Title
Neural Circuit for Locomotor Control, Brain State Regulation, and Decision-Making

Author
Lee, Andrew Moses

Publication Date
2013

Peer reviewed|Thesis/dissertation

Neural Circuits for Locomotor Control, Brain State Modulation, and Decision-making

by

Andrew Moses Lee

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Neuroscience

in the

GRADUATE DIVISION

Acknowledgements

I have had the pleasure of knowing a terrific collection of colleagues, classmates, and friends throughout graduate school. Without these people, I do not know if I would have survived graduate school, and it certainly would have been much less enjoyable.

I would, of course, like to thank my graduate school advisors for which there are now many. I began graduate school by working with Antonello Bonci. Anto was the first PI, who genuinely believed in my capabilities as a scientist, and it was his vote of confidence, which inevitably has carried me through graduate school. My entire thesis surrounds an observation that I made as a rotation student in his lab now four years ago. This observation has since become the basis of an elaborate program to understand the neural circuits underlying locomotor control, behavioral state modulation, and decision-making that now spans the work of at least three labs and continues to grow. Anto helped me see science in a completely different way. This valuable perspective will undoubtedly serve me well in the future.

After Anto left, I transitioned into the lab of Linda Wilbrecht. Linda has taught me how to think broadly about ideas in science. She has always been kind and available as a mentor. While my projects are somewhat orthogonal to the focus

of her lab, I am grateful for the fact that she let me have my small corner in the lab to pursue my ideas. She has been extremely generous in sharing her professional insights and helping me shape my career goals. I look forward to working with Linda in the future, and seeing how her lab changes and thrives at the University of California, Berkeley.

I also would like to thank Anatol Kreitzer, who I will be working with in the immediate future. Obviously prior work from the Kreitzer lab influenced my own experiments in the course of graduate school, and I am happy to be collaborating with members of his lab to better understand the neural basis of locomotor control. While I have only been working with Anatol briefly, I anticipate that I will gain a great number of insights into science in the upcoming months, and I appreciate all the members of his lab, who have in our brief time knowing each other adopted me into their family.

Perhaps, the people that I owe the most to in graduate school are the post-doc mentors, who I happened across in graduate school. All of them are not only my friends, but also some of the greatest sources of inspiration to me. Many of them are now faculty members in their own right, and I am sure they will serve as a guiding light for their own students.

The first person that I ran into as a rotation student in the Bonci lab was Dr. Garret Stuber, who trained me in the optogenetic methodologies and techniques,

which form the basis of almost every experiment within this thesis. Garret is a fantastic scientist, but has always been generous with his time and mentorship. His insights helped me formulate the questions that would serve to be the basis of my decision-making work. While he has long since left UCSF to begin his own successful career, his recent scientific work probing affective and motivational systems in little studied frontiers of the brain is an inspiration for me to try something different.

I also would like to thank Dr. Cris Niell, who not only trained me to be an in vivo electrophysiologist, but also a good friend. It is easy to underestimate everything that Cris taught me during the years that we have known each other. Cris gave me a haven in Eugene, Oregon during one of the most difficult periods of my graduate career. When I close my eyes and envision how I want to be as a young faculty member starting up his lab, I think of Cris.

Lung-Hao has always been a technical force of ingenuity, which is in contrast to his gentle and contemplative demeanor. Without Lung Hao, I have no idea how I would have survived through graduate school, let alone what I would have learned. Our conversations regarding the basis of decision-making were probably some of the most stimulating and insightful discussions that I have ever had in my life. They really challenged and changed the way that I thought of the brain.

I would also like to thank all of the former members of the Bonci lab, Wilbrecht lab, and the current members of the Kreitzer lab. Labs often become more than places where people work, but also families given the amount of time that people spend together. In particular, I would like to single out Michael Chiang, Nora Benavidez, and Claudia Wu as undergraduate helpers and research assistants, who have been instrumental to many of my projects.

I also would like to thank my thesis committee members…the majority of whom are also my graduate school mentors. Given that I have already singled out Anto, Linda, and Anatol, I would also like to thank Dr. Loren Frank and Dr. Michael Stryker. While we have only met personally a few times, I have always been struck by Loren's candor, his enthusiasm for teaching, and intellectual rigor. However, the quality that stands out the most in my mind when I think of Loren is his integrity, a quality that I feel is often under-valued in science given that it serves as the foundation of our profession. I also would like to thank Dr. Michael Stryker who is the chair of my committee. The scientific phenomena and initial results that form the basis of the second chapter of my thesis were originally described in the Stryker lab in collaboration with Cris Niell. His advice for that project has been invaluable, and I always have appreciated his input. It has been my privilege to learn from the members of my thesis committee, and I hope to continue to interact with them as a colleague and collaborator on future projects.

Of course, my accomplishments in graduate school are the result of years of investment on the part of my parents. I am especially grateful to my mother for all of the phone calls and care packages that kept me going though out graduate school. I also would like to dedicate this thesis to my father. My father helped cultivate my own interest in neuroscience from a young age as an amateur philosopher-thinker-scientist in his own right. Constantly cutting out articles in popular science magazines and collecting books, he helped nurture an early desire to find out the underlying underpinnings of the world and ourselves. He always taught me the importance of being earnest and thoughtful, never conforming to the beliefs of others, and always to think critically. Inevitably, these were the qualities that I drew upon most to get through graduate school.

Thanks…

Abstract of the Dissertation

Locomotion is a behavior essential for survival. It is important for guiding goal-directed approach towards desired outcomes and avoidance of aversive stimuli. To this end, a large number of processes in the brain are both regulated by and serve to inform the locomotor behavior of animals. Here, we attempt to define the neural circuits underlying locomotor control, the associated changes that locomotion has upon brain states, and the neurobiological basis of locomotor decisions. In Chapter 1, we describe what is known regarding the neural circuits guiding locomotor behaviors. We provide background also regarding the known mechanisms that guide changes in brain states and are associated with locomotion. We then touch upon recent literature attempting to understand how information is used to guide decision-making to better understand the specific problem of how locomotor decisions are made. In Chapter 2, we then present novel findings, identifying brainstem circuits that control locomotion and concurrently regulate visual processing of information in the cortex through the basal forebrain. These findings may apply to other networks beyond the visual system and form a general mechanism by which various brain regions are modulated by behavioral state. In Chapter 3, we demonstrate that these brainstem circuits are under the regulation of the basal ganglia. These studies identify a conserved, phylogenetically ancient pathway for guiding locomotion that may exist in all vertebrates and represent one of the earliest functions of the

basal ganglia system. In chapter 4, we leverage our understanding of the basal ganglia pathways for locomotor control to understand the processes of goal-directed decision-making. In chapter 5, we find that the ventral striatal shares a parallel organization to the dorsal striatum for implementing reinforcement learning to guide future locomotor decision-making. These studies into the basis of goal-directed locomotor behaviors may elucidate general principles for decision-making. Collectively, these results demonstrate control systems for locomotion are deeply interconnected with a diverse array of processes throughout the brain that guide goal-directed locomotor behaviors.

Table of Contents

List of Figures

Chapter 1

Introduction: Neural Substrates for Locomotion, Behavioral State Modulation, and Reward-based Decision-making:

Neural Substrates for Guiding Locomotor Behavior

In all vertebrates, neural circuits for coordinating the synergistic coordination of muscles to generate propulsive movement are located at the level of the spinal cord (Grillner 2003; Grillner, Hellgren et al. 2005; Grillner, Wallen et al. 2008). These circuits are often referred to as central pattern generators (CPGs) and guide locomotion whether it be swimming in fish, flight in birds, or locomotion in mammals. Numerous lines of evidence suggest that CPGs for locomotion are conserved throughout vertebrate evolution and even underlie human locomotion(Grillner, Wallen et al. 2008). Indeed, the pattern of flexion and extension in newborns is identical to those in quadruped mammals (Rhesus monkeys, cats, rats, mouse). Likewise, patients with partial spinal cord injury can regain locomotion through treadmill training, which allows sensory feedback to drive CPGs independent of descending commands from the brain(Grillner, Wallen et al. 2008).

In most cases, these locomotor CPGs are quiescent unless they are activated by descending commands from the brainstem via projections from reticulospinal neurons (Figure 1). Electrical stimulation within these brainstem areas give rise to walking, trotting, or galloping in tetrapods like cats depending on the strength of stimulation. Likewise, activation of this area in birds gives rise to walking and

flapping movements, and in fish in generates swimming at progressively higher speeds. In the midbrain, this command center for gait is called the mesencephalic locomotor region (MLR), and it is located at the mesopontine border at the caudal pole of the cuneiform nucleus and pedunculopontine nucleus (Orlovsky 1999). This area is composed of cholinergic, glutamatergic, and GABAergic populations of neurons in roughly equal proportion (Martinez-Gonzalez, Bolam et al. 2011). A subset of these neurons is believed to provide excitatory input onto the reticulospinal neurons that mediate locomotion. The serotonergic raphespinal and noradrenergic coeruleospinal systems also contribute as modulatory influences upon the spinal cord through G-protein-coupled metabotropic receptors (Grillner, Wallen et al. 2008).

In addition to forward locomotion, a number of control systems are required to execute goal-directed locomotor behaviors. The movements must be steered toward different targets of interest in the form of orienting movements. These orienting movements are coordinated by the optic tectum/superior colliculus and are mediated by tectospinal or reticulospinal pathways to motor outputs (Grillner, Wallen et al. 2008). These orienting movements are often accompanied by locomotor movements for approach and avoidance towards or away from particular portions of space (Dean, Redgrave et al. 1989). The turning movements associated with these orienting behaviors are executed by tectospinal pathways and asymmetric activity between reticulospinal populations that can induce torsion of the axial muscles along the body axis and changes in stride(Redgrave, Westby et al. 1993). The superior colliculus contains a detailed

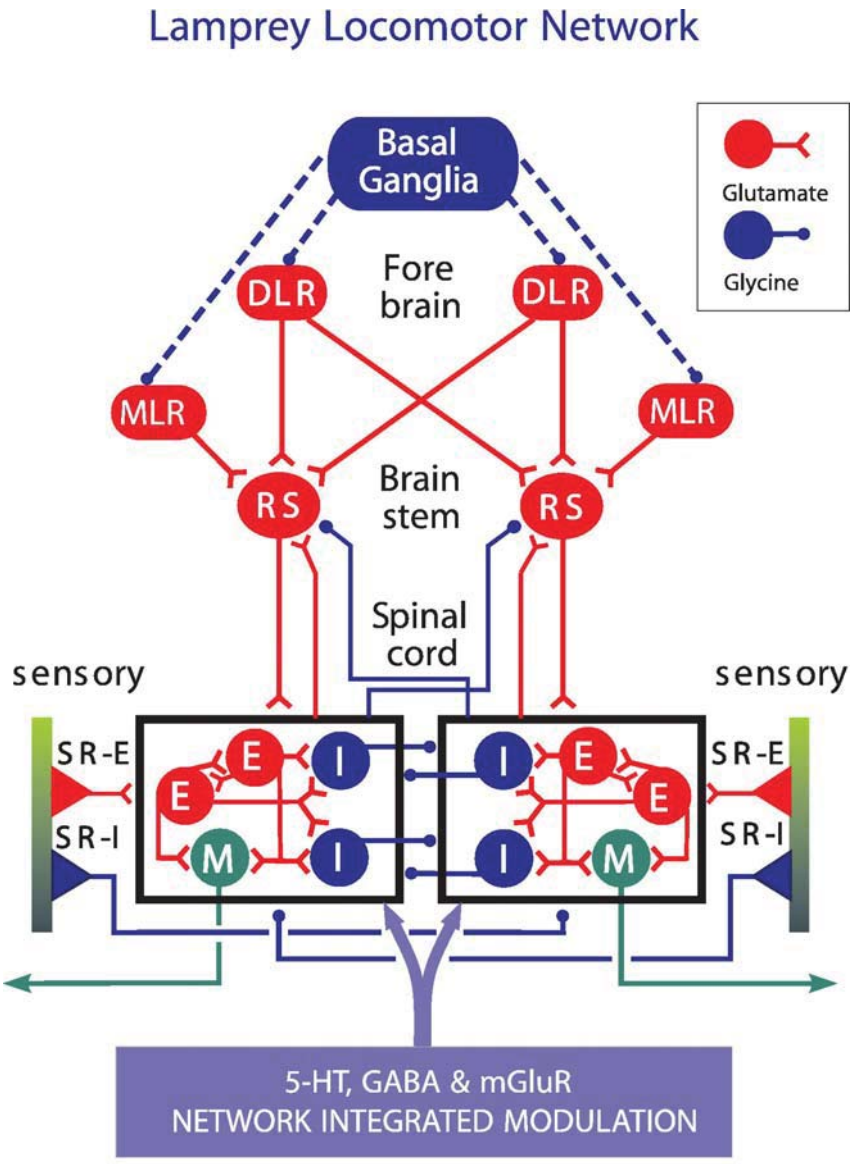

Fig. 1. Locomotor network based upon data from the lamprey.
Schematic representation of the forebrain, brainstem and spinal components of the neural circuitry that generates rhythmic locomotor activity. All neuron symbols denote populations rather than single cells. The reticulospinal (RS), glutamatergic neurons excite all classes of spinal interneurons and motoneurons. The excitatory interneurons (E) excite all types of spinal neurons, i.e. the inhibitory glycinergic interneurons (I) that cross the midline to inhibit all neuron types on the contralateral side and the motoneurons (M). The stretch receptor neurons are of two types; one excitatory (SR-E), which excites ipsilateral neurons and one inhibitory (SR-I), which crosses the midline to inhibit contralateral neurons. RS neurons receive excitatory synaptic input from the diencephalic and the mesencephalic locomotor regions (DLR and MLR), which in turn receive input from the basal ganglia as well as visual and olfactory input. In addition, metabotropic receptors are also activated during locomotion and are an integral part of the network (5-HT, GABA and mGluR).

Reproduced from Grillner S , Wallén P , Saitoh K , Kozlov A , Robertson B. Neural bases of goal-directed locomotion in vertebrates—An overview. Brain Research Reviews Volume 57, Issue 1 2008 2 – 12.

motor map of space that can elicit eye and orienting movement for different directions and amplitudes. In addition, postural systems and vestibular systems play a critical role in maintaining body orientation (Hikosaka, Takikawa et al. 2000; Deliagina and Orlovsky 2002; Deliagina, Orlovsky et al. 2006).

These brainstem systems appear to be sufficient for most locomotor control, provided that the basal ganglia and its thalamic innervation remain intact. Decorticate cats and rabbits display a large range of adaptive behaviors. They are proficient at searching for food, eating, remembering the location of previous food sources, attack behavior, and they go through states of activity and sleep (Bjursten, Norrsell et al. 1976). While primates obviously are more reliant on their cortex for motor behaviors, many of these behaviors may be mediated by cortical connections to the basal ganglia or brainstem. Selective lesions of the corticospinal tract do not impair posture or gait, but do lead to impairments in skilled hand and finger movements (Lawrence and Kuypers 1968). In the context of locomotion, the corticospinal tract may contribute more to precision walking in uneven terrain.

The Reticular Formation: Intersection Between Systems for Locomotor Control and Brain States

Locomotion is often accompanied by widespread changes in brain state. Brain state was initially studied as patterns of electroencephalogram (EEG) activity and has more recently shown to be reflected as well in the population spiking, neuronal correlations, and intracellular potentials. Perhaps, the best studied example of brain state modulation by locomotion is the theta oscillation in rodent

hippocampus. In contrast, periods of quiet immobility, often accompanied by a quiescent EEG that is punctuated by “sharp waves” and a 150-200Hz “ripple oscillation” (Buzsaki and Moser 2013). In turn, individual hippocampal pyramidal cells show phase relationships with the theta rhythm that vary across neurons or even within individual neurons on a moment to moment basis (Buzsaki and Moser 2013). Recent studies have also demonstrated that the rodent neocortex undergoes dramatic shifts in cortical states with a suppression of low-frequency power and a dramatic increase in gamma oscillations. These changes in cortical state are also accompanied by a dramatic increase in sensory responsiveness in visual cortex (Niell and Stryker 2010; Ayaz, Saleem et al. 2013). Moreover, these may be accompanied by changes from burst-to-tonic modes of firing in the thalamus (Niell and Stryker 2010). While locomotion is an experimentally convenient behavior to assay during experiments studying brain states in rodents, it is not the only behavior that regulates the state of theta in the hippocampus with various terms used to describe the list of behavior that accompany each state (“voluntary,” “active,” or “exploratory” for theta and “automatic” or “consummatory” for sharp waves) (Harris and Thiele 2011). Moreover, it is also likely that the list of behaviors for each state is specific to a species.

These data challenge the classical view that many brain states are a function of the sleep cycle with slow-wave sleep represented as a highly “synchronized” state of low frequency oscillations whereas wakefulness/REM sleep are represented as a more “desynchronized” state in which low frequency activity is suppressed. Instead, these recent studies have demonstrated a more complex

situation in which various states are also present in awake animals. In general, alert and actively behaving animals exhibit a highly desynchronized state while awake quiescent animals show synchronized fluctuations at lower frequencies(Harris and Thiele 2011). Thus, many of these brain states are profoundly regulated by the behavior of the animal with locomotion representing a particularly important behavioral state.

Interestingly, the reticular formation plays a critical role in regulating these brain states in addition to housing the motor programs for executing locomotion. The concept of a generalized activating system was first described by Moruzzi and Magoun and postulated to exist within the diffuse and widespread projections of the ascending reticular formation (Moruzzi and Magoun 1949). In anesthetized animals, stimulation of the reticular formation has been found “desynchronize” low frequency oscillations and induces a dramatic increase in gamma power of the local field potential and synchronized activity of the single units (Munk, Roelfsema et al. 1996). Activation of the reticular formation has also been demonstrated to shift the firing mode of the thalamus from a burst to tonic firing mode (Lu, Guido et al. 1993) as well as to increase theta oscillations in the hippocampus (Pignatelli, Beyeler et al. 2012). However, because many of these studies were performed in anesthetized animals, it was not possible to assess the animal’s behavior during stimulation. Thus, it is possible that stimulation may also have induced increases in locomotion had the animals not been anesthetized.

At the same time, pharmacological studies began to suggest that acetylcholine was involved in neocortical activation as described by the desynchronization of the low frequency oscillations in the cortex (Wikler 1952; Szerb 1967; Buzsaki, Bickford et al. 1988). Other studies implicated that the reticular formation may affect cortical activity through projections to the intralaminar nuclei, which send widespread efferents throughout the neocortex (Fox and Armstrong-James 1986). This influence over the intralaminar thalamus may be mediated through cholinergic neuromodulatory influences from the pedunculopontine nucleus, which resides within the reticular formation and is co-extensive with the MLR. (Hallanger and Wainer 1988). The reticular formation also sends dense projections to the nucleus basalis as well as other basal forebrain nuclei, which provide cholinergic input across the forebrain and may serve as an alternative means by which the reticular formation can affect cortical state (Saper and Chelimsky 1984). Circumscribed and restricted lesions of the basal forebrain resulted in slow delta waves during all behaviors , although the delta band was greater during periods of immobility and during consummatory behaviors as opposed to during locomotion (Buzsaki, Bickford et al. 1988). In fact, the desynchronizing influence of stimulation of the reticular formation can be in large part attenuated by blocking glutamatergic transmission in the basal forebrain (Dringenberg and Olmstead 2003). Numerous studies have identified dramatic increases in single unit activity during movement in the basal forebrain as well as in the reticular formation (Buzsaki, Bickford et al. 1988). Units within these area also undergo comparable increases in firing during wake to sleep transitions as well as during odor stimulation (Buzsaki, Bickford et al. 1988).

For years, this disparate role of the reticular formation in mediating locomotor behaviors and changes in behavioral/cognitive states has perplexed many who study the field. What possibly could be the relationship between locomotion and these “states of alertness,” mediated by the ascending reticular activating system have in common? We will return to this issue in Chapter 2.

Basal Ganglia Regulation of Neural Circuits for Locomotion

Locomotion represents one of a series of motor programs that have been conserved throughout vertebrate evolution that for the most part are generated at the brainstem and spinal cord levels. Anatomical studies in primates, rodents, (Buzsaki, Bickford et al. 1988) and lamprey have demonstrated that all of the components of the basal ganglia, including striatum, GPi, SNr, GPe, and STN and their connectivity, are conserved from lamprey to mammals. Lampreys, in particular, are an interesting species for comparative studies because they represent the oldest vertebrate species still existent today, diverging from a common ancestor around 560 million years ago. Therefore, the organization of systems for guiding locomotion is likely to be similar across various vertebrate species.

The basal ganglia provide tonic inhibition upon brainstem motor programs for locomotion and orienting movements, such as the MLR and tectum (Figure 2). The removal of tonic inhibition from basal ganglia output nuclei in the substantia nigra reticulate (SNr) and globus pallidus interna (GPi) are essential for release of these motor programs. These output nuclei fire at high rates to provide tonic inhibitory tone onto these motor regions. In the classic Albin-DeLong model

(Albin, Young et al. 1989; DeLong 1990; Gerfen 1992) of basal ganglia function, MSNs of the dorsal striatum that project to basal ganglia outputs in the substantia nigra reticulata (SNr) and medial globus pallidus (mGP) compose the 'direct pathway' (dMSNs) of the basal ganglia. Therefore, the function of the dMSNs is to disinhibit motor programs through basal ganglia outputs and promote motor responses. A second class of MSNs are the main input of an "indirect pathway" (iMSNs), projecting to the external globus pallidus (GPe) or lateral globus pallidus in rodents , which in turn sends inhibitory connections to basal ganglia outputs (SNr/mGP). Activation of iMSNs is predicted to have the opposite response, serving as a brake on movement by disinhibiting basal ganglia outputs and suppressing downstream motor programs (Kravitz, Freeze et al. 2010). dMSNs correspond to the D1R-expressing subtype of MSNs while iMSNs correspond to D2R-expressing MSNS. Although this scheme is highly oversimplified and not consistent with all of the available physiological and anatomical data, it has proved to be a valuable framework for a clinical understanding of numerous disorders and for guiding experimental investigations.

The basal ganglia's regulation of systems for guiding orienting and steering commands in the brainstem is a particularly interesting because it naturally poses an action selection or decision-making problem (Redgrave, Prescott et al. 1999). This is because while many motor programs are compatible with one another (for instance, it is possible to walk and chew at the same time), motor programs for turning left versus right in the tectum are fundamentally mutually exclusive due to features of the vertebrate body plan. This constraint creates a situation in which a

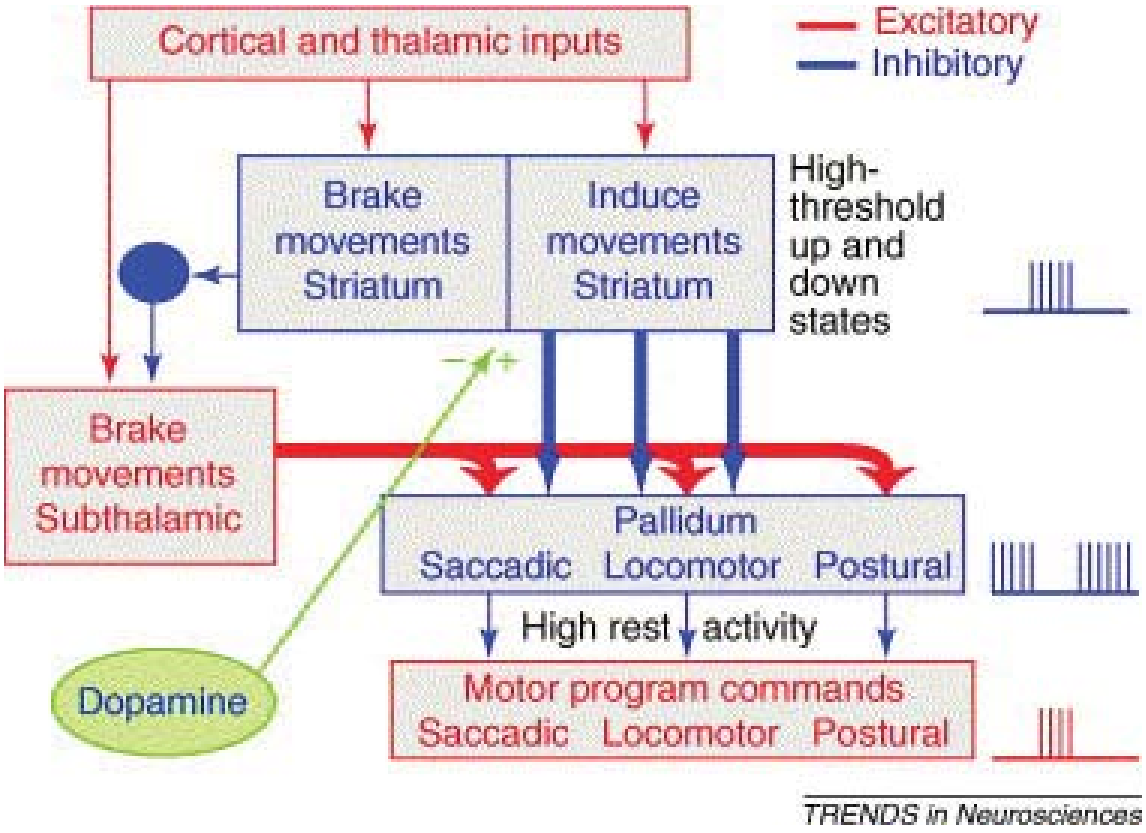

Figure 2. Basic building blocks of the basal ganglia, from a functional perspective. Inhibitory groups of neurons are indicated in blue, excitatory groups in red. The output stage of the basal ganglia is the pallidum, which includes several output nuclei (substantia nigra pars reticulata, and the dorsal and ventral pallidum) all characterized by GABAergic neurons with a high resting level of activity (~90Hz). These pallidal neurons target a large number of brainstem nuclei, in addition to their well-known thalamocortical targets (not illustrated). Subgroups of pallidal neurons project to the command neurons of several different motor programs (those for saccadic eye movements, locomotion and postural tone), and prevent them from being active under resting conditions. When the pallidal neurons in turn are inhibited by the striatum, a motor program can be released (compare spike trains to the right, illustrating correlated activity of striatal neurons, pallidal neurons and neurons of a motor program). The striatal neurons can thus induce activity and release or select a motor program (right half of the blue box representing the striatum). These neurons have a high threshold for activation, and their excitatory input is provided directly from the thalamus and extensively from different cortical regions (upper red box). A subgroup of striatal neurons with different properties is indicated on the left. Their activity is depressed by dopamine and they are instead connected to the globus pallidus pars externa, as indicated by a blue connecting neuron to the left, which in turn disinhibits the excitatory subthalamic neurons. These also receive input directly from the cortex, and they add excitation to the pallidal neurons – thus, they can brake or terminate a motor program (or part of a motor program).

Reproduced from Grillner S, Hellgren J, Ménard A, Saitoh K, Wikström M. Mechanisms for selection of basic motor programs – roles for the striatum and pallidum.Trends in Neurosciences Volume 28, Issue 7 2005 364 - 370

priority or value needs to be assigned to one motor program over the other. Thus, while many of us associate the term with a high order cognitive process, decision-making, the selection among mutually exclusive and discrete choices is a fundamental problem encountered by all organisms. In vertebrates, there is a substantial amount of evidence to suggest that the basal ganglia is involved in this process of decision-making (Redgrave, Prescott et al. 1999; Kable and Glimcher 2009).

Using Neural Circuits that Guide Locomotion to Understand Decision-making

Most experimentally tractable preparations for understanding decision-making involve some type of action to indicate a choice. Most of our understanding of the neural substrates for decision-making arises from either imaging or electrophysiological studies recording the activity of brain regions in the context of action selection, planning, or cancelation. Thus, some have argued that all forms of decision-making are an elaboration on simpler transformations of sensory information to motor outputs. Yet, decisions often feel as if they have a type of commitment to an abstract categorical proposition rather than a mere motor response.

While these issues are largely speculative and open for debate in the nascent field of decision-making, here, we describe decision-making as consisting of a two-step process. In the first step, values are assigned to particular actions through a process of learning. In the second step, these values for particular

responses then serve to determine the probability of selecting a particular motor response before being delivered to the motor system for implementation (Sugrue, Corrado et al. 2005; Kable and Glimcher 2009; Lee, Seo et al. 2012). This simple two step description is utilized in nearly all theoretical accounts of decision-making, whether they arise from economics, psychology, or computer science. The belief in a central valuation process is essential if diverse dimensions of options are to be taken into account when informing a decision. In this way, all information relevant for a decision, whether it be the sensory qualities, contextual environments, cognitive states, or internal drive states (hunger, thirst, etc) must converge to create an abstract measure of subjective value. Thus, decision-making fundamentally can be thought of as a process of dimensionality reduction, which may be mapped onto a single axis of value (Sugrue, Corrado et al. 2005; Kable and Glimcher 2009; Lee, Seo et al. 2012). From a formal standpoint, this mapping is required for all of our actions to be consistent. After this first step of valuation is achieved, this value must be converted to a single choice, which at least for experimental purposes must be made manifest as a motor behavior. Therefore, information from various neural systems must be routed to guide the endpoint of generating motor behavior.

While the idea of a common currency of value was initially proposed based upon theoretical arguments, there is increasing evidence that value signals exist in the brain (Montague, Hyman et al. 2004; Padoa-Schioppa and Assad 2006; Rangel, Camerer et al. 2008). These data suggest that these signals are present throughout the brain and have been identified in the striatum and basal ganglia (Sugrue, Corrado et al. 2005; Kable and Glimcher 2009; Lee, Seo et al. 2012).

We next will elaborate on traditional theories of reinforcement learning and review the evidence that these systems exist in the basal ganglia before discussing how these results may inform our understanding of neural substrates for locomotor decision-making.

Theories of Decision-making and Reinforcement Learning

Reinforcement learning provides a formal computational framework for understanding the decision-making process that yields concrete quantitative predictions regarding behavioral responses without the need to resort to semantic definitions or psychological labels (Kepecs 2013). The goal of reinforcement learning approaches is to describe the actions that an animal ought to take to maximize future rewards given certain environmental circumstances. In reinforcement learning, the world is parsed into a series of states that are experienced by the subject, and the subject must select among mutually-exclusive actions in these states, which may or may not be rewarded (Sutton RS 1998).

In this context, values are a reflection of the animal's internal estimate of total future rewards. Some formulations of reinforcement learning propose the existence of two separable processes that calculate two related types of value functions. One process is the Actor, which supports action selection, and the other process is the Critic, which evaluates outcomes to support learning that guides future actions (Sutton RS 1998; Joel, Niv et al. 2002; Niv and Schoenbaum 2008).

In this framework, the Actor calculates action value functions, *Q(s,a),* which are estimates of total future expected rewards for taking a particular action (*a*) in a given state (*s*) of the environment. The Actor uses these action values to determine the probability of selecting a particular response. When the value of an action is increased relative to the alternative, the probability of performing that action will be higher. In this way, the process of action selection executed by the Actor is formally equivalent to mapping action values onto a probability of making a particular response (Joel, Niv et al. 2002; Niv and Schoenbaum 2008).

However, learning is also required to acquire the appropriate values that guide future action selection. Reinforcement learning therefore posits that there is a Critic, which calculates state value functions *V(s)*. These state values are estimates of the total future rewards from a particular state of the animals' environment (*s*). In general, state value functions are used to evaluate the outcome of actions. Learning occurs whenever the actual reward that an animal experiences deviates from its prediction for the value of the state of the environment. This is driven by a reward prediction error, which serves as a teaching signal that updates current value estimates. Thus, rewards that are larger than expected should trigger learning in the form of a positive reward prediction error that increases both state and action value estimates. The higher action values should then increase the probability of the action being taken in that particular scenario in the future (Sutton RS 1998; Joel, Niv et al. 2002; Niv and Schoenbaum 2008).

While these abstract quantities of value in reinforcement learning seem needlessly abstract, they largely serve to formalize traditional psychological constructs such as motivation and reward, which have been validated by a large body of prior behavioral experiments. Action values represent an animals' motivation to perform an action. Rewards, in turn, increase the probability and vigor of actions that bring about their occurrence by enhancing the motivational value for particular actions.

Neural Substrates for Reinforcement Learning and Decision-making

In recent years, increasing evidence has pointed to a central role of the striatum and the extended basal ganglia system in playing a vital role in reward-based decision-making. Electrophysiological recordings in the striatum of primates and rodents have identified activity that parallels learning of rewarded responses (Barnes, Kubota et al. 2005; Pasupathy and Miller 2005). Other studies have shown that striatal signals correlate with the value of actions, vigor of responses (Lauwereyns, Watanabe et al. 2002; Samejima, Ueda et al. 2005), and chosen outcomes (Lau and Glimcher 2008). Indeed, striatal activity is essential for the acquisition and execution of goal-directed behaviors (Balleine, Delgado et al. 2007).

The basal ganglia are thought to regulate motor behavior in at least two distinct ways: (1) by engaging plasticity mechanisms to modify future behavior, and (2) by exploiting previously-learned associations to modify ongoing actions. Because learning ultimately impacts performance, it is useful to consider these features together, as two related aspects of basal ganglia circuit function. Traditionally,

the dorsal striatum is believed to play an important role in action selection and motor control while the ventral striatum is hypothesized to support motivated behavior and reinforcement (Montague, Hyman et al. 2004; O'Doherty, Dayan et al. 2004). Therefore, many attempts have been made to map the functions of the "Actor" and the "Critic" in reinforcement theory onto the dorsal striatum and ventral striatum respectively.

The striatum serves as the primary input nucleus of the basal ganglia and can be functionally segregated into dorsal and ventral subregions. In primates, the dorsal striatum is composed of the caudate and the putamen complex, which are separated by the internal capsule. In rodents, a contiguous dorsomedial and dorsolateral striatum serve similar functions. The term "ventral striatum" was first introduced by Heimer and Wilson (1975), who defined it as the portion of striatum associated with afferents from limbic structures such as the amygdala, hippocampus, midline thalamus, and certain regions of the prefrontal cortex (PFC) (Heimer L 1975). While this projection zone also encompasses ventral aspects of the caudate/putamen, the term ventral striatum is often used synonymously with the nucleus accumbens. The nucleus accumbens can be further divided into two distinct zones based upon cytoarchitectonic, anatomical, and histological features: the core (NAc-Core) and the shell (NAc-Shell).

Despite the functional distinctions that are assigned to them, these subregions share many organizational features. The principal neurons of the striatum consist of GABAergic medium spiny neurons (MSNs), which represent >95% of all striatal neurons. This population of MSNs can be further classified into two types:

1) MSNs expressing the D1 dopamine receptor (D1R) and the neuropeptides substance P and dynorphin, and 2) MSNs expressing the D2 dopamine receptor (D2R) with the neuropeptide enkephalin. Tonically-active cholinergic neurons and GABAergic interneurons compose the remaining fraction of neurons. GABAergic interneurons can further be classified into fast-spiking, corresponding to histochemically identified parvalbumin-postive, and low-threshold spiking, which can be sub-classified histochemically into somastostatin-, nitric-oxide-synthase, and calretinin-positive interneurons (Gerfen and Surmeier ; Rymar, Sasseville et al. 2004; Kreitzer and Malenka 2008).

One of the major roles of the striatum is to integrate sensory, motor, and associative information from various afferents to promote goal-directed behaviors. The striatum receives excitatory inputs from the cortex, thalamus, amygdala, and hippocampus, as well as dopaminergic input from the midbrain (Sesack and Grace ; Voorn, Vanderschuren et al. 2004). These inputs are segregated to various divisions of the striatum in a complex manner. Additionally, individual striatal neurons integrate converging afferent signals from multiple regions of the brain. Ultimately, striatal encoding derives from the precise and intricate temporal integration of information of these inputs (Figure 3).

In recent years, a number of electrophysiological studies have identified the existence of decision variables prominent in algorithms for reinforcement learning within the striatum. In these tasks, monkey subjects performed a decision-making task where their probability of responses dynamically adjusted

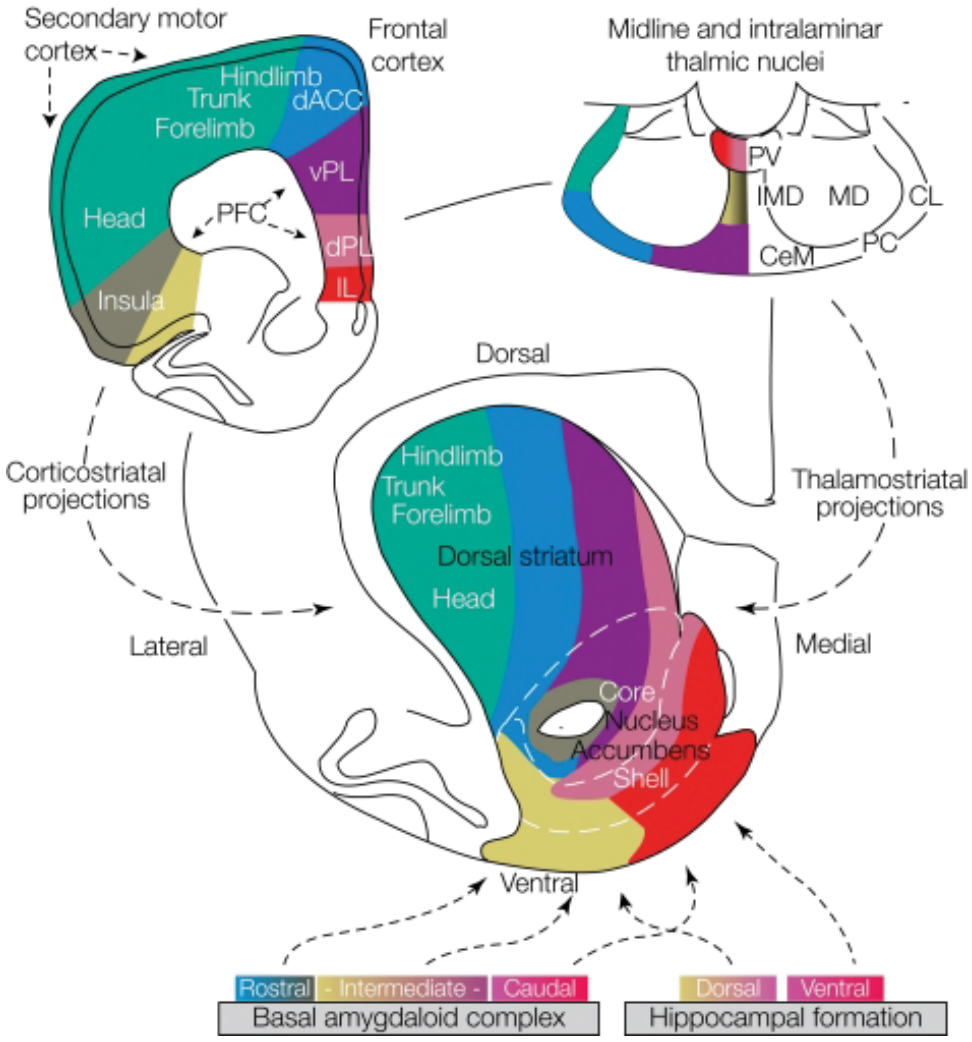

Figure 3. Cortical and thalamic inputs to the striatum distribute in dorsomedial-to-ventrolateral zones. The topographical arrangement of striatal afferents originating in the frontal cortex (upper left), midline and intralaminar thalamic nuclei (upper right), basal amygdaloid complex (lower left) and hippocampal formation (lower right) are illustrated. All these excitatory striatal afferent projections are strictly topographically organized. Thus, longitudinal striatal zones with a slightly oblique dorsomedial-to-ventrolateral orientation receive converging inputs from specific cortical areas that are, in turn, mostly interconnected through corticocortical fibers [29]. Frontal cortical areas and their corresponding striatal projection zones are shown in the same colors. The dorsolateral striatum receives somatotopically organized sensorimotor information [68] (green), the most ventromedial part of the striatum collects viscerolimbic cortical afferents (red and pink), and striatal areas between these extremes receive information from higher associational cortical areas 28 and 29 (blue and purple). Note that the topographical organization in the corticostriatal projections is the leading organizational principle, but thalamic and amygdaloid afferents nicely match this functional–anatomical organization. In the figure, the individual midline and intralaminar nuclei (upper right) are identified with different colors that match those used for the frontal cortical areas and the striatal zones to which they project. Thus, the midline paraventricular nucleus (red and pink), belonging to a group of viscerolimbic midline thalamic nuclei, projects to the ventromedial striatum [63]. At the other extreme, the posterior and lateral intralaminar thalamic nuclei (green and blue) are associated with primary motor functions and project to the dorsolateral aspects of the striatum. More ventrally and medially located intralaminar thalamic nuclei (purple and yellow), probably subserving polymodal sensory and cognitive functions, project onto longitudinally oriented striatal zones intermediate between the two extremes, matching the corticostriatal zones [63]. Similarly, the amygdalostriatal projections exhibit a mediolateral organization (color coding at bottom left). Caudal basal amygdaloid nuclei (red and pink), associated with viscerolimbic functions, project most medially, whereas nuclei of the rostral basal amygdaloid complex (yellow) send their fibers more laterally in the striatum [64]. Thus, amygdaloid fibers reach in a topographical way virtually the entire striatum, the most dorsolateral sensorimotor part being only very sparsely innervated. Frontal cortical, amygdaloid and midline and intralaminar projections to the striatum are arranged such that multiple interconnected networks exist between specific frontal cortical areas and distinct amygdaloid and thalamic nuclei that converge onto the same striatal region 63 and 69. Finally, the hippocampal formation (in particular the subiculum and the CA1 region) projects to the most ventral parts of the striatum, specifically to the medial, ventral and rostral shell, as well as to the immediately adjacent parts of the core. As indicated in the lower right corner, neurons of the dorsal (yellow and pink) and ventral (pink and red) hippocampus project laterally and medially, respectively [62]. Abbreviations: ac, anterior commissure; ACd, dorsal anterior cingulate cortex; AId, dorsal agranular insular cortex; AIv, ventral agranular insular cortex; CeM, central medial thalamic nucleus; CL, central lateral thalamic nucleus; IL, infralimbic cortex; IMD, intermediodorsal thalamic nucleus; MD, mediodorsal thalamic nucleus; PC, paracentral thalamic nucleus; PFC, prefrontal cortex; PLd, dorsal prelimbic cortex; PLv, ventral prelimbic cortex; PV, paraventricular thalamic nucleus; SMC, sensorimotor cortex. *Adapted from Voorn P, Vanderschuren L, Groenewegen h, Robbins T, Pennartz CM.Putting a spin on the dorsal–ventral divide of the striatu. Trends in Neuroscience. 2004*

to the block-wide estimates of the rewards earned for each response (Samejima, Ueda et al. 2005; Lau and Glimcher 2008). They then recorded in the caudate and fit the behavioral responses of the animal using reinforcement learning algorithms. These studies identified a neural signal that closely correlated with estimates of “action values” that they derived from the algorithms. The activity of these single units tracked the value associated with a particular behavioral responses, but did not necessarily reflect the eventual chosen motor response. “Action value” signals occurred early in trials just prior to the animal execution of a motor response and corresponded to roughly one third of responses in task-relevant regions of the striatum in multiple studies. For this reason, “action value”-correlated responses were hypothesized to be important for action selection as opposed to the implementation of the motor response, which may occur downstream of the striatum, and are consistent with notions that the striatum may serve as an “Actor” in specific types of reinforcement learning algorithms.

A second type of signal was later identified in the striatum, which may serve to facilitate learning and reinforcement (van der Meer and Redish 2011). These signals seemed to reflect the amount of reward that was associated with a chosen response. These neural representations may be analogous to the “chosen/state values,” which the “Critic” utilizes in some descriptions of reinforcement learning (Lau and Glimcher 2008). This activity was often present at the start of a trial and would ramp up in activity, just prior to the receipt of the outcome of the trial (Lau and Glimcher 2008; van der Meer and Redish 2011).

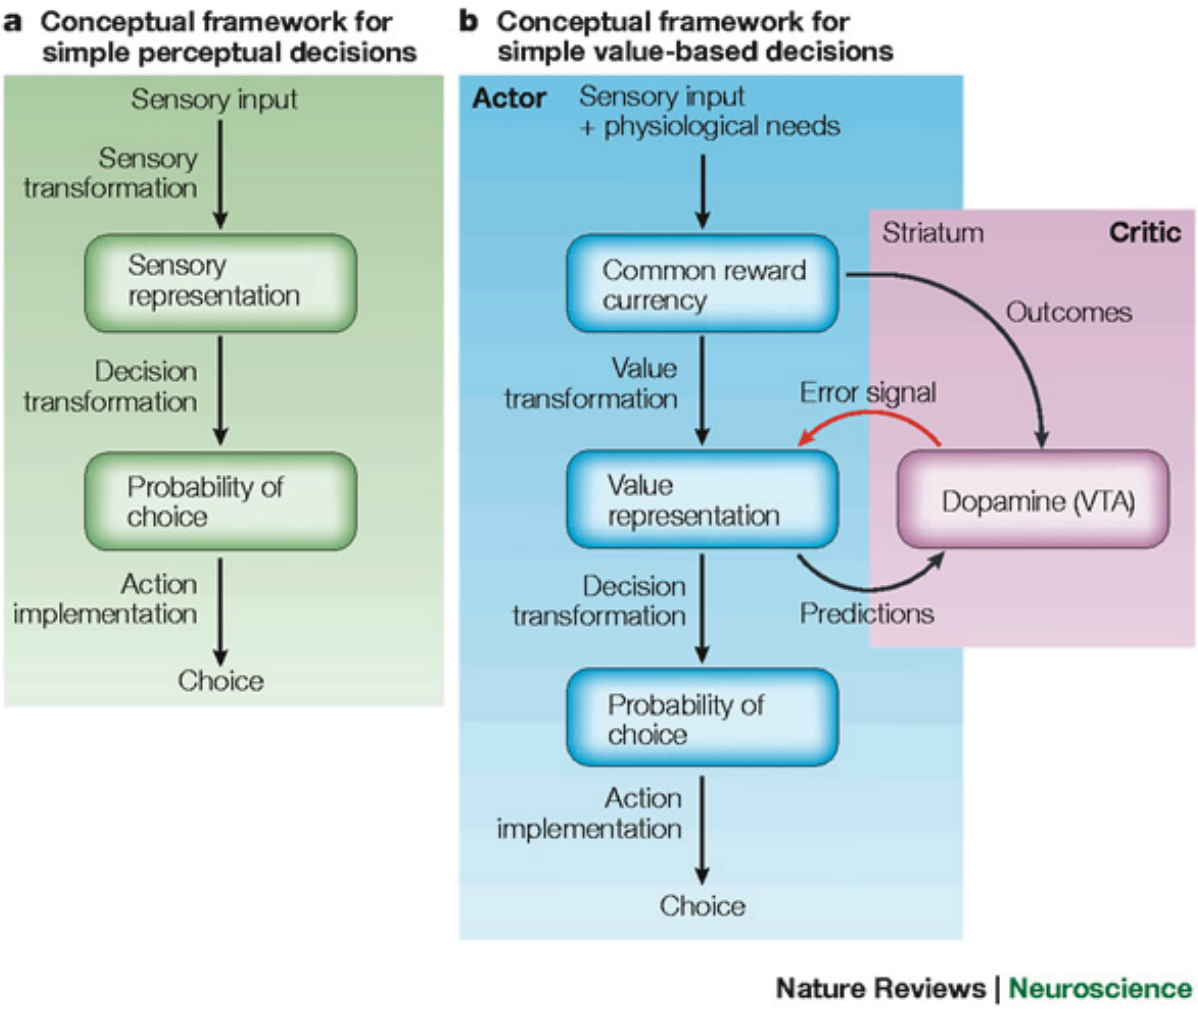

Figure 4: A conceptual framework that illustrates proposed processing stages for the formation of simple perceptual and value-based decisions. a | Perceptual decisions. A sensory transformation operates on primary sensory input to generate a representation of a higher-order stimulus dimension (for example, visual motion or auditory space). A decision transformation maps this sensory representation onto the probability of alternative operant responses. A final processing stage renders the actual binary decision, reducing the continuous probabilistic representation to a discrete plan for motor action. b | Value-based decisions. The absence of a dedicated sensory system for transducing rewards means that sensory input and physiological needs must first interact to identify 'rewards' in the animal's environment. For simplicity, we assume that this initial processing produces a 'common reward currency', which can be considered as the primary input to subsequent stages. This framework is considered to have an 'actor–critic' architecture. Within the actor component, the reward input is transformed into a higher-order representation of the value of different stimuli. Through the action of the critic, this mapping can be optimized to the environment. A decision transformation maps this value representation onto the probability of available behavioural responses. At the final stage, this representation is reduced to a single behavioural choice. VTA, ventral tegmental area, the midbrain origin of the dopaminergic neurons that contribute the 'error signal' to our proposed actor–critic architecture.

