Published in final edited form as:

Curr Opin Neurobiol. 2005 December ; 15(6): 721–729.

Beetles, boxes and brain cells: neural mechanisms underlying valuation and learning

C Daniel Salzman, Marina A Belova, and Joseph J Paton

Center for Neurobiology and Behavior, Department of Psychiatry, Columbia University, 1051 Riverside Drive, Unit 87, New York, NY, 10032, USA

Abstract

Sensory cues in the environment can predict the availability of reward. Through experience, humans and animals learn these predictions and use them to guide their actions. For example, we can learn to discriminate chanterelles from ordinary champignons through experience. Assuming the development of a taste for the complex and lingering flavors of chanterelles, we therefore learn to value the same action – picking mushrooms – differentially depending upon the appearance of a mushroom. One major goal of cognitive neuroscience is to understand the neural mechanisms that underlie this sort of learning. Because the acquisition of rewards motivates much behavior, recent efforts have focused on describing the neural signals related to learning the value of stimuli and actions. Neurons in the basal ganglia, in midbrain dopamine areas, in frontal and parietal cortices and in other brain areas, all modulate their activity in relation to aspects of learning. By training monkeys on various behavioral tasks, recent studies have begun to characterize how neural signals represent distinct processes, such as the timing of events, motivation, absolute (objective) and relative (subjective) valuation, and the formation of associative links between stimuli and potential actions. In addition, a number of studies have either further characterized dopamine signals or sought to determine how such signaling might interact with target structures, such as the striatum and rhinal cortex, to underlie learning.

Introduction

“Now someone tells me that *he* knows what pain is only from his own case!—Suppose everyone had a box with something in it: we call it a ‘beetle’. No one can look into anyone else’s box, and everyone says he knows what a beetle is only by looking at *his* beetle.—Here it would be quite possible for everyone to have something different in his box” [1].

This brief passage, taken from Wittgenstein’s seminal work *Philosophical Investigations* [1], helps to frame a special challenge for cognitive neuroscientists. Wittgenstein argues that the linguistic meaning of words that describe mental processes — sensations, perception, learning, memory, attention and decisions — do not derive from their referring to brain processes, because we learn to use these terms correctly without learning about what is inside the box. Despite this concern, cognitive neuroscientists (ourselves included) are rigorously pursuing investigations into how neural activity mediates cognitive functions — we’re looking inside the box, trying to describe how the ‘beetles’ relate to mental processes. In doing so, neuroscientists operationally define a particular process, such as learning — in essence creating a meta-language for mental terms — and then try to decipher the neural mechanisms underlying a more narrowly defined mental function.

Here, we review a series of studies published in 2004 and later that were conducted in alert, behaving monkeys. Through careful experimental manipulations, these studies have begun to delineate the neural mechanisms underlying distinctive functions related to learning. Several excellent recent reviews have focused on aspects of learning, such as the role of dopamine in learning and motivation [2–4], the relationship between dopamine neurons and memory formation [5], the connection between decision-making and learned neural representations of value [6•,7], the role of the amygdala and prefrontal cortex in aspects of reinforcement learning [8], and the connections between reward processing and learning [9,10], among others. Given the extensive attention to these topics, our goal here, inspired by the philosophical issues raised above, is to focus on how manipulating specific behavioral parameters or neural processing can help to disambiguate distinctive processes related to learning.

Computational models of reinforcement learning describe how subjects can learn to use sensory cues to take particular actions on the basis of their experience [11–14]. A conceptual framework for this type of learning is depicted in Figure 1. This framework is conceptually similar to other frameworks outlined recently [6•] and is intended less as a formal model than a heuristic device. Decisions to act are frequently based on learning that a particular stimulus predicts a reward if a particular action is pursued, a process that might, therefore, involve the formation of a representation of the value of stimulus–action contingencies. If this model of learning is correct, the representation of value should be updated with experience. Therefore, error signals — representing the difference between what is expected and what actually occurs — frequently are assumed to drive the adjustment of a value representation during learning. Extensive evidence now indicates that for learning about rewards, midbrain dopamine neurons carry such error signals [11]. Dopaminergic input to the basal ganglia, to frontal cortex, and to other structures, is, therefore, hypothesized to update representations of value. However, additional variables must also be represented to make optimal decisions to act. Representations of value clearly can be altered by a variety of factors, including a subjects' internal state — we can be thirsty, hungry, or lacking wealth — and the behavioral context in which stimuli are presented. Furthermore, learning about a stimulus–action contingency also involves knowing when to act so that one can, for example, predict the timing of salient events, such as upcoming reward. Thus, the representation of value might need to take in to account the temporal relationship between a stimulus and subsequent reinforcement (a temporally discounted representation of value), in addition to behavioral context and physiological state. To gain insights into the mechanisms underlying these aspects of learning, the studies we review below manipulate one or more parameters to induce changes in the context, predictability, timing or contingency of stimuli and actions, either while measuring neural activity or while manipulating neural processing through molecular interventions.

