

OPINION

A new perspective on the role of the orbitofrontal cortex in adaptive behaviour

Geoffrey Schoenbaum, Matthew R. Roesch, Thomas A. Stalnaker and Yuji K. Takahashi

Abstract | The orbitofrontal cortex (OFC) is crucial for changing established behaviour in the face of unexpected outcomes. This function has been attributed to the role of the OFC in response inhibition or to the idea that the OFC is a rapidly flexible associative-learning area. However, recent data contradict these accounts, and instead suggest that the OFC is crucial for signalling outcome expectancies. We suggest that this function — signalling of expected outcomes — can also explain the crucial role of the OFC in changing behaviour in the face of unexpected outcomes.

The orbitofrontal cortex (OFC) has long been associated with adaptive, flexible behaviour in the face of changing contingencies and unexpected outcomes. In 1868, John Harlow¹ described the erratic, inflexible, stimulus-bound behaviour of Phineas Gage, who reportedly suffered extensive damage to the orbital and midline prefrontal regions². Since then, owing to increasingly refined experimental work, the regulation of flexible behaviour has been attributed to the OFC, a set of loosely defined areas in the prefrontal regions that overlie the orbits.

In this Perspective, we first review data showing that the OFC is crucial for changing behaviour in the face of unexpected outcomes. Then we describe the two dominant hypotheses that have been advanced to explain this function, and follow this with data that directly contradict both of these. We focus on a region of the OFC for which there are strong anatomical and functional parallels across rodents and primates; this area encompasses the dorsal bank of the rhinal sulcus in rats, including lateral orbital regions and anterior parts of the agranular insular cortex. These cortical fields have a pattern of connectivity with the rostral basolateral amygdala, the ventral

striatum, the mediodorsal thalamus and sensory regions that is qualitatively similar to the pattern of connectivity of areas 11, 12 and 13 in the primate orbital prefrontal cortex^{3,4}. Furthermore, as we outline below, neurophysiological and lesion studies show remarkable similarities in the functions of these areas across species. These studies indicate that this part of the OFC is crucial for signalling information about expected outcomes and for using that information to guide behaviour. We suggest that signals from this part of the OFC that relate to expected outcomes also contribute to the detection of errors in reward prediction when contingencies are changing, thereby facilitating changes in associative representations in other brain areas and, ultimately, behaviour.

The OFC is crucial for flexible behaviour

The role of the OFC in flexible behaviour is typically assessed using tasks in which subjects must change an established behavioural response in order to adapt to new contingencies. FIGURE 1 shows an example in rats performing a simple odour discrimination task. Here, rats are trained to sample odour cues at a centrally located port and

then visit a nearby fluid well. They learn that one odour predicts that the well contains sucrose and that a different odour predicts that it contains quinine. Because rats like sucrose but want to avoid drinking quinine, they learn to discriminate between the two odour cues, drinking from the well after sampling the sucrose-predicting cue and withholding that response after sampling the quinine-predicting cue. After learning, their ability to rapidly reverse their responses can be assessed by switching the odour–outcome associations. OFC lesions cause an impairment in animals' ability to rapidly reverse responding in this setting, a deficit that has been shown in different laboratories over the past 50 years, across species, experimental designs and learning tasks^{5–16}. The reliability of these findings provides convincing evidence that the OFC performs a crucial function in flexible behaviour. Two hypotheses have dominated the discussion of what this function might be; they are reviewed below.

The OFC as response inhibitor. The first and perhaps most prevalent hypothesis is that the OFC is crucial for flexible behaviour generally — and reversal learning in particular — because it inhibits inappropriate responses (FIG. 1b). This idea was introduced by David Ferrier, who proposed, on the basis of experiments in dogs and monkeys, that the frontal lobes are crucial to suppressing motor acts in favour of attention and planning¹⁷. During the past few decades, response inhibition has become specifically associated with orbitofrontal function because it provides an attractive description of the general and specific behavioural effects of orbitofrontal damage¹⁸. For example, this explanation fits well with the constellation of symptoms caused by orbitofrontal damage in humans, which include impulsivity, perseveration and compulsive behaviours¹⁹, and it also provides an excellent description of reversal-learning deficits^{5–16}, as well as deficits in detour reaching tasks and stop signal tasks^{20,21}.

However, this hypothesis does not have good predictive power outside of these settings. In fact, it is easy to find behaviours that require response inhibition that are not affected by OFC damage. For example, in many of the reversal studies described

earlier, subjects with lesions of the OFC could successfully inhibit responses during initial learning even though they exhibited difficulties in inhibiting them after reversal. Similarly, the OFC is not required for inhibition of 'pre-potent' or innate response tendencies. For example, in reinforcer devaluation tasks, animals with orbitofrontal lesions have trouble changing cue-evoked responding, but they are readily able to withhold the selection or consumption of food that has been paired with illness or that has been fed to satiety^{12,22–26} (but see REF. 27). Another example comes from a study²⁸ in which

monkeys were allowed to choose between differently sized peanut rewards but had to select the smaller reward in order to receive the larger one. Monkeys with OFC damage inhibited their innate tendency to select the larger reward just as well as controls. So overall the OFC — or at least the region that seems the most analogous across species — does not seem to have a general role in response inhibition.