This is consistent with ideas that chosen value signals may play a prominent part in the evaluation of outcomes of choices. In rodent studies, these “chosen value” signals appear to be more prominent in the ventral striatum than in the dorsal striatum (Pennartz, Berke et al. 2009; van der Meer, Johnson et al. 2010; van der Meer and Redish 2011). These data are paralleled by results in human imaging studies that demonstrate the ventral striatum may be more active in tasks where a subject is passively rewarded in relationship to cues whereas the dorsal striatal signal were enhanced when a subject must perform an operant motor response.(Balleine and O'Doherty ; O'Doherty, Dayan et al. 2004; Pennartz, Berke et al. 2009; van der Meer, Johnson et al. 2010). The surprising degrees of consistency in studies across species provide evidence that the dorsal striatum may play an important role in action selection while the ventral striatum may be more relevant for reinforcement and learning.

Given that striatal neurons may encode the value of actions and states, how can this information me translated into the processes of action selection and reinforcement learning. The classical “two pathway” architecture of the basal ganglia provides an intuitive mechanism for the expression of reward-based motor performance. Within a given context, activation of dMSNs regulating previously-rewarded motor responses provides a motivational bias that increases the probability of such an action as well as invigorates the response. In contrast, activation of iMSNs may limit previously-unrewarded responses. Recent advances in our ability to target dMSNs and iMSNs with genetic tools have permitted tests of this hypothesis. Consistent with the classical model of basal ganglia function, prolonged, bilateral optogenetic activation of dMSNs promotes

forms of exploratory behavior such as locomotion, whereas activation of iMSNs suppress movement (Kravitz, Freeze et al. 2010). Recent studies in the auditory striatum have found that optogenetic stimulation and inhibition of corticostriatal terminals can significantly bias the responses of animals in a perceptual task following learning of a frequency-dependent sensory-response association (Znamenskiy and Zador). The magnitude of this decision bias could be predicted by the frequency tuning of the stimulated or inhibited neurons. This study lends support to theories that corticostriatal projections may provide a general mechanism for the control of motor decisions by sensory cortex. Advances in optical imaging and recording techniques have allowed for a better understanding of the encoding in these two populations of neurons during instrumental locomotor behaviors. In these studies, dMSNs and iMSNs were co-active during contraversive, but not ipsiversive locomotion (Cui, Jun et al. 2013). This argues for a more dynamic Go-NoGo organization where coordinated activity in striatal circuits is also involved in the process of shaping a locomotor response, as opposed to a simple model in which dMSNs or iMSNs merely mediate a generalized Go or NoGo command.

During learning of a motor response, synaptic plasticity is thought to adjust afferent drive in both pathways, in order to facilitate the selection of appropriate actions and the suppression of competing alternatives within a particular context. Evidence from optogenetic stimulation experiments in rodents has suggested that respective activation of dMSNs/iMSNs of the dorsomedial striatum is sufficient to induce conditioned place preference/aversion and reinforcement/ punishment of appetitive approach. (Kravitz, Tye et al.) Learning can also be

facilitated by evaluative systems associated with the ventral striatum. Recent optogenetic studies have bolstered the hypothesis that the ventral striatum exerts bi-directionally control over reinforcement learning. Activation of D1R-expressing NAc MSNs promotes reinforcement of responses whereas activation of D2R-expressing NAc MSNs extinguishes or blocks the expression of reinforced behaviors (Lobo, Covington et al. ; Bock, Shin et al.). These opposing behavioral effects have been documented in the NAc-shell (Lobo, Covington et al.) and -core(Bock, Shin et al.). Plasticity has also been observed in the nucleus accumbens after experience with cocaine reward (Britt, Benaliouad et al. ; Grueter, Brasnjo et al. ; Luscher and Pascoli ; Pascoli, Turiault et al. ; Thomas, Beurrier et al. 2001).

There is also strong evidence that dopamine neuromodulation within the striatum plays a critical role in mediating reinforcement. Initial evidence for a role of dopamine in reward was first established by Olds and Milner, who showed that animals will perform intracranial self-stimulation (ICSS) in the medial forebrain bundle (MFB), the rostral efferent pathway from dopaminergic neurons (Olds 1963). Further experiments demonstrated that this effect was mediated at least in part by dopamine release, primarily in the striatum. Electrophysiological experiments have demonstrated that some dopaminergic neurons fire during receipt of an unexpected reward in a manner that scales with size of reward (Schultz 1997; Bayer and Glimcher 2005). No excitation is observed if the reward is fully predicted by the cue, and the neurons demonstrate a pause in firing if the reward is withheld during the expected time of reward (Schultz, Dayan et al. 1997). These data are consistent with a subset of dopamine neurons encoding a

“reward prediction error” (RPE) signal that represents the difference between the actual and predicted magnitude of the reward received(Schultz 1997; Bayer and Glimcher 2005). This was recently tested in a study where dopamine neurons were selectively optogenetically stimulated and were able to produce cue-reward learning (Steinberg, Keiflin et al.) This error signal was hypothesized to serve as a “teaching signal,” which permits plasticity in the striatum. The timing of reward delivery is critical as rewards occurring prior to or after the expected time of reward receipt seem to also elicit transient increases in the firing of neurons (Mirenowicz and Schultz 1994). In this way, reward predictions can be thought of as occurring on a moment-to-moment basis. This type of timing specific RPE has been described in temporal difference (TD) machine learning algorithms (Niv and Schoenbaum 2008).

Results from whole cell recordings primarily from the dorsal striatum have substantiated the importance of dopamine in mediating plasticity in a cell-type specific manner. Some groups have found that the ability to induce long-term potentiation (LTP) at dMSNs is critically dependent on activation of D1 receptors (Shen, Flajolet et al. 2008), and many investigators have found that long-term depression (LTD) of synapses at iMSNs requires activation of the D2 receptor(Kreitzer and Malenka 2005; Kreitzer and Malenka 2007; Shen, Flajolet et al. 2008). Similar gating of plasticity by dopamine has been found using a spike-timing dependent plasticity (STDP) paradigm (Shen, Flajolet et al. 2008). Since dopamine is transiently released upon receipt of rewards, this implies that dopamine serves to gate plasticity of glutamatergic synapses onto MSNs.

Based upon these data, we would expect that information regarding external stimuli and internal states conveyed by glutamatergic inputs ought to be strengthened onto appropriate dMSNs by a diffuse temporal dopaminergic RPE. This increase in synaptic weights would manifest as reinforcement of Pavlovian approach and instrumental behaviors represented as the activity of behaviorally-relevant dMSNS. Such dMSN- specific plasticity has been observed after exposure to cocaine and in cocaine self-administration paradigms (Pascoli, Turiault et al. ; Bock, Shin et al. 2013). It is also possible that inputs onto iMSNs may simultaneously become depressed, disinhibiting rewarded responses. In this way, iMSNs may serve as a brake onto reward-seeking behavior (Lobo, Covington et al. ; Lobo and Nestler ; Bock, Shin et al. 2013) or restraint on inappropriate responses to facilitate behavioral switching (Dalley, Everitt et al. ; Hong and Hikosaka). The absence of reward during extinction or an aversive stimulus may induce a transient dip in dopamine to enhance synaptic connections onto iMSNs and possibly decrease synaptic weights on dMSNs, serving to punish and suppress concurrent behavior (Hong and Hikosaka). Together, these forms of plasticity may serve as the basis for reinforcement learning.

This form of dopamine-mediated plasticity may occur at numerous sites in parallel or in series. While many primate recording experiments occur after behavioral training is complete and performance is stable, many rodent experiments take place in the context of learning and task acquisition. These experiments provide support for a serial transfer of information from ventral to dorsal circuits during learning. Voltammetry experiments have found that

acquisition of phasic dopamine release in NAc-core to conditioned cues precedes similar responses in DLS by a few days in the course of acquiring cocaine-seeking behaviors (Willuhn, Burgeno et al.) Electrophysiological studies have also shown a dynamic change in representation across DMS and DLS during learning (Thorn, Atallah et al. ; Barnes, Kubota et al. 2005; Graybiel 2008). These data are consistent with the view that serial cascades of plasticity may occur in a ventromedial to dorsalateral progression across the striatum with recurring repetitions of contingencies and reward. Such a progressive set of changes in synaptic strength have been observed with recurrent exposure to drugs of abuse (Luscher and Malenka).

Goals of this Dissertation

As we've described in this introduction, the neural basis of locomotion represents a rich and understudied field. Here, we sought to understand locomotion from multiple perspectives. First, we hypothesized that the role of the reticular formation, in particular the pedunculopontine nucleus, in both locomotion and its ability to institute brain state changes consistent with alertness could be reconciled. Here, we postulate that many of the same mechanisms that underlie modulation in brain state occur concurrently with locomotion and that they are mediated by the same neural substrates. We provide evidence for this through optogenetic experiments, activating these brain regions and their ascending inputs to the basal forebrain in Chapter 2. We hypothesize that these mechanisms exist in order to support goal-directed spatial navigation involving visuomotor networks in tandem with spatial memory networks in order to support goal-directed locomotion. We next identify neural substrates that may mediate

goal-directed locomotion. We postulate that the basal ganglia may serve a vital role in this process in Chapter 3. We then demonstrate through optogenetic manipulations that distinct basal ganglia pathway bi-directionally modulated brainstem controllers for locomotion. Finally, we try to manipulate these systems in the context of goal-directed locomotion. Borrowing from David Marr's famous typology for computational studies of the brain, we attempt to investigate the process of decision-making at computational, algorithmic, and implementation levels simultaneously. Here, we focus on computational approaches to the problem of decision-making in Chapter 4 and 5. Using the locomotor system as a model system, we try map decision variables and algorithms critical in theoretical descriptions of decision-making onto neural substrates that can implement these functions. In many of these descriptions, the computation of values plays a critical role. Using optogenetic approaches, we attempt to manipulate the representations of value in various striatal systems using optogenetic tools. We demonstrate that these descriptions match the prediction of existing computational theories of reinforcement. These data complement previous recording studies identifying neural representations of values by demonstrating that neural activity can change behavior in a way that is consistent with these theories. The consistency of experimental findings and theory is remarkable. In conclusion, I posit that the basal ganglia's role in decision-making arose from selective pressures to decide among other simple motor behaviors to guide an organisms' locomotor path.

Chapter 2

The Mesencephalic Locomotor Region Regulates Cortical State

Abstract

Sensory processing is dependent upon behavioral state. Recent studies in mice have described striking increases in cortical visual responsiveness during locomotion, however, the central neural mechanisms that initiate these changes in cortical function are unknown. Here, we investigate the role of the mesencephalic locomotor region (MLR), as it is is known to initiate running and has also been described as part of the ascending reticular activating system. Above a certain threshold, optogenetic stimulation of the MLR in awake, head-fixed mice induces locomotion as well as changes in visual responses and cortical state. These changes in evoked responses and state were observed with MLR stimulation at a frequency below the threshold for overt movement, revealing that MLR effects on cortical processing are dissociable from locomotion. Stimulation of MLR projections to the basal forebrain also increases cortical responses in the absence of locomotion, suggesting a pathway that may link the MLR to cortex. These studies demonstrate that the MLR regulates in parallel changes in locomotor behavior and changes in cortical state .

Introduction

Neural networks are subject to modulation by behavioral state. State-dependent changes may be reflected in the representation of sensory stimuli in cortex. For example, during states of sleep, sensory responses are heavily attenuated while neural responses are often enhanced during states of alertness and attention. In mice, it has been shown that visual responses in the cortex dramatically increase while animals are locomoting as opposed to when they are standing quietly alert (Niell and Stryker 2010; Keller, Bonhoeffer et al. 2012; Ayaz, Saleem et al. 2013). This enhancement of visually evoked responses is accompanied by a shift in local field potential (LFP) from low frequencies to gamma oscillations (Niell and Stryker 2010; Keller, Bonhoeffer et al. 2012; Ayaz, Saleem et al. 2013). A recent study (Pollack et al 2013) has begun to elucidate the effects of neuromodulators on local cortical circuits that may mediate this effect, however, the central neural circuits that initiate these changes, and couple them with locomotor state, remain unknown.

While the general enhancement of responses across the visual field during locomotion is different from the restricted effects of spatial attention, there may be a general circuit principles that underlie both effects. In primates, studies have demonstrated that microstimulation in areas involved in orienting motor responses such as the superior colliculus (Cavanaugh and Wurtz 2004; Muller, Philiastides et al. 2005), frontal eye fields (Moore and Fallah 2001; Moore and Armstrong 2003; Armstrong, Fitzgerald et al. 2006), and lateral intraparietal cortex (Cutrell and Marrocco 2002) can enhance cortical responses similar to the

effects of spatial attention (Bisley 2011); thus, motor output and cortical sensory processing are coupled. While saccadic eye movements can be initiated by sufficiently high intensities of stimulation in each of these brain regions, these studies chose to use a subthreshold level of microstimulation in which no overt movements were made (Moore and Fallah 2001; Moore and Armstrong 2003; Muller, Philiastides et al. 2005; Armstrong, Fitzgerald et al. 2006). This critical choice to use stimulation parameters that were below the threshold for overt saccadic eye movements allowed the experimenters to dissociate changes in visual responses with stimulation from foveating eye movements.

We hypothesized that the locomotor-dependent changes in visual responses observed in mice may depend on a similar neural circuit motif in which brain regions coordinating specific motor programs may concurrently orchestrate changes in associated sensory processing. In many species, locomotion is mediated by the mesencephalic locomotor region (MLR), which is defined as the midbrain region in which electrical stimulation is sufficient to induce locomotion at short latencies (Shik, Severin et al. 1966; Grillner 2003). Anatomically, this region loosely coincides with the pedunculopontine tegmental nucleus and the cuneiform nucleus in mammals (Shik, Severin et al. 1966; Mori, Nishimura et al. 1978). Previous studies in decerebrate preparations have suggested that the MLR is able to regulate gait through descending projections, which can recruit the spinal cord central pattern generators via reticulospinal neurons to initiate locomotion (Shik, Severin et al. 1966; Mori, Nishimura et al. 1978; Grillner, Wallen et al. 2008).

The region around the MLR has also been described as part of the “ascending reticular activating system.” Electrical stimulation of this region can induce physiological correlates of alertness, such as desynchronization of low frequency oscillations (<10 Hz) of the EEG (Moruzzi and Magoun 1949) while lesions of this area can elicit a comatose state, abolishing arousal responses to typically salient sensory stimuli (Lindsley, Schreiner et al. 1950; French, Von Amerongen et al. 1952). Anatomical and functional studies have demonstrated that in addition to its descending projections to motor programs, the MLR also sends ascending projections to the thalamus and basal forebrain (Nauta W.J.H. 1958). In turn, activation of the basal forebrain is sufficient to induce changes in cortical state that are dependent in part on cholinergic neuromodulation (Sato, Hata et al. 1987; Buzsaki, Bickford et al. 1988; Rodriguez, Kallenbach et al. 2004; Hasselmo and Giocomo 2006; Goard and Dan 2009). Clinically, the pedunculopontine nucleus (PPN), an anatomical nucleus within the MLR, is a site for experimental deep brain stimulation (DBS) in patients with Parkinson’s disease and other disorders associated with postural and gait dysfunction (Hamani, Moro et al. ; Stefani, Lozano et al. 2007; Hamani, Moro et al. 2011). One of the side effects often reported in patients receiving low frequency DBS in the PPN often is the subjective feeling of “alertness;” whereas high frequency DBS has been reported to induce nonREM stages of sleep within minutes in some patients. These distinct findings may be interpreted to mean that high versus low frequency DBS may have opposing effects on the ouput of the PPN. Thus, numerous lines of scientific and clinical evidence point to the importance of the MLR in regulating behavioral state across species.

Based upon these functional and anatomical considerations, we hypothesized that ascending projections from the MLR to the basal forebrain may mediate changes in cortical processing while descending projections initiate locomotion. In this way, the same anatomical region that regulates motor behaviors may also provide a type of efference copy to the basal forebrain to regulate cortical state. To test this hypothesis, the MLR was identified in awake head-fixed rodents using optogenetic stimulation. A stimulation regime below the threshold for directly initiating locomotion was then applied while recording responses to visual stimuli. The subthreshold MLR stimulation was sufficient to induce increases in the gain of visual responses and gamma oscillations, like those normally associated with locomotion even in the absence of overt movement. To further isolate the neural pathway from the MLR involved in changes in cortical processing, inputs from the MLR to the basal forebrain were optogenetically stimulated. Stimulation of the MLR projections to the basal forebrain did not induce short latency locomotor responses, but partially recapitulated and occluded changes in cortical processing accompanying locomotion. We predict similar pathways from the brainstem to the basal forebrain may also operate in patients receiving DBS in the PPN.

Results:

Identification of the Physiologically Defined Mesencephalic Locomotor Region in Mice

The mesencephalic locomotor region (MLR) is defined physiologically as the mibrain region where stimulation can reliably induce locomotion at short latencies (Shik, Severin et al. 1966; Grillner 2003). This region was bilaterally targeted for infection with adeno-associated virus (AAV) expressing channelrhodopsin (ChR2) fused to yellow fluorescent protein (eYFP) under the CamK2α promoter. Following injection, one month was allowed to pass to permit expression of ChR2-eYFP prior to additional experiments.

Neurons expressing ChR2-eYFP in the MLR could be visualized upon histology in the brainstem (Figure 5A, B, and C). To verify that infected neurons could be driven, extracellular microelectrode recordings were performed in the infected region using a previously described experimental configuration where a mouse is free to either sit or run while its head is fixed on a spherical treadmill (Dombeck, Khabbaz et al. 2007; Harvey, Collman et al. 2009; Niell and Stryker 2010). A fiber optic was placed above the recording site, and 10ms pulse trains of blue light were delivered through a fiber optic. Short-latency single unit responses were elicited within 10ms to individual light pulses, indicating that we could drive neuronal activity within the area (Figure 5D). There was a range of variability across neurons in terms of how reliably light evoked responses could be elicited by trains. Optically evoked neural responses looked similar to spontaneous

activity from units and were reliably driven by light. While it is clear that light can drive the firing of neurons within the infected site, this activity may be recruited by direct activation of ChR2-expressing neurons or indirect synaptic activation of other neurons in the region. Animals locomotor behavior was next assessed following optical stimulation. Locomotor speed was registered by optical sensors that measured the displacement of the ball. Optical stimulation elicited robust locomotion at short latencies from the onset of stimulation, confirming that we were activating neurons within the MLR (Figure 5E, Supplemental Figures).

We were also able to titrate the intensity of laser stimulation to a point where stimulation at 20Hz reliably elicited locomotion while stimulation at 10Hz did not induce overt movement at 10Hz (Figure 5E). While some degree of locomotion was often present during trains of optical stimulation at 10Hz, it occurred with a large amount of variability and was not time-locked to the onset of stimulation trains. Thus, it was possible to identify epochs within our sessions with and without optical stimulation when locomotion was either present or absent at the lower 10Hz stimulation.

The activity of MLR neurons during periods of spontaneous locomotion or when the animal was at rest were also recorded. In general, there was an increase in the activity of MLR neurons with locomotion (Figure 5F and G). To quantify this increase in activity with locomotion, we plotted the correlation between speed and firing of MLR units (Figure 5G). The majority of neurons were positively correlated with locomotor speed while a much smaller fraction were inhibited or were not modulated by locomotion above chance. Thus, activation of units within

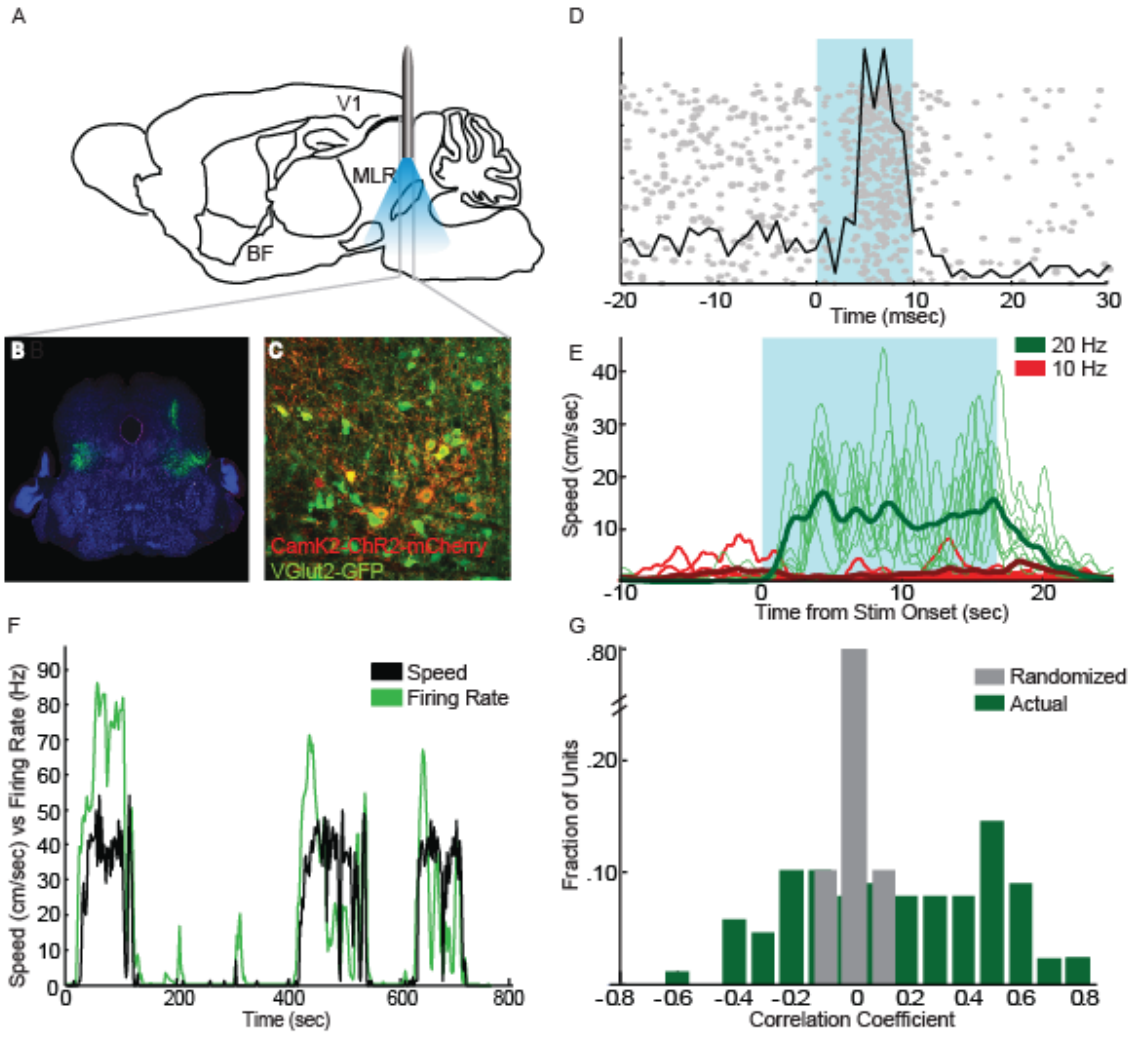

Figure 5. Activation of Neurons within the MLR Induces Locomotion, and MLR Single Units Correlate with Locomotor Parameters.
(A) Schematic of a sagittal section of the brain, depicting the MLRs location in the brainstem as well as the placement of a fiber optic stimulator. **(B)** Coronal section showing the location of a virus injection of AAV5-CamK2-ChR2-eYFP into the MLR **(C)** Confocal image at 40x demonstrating that CamK2-ChR2-eYFP infects neurons within a VGlut2-GFP mouse. **(D)** Single unit recordings from a putative light-activated neuron in the MLR aligned to 10ms light pulses in a 20Hz optical train. **(E)** Locomotor speed of an animal while being stimulated in the MLR at 10Hz (red) or 20Hz (green) in a head-fixed preparation. **(F)** Example of the firing rate of a single unit (green) and the locomotor speed of the animal (black) over the course of a recording session. **(G)** Correlation coefficients for the population of units recorded from the MLR (green) and when the locomotor data was shuffled in time (grey).

the MLR is sufficient to induce locomotion, and the activity of units within the MLR correlate with locomotor-related parameters.

The Effect of MLR stimulation on Cortical Oscillations and Gain of Visually-Evoked Responses in V1

We then proceeded to determine the effect of MLR stimulation on cortical processing in V1. A small craniotomy allowed us to insert a silicon multisite electrode into visual cortex in the same head-fixed preparation (Figure 6A). Visual stimuli were presented to the animals through a screen that was placed either directly in front of the animal or offset 45° from the animal's midline and in the visual field contralateral to the site of recording (Figure 6A). An increase in high-frequency gamma oscillations and a decrease in low-frequency power was observed in the local field potential (LFP) during periods of locomotion compared to periods when the animal was stationary similar to previous reports (Niell and Stryker 2010),. This shift in the high-frequency band occurred abruptly upon the initiation of locomotion and was present throughout bouts of movement, suggesting a transition into a different cortical state.

During optical stimulation in the MLR, an increase in gamma power was observed and a decrease in the low-frequency band within the local field potential (LFP) (Figure 6B). This pattern of LFP changes mimicked the effects of locomotion on cortical state even though the animal was stationary (Figure 6C) and could be observed across the population of animals (Figure 6D). The peak frequency of the gamma band was the same with or without stimulation, but the

gamma power dramatically increased similar to the transitions observed in different behavioral states. Moreover, the enhancement of gamma power with locomotion occluded the effects of MLR stimulation, as there was not a significant difference in gamma power with locomotion during periods of MLR stimulation (Figure 6C and D).

Single-unit visual responses were then performed in V1. Visual responses were evoked using a contrast-modulated white noise stimulus, which cycles from a gray screen up to full contrast and then back down to gray, with a ten second period (Niell and Stryker, 2008) (Figure 7A). This allowed us to quickly estimate a contrast-response function for isolated single-units. As previously reported, there was an increase in the firing rate evoked by visual stimulus across contrast levels with very little change in the spontaneous firing rate (Figure 7A). The change in the evoked responses can be represented by a change in the slope of the contrast-response function for the white noise stimuli indicating a multiplicative change in the responsiveness of V1 (Figure 7A).

Having qualitatively recapitulated previous findings observed with locomotion with our new visual stimulus paradigm, the neural responses of V1 during epochs with or without optogenetic MLR stimulation were recorded. Initially, only periods in which the animals was stationary were analyzed both with and without optical stimulation in the MLR. Similar to the effects of locomotion, MLR stimulation also significantly increased the visually evoked responses rapidly upon the onset of

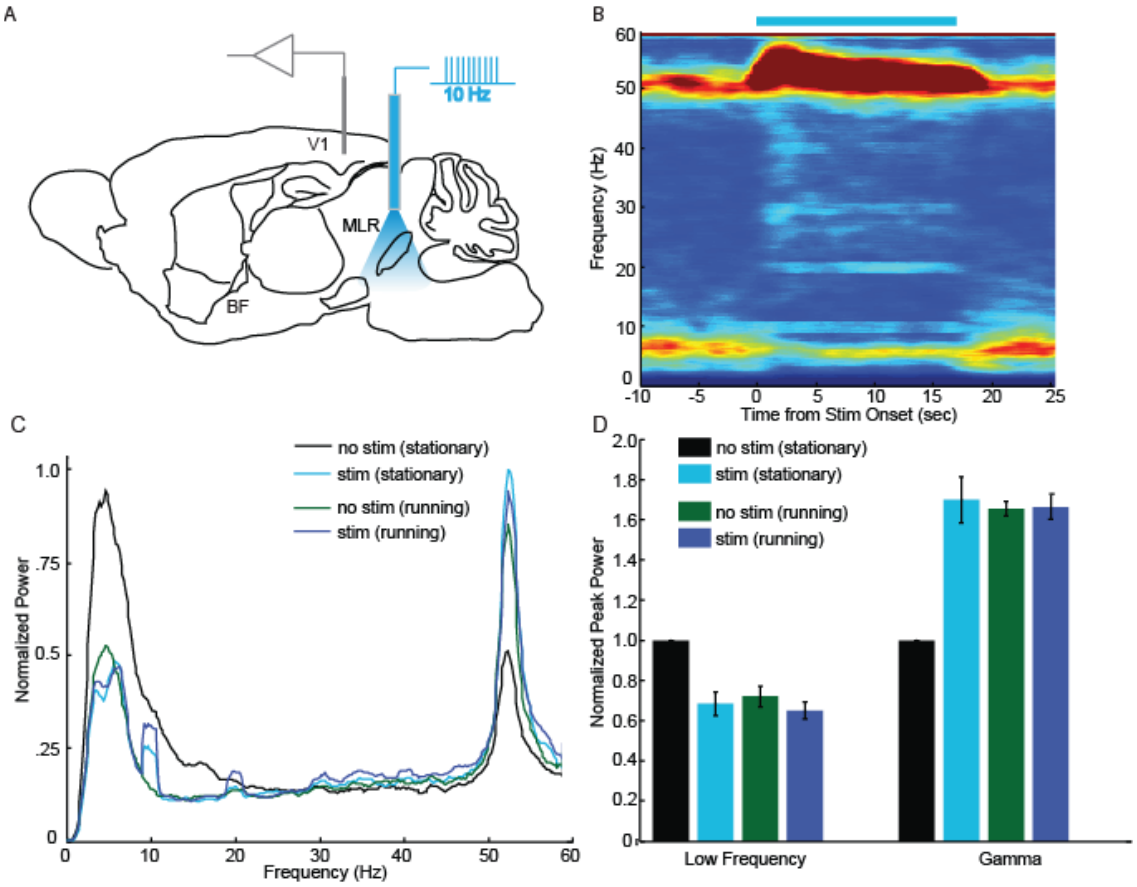

Figure 6. Optical Stimulation of the MLR Induces Changes in LFP Oscillations Similar to Locomotion.
(A) Schematic of the experimental setup with a recording array in V1 while simultaneously delivering 10Hz optical stimulation into the MLR. **(B)** LFP power across time aligned to the onset of optical stimulation. **(C)** LFP power across various frequencies in the presence/absence of optical stimulation and during conditions where the animal was stationary or running. **(D)** LFP power normalized to cases when the animal was stationary and not being stimulated for all experimental conditions.

stimulation (Figure 7B). This enhancement of visual responses could be observed across the population of recorded neurons (Figure 7C, D, E). However, there was no significant change in the spontaneous firing rate (Figure 7F). In this way, MLR stimulation below the threshold for overt movements qualitatively recapitulated all of the features of locomotion on responses in visual cortex. If the increases in visually evoked gain utilize the same mechanism as locomotion, then the effects of MLR stimulation on unit responses should occlude the effects of locomotion. Consistent with this occlusion of responses, we observed a small, but insignificant further increase in visual responses when the animal was running during optical stimulation (Figure 7G).

Stimulation of MLR terminal in the Basal Forebrain Increases in Gamma Oscillations and the Gain of Visually-Evoked Responses

Previous anatomical and functional studies have demonstrated that the MLR makes a dense projection to the basal forebrain in addition to its descending efferents (Dringenberg and Olmstead 2003; Martinez-Gonzalez, Bolam et al. 2011). To confirm the presence of these ascending projections, histology was obtained from mice injected into the region of the MLR with AAV-CamK2-ChR2-eYFP after two months of viral expression. After confirming injection at the site of the MLR, coronal sections in the region of the basal forebrain were recovered, and the presence of a dense terminal field in the basal forebrain was verified (Figure 8A, top panel). We then performed immunohistochemical staining against choline acetyltranserase (ChAT). Indeed, there were large numbers of

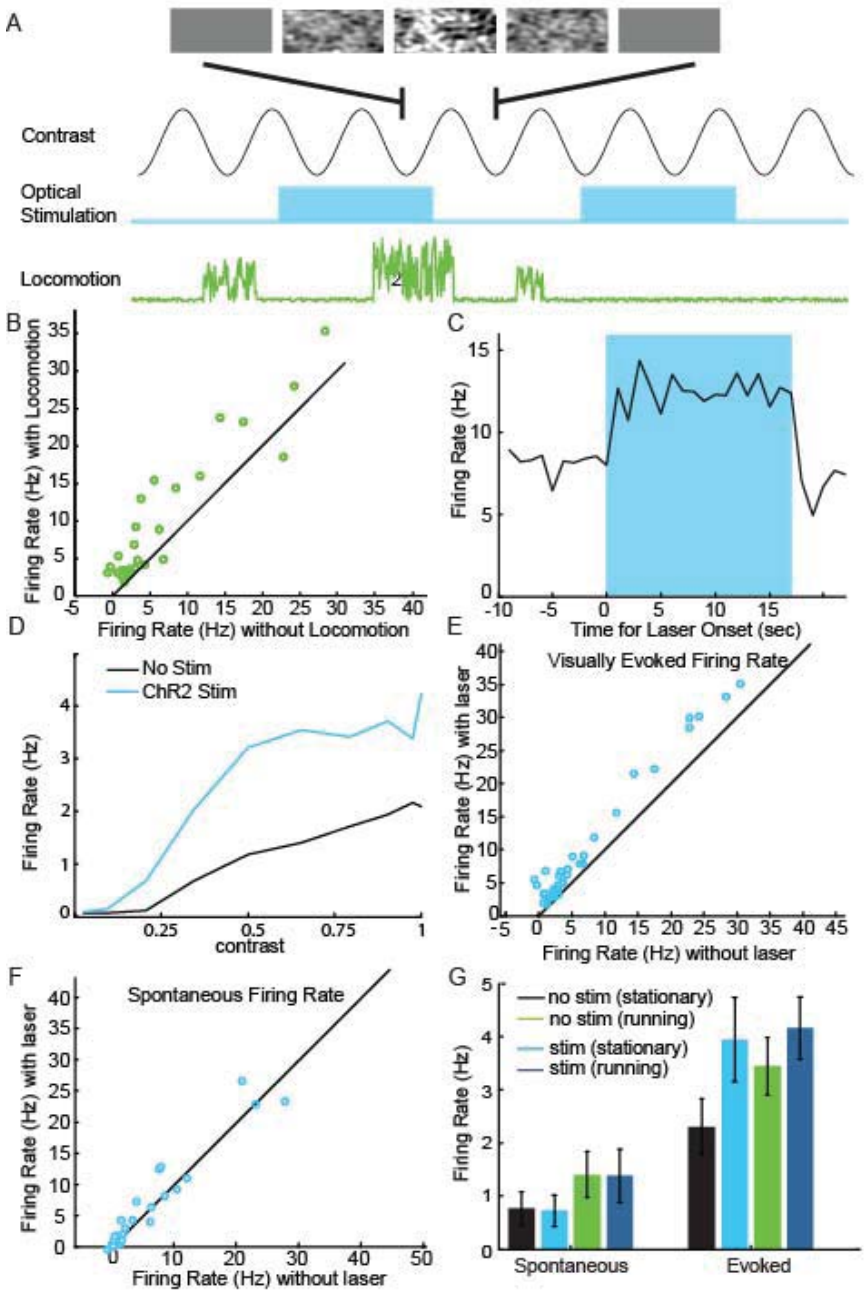

Figure 7. Optical Stimulation of the MLR increases the visually evoked responses of neurons in V1

(A) Schematic of experimental setup for the timing of the visual stimuli, optical stimulation, and spontaneous locomotion. **(B)** Time course of change in evoked firing rate averaged across various white noise contrast levels. **(C)** Example of single unit firing rate in the presence or absence of optical stimulation for various levels of contrast-modulated noise while an animal is stationary. **(D)** Visually evoked firing rate of all single units during periods when the animal is running and stationary. **(E)** Visually evoked firing rate of all single units during periods in the presence/absence of laser stimulation. **(F)** Population summary of spontaneous and visually evoked firing rates of single units when the animal is either running/stationary in the presence/absence of laser stimulation.

projections in the vicinity of cholinergic neurons of the nucleus basalis (Figure 8A, bottom panel) consistent with previous reports.

Single unit and local field potential recordings were then performed while stimulating MLR inputs to the basal forebrain at a high light intensity of 20mW. At the start of experiments, short-latency locomotor responses upon direct optical stimulation at the site of viral transfection verified expression of ChR2-eYFP within the MLR. We then moved the fiber optic stimulator to the basal forebrain to stimulate MLR terminals in the region. Anecdotally, optical stimulation also increased exploratory whisking and sniffing consistent with previous reports of basal forebrain stimulation (Berg, Friedman et al. 2005). Stimulation of basal forebrain terminals did occasionally induce an increase in locomotion, but the onset of locomotion was delayed relative to onset.

Next, LFP and single unit responses in V1 were assayed while presenting contrast-modulated white noise movies. Stimulation of MLR projections to the basal forebrain induced a shift in power to gamma frequencies while decreasing the power of low frequency oscillations even in the absence of locomotion (Figure 8B and C). However, the frequency where the peak in gamma power occurred did not change with stimulation. Although, we only analyzed periods in the absence of locomotion, these changes in LFP power were similar to the changes observed with running. In addition, locomotion occluded the effects of

stimulating MLR terminals in the basal forebrain on gamma and low frequency power as there was no additional increase in gamma power with running.

The single-unit responses as a function of contrast level were next analyzed with and without stimulation of MLR terminals into the nucleus basalis. Once again, only epochs when the animal was stationary were initially analyzed. Stimulation of MLR inputs to the basal forebrain enhanced the visually evoked firing rate of V1 neurons (Figure 8D). To quantify this, we plotted the visually evoked responses with and without stimulation of MLR efferents into the basal forebrain, and observed a significant increase in the visually evoked firing rate across the population (Figure 8E). In contrast to the changes in visually evoked responses, there was no change in the spontaneous firing rate of units with and without stimulation (Figure 8F). The multiplicative increase in the gain of visually evoked stimuli resembled the changes observed with locomotion despite the fact that we only analyzed periods when the animal was stationary. Once again, locomotion occluded the effects of stimulating MLR terminals on evoked single unit responses (Figure 8G).

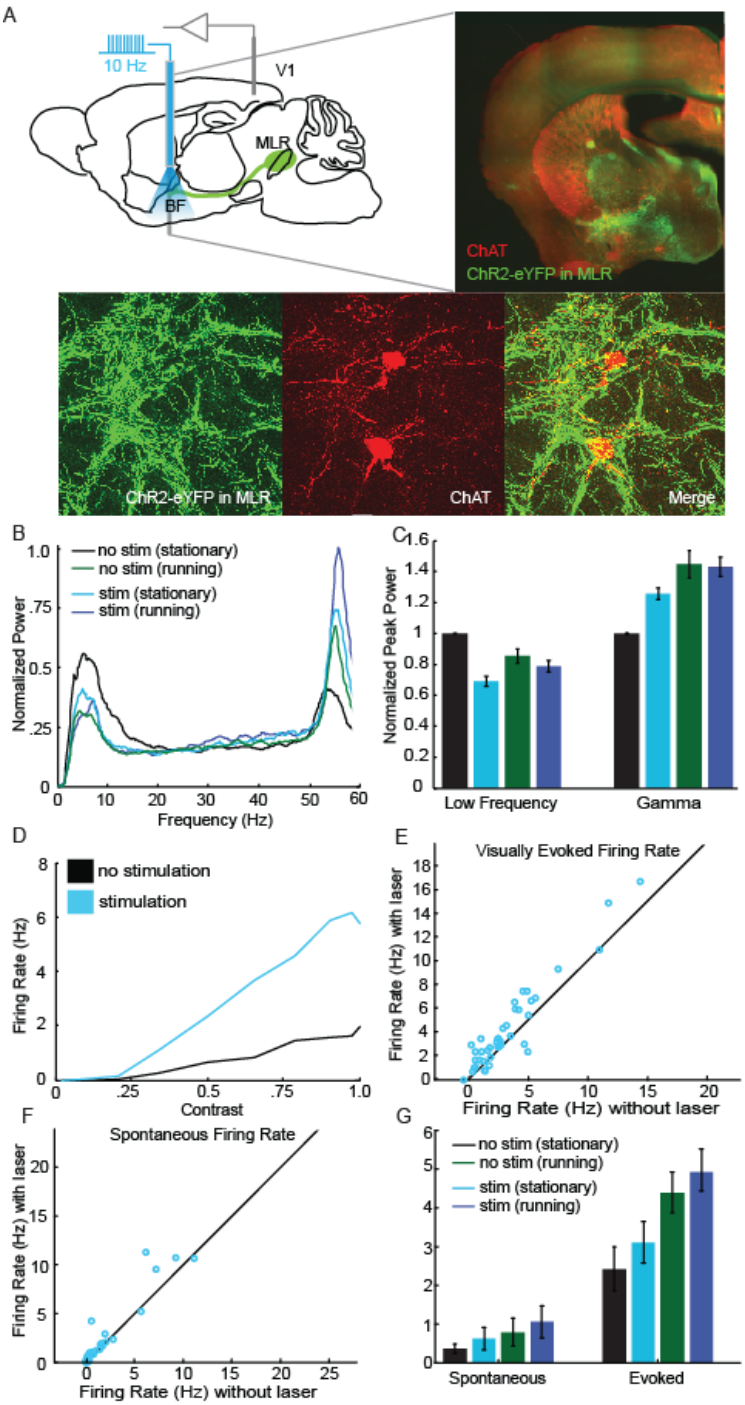

Figure 8. Optical Stimulation of MLR terminals in the basal forebrain mimics the effects of locomotion on changes in LFP and evoked firing.
(A) Schematic of the experimental setup with a recording array in V1 while simultaneously delivering 10Hz optical stimulation to terminals from the MLR, projecting to the basal forebrain. **(B)** Example of power spectrum in visual cortex in the presence/absence of optical stimulation during periods when the animal is stationary or locomoting. **(C)** Population summary of changes in low frequency and gamma power n the presence/absence of optical stimulation and during periods when the animal is stationary or locomoting. **(D)** Example of single unit firing rate in the presence or absence of optical stimulation for various levels of contrast-modulated noise while an animal is stationary. **(E)** Summary of the visually evoked firing rate of single units during stationary periods in the presence/absence of laser stimulation. **(F)** Summary of the spontaneous firing rate of neurons during stationary periods in the presence/absence of laser stimulation. **(G)** Summary of the change in spontaneous and visually evoked firing rate of neurons during periods when an animal is stationary in the presence/absence of laser stimulation.

Discussion:

The perceptual response elicited by sensory stimuli is highly dependent upon changes in behavioral state. Previous studies have demonstrated that locomotor activity could dramatically increase the visual responsiveness (Niell and Stryker 2010; Keller, Bonhoeffer et al. 2012; Ayaz, Saleem et al. 2013). Here, we provide evidence for a potential mechanism for these observed changes.

Our experimental design drew upon previous studies in primates, implicating brain regions involved in saccadic eye movements with spatial attention (Bisley 2011). Saccadic eye movments can be initiated by sufficiently high intensities of stimulation in the superior colliculus, frontal eye fields, and lateral intraparietal cortex. Microstimulation below the threshold for overt saccadic eye movements in these brain regions can recapitulate the spatially restricted effects of spatial attention on visually evoked cortical responses (Moore and Fallah 2001; Moore and Armstrong 2003; Muller, Philiastides et al. 2005; Armstrong, Fitzgerald et al. 2006). The choice to use subthreshold stimulation was critical in these studies to dissociate the effects of foveating eye movements from changes in sensory responses. We therefore reasoned that the locomotor-induced increases in visually evoked responses may rely upon a similar neural circuit motif in motor programs may can concurrently orchestrate changes in sensory processing related to the motor behavior. We optogenetically stimulated the MLR at intensities subthreshold for overt locomotion and found increases in evoked visual responses without changes in spontaneous firing rate that mimicked the effects of running.

However, optogenetic stimulation also afforded us benefits over traditional electrical microstimulation. One of the benefits of optogenetic stimulation is that we could identify the neuronal cell bodies that were being activated. This is in contrast to electrical stimulation, which may recruit activity from axons of passage as well as orthodromic and antidromic activation (Histed, Bonin et al. 2009). Second, we were able to identify that efferents to the basal forebrain from the MLR were sufficient to induce increases in visually evoked responses. Indeed, the changes in LFP with stimulation of terminals are characteristic of the cholinergic neuromodulatory influence of nucleus basalis stimulation on cortex (Alitto and Dan ; Sato, Hata et al. 1987; Buzsaki, Bickford et al. 1988; Rodriguez, Kallenbach et al. 2004; Hasselmo and Giocomo 2006; Goard and Dan 2009; Alitto and Dan 2012), and electrophysiological and microdialysis studies have identified neurons that correlate with locomotor parameters in the region of the basal forebrain (Buzsaki, Bickford et al. 1988; Kurosawa, Okada et al. 1993; Giovannini, Bartolini et al. 1998). Importantly, a recent study from Pollack et al (2013) demonstrates an important role of cholinergic modulation of local cortical circuits on state changes during locomotion. Together, these data are consistent with a model where the MLR provides the basal forebrain with an efference copy of locomotor signals to regulated cortical state through cholinergic neuromodulation. Despite evidence for this model, it is still possible that backpropagating action potentials traveling down collaterals of MLR neurons may contribute to the observed effects. This allows for the possibility that other brain regions and neuromodulatory centers in addition to basalis may also be recruited during optogenetic stimulation of MLR inputs to basal forebrain.