Distinguishing among neural representations of objective value, subjective value and motivation

Many studies have investigated how reward size modulates neural activity in anticipation of a reward [9], and several recently reviewed studies have described neural activity that either reflects or participates in updating, the value of an expected reward [6•,15,16•,17•]. Because multiple processes, such as subjective and objective valuation, motivation, arousal and attention, can co-vary with reward magnitudes, probabilities or contingencies [18•], a central challenge is to delineate exactly what is encoded in each brain area and how it is related to processing in other brain areas. Roesch and Olson [19•] recently set out to disentangle whether neural signals that are modulated by reward in orbitofrontal cortex (OFC) and post-arcuate premotor cortex (PMC) encode the value of an expected reward or monkeys' motivation. They recorded the activity of neurons in either OFC or PMC while monkeys performed a reward–penalty task. In this task, reward magnitude and a penalty 'time-out' were varied independently. Several behavioral measures indicated that motivation was increased by either a large reward

or a long penalty time. Large rewards increased activity across the population of neurons in both areas, but longer penalty times increased activity in PMC, and slightly decreased OFC activity. Although there was some overlap in properties across the two brain areas, these results strongly suggest that neural representations co-varying with reward magnitude can be disambiguated, with value (perhaps subjective value) preferentially encoded in OFC, whereas motivational level is encoded in PMC.

Although the study by Roesch and Olson [19•] provides evidence that helps disassociate neural signals related to value from those related to motivation, it does not distinguish between signals related to objective and subjective value. A recent study by McCoy and Platt [20•] specifically sought to discriminate between representations of objective and subjective value. They recorded the activity of neurons in the posterior cingulate cortex (CGp), a part of the brain implicated in visual orienting and reward processing [21,22]. Monkeys performed a task in which they chose to make a saccadic eye movement to either a 'certain' target, resulting in a fixed reward amount, or a 'risky' target, resulting in an uncertain reward amount that, on average across trials, equaled the certain target reward size. In this task, average objective reward value was constant at both targets. However, by increasing the variance of reward amounts at the risky target, the authors sought to manipulate how monkeys subjectively valued the two targets. Indeed, monkeys preferred the risky target, increasingly choosing it as risk (i.e. variance) increased. Moreover, the activity of CGp neurons was modulated by reward uncertainty. The authors suggest that CGp might participate in evaluating subjective values, and not objective values, to guide orienting eye movements, because the average objective value was the same for the two choices. Similar activity in the parietal lobe probably occurs later in trials than the authors describe in CGp [16•,17•]. If the authors' conclusions are correct, then CGp signals evaluating subjective value early in trials should develop at least as fast as monkeys learn the location of the risky target.

Neural representations of time and their relation to value

Learning requires knowing not only what to do but when to do it. Moreover, the value of a reward is affected by knowing when it will be received; rewards in the distant future are worth less than rewards coming tomorrow. Both of these facts make it crucial that time is represented with respect to actions and reinforcement. Information about timing is probably encoded in a network of neurons including the basal ganglia, cerebellum and the cerebral cortex [23]. Several recent studies measured signals in prefrontal [24,25] and parietal [26•,27] cortices showing that neurons might convey information about time relative to specific events in a task. For example, Tsujimoto and Sawaguchi [24] have investigated the temporal prediction of reward in the dorsolateral prefrontal cortex (DLPFC). Monkeys learned whether reward would be delivered after a short or long delay on the basis of the cues presented during the trial. About one-third of recorded neurons signaled the proximity of reward availability. Although these neural signals might encode the subjective value of the reward rather than purely timing, they still reflect an evaluation of pending reward with respect to time. Another study has examined the role of the OFC in coding the value of time in a similar task [25]. Monkeys were trained on a memory guided saccade task in which two parameters were varied: the proximity and magnitude of rewards. Neurons fired more strongly to cues predicting either a short delay or a large reward, suggesting that OFC neurons provide a time-discounted representation of reward value.