Of course, this does not imply that whatever role the OFC has is not important for inhibiting responses in some settings — it only implies that this is not the root function

of the OFC. Furthermore, it is of course possible that other prefrontal areas — perhaps even some that have been termed OFC but that lie outside of the area that we focus on here — mediate a more general inhibitory function. Indeed, this has recently been suggested²⁷. However, at least for the region of the OFC that is targeted by the studies described above, response inhibition does not seem to be a general function.

The OFC as rapidly flexible encoder of associative information. More recently it was suggested that the OFC is crucial

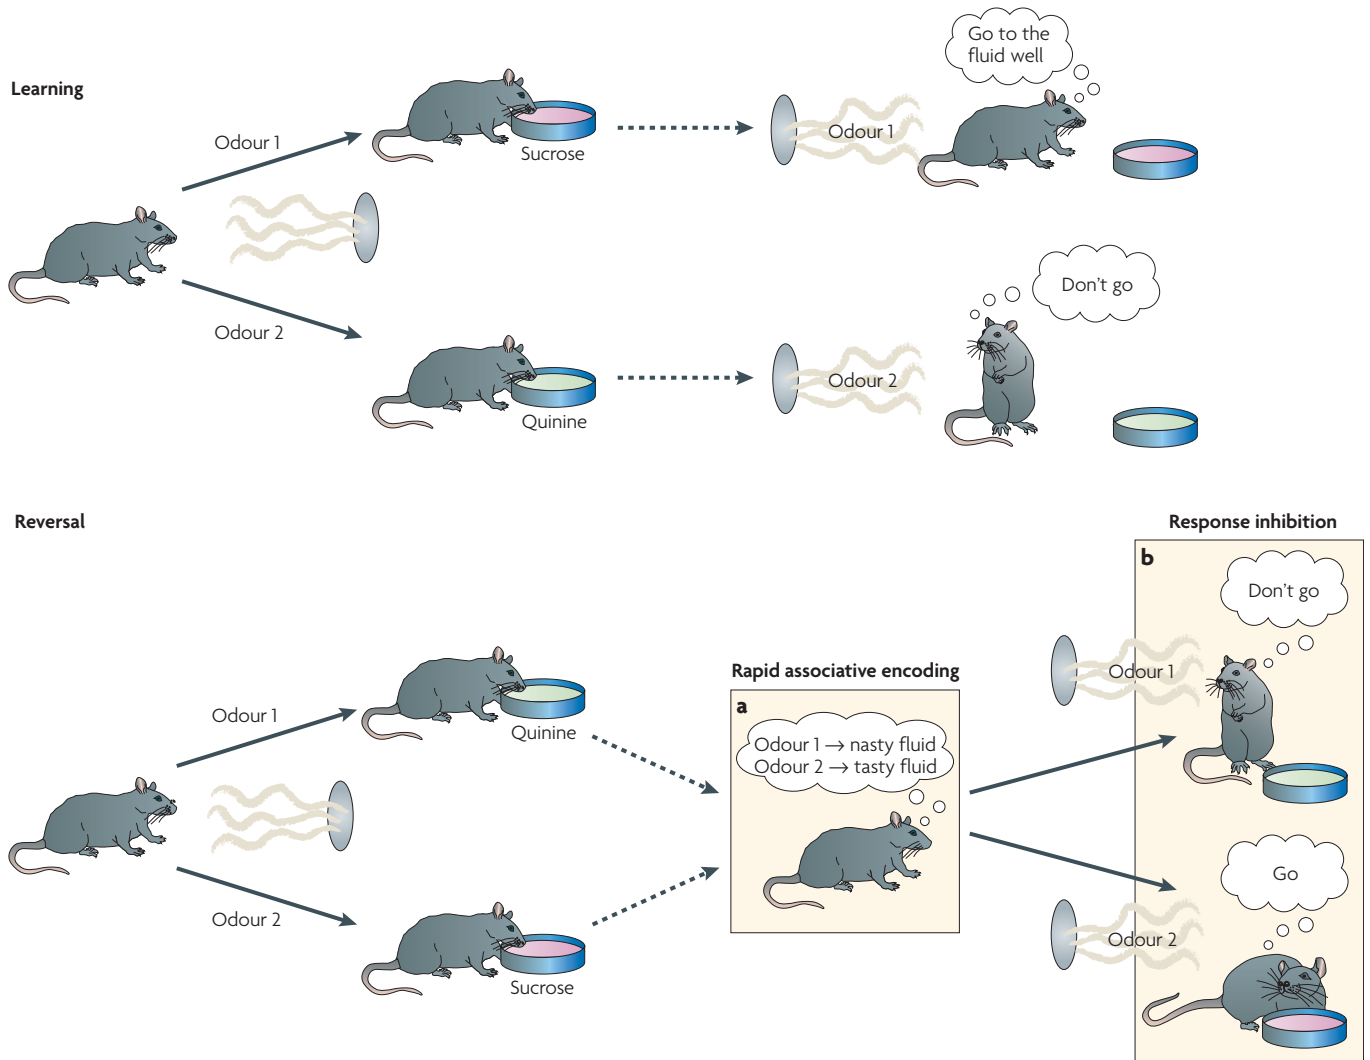


Figure 1 | Showing orbitofrontal involvement in flexible, adaptive behaviour in rats using an odour discrimination reversal task. The events in a simple odour discrimination reversal task that demonstrates the involvement of the orbitofrontal cortex in flexible, adaptive behaviour. Rats sample an odour from a port. After sampling, they decide whether to respond at a nearby fluid well. One odour predicts a tasty sucrose solution, whereas a second odour predicts an unpleasant-tasting quinine solution. As illustrated in the top panel, rats rapidly learn to discriminate between the odours, drinking from the fluid well after sampling

the sucrose-predicting odour but withholding that response after sampling the quinine-predicting odour. After this has been learned, the odour–outcome associations can be reversed, such that the sucrose-predicting odour comes to predict quinine and vice versa. This is illustrated in the bottom panel. Initially rats respond on the basis of what they have learned previously. However, normal rats rapidly recode these associations (a) and inhibit their old response patterns in favour of new ones (b); rats with damage to the orbitofrontal cortex fail to change their behaviour as rapidly.

for flexible behaviour because it is better than other brain areas at rapidly encoding new associations, particularly those between cues and outcomes²⁹. According to this hypothesis, the OFC supports flexible behaviour by signalling this associative information to other areas, thereby driving correct responses or, perhaps, inhibiting incorrect ones (FIG. 1a).

This idea was based on single-unit recording studies showing that cue-evoked neural activity in the OFC rapidly reflects new associative information during reversal learning. This was first reported by Rolls and colleagues³⁰, who noted that, in monkeys, single units in the OFC often reversed their firing response between two stimuli (that is, switched from firing in response to one stimulus to firing in response to another stimulus) when their associations with reward were reversed. This reversal of encoding was initially shown for different syringes used to deliver appetitive and aversive fluids³⁰ and subsequently shown for visual and olfactory discrimination^{31,32}.