A General Mechanism for Regulation of Brain States Associated with Locomotion

The MLR has been studied in numerous different contexts and with a variety of nomenclatures (Martinez-Gonzalez, Bolam et al. 2011), confounding attempts to identify a unitary function of this brain region. Here, we have chosen to describe this region by the functional definition of the MLR, the area in the midbrain where locomotion can be initiated at short latencies by electrical stimulation (Mori, Nishimura et al. 1978; Grillner, Wallen et al. 2008). However, the MLR is co-extensive with the PPN, the limits of which are histochemically defined by the presence of cholinergic neurons in the dorsal midbrain tegmentum (Martinez-Gonzalez, Bolam et al. 2011; Thankachan, Fuller et al. 2012). In the sleep literature, numerous studies have implicated the PPN in sleep-wake regulation (Rye 1997). The PPN has also been described as being a part of the reticular activating system, regulating behavioral signs and electrophysiological correlates of alertness (Moruzzi and Magoun 1949; French, Von Amerongen et al. 1952). An alternative nomenclature that is used to describe the MLR is the parabrachial region due to its close proximity to the brachium conjuctivum. In anesthetized animals, stimulation of the parabrachial region has been described to regulate behavioral states across the brain. In cortex, stimulation of the parabrachial region has been found to “desynchronize” low frequency oscillations and an increase in the gamma power of synchronized neural activity and local field potential in the cortex (Munk, Roelfsema et al. 1996). Other electrophysiological studies in thalamus have described a transition from burst to tonic firing modes

with parabrachial stimulation (Lu, Guido et al. 1993). In the hippocampus, increases in theta oscillations have been observed with stimulation in this region in anesthetized animals (Pignatelli, Beyeler et al. 2012). Because these studies had been conducted in anesthetized animals, it was not possible to assay the animals' behavior during stimulation. Therefore, it was not possible for experimenters to determine whether stimulation within the PPN or parabrachial region was sufficient to induce locomotion. (Moruzzi and Magoun 1949; French, Von Amerongen et al. 1952; Nauta W.J.H. 1958).

These seemingly disparate findings can be reconciled by a simple model in which the MLR initiates locomotion through descending pathways to the spinal cord while coordinating changes in forebrain brain state through its ascending projections (Hallanger and Wainer 1988). Here, we have identified that activation of projections from the MLR to the basal forebrain is sufficient to mimic the changes in visual cortical processing observed with locomotion. The MLR also projects to the medial septum, which contains central pattern generators for inducing hippocampal theta oscillations (Buzsaki and Moser 2013), and provides direct cholinergic neuromodulatory input to the thalamus (Erisir, Van Horn et al. 1997), facilitating burst to tonic transitions in firing (Curro Dossi, Pare et al. 1991; Steriade, Dossi et al. 1991). We speculate that these other ascending projections may mediate concomitant changes in hippocampus (Buzsaki and Moser 2013) and thalamus (Niell and Stryker 2010) that accompany locomotion.

Implication for self-reported increase in "alertness" during therapeutic PPN DBS for Parkinson's Disease

Currently, the PPN has been targeted as a site for deep brain stimulation (DBS) in Parkinson's patients to relieve the freezing of gait and postural instability, which are cardinal features of the disorder (Stefani, Lozano et al. 2007; Hamani, Moro et al. 2011). Surprisingly, however, numerous studies have described that patients often feel subjectively more "alert" upon the onset of low frequency (15-25Hz) DBS in the PPN (Stefani, Peppe et al. 2013). In sharp contrast to the "alerting" effect at low frequencies, some clinical sites have reported that patients enter a non-REM sleep state upon high frequency stimulation in the PPN during procedures to tune the efficacy of DBS (Arnulf, Ferraye et al. ; Stefani, Peppe et al. ; Arnulf, Ferraye et al. 2010; Stefani, Peppe et al. 2013). These unexpected effects have given rise to questions regarding whether various clinical sites are targeting different anatomical structures and the functional heterogeneity of the PPN region (Stefani, Peppe et al. 2013).

Our data can provide a potential explanation for theses findings, demonstrating that the MLR/PPN mediates both locomotion as well as changes in behavioral state that are naturally recruited in tandem. The 20Hz optical stimulation protocol utilized in these experiments in mice falls within this low frequency range, which patients self-report as increasing their state of "alertness" (Stefani, Peppe et al. 2013). Both the desynchronization of low frequency oscillations and the concomitant increase in gamma oscillations have been described as electrophysiological correlates of an "alert" behavioral state. An increase in the gain of sensory evoked responses is also consistent an "alert" state. Thus, it is reasonable to speculate that the subjective sense of "alertness" felt by patients

during low frequency DBS in the PPN may be due to the PPN's ascending efferents to cholinergic neuromodulatory centers in the basal forebrain. Likewise, the effect of high frequency stimulation may be interpreted as a disruption of ongoing activity that may be required to sustain alertness. These clinical studies are interesting in that they help address questions that are difficult to assay in animal models because patients can subjectively report their changes in perception. Likewise, animal models can complement these clinical observations, providing a mechanistic explanation of the changes in physiology that may be associated with clinical interventions.

Together, our studies provide an integrated view of the MLR's function in concurrently regulating both locomotion and brain state. This simple brainstem to basal forebrain circuit may account for other recent findings, demonstrating changes in extrastriate visual cortex (Andermann, Kerlin et al. 2011), auditory cortex (personal communication from Wehr lab and McCormick lab), and hippocampus associated with locomotion. Thus, while these systems are affected differently by locomotion, we hypothesize that this diversity may share a common mechanism mediated by ascending projections from the MLR. Hopefully continued optogenetic dissection of these circuits will reveal the full map of connections that can mediate the effects of behavioral state on higher brain functions.

Materials and Methods (chapter 2)

Methods:

Mice

Mice were maintained in the animal facility at University of Oregon, the Ernest Gallo Clinic and Research Center, and the University of California, San Francisco and used in accordance with protocols approved by the University of Oregon, Ernest Gallo Clinic and Research Center, and the University of California, San Francisco Institutional Animal Care and Use Committees. Experiments were performed on adult C57Bl/6 mice (age 2–6 months, both male and female).

Construct and virus preparation.

Plasmids encoding the DNA sequences for pAAV-CamK2α -ChR2(H124R)-eYFP were obtained from K. Deisseroth (Stanford University). Amplification and purification of plasmids was performed using a standard plasmid maxiprep kit (Qiagen) and confirmed by sequencing. EF1α-DIO-CHR2(H124R)-eYFP and EF1α-DIO-eYFP cassettes were packaged in AAV vectors and serotyped with AAV5 coat proteins by the viral core at University of North Carolina. The final concentration was 1–2 × 10^{12} viral particles per ml.

Stereotaxic AAV injection

Animals were anesthetized with either ketamine (150 mg per kg of body weight) and xylazine (50 mg per kg) or 2% isofluorane (vol/vol) gas anesthesia. Animals were placed in a stereotaxic frame and 26-gauge microinjection needles were inserted through a burr hole bilaterally into the MLR (coordinates from bregma: -

4.75 mm anterior, ±1.2 mm medial-lateral, –3.6 ventral). 0.5 µl injections were performed using a 1-µl Hamilton syringe through a hydraulic pump (Harvard Instruments) and took place over 5 min followed by 5 min of recovery. We then allowed at least a month before recording, to allow full ChR2 expression.

Surgical preparation

Surgical preparation and recordings were performed generally as described previously (Niell and Stryker, 2010). Briefly, in preparation for recording a metal headplate was affixed to the skull with cyanoacrylate and dental acrylic. After allowing several days for recovery, in a second surgery on the morning of recording, we performed a small craniotomy (~1mm) centered over V1, and made a burr hole over the site of the viral injection and/or the basal forebrain (coordinates from bregma: -0.5 mm anterior, ±1.75 mm medial-lateral, –3.75 ventral). The headplate opening was then filled with silicone elastomer, and the animal was allowed to recover for 2-3 hours before recording. After recovery, the animal was placed in the head-holder on the floating ball and the silicone plug was removed.

Optical Stimulation

To stimulate activity in the MLR, the tip of a 200-µm fiber optic was lowered into the burrhole until it was 0.5mm above the site of viral injection. The fiber optic was covered in furcation tubing to protect the fiber and to prevent light from escaping through the optic-patch cord. A 473-nm diode-pumped solid-state laser (200 mW) (model?) provided the light stimulation. The laser driver current was adjusted to induce overt locomotion by 5 msec pulses of light at 20Hz, but no locomotion at the same intensity and pulse duration but 10Hz frequency. For subsequent experiments, stimulation was delivered at

10Hz during visual presentation for a period of 17 secs, and 17 secs was allowed to elapse between trains of stimulation. For experiments involving stimulation of MLR terminals in BF, following confirmation that stimulation in the MLR could induce movement, we then moved the fiber to the burrhole over BF, lowered to a depth of XX, and used a similar stimulation protocol.

Extracellular single-unit recordings

We used silicon multisite electrodes, following the methods of Niell&Stryker 2010, to record single unit activity in cortex during locomotion and optogenetic stimulation. Briefly, the silicon probe (a1×32–25-5 mm-177, NeuroNexus Technologies) was lowered through the craniotomy using a stereotax-mounted Microdrive (Siskiyou Designs). The electrode was placed at an angle of ~45deg relative to the cortical surface, and inserted to a depth of up to 800um below the cortical surface to record cells across multiple layers.

The electrode was placed without regard for the presence of visually responsive units on individual sites, and all units stably isolated over the recording period were included in analysis. Following placement, the electrode was embedded in agarose to increase mechanical stability, and was allowed to settle in one position for 30 min to obtain stable single-unit recordings. Recordings continued for up to several hours, after which the animal was killed under deep anesthesia by cervical dislocation.The brain was removed immediately and fixed by immersion in 4% paraformaldehyde in PBS at 4°C overnight, after which 200 μm coronal sections were cut with a vibratome. The sections were mounted using Vectashield with DAPI (Vector Laboratories) and imaged on an Olympus BX6l microscope to confirm viral expression in the MLR.

Visual Stimuli

Visual stimuli were presented as described previously (Niell and Stryker, 2008). Briefly, stimuli were generated in Matlab using the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997) and displayed with gamma correction on a monitor (Planar SA2311W, 30 × 50 cm, 60 Hz refresh rate) placed 25 cm from the mouse, subtending □60–75° of visual space. Contrast-modulated noise movies (Niell and Stryker, 2008) were created by first generating a random spatiotemporal frequency spectrum in the Fourier domain with defined spectral characteristics. To drive as many simultaneously recorded units as possible, we used a spatial frequency spectrum that dropped off as A(f)~ 1/(f + fc), with fc = 0.05 cpd, and a sharp cutoff at 0.12 cpd, to approximately match the stimulus energy to the distribution of spatial frequency preferences. The temporal frequency spectrum was flat with a sharp low-pass cutoff at 5 Hz. This 3D (ωx, ωy, ωt) spectrum was then inverted to generate a spatiotemporal movie. To provide contrast modulation, this movie was multiplied by a sinusoidally varying contrast with 10 s period. Each movie was 5 min long and was repeated two or three times, for 10–15 min total presentation.

Data acquisition

Data acquisition was performed as described by Niell and Stryker (2010). Signals were acquired using a System 3 workstation (Tucker-Davis Technologies) and analyzed with custom software in Matlab (MathWorks). For local field potential (LFP) analysis, the extracellular signal was filtered from 1 to 300 Hz and sampled at 1.5 kHz. To obtain single-unit activity, the extracellular signal was filtered from

0.7 to 7 kHz and sampled at 25 kHz. Spiking events were detected on-line by voltage threshold crossing, and a 1 ms waveform sample on 4 neighboring recording sites was acquired around the time of threshold crossing. Single-unit clustering and spike waveform analysis were performed as described previously (Niell and Stryker, 2008), with a combination of custom software in Matlab and Klusta-Kwik (Harris et al., 2000). Quality of separation was determined based on the Mahalanobis distance and L-ratio (Schmitzer-Torbert et al., 2005) and evidence of a clear refractory period. Units were also checked for stability in terms of amplitude and waveform over the course of the recording time to ensure that they had not drifted or suffered mechanical damage.

Movement signals from the optical mice were acquired in an event-driven mode at up to 300Hz, and integrated at 100msec intervals. We then used these measurements to calculate the net physical displacement of the top surface of the ball. We determined the average speed during a stimulus presentation to classify the trial as stationary (<1cm/sec) or moving (>1 cm/sec).

Data analysis.

Data analysis was performed using custom routines written in Matlab. To analyze visual responsiveness using contrast-modulated white-noise movies, we binned the spikes of each neuron according to the contrast of the movie on the screen, to create a contrast-response function. The spontaneous firing rate was defined as the firing rate for zero contrast, while evoked firing rate was defined as firing for period within 80% of full contrast or greater. To calculate responsiveness as a

function of locomotor state or laser stimulation on/off, we simply performed this analysis after filtering the timepoint of each spike for the appropriate criteria.

For LFP analysis, the power spectrum was computed using multi-taper estimation in Matlab with the Chronux package (http://chronux.org/), (Mitra and Bokil, 2008). Spectra were normalized for presentation by applying a 1/f correction (Sirota et al., 2008). To compare power in different frequency ranges across states, we selected all the 1 second time windows meeting a certain criteria (laser on/off, moving/stationary), averaged the LFP spectrum, and found the peak of power within the appropriate frequency range**.** Frequencies at a multiple of the laser stimulation frequency were excluded from this analysis to avoid including harmonics.

Statistical significance was determined by Mann–Whitney U test, except where otherwise stated. For figures representing the median of data, error bars indicate SE of the median as calculated by a bootstrap. In other cases, error bars indicate SEM.

Chapter 3

Basal Ganglia Regulation of Brainstem Control Centers for Locomotion

Abstract

The basal ganglia plays an essential role in goal-directed locomotion. In turn, dysfunction in the basal ganglia leads to freezing of gait and akinenia that serve as cardinal symptoms of disorders such as Parkinson's disease. Here, we test the hypothesis that the basal ganglia are able to regulate locomotion by bi-directionally regulating activity in the mesencephalic locomotor region (MLR). Optogenetic activation of the direct pathway increased activity within the MLR while activation of the indirect pathways suppressed activity within this region. As we have demonstrated previously (Chapter 2), stimulation of the MLR is sufficient to initiate locomotion at short-latencies, and the intensity of stimulation in the MLR dictates the speed at which animals run. Together, these data provide a circuit through which the basal ganglia can exert bi-directional control over locomotion and mediate goal-directed locomotion.

Introduction

The basal ganglia play a critical role in the selection of motor programs (Grillner, Hellgren et al. 2005). In general, most descriptions of the basal ganglia emphasize its regulation of motor thalamus, which can in turn relay information to motor cortex (Albin, Young et al. 1989). However, the basal ganglia's outputs to the brainstem represent a phylogenetically older pathway that is vital in the regulation of basic motor programs for locomotion and posture (Grillner 2003; Grillner, Hellgren et al. 2005; Grillner, Wallen et al. 2008). In particular, it has been proposed that the basal ganglia regulate locomotion through its output onto the mesencephalic locomotor region (MLR). The MLR is functionally defined as the mesopontine region where electrical stimulation is sufficient to induce locomotor responses with short latency (Shik, Orlovskii et al. 1966; Grillner 2003) . Anatomically, the MLR overlaps with the pedunculopontine tegmentum (PPTg) and cuneiform nucleus in the midbrain (Shik, Orlovskii et al. 1966; Grillner 2003). The ability of the MLR stimulation to initiate locomotion has been attributed to its descending excitatory pontomedullary inputs onto reticulospinal neurons, which can recruit spinal cord central pattern generators, mediating the bilateral alternating limb movements associated with gait (Grillner, Wallen et al. 2008).

Activity within the MLR has been proposed to be mediated by efferents from basal ganglia output nuclei such as the substantia nigra reticulate (SNr) and globus pallidus interna (GPi). The SNr and GPi fire at high frequencies and are thought to provide tonic GABAergic inhibition onto neurons in the MLR (Kang and Kitai 1990; Granata and Kitai 1991; Saitoh, Hattori et al. 2003). It has been

proposed that direct pathway activation can promote the initiation of locomotion by suppressing activity in basal ganglia output nuclei to release inhibition upon the MLR (Grillner, Hellgren et al. 2005). Conversely, indirect pathway activation ought to suppress locomotion by increasing the tonic inhibition onto the MLR from basal ganglia outputs (Grillner, Hellgren et al. 2005). Based upon these anatomical and physiological studies, the MLR has been proposed to be a critical relay from the basal ganglia to spinal cord pattern generators mediating goal-directed locomotion (Mena-Segovia, Bolam et al. 2004; Grillner, Hellgren et al. 2005; Grillner, Wallen et al. 2008).

In this framework, the basal ganglia's regulation of the MLR is important for understanding both the pathophysiology and potential therapies for basal ganglia disorders such as Parkinson's Disease (PD). The depletion of dopamine within the striatum following Parkinson's disease is believed to increase the activity of the indirect pathway while suppressing activity in the direct pathway (Albin, Young et al. 1989; DeLong 1990; Gerfen 1992). The net result should be an increase in tonic inhibition onto the MLR, providing a neural circuit description for the development of axial symptoms such as freezing of gait and postural instability associated with PD.

Therapeutically, the PPN, which is found within the functionally defined MLR, has been targeted as a site for low frequency deep brain stimulation (DBS) to treat the axial symptoms in PD (Hamani, Moro et al. 2011). While these axial symptoms are not effectively treated by DBS in more traditional sites, such as the subthalamic nuclei (STN) or the globus pallidus interna (GPi) (St George, Nutt et

al. 2010), recent studies have shown that PPN DBS alone (Mazzone, Lozano et al. 2005; Plaha and Gill 2005; Stefani, Lozano et al. 2007; Moro, Hamani et al. 2010) or in conjunction with more traditional DBS sites (Stefani, Lozano et al. 2007; Follett and Torres-Russotto 2012) can improve axial symptoms in PD patients. In this way, the PPN DBS has been proposed to potentially compensate for a hypoactivity of neurons within the MLR.

While anatomical and physiological studies in vitro have provided evidence for this model of the basal ganglia's interaction with the MLR, no study to date has directly tested the effect of direct and indirect pathway stimulation on neurons of the MLR . Here, we sought to test directly the hypothesis that the basal ganglia bi-directionally regulates activity in the MLR through the direct and indirect pathways using optogentic stimulation. As a population, we found that activation of the direct pathway increased the activity of most neurons within the MLR in addition to initiating locomotion. In contrast, indirect pathway decreased activity within the MLR in addition to suppressing locomotion. Thus, activity in the MLR was a reflection of the locomotor behavior induced by stimulation of distinct basal ganglia pathways. Based upon these optogenetic stimulation experiments, we were able to classify neurons within the MLR as either being direct pathway- or indirect pathway-activated/suppressed.

Results

To investigate the basal ganglia brainstem neural circuit for locomotion, we utilized a head-fixed preparation that allowed the mouse to run freely on a spherical treadmill. This set-up enabled us to measure parameters of locomotion (locomotor starts, immobility, locomotor speed), while allowing us to acutely perform stereotaxically targeted recordings for hours-long acute recordings in vivo. Given that prior studies that had demonstrated that direct pathway activation can increase lococmotion while the indirect can suppress locomotion, we sought to determine the effects of each basal ganglia pathway on neural activity within the MLR. Prior to recording, D1-Cre or A2A-Cre BAC transgenic mice were bilateral injected with AAV5-EF1a-DIO-ChR2-YFP into the dorsomedial striatum (+0.5 RC, +/- 1.5 ML, -3.0 DV) to target optogenetic stimulation to the direct or indirect pathway, respectively. Indwelling fibers coupled to ferrules were implanted above the site of injection in the same surgery. After waiting 2-4 weeks for viral expression in dMSNs or iMSNs, mice were habituated to head-fixation on the trackball over several days. After mice were comfortable running freely on the trackball, recording experiments were performed.

Numerous previous studies have demonstrated that the MLR is co-extensive with the pedunculopontine and cuneiform nucleus, which we have confirmed previously as well. We therefore recorded from stereotaxic coordinates for the posterior pedunculopontine tegmentum (AP: -4.75, ML: +/- 1.5, DV: -3.6) (Figure 9a) and obtained single-unit recordings. We found that activity in a large proportion of neurons within the MLR/PPTg region were correlated with the

locomotor speed of the animals (Figure 9b). As a population, the majority of neurons were positively correlated with locomotion (Figure 9c) consistent with previous studies (Norton, Jo et al. 2011). This allowed us to classify units as either locomotor-active or locomotor-suppressed.

We then stimulated either dMSNs or iMSNs through the indwelling fiber optic while recording in the MLR/PPTg. We found that a large fraction of MLR neurons were activated by direct pathway activation. While there was a small population of neurons that was suppressed by direct pathway activity, the majority of neurons increased in firing. This pattern of activity largely mirrored the animals' behavior and is consistent with the idea that MLR and direct pathway activation can induce locomotion (Figure 9 A,C). In general, the activity of MLR neurons preceded the onset of locomotion. In contrast, we found that the activity of a large fraction of MLR neurons was suppressed by indirect pathway activation with only a small population that was activated (Figure 9 B,C). Once again, this decrease in the activity of the MLR population was consistent with the ability of indirect pathway to suppress the locomotor activity of animals.

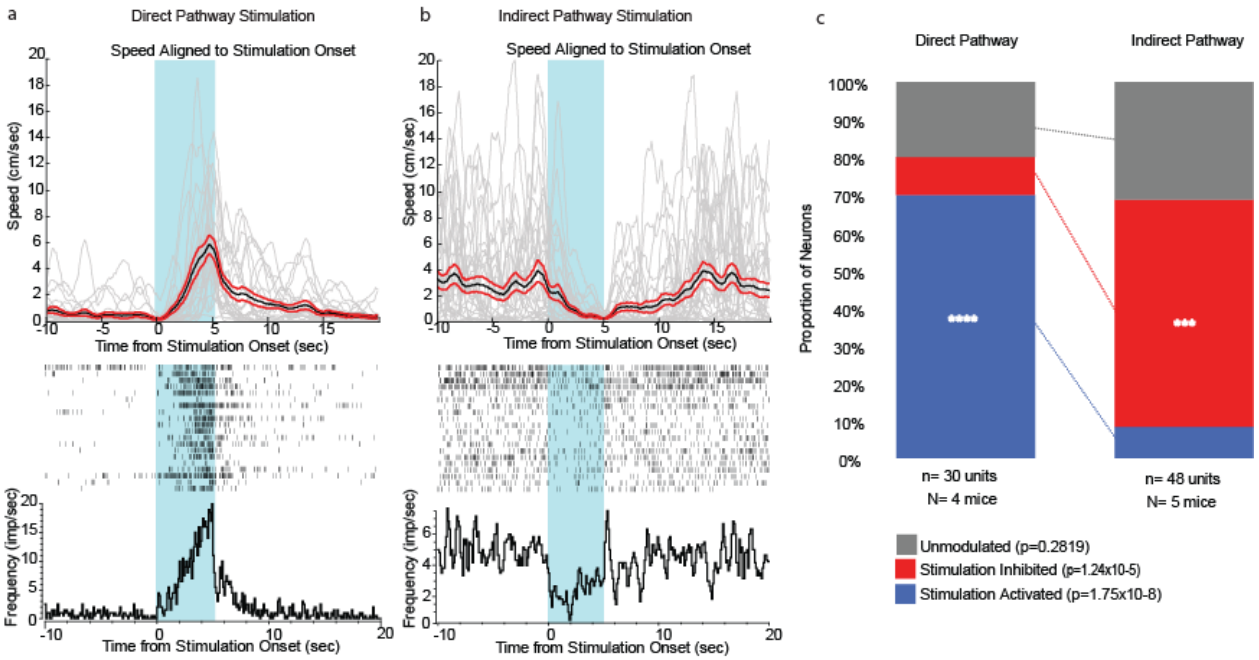

Figure 9. Modulation of locomotor speed and MLR activity by optogenetic stimulation of direct and indirect pathways of the basal ganglia.
(A) Example of locomotor speed following bilateral optical stimulation of direct pathway in a D1-Cre mouse after infection with AAV5-EF1a-DIO-ChR2-YFP (top) on a spherical treadmill. Response of a single unit in the MLR following unilateral stimulation of direct pathway (bottom). The blue shaded region represent the duration of the stimulation train. **(B)** Example of locomotor speed following optical stimulation of indirect pathway in an A2A-Cre mouse after infection with AAV5-EF1a-DIO-ChR2-YFP (top). Response of single unit in the MLR following stimulation of direct pathway (bottom). The blue shaded region represent the duration of the stimulation train. **(C)** Summary of proportion of MLR units activated (blue) / inhibited (red) / uncorrelated (grey) with direct pathway stimulation (left) or indirect pathway stimulation (right).

Discussion

Here, we demonstrate that the basal ganglia can bi-directionally regulate the activity of the MLR. Stimulation of the direct pathway can increase activity within the MLR while also promoting the initiation of locomotion. Indirect pathway activation can inhibit activity within the MLR while also suppressing locomotion. Together, these data indicate the MLR is bi-directionally regulated by basal ganglia outflow. Given direct pathway activation can increase the activity of neurons within the MLR and activation of the MLR itself is sufficient to initiate locomotion, these data provide evidence for a basal ganglia-brainstem pathway for initiating locomotion.

These data are consistent with the classical model of basal ganglia function where direct pathway activation can promote movement by suppressing basal ganglia output from the substantia nigra reticulate (SNr) and globus pallidus interna (GPi) to disinhibit downstream motor programs (Albin, Young et al. 1989; DeLong 1990; Gerfen 1992; Grillner, Hellgren et al. 2005). Conversely, indirect pathway activation ought to suppress motor behavior by enhancing basal ganglia tonic inhibition onto motor programs (Kang and Kitai 1990; Granata and Kitai 1991; Saitoh, Hattori et al. 2003; Grillner, Hellgren et al. 2005). Previous anatomical studies and electrophysiological studies have demonstrated that this region receives a dense projection from basal ganglia output nuclei arising from the SNr and EP. (Grofova I and Zhou M 1998). In turn, the ability of the MLR stimulation to initiate locomotion has been attributed to its descending excitatory pontomedullary inputs onto reticulospinal neurons (Grillner, Hellgren et al. 2005;

Grillner, Wallen et al. 2008), which can recruit spinal cord central pattern generators, mediating the bilateral alternating limb movements associated with gait (Grillner, Wallen et al. 2008). Indeed, evidence from decerebrate cat preparations provide strong evidence that only the descending projections from the MLR are required for locomotion in many species(Bjursten, Norrsell et al. 1976). In rodents, the corticospinal tract does not innervate motoneurons within the spinal cord in rodents, but only interneurons, it is largely believed that the corticospinal system may only be important for fine motor movements. While the organization of the motor system is likely to be very different in primates, there is a large body of evidence that central pattern generators under the control of brainstem locomotor centers may still mediate locomotion independent of the corticospinal tract (Lawrence and Kuypers 1968; Lawrence and Kuypers 1968). Thus, the tectospinal and reticulospinal systems in the brainstem are likely the key players in driving locomotor behaviors (Lawrence and Kuypers 1968; Hikosaka, Takikawa et al. 2000; Grillner, Hellgren et al. 2005). Together, our data indicate that the MLR may serve as a critical relay for conveying information from the basal ganglia to spinal cord pattern generators mediating goal-directed locomotion (Mena-Segovia, Bolam et al. 2004; Grillner, Hellgren et al. 2005; Grillner, Wallen et al. 2008).

Here, we demonstrate that MLR are neurons can be reliably driven and suppressed by basal ganglia pathways. This ability to drive firing is in sharp contrast to many theories that propose that the basal ganglia may merely gate motor behaviors. This theory largely stems from the belief that the main output nuclei of the basal ganglia is the motor ventral thalamus in which disinhibition

from basal ganglia nuclei is likely to merely gate the influence of concurrent excitatory input, but insufficient to drive the firing of thalamocortical neuron (Albin, Young et al. 1989). However, neurons within the MLR are tonically active in the absence of any exogenous inputs, differing dramatically in their intrinsic excitability from thalamocortical neurons. Because of their unique physiology, the coordinated release of basal ganglia tonic inhibition onto MLR neurons can be sufficient to drive these neurons to fire (Kang and Kitai 1990; Granata and Kitai 1991; Saitoh, Hattori et al. 2003). In this way, locomotion may represent a distinct motor behavior which is driven directly by a basal ganglia-brainstem system and less dependent on cortical influences.

Whether or not the MLR is the sole output of basal ganglia control over locomotion is still unclear. The substantia nigra reticulata (SNr) and medial globus pallidus (mGP) are the two major basal ganglia output nuclei in rodents. Both of these brain regions send efferents not only to the MLR, but also to the territories in which reticulospinal neurons reside (Kang and Kitai 1990; Granata and Kitai 1991; Saitoh, Hattori et al. 2003). Moreover, a second locomotor initiation program in the diencephalon around the region of the zona incerta is also believed to exist, which may represent a parallel route to access descending control of spinal cord central pattern generators via reticulospinal neurons (Grillner 2003; Grillner, Hellgren et al. 2005). These pathways are likely to work in parallel with the MLR to mediate basal ganglia control of locomotion. Also, additional efforts will be needed to demonstrate that it is the change in tonic inhibitory output from the basal ganglia, which mediates these changes in

locomotion. Alternatively, other recurrent pathways may also mediate the changes in activity within the MLR.

Moreover, while it is clear that activity within the MLR is sufficient to induce locomotion, it is unclear whether MLR activity is required for all forms of spontaneous voluntary locomotion. Inactivation of the region often does not affect baseline locomotion in rodents. Instead, these studies point to more subtle disruptions in Pavlovian approach for the receipt of rewards (Wanat, Bonci et al. 2013). These data are consistent with descriptions of the basal ganglia, which highlight its role in motivation and invigorating motor behaviors. It is also possible that the role of the MLR may differ from species to species. For example, the primate literature seems to suggest that lesions of the PPN, which is generally thought to be the primate homologue of the MLR, can induce akinesia (Munro-Davies, Winter et al. 1999; Nandi, Aziz et al. 2002), and case reports from patients suggest that bilateral PPN lesions can induce freezing of gait (Kuo, Kenney et al. 2008).

Implications for Parkinson's Disease

These results identify a neural pathway that may explain the gait dysfunction and postural instability that are key symptoms of many basal ganglia disorders such as Parkinson's Disease (PD). The loss of dopaminergic innervation to the striatum during Parkinson's disease is believed to increase the activity of the indirect pathway while suppressing activity in the direct pathway. The pathological imbalance of activity between these two pathways are predicted to increase inhibition onto motor output nuclei, mediating the cardinal symptoms of

the disorder. Two common symptoms of PD are freezing of gait (FOG) and postural instability, which are often described as axial symptoms. These axial symptoms can be in part explained by a model where excessive inhibition onto the MLR, prevents the initiation of gait and changes in posture that accompany locomotion. Here, we demonstrate that an enhancement of indirect pathway activity suppresses MLR activity and can mimic a Parkinsonian state consistent with the predictions of the classical model.

If loss of dopamine in the striatum increases indirect pathway activity relative to the direct pathway, our data suggest that the PPN/MLR of patients should be hypoactive. Direct activation of the PPN may therefore reverse this pathological underactivity. Recently, low frequency deep brain stimulation (DBS) in the PPN has been utilized as an experimental therapy for PD to treat the axial symptoms of PD that are not usually affected by high frequency DBS in more traditional sites such as the subthalamic nuclei and globus pallidus (Hamani, et al. 2011; St. George, et al. Neurology. 2010). PPN DBS alone—or in combination with more traditional sites—may have significant benefits for PD patients, particularly for those with gait deficits and postural problems (Mazzone, et al. 2005; Plaha and Gill, et al. 2005; Stefani, et al. 2007; Moro, et al. 2010; Stefani, et al. 2007; Follett and Torres-Russotto 2012). Beyond PD, PPN DBS has also been proposed as a treatment in other parkinsonian syndromes such as progressive supranuclear palsy (PSP) and multiple system atrophy (MSA) associated with postural instability and freezing of gait (Hamani et al., 2011). Our data argue that DBS within the PPN may serve to correct a pathological hypoactivity of the PPN in PD

mediated by the indirect pathway, rather than merely superimposing an exogenous activation onto the brainstem locomotor system in PD patients.

One concern regarding PPN DBS is its variable degree of effectiveness in different patients (St George, Nutt et al. 2010; Stefani, Peppe et al. 2013). This has been attributed to the variability in targeting the PPN by various clinical sites (Stefani, Peppe et al. 2013). However, it may also be the result of confusion arising from the heterogeneity of cell-types within the PPN. The boundaries of the PPN are defined by the cholinergic neurons within the midbrain tegmentum. However, only ~30% of neurons within this region are cholinergic. The largest population of neurons is actually the glutamatergic neurons, which compose of ~40% of neurons within the nucleus, with GABAergic neurons make up the remaining ~30% (Wang and Morales 2009). It still remains unclear whether activation of the cholinergic neurons or some other neuronal subtype can drive locomotion and hence the benefits observed in PD patients (Thankachan, et al. 2012). If the some other neuronal subtype were to drive locomotion, then these neurons would be a better indicator of the boundaries of the functionally defined MLR and may mediate the therapeutic benefits observed with PPN DBS. Thus, better defining the cell-type within the MLR that mediates locomotion and how it interacts with the basal ganglia will improve our understanding of the mechanisms of PPN DBS and the pathophysiology underlying the axial symptoms of PD.

Potential Basal Ganglia Regulation of Behavioral and Brain States

In addition to its role in locomotion, the MLR/PPTg region has been attributed with numerous other functions that can be mapped onto its variety of efferents. We have demonstrated that the MLR inputs to the basal forebrain can dramatically regulate cortical state, enhancing the evoked visual responses in cortex and shifting LFP power from low frequency oscillations toward the gamma ranges. Other studies have implicated the MLR/PPTg in generating theta oscillations within the hippocampus through its direct connections to the medial septum and the posterior lateral hypothalamus (Pignatelli, Beyeler et al. 2012; Buzsaki and Moser 2013), and shifting the mode activity in the thalamus from burst to tonic mode of firing through its cholinergic projections to the thalamus (Lu, Guido et al. 1993). The PPTg is also the densest glutamatergic input into the dopamine neurons of the substantia nigra pars compacta (Good and Lupica 2010), and the cholinergic input is required for transitioning dopamine neuron firing from a burst to tonic mode of activity (Drenan, Grady et al. 2008; Patel, Rossignol et al. 2012). Interestingly, many of these changes in brain state that are observed in responses to PPTg stimulation are the exact same types of changes that have been reported during transitions from a stationary state to active locomotion. Thus, we propose that a unifying theory of the MLR/PPTg function is that this region primarily serves to regulate locomotion and all of the brain state changes associated with locomotion. However, as we previously described (See Chapter 2), these brain state changes in various brain regions can be dissociated from locomotion under special circumstances, and may also be recruited in states of “alertness” or during specific phases of sleep such as REM sleep.

We hypothesize that the basal ganglia, in addition to regulating locomotion and motor behaviors, will also bi-directionally regulate transitions between processing states across the brain as well. In this way, we would expect that direct pathway activation should have powerful effects on modulating the state of the cortex, hippocampus, thalamus, and dopaminergic midbrain while indirect pathway should have opposing effects. The regulation of brain states perhaps points to a less recognized function of the basal ganglia in regulating behavioral states associated with alertness, sensory gating, and sleep.

The existence of such a pathway potentially explains a large set of experimental findings (Lazarus, Chen et al. 2013). For example, bilateral lesions of the dorsal striatum are known to reduce total time spent awake (Qiu, Liu et al. 2012). Lesions in the downstream GPe lead to insomnia in rodents whereas SNr lesions enhance wakefulness (Qiu, Vetrivelan et al. 2010). The striatum contains the highest density of dopamine and adenosine receptors in the brain. Excessive sleepiness in PD and other sleep disorders, such as narcolepsy are commonly treated with modafanil, which acts by enhancing extracellular levels of dopamine through a process that is D1 or D2 receptor-mediated (Qu, Huang et al. 2008). Adenosine receptors are known to promote sleep by acting through the A_1 and A_{2A} receptors. While little is known about the action of the A_1 receptor on striatonigral neurons, studies suggest that activation of the A_{2A} receptor on indirect pathway striatal neurons (in fact we use the A_{2A}-Cre line to selectively target indirect pathway neurons) can induce sleep (Huang, Qu et al. 2005).

Likewise, caffeine, which acts as an antagonist for the A_1 and A_{2a} receptors is known to promote wakefulness (Huang, Qu et al. 2005).

The regulation of the MLR/PPTg by the basal ganglia can potentially explain the sleep disturbances and deficits in alertness from which many patients with PD suffer. The classical model of basal ganglia function proposes that the loss of dopamine in PD ought to increase activity in the indirect pathway and suppress activity in downstream basal ganglia outputs. If the MLR is associated with regulating brain states associated alertness in addition to its control over locomotion, then we would predict that PD patients ought to have some form of deficit in behavioral states. Indeed, it has recently been reported that there is an unprecedentedly high prevalence of parkinsonism and cognitive impairment among a population of patients, who experience a disorder known as REM-behavior disorder (Boeve 2013). The PPTg/MLR region is thought to be both required and sufficient for the maintenance for both the electophysiological EEG correlates and atonia associated with REM sleep (Saper, Fuller et al. 2010). In fact, because the onset of REM-behavior disorder is earlier and strongly predictive of the development of PD, clinicians have been advised to screen patients with the sleep disorder for parkinsonian symptoms regularly. Consistent with a role for the MLR/PPTg in arousal, PD patients with low frequency DBS in the PPN/MLR region often self-report that they feel alert upon PPN stimulation and an improvement in cognitive function. High frequency DBS can induce a non-REM sleep state immediately after stimulation (Arnulf, Ferraye et al. 2010). There have also been documented disruptions in sleep architecture in patients with frequent nocturnal awakenings, increased wakefulness after sleep onset,

and a remarkable reduction of REM sleep time (Romigi, Placidi et al. 2008; Alessandro, Ceravolo et al. 2010). Patients being treated with DBS in the PPN experience a normalization of sleep structure with increased time spent in REM sleep (Romigi, Placidi et al. 2008; Alessandro, Ceravolo et al. 2010). Thus, investigating the basal ganglia's interactions with the MLR has the possibility of revealing understudied functions of the basal ganglia.

Materials and Methods (chapter 3)

Mice

Animals were maintained in the animal facility at the University of California, San Francisco and used in accordance with protocols approved by the University of California, San Francisco Institutional Animal Care and Use Committees.

Construct and virus preparation.

Plasmids encoding the DNA sequences for pAAV- EF1α-DIO-CHR2(H124R)-eYFP and EF1α-DIO-eYFP cassettes were packaged in AAV vectors and serotyped with AAV5 coat proteins by the viral core at University of North Carolina. The final concentration was 1–2 × 10^{12} viral particles per ml.

Implantable chronic optical fibers and optic cables construction.

Optical stimulators were constructed according to published protocols[51]. Briefly, optical fibers were constructed by attaching a 200-μm, 0.37 NA optical fiber (Thor Labs) with epoxy resin into a metal ferrule that had previously been cut and scored. Fiber-ferrule units were then cut and polished. Only implants with efficiency greater than 70% and comparable efficiencies (±10%) were used. Optical-patch cables were constructed from 62.5-μm core diameter optic fiber (Thor labs) that were connected to a ferrule on one end and a fiber-optic connector for physical contact (FCPC) connector on the other end. Cables were covered in furcation tubing to protect the fiber and to prevent light from escaping through the optic-patch cord. The ferrule at the end of the optic patch cord was

fitted with a zirconium sleeve to interface with the chronic implant. The FCPC connector was coupled to a 473-nm diode-pumped solid-state laser (200 mW). The laser driver current was adjusted to yield 20-mW output from the patch cable.

Stereotaxic AAV injection and Implantation of Optical Stimulator

Animals were anesthetized with either ketamine (150 mg per kg of body weight) and xylazine (50 mg per kg) or 2% isofluorane (vol/vol) gas anesthesia. For MLR injections, animals were placed in a stereotaxic frame and 26-gauge microinjection needles were inserted through a burr hole bilaterally into the MLR (coordinates from bregma: -4.75 mm anterior, ±1.2 mm medial-lateral, –3.6 ventral). 0.5 µl injections were performed using a 1-µl Hamilton syringe through a hydraulic pump (Harvard Instruments) and took place over 5 min followed by 5 min of recovery. For striatal injections, animals were placed in a stereotaxic frame and 26-gauge microinjection needles were inserted through a burr hole bilaterally into the dorsomedial striatum (coordinates from bregma: 0.75 mm anterior, ±1.5 mm medial-lateral, –2.5 ventral). 1 µl injections were performed through a hydraulic pump (Harvard Instruments) and took place over 10 min followed by 10 min of recovery.

For implantation of optical stimulator, The length of the optic fiber protruding from the implant was cut to be 2 mm. The tip of the fiber optic from the implant was then inserted through the same burr hole as was used previously for the virus injection and was lowered 2 mm ventral to the dura. The implant was cemented

to the skull using dental cement. Mice were monitored until recovery from surgery and then returned to their home cage where they were housed individually.

Acute Optical Stimulation

At least one month after injection of the MLR, a burr hole was drilled above the injection site. An optical implant constructed from an optic fiber threaded through a ferrule. was lowered into the burrhole. The tip of a 200-μm fiber optic was lowered until the probe was 0.1mm above the site of viral injection.

Extracellular multisite electrophysiology.

Electrophysiological recordings were performed on adult C57Bl/6 mice (age 2–6 months, both male and female). In preparation for recording, mice were anesthetized with a surgical level of isoflurane (4% induction, 2% maintenance, in O2). After a scalp incision, the fascia was cleared from the surface of the skull, and a metal headplate was affixed to the skull with vetbond and dental acrylic (Niell and Stryker, 2010). The headplate provided stability for mounting the mouse, and an opening to allow access to the skull.

On the day of recording in the MLR, a small craniotomy of was performed over the MLR, at 1.2 mm lateral and 4.75 mm posterior to bregma. The exposed craniotomy was covered with 2% agarose in 0.9% saline to prevent drying and provide mechanical support.Recordings were made with silicon multisite electrodes (a1×32–25-5 mm-177, NeuroNexus Technologies). Recordings in the MLR were performed with a small amount of the lipophilic vital dye DiI or DiO (Invitrogen). The electrode was inserted through the craniotomy and overlying

agarose using a microdrive (Siskiyou Designs), and lowered down to the MLR. The MLR contained large units with narrow waveforms. Beyond stereotaxically localizing the MLR, the electrode was placed without regard for the presence of responsive units on individual sites, and all units stably isolated over the recording period were included in analysis. Upon locating the MLR, the electrode was further embedded in agarose to increase mechanical stability, and the electrode was allowed to settle in one position for 30 min to obtain stable single-unit recordings.

At the end of recording, the animal was killed under deep anesthesia by cervical dislocation. The brain was removed immediately and fixed by immersion in 4% paraformaldehyde in PBS at 4°C overnight, after which 200 μm coronal sections were cut with a vibratome. The sections were mounted using Vectashield with DAPI (Vector Laboratories) and imaged to reconstruct optical fiber penetrations into the MLR.

Data acquisition.

Data acquisition was performed as described by Niell and Stryker (2008). Signals were acquired using a System 3 workstation (Tucker-Davis Technologies) and analyzed with custom software in Matlab (MathWorks). To obtain single-unit activity, the extracellular signal was filtered from 0.7 to 7 kHz and sampled at 25 kHz. Spiking events were detected on-line by voltage threshold crossing, and a 1 ms waveform sample on 4 neighboring recording sites was acquired around the time of threshold crossing. Single-unit clustering and spike waveform analysis were performed as described previously (Niell and Stryker, 2008), with a

combination of custom software in Matlab and Klusta-Kwik (Harris et al., 2000). Quality of separation was determined based on the Mahalanobis distance and L-ratio (Schmitzer-Torbert et al., 2005) and evidence of a clear refractory period. Units were also checked for stability in terms of amplitude and waveform over the course of the recording time to ensure that they had not drifted or suffered mechanical damage. Units that were found by histology to be outside the LGN, generally above or below resulting from the length of the electrode, were excluded from subsequent analysis.

Data analysis.

Data analysis was performed using custom routines written in Matlab. Statistical significance was determined by Mann–Whitney U test, except where otherwise stated. For figures representing the median of data, error bars indicate SE of the median as calculated by a bootstrap. In other cases, error bars indicate SEM.

Chapter 4

Transient Stimulation of Distinct Subpopulations of Striatal Neurons Mimics Changes in the Action Value for a Locomotor Decision

Abstract

In changing environments animals must adaptively select actions to achieve their goals. In tasks involving goal-directed action selection, striatal neural activity has been shown to represent the value of competing actions. Striatal representations of action value could potentially bias responses toward actions of higher value. However, no study to date has demonstrated the direct impact of distinct striatal pathways in goal-directed action selection. Here we show in mice that transient optogenetic stimulation of dorsal striatal dopamine D1 and D2 receptor expressing neurons during decision-making introduces opposing biases in the distribution of choices. The effect of transient stimulation on choice is dependent on recent reward history and mimics an additive change in the action value. While stimulation prior to and during movement initiation produces a robust bias in choice behavior, this bias is significantly diminished when stimulation is delayed after response initiation. Together, our data demonstrate the role of striatal activity in goal-directed action selection.

Introduction

Animals are faced with the challenge of optimally selecting actions in changing environments. Theoretically, these decisions can be implemented by estimating the value of different actions and then choosing the action of greatest value. Recently, the striatum has been implicated as an important neural structure that may mediate this process of action selection. Electrophysiological recordings in primate and rodent striatum have identified signals correlating with the expected outcomes of actions and measures of motivation for particular responses(Lauwereyns, Watanabe et al. 2002; Cromwell and Schultz 2003). Other studies have shown that the striatum encodes representations of the value of actions in free choice tasks(Samejima, Ueda et al. 2005; Lau and Glimcher 2008). Striatal activity also parallels the learning of rewarded responses based upon previous experience(Barnes, Kubota et al. 2005; Pasupathy and Miller 2005) and is essential for the acquisition and execution of goal-directed behaviors(Balleine, Delgado et al. 2007).

Two dorsal striatal pathways have been proposed to mediate opposing influences on the selection of actions(Albin, Young et al. 1989; DeLong 1990; Gerfen 1992). Activity within the ‘direct pathway’ has been shown to facilitate actions while activity in the ‘indirect pathway’ has been demonstrated to inhibit behaviors(Kravitz, Freeze et al. 2010; Lobo and Nestler 2011). The direct pathway is comprised of medium spiny neurons (MSNs) that express the dopamine D1 receptor (D1R) while the indirect pathway is comprised of MSNs that express the dopamine D2 receptor (D2R)(Kreitzer and Malenka 2008;

Shuen, Chen et al. 2008; Kreitzer and Berke 2011). It has been proposed that the D1R and D2R-expressing MSNs encode a representation of the value of actions and generate a response bias towards actions of higher value(Hikosaka, Nakamura et al. 2006; Isoda and Hikosaka 2011). While this model is consistent with numerous studies, other studies seem to suggest that these populations of neurons may merely play a permissive role in learning associations(Brainard and Doupe 2000) or only modulate the vigor of actions without affecting the animals' choice behavior(Desmurget and Turner 2010; Turner and Desmurget 2010). No study to date has definitely demonstrated the role of these two populations of neurons in the context of reward-based decision-making(Kreitzer and Berke 2011).

