015) Cognitive Network Neuroscience. Journal of cognitive neuroscience.

Meng X-L, Rosenthal R, Rubin DB (1992) Comparing correlated correlation coefficients. Psychological bulletin 111:172.

Menon V (2011) Large-scale brain networks and psychopathology: a unifying triple network model. Trends in cognitive sciences 15:483-506.

Meunier D, Lambiotte R, Bullmore ET (2010) Modular and hierarchically modular organization of brain networks. Frontiers in neuroscience 4:200.

Mišić B, Sporns O (2016) From regions to connections and networks: new bridges between brain and behavior. Current Opinion in Neurobiology 40:1-7.

Mitchell AS, Sherman SM, Sommer MA, Mair RG, Vertes RP, Chudasama Y (2014) Advances in understanding mechanisms of thalamic relays in cognition and behavior. Journal of Neuroscience 34:15340-15346.

Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G (2015) Human-level control through deep reinforcement learning. Nature 518:529.

Mohler CW, Goldberg ME, Wurtz RH (1973) Visual receptive fields of frontal eye field neurons. Brain research 61:385-389.

Moussa MN, Vechlekar CD, Burdette JH, Steen MR, Hugenschmidt CE, Laurienti PJ (2011) Changes in cognitive state alter human functional brain networks. Frontiers in human neuroscience 5:83.

Mucha PJ, Richardson T, Macon K, Porter MA, Onnela J-P (2010) Community structure in time-dependent, multiscale, and multiplex networks. Science 328:876-878.

Muldoon SF, Bassett DS (2014) Why network neuroscience? Compelling evidence and current frontiers. Comment on``Understanding brain networks and brain organization"by Luiz Pessoa. Physics of life reviews 11:455-457.

Myers CE, Shohamy D, Gluck MA, Grossman S, Onlaor S, Kapur N (2003) Dissociating medial temporal and basal ganglia memory systems with a latent learning task. Neuropsychologia 41:1919-1928.

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

Nelson CA (1995) The ontogeny of human memory: A cognitive neuroscience perspective. Developmental psychology 31:723.

Newman ME (2004) Fast algorithm for detecting community structure in networks. Physical review E 69:066133.

Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. Nature 341:52.

Nyberg L, Persson J, Habib R, Tulving E, McIntosh AR, Cabeza R, Houle S (2000) Large scale neurocognitive networks underlying episodic memory. Journal of Cognitive Neuroscience 12:163-173.

O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. Neuron 38:329-337.

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'KeefeIII J (1982) Place navigation impaired in rats with hippocampal lesions. Nature 297:24.

Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. Nature 441:223-226.

Pagnoni G, Zink CF, Montague PR, Berns GS (2002) Activity in human ventral striatum locked to errors of reward prediction. Nature neuroscience 5:97.

Palla K, Caron F, Teh YW (2016) Bayesian nonparametrics for Sparse Dynamic Networks. arXiv preprint arXiv:160701624.

Paller KA, Wagner AD (2002) Observing the transformation of experience into memory. Trends in cognitive sciences 6:93-102.

Park H-J, Friston K (2013) Structural and functional brain networks: from connections to cognition. Science 342:1238411.

Pauli WM, O'Reilly RC, Yarkoni T, Wager TD (2016) Regional specialization within the human striatum for diverse psychological functions. Proceedings of the National Academy of Sciences 113:1907-1912.

Pavlov IP (1927) Conditioned reflexes. An Investigation of the physiological activity of the cerebral cortex.

Petersen SE, Sporns O (2015) Brain networks and cognitive architectures. Neuron 88:207-219.

Pizzagalli DA, Evins AE, Schetter EC, Frank MJ, Pajtas PE, Santesso DL, Culhane M (2008) Single dose of a dopamine agonist impairs reinforcement learning in humans: behavioral evidence from a laboratory-based measure of reward responsiveness. Psychopharmacology 196:221-232.