These findings, together with OFC-dependent reversal learning impairments, are consistent with the idea that the OFC supports flexible behaviour because it stores the new associative information more rapidly or in a more functional way than other brain areas. However, like the response inhibition hypothesis, this idea is not consistent with more recent evidence from studies that directly tested its predictions.

For example, only 25% of cue-selective neurons in the OFC reverse encoding³³. During reversal learning, cue-selective neurons in many other brain regions reverse firing much more rapidly and in greater proportions than do OFC neurons. In the rat basolateral amygdala, for instance, 55–60% of cue-selective neurons reverse firing³², and similar findings have been reported in the monkey amygdala³⁴. Moreover, there is no relationship between cue preference or the degree of cue selectivity before and after reversal in OFC neurons. Thus, reversal of encoding in the OFC is neither uniquely rapid nor particularly pervasive.

These results are contrary to the basic version of this proposal in which the OFC is simply quicker or more efficient than other areas at storing new associations. An alternative version of this proposal is that storage in the OFC might be more functional — owing to connectivity or for some other reason — than similar information storage in the amygdala or elsewhere. Importantly, this version would not require that associations be stored more rapidly in the OFC, but it would predict a close relationship between reversal of encoding in the OFC and reversal of behaviour. In other words, reversal of encoding in OFC neurons should be associated with rapid reversal learning. But we have observed precisely the opposite relationship; that is, rats reverse their behaviour more slowly when there is rapid reversal of encoding in the OFC³³. This, together with other evidence presented below³⁵, is inconsistent with the rapidly flexible encoding hypothesis.

The OFC signals outcome expectancies

If the OFC does not have a special role in response inhibition and is not better than other regions at encoding new associative information, then what is its crucial function in flexible behaviour? We suggest that this function involves signalling outcome expectancies. Specifically, we suggest that the OFC signals the predicted characteristics — such as sensory properties (size, shape, texture and flavour) and perhaps even contextual cues such as particular timing or likelihood — and unique value of specific outcomes that an animal expects given particular circumstances and cues in the environment³.

Evidence from neuronal recording studies.