To directly investigate the role of striatal activity in action selection, we created a probabilistic switching task in which mice chose to enter a reward port located to their left or right side. To inform their choices in the task, mice relied on recent reward history to assess whether water will be delivered from one of two reward ports. Critically, we selected a left versus right choice design based upon previous studies showing that 1) unilateral striatal manipulations can affect lateralized body movements(Schwarting and Huston 1996; Kravitz, Freeze et al. 2010), 2) neurons within the striatum encode actions for movements contralateral or ipsilateral to the recording site(Kim, Sul et al. 2009; Kubota, Liu et al. 2009; Thorn, Atallah et al. 2010), and 3) brainstem motor programs mediating orienting and approach behaviors(Felsen and Mainen 2008) are regulated by the basal ganglia(Grillner, Hellgren et al. 2005; Hikosaka, Nakamura et al. 2006; Felsen and Mainen 2008; Grillner, Wallen et al. 2008). From these studies, we reasoned

that unilateral manipulations to each striatal hemisphere could affect the selection of spatially lateralized responses in the context of this task. We developed a computational model that assigns a value for each action given a particular history of rewards to predict the distribution of left versus right choices. We then used optogenetic techniques to examine the impact of transient unilateral stimulation of the D1R- and D2R-expressing striatal neurons during an epoch in the task when animals were choosing to approach a left or a right port.

Here, we show that transient activation of D1R-expressing striatal neurons biased choices toward the port contralateral to the side of stimulation and transient activation of D2R-expressing striatal neurons biased choices toward the ipsilateral port. However, rather than giving rise to a stereotyped and consistent motor response, optical stimulation induced a bias in the likelihood of choice responses that was dependent on the previous history of rewards for each choice. In the context of our model, we show that optical stimulation mimics a fixed additive shift in action value(Sugrue, Corrado et al. 2005; Gold and Shadlen 2007). The effect of stimulation only was effective in a narrow temporal window prior to and during the early initiation of movements. Striatal activation also changed the latency to movement initiation based upon the relative action value. Together, these data are consistent with the hypothesis that the striatum encodes the value of actions and demonstrates that activity within these pathways can be used to bias the selection and vigor of competing goal-directed actions.

Results

To study action selection during decision making in mice we trained adult BAC transgenic mice (n=28) expressing Cre recombinase under either dopamine D1R or D2R regulatory elements on a spatial two-alternative forced-choice probabilistic switching task. In brief, the task required animals to initiate a trial with a nose poke into a central port followed by movement to a left or a right port to obtain reward. The rewarded port delivered a water reward for 75% of correct responses, and this port was periodically switched across blocks. The length of each block was randomly distributed between 7-23 trials and the switch only took place after a rewarded trial (**Fig. 10a**). A 'Go' cue signaled when mice could approach either choice port. Thus, the only information provided to guide the animals' choice behavior was the expectation of reward based upon the outcome of previous trials. After initial training, mice took on average 2.20±0.04 (right to left) and 2.22±0.05 (left to right) trials (± s.e.m., n=28) to switch their behavioral responses following reversals in action-outcome contingencies between blocks (**Fig. 10b**). When we analyzed the effects of reward history in the previous two trials on upcoming choice, we found that mice implemented a 'win-stay, lose-shift' strategy in which rewards served as evidence to continue responding at a port and the lack of reward served as evidence to switch (**Fig. 10c**). Animals' choice probability generally tracked the reward probability at each port for various reward histories (**Fig. 11a**).

Estimating the value of actions based upon previous reward history

Next, we generated a quantitative model to describe the animals' behavior in the

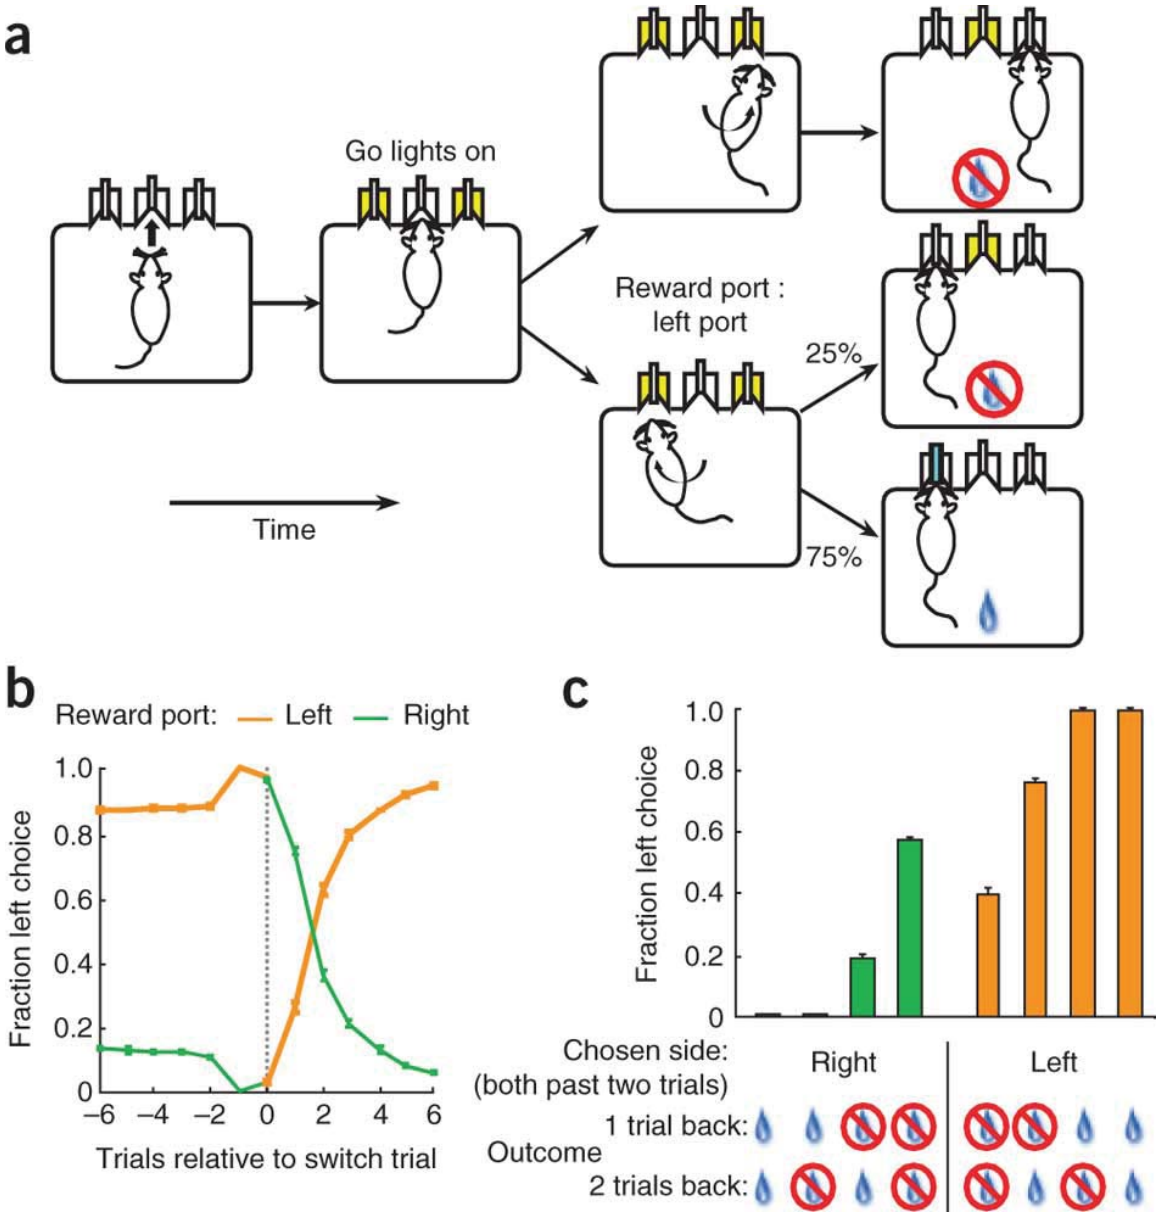

Figure 10: Design of the probabilistic switching task and mouse performance.
(**a**) Sequence of events in a probabilistic switching task. Mice learned to initiate a trial at the center port and to choose a left or right peripheral port for water reinforcement. Only one peripheral port was rewarded at a time. In 25% of trials, neither port was rewarded. The rewarded port was switched only after a rewarded trial. (**b**) Fraction of choices for left port (*n* = 28 subjects) for trials before and after a switch of the rewarded port (at trial 0). (**c**) Fraction of choices for the left port from one subject for reward histories in which two consecutive choices to either the left or the right port were made during the previous two trials. Data from mixed choice histories are not shown for brevity. All error bars represent s.e.m.

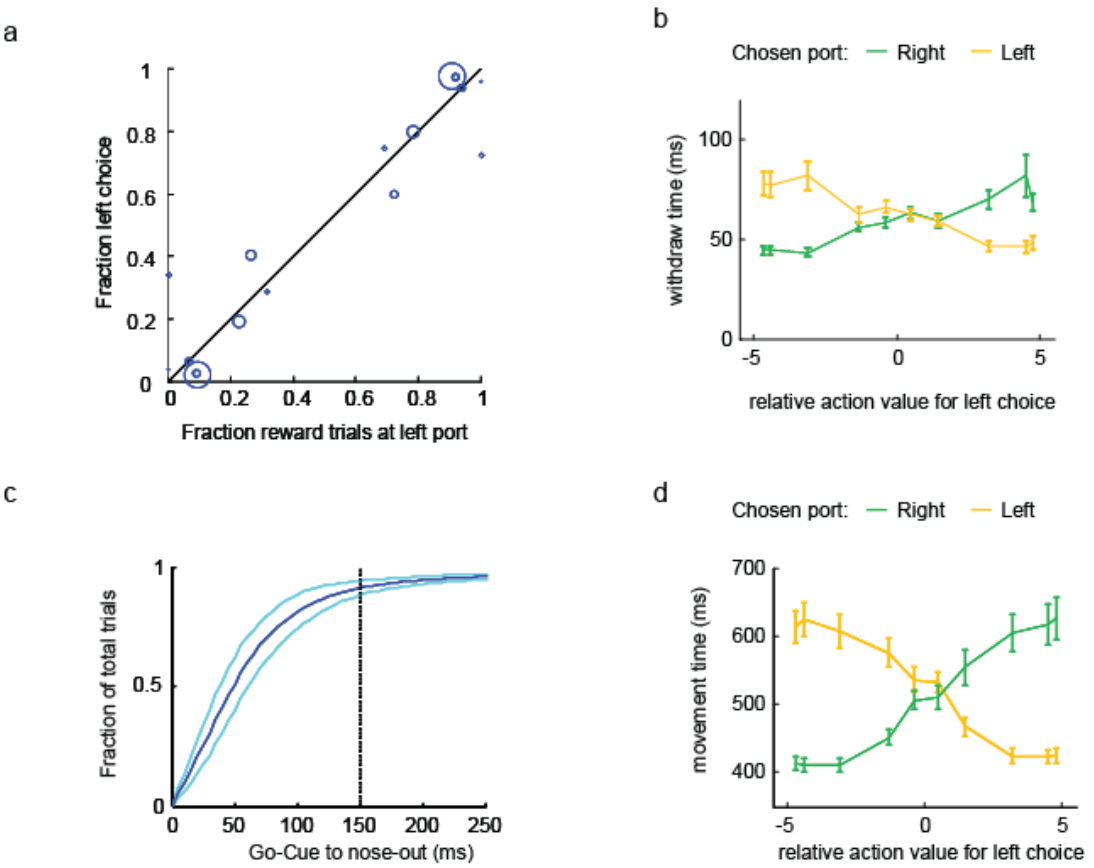

Figure 11. Characterization of responses within the task

(a) The relationship of fraction of left choices and fraction of reward trials at left port for all 16 possible reward histories in the previous two trials averaged across all subjects. The frequency of trials of a given reward history is indicated by the relative size of the circle. **(b)** Average median withdrawal time measured from go-signal to nose-withdrawal at center port (n=28). Withdrawal time is shorter when the action value for the chosen port is higher. **(c)** Cumulative distribution of reaction time from Go-Cue to nose-withdrawal from center port for all subjects (mean ±s.d., n=28). The dotted line indicates the onset of optical stimulation at 150 ms latency. **(d)** Average movement time from nose-withdrawal at center port to reward port (n=28). Movement time is shorter when the action value for the chosen port is higher. All error bars represent s.e.m.

task. Prior electrophysiological studies have found that the activity of striatal neurons correlates with estimates of trial-to-trial values of actions(Lau and Glimcher 2005; Samejima, Ueda et al. 2005; Lau and Glimcher 2008). These estimates of value assume the softmax decision rule(Sutton 1998), which describes the tendency of decision-makers to have more variable responses when the alternatives are more similar in value. In the case of two alternatives under the softmax rule, the contribution of previous reward history to the value of each action can be estimated by multivariate logistic regression(Lau and Glimcher 2005). We fit a regression model where the probability of choices at each port in the upcoming trial was determined by the animals' previous reward history (see methods).

The regression analysis demonstrated that the contribution of prior rewards declined with the passage of trials (**Fig. 12a**). Rewards within the previous three trials had a significant effect on choices in the upcoming trial, serving as evidence for the animal to stay at the rewarded port as indicated by the positive regression coefficient (**Fig. 12a**) while the lack of reward at the chosen port in the previous trial significantly promoted switching in the following trial as represented by the negative regression coefficient (**Fig. 12a**). The large regression coefficients for the previous two trials validated our initial attempts to analyze choice behavior by segregating trials based upon the animals' reward history in the previous two trials (**Fig. 10c**). Based upon the model, we generated dynamic estimates of action-values and probabilities for responding at each port based on the animals' recent reward history (**Fig. 12a,b,e**). These action value estimates were the sum of regression coefficients corresponding to the previous reward

history for each side (**Fig. 12a**) determining the distribution of choices of the subject (**Fig. 12b**). The choice probabilities predicted by an identically generated regression model using 70% of the data recapitulated the actual distribution of choices in the remaining 30% of the data, confirming the predictive validity of our model (**Fig. 12e,f).** .

Selective optogenetic stimulation of D1R-expressing and D2R-expressing striatal neurons biases choice

To independently study the activity of D1R-expressing and D2R-expressing striatal neurons in our task, we bilaterally injected an adeno-associated virus (AAV) into the dorsal striatum that enabled Cre-dependent viral expression of channelrhodopsin (ChR2) and yellow fluorescent protein (eYFP) in transgenic mice expressing Cre under regulatory elements for either the dopamine D1R or D2R. We then chronically implanted an optic fiber in or just above the dorsal medial striatum of each hemisphere. Our Cre-dependent strategy targeted ChR2-expression to the direct or indirect pathways of the basal ganglia (**Fig. 13**). In D2-cre mice, 37.7% of putative MSNs expressed eYFP (consistent with indirect pathway targeting) and additionally 38.7% of the choline acetyltransferase (ChAT)-expressing neurons co-expressed ChR2-eYFP (**Fig. 14**). Estimates suggest 80-97.7% of neurons in the striatum are MSNs and 0.3-2% are cholinergic(Kawaguchi, Wilson et al. 1995; Rymar, Sasseville et al. 2004). We therefore estimate in D2-Cre mice we transduced a ratio of MSNs to cholinergic neurons of roughly 50:1. By simultaneously recording and optically stimulating striatal neurons, we were able to confirm optically driven striatal neuronal firing in our system was time-locked to

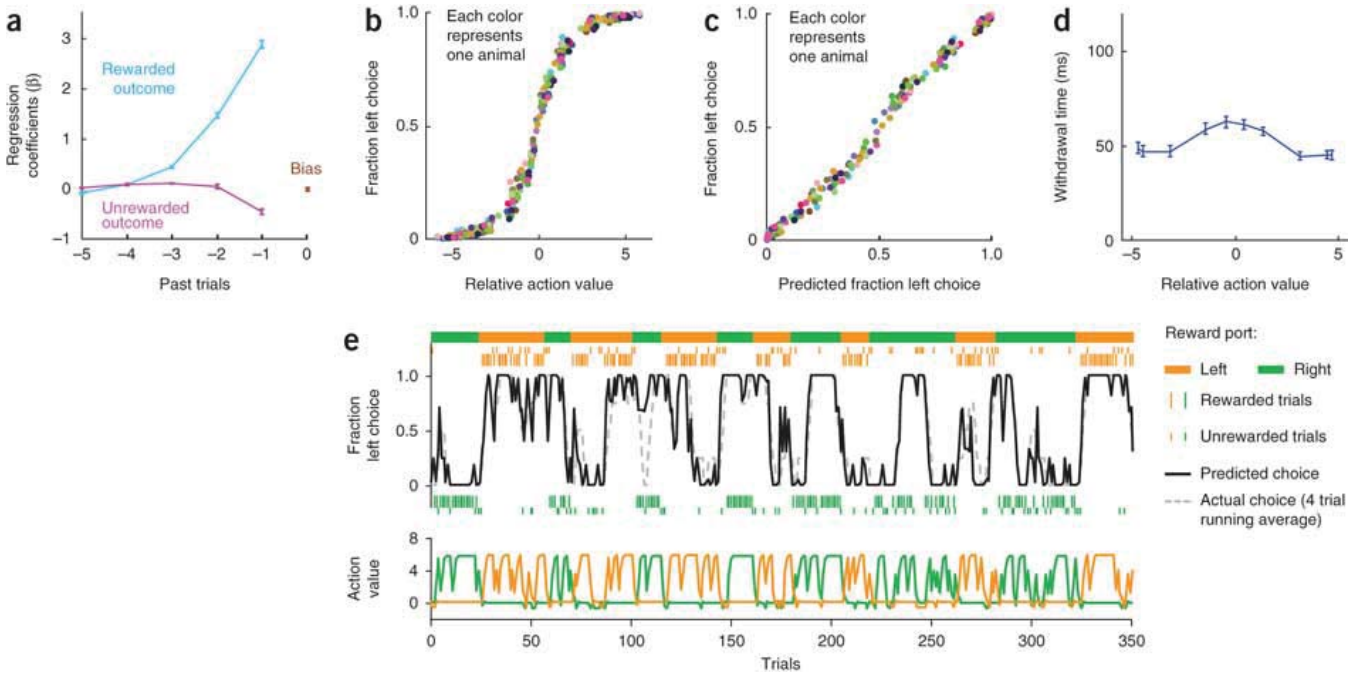

Figure 12: Modeling action value and choice.
(**a**) Contributions of intrinsic bias (brown) and rewarded (cyan) and unrewarded (magenta) outcomes in the previous five trials on choices in the current trial derived from logistic regression (*n* = 28 subjects, 22,726 ± 2,881 trials per subject). (**b**) The fraction of choices for the left port plotted against the relative action value. Data from each subject were grouped into ten bins and represented in a distinct color. (**c**) The actual fraction of choices for the left port plotted against the fraction of left choices predicted by the regression model. (**d**) Average median withdrawal time measured from Go light signal to nose withdrawal from center port (*n* = 28). Withdrawal time was shorter when the relative action value for either port was higher. (**e**) Example data from 14 trial blocks. Top, right reward blocks are represented in green and left reward blocks in orange. The dashed line indicates the subject's probability of choosing the left port averaged across four trials and the black line indicates the predicted probability of choice on the basis of action value estimates (bottom). Long ticks correspond with rewarded trials and short tics represent unrewarded trials. All error bars represent s.e.m.

Past	$\beta_{Intrinsi}$	β_{Reward}	$\beta_{no\text{-}Reward}$
1	-	49.06 ± 1.5 %	8.25 ± 0.79
2	-	24.3 ± 0.52 %	2.85 ± 0.43
3	-	6.95 ± 0.31 %	1.53 ± 0.15
4	-	1.21 ± 0.20 %	1.32 ± 0.20
5	-	1.66 ± 0.23 %	0.79 ± 0.13
sum	2.09	83.17 %	14.74 %

Supplementary table 1.
Relative contribution of each regression coefficient as a proportion of the unsigned sum of all regression coefficients. (mean ± s.e.m. , n=28)

Figure 13
Supplementary Figure 2

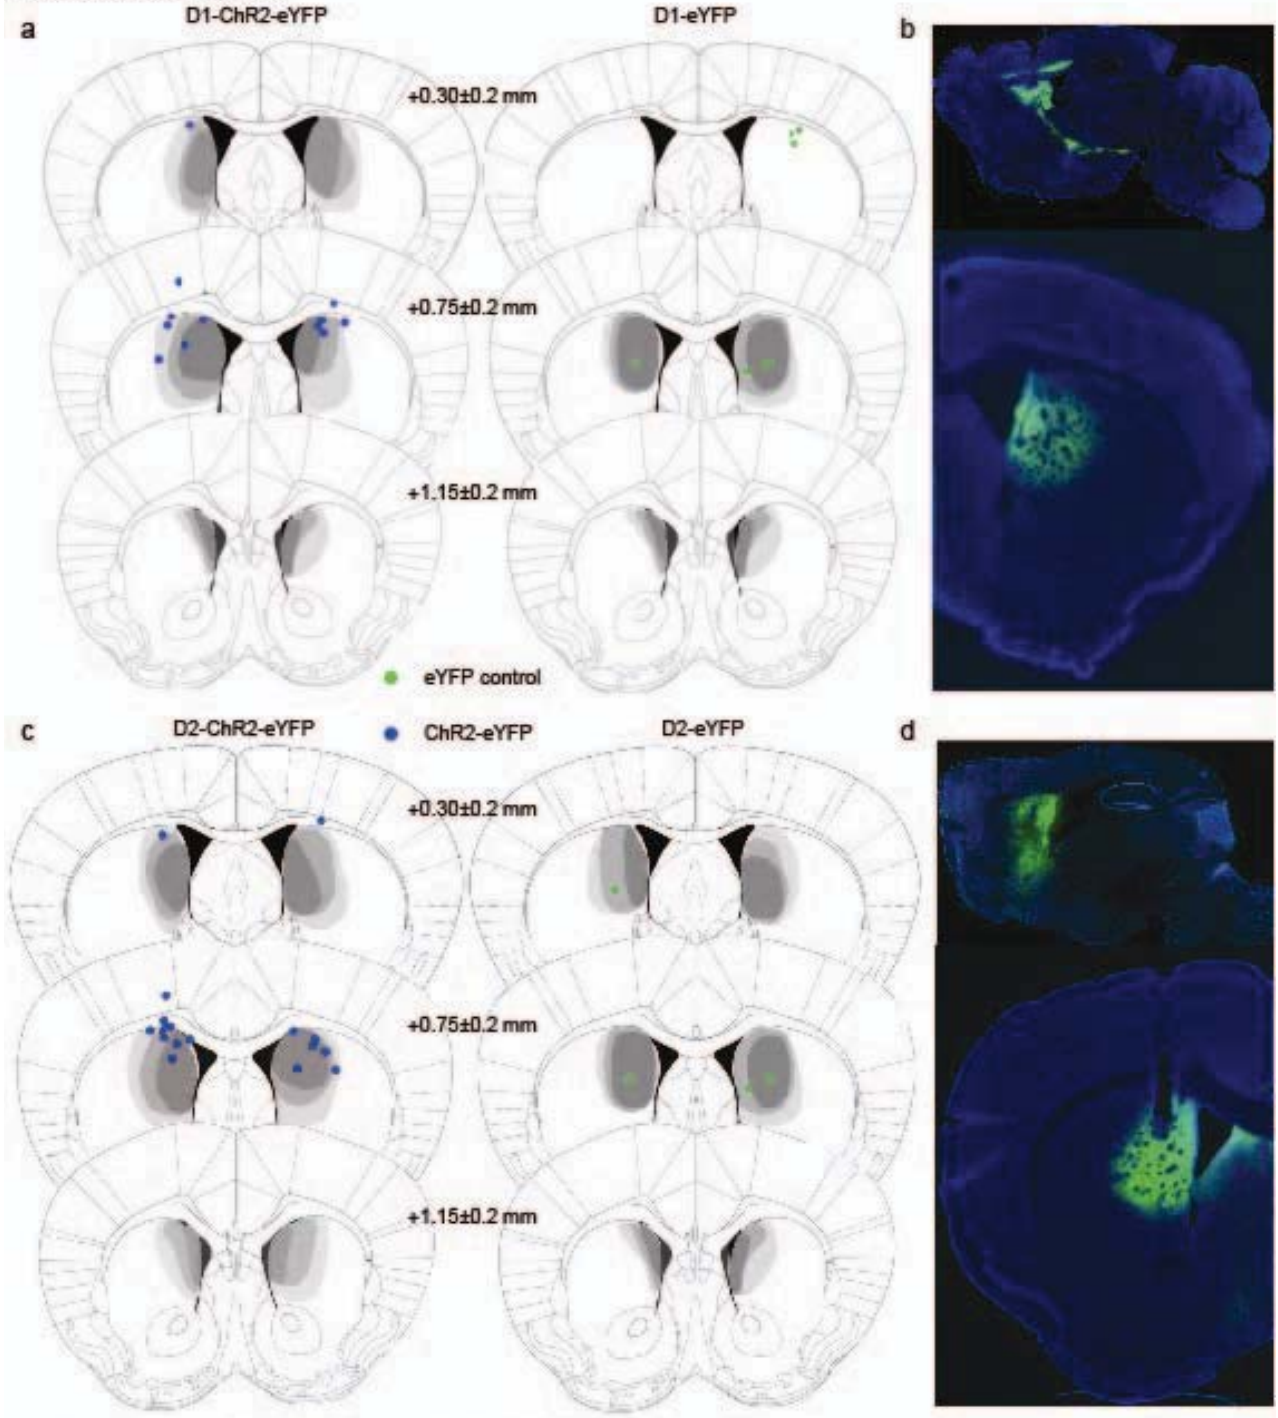

Supplementary Figure 2. Anatomy of stimulation sites

(a) Coronal series demonstrating the extent of infection (grey) and placements of fiber optic tips (dots) for D1-Cre injected with AAV2/5-EF1-DIO-ChR2-eYFP (left series) or AAV2/5-EF1-DIO-eYFP (right) series. Light grey represents the extent of the largest injection while dark grey represents the smallest extent. (b) Coronal (top panel) and sagittal histological sections (bottom panel) from two representative D1-Cre animals expressing ChR2-eYFP. (c) Coronal series demonstrating the extent of infection (grey) and placements of fiber optic tips (dots) for D2-Cre injected with AAV2/5-EF1-DIO-ChR2-eYFP (left series) or AAV2/5-EF1-DIO-eYFP (right) series. (d) Coronal (top panel) and sagittal histological sections (bottom panel) from two representative D2-Cre animals expressing ChR2-eYFP.

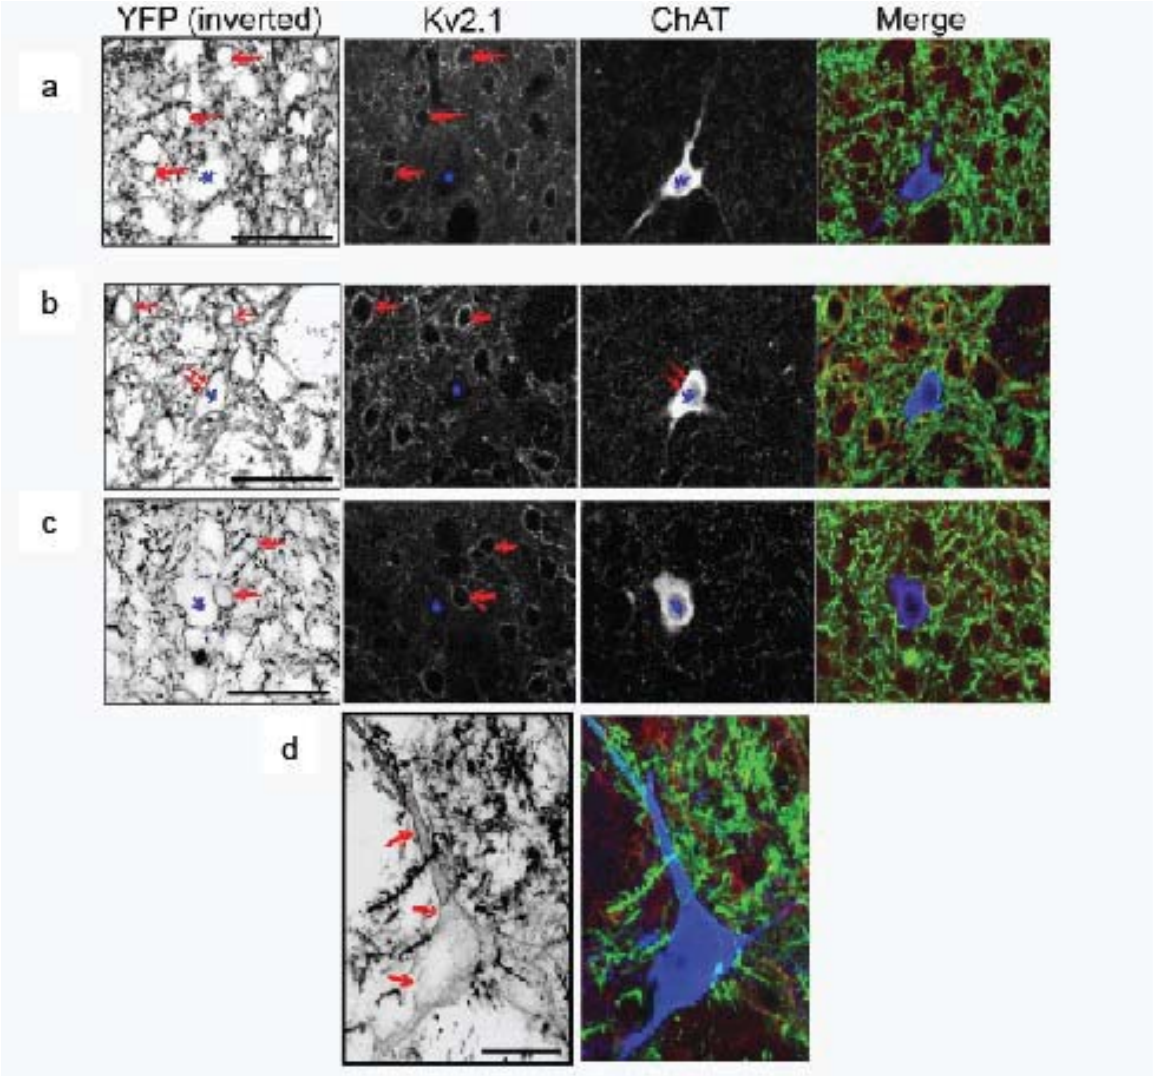

Figure 14. Expression of ChR2 in Striatal MSNs and ChAT+ Interneuons
(a) Single plane confocal images of medium spiny neurons and cholinergic interneurons in the dorsomedial striatum from a D1-Cre animal injected with AAV-EF1α-DIO-ChR2-eYFP (green in merged image). Histological slices were labeled using immunohistochemistry for the intracellular C-terminus of Kv2.1 channel, a marker of MSNs, 49 and choline acetyltransferase (ChAT, blue in merged image) respectively. ChR2-eYFP expression in D1-Cre animals was found to colocalize with Kv2.1 expression as indicated by the red arrows in panel (281/476) **(a)**, but never with ChAT positive neurons (0/68 neurons). **(b)** Single plane confocal images of medium spiny neurons and cholinergic interneurons in the dorsomedial striatum from a D2-Cre animal injected with AAV-EF1α-DIO-ChR2-eYFP (green in merged image). Histological slices were labeled using immunohistochemistry for the intracellular C-terminus of Kv2.1 channel, a marker of MSNs, 49 and choline acetyltransferase (ChAT, blue in merged image) respectively. ChR2-eYFP expression in D2-Cre animals was found to colocalize with Kv2.1 expression (430/1141 neurons) as indicated by the red arrows in many cells (panel b and c). In D2-Cre animals, ChR2-eYFP expression was found to colocalize with a subset of ChAT immunostained neurons (29/75) (blue dot, panel **b**), but not all ChAT immunostained neurons (blue dot, panel **c**) despite the presence of expression of ChR2-eYFP in nearby cells as seen by 3D-reconstruction **(d)**.

stimulation (**Fig. 15**). Optically evoked activity could be due to direct ChR2 activation or potentially indirect activation of neurons downstream of ChR2-expressing cells.

We then sought to determine whether activation of specific neural populations in the striatum could affect the choice behavior of animals(Salzman, Britten et al. 1990). Optical stimulation was delivered at a decision point within the task when the 'Go' cue signals the animal to make their choice (**Fig. 16a**). At this point in the task, the choice ports are equidistant in egocentric space to the left and right of the mouse, and this period marks the time when animals must select and initiate an action to acquire a potential reward (**Fig. 10a**). Stimulation was delivered at 5 Hz, 10 Hz, or 20 Hz for 500 ms in 6% of total trials interspersed at random. The stimulation parameters used here reflect physiologically relevant activity found in awake, freely moving striatal recordings in mice(Kubota, Liu et al. 2009; Jin and Costa 2010). The presence of stimulation and non-stimulation trials in the same session allowed us to make highly controlled within subject comparisons to determine the effect of stimulation. Stimulation sessions were interspersed with training sessions without stimulation, and the hemisphere that was stimulated alternated with every stimulation session. There were no differences in animals' choice behavior across sessions with identical stimulation parameters, therefore, data from these sessions were pooled for analysis. To rule out any possible effect of the optical stimulation or viral expression, we also injected animals with a control virus expressing Cre-dependent eYFP. These control animals were subject to identical training regimens and stimulation parameters as experimental animals.

Supplementary Figure 4

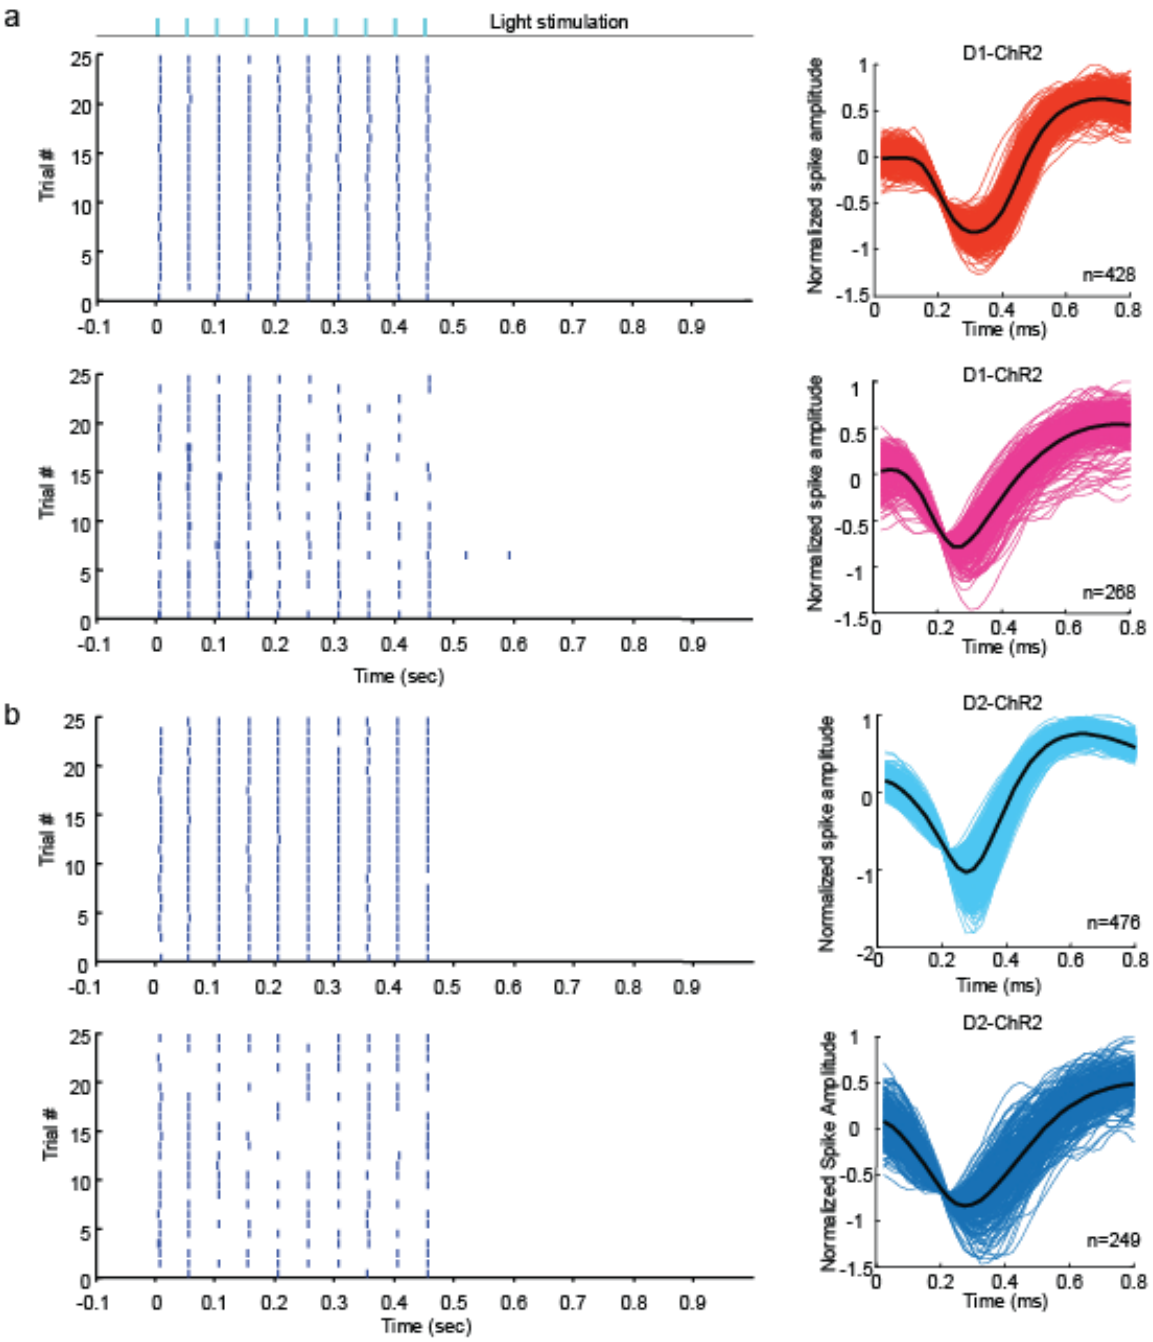

Figure 15. Optical stimulation induces spiking in ChR2-transduced striatum

(a) Spike raster of two representative single units for light-evoked activity in the striatum of D1-ChR2 mice. Stimulation was delivered with 5ms pulses at 20 Hz for 500 ms. The waveforms of the units are shown in the right panels. (b) Spike raster of two representative single units for light-evoked activity in the striatum of D2-ChR2 mice. Right: recorded spike waveforms.

Given that the animals' responses were dependent upon their reward history, we analyzed choice in stimulation and non-stimulation trials based on the history of reward and choice over the previous two trials. Activation of D1R-expressing neurons with optical stimulation induced a bias in choice for the port contralateral to the hemisphere in which the light was delivered. The stimulation-induced bias was greater after unrewarded trials when the animals' responses tended to be more variable (**Fig. 16b**). A small or insignificant bias was seen following rewarded trials (**Fig. 16b**).

Striatal stimulation in D2-Cre mice induced a bias to the port ipsilateral to the site of stimulation (**Fig. 16c**), opposite to the direction observed in D1-Cre mice. Again stimulation induced bias was greater after unrewarded trials when the animals' responses were more variable (**Fig. 16c**).

To further investigate the relationship between stimulation and reward history, we plotted the probability of a left choice for a given reward history in trials with and without stimulation (**Fig. 16d,e**). The plot revealed 'bowing' of data points off the unity line in opposite directions for D1- or D2-Cre animals. This bowing can be quantified as the odds of choice without stimulation scaled by a fixed factor called the "odds ratio".

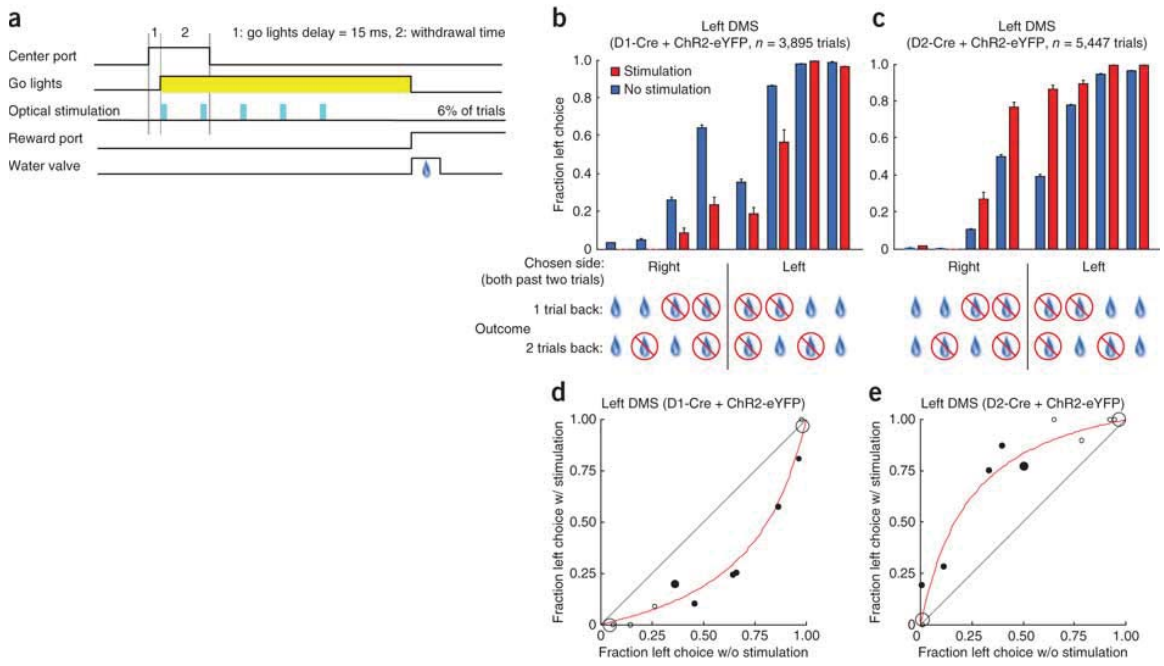

Figure 16: Optical stimulation induces opposing biases in a mouse's choice.

(**a**) Timing of optical stimulation in the task. In 6% of trials, optical stimulation was delivered to the dorsal striatum during a 500-ms period starting at the same time as the Go light cues. Stimulation occurred at 5, 10 or 20 Hz, delivering 3, 5 or 10 pulses, respectively, of 5-ms light stimulation. (**b**,**c**) Examples showing the effect of 10-Hz stimulation in the left dorsomedial striatum (DMS) of a D1-Cre mouse (**b**) and a D2-Cre mouse (**c**) expressing ChR2-eYFP. Individual bars represent the fraction of left choices for various reward histories in trials in which the mouse previously made two consecutive responses at the same port. Red bars indicate stimulation trials and blue bars represent trials without stimulation. (**d**,**e**) Fraction of left choices with and without stimulation for all possible combinations of choices and outcomes in the previous two trials with more than five total occurrences. (**d**) Data from the D1-Cre mouse shown in **b**. The frequency of trials with a given reward history are indicated by the relative size of the circle. Filled circles represent a significant change in fraction of left choice with stimulation ($P < 0.05$, Fisher's exact test). The red curve relates the probabilities of choice with and without stimulation for a fixed odds ratio (odds ratio = $e^{-1.33 \pm 0.20}$). (**e**) Data from the D2-Cre mouse shown in **c** (odds ratio = $e^{1.45 \pm 0.18}$). All error bars represent s.e.m.

Optical stimulation mimics a change in the animals' valuation of actions

In mice performing the switching task without stimulation, recent reward history exerted a strong effect on an animal's upcoming choice (**Fig. 10**). We quantified these effects by generating estimates of the value of actions (**Fig. 12b**), which could be used to guide the animals' selection of future actions(Lau and Glimcher 2005; Samejima, Ueda et al. 2005; Lau and Glimcher 2008). Similar to rewards, striatal stimulation also had the ability to bias animals' choices. Given that neurons in striatum may encode the value of actions(Samejima, Ueda et al. 2005; Lau and Glimcher 2008), we hypothesized that striatal stimulation may mimic the effects that a change in valuation of actions (based on reward history) would have on the animals' upcoming choice.

Using our previous estimates of the relative action value for various reward and choice histories (**Fig.12b**), we plotted the probability of choices made with and without stimulation for individual subjects in the same sessions (**Figs. 17 and 18, Figs. 20 and 21**). This analysis revealed that striatal stimulation shifted the sigmoid choice probability curve along the relative action value axis. In this way, stimulation can be interpreted as mimicking a change in the relative valuation for selecting the left versus right port by a fixed factor. In individual animals, optical stimulation of D1R-expressing striatal neurons appeared to increase the value of the contralateral port (**Fig. 17a**) while stimulation of the D2R-expressing neurons resembled a decrease in the value for the contralateral port (**Fig. 18a**). In both individual and group data, the magnitude of the shift in relative

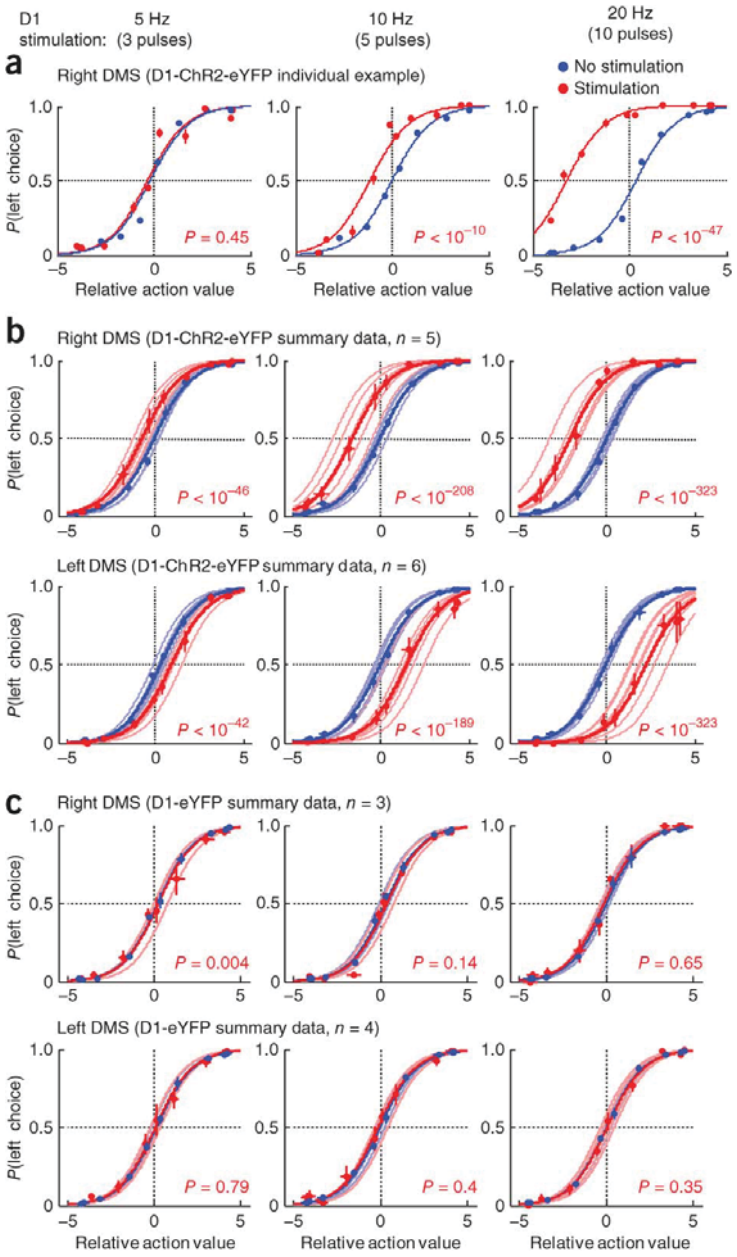

Figure 17: Dorsal striatal D1R-expressing neuron activation mimics an increase in relative action value for contralateral choice.