Neuronal responses in the lateral intraparietal area (LIP) also reflect the influence of learned time intervals, as neural firing rates change when monkeys anticipate the timing of an instructed saccade [26•]. Janssen and Shadlen [26•] trained monkeys to make eye movements after a delay; the probability distribution of delay lengths was fixed throughout blocks of trials and corresponded to either a bimodal or a unimodal time schedule. Many neurons in area LIP

modulated their firing rates in a manner resembling the theoretically predicted anticipation functions associated with the two schedules. This strongly suggests that the parietal cortex is a part of complex neural circuitry for representing the time structure of environmental cues over a range of seconds. However, the fact that activity in LIP is also correlated with the value of target choices [16•,17•,28], suggests that LIP receives information about both time and reward value of possible actions. One potentially unifying hypothesis is that LIP represents the salience of visual stimuli that are the targets of eye movements or visual attention [29–31], and that this salience representation might also be modified by timing anticipation, valuation and decisions to make eye movements [32].

Neural activity during learning in multiple brain areas

An increasing number of studies have investigated how neural activity changes as a function of trial number during learning either about value or about stimulus–action contingencies [33,34•,35•,36–41]. Many of the neurons described reside in either the basal ganglia or the frontal cortex. These structures are interconnected in cortico–basal ganglionic loops, in which the frontal cortex might directly influence processing in the striatum, and the striatum might indirectly influence cortical processing via a pathway through the globus pallidus and thalamus [42]. Three recent studies have described how activity in one or more of these areas relates to learning, supplying new details regarding how processing might occur within these neural circuits.

During separate sessions, Brasted and Wise [34•] recorded the activity of cells in the dorsal premotor cortex (PMd) and the striatum (predominantly in the putamen) while monkeys learned visuomotor associations. Neurons in both brain areas changed their activity at approximately the same time with respect to learning. Although neural activity continued to change after monkeys reached asymptotic performance, and, on average, neurons changed activity after monkeys learned, a significant portion of neurons changed activity at a similar time to behavioral learning. These data suggest that the striatum and PMd function in tandem during learning. By contrast, Pasupathy and Miller [35•] found that changes in activity in another part of striatum — the caudate — precede changes in prefrontal cortex (PFC, areas 9 and 46) during learning of a different visuomotor task. In their task, the associative rules of the visual cues were reversed multiple times within a session. The strength of signals encoding the direction of the eye movement in response to the cue increased more rapidly in the caudate than in the PFC (Figure 2). However, activity in PFC was more closely correlated to behavior than activity in the caudate; the authors suggest that perhaps caudate neural activity gradually trains PFC activity to support behavioral learning. In a third study, Watanabe and Hikosaka [43•] trained monkeys to perform a visually guided saccade task in which only one of two target locations were rewarded within a block of trials, with the rewarded location reversed without warning from block to block. Monkeys learned the reversal after only one trial whether or not the first trial after reversal was rewarded, indicating that monkeys had learned a ‘reversal set rule’ regarding the position–reward contingency. Strikingly, caudate neural activity also reflected this learning. The authors suggest that the rapid changes in neural activity could be driven by inputs from prefrontal cortex, or from dopamine neurons carrying error signals.

Among these three studies, there are clear differences in the timing of changes in neural activity in the basal ganglia and frontal cortex. Variation in recording sites, analytic methods, and task design (tasks differed in operant responses, the inclusion of cue reversals and the use of spatial cues) could, in principle, account for these differences. Future studies that use a variety of complex tasks might reveal distinctive physiological properties besides differences in timing in these areas. Furthermore, establishing that neural activity is causally related to learning requires developing experimental techniques that perturb neural activity in functionally

characterized neural circuits and examining the effects of such interventions on neural activity in other brain areas and on behavior indicative of learning.

Midbrain dopamine neurons: mechanisms of action at target structures and the influence of context on physiological responses

Dopamine (DA) neurons carry a reward prediction error signal that reflects the difference between what is expected and what is actually received [3,11,44]. Because dopamine neurons have a low baseline firing rate, their signals represent positive prediction errors (increases in firing after an unexpected reward) more precisely than negative prediction errors (inhibition after omission of an expected reward) [45]. Recent work has further characterized the properties of DA error signals. Tobler *et al.* [46••] show that DA neurons adaptively encode the relative, as opposed to the absolute, values of reward available in a particular context. For example, DA neurons responded with a phasic increase or decrease in firing to the same sized reward depending upon whether it was shown in a block of trials containing smaller or larger rewards, respectively (Figure 3). Using a task that varied the probability of reward depending upon the number of trials since the last reward, Nakahara *et al.* [47•] found that DA responses also reflected this context. Thus, information about the range of rewards available, in addition to reward probability over time, can modulate DA neural responses, resulting in an adaptive representation of error signals. It remains unknown whether this information is provided directly to DA neurons, or provided indirectly via a value representation that takes context into account. Relatively less attention has been focused on how DA error signals are processed and used by their target structures. In an elegant set of experiments, Morris *et al.* [48••] describe and compare the signals carried by both dopamine neurons in the midbrain and tonically active neurons (TANs) in the striatum during a probabilistic instrumental conditioning task. The authors find that whereas DA neurons indeed carry a reward prediction error signal, TANs pause their activity in response to any reward-related event and are unaffected by changes in reward probability. This pause is coincident with reward prediction errors carried by DA neurons (Figure 4). TANs are probably cholinergic interneurons, which inhibit striatal medium spiny projection neurons (MSNs). Midbrain DA neurons also target MSNs. Therefore, the pause in TAN activity following reward-related events could function to facilitate the transmission of teaching signals from DA neurons to MSNs by reducing inhibitory input onto MSNs. Indeed, this might be necessary to compensate for the relatively temporally imprecise nature of information carried by extracellular DA levels [49].