Evidence in support of this proposal comes from neuronal recording studies that demonstrate anticipatory firing in OFC neurons. Although these anticipatory signals are also apparent before other events, they are particularly strong in the OFC before the animal experiences primary rewarding or punishing events. For example, during discrimination learning, distinct populations of neurons in the rat OFC develop selective firing in anticipation of delivery of sucrose or quinine³⁶. Often, these neurons initially fire in response to one or the other outcome, then come to fire in anticipation of that outcome and finally also fire in response to cues that predict the outcome (FIG. 2a). Unlike reward-responsive dopamine neurons, which also transfer activity from outcomes to predictive events^{37–41}, these OFC neurons do not stop firing in response to the reward

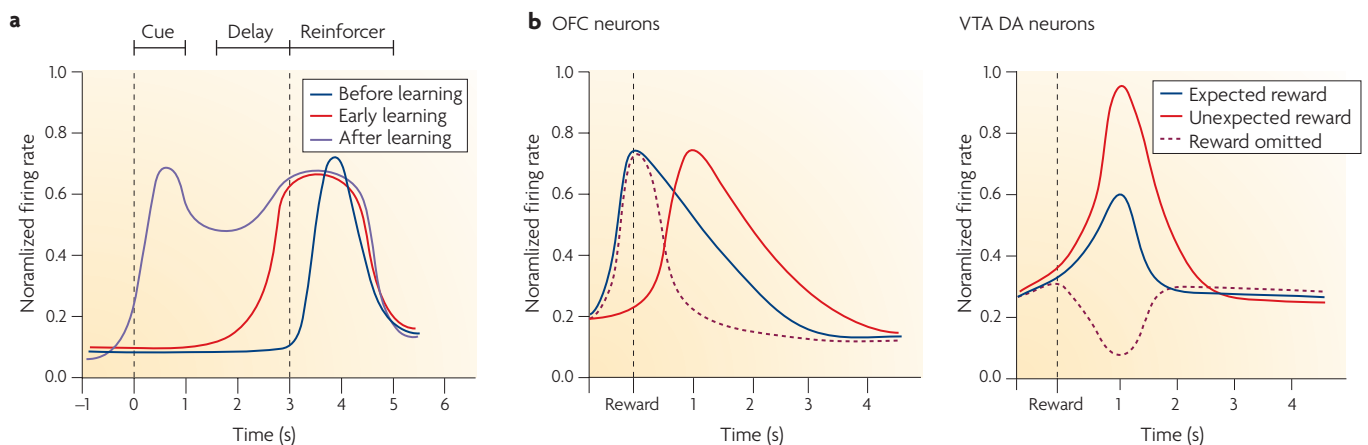


Figure 2 | Stylized plots of reward-responsive neurons in the orbitofrontal cortex. a | In discrimination learning tasks, neurons in the rat orbitofrontal cortex (OFC) initially fire in response to either the rewarding reinforcer (for example, sucrose) or the aversive reinforcer (for example, quinine) (dark blue line). After a number of trials the neurons also start to fire in anticipation of the reinforcer (red line). Finally, they also come to fire in response to cues that predict the reinforcer (purple line). **b** | These OFC

neurons do not stop firing in response to the reward, even after many trials. Moreover, their firing is not stronger in response to an unexpected reward or weaker in response to the omission of a reward (left panel). In this respect, OFC neurons are unlike dopamine (DA) neurons in the ventral tegmental area (VTA), which are also reward responsive but which fire more strongly in response to unexpected rewards and decrease firing when an expected reward is not delivered (right panel). Based on data from REFS 36–43.

after learning⁴², nor do they show stronger (or weaker) activity when reward is delivered (or omitted) unexpectedly⁴³ (FIG. 2b). Therefore, their activity is not well described as a prediction error signal (see also below); rather, activation of these neurons during progressively earlier periods in the trial could be considered as a representation of the expected outcome.

Similar findings have been reported in monkeys and humans^{44–50}. For example, Schultz and colleagues showed that OFC neurons exhibit outcome-expectant activity during a visual delayed response task⁵⁰. The strength of anticipatory firing depended on which particular food reward the monkey expected to receive: it was stronger when preferred foods were expected. Thus, anticipatory activity reflects both the physical and/or sensory attributes of the expected food reward and the subjective value that the monkey places on the item.

“ how does signalling of expected outcomes account for the role of the OFC in modifying behaviour when contingencies have changed? One possibility is that these signals are also crucial for ... calculating the teaching signals that are thought to drive learning ”

Anticipatory firing has also been observed in other areas^{34,36,44,47,51–53}, but it occurs first in the OFC^{36,44}. Furthermore, in other areas — the basolateral amygdala in particular — it is reduced to near chance levels by (even unilateral) lesions of the OFC⁵⁴. These results suggest that, to some extent, OFC neurons construct predictions on the basis of afferent input rather than simply passing on information received from other brain areas. Indeed, neural activity in the OFC of an animal performing a task is striking in that it typically anticipates trial events rather than being triggered by them⁵⁵. Each predictable event holds value because it provides information about the probability or proximity of reward. Thus, signalling expected outcomes could be considered a more general property of the OFC if ‘outcomes’ are broadly construed to include not only primary rewarding or punishing events but also events that have acquired value or meaning because they provide information about the likelihood of receiving such primary reinforcers in the future.