Fraction of choices for the left port on trials with different relative action value estimates in D1-Cre mice in the presence (red) or absence (blue) of optical stimulation. (**a**) Representative data from one mouse transduced with AAV-EF1α-DIO-ChR2-eYFP with optical stimulation in the right hemisphere. Logistic regression was used to fit the data from trials with (red line) and without stimulation (blue line). A leftward shift in the logistic curve represents a bias for the left reward port. (**b**) Summary data for probability of choice for the left port and relative action value pooled from all D1-Cre mice expressing ChR2-eYFP and stimulated on either the right hemisphere (top) or left hemisphere (bottom). A logistic regression fit for the no stimulation (thick blue line) and stimulation (thick red line) trial data was produced for data pooled from all mice. Curves representing the estimated probability of choice for given relative action values in individual animals are plotted in light blue (no stimulation) or light red (with stimulation). Reported *n* refers to number of stimulation sites (one per hemisphere). (**c**) Summary data from all D1-Cre mice expressing eYFP alone and stimulated on either the right hemisphere (top) or left hemisphere (bottom). *P* values reported for *t* tests: H_0:β_{stim} = 0 (distance between thick red and blue lines). All error bars represent s.e.m.

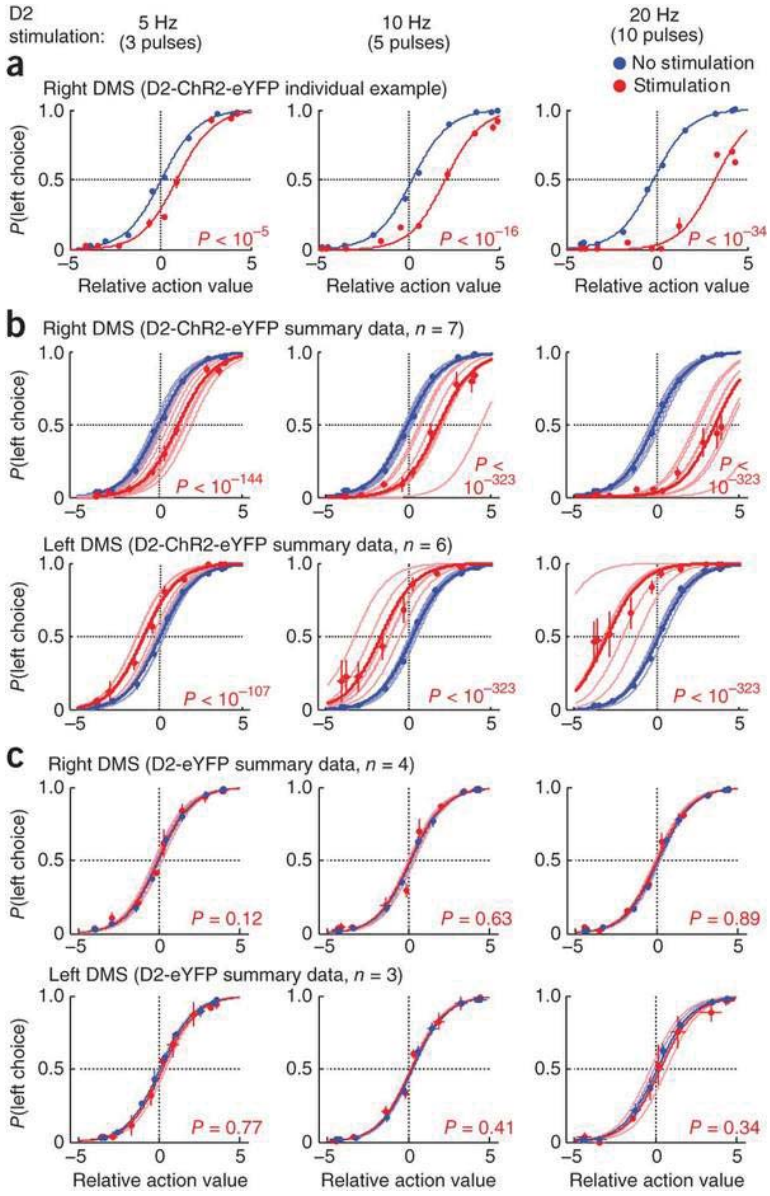

Figure 18: Dorsal striatal D2R-expressing neuron activation mimics a decrease in relative action value for contralateral choice.

Fraction of choices for the left port on trials with different relative action value estimates in D2-Cre mice in the presence (red) or absence (blue) of optical stimulation. (**a**) Representative data from one mouse transduced with AAV-EF1α-DIO-ChR2-eYFP with optical stimulation in the right hemisphere. Logistic regression was used to fit the data from trials with (red line) and without stimulation (blue line). A rightward shift in the logistic curve represents a bias for the right reward port. (**b**) Summary data for the probability of choice for the left port and relative action value pooled from all D2-Cre mice expressing ChR2-eYFP and stimulated on either the right hemisphere (top) or left hemisphere (bottom). A logistic regression fit for the no stimulation (thick blue line) and stimulation (thick red line) trial data was produced for data pooled from all mice. Curves representing the estimated probability of choice for given relative action values in individual mouse are plotted in light blue (no stimulation) or light red (with stimulation). Reported *n* refers to number of stimulation sites (one per hemisphere). (**c**) Summary data from all D2-Cre mice expressing eYFP alone and stimulated on either the right hemisphere (top) or left hemisphere (bottom). *P* values reported for *t* tests: H_0:β_{stim} = 0 (distance between thick red and blue lines). All error bars represent s.e.m.

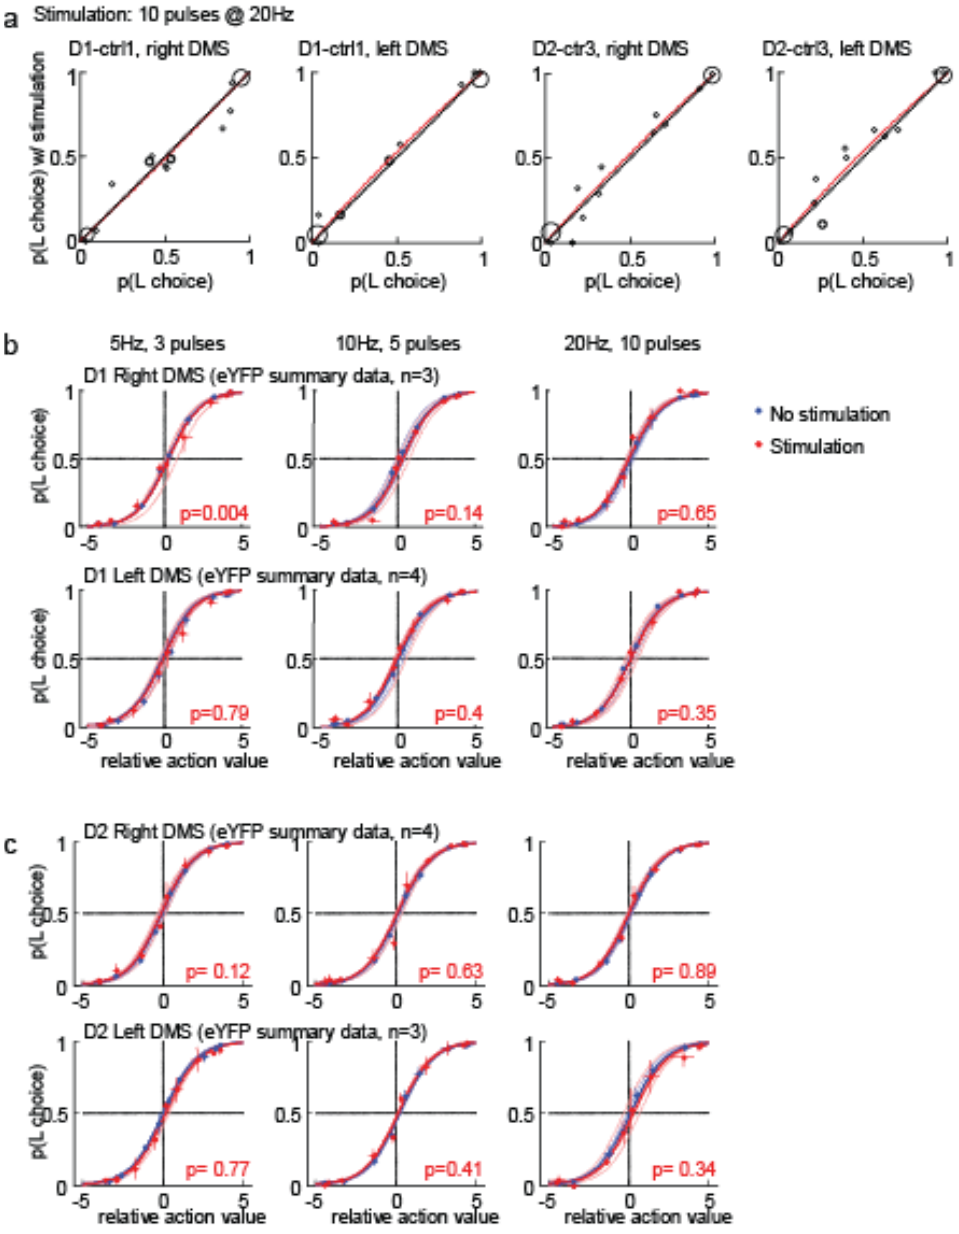

Figure 19. Optical stimulation of control animals caused no bias

(a) Representative data from individual D1-cre or D2-cre control subjects transduced with AAV-EF1α-DIO-eYFP showing fraction of left choices for reward histories in the previous two trials with occurrence>5 trials (out of 16 possibilities). The frequency of trials of a given reward history is indicated by the relative size of the circle. **(b)** Fraction of left choices on trials with different relative action value estimates in D1-cre and **(c)** D2-cre animals in the presence (red) or absence (blue) of optical stimulation (protocol in Fig. 3a). Data from all subjects transduced with AAV-EF1α-DIO-eYFP with optical stimulation in the right and left hemisphere. Logistic regression was used to fit the data from trials with (red line) and without stimulation (blue line). Curves representing the estimated probability of choice for given relative action values in individual animals are plotted in light blue (no stimulation) or light red (with stimulation). Reported n refers to number of stimulation sites. P values reported for t-tests: H0:ßstim=0 (red lines). All error bars represent s.e.m.

action value consistently increased with the frequency of stimulation and the direction of the shift induced by stimulation showed a consistent relationship with the stimulated hemisphere (**Figs. 17b, 18b and 19a**; individual data can be seen in **Figs. 20 and 21**). Importantly, control animals did not demonstrate a bias in their distribution of choices between trials with or without optical stimulation. In either D1- or D2-Cre control animals expressing only eYFP without ChR2, no shift in action value was observed (**Figure 17c,18c**).

Optical stimulation altered withdrawal time in manner dependent on reward and choice history

In addition to its role in action selection, activity in the striatum correlates with reaction time(Lauwereyns, Watanabe et al. 2002; Watanabe and Hikosaka 2005). In our task, the time a mouse takes to withdraw from the center port after trial initiation can be used to measure the latency to initiate movement. Without optical stimulation, the withdrawal time for a particular response was significantly shorter when the relative action value for that response was high (**Fig. 12d**). Stimulation of D1R-expresing neurons decreased center port withdrawal time when the value of the contralateral port was greater (**Fig. 22b**). However, when the ipsilateral port was valued more highly, stimulation of D1R-expressing neurons slowed withdrawal time (**Fig. 22b**). In contrast to D1R-expressing neurons, striatal optogenetic stimulation in the D2-Cre mice increased center port withdrawal time when the value of the contralateral port was greater (**Fig. 22b**). However, when the ipsilateral port was valued more highly, stimulation of D2-expressing neurons sped withdrawal times (**Fig. 22b**). Changes in withdrawal time with optical stimulation were still apparent when we controlled for the

eventual choice response of the animals (**Fig. 23a,b**). Importantly, we find that after rewarded trials when optical stimulation is not likely to affect port choice, stimulation can still be observed to affect withdrawal time.

Optical stimulation can bias choice behavior prior to response initiation and the effect is diminished after a 150 ms delay

Previous reports have found that phasic striatal activity prior to response initiation correlated with trial-by-trial estimates of action value(Samejima, Ueda et al. 2005; Lau and Glimcher 2008). In the experiments described above, light pulses were delivered during a 500ms epoch coinciding with onset of a 'Go' cue after which animals initiated their motor responses (**Fig. 16a**).

We next performed two further experiments to investigate the impact of striatal activity prior to and after response initiation. Response initiation in our task can be measured as the latency to withdraw from the center port. The median withdrawal times for animals in this study were approximately 50 ms, and the majority of response withdrawal times can be found within a window spanning approximately 30 to 140 ms (50±25% confidence interval; **Fig. 11**).

To determine whether stimulation prior to movement initiation was sufficient to bias the animals' choice behavior, a new cohort of mice were trained in a protocol in which two optical pulses separated by 50 ms (20Hz) were delivered prior to an auditory 'Go' cue (**Fig. 24a**). Using this second protocol, we confirmed that it is possible to bias the animals' choice behavior with striatal activation of

Figure 20

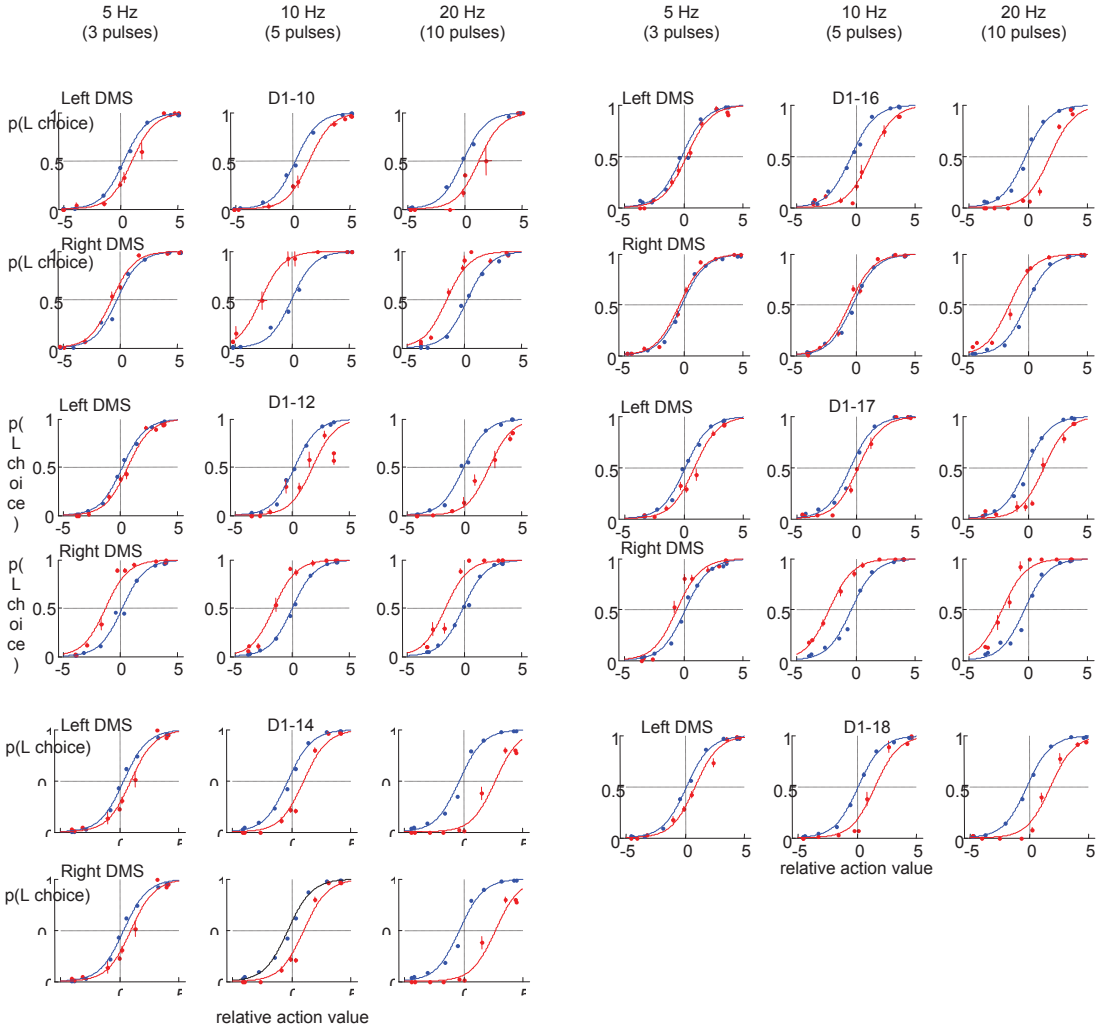

Figure 20. Gallery of individual D1-ChR2 animals' data

Fraction of choices for the left port on trials with different relative action value estimates in D1-cre animals in the presence (red) or absence (blue) of optical stimulation (protocol in Fig. 3a). All subjects were transduced with AAV-EF1α-DIO-ChR2-eYFP. Optical stimulation in the right and left hemisphere are shown separately. Some mice only demonstrated ChR2 expression unilaterally and stimulation sessions were only used from the transduced side. Logistic regression was used to fit the data from trials with (red curve) and without stimulation (blue curve). All error bars represent s.e.m.

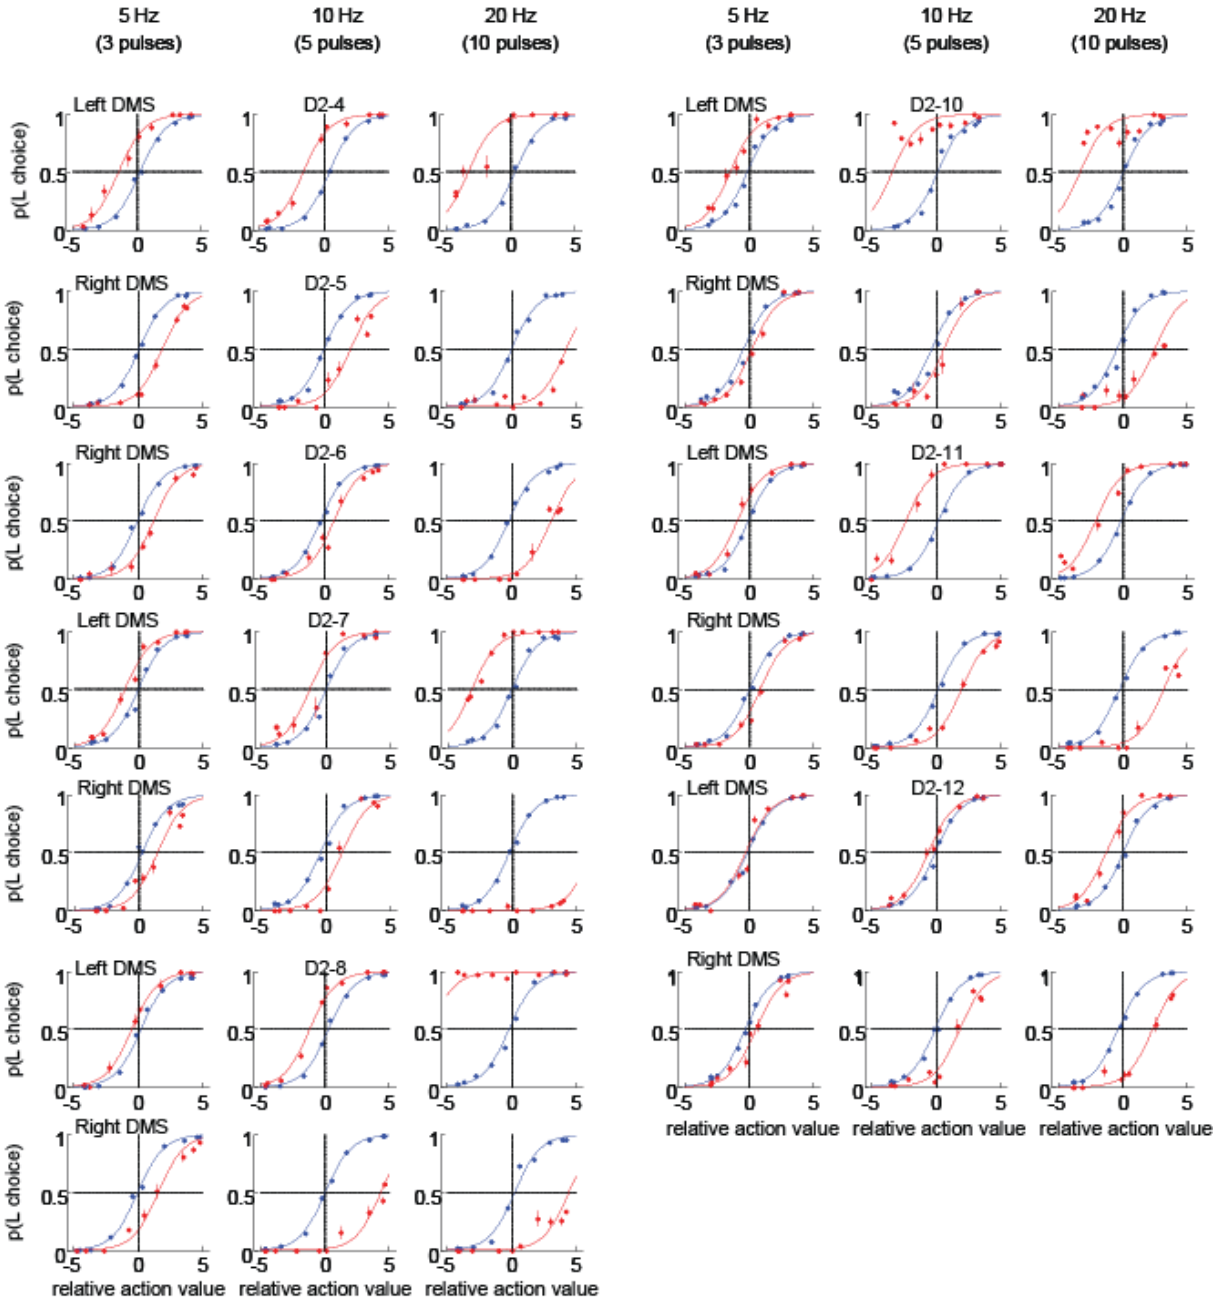

Figure 21. Gallery of individual D2-ChR2 animals' data
Fraction of choices for the left port on trials with different relative action value estimates in D2-cre animals in the presence (red) or absence (blue) of optical stimulation (protocol in Fig. 3a). Data from all subjects transduced with AAV-EF1α-DIO-ChR2-eYFP with optical stimulation in the right and left hemisphere. Logistic regression was used to fit the data from trials with (red curve) and without stimulation (blue curve). All error bars represent s.e.m.

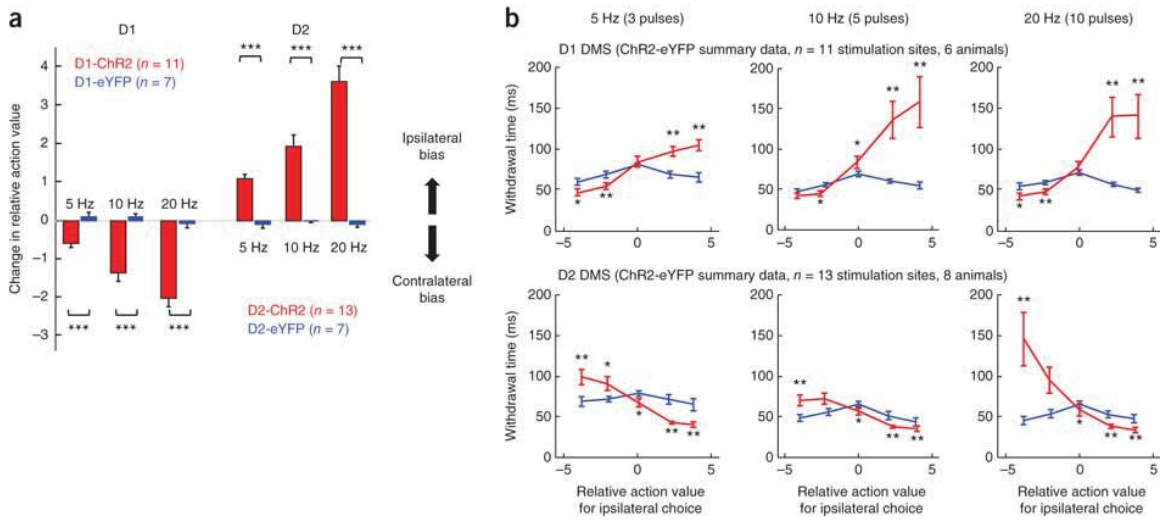

Figure 22: Comparison of D1R- and D2R-expressing neuron stimulation.

(**a**) Estimated change in the relative action value for choosing the port ipsilateral versus contralateral to the site of stimulation averaged across individuals in a group. Positive changes in relative action values correspond to an ipsilateral bias and negative changes correspond to a contralateral bias. Estimates of relative action change were derived from logistic regression analysis (Online Methods) for 5-, 10- and 20-Hz stimulation sessions. Reported *n* refers to number of stimulation sites. The D1-Cre data set consisted of six stimulated mice expressing ChR2-eYFP and four control mice expressing only eYFP. The D2-Cre data set consisted of eight stimulated mice expressing ChR2-eYFP and five control mice expressing only eYFP. (**b**) The median time taken to withdraw from the center port averaged across individual D1-Cre mice (top) and D2-Cre mice (bottom) expressing ChR2-eYFP in trials without stimulation (blue) or with stimulation (red) and across different relative action values for choosing the port ipsilateral versus contralateral to the site of stimulation. Positive relative action values correspond to trials in which the value of the port ipsilateral to the site of stimulation is greater than the contralateral port. Median times to withdraw are plotted for 5- (left), 10- (middle) and 20-Hz (right) stimulation sessions. All error bars represent s.e.m. **P* < 0.05, ***P* < 0.01, Wilcoxon signed-rank test. ****P* < 0.001, Wilcoxon rank-sum test.

both D1R-expressing and D2R-expressing striatal neurons prior to movement initiation (**Fig. 24b,c,d; Supplementary Fig. 25a,b**).

In a third protocol, we examined the effect of delaying stimulation until after movement initiation. In a subset of mice trained in the original protocol, 10 Hz optical stimulation was delivered at either 0 ms or 150 ms following the presentation of the 'Go' cue light (**Fig. 26a**). The bias induced by optical stimulation delayed by 150 ms was significantly weaker than the bias induced without delay (**Fig. 26b,c,d**). These data suggest that the impact of striatal activity on choice behavior decayed with time after the onset of the 'Go' cue.

Prolonged but not transient stimulation outside of the task context induced lateralized orienting responses and contralateral movement

Following the switching task, animals also underwent two other experiments to grossly measure the impact of transient and prolonged stimulation on the animals' locomotor behavior outside the context of the task. In a 10 cm diameter cylindrical environment, transient bursts of 20 Hz stimulation (identical to the maximum burst used in our switching task) did not produce a significant change in head or body orientation. However, prolonged stimulation for 60 seconds at 5Hz, 10Hz, and 20Hz of the D1-expressing striatal neurons induced a graded increase in contralateral rotations and prolonged stimulation of D2R-expressing neurons induced a decrease in contralateral rotations in an open field. Together, these data are consistent with previous reports(Kravitz, Freeze et al. 2010) and the known

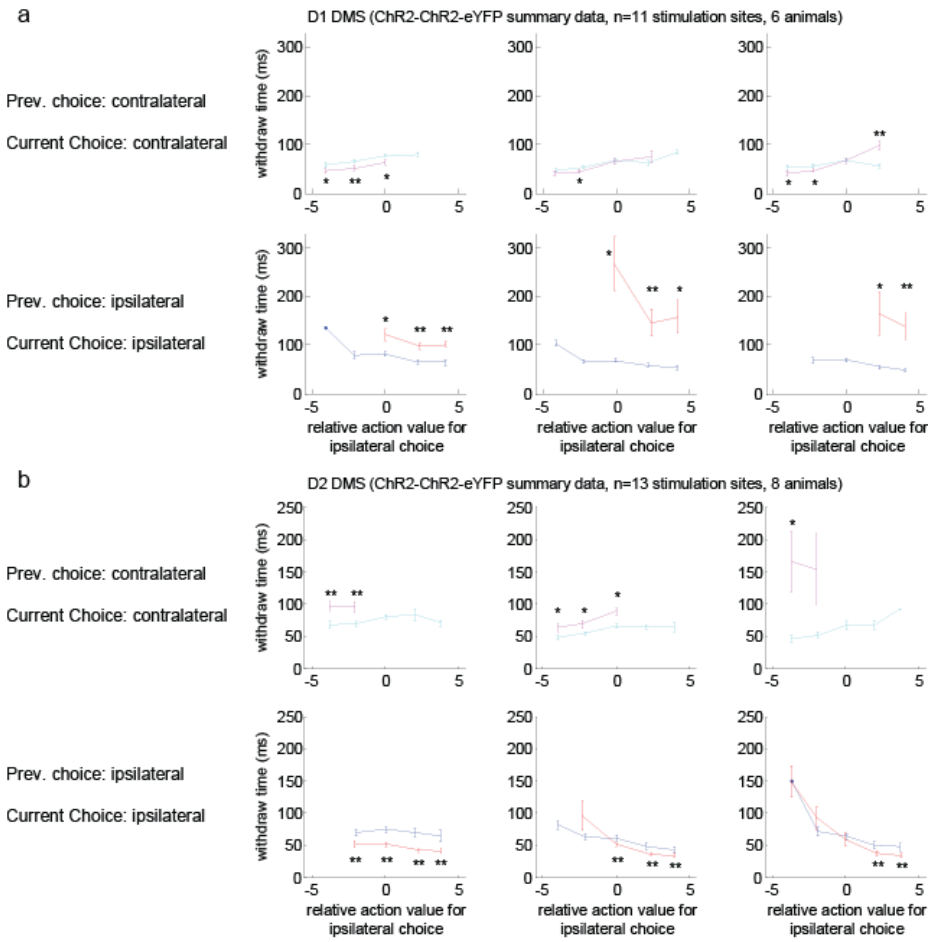

Figure 23. Optical stimulation of the dorsomedial striatum altered time to withdraw from the center port on trials when animals did not switch sides

On trials when animals did not switch sides relative to the previous trial (stay trials), we plot time the withdrawal time (time until withdrawal from center port after the go signal) for different relative action value estimates. In the task, animals can either stay at the port contralateral to optical stimulation (top panels in **(a)** and **(b)**) or stay at the port ipsilateral to stimulation (bottom panel in **(a)** and **(b)**). In stay trials at the port contralateral to stimulation for D1-Cre animals **(a)** or D2-Cre **(b)**, pink lines represent trials with stimulation and cyan lines represent trials without stimulation (protocol in Fig. 3a). In stay trials at the port ipsilateral to stimulation, red lines represent trials with stimulation and blue lines represent trials without stimulation. Note we plot only data points where more than 50% of subjects have 5 or more trials. All data is from subjects transduced with AAV-EF1α-DIO-ChR2-eYFP. All error bars represent s.e.m. *: p value for paired t-test<0.05; **: p value for paired t-test<0.01.

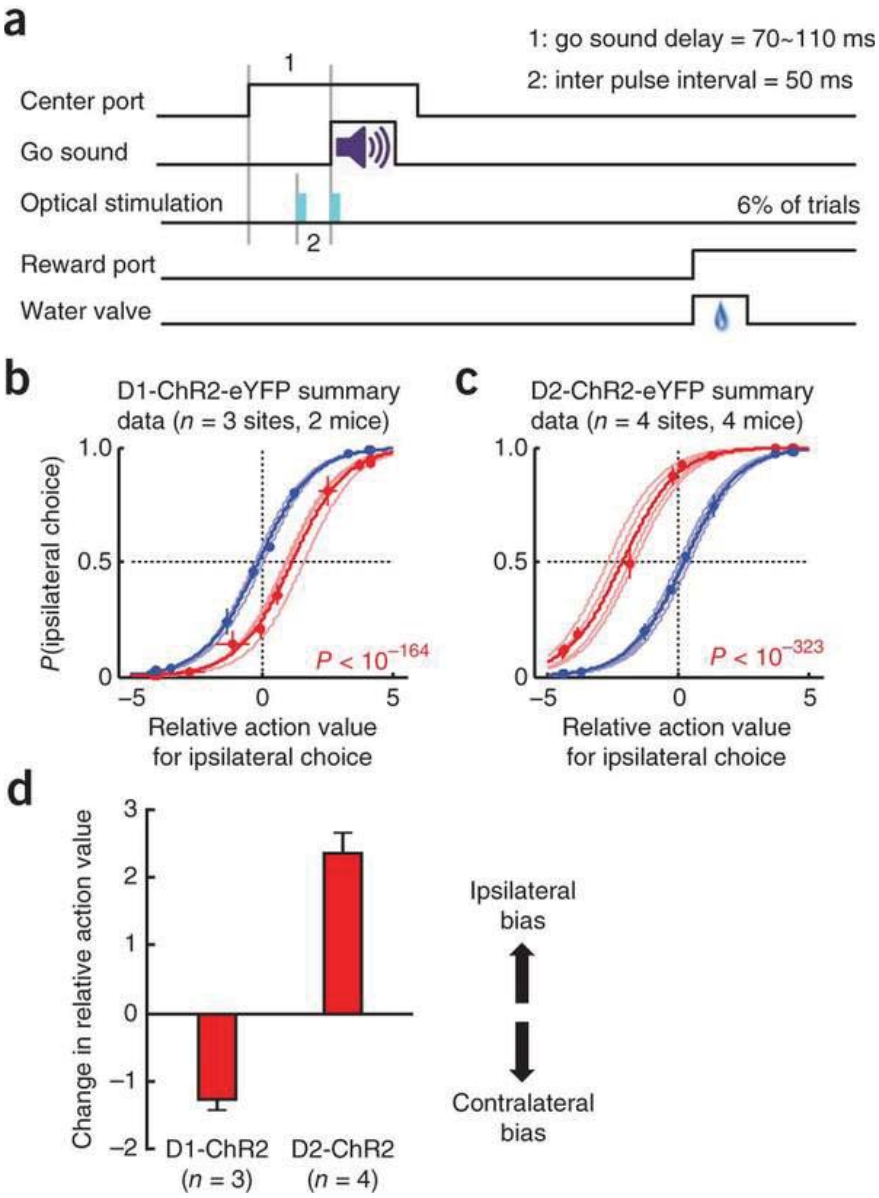

Figure 24: Significant bias was induced by stimulation limited to a 50-ms period before a Go cue.

(**a**) The timing of events in sessions in which a Go cue sound was delayed 70–110 ms after initiation of trial with center port entry. In these trials, two 5-ms optical pulses separated by 50 ms were delivered just prior and at the onset of the Go cue and before withdrawal from the center port. (**b**,**c**) Summary data for probability of choice for the ipsilateral port on trials with different relative action value estimates pooled from D1-Cre mice (**b**) and D2-Cre mice (**c**) expressing ChR2-eYFP with confirmed eYFP expression sites. Positive relative action values correspond to trials in which the value of the port ipsilateral to the site of stimulation was greater than the contralateral port. A logistic regression fit for the no stimulation (thick blue line) and stimulation (thick red line) trial data was produced for data pooled from all subjects. Curves representing the estimated probability of ipsilateral choice for given relative action values in individual animals are plotted in light blue (no stimulation) or light red (with stimulation).*P* values reported for *t* tests: H_0:β_{stim} = 0 (distance between thick red and blue lines). (**d**) Change in estimated relative action value for choosing the port ipsilateral versus contralateral to the site of stimulation averaged across individuals stimulated with two optical pulses before movement initiation. All error bars represent s.e.m.

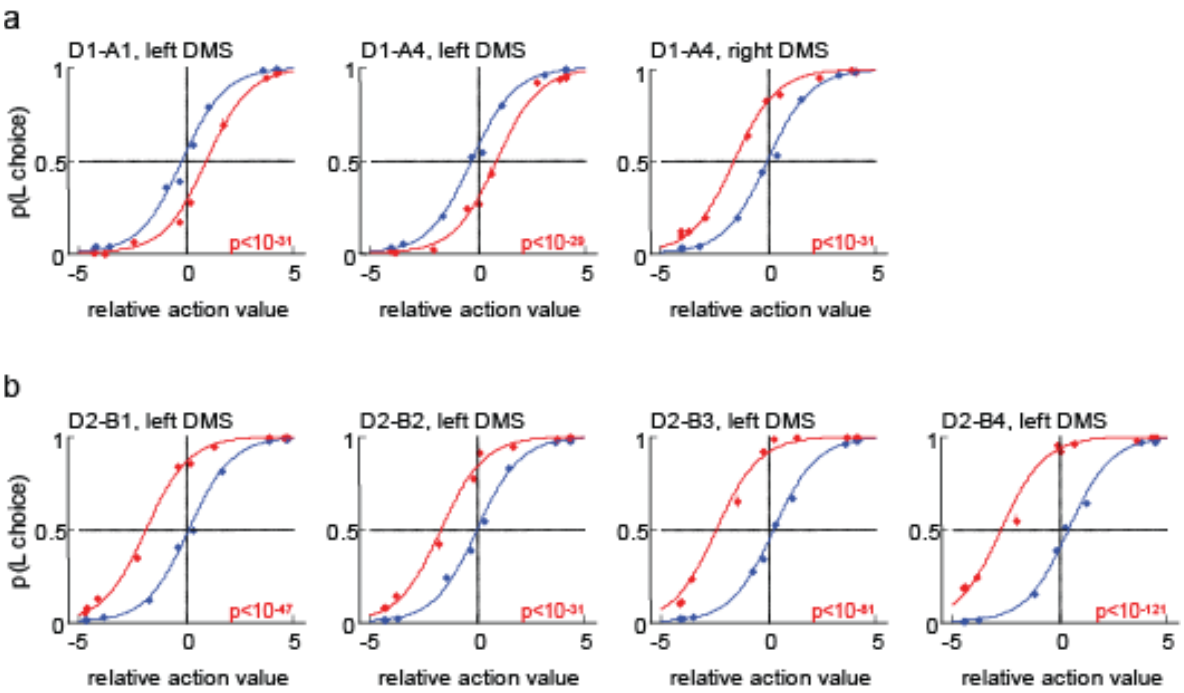

Figure 25. Stimulation prior to 'go-signal' and movement initiation mimics a change in the relative value for approaching each port

Fraction of choices for the left port on trials with different relative action value estimates in individual (a) D1-cre or (b) D2-cre animals in the presence (red) or absence (blue) of optical stimulation (protocol shown in Fig 7a). All subjects were transduced with AAV-EF1α-DIO-ChR2-eYFP and received optical stimulation unilaterally in the right and/or left hemisphere. Logistic regression was used to fit the data from trials with (red curve) and without stimulation (blue curve). All error bars represent s.e.m.

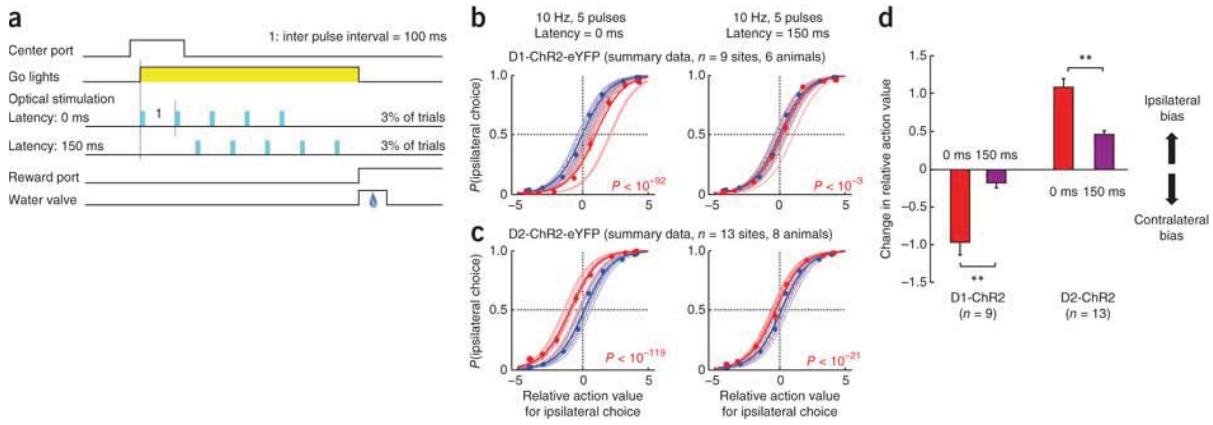

Figure 26: Bias induced by stimulation diminished when stimulation was delayed 150 ms after a Go cue.

(**a**) The timing of events in sessions in which stimulation occurred at latencies of either 0 or 150 ms relative to Go light cue onset. (**b**,**c**) Summary data for probability of choice for the ipsilateral port and relative action value pooled from D1-Cre mice (**b**) and D2-Cre mice (**c**) expressing ChR2-eYFP with confirmed eYFP expression sites stimulated with 0-ms latency (left) or 150-ms latency (right) relative to Go light onset. A logistic regression fit for the no stimulation (thick blue line) and stimulation (thick red line) trial data was produced for data pooled from D1-Cre or D2-Cre mice. Curves representing the estimated probability of choice for given relative action values in individual animals are plotted in light blue (no stimulation) or light red (with stimulation). *P* values reported for *t* tests: H_0:β_{stim} = 0 (distance between thick red and blue lines). (**d**) Estimated change in relative action value for choosing the port ipsilateral versus contralateral to the side of stimulation averaged across individuals stimulated at 10 Hz with 0- and 150-ms latency in the same session. Estimates of relative action value change were derived from logistic regression analysis (Online Methods). Reported *n* refers to number of stimulation sites. ***P*< 0.005, Wilcoxon signed-rank test. All error bars represent s.e.m.

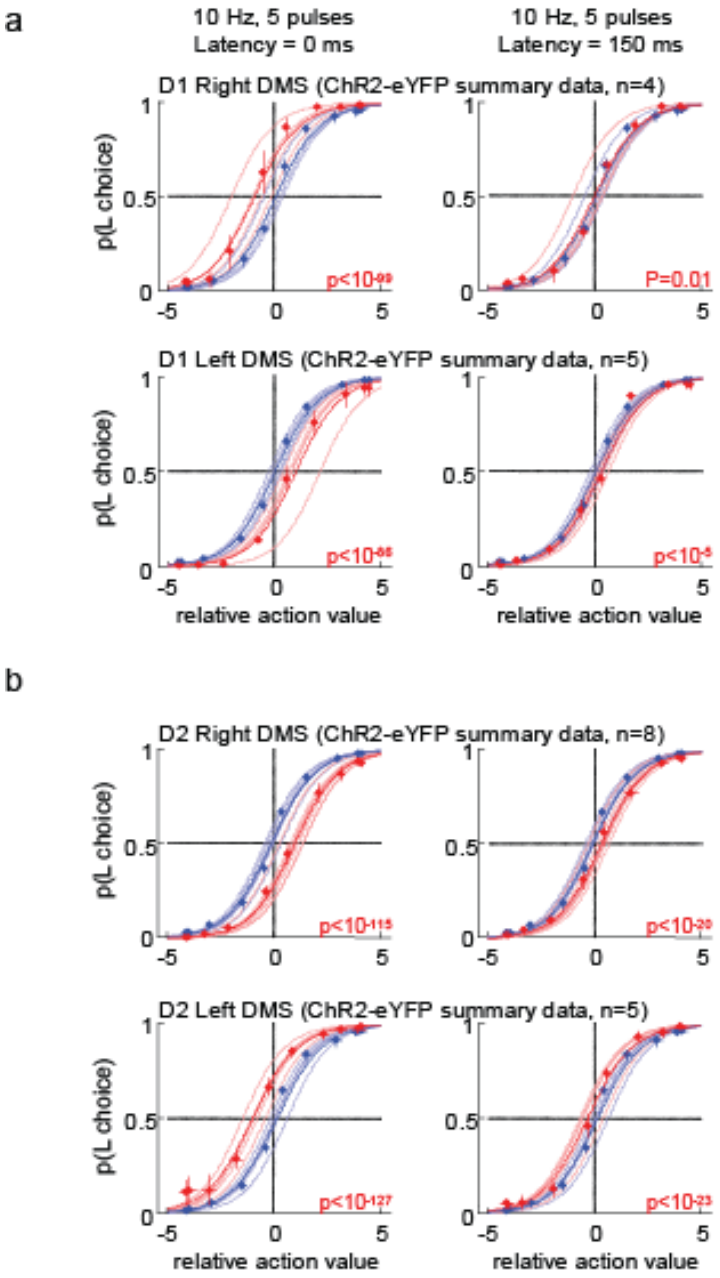

Figure 27. Delayed striatal stimulation mimics a change in the relative value for choosing the left versus right port, but is diminished in magnitude

(a) Summary data for probability of left choice and relative action value pooled from D1-Cre subjects expressing ChR2-eYFP and stimulated on the right hemisphere (top) or left hemisphere (bottom) at latency of 0 ms (left) or 150 ms (right) (protocol in Fig. 7c). A logistic regression fit for the no stimulation (thick blue line) and stimulation (thick red line) trial data was produced for the pooled data. Curves representing the estimated probability of choice for given relative action values in individual animals are plotted in light blue (no stimulation) or light red (with stimulation). **(b)** Summary data for probability of left choice and relative action value pooled from D2-Cre subjects expressing ChR2-eYFP and stimulated at latency of 0 ms (left) or 150 ms (right). Reported n refers to number of stimulation sites. P values reported for t-tests: H0:ßstim=0 (red lines). All error bars represent s.e.m.