Both DA and some components of acetylcholine (ACh) signaling at MSN synapses are mediated by relatively slower signaling mechanisms [50,51]. How then can error signals guide rapid and precise changes in neural activity? One possibility is that DA neurons might transmit an additional faster signal to MSNs. Recent studies by Rayport and co-workers [52] suggest that glutamate might be co-transmitted along with DA. In principle, glutamatergic activity might relay the reward prediction error signal quickly, gated in part by a transient decrease in cholinergic tone. Still, DA could serve another important function during learning. In rats, Bamford *et al.* [53••] show that DA acts through D2 receptors to filter out the least active among motor area cortico-striatal synapses. Taken together, these studies suggest that extracellular DA might mediate reinforcement of strong corticostriatal inputs relatively slowly, whereas glutamate faithfully transmits reward prediction error signals rapidly. Further experiments in which glutamate signaling by DA neurons is perturbed during performance of conditioning tasks, perhaps by interfering with the action of glutamate transporters, could provide further insights into these mechanisms.

Midbrain dopamine neurons project not only to the striatum but also to a diffuse range of brain areas including prefrontal and rhinal cortices. Perhaps the strongest evidence that DA signaling is required for learning about the value of sensory stimuli comes from a recent study by Liu

and colleagues [54••] that focused on the dense dopamine projection to rhinal cortices [55–58]. Representing a significant technical advance, these investigators injected DNA constructs designed to express the antisense of DA D2 receptor mRNA into monkeys' rhinal cortex. In the mouse, this process significantly reduces D2 ligand binding [59]. Monkeys were then trained on a reward schedule task in which visual cues indicate how many trials must be completed before reward delivery. As is the case with rhinal cortex lesions, a decrease in D2 receptor expression in the rhinal cortex blocked learning of the predictive visual cues. However, unlike after lesions, monkeys with the D2 antisense treatment could learn to use the visual cues again after recovering for a few weeks. This represents the first use of molecular biological techniques to knock down expression of a particular receptor type in the awake, behaving monkey; this class of tools should prove especially powerful in understanding the roles of different receptor classes in carefully defined cognitive functions. Ultimately, electrophysiological studies carried out in parallel with such molecular interventions will enable scientists to connect functionality at the molecular level with physiology and behavior in the primate.

Conclusions

Activity in neurons located in the frontal, parietal and rhinal cortices, the basal ganglia and other structures underlies learning about the values of stimuli and of potential actions in response to stimuli. By manipulating reward magnitude, reward probability and task event timing, the studies reviewed here demonstrate how distinctive coding by neurons might underlie specifically defined cognitive functions. Additional studies have begun to investigate the relative timing of activity in different brain structures during learning, and to establish the significance of dopamine in aspects of learning through manipulations at the molecular level. By combining sophisticated behavioral tasks — that operationally define particular mental functions — with neurophysiological recording and molecular manipulations, cognitive neuroscientists are able to proceed in describing underlying brain processes related to learning — the beetles in Wittgenstein's box.

Acknowledgements

We thank ME Goldberg, J Gottlieb and SE Morrison for helpful comments on the manuscript. This work was supported by the Keck foundation, by grants from the National Institutes of Health (MH K01 01724), and the Klingenstein, Sloan, and National Alliance for Research on Schizophrenia and Depression foundations, and by a Charles E. Culpeper Scholarship award from Goldman Philanthropic Partnerships to CD Salzman. JJ Paton received support from National Institute of Child Health and Human Development and National Eye Institute institutional training grants.