In this regard, such signalling would be conceptually similar to the ‘state value’ defined by temporal difference reinforcement learning models⁵⁶, which reflects the value of all current and future rewards predicted by cues and circumstances that are present in the environment. Indeed, even the neural response to an actual reward might reflect both the reward’s current hedonic properties and its future nutritive (or other) consequences (for example, a reward can also serve as a sensory cue that predicts a subsequent reward).

Evidence from behavioural studies. Of course, the neural firing patterns described above are only correlates of behaviour. If the OFC is crucial for signalling outcome expectancies, then manipulations that disrupt or alter the output from this brain area should preferentially disrupt behaviours that depend on information about outcomes. This turns out to be an excellent description of most OFC-dependent behaviours. For example, as rats learn in the odour discrimination task described above, they begin to respond faster (that is, approach the well faster) after sampling a positive predictive odour (that is, when they expect sucrose) than after sampling a negative predictive odour (that is, when they expect quinine)¹⁶. This behavioural differentiation emerges at the time that neurons in the OFC begin to show differential firing in anticipation of sucrose or quinine³⁶ and, unlike changes in choice performance (the decision to go or not to go to the well), changes in the latency of the behavioural response are abolished by OFC lesions¹⁶.

A second and perhaps more conclusive example of the crucial role of the OFC in using information about expected outcomes comes from studies using Pavlovian reinforcer devaluation. In these studies, an animal is trained to associate a cue with a particular reward. Subsequently, the value of the reward is reduced by pairing it with illness or by feeding it to satiety. The animals’ ability to access and use that new value to guide the learned responding is then assessed by presenting the cue alone. Animals normally show a reduced response to the predictive cue, reflecting their ability to access and use cue-evoked representations of the reward and its current value. However, monkeys and rats with OFC lesions fail to show this effect of devaluation^{12,22,57}. The deficit caused by OFC damage is evident even if the OFC is intact during the initial training and devaluation^{24,25}; therefore, the OFC is necessary for mobilizing information about the value of

the expected outcome to guide or influence responding. This observation is consistent with data from rats, monkeys and humans showing that neural activity in the OFC changes as a result of satiation^{45,58–60}. The role of the OFC in encoding and using information about expected outcomes in this setting differs from that of the amygdala or the mediodorsal thalamus, which are necessary during the earlier phases of the reinforcer devaluation task^{25,61–64}.

A third example comes from Pavlovian-to-instrumental transfer experiments. These involve a task in which animals are trained to independently associate a cue and a behaviour (for instance, pressing a lever) with a reward. Subsequently, they will increase this behaviour in the presence of the cue. This increased responding is thought to reflect the effect of Pavlovian information, triggered by cue presentation, on the behaviour (that is, lever pressing), which is controlled by instrumental representations. Consistent with the proposal that the OFC is crucial for signalling outcome information triggered by cues, Balleine and colleagues reported that outcome-specific (but not general) Pavlovian-to-instrumental transfer is abolished by OFC lesions⁶⁵.

These examples indicate that one fundamental role of neural activity in the OFC is to drive behaviour on the basis of the specific features of expected outcomes. Other examples of OFC-dependent behaviours corroborate this account. These include delayed discounting^{66–68}, conditioned reinforcement and other second-order behaviours^{26,69–71}, Pavlovian approach responses¹³, the enhancement of discriminative responding by different outcomes⁷², and even cognitive and affective processes, such as regret and counterfactual reasoning⁷³. In each of these examples, normal performance can be plausibly argued to require the ability to signal, in real-time, information about the characteristics and values of outcomes predicted by cues and circumstances in the environment.

Of note, the brain areas that mediate these types of value-guided behaviour are also typically important for reversal learning^{12,16,22,65}. One exception to this general rule is described in a recent study by Bachevalier and colleagues that showed that reversal learning was unimpaired in monkeys with lesions of areas 11 and 13 (REF. 27); the same monkeys were previously reported to exhibit deficits in behavioural changes after reinforcer devaluation. Although this report indicates that reversal learning can be mediated by OFC-independent mechanisms, it does not necessarily contradict earlier

findings more consistent with colocalization of reversal learning and devaluation in the OFC. OFC-dependent reversal deficits are transient and disappear quickly with practice^{74,75}. Indeed, we can completely eliminate OFC-dependent reversal deficits by damaging other brain areas, whereas deficits in behaviours (latencies to respond) that are probably crucially dependent on value representations persist^{35,76}. This is relevant because these monkeys were tested ~18 months after surgery and had undergone extensive training in various tasks before reversal testing²⁷. In any event, although this report²⁷ raises important questions about reversal learning and about whether this process depends on different brain regions, it does not necessarily contradict the basic idea presented in the next section, that adaptive behaviour in general — epitomized by reversal deficits — can be supported by the same information as that used in devaluation or in other value-guided tasks; it merely confirms that the behaviour does not have to be, which is a topic that we address later.