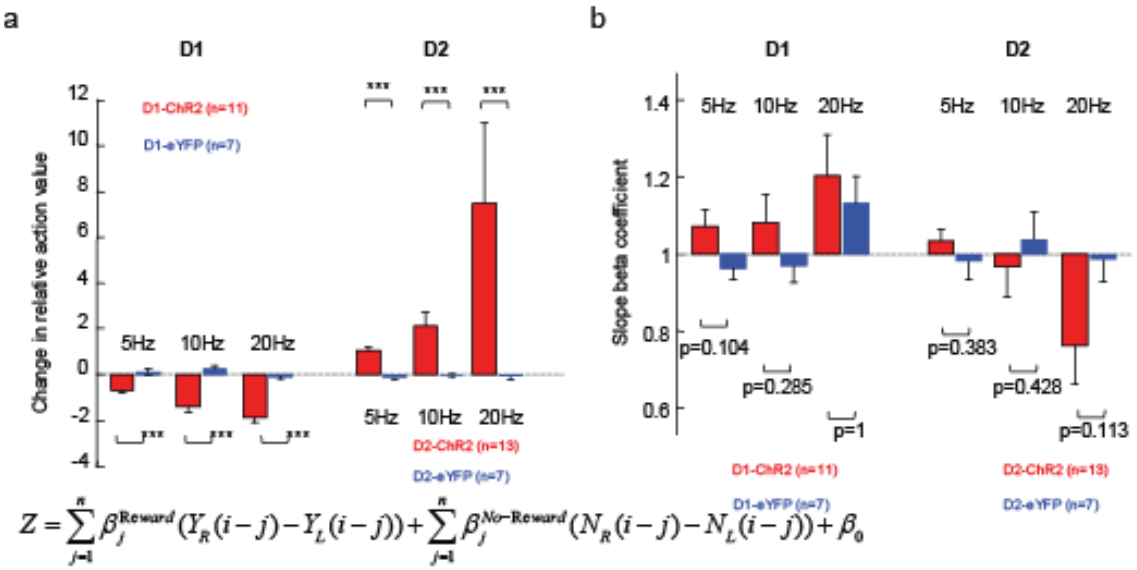

$$Z = \sum_{j=1}^{n} \beta_j^{\text{Reward}} (Y_R(i-j) - Y_L(i-j)) + \sum_{j=1}^{n} \beta_j^{\text{No-Reward}} (N_R(i-j) - N_L(i-j)) + \beta_0$$

$$\log\left(\frac{P_L(i)}{P_R(i)}\right) = \beta_{slope} * Z + \beta_{stimulation} \cdot X_{stim}(i)$$

Figure 28. Optical stimulation mimics a change in the estimated relative action value (shift), but does not significantly affect reward sensitivity (slope)

We analyzed the effect of optical stimulation using another model which allows the stimulation to cause changes in animals' sensitivity to reward (slope) as well as a shift in relative action value. We first define the relative action value (Z) derived from logistic regression of trials without stimulation as:

$$Z = \sum_{j=1}^{n} \beta_j^{\text{Reward}} (Y_R(i-j) - Y_L(i-j)) + \sum_{j=1}^{n} \beta_j^{\text{No-Reward}} (N_R(i-j) - N_L(i-j)) + \beta_0$$

We then fit the behavior data from stimulation trials with the following logistic regression to allow a change in sensitivity to reward (slope) as well as a shift in relative action value simultaneously.

$$\log\left(\frac{P_R(i)}{1-P_R(i)}\right) = \beta_{slope} * Z + \beta_{stimulation} \cdot X_{stim}(i)$$

(a) Estimated change in the relative action value (shift) for choosing the port ipsilateral versus contralateral to the side of stimulation averaged across individuals within a group. Positive changes in relative action values correspond with an ipsilateral bias while negative changes correspond to a contralateral bias. Increasing stimulation frequency mimicked a change in the relative for choosing the contralateral port with D1 stimulation while D2 stimulation mimicked a change in the relative value for choosing the ipsilateral port. **(b)** Estimated change in reward sensitivity (slope) averaged across individuals within a group. No significant change in reward sensitivity was observed in stimulation trials. Reported n refers to number of stimulation sites. ***: p<0.001, P values reported for Wilcoxon rank-sum test. All error bars represent s.e.m.

organization of distinct striatal pathways in the basal ganglia (Albin, Young et al. 1989; DeLong 1990; Kreitzer and Berke 2011).

Discussion

Our data suggest that activity in distinct populations of striatal neurons exert opposing biases on the selection of goal-directed responses. Activation of D1R-expressing striatal neurons increased the occurrence of choices for the port contralateral to the side of stimulation (**Figs. 16, 17, and 22**) while stimulation of D2R-expressing striatal neurons increased the occurrence of ipsilateral choices (**Figs. 16, 18, and 22**).The effect of stimulating each population of neurons was not deterministic, but dependent on the animals recent reward history (**Fig. 16**). Upon closer inspection, the magnitude of this bias mimicked an additive change in the relative value of actions estimated using a simple model based upon the animals history of rewards and choices (**Fig. 2b, Figs. 16-18, 22**). Stimulation also altered the latency to movement initiation as measured by the withdrawal time in manner that was dependent on the relative action value (**Figs. 18b**).

Opposing biases induced by stimulation of distinct striatal populations are dependent on the history of rewarded actions and timing

Qualitatively, the effects of stimulating D1R and D2R-expressing neurons match existing accounts of the opposing functions of the direct and indirect pathway within the basal ganglia. This lateralized effect may be due to the presence of dense ipsilateral descending connectivity from basal ganglia nuclei and the role that downstream efferent structures play in controlling contralateral movement (Grillner, Hellgren et al. 2005; Felsen and Mainen 2008; Grillner, Wallen et al.

2008). The ability of D2R-expressing neurons to promote choices to the ipsilateral port is consistent with a larger literature suggesting that action selection involves inhibition of competing alternatives through the indirect pathway in a manner that allows or even facilitates focal promotion of desired actions by the direct pathway(Redgrave, Prescott et al. 1999). The proposed occurrence of targeted inhibition from basal ganglia pallidal outputs has been suggested to coordinate agonist and antagonist musculature involved in limb and visuomotor movements in primates (Mink 1996)(Jiang, Stein et al. 2003). Competition between the opposing actions of orienting to the left and right may also be regulated within the downstream targets of the basal ganglia. The direction of the bias observed in our study is largely consistent with previous reports in which D1R and D2R expressing neurons are selectively activated using optogenetic techniques(Kravitz, Freeze et al. 2010) as well as the effect of pharmacological manipulations to the dorsal striatum in rodents(Schwarting and Huston 1996).

Optogenetic stimulation did not induce a uniform effect across all trials, but was dependent on the animals' previous reward history. Stimulation induced a larger bias on choice when animals had greater variability in their responses following unrewarded trials (**Figs. 16, 3-5**). However, striatal stimulation had a weak or insignificant effect on choice behavior after recently rewarded trials when animals were likely to return to a port where water had just been delivered. Thus, the effect of stimulation could be overruled if the alternative response had a high incentive value after being recently rewarded.

The effect of optical stimulation was also time dependent. Stimulation was effective when limited to two 5 ms pulses within an epoch of the task prior to the animals' movement initiation (**Fig. 24**), and it became significantly weaker if delayed by 150 ms after trial initiation (**Fig. 26**). Together, these data suggest that striatal activation may need to take place in a "decision window" to alter the action selection of the animal. This is consistent with a number of recording studies in the striatum of rodents(Kim, Sul et al. 2009; Thorn, Atallah et al. 2010) and in primates(Samejima, Ueda et al. 2005; Watanabe and Hikosaka 2005; Lau and Glimcher 2008). Finally, outside of the task, transient stimulation at the maximal experimental level did not induce head or body orientation. The interaction of stimulation with reward history, the presence of an effective decision window, and the lack of similar motor output outside of the task, demonstrate that striatal stimulation did not dictate a motor action deterministically impacting choice. We conclude striatal activity alone is not sufficient to drive the motor responses of animals in this task, but instead may also require the temporally coincident activity of other neural structures to orchestrate a complex process of action selection.

Stimulation bias mimics a change in the value of actions

Given that the magnitude of the stimulation bias was dependent on whether previous actions had been rewarded, we hypothesized that the striatal activation may act similarly to the influence of rewards over choice. Within the context of this task, animals are required to adaptively and flexibly switch their actions across blocks due to changing reward contingencies. At a theoretical level, this process of goal-directed action selection can be modeled as a dynamic

comparison between the value of actions and bias for selection of the option of highest value(Rangel, Camerer et al. 2008). More specifically, we conjectured that the additional activity induced by optical stimulation mimicked the effects resulting from a change in the value of competing actions.

To demonstrate this directly, we estimated the value of actions based upon various reward and choice histories assuming the softmax decision rule(Schultz, Dayan et al. 1997; Sutton 1998; Lau and Glimcher 2008). We then re-examined quantitative features of the bias introduced by striatal activation and found that stimulation mimicked a fixed additive shift in the value of the contralateral choice (**Figs 17, 21, and 22**) without altering sensitivity to reward (**Fig. 28)**. In this way, activation of D1R-expressing neurons mimics an increase in the value of the contralateral choice while activation of D2R-expressing neurons mimics a decrease of value of the contralateral choice. The nonlinear features of the softmax rule were sufficient to explain the gross tendency for stimulation to have a larger effect in a range where responses are most variable (**Figs. 16, 17, and 5).**

While many descriptions of striatal function focus on its role in action selection, others have suggested that the basal ganglia merely regulates the “vigor of responses” without altering which response is selected(Desmurget and Turner 2010; Turner and Desmurget 2010). Here, we show that in addition to biasing the animals’ action selection, striatal activity can also alter the “vigor of responses” by speeding up or slowing down the initiation of movements. We found that stimulation of D1R- expressing neurons reduced movement latencies in trials

when the value of the contralateral port was greater and congruent with the direction of bias caused by stimulation (**Fig. 22b**). This is consistent with the canonical view that the direct pathway promotes movements(Albin, Young et al. 1989; DeLong 1990). However, latencies following stimulation of D1R-expressing neurons were dramatically slower in trials when the ipsilateral port was valued more highly. The slowed response is perhaps due to the incongruence between bias caused by stimulation and intrinsic valuation in the action selection systems. Similarly, stimulation of D2R-expresing neurons slowed movements when the port contralateral to the site of stimulation was of greater value and it sped movements when the port ipsilateral to the stimulation site was of greater value. This data is consistent with the idea that action selection can be facilitated by the suppression of alternate reponses(Mink 1996). Our data suggest the balance of activity within striatal populations reflects the relative value of approaching each port impacts the speed of movement initiation. These data are consistent with evidence that striatal activity correlates with reaction times(Lauwereyns, Watanabe et al. 2002; Watanabe and Hikosaka 2005)

Our data is consistent with electrophysiological studies suggesting neural activity in the striatum more often represents the value of actions than pure motor variables(Samejima, Ueda et al. 2005; Lau and Glimcher 2008). Striatal activity has been associated with response bias for rewarded actions (Lauwereyns, Watanabe et al. 2002) and successful switching following action contingency reversals(Watanabe and Hikosaka 2005; Kimchi and Laubach 2009). In these studies, neurons primarily encode a bias for contralateral responses when decisions are reported as saccades or locomotor approach(Lauwereyns,

Watanabe et al. 2002; Watanabe and Hikosaka 2005; Thorn, Atallah et al. 2010). Lesions of the dorsomedial striatum can also impair the learning of the contingencies between actions and their outcomes in various reversal tasks (Balleine, Delgado et al. 2007; Ragozzino, Lassandro et al. 2007; Ragozzino 2007). Taken together, these studies are consistent with our interpretation that striatal stimulation biases both the selection and vigor of actions based upon the value of choices.

We have attempted to best reproduce physiological conditions using optogenetics. The stimulation patterns used in our experiments parallel data from awake *in vivo* striatal recordings in mice (Kubota, Liu et al. 2009; Jin and Costa 2010). Other labs have also identified striatal activity in a similar decision point within a reward-based spatial task in rodents (Kim, Sul et al. 2009). The bias induced by stimulation scaled over a range of frequency parameters, supporting the robustness of our results. However, questions still remain regarding whether optogenetic stimulation mimics physiological patterns *in vivo* given the large number of neurons that are synchronously recruited. In the D2-Cre mice, we also infect a small proportion of cholinergic neurons, which may contribute to our behavioral effects. An additional feature and caveat of our study is that our optical stimulation was delivered unilaterally in the context of a task where competing responses are lateralized to the left and right. This is both feature and caveat because we predict our method will likely not cause behavioral bias in alternative task designs where responses take other forms, such as up vs. down. Lastly, our results do not exclude the possibility that D1R-expressing and D2R-expressing neurons may serve other roles outside of action selection. Numerous

labs have found evidence that striatal activity also correlates with a process that evaluates the outcomes of actions in rodent (Kim, Sul et al. 2009; Kravitz, Tye et al. 2012) and in primates(Lau and Glimcher 2008), and our data do not exclude these aspects of striatal function.

Striatum in goal-oriented action selection

Striatal neurons receive massively convergent input from cortical and thalamic sources. This integration of diverse information may be used to generate a representation of action value that can be utilized to mediate goal-oriented action selection. In this framework, the updating of action values may correspond with dopamine-dependent plasticity of inputs into the striatum in concert with shifts in goal and reward-related activity from the cortex(Schultz, Dayan et al. 1997). Striatal neurons may alter action selection by regulating tonic inhibition from the globus pallidus and the substantia nigra pars reticulata onto the brainstem and motor thalamus(Ikeda and Hikosaka 2003; Lo and Wang 2006). This provides a mechanism by which rewards modulate the responses of premotor structures for particular actions(Ikeda and Hikosaka 2003; Kable and Glimcher 2009). In this way, the cortico-basal ganglia system may instantiate the computations necessary for goal-oriented, highly flexible behavior(Redgrave, Prescott et al. 1999; Hikosaka and Isoda 2010).

Chapter 5

Activation of Distinct Ventral Striatal Subpopulations Update the Values of Choices for Future Locomotor Decisions

Abstract

Goal-directed behaviors require animals to evaluate the outcome of their actions to guide future motor behavior. Some reinforcement theory models propose the existence of a 'Critic' that evaluates the outcomes resulting from a choice to determine if the resulting outcome is better/worse than expected to guide the selection of future motor responses. Here, we provide evidence that activity within the ventral striatum can serve the function of the "Critic" proposed in theories of reinforcement learning. We found that transient optogenetic stimulation of ventral striatal dopamine D1 or D2-expressing neurons during the evaluation of an outcome in a decision-making task had opposing effect on how animals learned from their previous history of rewards and choices to impact their future choices. In particular, the effect of stimulation mimicked a fixed additive change in the value of the chosen responses on the distribution of choices on the following trial, suggesting a well-defined mapping between ventral striatal activity onto the probability of future choices. This ability of ventral striatal stimulation to influence the evaluation of choices was limited to a specific window of the task beyond which there were no change in future responses. Moreover, ventral striatal activity during the selection of responses did not measurably impact choice behavior. These data are in remarkable consistency with existing computational theories of decision-making that propose a precise mapping of ventral striatal activity onto the function of the 'Critic.'

Introduction

To survive in rapidly changing environments, animals must select responses among competing alternatives in the pursuit of rewards. In many situations, the evaluation of outcomes from previous actions can driving learning that guides the choice of future responses (Sutton RS 1998; Lee, Seo et al. 2012). One computational solution to adapt behavior in the face of changing contingencies is the Actor/Critic decision-making and learning architecture (Joel, Niv et al. 2002; Takahashi, Schoenbaum et al. 2008). In this framework, an "Actor" selects a response among existing choices and the outcome of the choice is evaluated by a 'Critic.' In particular, the "Critic" evaluates the outcomes in conjunction with the choice to determine if the resulting outcome is better/worse than expected. Based upon this feedback, the 'Critic' can instruct the "Actor" whether an action should be reinforced/suppressed to reflect its motivational value (Sutton RS 1998; van der Meer and Redish 2011).

Many neurobiological accounts of decision-making propose that the ventral striatum is the neural instantiation of the 'Critic' (Houk, Adams et al. 1995; O'Doherty, Dayan et al. 2004; Takahashi, Schoenbaum et al. 2008; van der Meer and Redish 2011; Lee, Seo et al. 2012). Numerous studies have identified signals in the ventral striatum consistent with the value of chosen options (Knutson, Taylor et al. 2005), and this activity is most prominent during the outcome as opposed to the selection of the response (Khamassi, Mulder et al. 2008; Kim, Sul et al. 2009; van der Meer, Johnson et al. 2010; van der Meer and Redish 2011). This representation of the chosen value may represent the evaluation of outcomes to promote learning. While these studies are consistent with a role for the NAc as a 'Critic,' the role of the ventral striatum in critiquing as playing a key role in evaluation is far from dominant. Early conceptualizations of the ventral striatum envisioned this region as a limbic-motor interface for linking motivation to actions (Mogenson, Jones et al. 1980), and more recent re-interpretations upon this idea have

focused on the accumbens role in invigorating reward-seeking responses (Nicola 2007; Salamone, Correa et al. 2009) and attributing motivational salience to reward-predicting cues (Berridge 2007). These proposals suggest that the ventral striatum influences action selection at the time an action is initiated or interrupted. Moreover, electrophysiological recordings do not always yield unambiguous data regarding whether ventral striatal neurons are encoding parameters related to learning that can guide future actions or motivational features that can guide performance (Knutson, Taylor et al. 2005; Ito and Doya 2009).

The heterogeneity of neural representations in the ventral striatum is paralleled by the diversity of cell types in the region. Spiny projection neurons in the ventral striatum either express the D1 receptor (D1R) or the D2 receptor (D2R) with only a small population expressing both (add numbers and citation). These projection neurons represent the sole output of the ventral striatum. A small population of cholinergic interneurons that represents less than 1% of ventral striatal neurons also expresses D2Rs(cite). Recent studies harnessing the power of genetic techniques have suggested that the nucleus accumbens (NAc) may have an opposing process functional organization mediated by distinct classes of projection neurons similar to the direct and indirect pathway found in the dorsal striatum preferences. Recent studies show activation of D1R-expressing NAc MSNs promotes conditioned place preference whereas activation of D2R-expressing NAc MSNs blocks preferences (Lobo, Covington et al. 2010; Lobo and Nestler 2011). Opposing effects on preference obtained in these studies strongly suggest that there may be a functional dichotomy between D1R- and D2R-expressing neurons that needs to be taken into future accounts of ventral striatal function (Lobo and Nestler 2011).

Here, we aimed to determine the role of ventral striatal activity at different times in a rapid decision-making context where animals must assign a value to distinct choices and

repeatedly update subsequent choices. The mice were then trained in a two alternative choice probabilistic switching task in which we have previously demonstrated that mice assign values to different options based upon the outcome of their choices, and they use this value to guide their choice behavior. We leveraged optogenetic techniques in transgenic mice to selectively activate either D1R- or D2R-expressing ventral striatal neurons and delivered temporally specific stimulation either during a period of time when the subjects were evaluating the outcome of their choices or selecting choice responses. Their behavior was fit to a standard decision-making model in the presence and absence of stimulation.

We found that stimulation delivered during the epoch of the task in which subjects were evaluating their choices altered how animals learned from their previous history of rewards and choices to impact their future choices. Consistent with previous studies (Lobo, Covington et al. ; Lobo and Nestler), stimulation of D1R-expressing neurons in the ventral striatum produced a tendency to stay at the previously visited port. Activation of D2R-expressing neurons in the ventral striatum induced a tendency to switch away from the stimulated port on the following trial. Strikingly, the distribution of choices after stimulation mimicked a fixed change in the estimated value of previous choices, suggesting that ventral striatal activity can be mapped onto the probability of future choice behavior. This ability of ventral striatal activity to influence the evaluation of choices was limited to a specific time window around when the outcome of a choice was expected. Moreover, ventral striatal activity during the selection of responses did not seem to impact choice behavior.. These data are in remarkable consistency with existing computational theories of decision-making that propose a precise mapping of ventral striatal activity onto the function of the ‘Critic’.

Results

To identify the role of activity in D1 receptor (D1R) and D2 receptor (D2R)-expressing ventral striatal neurons in various aspects of decision-making, we stimulated each population of neurons during different epochs of a decision-making task. We selectively expressed channelrhodopsin (ChR2) in D1R- or D2R-expressing populations of neurons to independently study the role of activity in these two populations of spiny projection neurons using a Cre-dependent strategy (see Experimental Procedures for details). Another set of subjects in which yellow fluorescent protein (eYFP) was expressed in D1R- or D2R-expressing neurons alone without ChR2 served as controls. We chronically implanted optic fibers above the site of injection in each hemisphere. Histology confirmed the presence of viral expression.

After implantation surgery, D1-Cre or D2-Cre mice were then trained on a spatial two-alternative forced-choice (2AFC) probabilistic switching task, which we previously developed (Figure 10, see Chapter 4). The task required mice to initiate a trial into a central port followed by the presence of a 'Go' cue that signaled the mice to report their choice by approaching either the left or right peripheral port where they can receive water rewards. Optical stimulation was then delivered in different sessions at either an epoch of the task when the subject was either: 1) receiving information about the outcome of its choice at the peripheral port or 2) just about to indicate its next choice after new trial initiation at the central port. All animals underwent stimulation during both conditions, and the order of whether stimulation was delivered at either the center or peripheral ports was pseudorandomly assigned at the start of each experiment. Stimulation was interspersed with sessions when the animals were trained in the task without stimulation, and we alternated the hemisphere that was stimulated in every session.

Modeling the Choice Behavior and Evaluation of Choices

As we have shown previously, the animals' choice behavior in this task is guided by its expectation of reward based upon the outcome of previous choices (Figure 10, chapter 4). Given the structure of the switching task, previously rewarded choices can provide evidence that the subject should return to the same port. On the other hand, an unrewarded event at a previously rewarded port can be interpreted in two ways: 1) a correct response was made and the trial was unrewarded (due to 75% probabilistic reward delivery) or 2) the choice was now incorrect because the task contingencies have now changed. Given this structure, mice implement a win-stay, lose-shift strategy in which rewards serve as evidence to stay at a port while lack of reward serves as evidence to switch (Figure 10, chapter 4).

We then fit a quantitative multivariate logistic regression model to describe the contribution that previous outcomes at each port have on guiding future choices (see Experimental procedures for details). Similar to our previous studies, we found that the contribution of prior rewards declined with additional trials back in time and provided evidence for the animal to stay and the rewarded port as indicated by the positive regression coefficients (Figure 12, chapter 4). Lack of reward only had a small effect and promoted switching only if it occurred in the immediately previous trial as indicated by the negative regression coefficient (Figure 12, chapter 4). We then computed dynamic estimates of value of each choice that were defined as the sum of regression coefficients corresponding to the previous reward history for both the left and right port (Figure 12, chapter 4). The predictive validity of our fit was demonstrated by generating a regression model using 70% of the data and demonstrating that it could recapitulate the actual distribution of choice in the remaining 30% of trials.

Stimulation of Distinct Ventral Striatal Populations during Evaluation of the Choices Promotes Staying/Switching on the Following Trial

We first examined the effects of optical stimulation of either D1R or D2R-expressing neurons of the ventral striatum on the evaluation of choices. Stimulation was delivered on 6% of trials randomly distributed during a session following entry into the peripheral choice port and lasted either 500ms or terminated upon a nosepoke into a central port, indicating the reinitiation of a new trial. Each subject was stimulated across at least four sessions at 20Hz. The fact that stimulation trials were interspersed with non-stimulation trials interspersed allowed for measurement of the effects of ventral striatal activity on behavior independent from changes in satiety and motivation, which may change over the course of a session. Stimulation sessions were interleaved with sessions when the animal was trained without stimulation, and the hemisphere of stimulation was alternated across stimulation sessions. Control animals expressing virally delivered Cre-dependent eYFP were subject to identical training schedules and stimulation protocols as experimental animals.

Following stimulation at the choice port, we monitored choice on the subsequent trial and compared choices made after similar histories of reward and choice, but without stimulation. Differences between the post-stimulation outcome and expected outcome based on reward history revealed a bias induced by stimulation that differed depending on the cell type we stimulated. Stimulation of D1R-expressing neurons at the previously chosen peripheral port induced a bias toward returning to or ‘staying’ with this port on the following trial (Figure 29a). In contrast, stimulation of D2R-expressing neurons at the previously chosen peripheral port induced a bias toward ‘switching away’ from this choice on the subsequent trial (Figure 29 b). This stimulation-induced staying/switching bias on the following trial was not related to the hemisphere of stimulation and was not

necessarily larger for the port contralateral versus the ipsilateral to the hemisphere of stimulation. In contrast to optogenetic stimulation in the dorsal striatum, ventral striatal stimulation does not appear to affect choice behavior in any specific spatial or egocentric reference frame (Figure 29 a,b, 30 c). The magnitude of the staying or switching bias was always greater after unrewarded trials when the animals' responses were more variable (Figure 29 a-d).

To further explore the relationship between stimulation and reward-dependent action selection, we plotted the probability of left choice for different histories following trials with and without stimulation. The plot revealed a 'bowing' of data point off the unity line in a manner dependent upon peripheral port had chosen during the stimulation trial (Figure 29 c,d). The direction of the 'bowing' was in opposite direction for D1- and D2-Cre animals based upon the port where they received stimulation, indicative of their opposing biases on choices during the following trial (Figure 29 c,d). This relationship could be captured by a function where the odds of choosing the left port over the right with stimulation are scaled by a fixed factor known as the odds ratio when compared to the odds without stimulation depending upon which port was previously chosen. In this way, the odds ratio is a means of quantifying the strength of the bias induced by stimulation.

Stimulation of Distinct Striatal Populations during the Evaluation of Choices Mimics a Change in the Value of the Choice on the Following Trial

We then analyzed our data based upon the estimates of value for the left versus the right choice generated from our fits of the impact of previous reward and choice history on the upcoming distribution of choices. Given that stimulation in the ventral striatum could induce staying/switching biases on the following trial similar to the

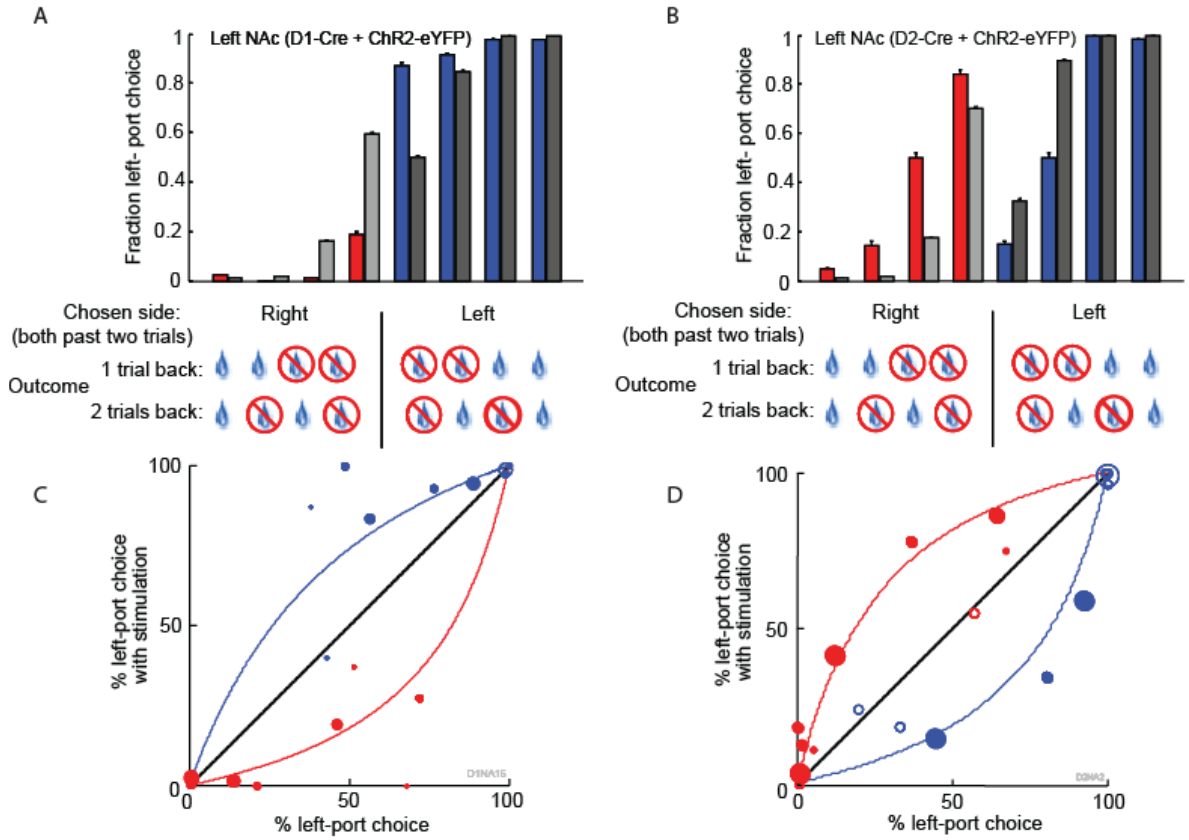

Figure 29: Effect of optical stimulation at outcome evaluation on the distribution of choices for various reward histories.

(**a**,**b**) Examples showing the effect of 20-Hz stimulation in the left nucleus accumbens (NAc) of a D1-Cre mouse (**a**) and a D2-Cre mouse (**b**) expressing ChR2-eYFP on choice behavior on the following trial. Individual bars represent the fraction of left choices for various reward histories in trials in which the mouse previously made two consecutive responses at the same port. Light grey bars indicate trials when the animals previously visited the right port while dark grey bars indicate trials when the animals previously visisted the left port. Red bars indicate stimulation trials at a right port on the previous trialand blue bars represent stimulation stimulation trials at the left port on the previous trial. (**c**,**d**) Fraction of left choices with and without stimulation at either the left (blue) or right (red) port for all possible combinations of choices and outcomes in the previous two trials with more than five total occurrences. (**c**) Data from the D1-Cre mouse shown in **a**. The frequency of trials with a given reward history are indicated by the relative size of the circle. Filled circles represent a significant change in fraction of left choice with stimulation ($P < 0.05$, Fisher's exact test). The blue/red curve relates the probabilities of choice with and without stimulation for a fixed odds ratio. (**d**) Data from the D2-Cre mouse shown in **b**. All error bars represent s.e.m.

presence/absence of water rewards, we hypothesized that ventral striatal activation at the time of outcome evaluation may mimic the effects of a change in the value of the chosen port. This is consistent with the existence of neurons that encode the value of choices in the ventral striatum according to previous studies.

Using estimates of the relative value of the choices, we plotted thee probability of left choices with and without stimulation for various reward and choice histories. Given that stimulation trials were interleaved with non-stimulation trials, we chose to analyze the pooled trials across days. We found that ventral striatal stimulation shifted the sigmoid choice probability curve along the relative value axis based upon the port that the animal chose concurrent with stimulation, mimicking a change in the difference in value for the left versus the right port (Figure 30a,b). Optical stimulation of D1R-expressing ventral striatal neurons resembled a fixed increase in the value of the previously chosen port (Figure 30a) while stimulation of the D2R-expressing neurons imitated a fixed decrease in the value of the previously chosen port (Figure 30b). We say that this change was "fixed" because it interacted with, but was found to simply add (or subtract) to the value of the choice that varied across the range of trial histories.

While the bias was on average present during each session across all of the animals, the effect was larger in early sessions in D1-Cre animals and slight larger after the first session in D2-Cre animals (Figure 30d). Based upon the relative weights of previous rewarded and unrewarded choices, it appeared as if the overall strategy of mice on trials without stimulation was not overall different from days in which the animals received no optical stimulation at all.

Control animals expressing virally delivered Cre-dependent eYFP in either D1- or D2-Cre neurons of the ventral striatum also showed a small tendency to switch away from

previous choices when stimulated (Figure 30e). This indicates that the optical stimulation possibly may itself be slightly aversive in the ventral striatum (Figure 30e). However, this effect was not statistically significant in most subjects, and the effects of optical stimulation on choices in the following trial in animals expressing ChR2 in the ventral striatum were significantly larger (Figure 30e).

We next determined the effect of ventral striatal stimulation on evaluating the outcome of choices on the upcoming choice. After activation of D1 neurons, a staying bias was present one trial later. Thiscould be due to the fact that the effect of D1 stimulation was relatively small compared to the receipt of rewards. In contrast, the effect of D2 stimulation was significant for the subsequent trials before decaying to a level not significantly different from non stimulation behavior (Figure 30f).

We also analyzed whether stimulation during the outcomes of choices could invigorate responding on the following trial. In our early analysis of trials without stimulation, we found that the latency with which animals reinitiate trials following the presentation of the outcome, their latency to withdraw from the central port, and the latency from withdrawal to entry in the peripheral choice port were all indications of the value of the upcoming choice. We analyzed these three types of latencies to determine if stimulation at the time of the outcome could invigorate response latencies in the following trial. In general, the re-initiation times were slowed after stimulation of both D1R-expressing or D2-R expressing populations, perhaps indicating a nonspecific effect of stimulation on movement. However, both withdrawal time from the center port and latency to enter one of the peripheral choice ports on the following trial were

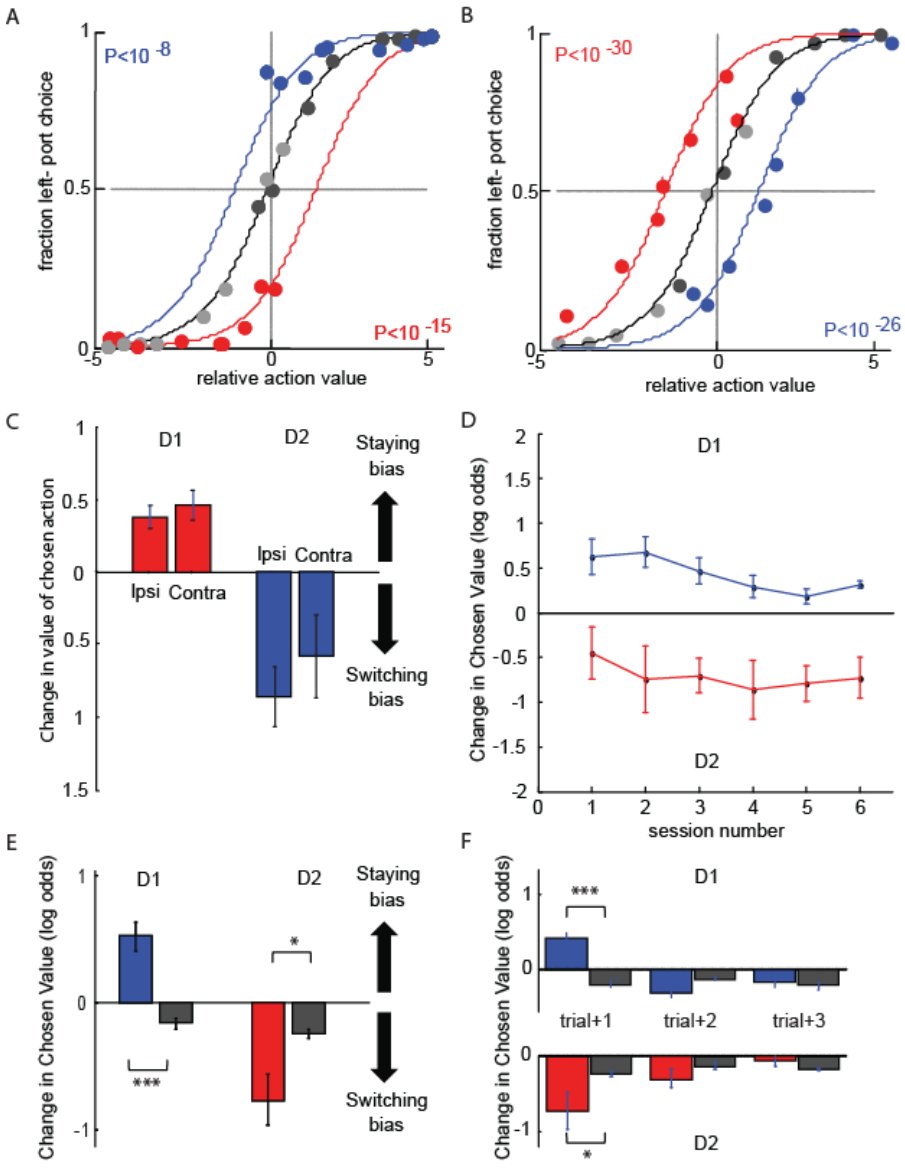

Figure 30: Effect of optical stimulation at outcome evaluation on estimates of chosen values within the task.

Fraction of choices for the left port on trials with different relative action value estimates in D1-Cre mouse (**a**) and a D2-Cre mouse (**b**) in the presence (red) or absence (blue) of optical stimulation. (**a,b**) Representative data from one mouse transduced with AAV-EF1α-DIO-ChR2-eYFP with optical stimulation in the left hemisphere. Logistic regression was used to fit the data from trials with stimulation after a previously choosing the left port (blue line) or right port (red line) and without stimulation (black line). A leftward shift in the logistic curve represents a bias for the left reward port. (**c**) Summary data for the magnitude of change in chosen value following stimulation at the port ipsilateral or contralateral to the port of stimulation, indicating that there is not an obvious spatial bias in D1-Cre (red) or D2-Cre animals, indicating no consistent spatial bias in D1-Cre (red) or D2-Cre animals. (**d**) Summary data for the magnitude of change in chosen value following stimulation across stimulation sessions. **(e)** Comparison of the change in chosen value between stimulated animals (red) expressing ChR2-eYFP and control animals (blue) expressing eYFP. **(e)** Comparison of the change in chosen value between stimulated animals (red) expressing ChR2-eYFP and control animals (blue) expressing eYFP, **(e)** Change in chosen value for one, two, and three trials following stimulation in stimulated animals (red) expressing ChR2-eYFP and control animals (blue) expressing eYFP,

unaffected by stimulation. These two latencies for movement occur after the stimulation has been terminated, which may be why these motor responses were unaffected.

Stimulation of Distinct Striatal Populations during Action Selection Does Not Impact the Distribution of Upcoming Responses or Responses Made on the Following Trial

We next analyzed the effect of stimulation in ventral striatum during an epoch in the task when the animal is selecting their responses, rather than outcome evaluation. Optical stimulation was delivered at the time of decision within the task when the 'Go' cue signals the subject to make a choice to either approach the left or right peripheral ports (Figure 31a). We first analyzed the effect of stimulation based upon different types of reward history. In general, we found no consistent effect of optogenetic stimulation for different reward histories in either D1-Cre (Figure 31b, d) or D2-Cre (Figure 31c,e). We also analyzed the data as a function of our estimates of action value to determine if the animals' choices were altered by stimulation at the central port (Figure 5d and e). We found that across the population, there was no significant effect of stimulation during the time of action selection on choice bias. The effect of stimulation in control animals virally expressing Cre-dependent eYFP in either D1- or D2-Cre expressing neurons in the dorsal striatum also did not demonstrate any significant bias in any of these parameters (Fig 31). Thus, we did not observe any systematic change in choice behavior if ventral striatum stimulation occurred during action selection.

This result was in sharp contrast to the robust changes in behavior choice following stimulation delivered at one of the peripheral ports that we observed in the same set of animals. We therefore considered a second alternative in which animals' response during stimulation may be dependent upon the choices that they made in the previous trial, perhaps indicating that the temporal window when ventral striatal activity can alter

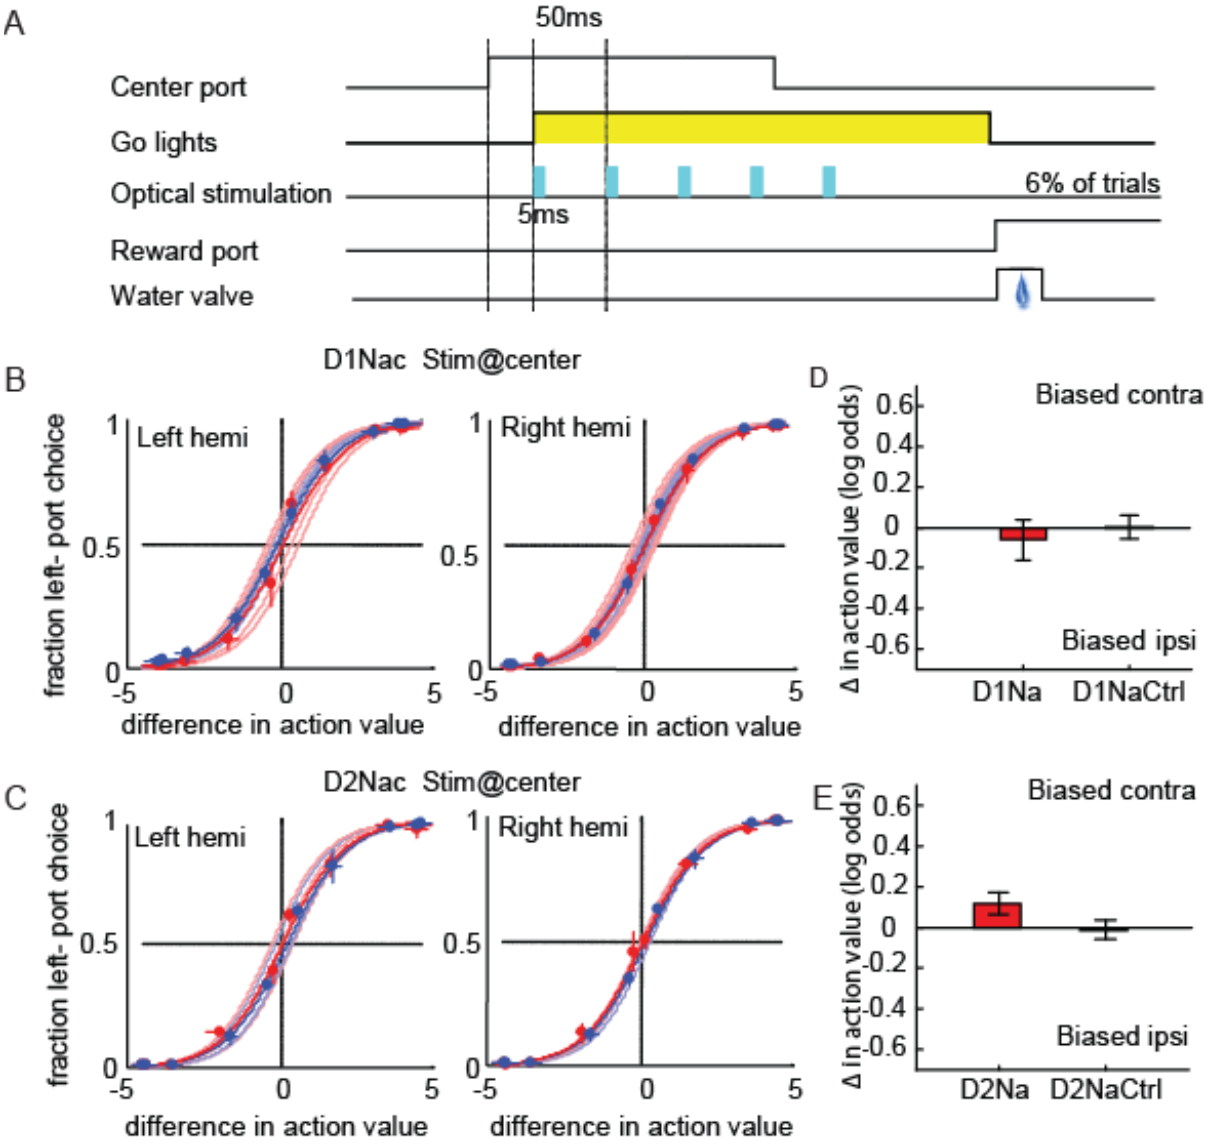

Figure 31: Effect of optical stimulation during action selection on estimates of action values within the task

Timing of optical stimulation in the task. In 6% of trials, optical stimulation was delivered at 20Hz to the ventral striatum during a 500-ms period starting at the same time as the Go light cues. (**b**,**c**) Examples showing the effect of 20-Hz stimulation in the left and right ventral striatum of a D1-Cre mouse (**b**) and a D2-Cre mouse (**c**). (**d**,**e**) There was no significant change in action value with either D1(**d**) or D2 (**e**) ventral striatal stimulation.

animals' evaluation of previous choices may still be "open" as they reinitiate a trial. We therefore analyzed the choice responses after center port stimulation with regard to the port the animals had returned from in the previous trial (Figure 32a). There was no significant change in the bias when animals were stimulated at the center port based upon the port where they had previously chosen (Figure 32c). This suggests the time window in which ventral striatal activity can influence future decision based upon past choice seems to have closed prior to initiation of the next trial.

We then wanted to consider a third alternative, in which optical stimulation from center port withdrawal to choice port entry may affect decision-making on the following trial based upon the upcoming response. In this way, we wanted to try to better define the start of the temporal window when ventral striatal activity can alter the evaluation of the choice as reflected by future responses. In many in vivo recording studies, the epoch in which an animal is approaching a potential reward site has frequently been identified as a period when ventral striatal neurons "ramp up" their firing rate, reaching a peak at the expected outcome delivery (Kim, Sul et al. 2009; van der Meer, Johnson et al. 2010; van der Meer and Redish 2011). We therefore analyzed the choice responses on the trial following center port stimulation as a function of recent reward and choice history (Figure 32b). We found that there was no significant change in the bias for choices made on the trial following optical stimulation from center port withdrawal to choice port entry (Figure 32d). Therefore, these analyses revealed that the temporal window for the evaluation of a choice in our task is limited to a period from choice port entry to re-initiation of the following trial.

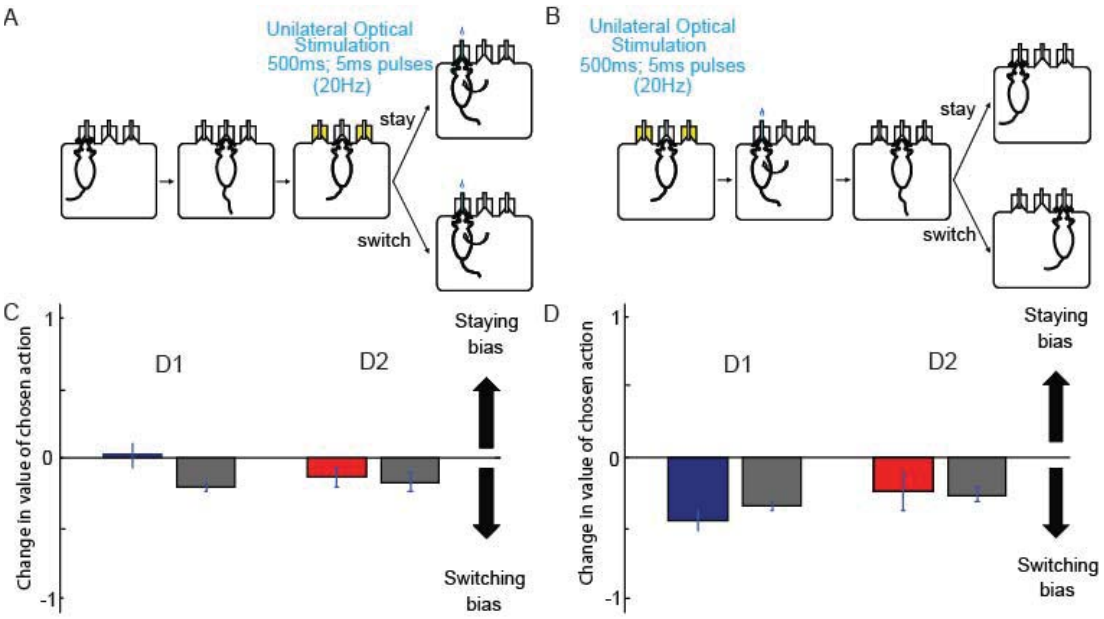

Figure 32: Effect of optical stimulation during action selection does not change responses for previously chosen or upcoming actions
(**a,b**) Series of events in the switching task. (**c,d**) Chosen values were measured relative to either the choices performed prior to stimulation of either D1 neurons (blue) or D2 neurons (red) (**c**) or following stimulation (**d**) at the center port.