References and recommended reading

Papers of particular interest, published within the annual period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Wittgenstein, W. Section 293. In: Anscombe, G.; Rhees, R., editors. *Philosophical Investigations*. 2. Blackwell; 1958.
 2. Wise RA. Dopamine, learning and motivation. *Nat Rev Neurosci* 2004;5:483–494. [PubMed: 15152198]
 3. Montague PR, Hyman SE, Cohen JD. Computational roles for dopamine in behavioural control. *Nature* 2004;431:760–767. [PubMed: 15483596]
 4. Ungless MA. Dopamine: the salient issue. *Trends Neurosci* 2004;27:702–706. [PubMed: 15541509]

5. Lisman JE, Grace AA. The hippocampal-VTA loop: controlling the entry of information into long-term memory. *Neuron* 2005;46:703–713. [PubMed: 15924857]
- 6••. Sugrue LP, Corrado GS, Newsome WT. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat Rev Neurosci* 2005;6:363–375. [PubMed: 15832198]The authors present an up-to-date review about the neural signals underlying valuation and decision making with a focus on several recent studies in LIP and PFC
7. Glimcher PW, Rustichini A. Neuroeconomics: the consilience of brain and decision. *Science* 2004;306:447–452. [PubMed: 15486291]
8. Holland PC, Gallagher M. Amygdala-frontal interactions and reward expectancy. *Curr Opin Neurobiol* 2004;14:148–155. [PubMed: 15082318]
9. Schultz W. Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. *Curr Opin Neurobiol* 2004;14:139–147. [PubMed: 15082317]
10. Schultz W, Tremblay L, Hollerman JR. Changes in behavior-related neuronal activity in the striatum during learning. *Trends Neurosci* 2003;26:321–328. [PubMed: 12798602]
11. Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science* 1997;275:1593–1599. [PubMed: 9054347]
12. Montague PR, Berns GS. Neural economics and the biological substrates of valuation. *Neuron* 2002;36:265–284. [PubMed: 12383781]
13. Sutton, R.; Barto, A. Reinforcement Learning. MIT Press; 1998.
14. Dayan, P.; Abbott, LF. Theoretical Neuroscience. MIT Press; 2001. Classical conditioning and reinforcement learning; p. 331–358.
15. Barraclough DJ, Conroy ML, Lee D. Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci* 2004;7:404–410. [PubMed: 15004564]
- 16•. Dorris MC, Glimcher PW. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 2004;44:365–378. [PubMed: 15473973]This study examines neural activity recorded from LIP while monkeys play a strategic computer game. The authors argue that LIP activity reflects the relative subjective desirability of actions in the game
- 17•. Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. *Science* 2004;304:1782–1787. [PubMed: 15205529]This study describes the activity of LIP neurons while monkeys chose to make eye movements to targets that supplied rewards at different rates. They compared the neural representation of the value of choices both to monkeys' performance and to predictions derived from a computational model. LIP neural activity reflects the relative value of competing actions, matching predictions from their model and the monkeys' behavior. Although the authors do not argue that LIP computes the value of choices, their data demonstrates that LIP has access to information about value
- 18••. Maunsell JH. Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci* 2004;8:261–265. [PubMed: 15165551]This is an interesting review that emphasizes the challenges facing cognitive neuroscientists when trying to determine exactly what cognitive process neural signals represent
- 19•. Roesch MR, Olson CR. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 2004;304:307–310. [PubMed: 15073380]This study shows that neural activity reflecting different anticipated reward magnitudes might actually represent either the value of the reward or the degree of motivation that the monkeys have. They used a task that modulated reward magnitude and penalty time to vary both reward value and motivation independently. Although there was some overlap in response properties, in general, orbitofrontal cortex preferentially represented reward value, and premotor cortex activity reflected motivation
- 20•. McCoy AN, Platt ML. Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat Neurosci* 2005;8:1220–1227. [PubMed: 16116449]This work strives to disambiguate neural signals related to objective and subjective value by recording neural activity while monkeys performed a visual gambling task. Monkeys chose between making an eye movement to a target with a certain reward amount, or to a target with a variable reward amount but the average reward of which was the same as the other target. Monkeys chose the uncertain reward target more often, and neural activity in the posterior cingulate cortex reflected this subjective evaluation of reward