Expectancies and flexible behaviour

The OFC-dependent behaviours described above differ from flexible behaviour that can be assessed by reversal learning or related tasks because they generally do not involve changes in established associations or response contingencies. For example, OFC inactivation impairs changes in conditioned responding after devaluation, even though nothing about the underlying associations has changed and only the value of the outcome has been altered. Similarly, OFC damage disrupts Pavlovian-to-instrumental transfer even when lesions are made after rats have learned the underlying cue–reward and response–reward associations. These data suggest that signals from the OFC that relate to expected outcomes influence judgment and decision making when contingencies are stable.

But how does signalling of expected outcomes account for the role of the OFC in modifying behaviour when contingencies have changed? One possibility is that these signals are also crucial for updating associative representations in other brain regions in the face of unexpected outcomes. The OFC might play this part by contributing information that is necessary for calculating the teaching signals that are thought to drive learning in such situations⁵⁶. These teaching signals (also known as prediction errors) result from differences between the value of the expected outcome and that of the actual outcome. There is strong evidence that these

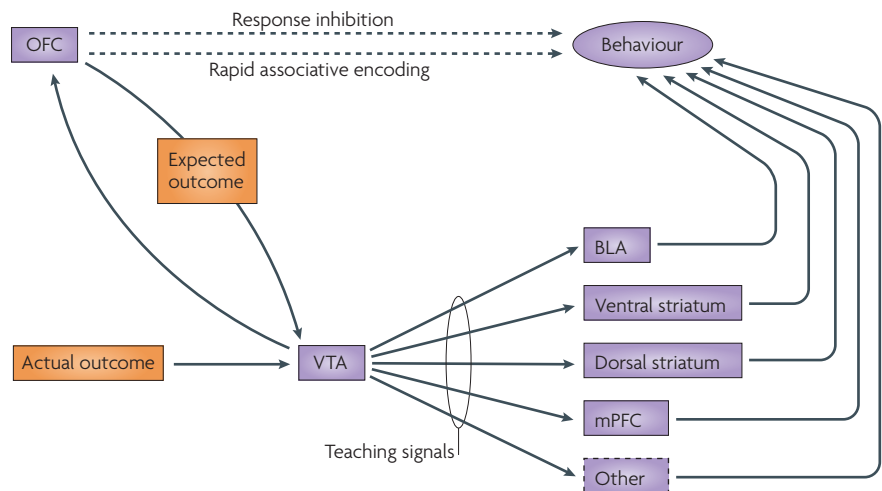


Figure 3 | The proposed role of the orbitofrontal cortex in flexible, adaptive behaviour. A summary of several proposals reviewed in the text for explaining the role of the orbitofrontal cortex (OFC) in facilitating reversal learning. At the top is the historical proposal that output from the OFC directly inhibits ‘pre-potent’ responses and is therefore crucial for reversal learning. Below is the more recent idea that the OFC functions as a highly flexible associative reference table to directly guide correct responding and so is vital for reversal learning. At the bottom is our alternative proposal that the OFC is crucial for reversal learning because it drives associative learning in other structures. According to this proposal, the OFC does this by supporting the generation of teaching signals (for example, by the ventral tegmental area (VTA)) when actual outcomes do not match signals from the OFC regarding expected outcomes. BLA, basolateral amygdala; mPFC, medial prefrontal cortex.

prediction errors are signalled by phasic activity in dopamine neurons in the substantia nigra and ventral tegmental area^{37–41}, and activity consistent with error signalling has also been observed in other brain areas^{77–80}. Although OFC neurons do not signal prediction errors themselves⁴³, the OFC might still contribute to the signalling of these prediction errors by providing a source of information (although not the only source; see below) about the expected outcome value; this information could then be used to compute prediction errors (FIG. 3). The loss of this function would explain why OFC damage disrupts behavioural flexibility under conditions of changing contingencies.