Poldrack RA (2011) Inferring mental states from neuroimaging data: from reverse inference to large-scale decoding. Neuron 72:692-697.

Poldrack RA, Packard MG (2003) Competition among multiple memory systems: converging evidence from animal and human brain studies. Neuropsychologia 41:245-251.

Poldrack RA, Clark J, Pare-Blagoev E, Shohamy D, Moyano JC, Myers C, Gluck M (2001) Interactive memory systems in the human brain. Nature 414:546-550.

Port RL, Beggs AL, Patterson MM (1987) Hippocampal substrate of sensory associations. Physiology & behavior 39:643-647.

Porter MA, Onnela J-P, Mucha PJ (2009) Communities in networks. Notices of the AMS 56:1082-1097.

Posner MI, DiGirolamo GJ (2000) Cognitive neuroscience: origins and promise. Psychological bulletin 126:873.

Posner MI, Petersen SE, Fox PT, Raichle ME (1988) Localization of cognitive operations in the human brain. Science 240:1627-1631.

Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE (2012) Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage 59:2142-2154.

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL (2011) Functional network organization of the human brain. Neuron 72:665-678.

Pylyshyn ZW (1980) Computation and cognition: Issues in the foundations of cognitive science. Behavioral and Brain Sciences 3:111-132.

Raichle ME (2015) The brain's default mode network. Annual review of neuroscience 38:433-447.

Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL (2001) A default mode of brain function. Proceedings of the National Academy of Sciences 98:676-682.

Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. Classical conditioning II: Current research and theory 2:64-99.

Robbins TW, Brown VJ (1990) The role of the striatum in the mental chronometry of action: a theoretical review. Reviews in the Neurosciences 2:181-214.

Robbins TW, Everitt BJ (1996) Neurobehavioural mechanisms of reward and motivation. Current opinion in neurobiology 6:228-236.

Roth MM, Dahmen JC, Muir DR, Imhof F, Martini FJ, Hofer SB (2016) Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. Nature neuroscience 19:299.

Roy M, Shohamy D, Wager TD (2012) Ventromedial prefrontal-subcortical systems and the generation of affective meaning. Trends in cognitive sciences 16:147-156.

Rutledge RB, Dean M, Caplin A, Glimcher PW (2010) Testing the reward prediction error hypothesis with an axiomatic model. The Journal of Neuroscience 30:13525-13536.

Sabour S, Frosst N, Hinton GE (2017) Dynamic routing between capsules. In: Advances in Neural Information Processing Systems, pp 3859-3869.

Sadacca BF, Jones JL, Schoenbaum G (2016) Midbrain dopamine neurons compute inferred and cached value prediction errors in a common framework. Elife 5.

Satterthwaite TD, Wolf DH, Loughead J, Ruparel K, Elliott MA, Hakonarson H, Gur RC, Gur RE (2012a) Impact of in-scanner head motion on multiple measures of

functional connectivity: relevance for studies of neurodevelopment in youth. Neuroimage.

Satterthwaite TD, Elliott MA, Gerraty RT, Ruparel K, Loughead J, Calkins ME, Eickhoff SB, Hakonarson H, Gur RC, Gur RE (2012b) An Improved Framework for Confound Regression and Filtering for Control of Motion Artifact in the Preprocessing of Resting-State Functional Connectivity Data. NeuroImage.

Schacter DL, Tulving E (1994) Memory systems 1994.

Schapiro AC, Kustner LV, Turk-Browne NB (2012) Shaping of object representations in the human medial temporal lobe based on temporal regularities. Current Biology 22:1622-1627.

Schoenbaum G, Takahashi Y, Liu TL, McDannald MA (2011) Does the orbitofrontal cortex signal value? Annals of the New York Academy of Sciences 1239:87-99.

Schonberg T, O'Doherty JP, Joel D, Inzelberg R, Segev Y, Daw ND (2010) Selective impairment of prediction error signaling in human dorsolateral but not ventral striatum in Parkinson's disease patients: evidence from a model-based fMRI study. Neuroimage 49:772-781.

Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593-1599.

Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. Journal of neurology, neurosurgery, and psychiatry 20:11.

Shadlen MN, Shohamy D (2016) Decision making and sequential sampling from memory. Neuron 90:927-939.

Sharp ME, Foerde K, Daw ND, Shohamy D (2015) Dopamine selectively remediates ‘model-based’reward learning: a computational approach. Brain:awv347.

Sharpe MJ, Chang CY, Liu MA, Batchelor HM, Mueller LE, Jones JL, Niv Y, Schoenbaum G (2017) Dopamine transients are sufficient and necessary for acquisition of model-based associations. Nature neuroscience 20:735.

Sheline YI, Price JL, Yan Z, Mintun MA (2010) Resting-state functional MRI in depression unmasks increased connectivity between networks via the dorsal nexus. Proceedings of the National Academy of Sciences 107:11020-11025.

Shepard RN (1967) Recognition memory for words, sentences, and pictures. Journal of verbal Learning and verbal Behavior 6:156-163.

Shine JM, Koyejo O, Bell PT, Gorgolewski KJ, Gilat M, Poldrack RA (2015) Estimation of dynamic functional connectivity using Multiplication of Temporal Derivatives. NeuroImage 122:399-407.

Shine JM, Bissett PG, Bell PT, Koyejo O, Balsters JH, Gorgolewski KJ, Moodie CA, Poldrack RA (2016) The Dynamics of Functional Brain Networks: Integrated Network States during Cognitive Task Performance. Neuron 92:544-554.

Shiner T, Seymour B, Wunderlich K, Hill C, Bhatia KP, Dayan P, Dolan RJ (2012) Dopamine and performance in a reinforcement learning task: evidence from Parkinson's disease. Brain:aws083.

Shohamy D, Wagner AD (2008) Integrating memories in the human brain: hippocampal-midbrain encoding of overlapping events. Neuron 60:378-389.

Shohamy D, Adcock RA (2010) Dopamine and adaptive memory. Trends in cognitive sciences 14:464-472.

Shohamy D, Turk-Browne NB (2013) Mechanisms for widespread hippocampal involvement in cognition. Journal of Experimental Psychology: General 142:1159.

Shohamy D, Daw ND (2015) Integrating memories to guide decisions. Current Opinion in Behavioral Sciences 5:85-90.

Shohamy D, Myers C, Onlaor S, Gluck M (2004a) Role of the basal ganglia in category learning: how do patients with Parkinson's disease learn? Behavioral neuroscience 118:676.

Shohamy D, Myers C, Grossman S, Sage J, Gluck M, Poldrack R (2004b) Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. Brain 127:851-859.

Siebenhühner F, Weiss SA, Coppola R, Weinberger DR, Bassett DS (2013) Intra-and inter-frequency brain network structure in health and schizophrenia. PloS one 8:e72351.

Siegel M, Buschman TJ, Miller EK (2015) Cortical information flow during flexible sensorimotor decisions. Science 348:1352-1355.

Sizemore AE, Bassett DS (2017) Dynamic graph metrics: Tutorial, toolbox, and tale. NeuroImage.

Skinner BF (1948) 'Superstition'in the pigeon. Journal of experimental psychology 38:168.

Smith SM (2002) Fast robust automated brain extraction. Human brain mapping 17:143-155.

Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE (2004) Advances in functional

and structural MR image analysis and implementation as FSL. Neuroimage 23:S208-S219.

Smith SM, Miller KL, Moeller S, Xu J, Auerbach EJ, Woolrich MW, Beckmann CF, Jenkinson M, Andersson J, Glasser MF (2012) Temporally-independent functional modes of spontaneous brain activity. Proceedings of the National Academy of Sciences 109:3131-3136.

Smolensky P (1990) On the proper treatment of connectionism. In: Philosophy, Mind, and Cognitive Inquiry, pp 145-206: Springer.