21. Olson CR, Musil SY, Goldberg ME. Single neurons in posterior cingulate cortex of behaving macaque: eye movement signals. *J Neurophysiol* 1996;76:3285–3300. [PubMed: 8930273]
22. McCoy AN, Crowley JC, Haghighian G, Dean HL, Platt ML. Saccade reward signals in posterior cingulate cortex. *Neuron* 2003;40:1031–1040. [PubMed: 14659101]
23. Buhusi CV, Meck WH. What makes us tick? Functional and neural mechanisms of interval timing. *Nat Rev Neurosci* 2005;6:755–765. [PubMed: 16163383]
24. Tsujimoto S, Sawaguchi T. Neuronal activity representing temporal prediction of reward in the primate prefrontal cortex. *J Neurophysiol* 2005;93:3687–3692. [PubMed: 15634707]
25. Roesch MR, Olson CR. Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J Neurophysiol* 2005;94:2457–2471. [PubMed: 15958600]
- 26•. Janssen P, Shadlen MN. A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 2005;8:234–241. [PubMed: 15657597] Inspired by the recognition that humans and animals routinely perform tasks that require interval timing, this study demonstrates that activity in LIP flexibly reflects monkeys' knowledge of time intervals during a task requiring eye movements
27. Leon MI, Shadlen MN. Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* 2003;38:317–327. [PubMed: 12718864]
28. Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. *Nature* 1999;400:233–238. [PubMed: 10421364]
29. Bisley JW, Goldberg ME. Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 2003;299:81–86. [PubMed: 12511644]
30. Itti L, Koch C. Computational modelling of visual attention. *Nat Rev Neurosci* 2001;2:194–203. [PubMed: 11256080]
31. Gottlieb J. Parietal mechanisms of target representation. *Curr Opin Neurobiol* 2002;12:134–140. [PubMed: 12015228]
32. Shadlen MN, Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 2001;86:1916–1936. [PubMed: 11600651]
33. Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. Single neurons in the monkey hippocampus and learning of new associations. *Science* 2003;300:1578–1581. [PubMed: 12791995]
- 34•. Brasted PJ, Wise SP. Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. *Eur J Neurosci* 2004;19:721–740. [PubMed: 14984423] The authors compare the time course of changes of neural activity in the dorsal premotor cortex and striatum (primarily the putamen) during performance of a visuomotor associative learning task. On average, neurons in both areas exhibited changes at approximately the same time with respect to behavioral learning. Recordings in the two brain areas were conducted in separate sessions, but the authors alternated between brain areas every few days to ensure that any differences in the two brain areas could not be attributed to long-term changes in the monkeys' performance
- 35•. Pasupathy A, Miller EK. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 2005;433:873–876. [PubMed: 15729344] This technically challenging study employed simultaneous recording of single neurons from dorsolateral prefrontal cortex and the caudate in monkeys performing a visuomotor associative learning task involving frequent stimulus–action contingency reversals. During the learning of such reversals, directional signals in the striatum strengthened more rapidly than similar signals in prefrontal cortex, which were actually more closely correlated with the gradual improvement in monkeys' behavior. The authors suggest that rapid changes in activity in the striatum might gradually train neurons in the prefrontal cortex, rather than the other way around
36. Tremblay L, Hollerman JR, Schultz W. Modifications of reward expectation-related neuronal activity during learning in primate striatum. *J Neurophysiol* 1998;80:964–977. [PubMed: 9705482]
37. Asaad WF, Rainer G, Miller EK. Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 1998;21:1399–1407. [PubMed: 9883732]
38. Lauwereyns J, Watanabe K, Coe B, Hikosaka O. A neural correlate of response bias in monkey caudate nucleus. *Nature* 2002;418:413–417. [PubMed: 12140557]
39. Chen L, Wise S. Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *J Neurophysiol* 1995;73:1101–1121. [PubMed: 7608758]