Predictions and implications. Our hypothesis makes a couple of testable predictions. The first is that changes in associative representations in downstream regions, such as the amygdala, should be at least partly dependent on the OFC. Consistent with this, associative encoding in the basolateral amygdala is markedly less flexible in rats with lesions of the OFC⁵⁴. Moreover, lesions or inactivation of the basolateral amygdala abolish the reversal deficit caused by OFC lesions³⁵. This result suggests that this downstream (amygdala) inflexibility underlies the deficit in reversal behaviour that results

from OFC lesions. As recoding of associations in the OFC occurs more slowly than does recoding in the basolateral amygdala³², these results cannot be easily explained by the ‘rapidly flexible encoding’ hypothesis of OFC function. However, they are fully consistent with the proposal that signalling of old associations by OFC neurons drives reversal learning by supporting error signalling in some other region (for example, the ventral tegmental area; see below). These error signals would then facilitate the encoding of new associations by downstream areas such as the amygdala. The better the OFC signals the old associations, the stronger the teaching signals and the faster the recoding process. This proposal also fits with the finding, discussed earlier, that better reversal behaviour occurs when cue selectivity in the OFC fails to reverse³³.

The abolition of the OFC-dependent reversal deficit by lesions or inactivation of the basolateral amygdala also implies that the OFC–amygdala system is not actually required for normal reversal learning. This somewhat unconventional statement makes sense when one considers that reversal learning probably involves changes in encoding of various types of associative information maintained in different brain regions; the speed of reversal learning will therefore depend in part on the task and the type of

associative strategy it emphasizes, as well as on the rate at which information in different brain regions is updated (with the 'slowest' regions determining the speed of reversal learning). Normally, the basolateral amygdala changes its representations very rapidly^{32,34}, and so is probably not the rate-limiting step. Consistent with this, damage to the basolateral amygdala does not affect reversal learning, at least not in the sorts of task often used to assess OFC-dependent reversal learning^{16,27,81}. However, when the OFC is lesioned, representations in the basolateral amygdala are slow to change, and this area becomes an impediment to the speed of reversal learning. In other words, damage to the OFC introduces an artificial bottleneck that slows learning. When the basolateral amygdala is damaged or inactivated in an OFC-lesioned

rat, reversal might go back to being controlled by whatever area normally governs its rate.

A second prediction of our hypothesis is that the OFC should be important for learning in other situations, when error signals are crucially dependent on information about expected outcomes. Results from a Pavlovian over-expectation task are consistent with this prediction. In this task, rats initially learn that several Pavlovian cues are each independent predictors of a reward. Then, two of these cues are presented together (compound training), followed by the same reward. When either of these cues is subsequently presented alone, rats exhibit a spontaneous reduction in their conditioned response. This reduced response is thought to result from negative prediction errors induced by violation of the animal's 'summed' expectation for the reward. Reversible inactivation of the OFC during compound training prevents summation and the later reduction in response to the individual cues⁴³. This result cannot be explained as a simple deficit in using associative information encoded by the OFC, because the OFC is fully functional when the use of new information is assessed. Instead, it suggests that the OFC is crucial for the prediction-error-induced learning. Notably, disconnecting the OFC from the ventral tegmental area during compound training in the over-expectation task also abolishes the response reduction⁴³. This indicates that the role of the OFC in supporting learning in this setting requires interaction with dopamine neurons in the ventral tegmental area.

Of course, the OFC is not required for all learning; clearly, other brain areas can acquire new representations without input from the OFC. Indeed, much learning — including initial discrimination learning and simple Pavlovian learning — is unaffected by OFC damage²². One possible explanation for this preserved learning is that these situations, which involve new learning, do not require expectancies for the production of error signals. If anything, expectancies would actually impede learning by reducing the size of the teaching signal as learning progressed. So it is not surprising that the OFC is not necessary for all learning.

Moreover, it is likely that the OFC is not the only source of information about expected outcomes. Below, we suggest that this role of the OFC in driving learning may be constrained by the circuits in which it is embedded and by the particular role that these circuits have in particular categories of associative learning. As we describe, outside these domains learning may not depend on signalling from the OFC.

Conclusions and future directions

Recent evidence contradicts long-standing ideas that the role of the OFC in adapting behaviour in the face of changing contingencies and unexpected outcomes reflects response inhibition or rapid flexibility of associative encoding. Instead, we suggest that it reflects contributions of the OFC to changing associative representations in other brain regions, mediated indirectly through the support of prediction-error signalling by systems such as the midbrain dopamine system. This is consistent with evidence that the OFC is crucial for signalling outcome expectancies.

This proposal is not unlike ideas about the role of the ventral striatum in actor-critic models of reinforcement learning⁸². According to these theories, the ventral striatum serves as the critic, providing information necessary to compute a 'state value', which is, in turn, required by downstream areas such as the midbrain to calculate reward prediction errors. These errors serve as teaching signals to facilitate the acquisition of rules or policies in other regions — such as the dorsolateral striatum⁸³ — that serve as the actor to drive behaviour. This idea of state value is comparable to the outcome expectancies that we suggest are signalled by the OFC. Thus, another way to view our proposal is that the OFC functions as a critic.