Solomon PR, Vander Schaaf ER, Thompson RF, Weisz DJ (1986) Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. Behavioral neuroscience 100:729.

Sporns O (2003) Graph theory methods for the analysis of neural connectivity patterns. In: Neuroscience databases, pp 171-185: Springer.

Sporns O (2010) Networks of the Brain: MIT press.

Sporns O (2011) The human connectome: a complex network. Annals of the New York Academy of Sciences 1224:109-125.

Sporns O (2013) Network attributes for segregation and integration in the human brain. Current opinion in neurobiology 23:162-171.

Squire LR (2009) Memory and brain systems: 1969–2009. The Journal of Neuroscience 29:12711-12716.

Squire LR, Wixted JT (2011) The cognitive neuroscience of human memory since HM. Annual review of neuroscience 34:259.

Sreenivas S, Boehm S, Linden DEJ (2012) Emotional faces and the default mode network. Neuroscience letters 506:229-234.

Sripada RK, Wang X, Sripada CS, Liberzon I (2012) Altered resting-state amygdala functional connectivity in men with posttraumatic stress disorder. Journal of psychiatry & neuroscience: JPN 37:241.

Staresina BP, Davachi L (2009) Mind the gap: binding experiences across space and time in the human hippocampus. Neuron 63:267-276.

Steinberg EE, Keiflin R, Boivin JR, Witten IB, Deisseroth K, Janak PH (2013) A causal link between prediction errors, dopamine neurons and learning. Nature neuroscience 16:966.

Sun FT, Miller LM, D'esposito M (2004) Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. Neuroimage 21:647-658.

Sutton RS, Barto AG (1998) Reinforcement learning: An introduction: MIT press.

Takita M, Izaki Y, Jay TM, Kaneko H, Suzuki SS (1999) Induction of stable long-term depression in vivo in the hippocampal–prefrontal cortex pathway. European Journal of Neuroscience 11:4145-4148.

Tambini A, Ketz N, Davachi L (2010) Enhanced brain correlations during rest are related to memory for recent experiences. Neuron 65:280-290.

Telesford QK, Lynall M-E, Vettel J, Miller MB, Grafton ST, Bassett DS (2016) Detection of functional brain network reconfiguration during task-driven cognitive states. Neuroimage 142:198-210.

Thorndike E (1898) Some experiments on animal intelligence. Science:818-824.

Timmer MH, Sescousse G, van der Schaaf ME, Esselink RA, Cools R (2017) Reward learning deficits in Parkinson's disease depend on depression. Psychological medicine 47:2302-2311.

Tolman EC (1948) Cognitive maps in rats and men. Psychological review 55:189.

Tulving E, Markowitsch HJ (1997) Memory beyond the hippocampus. Current opinion in neurobiology 7:209-216.

Valenstein E, Bowers D, Verfaellie M, Heilman KM, Day A, Watson R (1987) Retrosplenial amnesia. Brain 110:1631-1646.

van Kesteren MT, Fernández G, Norris DG, Hermans EJ (2010) Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. Proceedings of the National Academy of Sciences 107:7550-7555.

van Kesteren MT, Ruiter DJ, Fernández G, Henson RN (2012) How schema and novelty augment memory formation. Trends in neurosciences 35:211-219.

Vann SD, Aggleton JP (2002) Extensive cytotoxic lesions of the rat retrosplenial cortex reveal consistent deficits on tasks that tax allocentric spatial memory. Behavioral neuroscience 116:85.

Vehtari A, Gelman A, Gabry J (2017) Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. Statistics and Computing 27:1413-1432.

Vickery TJ, Chun MM, Lee D (2011) Ubiquity and specificity of reinforcement signals throughout the human brain. Neuron 72:166-177.

Vincent JL, Snyder AZ, Fox MD, Shannon BJ, Andrews JR, Raichle ME, Buckner RL (2006) Coherent spontaneous activity identifies a hippocampal-parietal memory network. Journal of neurophysiology 96:3517-3531.