Discussion

Decision-making depends upon the ability of motivation to guide performance and learning driven by past experiences. This distinction between performance and learning is embodied in 'Actor-Critic' computational models. 'Actor-Critic' models assume that task-relevant information can discretely be described as states, which generically correspond to any discrete external event, such as a cue light, or a particular context/spatial location (Sutton RS 1998; Takahashi, Schoenbaum et al. 2008; Lee, Seo et al. 2012). Because representations of states are internal to the agent, they can also correspond to internal perceptual decisions, features stored in working memory, or recent histories of choices and rewards (Nakahara, Itoh et al. 2004; van der Meer and Redish 2011). In our task, entry into each port is analogous to a state. Values are then associated with these states and correspond to an estimate of the anticipated future reward for different states. Because entry into the ports is associated with rewards, the choice to enter into one of the peripheral ports can be assigned a state value. We estimated these state values by using logistic regression analysis to understand how previous rewards and choices affect future decision-making based upon the structure of the task.

Here, we find an agreement between experimental results and theoretical desciptions of the ventral striatum's role as a 'Critic,' promoting learning to guide future decision-making. Stimulation of D1R-expressing neurons at the outcome evaluation produced a tendency to stay at the previously visited port on the following trial whereas activation of D2R-expressing neurons produced a tendency to switch. To better understand the nature of these manipulations on the value that our subjects placed upon prior choices, we then attempted to fit the animals' choice behavior to a quantitative model. We found activation of either D1R- or D2R-expressing neurons mimicked a fixed change in the value of the previous choice (state) on the distribution of choices that the animal made

on the following trial. This is consistent with previous electrophysiological data, suggesting that the ventral striatum represents the value of choice (Knutson, Taylor et al. 2005), which is most prominent during the outcome as opposed to the selection of the response (Khamassi, Mulder et al. 2008; Kim, Sul et al. 2009; van der Meer, Johnson et al. 2010; van der Meer and Redish 2011).

Actor-critic models also offer a description of how the 'Critic' is able to inform the 'Actor' regarding, which actions are required to achieve highly valued states. The 'Critic' computes the difference between the estimated value of a state and its actual outcome in the form of a reward prediction error. This prediction error then trains a separate actor regarding the value associated with particular actions to produce a bias for more highly valued motor responses. In this way, the 'Critic' facilitates learning whereas the 'Actor' invigorates the performance of actions of high value. This process continues iteratively until there is no discrepancy between the estimated and actual delivered reward. In relation to our experimental results, ventral striatal stimulation can be interpreted as generating a reward prediction error signal from 'the 'Critic' to update the distribution of choices of the 'Actor' on the following trial. The fixed change in value induced by stimulation of either D1R- or D2R-expressing populations of neurons implies that the reward prediction error is incrementally changing value in a uniform way across various reward and choice histories.

These data are complementary and shed light on previous imaging and electrophysioglogical experiments, which have found a diverse collection of representations in the ventral striatum. Many studies report the presence of "chosen value" signals, which do not encode any particular action or transition between states, but rather correlate with magnitude of a predicted upcoming rewards (Kim, Sul et al. 2009; Roesch, Singh et al. 2009). This "chosen value" encoding within neurons often

becomes more prevalent prior to reward delivery, ramping up as the subject approaches or waits for an anticipated outcome (Kim, Sul et al. 2009; Roesch, Singh et al. 2009; van der Meer, Johnson et al. 2010). Many studies have shown that this activity is independent of a particular behavioral response (Kim, Huh et al. 2007; Kim, Sul et al. 2009; Kimchi and Laubach 2009) or are more closely temporally associated with the approach of the outcome. Our data are consistent with a model in which the ventral striatum represents the value of choices to evaluate outcomes.

Another possible interpretation consistent with our stimulation results is that the ventral striatum may represent a reward prediction error teaching signal, generated by the 'Critic.' While this is possible and evidence for this form of activity can be observed in fMRI studies, most electrophysiological studies do not observe a large population of neurons that appear to directly represent signals consistent with a reward prediction error (Pennartz, Berke et al. 2009; van der Meer, Johnson et al. 2010; van der Meer and Redish 2011). Instead, these recording data along with our own stimulation studies instead suggest a model where the ventral striatum represents chosen values to promote learning and guide future behavior. The ability to directly manipulate neural activity in the ventral striatum is critical to establish the function of ventral striatal neurons given that this ramping form of activity could instead be interpreted as a motivational signal that energizes motor responses to approach or withhold a motor response prior to a reward.

These data are also consistent with a range of recent papers that have demonstrated that the ventral striatum has a similar opposing process function organization to the dorsal striatum (Lobo, Covington et al. 2010; Lobo and Nestler 2011; Kravitz and Kreitzer 2012; Kravitz, Tye et al. 2012). Interestingly, these previous experiments have demonstrated that D1-expressing neurons of the ventral striatum promote place

preference for cocaine rewards while D2-expressing neurons may facilitate de-valuation of preference, but that stimulation alone does not induce a preference. Previous functional studies have also implicated the existence of a direct and indirect pathway through the ventral striatopallidal system (Maurice, Deniau et al. 1997; Maurice, Deniau et al. 1999). While we propose here that the ventral striatum may play a key role in evaluating the outcome of choices in the context of decision-making, it is almost certainly true that the ventral striatum also support numerous other functions involved in reward-seeking, motivation, and palatability as well. Numerous previous studies in Pavlovian and operant tasks have found a large number of neurons that fire at the time of reward-predictive cues, motor responses, and reward consummation (Taha and Fields 2005; Ambroggi, Ishikawa et al. 2008; Ishikawa, Ambroggi et al. 2008; Ambroggi, Ghazizadeh et al. 2011; van der Meer and Redish 2011). Although, it is interesting to speculate that some of this activity may also be involved in encoding the value of a particular task-relevant state and the formation of a reward prediction error.

From an anatomical perspective, the ventral striatum is ideal for playing the role of the ‘Critic’ because it receives inputs from brain regions such as the amygdala, hippocampal, thalamic, and prefrontal cortex that are known to encode task relevant information and value information. The ventral striatum and dopaminergic neurons of the midbrain act as an interconnected system (Nauta, Smith et al. 1978; Alexander and Crutcher 1990). Midbrain dopamine neurons densely innervate both the dorsal and ventral striatum. These dopamine neurons are known to encode a reward prediction error (RPE) signal and are likely to function as part of the ‘Critic’ in concert with ventral striatum (Sutton RS 1998; Niv and Schoenbaum 2008). In fact, there is evidence that the ventral striatum may also provide dopamine neurons with information required to generate RPEs (Haber, Fudge et al. 2000). Moreover, activation of ventral striatal subpopulation may update state values through the ventral striatopallidallal-

thalamocortical loops (Alexander and Crutcher 1990; Everitt and Robbins 2005). These anatomical pathways may serve as the neural substrates for the 'Critic' from a computational perspective of reinforcement learning theory as well (Joel, Niv et al. 2002). One way in which learning can be facilitated by ventral striatal activation is through the updating of synaptic weights. This change in weights may be mediated by instructive reward prediction errors generated by dopamine release and subsequent plasticity at corticostriatal synapses (Shen, Flajolet et al. 2008; Pennartz, Berke et al. 2009; Gerfen and Surmeier 2011). Alternatively, information can be stored in working memory signals otherwise known as "eligibility traces" related to chosen states or actions (Sutton RS 1998). Indeed, previous experiments have identified these signals in the prefrontal and posterior parietal cortex of monkeys (Seo, Barraclough et al. 2009) and the frontal cortex (Fecteau and Munoz 2003; Barraclough, Conroy et al. 2004; Sul, Kim et al. 2010) and striatum of rodents (Kim, Sul et al. 2009). These regions are also recruited during serial reversal tasks in both prefrontal cortex and striatal systems.

If the ventral striatum function as a component within the 'Critic,' the question then naturally arises: What is the corresponding neural substrate for the 'Actor'? A likely candidate for the 'Actor' is the dorsal striatum (O'Doherty, Dayan et al. 2004). Using the same probabilistic switching task and model, we previously demonstrated that stimulation of D1 and D2-expressing neurons of the dorsal striatum during an epoch of the task when animals are selecting their responses can bias choice behavior (Tai, Lee et al. 2012). The bias in choice behavior induced by stimulation mimics a fixed additive shift to the action value (Tai, Lee et al. 2012). These results are consistent with notions that the dorsal striatum functions as an 'Actor' in computational theories of decision-making that represents the value of particular actions to generate a bias for actions of higher value and invigorates the performance of these responses. Together, these studies provide evidence for an 'Actor-Critic' system within the dorsal and ventral

striatum, respectively (Sutton RS 1998; Niv and Schoenbaum 2008). Ventral striatal systems can serve as a neural substrate for reinforcement and outcome evaluation while dorsal striatal systems can mediate action selection. In the ventral striatum, activity may represent “chosen values,” which guide learning and reinforcement, whereas activity in the dorsal striatum may primarily represent “action values,” which guide performance and decision-making (Kim, Sul et al. 2009). Together these studies imply that, both systems use a “common currency” of value in executing their respective computations (Sugrue, Corrado et al. 2005).

Materials and Methods (chapter 4 & 5)

Animals

C57BL/6J BAC transgenic mice expressing Cre recombinase under the regulatory elements for the D1 and D2 receptor (D1-Cre and D2-Cre ER43) were obtained from MMRC and bred in our colony. All animals used in this study were adults (25-30g) and group housed under a reverse 12 h light/dark cycle (light onset at 10:00 A.M.) until surgery. Mice were given food and water *ad libitum* prior to water deprivation in preparation for training. All procedures were approved by the Ernest Gallo Clinic and Research Center Animal Care and Use Committee.

Construct and virus preparation

Plasmids encoding the DNA sequences for pAAV-EF1α -DIO-ChR2(H124R)-eYFP or pAAV-EF1α-DIO-eYFP were obtained from the laboratory of Karl Deisseroth. Amplification and purification of plasmids was performed using a standard plasmid maxiprep kit (Qiagen) and confirmed by sequencing. EF1α-DIO-CHR2(H124R)-eYFP and EF1α-DIO-eYFP cassettes were packaged in AAV vectors and serotyped with AAV5 coat proteins by the viral core at University of North Carolina. The final concentration was 1-2x10^{12} viral particles ml^{-1}.

Implantable chronic optical fibers and optic cables construction

Optical fibers were constructed by attaching a 200 □m, 0.37 NA optical fiber (Thor Labs) with epoxy resin into a metal ferrule that had previously been cut and

scored. Fiber-ferrule units were then cut and polished. After construction, implants were tested to determine their efficiency. Only implants with efficiency greater than 70% and comparable efficiencies (+/- 10%) were used in the study. Optical-patch cables were constructed from 62.5 □m core diameter optic fiber (Thor labs) that were connected to a ferrule on one end and an FCPC connector on the other end. Cables were covered in furcation tubing to protect the fiber and to prevent light from escaping through the optic-patch cord. The ferrule at the end of the optic patch cord was fitted with a zirconium sleeve to interface with the chronic implant. The FCPC connector was coupled to a 473-nm DPSS laser (200□mW).The laser driver current was adjusted to yield 20mW output from the patch cable.

Stereotaxic AAV injection and optical implant surgery

Animals were anesthetized with either 150mg/kg ketamine and 50mg/kg xylazine or 2% isofluorane gas anesthesia. Animals were placed in a stereotaxic frame and 26-gauge microinjection needles were inserted through a burr hole bilaterally into the dorsomedial striatum (coordinates from Bregma: 1.30 AP; +/- 1.2 ML; -4.1 DV) of D1/D2-Cre mice to deliver 0.5□L of either AAV-EF1-DIO-ChR2(H124R)-eYFP or AAV-EF1-DIO-eYFP (~1012 IU/mL). Injections were performed using a 1-□L Hamilton syringe through a hydraulic pump (Harvard Instruments) and took place over 10 min followed by 10 min of recovery. The length of the optic fiber protruding from the implant was cut to be 2mm. The tip of the fiber optic from the implant was then inserted through the same burr hole as was used previously for the virus injection and was lowered 2mm ventral to the dura. The implant was cemented to the skull using dental cement. Mice were

then returned to their home cage and monitored until recovery from surgery.

Anesthetized *in vivo* optrode recording

Recordings under conditions of anesthesia were made with a custom silicon probe (model A1x16-5mm50-413, NeuroNexus Technologies) that was attached with a 200□m diameter core optic fiber connected to a FCPC connector with epoxy resin. One month following injection of AAV-EF1-DIO-ChR2(H124R)-eYFP into the dorsomedial striatum, animals were placed under ketamine and xylazine anesthesia (150mg/kg ketamine and 50mg/kg xylazine i.p.). A craniotomy was performed above the injection site and the optrode was lowered. Data was acquired using commercial systems (Plexon). Optical stimulation was delivered during recording as continuous trains of stimulation to assess the fidelity of optically induced firing as well as using the same stimulation protocols that we used during our behavioral paradigm. After each recording, the probe was moved to a new recording tract within the same animal. Following recordings, animals were euthanized for assessment of track location and viral expression. Single units were identified with principal component analysis (Offline Sorter, Plexon). The data were then imported into Matlab (MathWorks) for subsequent analysis.

Probabilistic switching (two-alternative spatial choice) task

Mice were trained on a two-alternative spatial choice task in which the location of a water reward was periodically switched at random intervals. The initiation port was located in the middle of one wall, and two choice ports were located 63.5 mm left and right of the initiation port (center-to-center; **Fig. 1a**). An infrared photodiode/phototransistor pair placed on either side of the port to report the

times of port entry and exit (Island Motion, Tappan, NY). The water valves (Neptune Research) were calibrated to deliver a volume of water (2 µl) for rewarded choices.

Mice initiated each trial by entering the center port, triggering “go lights” instructing animals that water was potentially available. Mice then chose a left or right peripheral port for water reinforcement (**Fig.1a**). Only one peripheral port was rewarded at a time, on 25% of trials, neither port was rewarded. The length of trial blocks was dependent on the number of rewards obtained in each block, and this number of rewards was randomly set between 7-23 rewards. After the set amount of rewards was obtained, the rewarded side was switched to the opposite port. This structure prevented the subjects from predicting the timing of the block switch.

Logistic regression analysis of behavior choices

The contribution of past rewards or lack of rewards on the subject’s current choice was analyzed on a trial-by-trial basis using the following logistic regression model (Lau and Glimcher 2005):

$$\log(\frac{P_R(i)}{1-P_R(i)}) = \sum_{j=1}^{n} \beta_j^{\mathrm{Reward}} (Y_R(i-j) - Y_L(i-j)) + \sum_{j=1}^{n} \beta_j^{No-\mathrm{Reward}} (N_R(i-j) - N_L(i-j)) + \beta_0$$

where $P_R(i)$ is the probability of selecting the right port in the *i*-th trial. The variables $Y_R(i)$ or $Y_L(i)$ represent whether a reward is delivered (1 or 0) at the right or left port in the *i*-th trial, respectively. And $N_R(i)$ or $N_L(i)$ represent the lack of reward (1 or otherwise 0) at the chosen right or left port in the *i*-th trial,

respectively. The variable n indicates the number of past trials that were included in the model (n=5). The regression coefficients $ß_j^{Reward}$ and $ß_j^{No\text{-}Reward}$ represent the contribution of past rewards and lack of rewards, respectively, and $ß_0$ indicates the intrinsic bias of the animal.

We modeled the contribution of optical stimulation on the subject's current choice as a dummy variable characterized by the $ß_{stim}$ coefficient:

$$\log(\frac{P_R(i)}{1-P_R(i)}) = \sum_{j=1}^{n} \beta_j^{\mathrm{Re}\,ward}(Y_R(i-j) - Y_L(i-j)) + \sum_{j=1}^{n} \beta_j^{No-\mathrm{Re}\,ward}(N_R(i-j) - N_L(i-j)) + \beta_0 + \beta_{stim} \cdot X_{stim}(i)$$

The variables $X_{stimulation}(i)$ represents whether stimulation is delivered (1) or not (0) in the *i*-th trial. These estimated coefficients are the shifts in action value characterized in figure 6.

Optical stimulation in the behavior task

Optical stimulation was delivered at the start of the 'Go' signal. Stimulation was delivered at 20 Hz for 500 ms with the frequency pseudo-randomly chosen prior to each session. Stimulation trials occurred in 6% of total trials. Stimulation sessions were performed every other day interspersed by training sessions. The hemisphere that was stimulated was alternated across stimulation sessions. The infrequent occurrence of stimulation trials was to prevent any plastic or compensating adaptations from occurring during the course of a session, and the relatively long interval of days between stimulation sessions was to prevent any systemic biases in responses from arising across stimulation sessions.

A subset of mice that underwent stimulation at 20 Hz underwent a second experiment in which two types of stimulation trials were present. In one set of

stimulation trials, stimulation at 10Hz occurred in 3% of trials at the same time that the ‘Go’ light appeared 15ms after initiation of a trial by a center poke. In another set of stimulation trials, stimulation once again occurred at 10Hz in 3%, but was delayed by 150ms from the appearance of the ‘Go’ light, occurring in >90% of trials after movement initiation. This was to determine whether stimulation had to occur within a specific time-window relative to movement initiation.

A new cohort of mice (n=6) were a variant of the stimulation experiment in which two optical pulses (IPI=50ms) were delivered prior to and coincident with a ‘Go’ tone which was delayed by 70-100ms after initiation of a trial by a center poke. This was to confine stimulation to an epoch of the task prior to the initiation of the animals movement as determined by the time of withdrawal from the center poke. Stimulation trials occurred in 6% of total trials. Stimulation sessions were performed every other day interspersed by training sessions.

Statistical analysis

Fisher’s exact test was used to determine whether the probability of choices with and without stimulation were significantly different. Logistic regression was used to fit data for trials with different reward histories with and without stimulation. T-test was used to determine whether the change in relative action value caused by striatal activation ($ß_{stim}$) was significantly different from zero for a given stimulation condition. Wilcoxon sign rank and rank-sum tests were used to determine whether the changes in relative action value between different stimulation conditions or between groups of subject expressing ChR2-eYFP and

eYFP were significantly different. Standard errors for probabilities of choice were calculated based upon binomial statistics. α was set at 0.05.

Histology and reconstruction of optical stimulation sites

Viral expression of ChR2-eYFP and eYFP was confirmed by histology after stimulation experiments (**Supplementary Fig. 2a,b**). Animals were perfused with saline and 4% paraformaldehyde, and brain tissue was fixed for subsequent coronal sectioning. These sections were then stained to identify cell bodies (Neurotrace) alongside AAV driven expression of YFP. In the cases reported, fiber implant tracks could be identified and were found to be located in the dorsal striatum (**Supplementary Fig. 2c**).

To determine whether ChR2-eYFP was expressed in medium spiny neurons as well as cholinergic neurons of the striatum, sections (coronal, 50 □m) were permeabilized in 50% alcohol for 10 min, rinsed in PBS, blocked in 10% normal donkey serum (NDS) for 30 min and incubated for 48 hrs in a mixture of primary antibodies: rabbit polyclonal anti-GFP (for YFP; 1:10,000, Abcam Inc., Cambridge, MA); goat polyclonal anti-ChAT (1:500; Millipore, Temecula, CA); and mouse monoclonal anti-Kv2.1 (1:200; UC Davis and NIMH NeuroMab, Davis, CA) at 4^{o}C (Ariano, Cepeda et al. 2005). Sections were then rinsed in PBS and incubated in 2% NDS for 10 min, then incubated for 4-6 hours in mixture of secondary antibodies (all made in donkey): Alexa Fluor® 488 anti-rabbit (1: 300; Life Technologies; Carlsbad, CA); CF™647 anti-goat and CF™555 anti-mouse (1, 300; Biotium, Inc. Hayward, CA). Sections were rinsed in PBS, mounted and coverslipped with VECTASHIELD® mounting media

(VECTOR LABORATORIES, INC. Burlingame, CA). Multi-channels images of ChaT-positive cells were acquired using Zeiss LSM 510 META laser confocal microscope (Zeiss, Thornwood, NY), using 63x/1.4 NA PlanApo objective, 488 543 and 633 nm excitation lines and factory recommended detector settings.

Chapter 6

Significance of dissertation and remaining questions

Locomotion is a fundamental behavior shared by all animals, guiding goal-directed approach towards desired outcomes and avoidance of aversive stimuli. To this end, a large number of neural processes are regulated by and serve to guide the locomotor behaviors in animals. Here, we made efforts to define the neural circuits underlying locomotor control, changes in brain states associated with locomotion, and locomotor decisions.

In Chapter 2, we identified brainstem circuits that concurrently initiate locomotion and concurrently regulate cortical visual processing, perhaps through a pathway involving the basal forebrain. These findings may elucidate a general mechanism by which various brain networks are modulated by behavioral state. In Chapter 3, we demonstrated that these brainstem circuits are under the regulation of the basal ganglia. These studies identified phylogenetically conserved pathway for guiding locomotion. In chapter 4, we leveraged our understanding of the basal ganglia pathways for locomotor control to understand the processes of goal-directed decision-making, and in chapter 5, we found that the ventral striatal shares a parallel organization to the basal ganglia in implementing reinforcement learning to guide future locomotor decision-making. Collectively, these results demonstrate control systems for locomotion are deeply interconnected with a diverse array of processes throughout the brain, perhaps, with the ultimate purpose of guiding goal-directed locomotor behaviors.

Remaining Questions:

While the investigations described here primarily focus on the neural basis of goal-directed locomotion as well as the regulation of brain states associated with movement, we believe these studies can serve to guide future physiological investigations into diverse fields of interest. Here, I potential open-ended questions and projects:

- Brain states associated with locomotion in rodents may be analogous to behavioral states of alertness in human subjects
- A general mechanism for brain state regulation by locomotion
- Elaborating the Networks for Goal-Directed Locomotion and Navigation in Rodents
- Using reward-based tasks and reinforcement theory to understand the function of cortico-striatal loops
- Methods for dissecting the cognitive and motivational processes underlying locomotor decision-making
- A “common curreny” for decision-making

<u>Brain States Associated with Locomotion in Rodents May Be Analogous to Behavioral States of Alertness in Human Subjects</u>

In chapter 2, we hypothesize that the MLR/PPTg within the reticular formation may mediate changes in processing states that are associated with locomotion through its widespread projections in rodents. Interestingly, we have observed that MLR/PPTg activation can shift LFP power from low frequencies to the high frequency gamma oscillations, which are accompanied by an increase in the gain

of visual responsiveness. These same characteristics also occur during locomotion in mice.

However, is the PPTg/MLR's regulation of cortical state and behavioral state unique to mice? Clinical studies involving the pontine midbrain can help elucidate these questions. As we described earlier, the PPTg is a site for DBS in Parkinson's patients. Patients often self-report that they feel "alert" upon the onset of low frequency DBS, and instances of high frequency DBS are known to induce behavioral signs of sleep and a rhythmic state of slow wave oscillations, which is often associated with non-REM sleep (Alessandro, Ceravolo et al. 2010; Arnulf, Ferraye et al. 2010). One explanation for these divergent effects is that low frequency DBS may serve to activate regions and associated with the implanted electrode while high frequency DBS may shut down activity. Thus, it seems reasonable to speculate that the behavioral states associated with locomotion in mice may have parallels to behavioral states for regulating alertness and awake states in human subjects. PET scans have demonstrated that the reticular formation is active in human subjects engaged in a task where they are utilizing visual information to guide behavioral choices, but not in situations where they are receiving the same stimuli, but are not required to act upon it (Kinomura, Larsson et al. 1996). Thus, it may be that these brainstem systems for brain state regulation are recruited to serve different behavioral demands across species (Harris and Thiele 2011). However, these regulatory systems may regulate brain states in a similar fashion through their ascending projections from the reticular formation to cholinergic systems of the basal forebrain.

There is a body of work to suggest the human PPTg has a role to play in gating of external visual stimuli in human patients analogous to the enhancement in sensory-evoked cortical responses in mice. Lesions of the PPTg in human patients can induce peduncular or pontine hallucinosis, which is marked by vivid visual hallucinations that typically occur in dark environments and last for minutes (Benke 2006). These hallucinations are often very realistic and can involve familiar people and environments. Patients with PD, narcolepsy-cataplexy syndrome, and Lewy Body Dementia, and temporal lobe epilepsy are also prone to complex visual hallucinations, and there is some suggestion that these disorders may have a common pathogenesis to lesions of the brainstem (Benke 2006). In fact, PD and Lewy Body Dementia are often associated with a-synuclein plaques and neurodegeneration within the pontine midbrain, which precedes the loss of midbrain dopamine neurons (Boeve 2013). The concomitant dysfunction in sleep that may be related to the PPTg's role as an REM-On sleep center and a generator of ponto-geniculo-occipital (PGO) waves. In fact, recent clinical studies have reported an unusually high incidence of sleep disorders and visual hallucinations in early or pro-dromal stages of these synucleinopathies, which are now being increasingly associated with lesions of these brainstem regions.

Many theories of cholinergic function propose that acetylcholine may serve as a general mechanism to enhance representations of the external world at the expense of internal models(Hasselmo 2006). The presence of visual hallucinations following lesions of the cholinergic brainstem are consistent with a

role for acetylcholine in gating external sensory representations. Future clinical studies may help to further elucidate the mechanism of the changes in subjective "alertness" that result from PPTg activation. In our studies, we suggested that ascending pathways to the cholinergic basal forebrain may mediate the changes in LFP and gain that we observed with MLR/PPTg stimulation. It would be interesting if analogous changes in cortical state also occurred in patients receiving DBS in the PPTg, which can be measured by non-invasive EEG. Parkinson's patients often are prescribed anti-cholinergic drugs as second line therapy after the side effects associated with dopamine replacement therapy are no longer tolerable. As a routine procedure, patients often undergo a period of time in which they are receiving DBS, but are taken off their medication to determine the lowest possible dose of medications that they require. It would be interesting to survey patients before and after their use of anti-cholinergics with and without PPN DBS to elucidate whether subjective reports of alertness, resulting from DBS, are perhaps mediated by cholinergic pathways. Moreover, EEG represents a non-invasive means of identifying changes in brain states that may correlate with these reported changes in alertness.

A General Mechanism for Brain State Regulation by Locomotion

Locomotion is often accompanied by changes in information processing states across the brain. In Chapter 2, we proposed that these diverse changes may be mediated by activation of the MLR/PPTg region. Interestingly, even within the cortex, there seems to be markedly divergent changes in neuronal encoding during locomotor states. For example, while sensory evoked neural responses in the visual cortex are enhanced by locomotion, there have been recent reports

that sensory signals are suppressed in the auditory cortex, mainly in superficial layers of cortex (personal communication from Wehr/McCormick lab). Interestingly, activity in superficial layers is also suppressed during stimulation of the MLR/PPTg in anesthetized animals(Sakata and Harris 2012). There is also a larger literature to suggest that the MLR/PPTg may serve an important role in the gating of behavioral responses to auditory stimuli (Swerdlow, Geyer et al. 2001). It would be interesting to determine if similar pathways from the MLR/PPTg regulate auditory processing, utilizing similar methods to what has been described here.

In addition, we can make predictions regarding how locomotion may alter physiological responses based upon previous results involving MLR/PPTg stimulation. It has been demonstrated that electrical stimulation of the MLR/PPTg region in cats can induce a transition from a burst to tonic mode firing within the lateral geniculate nucleus through cholinergic neuromodulation. This change in thalamic state allows visual stimulation to more faithfully and linearly represent sensory information from the periphery (Lu, Guido et al. 1993). In turn, the PPTg makes a very dense projection to LGN, consisting of up to 40% of all synaptic input into the region. In contrast, retinal input may only consist of 20% of all synaptic inputs with the remaining 40% arriving from cortex(Erisir, Van Horn et al. 1997). Based upon this, we would predict that locomotion may facilitate a transition in the thalamus from a burst to a tonic mode of firing. Indeed, preliminary reports from colleagues have identified that this transition from burst to tonic does indeed occur when animals are locomoting (personal communication from Niell lab).

Although the borders of the PPTg, which resides within the MLR, are classically defined by the extent of its cholinergic cell population, only ~30% of neurons within the PPTg are actually cholinergic. The largest population of neurons within the PPTg is the glutamatergic neurons, which compose of ~40% of neurons within the nucleus, with GABAergic neurons make up the remaining ~30% (Wang and Morales 2009). In order to clearly delineate these cell populations and understand their relationship with basal ganglia output nuclei, further studies will need to be performed to understand how the neurotransmitter identity of these neurons overlays on there actions through ascending pathways to the forebrain and diencephalon and descending pathways to motor outputs. Based on our preliminary data, we expect that activation of glutamatergic neurons in the PPTg, but not cholinergic or GABAergic neurons, will increase locomotion. However, it is possible that other subtypes of neurons play distinct role in other aspects of behavior, such as muscle tone. We also expect that during normal (spontaneous) locomotion, the neurons that become active are glutamatergic or cholinergic. It still remains unclear whether activation of the cholinergic neurons or some other neuronal subtype can drive locomotion and hence the benefits observed in PD patients (Thankachan, et al. 2012). If the some other neuronal subtype were to drive locomotion, then these neurons would be a better indicator of the boundaries of the functionally defined MLR and may mediate the therapeutic benefits observed with PPN DBS. Thus, better defining the cell-type within the MLR that mediates locomotion and how it interacts with the basal ganglia will improve our understanding of the mechanisms of PPN DBS and the pathophysiology underlying the axial symptoms of PD. Further efforts will also be

needed to define the cell-types within the MLR/PPTg region that may be involved in regulating brain states.

Elaborating the Neural Circuits for Guiding Goal-directed Locomotion and Navigation in Rodents

In the course of this thesis, I have identified neural circuits that mediate locomotor behaviors, the selection of locomotor responses, and the evaluation of outcomes, resulting from these responses. In this way, we have begun a project to outline the brain networks that guide goal-directed locomotion and navigation in rodents. Here, we discuss future directions, which can serve to guide future efforts to develop a more complete description of locomotor decisions.

We believe future attempts to identify neural circuits for guiding locomotor decisions would greatly benefit from ongoing investigations into the neural representations of places, spatial representations, and contexts within the hippocampus and entorhinal cortex. Indeed, there is a long tradition that has sought to understand the basis of locomotor decisions in maze learning contexts in rodents. In Chapter 4 and 5, we utilized a common 3-port spatial task design to investigate mechanisms of action selection and reinforcement learning. In the context, the spatial location of the ports represented states that were being reinforced and selected. These representations for spatial location are undoubtedly encoded in the hippocampus among other places in the context of our task. It is therefore interesting to contemplate how this information regarding spatial location can be routed to systems for action selection and reinforcement learning.

Recent studies have demonstrated that subicular inputs, representing the output of the hippocampus, into the ventral striatum are sufficient to drive conditioned place preference and intracranial self-stimulation reinforcement (Britt, Benaliouad et al.). Interestingly, ventral hippocampal inputs have a preference for forming connections onto D1-expressing neurons within the ventral striatum as opposed to D2-expressing neurons in contrast to other afferent inputs (MacAskill, Little et al. 2012). D1 activation, in turn, promoted reinforcement in the context of our decision-making task. This connectivity may serve as the basis of how information regarding spatial location, which serves as the state in a navigation task, can be delivered to a 'Critic' that functions in outcome evaluation. These studies can provide a more mechanistic view regarding various components involved in locomotor decision-making.

Previous studies have demonstrated that stimulation in the MLR/PPTg can induce synchrony in the range of gamma frequency in the firing of neurons across hemispheres of cortex in response to visual stimuli in the binocular zone of visual cortex (Munk, Roelfsema et al. 1996). If the MLR is also recruited during locomotion, then we would predict that there should be interhemispheric synchrony also when animals are running, and possibly across extrastriate areas of cortex representing the same visual field. In general, gamma oscillations have been hypothesized to create a temporal structure, which can give rise to co-incident input that can effectively drive downstream areas. While it is unclear whether this synchrony is a cause or consequence of enhanced excitatory drive, the changes in brain states that are associated with locomotion may represent a

physiological context in which the role of synchrony in the gamma range and questions regarding the function of brain states may be studied.

As we mentioned earlier, the MLR/PPTg region appears to be critical for regulation of brain states. Perhaps, the most fruitful efforts to understand these changes in brain states and synchrony associated with locomotion is in the context of neural circuits sub-serving spatial navigation in the hippocampus and entorhinal cortex. In general, spatial navigation is thought to employ a combination of at least two strategies: the use of visuosaptial landmark cues and path integration (Buzsaki and Moser 2013). Path integration requires utilizes the speed and direction of movement to iteratively compute the position of animals. The ability of these parameters to affect computations in the entorhinal cortex and hippocampus is apparent in the way that running speed both modulates firing rates and the spectral properties of theta and gamma oscillations in the hippocampus and entorhinal cortex (Buzsaki and Moser 2013). Interestingly, PPTg stimulation has been shown to induce theta oscillations in the hippocampus (Pignatelli, Beyeler et al. 2012).

Previous studies have suggested that locomotion can facilitate a form of coupling in information transmission between the temporoammonic pathway spanning the entorhinal cortex and hippocampus at high gamma frequencies while suppressing the flow information through the tri-synaptic (entorhinal cortex -> dentate gyrus -> CA3 -> CA1) through the hippocampus that is synchronized at lower gamma frequencies(Colgin, Denninger et al. 2009; Colgin and Moser 2010). This differential routing of information has been proposed to divert

resources for the purposes of navigation and encoding memories during periods of locomotion in contrast to periods when animals are quiescent and possibly accessing memories for recalled or updating. It is possible that this rapid modulation may be mediated by cholinergic input from the medium septum, which receives both direct and indirect input from the MLR/PPTg (Pignatelli, Beyeler et al. 2012). In this way, signals regarding speed and velocity have a means of being encoded within the hippocampus in the structure of theta and gamma oscillations. This speed related information may aid the process of path integration to define rough displacements and the generation of place and grid fields in the absence of visuospatial cues (Buzsaki and Moser 2013). In fact, increases in locomotor speed are often correlated with increases the frequency of both theta and gamma oscillations. These correlations may occur because as animal are moving faster, they are encountering adjacent place fields in quicker succession in time. For the neurons to preserve their phase relationships to theta, the period of oscillations may have to be temporally compressed, leading to a higher frequency oscillation. While this is largely speculative, one prediction is that the stimulation of the MLR/PPTg region subthreshold for overt locomotion should be able to enhance the power of theta/gamma oscillations and aid in routing information through the temporoammonic pathway at the expense of the tri-synaptic pathway. These predictions are readily testable using preparations developed in this thesis. While these future studies would investigate these neural pathways in the context of spatial memory, it is likely that these studies would be informative to a general understanding of episodic memory.

Analogies have been made between the ability to mentally time travel to recall and learn-first-person experiences in the context of both space and subjective time to plan future actions. Because of the clear parallels between allocentric navigation and path integration, it has been proposed that support more abstract forms of mental travel may have their evolutionary origin in neural circuits and computations that support locomotor systems that support physical travel through an environment (Buzsaki and Moser 2013). Interestingly, PPTg DBS is known to enhance performance of patients in a series of tasks that test cognitive function and memory, and degeneration of these brainstem cholinergic systems are known to underlie the severe and debilitating effects of Lewy Body Dementia and other neurodegenerative disorders.

Using Reward-based Tasks and Reinforcement Theory to Understand the Function of Cortico-striatal Loops

In general, instrumental behavior requires information about past outcomes to guide future actions. Given the common distinction between the ventral striatum's role in reward and the dorsal striatum's role in movement, we would expect information from the ventral striatum to eventually be transferred to the dorsal striatum. As we described in Chapter 5, stimulation of distinct pathways in the nucleus accumbens mimicked a change in the value of the chosen port where the animal had been stimulated. This was manifested as a change in value of a left or right locomotor response on the following trial. In turn, stimulation of the dorsal striatum mimicked a change in the value of locomotor actions in its effect on the distribution of choices of animals.

It is natural to hypothesize that information regarding the value is being transferred from the ventral striatum during the evaluation of outcomes from a previous trial to affect dorsal striatal activity and the selection of responses on the following trial. In this way, our results have attempted to book-end both the neural substrates and timing of processes for outcome evaluation and action selection in the context of our task. In theory, if we could track the downstream effects of ventral striatal stimulation in either D1- or D2-expressing neurons on neural activity throughout the brain, we could identify the neural pathways that mediate this process through which outcomes become manifest as future reinforced behavior. It would, therefore, be interesting to perform recordings in the dorsal striatum following trials in which animal had been stimulated in the ventral striatum during outcome evaluation. The prediction would be that the value representation for particular actions in the dorsal striatum would consistently change following stimulation ventral striatal stimulation during the outcome phase. These experiments are technically analogous to circuit mapping experiments in Chapter 3 where we demonstrated that MLR activity is bi-directionally regulated by the direct and indirect pathways of the basal ganglia.

While there is a large body of data that anatomically support the idea that there is a series of basal ganglia loops that can facilitate this transfer of information, there is very little physiological data to date to support this hypothesis. So where would the hypothetical pulses of ventral striatal activity travel to in the brain to affect future decisions? Anatomical studies have suggested that there is an extended anatomical pathway that loops and exits into the dorsal striatum. The dorsomedial striatum in turn serves as the entry way to a cognitive loops followed

by entry into a dorsolateral motor loop, which together consist of a system of open spiraling pathways in an extended basal ganglia system (Balleine and O'Doherty ; Alexander and Crutcher 1990). This is in contrast to a model in which dorsal and ventral act in parallel to control performance and learning, respectively (Lobo and Nestler ; O'Doherty, Dayan et al. 2004). These loops may take the form of a serial organization of striato-pallidal-thalamo-cortical pathways. It has also been noted that the projections from the NAc-shell may influence VTA dopaminergic neurons, which in turn project to the NAc core (Luscher and Malenka ; Everitt and Robbins 2005; Niv and Schoenbaum 2008). This spiraling set of connectivity continues from the NAc-core to medial SNc to the DMS to lateral SNc and then to the DLS. The result is a looped medial to lateral organization progressing from limbic to motor basal ganglia subcircuits (Everitt and Robbins 2005). First described in rats by Nauta in 1978(Nauta, Smith et al. 1978), a better defined spiraling stratonigro-striatal organization has been identified in primates (Haber, Fudge et al. 2000). This set of anatomical connections, spanning affective to motor components of the basal ganglia circuit, may provide a substrate for the process by which motivation and incentives give rise to instrumental, and eventually, habitual behaviors through the process of reinforcement learning (Everitt and Robbins 2005).

Methods for Dissecting the Cognitive and Motivational Processes Underlying Locomotor Decision-making

In many areas of perceptual and reward-based decision-making, it has become common for people to utilize a task design in which rodents indicate their choices by locomoting in an environment. This approach underlies the three-port task

design, which is frequently attributed to being developed at Cold Spring Harbor and has since been adopted by many others in both rats and mice. In these set ups, rodents initiate trials with a nose poke into a central port prior to reporting a decision as a movement either to the left or right for water rewards. These tasks can be paired with neural recordings or other methods of measuring and manipulating activity within the brain. The locomotor system that we have defined in rodents has many parallels with the neural circuits that have been traditionally used to study the occulomotor system in primates. In particular, the orienting behaviors, which rodents use to define their direction of locomotion, share a common neural substrate in the superior colliculus as saccadic eye movements system is perhaps the most studied and best understood motor system in terms of its anatomy and physiological response in various operant tasks (Felsen and Mainen 2008; Felsen and Mainen 2012). This allows for a rich set of parallels to be drawn between rodent systems for locomotor decision-making and the older and more extensive field of occulomotor research in primates. Thus, numerous recent studies have been motivated by finding rodent homologues of regions analogous to the frontal eye fields and lateral intraparietal cortex (Erlich, Bialek et al. 2011; Brunton, Botvinick et al. 2013). Collectively, these studies will hopefully allow a firm and rigorous set of mappings between primate and rodent systems for decision-making and reward evaluation. Moreover, these recent advances in behavioral experiments allow for excellent stimulus control, quantitative behavioral measurements, and repeatability in rodent subjects that rivals the degree of control previously only achieved in primate decision-making tasks (O'Connor, Huber et al. 2009). However, the enormous benefit of working with rodents is the ability to make genetically targeted manipulations and record from

large populations of neurons with cell-type specificity using novel optical/ electrophysiological methods. Since their creation, these sophisticated task designs have been utilized to probe many different sensory systems, ranging from the auditory, olfactory, and visual. These studies are revealing the intricate neural processes, underlying a commitment to a perceptual decision (Carandini and Churchland 2013).

More recently, experimenters have moved beyond perceptual tasks in which the sensory evidence is being manipulated and have started to investigate the processes underlying reward-based decision-making by altering either the magnitude or schedule of rewards delivered in their tasks. However, because all of these decisions are essentially being reported as a locomotor response to move either to the left or right, all of the information within these tasks eventually must be routed to the systems that guide goal-directed locomotion. In this way, decision-making in this context represents an information routing problem whereby sensory or reward signals are routed from the periphery through the brain before being mapped onto the value for a right or left locomotor response.

Here, we provide a neural circuit description of how this may occur. The dorsal striatum receives dense input from the entire cortical mantle, limbic system, and thalamus. In this way, the striatum can serve as a “switchboard” in which various sensory modalities and cognitive systems are able to directly input into to gain access to behavioral responses. In this way, striatal inputs from cortex, hippocampus, and amygdala upon the spines of medium spiny neurons may serve as predictive representations of “states of the world” that can serve to

guide behavior. In the context of the Cold Spring Harbor task, these behavioral responses are executed as locomotor decisions to approach a left/right port. We therefore believe that all of this information from these various afferent regions is being converted into a common currency of value, either to support immediate action selection or learning through reinforcement that guides future behavior. Inevitably, for perceptual, cognitive, and motivational signals to be selected for, they must adaptively impact organisms' current or future behavior and interaction. Thus, it seems reasonable that diverse systems should eventually converge upon regions of the brain that guide immediate action selection or learning for guiding future behavior.

Corticostriatal plasticity in this framework can serve as a means of ensuring the proper routing of information regarding "states of the world" to behavioral responses to meet the contingencies set by tasks. Likewise, these 'states of the world" may extend beyond sensory information, but may also encompass abstract rules and stimulus sets encoded by prefrontal regions or other associative areas of the brain that must adaptively form connections with correct behavioral responses mediated by the striatum. In the absence of this plasticity, there is likely to be no preprogrammed map between these states onto behavior. This has been demonstrated by numerous studies, which have shown that there is no inherent biases between cortical, amygdalar, or thalamic inputs onto direct and indirect pathway striatal neurons in naïve animals (MacAskill, Little et al. 2012; Kress, Yamawaki et al. 2013).

The current genetic tools available to dissect neural circuits in rodents utilized in the context of locomotor decision-making are likely going to open new windows into process within the brain. New methods of recording from single as well as populations of neurons will help to identify neural representations of both simple as well as higher order sensory phenomena. Opto/pharmacogenetic manipulations in the context of these tasks can then parse the role of these signals in generating behaviors with a higher degree of temporal and genetic specificity than has previously been available. Together, these efforts can improve our understanding of the neural substrates for guiding various forms of decision-making.

A “Common Currency” for Decision-Making

In chapter 4 and 5, we identified that stimulation of dorsal and ventral striatum could influence immediate action selection or reinforcement learning. Perhaps, the most intriguing aspect of our results is that stimulation seemed to mimic a change in the value of actions either during the current trial with dorsal striatum stimulation or the following trial with ventral striatal stimulation. In this way, a fixed amount of neural activity that we introduce into the brain can be equated to a well-defined change in the probability of the animal’s choice behavior across various values for left and right choice, whether it is in system for action selection or outcome evaluation. Here, we described this shift as a change in the action value because the rewards are capable of directing the animals’ probability of choices.

Our data bear a remarkable resemblance to data to from Michael Shadlen and Bill Newsome's lab in the context of perceptual tasks (Gold and Shadlen 2007). In these experiments, primate subjects were trained in a motion discrimination task in which animals report their decisions as a saccadic eye movement. During this experiment, electrical stimulation was delivered to area MT (involved in the perception of motion stimuli) while recording from LIP, which plays a vital role in motor planning for saccadic eye movements (Figure 33). They found that stimulation of MT dramatically changed the activity of decision-related activity in LIP, and this activity was predictive of an animals' behavioral choice. In these tasks, the psychometrics curve represents the way in which various sensory streams of input are converted to a choice behavioral response among two alternatives. (Salzman, Britten et al. 1990). This result was remarkable because it implied that neuronal activity resulting from microstimulation could be directly related to the probability of a choice response in a consistent manner across all task relevant stimuli conditions. In this way, sensory information in the context of a decision-making task was being routed to action selection systems for defining the likelihood of a choice before being implemented as a motor response. Similar results were also identified in a recent study using a rodent auditory perceptual task where corticostriatal inputs were optogenetically stimulated (Znamenskiy and Zador).

So what exactly do these shifts with stimulation in these sigmoid curves mean? Clearly, this result generalizes across diverse regions of the brain serving different functions and in completely unrelated types of decision-making tasks? There must be some connection between all of these results. Mathematically,

these sigmoidal curves can be represented by a logistic equation curve in which the task relevant parameter is linearly related to the log probability of a choice. In the case of psychometrics curves, the stimulus intensity is being linearly mapped onto a log probability of a choice. In our reward-based tasks, the definition of our action value in our tasks is essentially the extent to which we predict the log probability or log likelihood of choice of the animals' behavior based upon their previous history of rewards and choice. Therefore, the units of action values are log likelihood of choice, and a shift in these curves is equivalent to an additive change by a fixed log likelihood of choice (Gold and Shadlen 2007). In this way, diverse neural representations can be linearly transformed to be represented along a single axis of log probability of choice. Together, these data seem to imply an astounding principle, namely that activity introduced by either electrical/optical stimulation in diverse, but task-relevant areas of the brain essentially can be converted into the common units of log probability of a choice. Thus, all task-relevant activity in the brain can be converted into a common currency defined by its relationship to behavior.