40. Schoenbaum G, Chiba AA, Gallagher M. Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J Neurosci* 1999;19:1876–1884. [PubMed: 10024371]
41. Hadj-Bouziane F, Boussaoud D. Neuronal activity in the monkey striatum during conditional visuomotor learning. *Exp Brain Res* 2003;153:190–196. [PubMed: 14610634]
42. Alexander GE, DeLong MR, Strick PL. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci* 1986;9:357–381. [PubMed: 3085570]
- 43•. Watanabe K, Hikosaka O. Immediate changes in anticipatory activity of caudate neurons associated with reversal of position-reward contingency. *J Neurophysiol* 2005;94:1879–1887. [PubMed: 15872072] This interesting study demonstrates that neural activity in the caudate can change within one trial after a reversal in position–reward contingencies. The rapid change in activity reflects the fact that the neurons, like the monkeys, might learn about one trial type from their experience with another trial type, which the authors call learning a ‘reversal set rule’
44. Schultz W, Dickinson A. Neuronal coding of prediction errors. *Annu Rev Neurosci* 2000;23:473–500. [PubMed: 10845072]
45. Bayer HM, Glimcher PW. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 2005;47:129–141. [PubMed: 15996553]
- 46••. Tobler PN, Fiorillo CD, Schultz W. Adaptive coding of reward value by dopamine neurons. *Science* 2005;307:1642–1645. [PubMed: 15761155] This study demonstrates that the reward prediction error signals carried by midbrain DA neurons are adaptive to behavioral context, thereby providing a more efficient representation of information across a range of experimental conditions. By manipulating reward magnitudes, probabilities and relative reward amounts, the authors show that the reward prediction error signal encoded by DA neurons is flexible and adaptive
- 47•. Nakahara H, Itoh H, Kawagoe R, Takikawa Y, Hikosaka O. Dopamine neurons can represent context-dependent prediction error. *Neuron* 2004;41:269–280. [PubMed: 14741107] This paper demonstrates that DA error signals might be influenced by the expected reward probability that varies during task performance. Along with the paper by Tobler *et al.* [46••], this study indicates that DA neurons adapt their firing rate depending upon behavioral context.
- 48••. Morris G, Arkadir D, Nevet A, Vaadia E, Bergman H. Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. *Neuron* 2004;43:133–143. [PubMed: 15233923] This study compared the timing of responses of TANs in the striatum and DA neurons in the midbrain during performance of a conditioning task. DA neurons encoded a well described reward prediction error signal, and TANs respond to all motivationally significant events with a pause in activity that was coincident with the DA error signal, and synchronized across TANs. The authors speculate that this pause in TAN activity could act as a temporal gate for DA signaling in the striatum. The experiments were conducted so that the dynamics of operant responses did not vary across trial types, removing a major potential confound.
49. Benoit-Marand M, Jaber M, Gonon F. Release and elimination of dopamine in vivo in mice lacking the dopamine transporter: functional consequences. *Eur J Neurosci* 2000;12:2985–2992. [PubMed: 10971639]
50. Greengard P. The neurobiology of slow synaptic transmission. *Science* 2001;294:1024–1030. [PubMed: 11691979]
51. Zhou FM, Wilson C, Dani JA. Muscarinic and nicotinic cholinergic mechanisms in the mesostriatal dopamine systems. *Neuroscientist* 2003;9:23–36. [PubMed: 12580337]
52. Chuhma N, Zhang H, Masson J, Zhuang X, Sulzer D, Hen R, Rayport S. Dopamine neurons mediate a fast excitatory signal via their glutamatergic synapses. *J Neurosci* 2004;24:972–981. [PubMed: 14749442]
- 53••. Bamford NS, Zhang H, Schmitz Y, Wu NP, Cepeda C, Levine MS, Schmauss C, Zakharenko SS, Zablów L, Sulzer D. Heterosynaptic dopamine neurotransmission selects sets of corticostriatal terminals. *Neuron* 2004;42:653–663. [PubMed: 15157425] A technical tour de force that combined two-photon imaging of cortico–striatal terminal destaining, electrochemical recordings of striatal DA release, pharmacology, intracellular electrophysiology, and electrical stimulation to examine how DA acts on corticostriatal terminals. Results from this study suggest that DA acts through D2 receptors to filter out the least active among motor area cortico–striatal synapses. The authors

suggest that this mechanism could aid in establishing appropriate behaviors during reinforcement learning

- 54••. Liu Z, Richmond BJ, Murray EA, Saunders RC, Steenrod S, Stubblefield BK, Montague DM, Ginns EI. DNA targeting of rhinal cortex D2 receptor protein reversibly blocks learning of cues that predict reward. *Proc Natl Acad Sci USA* 2004;101:12336–12341. [PubMed: 15302926] This innovative study demonstrates the importance of dopaminergic inputs to rhinal cortex on a monkey's ability to use visual stimuli that cue proximity to reward. Dopamine D2 receptor antisense was targeted to the rhinal cortex. This intervention blocked monkeys' ability to learn the significance of visual cues that indicated how many trials must be completed before reward receipt
55. Akil M, Lewis DA. The dopaminergic innervation of monkey entorhinal cortex. *Cereb Cortex* 1993;3:533–550. [PubMed: 7907902]
56. Goldsmith SK, Joyce JN. Dopamine D2 receptors are organized in bands in normal human temporal cortex. *Neuroscience* 1996;74:435–451. [PubMed: 8865195]
57. Richfield EK, Young AB, Penney JB. Comparative distributions of dopamine D-1 and D-2 receptors in the cerebral cortex of rats, cats, and monkeys. *J Comp Neurol* 1989;286:409–426. [PubMed: 2506254]
58. Berger B, Trottier S, Verney C, Gaspar P, Alvarez C. Regional and laminar distribution of the dopamine and serotonin innervation in the macaque cerebral cortex: a radioautographic study. *J Comp Neurol* 1988;273:99–119. [PubMed: 3209731]
59. Zhou LW, Zhang SP, Weiss B. Intrastratial administration of an oligodeoxynucleotide antisense to the D2 dopamine receptor mRNA inhibits D2 dopamine receptor-mediated behavior and D2 dopamine receptors in normal mice and in mice lesioned with 6-hydroxydopamine. *Neurochem Int* 1996;29:583–595. [PubMed: 9113126]

Figure 1.