Vo K, Rutledge RB, Chatterjee A, Kable JW (2014) Dorsal striatum is necessary for stimulus-value but not action-value learning in humans. Brain 137:3129-3135.

Wang J-H, Zuo X-N, Gohel S, Milham MP, Biswal BB, He Y (2011) Graph theoretical analysis of functional brain networks: test-retest evaluation on short-and long-term resting-state functional MRI data. PloS one 6:e21976.

Wang L, Li Y, Metzak P, He Y, Woodward TS (2010) Age-related changes in topological patterns of large-scale brain functional networks during memory encoding and recognition. Neuroimage 50:862-872.

Wang L, Zhu C, He Y, Zang Y, Cao Q, Zhang H, Zhong Q, Wang Y (2009) Altered small-world brain functional networks in children with attention-deficit/hyperactivity disorder. Human brain mapping 30:638-649.

Watanabe S (2013) A widely applicable Bayesian information criterion. Journal of Machine Learning Research 14:867-897.

Wesierska M, Dockery C, Fenton AA (2005) Beyond memory, navigation, and inhibition: behavioral evidence for hippocampus-dependent cognitive coordination in the rat. Journal of Neuroscience 25:2413-2419.

Westerink B, Kwint H (1996) The pharmacology of mesolimbic dopamine neurons: a dual-probe microdialysis study in the ventral tegmental area and nucleus accumbens of the rat brain. Journal of Neuroscience 16:2605-2611.

White NM (1989a) Reward or reinforcement: what's the difference? Neuroscience & Biobehavioral Reviews 13:181-186.

White NM (1989b) A functional hypothesis concerning the striatal matrix and patches: Mediation of S□ R memory and reward. Life sciences 45:1943-1957.

White NM, McDonald RJ (2002) Multiple parallel memory systems in the brain of the rat. Neurobiology of learning and memory 77:125-184.

Wimmer G, Shohamy D (2011) The striatum and beyond: hippocampal contributions to decision making. Attention & Performance XXII:281-310.

Wimmer GE, Shohamy D (2012) Preference by Association: How Memory Mechanisms in the Hippocampus Bias Decisions. Science 338:270-273.

Wimmer GE, Daw ND, Shohamy D (2012) Generalization of value in reinforcement learning by humans. European Journal of Neuroscience 35:1092-1104.

Woolrich MW, Ripley BD, Brady M, Smith SM (2001) Temporal autocorrelation in univariate linear modeling of FMRI data. Neuroimage 14:1370-1386.

Xu L, Zhang H, Hui M, Long Z, Jin Z, Liu Y, Yao L (2014) Motor execution and motor imagery: a comparison of functional connectivity patterns based on graph theory. Neuroscience 261:184-194.

Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. Nature methods 8:665.

Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of neurophysiology 106:1125-1165.

Yin HH, Knowlton BJ (2006) The role of the basal ganglia in habit formation. Nature Reviews Neuroscience 7:464.

Yin HH, Knowlton BJ, Balleine BW (2006) Inactivation of dorsolateral striatum enhances sensitivity to changes in the action–outcome contingency in instrumental conditioning. Behavioural brain research 166:189-196.

Yuste R (2015) From the neuron doctrine to neural networks. Nature reviews neuroscience 16:487.

Zeithamova D, Dominick AL, Preston AR (2012) Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. Neuron 75:168-179.

Zhang Z, Telesford QK, Giusti C, Lim KO, Bassett DS (2016) Choosing wavelet methods, filters, and lengths for functional brain network construction. PloS one 11:e0157243.

Zola-Morgan S (1995) Localization of brain function: The legacy of Franz Joseph Gall (1758-1828). Annual Review of Neuroscience 18:359-383.

Zuo X-N, Kelly C, Adelstein JS, Klein DF, Castellanos FX, Milham MP (2010) Reliable intrinsic connectivity networks: test–retest evaluation using ICA and dual regression approach. Neuroimage 49:2163-2177.