So why should information be related to the log probability of a behavioral choice? What is the significance of this quantity? To fully appreciate this, one needs to turn to the formal framework of Bayesian statistics. In the Bayesian interpretation, probabilities and odds represent measures of a degree of belief given certain forms of evidence. In turn, these measures of belief prior to a piece of evidence are be updated after accounting for an additional piece of evidence forming a posterior belief. Imagine a situation in which you have a categorical choice between two option, $\mathbf{A_1}$ and $\mathbf{A_2}$. Prior to receiving a novel piece of

evidence **B**, our measure of belief for $\mathbf{A_1}$ verus $\mathbf{A_2}$ can be mathematically formalized to be the prior odds of event $\mathbf{A_1}$ being true versus $\mathbf{A_2}$, **O** ($\mathbf{A_2 : A_2}$). After knowing a piece of evidence B, we must update our beliefs to come to a new estimate, which is represented by our posterior odds **O** ($\mathbf{A_2 : A_2| B}$) to take into account our new knowledge. The extent to which are prior understanding needs to be changed given this additional piece of evidence can be described by a likelihood ratio **Λ(A1:A2|B)**. This process of belief updating is made formal by the odds formulation of Bayes' Rule:

$$O(A_1 : A_2|B) = O(A_1 : A_2) \cdot \Lambda(A_1 : A_2|B)$$

However, Bayes' Rule is often easier to understand in another format, in which we take the logarithm of both sides to convert the operation of multiplication between the prior and the likelihood ratio into addition.

$$log\ O(A_1 : A_2|B) = \ log\ O(A_1 : A_2) + \ log\ \Lambda(A_1 : A_2|B)$$

In this way, all evidence bearing upon a proposition is now additive accumulated in units of log odds of the proposition. If we think of a commitment to a proposition $\mathbf{A_1}$ over $\mathbf{A_2}$ as being equivalent to the commitment to a motor response through the process of action selection, then all of our evidence will now be accumulated in units of log probability of a choice response. It so happens that there exists a generalization to Bayes' Law that can take into account the values/utilities of choices, which forms the basis of statistical decision-making. In this framework, values/utilities also can be converted into units of log likelihoods for the commitment to a choice. Therefore, if all decision-making processes can be thought of as accumulating evidence, a measure of belief, for the purpose of committing to a categorical proposition, it is not

surprising that all of the components relevant to a decision are converted in a common Bayesian currency of log probability of choice.

These data imply a remarkable consilience that has been achieved if this general theory of decision-making is true. Not only would this be a remarkable convergence between theory and experimental data, but this would also mark a union between two disciplines in neuroscience that have traditionally been treated separately, the field of cognitive science with its emphasis on the neural basis of perception and the study of motivation. Moreover, the remarkable resemblance of different data sets across species bolsters support for the idea that studying decision-making in rodents can tell us something about how these processes may operate in primates and human subjects. The similarity in computation also serves as a testament to the conserved and possibly essential nature of decision-making processes in the brain.

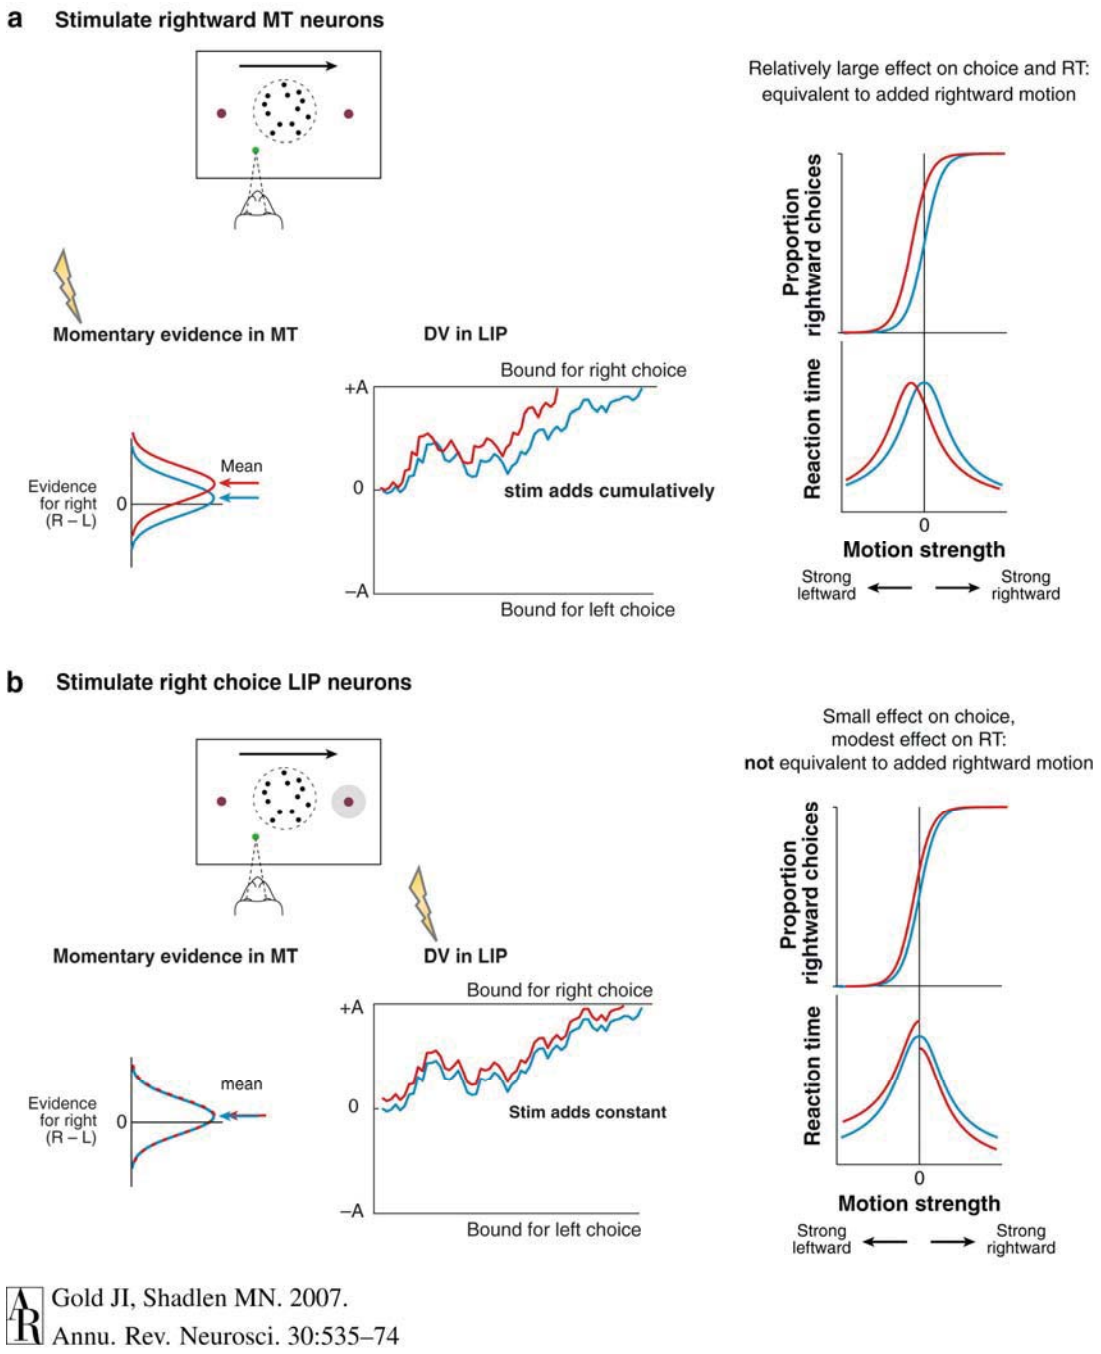

Gold JI, Shadlen MN. 2007.
Annu. Rev. Neurosci. 30:535–74

Figure 33. Effects of microstimulation in MT and LIP.
In both areas microstimulation (*red curves*) causes a change in both choice and RT. The schematic shows the consequences of adding a small change in spike rate to the evidence or to the DV. The graphs on the right are theoretical results obtained using the bounded diffusion model. They resemble the pattern of data in Hanks et al. (2006). (*a*) MT microstimulation mimics a change in stimulus strength (evidence). (*b*) LIP microstimulation mimics an additive offset to the DV (or, equivalently, the height of the bounds). *Adapted from Gold JI and Shadlen MN. Annual Reviews. 2007.*

References

Albin, R. L., A. B. Young, et al. (1989). "The Functional-Anatomy of Basal Ganglia Disorders." Trends in Neurosciences **12**(10): 366-375.

Albin, R. L., A. B. Young, et al. (1989). "The functional anatomy of basal ganglia disorders." Trends Neurosci **12**(10): 366-375.

Alessandro, S., R. Ceravolo, et al. (2010). "Non-motor functions in parkinsonian patients implanted in the pedunculopontine nucleus: focus on sleep and cognitive domains." J Neurol Sci **289**(1-2): 44-48.

Alexander, G. E. and M. D. Crutcher (1990). "Functional architecture of basal ganglia circuits: neural substrates of parallel processing." Trends Neurosci **13**(7): 266-271.

Alitto, H. J. and Y. Dan "Cell-type-specific modulation of neocortical activity by basal forebrain input." Front Syst Neurosci **6**: 79.

Alitto, H. J. and Y. Dan (2012). "Cell-type-specific modulation of neocortical activity by basal forebrain input." Front Syst Neurosci **6**: 79.

Ambroggi, F., A. Ghazizadeh, et al. (2011). "Roles of nucleus accumbens core and shell in incentive-cue responding and behavioral inhibition." J Neurosci **31**(18): 6820-6830.

Ambroggi, F., A. Ishikawa, et al. (2008). "Basolateral amygdala neurons facilitate reward-seeking behavior by exciting nucleus accumbens neurons." Neuron **59**(4): 648-661.

Andermann, M. L., A. M. Kerlin, et al. (2011). "Functional specialization of mouse higher visual cortical areas." Neuron **72**(6): 1025-1039.

Ariano, M. A., C. Cepeda, et al. (2005). "Striatal potassium channel dysfunction in Huntington's disease transgenic mice." J Neurophysiol **93**(5): 2565-2574.

Armstrong, K. M., J. K. Fitzgerald, et al. (2006). "Changes in visual receptive fields with microstimulation of frontal cortex." Neuron **50**(5): 791-798.

Arnulf, I., M. Ferraye, et al. "Sleep induced by stimulation in the human pedunculopontine nucleus area." Ann Neurol **67**(4): 546-549.

Arnulf, I., M. Ferraye, et al. (2010). "Sleep induced by stimulation in the human pedunculopontine nucleus area." Ann Neurol **67**(4): 546-549.

Ayaz, A., A. B. Saleem, et al. (2013). "Locomotion controls spatial integration in mouse visual cortex." Curr Biol **23**(10): 890-894.

Balleine, B. W., M. R. Delgado, et al. (2007). "The role of the dorsal striatum in reward and decision-making." J Neurosci **27**(31): 8161-8165.

Balleine, B. W. and J. P. O'Doherty "Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action." Neuropsychopharmacology **35**(1): 48-69.

Barnes, T. D., Y. Kubota, et al. (2005). "Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories." Nature **437**(7062): 1158-1161.

Barraclough, D. J., M. L. Conroy, et al. (2004). "Prefrontal cortex and decision making in a mixed-strategy game." Nat Neurosci **7**(4): 404-410.

Bayer, H. M. and P. W. Glimcher (2005). "Midbrain dopamine neurons encode a quantitative reward prediction error signal." Neuron **47**(1): 129-141.

Benke, T. (2006). "Peduncular hallucinosis: a syndrome of impaired reality monitoring." J Neurol **253**(12): 1561-1571.

Berg, R. W., B. Friedman, et al. (2005). "Activation of nucleus basalis facilitates cortical control of a brain stem motor program." J Neurophysiol **94**(1): 699-711.

Berridge, K. C. (2007). "The debate over dopamine's role in reward: the case for incentive salience." Psychopharmacology (Berl) **191**(3): 391-431.
Bisley, J. W. (2011). "The neural basis of visual attention." J Physiol **589**(Pt 1): 49-57.
Bjursten, L. M., K. Norrsell, et al. (1976). "Behavioural repertory of cats without cerebral cortex from infancy." Exp Brain Res **25**(2): 115-130.
Bock, R., J. H. Shin, et al. (2013). "Strengthening the accumbal indirect pathway promotes resilience to compulsive cocaine use." Nat Neurosci **16**(5): 632-638.
Boeve, B. F. (2013). "Idiopathic REM sleep behaviour disorder in the development of Parkinson's disease." Lancet Neurol **12**(5): 469-482.
Brainard, M. S. and A. J. Doupe (2000). "Interruption of a basal ganglia-forebrain circuit prevents plasticity of learned vocalizations." Nature **404**(6779): 762-766.
Britt, J. P., F. Benaliouad, et al. "Synaptic and behavioral profile of multiple glutamatergic inputs to the nucleus accumbens." Neuron **76**(4): 790-803.
Brunton, B. W., M. M. Botvinick, et al. (2013). "Rats and humans can optimally accumulate evidence for decision-making." Science **340**(6128): 95-98.
Buzsaki, G., R. G. Bickford, et al. (1988). "Nucleus basalis and thalamic control of neocortical activity in the freely moving rat." J Neurosci **8**(11): 4007-4026.
Buzsaki, G. and E. I. Moser (2013). "Memory, navigation and theta rhythm in the hippocampal-entorhinal system." Nat Neurosci **16**(2): 130-138.
Carandini, M. and A. K. Churchland (2013). "Probing perceptual decisions in rodents." Nat Neurosci **16**(7): 824-831.
Cavanaugh, J. and R. H. Wurtz (2004). "Subcortical modulation of attention counters change blindness." J Neurosci **24**(50): 11236-11243.
Colgin, L. L., T. Denninger, et al. (2009). "Frequency of gamma oscillations routes flow of information in the hippocampus." Nature **462**(7271): 353-357.
Colgin, L. L. and E. I. Moser (2010). "Gamma oscillations in the hippocampus." Physiology (Bethesda) **25**(5): 319-329.
Cromwell, H. C. and W. Schultz (2003). "Effects of expectations for different reward magnitudes on neuronal activity in primate striatum." J Neurophysiol **89**(5): 2823-2838.
Cui, G., S. B. Jun, et al. (2013). "Concurrent activation of striatal direct and indirect pathways during action initiation." Nature **494**(7436): 238-242.
Curro Dossi, R., D. Pare, et al. (1991). "Short-lasting nicotinic and long-lasting muscarinic depolarizing responses of thalamocortical neurons to stimulation of mesopontine cholinergic nuclei." J Neurophysiol **65**(3): 393-406.
Cutrell, E. B. and R. T. Marrocco (2002). "Electrical microstimulation of primate posterior parietal cortex initiates orienting and alerting components of covert attention." Exp Brain Res **144**(1): 103-113.
Dalley, J. W., B. J. Everitt, et al. "Impulsivity, compulsivity, and top-down cognitive control." Neuron **69**(4): 680-694.
Dean, P., P. Redgrave, et al. (1989). "Event or emergency? Two response systems in the mammalian superior colliculus." Trends Neurosci **12**(4): 137-147.
Deliagina, T. G. and G. N. Orlovsky (2002). "Comparative neurobiology of postural control." Curr Opin Neurobiol **12**(6): 652-657.
Deliagina, T. G., G. N. Orlovsky, et al. (2006). "Neural bases of postural control." Physiology (Bethesda) **21**: 216-225.
DeLong, M. R. (1990). "Primate models of movement disorders of basal ganglia origin." Trends Neurosci **13**(7): 281-285.

Desmurget, M. and R. S. Turner (2010). "Motor sequences and the basal ganglia: kinematics, not habits." J Neurosci **30**(22): 7685-7690.

Dombeck, D. A., A. N. Khabbaz, et al. (2007). "Imaging large-scale neural activity with cellular resolution in awake, mobile mice." Neuron **56**(1): 43-57.

Drenan, R. M., S. R. Grady, et al. (2008). "In vivo activation of midbrain dopamine neurons via sensitized, high-affinity alpha 6 nicotinic acetylcholine receptors." Neuron **60**(1): 123-136.

Dringenberg, H. C. and M. C. Olmstead (2003). "Integrated contributions of basal forebrain and thalamus to neocortical activation elicited by pedunculopontine tegmental stimulation in urethane-anesthetized rats." Neuroscience **119**(3): 839-853.

Erisir, A., S. C. Van Horn, et al. (1997). "Relative numbers of cortical and brainstem inputs to the lateral geniculate nucleus." Proc Natl Acad Sci U S A **94**(4): 1517-1520.

Erlich, J. C., M. Bialek, et al. (2011). "A cortical substrate for memory-guided orienting in the rat." Neuron **72**(2): 330-343.

Everitt, B. J. and T. W. Robbins (2005). "Neural systems of reinforcement for drug addiction: from actions to habits to compulsion." Nat Neurosci **8**(11): 1481-1489.

Fecteau, J. H. and D. P. Munoz (2003). "Exploring the consequences of the previous trial." Nat Rev Neurosci **4**(6): 435-443.

Felsen, G. and Z. F. Mainen (2008). "Neural substrates of sensory-guided locomotor decisions in the rat superior colliculus." Neuron **60**(1): 137-148.

Felsen, G. and Z. F. Mainen (2012). "Midbrain contributions to sensorimotor decision making." J Neurophysiol **108**(1): 135-147.

Follett, K. A. and D. Torres-Russotto (2012). "Deep brain stimulation of globus pallidus interna, subthalamic nucleus, and pedunculopontine nucleus for Parkinson's disease: which target?" Parkinsonism Relat Disord **18 Suppl 1**: S165-167.

Fox, K. and M. Armstrong-James (1986). "The role of the anterior intralaminar nuclei and N-methyl D-aspartate receptors in the generation of spontaneous bursts in rat neocortical neurones." Exp Brain Res **63**(3): 505-518.

French, J. C., F. K. Von Amerongen, et al. (1952). "An activating system in brain stem of monkey." AMA Arch Neurol Psychiatry **68**(5): 577-590.

Gerfen, C. R. (1992). "The neostriatal mosaic: multiple levels of compartmental organization." Trends Neurosci **15**(4): 133-139.

Gerfen, C. R. (1992). "The neostriatal mosaic: multiple levels of compartmental organization." J Neural Transm Suppl **36**: 43-59.

Gerfen, C. R. and D. J. Surmeier "Modulation of striatal projection systems by dopamine." Annu Rev Neurosci **34**: 441-466.

Gerfen, C. R. and D. J. Surmeier (2011). "Modulation of striatal projection systems by dopamine." Annu Rev Neurosci **34**: 441-466.

Giovannini, M. G., L. Bartolini, et al. (1998). "Acetylcholine release from the frontal cortex during exploratory activity." Brain Res **784**(1-2): 218-227.

Goard, M. and Y. Dan (2009). "Basal forebrain activation enhances cortical coding of natural scenes." Nat Neurosci **12**(11): 1444-1449.

Gold, J. I. and M. N. Shadlen (2007). "The neural basis of decision making." Annual Review of Neuroscience **30**: 535-574.

Gold, J. I. and M. N. Shadlen (2007). "The neural basis of decision making." Annu Rev Neurosci **30**: 535-574.

Good, C. H. and C. R. Lupica (2010). "Afferent-specific AMPA receptor subunit composition and regulation of synaptic plasticity in midbrain dopamine neurons by abused drugs." J Neurosci **30**(23): 7900-7909.

Granata, A. R. and S. T. Kitai (1991). "Inhibitory substantia nigra inputs to the pedunculopontine neurons." Exp Brain Res **86**(3): 459-466.

Graybiel, A. M. (2008). "Habits, rituals, and the evaluative brain." Annu Rev Neurosci **31**: 359-387.

Grillner, S. (2003). "The motor infrastructure: from ion channels to neuronal networks." Nat Rev Neurosci **4**(7): 573-586.

Grillner, S., J. Hellgren, et al. (2005). "Mechanisms for selection of basic motor programs--roles for the striatum and pallidum." Trends Neurosci **28**(7): 364-370.

Grillner, S., P. Wallen, et al. (2008). "Neural bases of goal-directed locomotion in vertebrates--an overview." Brain Res Rev **57**(1): 2-12.

Grueter, B. A., G. Brasnjo, et al. "Postsynaptic TRPV1 triggers cell type-specific long-term depression in the nucleus accumbens." Nat Neurosci **13**(12): 1519-1525.

Haber, S. N., J. L. Fudge, et al. (2000). "Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum." J Neurosci **20**(6): 2369-2382.

Hallanger, A. E. and B. H. Wainer (1988). "Ascending projections from the pedunculopontine tegmental nucleus and the adjacent mesopontine tegmentum in the rat." J Comp Neurol **274**(4): 483-515.

Hamani, C., E. Moro, et al. "The pedunculopontine nucleus as a target for deep brain stimulation." J Neural Transm **118**(10): 1461-1468.

Hamani, C., E. Moro, et al. (2011). "The pedunculopontine nucleus as a target for deep brain stimulation." J Neural Transm **118**(10): 1461-1468.

Harris, K. D. and A. Thiele (2011). "Cortical state and attention." Nat Rev Neurosci **12**(9): 509-523.

Harvey, C. D., F. Collman, et al. (2009). "Intracellular dynamics of hippocampal place cells during virtual navigation." Nature **461**(7266): 941-946.

Hasselmo, M. E. (2006). "The role of acetylcholine in learning and memory." Curr Opin Neurobiol **16**(6): 710-715.

Hasselmo, M. E. and L. M. Giocomo (2006). "Cholinergic modulation of cortical function." J Mol Neurosci **30**(1-2): 133-135.

Heimer L, W. R. (1975). "The subcortical projections of the allocortex: similarities in the neural associations of the hippocampus, the piriform cortex, and the neocortex. I." Perspectives in Neurobiology. **Golgi Centennial Symposium, New York, NY**: 177-193.

Hikosaka, O. and M. Isoda (2010). "Switching from automatic to controlled behavior: cortico-basal ganglia mechanisms." Trends Cogn Sci **14**(4): 154-161.

Hikosaka, O., K. Nakamura, et al. (2006). "Basal ganglia orient eyes to reward." J Neurophysiol **95**(2): 567-584.

Hikosaka, O., Y. Takikawa, et al. (2000). "Role of the basal ganglia in the control of purposive saccadic eye movements." Physiol Rev **80**(3): 953-978.

Histed, M. H., V. Bonin, et al. (2009). "Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation." Neuron **63**(4): 508-522.

Hong, S. and O. Hikosaka "Dopamine-mediated learning and switching in cortico-striatal circuit explain behavioral changes in reinforcement learning." Front Behav Neurosci **5**: 15.

Houk, J. C., J. L. Adams, et al. (1995). "A model of how the basal ganglia generate an use neural signals that predict reinforcement." 249-270.
Huang, Z. L., W. M. Qu, et al. (2005). "Adenosine A2A, but not A1, receptors mediate the arousal effect of caffeine." Nat Neurosci **8**(7): 858-859.
Ikeda, T. and O. Hikosaka (2003). "Reward-dependent gain and bias of visual responses in primate superior colliculus." Neuron **39**(4): 693-700.
Ishikawa, A., F. Ambroggi, et al. (2008). "Dorsomedial prefrontal cortex contribution to behavioral and nucleus accumbens neuronal responses to incentive cues." J Neurosci **28**(19): 5088-5098.
Isoda, M. and O. Hikosaka (2011). "Cortico-basal ganglia mechanisms for overcoming innate, habitual and motivational behaviors." Eur J Neurosci **33**(11): 2058-2069.
Ito, M. and K. Doya (2009). "Validation of decision-making models and analysis of decision variables in the rat basal ganglia." J Neurosci **29**(31): 9861-9874.
Jiang, H., B. E. Stein, et al. (2003). "Opposing basal ganglia processes shape midbrain visuomotor activity bilaterally." Nature **423**(6943): 982-986.
Jin, X. and R. M. Costa (2010). "Start/stop signals emerge in nigrostriatal circuits during sequence learning." Nature **466**(7305): 457-462.
Joel, D., Y. Niv, et al. (2002). "Actor-critic models of the basal ganglia: new anatomical and computational perspectives." Neural Netw **15**(4-6): 535-547.
Kable, J. W. and P. W. Glimcher (2009). "The neurobiology of decision: consensus and controversy." Neuron **63**(6): 733-745.
Kang, Y. and S. T. Kitai (1990). "Electrophysiological properties of pedunculopontine neurons and their postsynaptic responses following stimulation of substantia nigra reticulata." Brain Res **535**(1): 79-95.
Kawaguchi, Y., C. J. Wilson, et al. (1995). "Striatal interneurones: chemical, physiological and morphological characterization." Trends Neurosci **18**(12): 527-535.
Keller, G. B., T. Bonhoeffer, et al. (2012). "Sensorimotor mismatch signals in primary visual cortex of the behaving mouse." Neuron **74**(5): 809-815.
Kepecs, A. (2013). "The uncertainty of it all." Nat Neurosci **16**(6): 660-662.
Khamassi, M., A. B. Mulder, et al. (2008). "Anticipatory reward signals in ventral striatal neurons of behaving rats." Eur J Neurosci **28**(9): 1849-1866.
Kim, H., J. H. Sul, et al. (2009). "Role of striatum in updating values of chosen actions." J Neurosci **29**(47): 14701-14712.
Kim, H., J. H. Sul, et al. (2009). "Role of Striatum in Updating Values of Chosen Actions." Journal of Neuroscience **29**(47): 14701-14712.
Kim, Y. B., N. Huh, et al. (2007). "Encoding of action history in the rat ventral striatum." J Neurophysiol **98**(6): 3548-3556.
Kimchi, E. Y. and M. Laubach (2009). "The dorsomedial striatum reflects response bias during learning." J Neurosci **29**(47): 14891-14902.
Kimchi, E. Y. and M. Laubach (2009). "Dynamic encoding of action selection by the medial striatum." J Neurosci **29**(10): 3148-3159.
Kinomura, S., J. Larsson, et al. (1996). "Activation by attention of the human reticular formation and thalamic intralaminar nuclei." Science **271**(5248): 512-515.
Knutson, B., J. Taylor, et al. (2005). "Distributed neural representation of expected value." J Neurosci **25**(19): 4806-4812.
Kravitz, A. V., B. S. Freeze, et al. (2010). "Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry." Nature **466**(7306): 622-626.

Kravitz, A. V. and A. C. Kreitzer (2012). "Striatal mechanisms underlying movement, reinforcement, and punishment." Physiology (Bethesda) **27**(3): 167-177.

Kravitz, A. V., L. D. Tye, et al. (2012). "Distinct roles for direct and indirect pathway striatal neurons in reinforcement." Nat Neurosci **15**(6): 816-818.

Kravitz, A. V., L. D. Tye, et al. (2012). "Distinct roles for direct and indirect pathway striatal neurons in reinforcement." Nat Neurosci.

Kreitzer, A. C. and J. D. Berke (2011). "Investigating striatal function through cell-type-specific manipulations." Neuroscience.

Kreitzer, A. C. and R. C. Malenka (2005). "Dopamine modulation of state-dependent endocannabinoid release and long-term depression in the striatum." J Neurosci **25**(45): 10537-10545.

Kreitzer, A. C. and R. C. Malenka (2007). "Endocannabinoid-mediated rescue of striatal LTD and motor deficits in Parkinson's disease models." Nature **445**(7128): 643-647.

Kreitzer, A. C. and R. C. Malenka (2008). "Striatal plasticity and basal ganglia circuit function." Neuron **60**(4): 543-554.

Kress, G. J., N. Yamawaki, et al. (2013). "Convergent cortical innervation of striatal projection neurons." Nat Neurosci **16**(6): 665-667.

Kubota, Y., J. Liu, et al. (2009). "Stable encoding of task structure coexists with flexible coding of task events in sensorimotor striatum." J Neurophysiol **102**(4): 2142-2160.

Kuo, S. H., C. Kenney, et al. (2008). "Bilateral pedunculopontine nuclei strokes presenting as freezing of gait." Mov Disord **23**(4): 616-619.

Kurosawa, M., K. Okada, et al. (1993). "Extracellular release of acetylcholine, noradrenaline and serotonin increases in the cerebral cortex during walking in conscious rats." Neurosci Lett **161**(1): 73-76.

Lau, B. and P. W. Glimcher (2005). "Dynamic response-by-response models of matching behavior in rhesus monkeys." J Exp Anal Behav **84**(3): 555-579.

Lau, B. and P. W. Glimcher (2008). "Value representations in the primate striatum during matching behavior." Neuron **58**(3): 451-463.

Lauwereyns, J., K. Watanabe, et al. (2002). "A neural correlate of response bias in monkey caudate nucleus." Nature **418**(6896): 413-417.

Lawrence, D. G. and H. G. Kuypers (1968). "The functional organization of the motor system in the monkey. I. The effects of bilateral pyramidal lesions." Brain **91**(1): 1-14.

Lawrence, D. G. and H. G. Kuypers (1968). "The functional organization of the motor system in the monkey. II. The effects of lesions of the descending brain-stem pathways." Brain **91**(1): 15-36.

Lazarus, M., J. F. Chen, et al. (2013). "Role of the basal ganglia in the control of sleep and wakefulness." Curr Opin Neurobiol.

Lee, D., H. Seo, et al. (2012). "Neural basis of reinforcement learning and decision making." Annu Rev Neurosci **35**: 287-308.

Lindsley, D. B., L. H. Schreiner, et al. (1950). "Behavioral and EEG changes following chronic brain stem lesions in the cat." Electroencephalogr Clin Neurophysiol **2**(4): 483-498.

Lo, C. C. and X. J. Wang (2006). "Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks." Nat Neurosci **9**(7): 956-963.

Lobo, M. K., H. E. Covington, 3rd, et al. "Cell type-specific loss of BDNF signaling mimics optogenetic control of cocaine reward." Science **330**(6002): 385-390.

Lobo, M. K., H. E. Covington, 3rd, et al. (2010). "Cell type-specific loss of BDNF signaling mimics optogenetic control of cocaine reward." Science **330**(6002): 385-390.

Lobo, M. K. and E. J. Nestler "The striatal balancing act in drug addiction: distinct roles of direct and indirect pathway medium spiny neurons." Front Neuroanat **5**: 41.

Lobo, M. K. and E. J. Nestler (2011). "The striatal balancing act in drug addiction: distinct roles of direct and indirect pathway medium spiny neurons." Front Neuroanat **5**: 41.

Lu, S. M., W. Guido, et al. (1993). "The brain-stem parabrachial region controls mode of response to visual stimulation of neurons in the cat's lateral geniculate nucleus." Vis Neurosci **10**(4): 631-642.

Luscher, C. and R. C. Malenka "Drug-evoked synaptic plasticity in addiction: from molecular changes to circuit remodeling." Neuron **69**(4): 650-663.

Luscher, C. and V. Pascoli "[Optogenetic reversal of cocaine-evoked synaptic potentiation normalizes behavioural sensitization]." Med Sci (Paris) **28**(4): 353-355.

MacAskill, A. F., J. P. Little, et al. (2012). "Subcellular connectivity underlies pathway-specific signaling in the nucleus accumbens." Nat Neurosci **15**(12): 1624-1626.

Martinez-Gonzalez, C., J. P. Bolam, et al. (2011). "Topographical organization of the pedunculopontine nucleus." Front Neuroanat **5**: 22.

Maurice, N., J. M. Deniau, et al. (1999). "Relationships between the prefrontal cortex and the basal ganglia in the rat: physiology of the cortico-nigral circuits." J Neurosci **19**(11): 4674-4681.

Maurice, N., J. M. Deniau, et al. (1997). "Position of the ventral pallidum in the rat prefrontal cortex-basal ganglia circuit." Neuroscience **80**(2): 523-534.

Mazzone, P., A. Lozano, et al. (2005). "Implantation of human pedunculopontine nucleus: a safe and clinically relevant target in Parkinson's disease." Neuroreport **16**(17): 1877-1881.

Mena-Segovia, J., J. P. Bolam, et al. (2004). "Pedunculopontine nucleus and basal ganglia: distant relatives or part of the same family?" Trends Neurosci **27**(10): 585-588.

Mink, J. W. (1996). "The basal ganglia: focused selection and inhibition of competing motor programs." Prog Neurobiol **50**(4): 381-425.

Mirenowicz, J. and W. Schultz (1994). "Importance of unpredictability for reward responses in primate dopamine neurons." J Neurophysiol **72**(2): 1024-1027.

Mogenson, G. J., D. L. Jones, et al. (1980). "From motivation to action: functional interface between the limbic system and the motor system." Prog Neurobiol **14**(2-3): 69-97.

Montague, P. R., S. E. Hyman, et al. (2004). "Computational roles for dopamine in behavioural control." Nature **431**(7010): 760-767.

Moore, T. and K. M. Armstrong (2003). "Selective gating of visual signals by microstimulation of frontal cortex." Nature **421**(6921): 370-373.

Moore, T. and M. Fallah (2001). "Control of eye movements and spatial attention." Proc Natl Acad Sci U S A **98**(3): 1273-1276.

Mori, S., H. Nishimura, et al. (1978). "Controlled locomotion in the mesencephalic cat: distribution of facilitatory and inhibitory regions within pontine tegmentum." J Neurophysiol **41**(6): 1580-1591.

Moro, E., C. Hamani, et al. (2010). "Unilateral pedunculopontine stimulation improves falls in Parkinson's disease." Brain **133**(Pt 1): 215-224.

Moruzzi, G. and H. W. Magoun (1949). "Brain stem reticular formation and activation of the EEG." Electroencephalogr Clin Neurophysiol **1**(4): 455-473.

Muller, J. R., M. G. Philiastides, et al. (2005). "Microstimulation of the superior colliculus focuses attention without moving the eyes." Proc Natl Acad Sci U S A **102**(3): 524-529.

Munk, M. H., P. R. Roelfsema, et al. (1996). "Role of reticular activation in the modulation of intracortical synchronization." Science **272**(5259): 271-274.

Munro-Davies, L. E., J. Winter, et al. (1999). "The role of the pedunculopontine region in basal-ganglia mechanisms of akinesia." Exp Brain Res **129**(4): 511-517.

Nakahara, H., H. Itoh, et al. (2004). "Dopamine neurons can represent context-dependent prediction error." Neuron **41**(2): 269-280.

Nandi, D., T. Z. Aziz, et al. (2002). "Reversal of akinesia in experimental parkinsonism by GABA antagonist microinjections in the pedunculopontine nucleus." Brain **125**(Pt 11): 2418-2430.

Nauta W.J.H. , K. H. G. J. M. (1958). Reticular formation of the brain. Boston, Little, Brown, and Co,.

Nauta, W. J., G. P. Smith, et al. (1978). "Efferent connections and nigral afferents of the nucleus accumbens septi in the rat." Neuroscience **3**(4-5): 385-401.

Nicola, S. M. (2007). "The nucleus accumbens as part of a basal ganglia action selection circuit." Psychopharmacology (Berl) **191**(3): 521-550.

Niell, C. M. and M. P. Stryker (2010). "Modulation of visual responses by behavioral state in mouse visual cortex." Neuron **65**(4): 472-479.

Niv, Y. and G. Schoenbaum (2008). "Dialogues on prediction errors." Trends Cogn Sci **12**(7): 265-272.

Norton, A. B., Y. S. Jo, et al. (2011). "Independent neural coding of reward and movement by pedunculopontine tegmental nucleus neurons in freely navigating rats." Eur J Neurosci **33**(10): 1885-1896.

O'Connor, D. H., D. Huber, et al. (2009). "Reverse engineering the mouse brain." Nature **461**(7266): 923-929.

O'Doherty, J., P. Dayan, et al. (2004). "Dissociable roles of ventral and dorsal striatum in instrumental conditioning." Science **304**(5669): 452-454.

Olds, J. (1963). "Self-stimulation experiments." Science **140**(3563): 218-220.

Orlovsky, G. N., Deliagina T., Grillner S. (1999). "Neuronal Control of Locomotion. From Mollusc to Man."

Padoa-Schioppa, C. and J. A. Assad (2006). "Neurons in the orbitofrontal cortex encode economic value." Nature **441**(7090): 223-226.

Pascoli, V., M. Turiault, et al. "Reversal of cocaine-evoked synaptic potentiation resets drug-induced adaptive behaviour." Nature **481**(7379): 71-75.

Pasupathy, A. and E. K. Miller (2005). "Different time courses of learning-related activity in the prefrontal cortex and striatum." Nature **433**(7028): 873-876.

Patel, J. C., E. Rossignol, et al. (2012). "Opposing regulation of dopaminergic activity and exploratory motor behavior by forebrain and brainstem cholinergic circuits." Nat Commun **3**: 1172.

Pennartz, C. M., J. D. Berke, et al. (2009). "Corticostriatal Interactions during Learning, Memory Processing, and Decision Making." J Neurosci **29**(41): 12831-12838.

Pennartz, C. M. A., J. D. Berke, et al. (2009). "Corticostriatal Interactions during Learning, Memory Processing, and Decision Making." The Journal of Neuroscience **29**(41): 12831-12838.

Pignatelli, M., A. Beyeler, et al. (2012). "Neural circuits underlying the generation of theta oscillations." J Physiol Paris **106**(3-4): 81-92.

Plaha, P. and S. S. Gill (2005). "Bilateral deep brain stimulation of the pedunculopontine nucleus for Parkinson's disease." Neuroreport **16**(17): 1883-1887.

Qiu, M. H., W. Liu, et al. (2012). "The role of nucleus accumbens core/shell in sleep-wake regulation and their involvement in modafinil-induced arousal." PLoS One **7**(9): e45471.

Qiu, M. H., R. Vetrivelan, et al. (2010). "Basal ganglia control of sleep-wake behavior and cortical activation." Eur J Neurosci **31**(3): 499-507.

Qu, W. M., Z. L. Huang, et al. (2008). "Dopaminergic D1 and D2 receptors are essential for the arousal effect of modafinil." J Neurosci **28**(34): 8462-8469.

Ragozzino, A., F. Lassandro, et al. (2007). "Value of MRI in three patients with major vascular injuries after laparoscopic cholecystectomy." Emerg Radiol **14**(6): 443-447.

Ragozzino, M. E. (2007). "The contribution of the medial prefrontal cortex, orbitofrontal cortex, and dorsomedial striatum to behavioral flexibility." Ann N Y Acad Sci **1121**: 355-375.

Rangel, A., C. Camerer, et al. (2008). "A framework for studying the neurobiology of value-based decision making." Nature Reviews Neuroscience **9**(7): 545-556.

Redgrave, P., T. J. Prescott, et al. (1999). "The basal ganglia: a vertebrate solution to the selection problem?" Neuroscience **89**(4): 1009-1023.

Redgrave, P., T. J. Prescott, et al. (1999). "Is the short-latency dopamine response too short to signal reward error?" Trends Neurosci **22**(4): 146-151.

Redgrave, P., G. W. Westby, et al. (1993). "Functional architecture of rodent superior colliculus: relevance of multiple output channels." Prog Brain Res **95**: 69-77.

Rodriguez, R., U. Kallenbach, et al. (2004). "Short- and long-term effects of cholinergic modulation on gamma oscillations and response synchronization in the visual cortex." J Neurosci **24**(46): 10369-10378.

Roesch, M. R., T. Singh, et al. (2009). "Ventral striatal neurons encode the value of the chosen action in rats deciding between differently delayed or sized rewards." J Neurosci **29**(42): 13365-13376.

Romigi, A., F. Placidi, et al. (2008). "Pedunculopontine nucleus stimulation influences REM sleep in Parkinson's disease." Eur J Neurol **15**(7): e64-65.

Rye, D. B. (1997). "Contributions of the pedunculopontine region to normal and altered REM sleep." Sleep **20**(9): 757-788.

Rymar, V. V., R. Sasseville, et al. (2004). "Neurogenesis and stereological morphometry of calretinin-immunoreactive GABAergic interneurons of the neostriatum." J Comp Neurol **469**(3): 325-339.

Saitoh, K., S. Hattori, et al. (2003). "Nigral GABAergic inhibition upon cholinergic neurons in the rat pedunculopontine tegmental nucleus." Eur J Neurosci **18**(4): 879-886.

Sakata, S. and K. D. Harris (2012). "Laminar-dependent effects of cortical state on auditory cortical spontaneous activity." Front Neural Circuits **6**: 109.

Salamone, J. D., M. Correa, et al. (2009). "Dopamine, behavioral economics, and effort." Front Behav Neurosci **3**: 13.

Salzman, C. D., K. H. Britten, et al. (1990). "Cortical microstimulation influences perceptual judgements of motion direction." Nature **346**(6280): 174-177.

Samejima, K., Y. Ueda, et al. (2005). "Representation of action-specific reward values in the striatum." Science **310**(5752): 1337-1340.

Saper, C. B. and T. C. Chelimsky (1984). "A cytoarchitectonic and histochemical study of nucleus basalis and associated cell groups in the normal human brain." Neuroscience **13**(4): 1023-1037.

Saper, C. B., P. M. Fuller, et al. (2010). "Sleep state switching." Neuron **68**(6): 1023-1042.

Sato, H., Y. Hata, et al. (1987). "A functional role of cholinergic innervation to neurons in the cat visual cortex." J Neurophysiol **58**(4): 765-780.

Schultz, W. (1997). "Dopamine neurons and their role in reward mechanisms." Curr Opin Neurobiol **7**(2): 191-197.

Schultz, W., P. Dayan, et al. (1997). "A neural substrate of prediction and reward." Science **275**(5306): 1593-1599.

Schultz, W., P. Dayan, et al. (1997). "A Neural Substrate of Prediction and Reward." Science **275**(5306): 1593-1599.

Schwarting, R. K. and J. P. Huston (1996). "The unilateral 6-hydroxydopamine lesion model in behavioral brain research. Analysis of functional deficits, recovery and treatments." Prog Neurobiol **50**(2-3): 275-331.

Seo, H., D. J. Barraclough, et al. (2009). "Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game." J Neurosci **29**(22): 7278-7289.

Sesack, S. R. and A. A. Grace "Cortico-Basal Ganglia reward network: microcircuitry." Neuropsychopharmacology **35**(1): 27-47.

Shen, W., M. Flajolet, et al. (2008). "Dichotomous dopaminergic control of striatal synaptic plasticity." Science **321**(5890): 848-851.

Shik, M. L., G. N. Orlovskii, et al. (1966). "[Organization of locomotor synergism]." Biofizika **11**(5): 879-886.

Shik, M. L., F. V. Severin, et al. (1966). "[Control of walking and running by means of electric stimulation of the midbrain]." Biofizika **11**(4): 659-666.

Shuen, J. A., M. Chen, et al. (2008). "Drd1a-tdTomato BAC transgenic mice for simultaneous visualization of medium spiny neurons in the direct and indirect pathways of the basal ganglia." J Neurosci **28**(11): 2681-2685.

St George, R. J., J. G. Nutt, et al. (2010). "A meta-regression of the long-term effects of deep brain stimulation on balance and gait in PD." Neurology **75**(14): 1292-1299.

Stefani, A., A. M. Lozano, et al. (2007). "Bilateral deep brain stimulation of the pedunculopontine and subthalamic nuclei in severe Parkinson's disease." Brain **130**(Pt 6): 1596-1607.

Stefani, A., A. Peppe, et al. "The serendipity case of the pedunculopontine nucleus low-frequency brain stimulation: chasing a gait response, finding sleep, and cognition improvement." Front Neurol **4**: 68.

Stefani, A., A. Peppe, et al. (2013). "The serendipity case of the pedunculopontine nucleus low-frequency brain stimulation: chasing a gait response, finding sleep, and cognition improvement." Front Neurol **4**: 68.

Steinberg, E. E., R. Keiflin, et al. "A causal link between prediction errors, dopamine neurons and learning." Nat Neurosci.

Steriade, M., R. C. Dossi, et al. (1991). "Fast oscillations (20-40 Hz) in thalamocortical systems and their potentiation by mesopontine cholinergic nuclei in the cat." Proc Natl Acad Sci U S A **88**(10): 4396-4400.

Sugrue, L. P., G. S. Corrado, et al. (2005). "Choosing the greater of two goods: neural currencies for valuation and decision making." Nat Rev Neurosci **6**(5): 363-375.

Sul, J. H., H. Kim, et al. (2010). "Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making." Neuron **66**(3): 449-460.

Sutton RS, B. A. (1998). "Reinforcement Learning: An Introduction."

Sutton, R. S. B., A. G. (1998). Reinforcement Learning: An Introduction Cambridge, MA, MIT Press.

Swerdlow, N. R., M. A. Geyer, et al. (2001). "Neural circuit regulation of prepulse inhibition of startle in the rat: current knowledge and future challenges." Psychopharmacology (Berl) **156**(2-3): 194-215.

Szerb, J. C. (1967). "Cortical acetylcholine release and electroencephalographic arousal." J Physiol **192**(2): 329-343.

Taha, S. A. and H. L. Fields (2005). "Encoding of palatability and appetitive behaviors by distinct neuronal populations in the nucleus accumbens." J Neurosci **25**(5): 1193-1202.

Tai, L. H., A. M. Lee, et al. (2012). "Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value." Nat Neurosci **15**(9): 1281-1289.

Takahashi, Y., G. Schoenbaum, et al. (2008). "Silencing the critics: understanding the effects of cocaine sensitization on dorsolateral and ventral striatum in the context of an actor/critic model." Front Neurosci **2**(1): 86-99.

Thankachan, S., P. M. Fuller, et al. (2012). "Movement- and behavioral state-dependent activity of pontine reticulospinal neurons." Neuroscience **221**: 125-139.

Thomas, M. J., C. Beurrier, et al. (2001). "Long-term depression in the nucleus accumbens: a neural correlate of behavioral sensitization to cocaine." Nat Neurosci **4**(12): 1217-1223.

Thorn, C. A., H. Atallah, et al. "Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning." Neuron **66**(5): 781-795.

Thorn, C. A., H. Atallah, et al. (2010). "Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning." Neuron **66**(5): 781-795.

Turner, R. S. and M. Desmurget (2010). "Basal ganglia contributions to motor control: a vigorous tutor." Curr Opin Neurobiol **20**(6): 704-716.

van der Meer, M. A., A. Johnson, et al. (2010). "Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task." Neuron **67**(1): 25-32.

van der Meer, M. A. and A. D. Redish (2011). "Theta phase precession in rat ventral striatum links place and reward information." J Neurosci **31**(8): 2843-2854.

van der Meer, M. A. and A. D. Redish (2011). "Ventral striatum: a critical look at models of learning and evaluation." Curr Opin Neurobiol **21**(3): 387-392.

Voorn, P., L. J. Vanderschuren, et al. (2004). "Putting a spin on the dorsal-ventral divide of the striatum." Trends Neurosci **27**(8): 468-474.

Wanat, M. J., A. Bonci, et al. (2013). "CRF acts in the midbrain to attenuate accumbens dopamine release to rewards but not their predictors." Nat Neurosci **16**(4): 383-385.

Watanabe, K. and O. Hikosaka (2005). "Immediate changes in anticipatory activity of caudate neurons associated with reversal of position-reward contingency." J Neurophysiol **94**(3): 1879-1887.

Wikler, A. (1952). "Pharmacologic dissociation of behavior and EEG "sleep patterns" in dogs; morphine, n-allylnormorphine, and atropine." Proc Soc Exp Biol Med **79**(2): 261-265.

Willuhn, I., L. M. Burgeno, et al. "Hierarchical recruitment of phasic dopamine signaling in the striatum during the progression of cocaine use." Proc Natl Acad Sci U S A **109**(50): 20703-20708.

Znamenskiy, P. and A. M. Zador "Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination." Nature **497**(7450): 482-485.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses, dissertations, and manuscripts. Copies of all UCSF theses, dissertations, and manuscripts will be routed to the library via the Graduate Division. The library will make all theses, dissertations, and manuscripts accessible to the public and will preserve these to the best of their abilities, in perpetuity.

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis, dissertation, or manuscript to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.

Author Signature

9/10/13
Date