Conceptual scheme for learning about the value of actions. The boxes represent functions supported by neural systems that might be distributed across brain areas. This scheme posits that sensory input and information about an animal's internal state converge with error signals to form a representation of the value of potential actions in response to sensory stimuli (stimulus-action value representation). Decisions and actions are based on this value representation. Signals related to reward value have been described in multiple brain areas. Midbrain dopamine neurons are frequently described as supplying reward prediction error signals that can drive learning, but in principle other brain structures might also supply error signals depending upon the task and the nature of reinforcement (e.g. for punishments). Decisions and actions might also be implemented by different brain structures depending upon the task. Table 1 shows the various experimental manipulations made in studies reviewed in this article, and the brain areas studied. These manipulations seek to determine which factors are important for neural responses and behavior. These parameters might also have their own representative systems in the brain.

Figure 2.

Changes in directional signals in the caudate precede changes in PFC during visuomotor association learning. **(a,b)** Proportion of neural response variance explained by saccade direction (PEV_{dir} – see color map bar), the operant response on the task, as a function of time (x-axis) and correct trial number relative to reversal (y-axis), across the entire population of cells recorded in PFC (a) and the caudate (Cd) (b). Black dots indicate the first time point for each trial that half of the maximum directional signal across all trials is reached. **(c)** Best fit sigmoids to rise times marked by black dots in a and b. Rise times change more rapidly and move earlier in the trial in the Cd than in the PFC. Reprinted by permission from Macmillan Publishers Ltd: Nature [35•], copyright (2005).

Figure 3.

DA responses to reward delivery scale with the range of rewards predicted by a cue. In this task, each of two visual stimuli was followed by presentation of either larger or smaller volumes of liquid reward, with equal probability. Rasters, with peristimulus time histograms (PSTHs) shown above, are displayed for four conditions (left image: small or medium reward; right image: medium or large reward). Each line in the raster indicates a trial, and each dot indicates an action potential. This cell responded differently to the medium reward amount, depending on whether the other trials in the block contained small or large rewards. DA neurons, therefore, adaptively encode the relative, as opposed to the absolute, amounts of reward available in a particular context. Reprinted with permission from [46••] Copyright (2005) AAAS.

Figure 4.

TANs pause when DA neurons transmit a reward prediction error signal. **(a)** covariation matrix for a simultaneously recorded TAN (rows) and DA neuron (columns), aligned on predictive cue presentation during an instrumental conditioning task. Each pixel corresponds to the percent covariance of spiking within the 10 ms bins indicated by the axes, across the two neurons. The mean PSTH for the TAN is shown along the y-axis, with the mean PSTH for the DA neuron along the x-axis (cue onset, white arrows). The dotted circle highlights the period of maximum negative correlation between firing in the TAN, and firing in the DA neuron, occurring approximately 120–250 ms post cue presentation in both neurons. **(b)** Normalized population average of DA (gray line) and TAN (black line) responses to cue presentation. Reprinted from [48•], Copyright (2004), with permission from Elsevier.

Table 1

Brain areas investigated	Parameters varied			
	Reward probability	Reward magnitude(s)	Timing	Gene expression
Striatum	[34•,35•,43•,48••]			
Midbrain dopamine neurons	[46••,47•,48••]	[46••]	[47•]	[54••]
Rhinal cortices			[54••]	
Prefrontal cortex	[35•]		[24]	
Orbitofrontal cortex		[19•,25]	[19•,25]	
Premotor cortex	[34•]			
Lateral intraparietal area			[26•]	
Posterior cingulate	[20•]	[20•]		

Summary of studies reviewed in this article categorized by brain area and experimental manipulation. Note that manipulations of reward probability encompass tasks in which a particular action cued by a stimulus switches from being rewarded 100% of the time to being unrewarded (e.g. [35•,43•]). Timing refers to manipulations of the time until reward delivery [24,25], the time until the next trial [19•], the reward probability that is modulated as a function of completed trial number [47•,54••], or the distribution of time intervals that a monkey learns [26•].