Importantly, we do not argue that the OFC supplants the ventral striatum in this regard. Rather, we suggest that the OFC might act in parallel or in series with the ventral striatum and perhaps other regions, such as the amygdala or the medial prefrontal cortex, in signalling this sort of information. Similarly to the fact that there are different systems for different kinds of memory, there are likely to be multiple critics, each providing a particular type of information relevant to computing state value.

For example, it is now clear that even simple Pavlovian conditioning results in the formation of multiple associations in the brain, each capturing different aspects of the cue–outcome pairing⁸⁴. As a result, cues become able to evoke a representation of the outcome, including its particular sensory features and its unique, changing value. In addition, cues become linked to the general emotion associated with the outcome, which that outcome shares with other, similar outcomes. Analogous distinctions have been drawn for goal-directed versus habitual behaviour in instrumental learning⁸⁵.

These different representations support different types of behaviour and depend on

Glossary

Conditioned reinforcement

The process by which a Pavlovian-acquired value can reinforce instrumental action.

Delayed discounting task

A task that assesses the effect of delaying a reward on a subject's choice behaviour.

Detour reaching task

A task that assesses the ability of a subject to learn to reach around a clear barrier in order to retrieve a reward.

Instrumental

Related to a training situation in which associations are arranged between actions and outcomes.

Pavlovian

Related to a training situation in which associations are arranged between cues and outcomes.

Pavlovian approach responses

Approaches to a conditioned stimulus.

Reinforcer devaluation task

A task that assesses the ability of a subject to modify conditioned responding after the expected reward is devalued, typically by pairing it with illness or by feeding it to satiety.

Stop signal task

A task that assesses the ability of a subject to stop or inhibit an ongoing reward-seeking behaviour in response to a 'stop' signal.

Temporal difference reinforcement learning models

A family of non-trial-based reinforcement learning models in which the difference between the expected and actual values of a particular state in a sequence of behaviours is used as a teaching signal to facilitate the acquisition of associative rules or policies to direct future behaviour.

Visual delayed response task

A task that assesses the ability of a subject to respond to visual cues to obtain a reward that is presented after a short delay.

somewhat non-overlapping neural circuits. As described above, signalling of information about specific outcomes depends on a circuit that includes the OFC and the basolateral amygdala. By contrast, signalling of general affective information depends on other areas, such as the nucleus accumbens in the ventral striatum⁸⁶.

We suggest that each of these regions might function as a relatively independent critic, depending on what type of expectancy is being violated. Input from the OFC may be particularly important when the predicted outcome changes. This would be consistent with recent data showing that the OFC is necessary for learning when one outcome is substituted for another²⁶. On the other hand, input from regions that signal general affective information might be crucial for learning when the general affective value of the outcome, rather than its particular features, are altered — for example, when an animal's internal motivational state changes after learning.

Similarly, the OFC might be particularly important in signalling Pavlovian information, providing the crucial predictions and playing a lesser or perhaps inconsequential part in signalling information about expected outcomes derived from instrumental or action–outcome associations. This proposal has been advanced by Balleine and colleagues on the basis of observations that OFC damage affects changes in Pavlovian responding but not changes in instrumental responding after reinforcer devaluation⁶⁵. Although neurophysiological and imaging results show action–outcome correlates in neural activity in the OFC, these could reflect sensory information that differs between different responses^{87–89}. If the OFC has a crucial role in signalling only cue-evoked or Pavlovian information, then some other area would have to provide instrumental information. Leading candidates would be parts of the medial prefrontal cortex and the dorsomedial striatum, which, Balleine and colleagues have shown, are crucial for goal-directed behaviour^{90,91}.

Inputs from multiple critics might converge on midbrain areas through direct or indirect connections. For example, the ventral striatum receives input from the OFC and the medial prefrontal cortex; thus, the ventral striatum could function as a super-critic, integrating information and sending it off to be used in error signalling. Alternatively, this information could reach the midbrain through direct projections. Direct and indirect pathways might both be used, with each influencing learning in

subtly different ways. Of course, it is also possible — indeed likely — that different types of information interact with different error-signalling regions. There are several regions (other than the midbrain dopamine neurons) that seem to signal prediction errors.

Importantly, proposals regarding the differing contributions of the OFC and other brain areas to error encoding are testable. Functional MRI and single-unit studies combined with lesions or inactivation and with behavioural tasks that manipulate expectancies based on the specific and general values of actions and Pavlovian cues can directly confirm or invalidate our hypothesis. Although substantial experimental evidence supports the ideas laid out in this Perspective, further work holds the potential to sketch out this circuit and its functions in more detail, hopefully leading to the creation of a truly useful framework for how brain circuits implement the simple associative-learning processes that help us to navigate our ever-changing world.

Geoffrey Schoenbaum, Thomas A. Stalnaker and Yuji K. Takahashi are at the Department of Anatomy and Neurobiology, University of Maryland School of Medicine, Baltimore, Maryland 21201, USA.

Matthew R. Roesch is at the Department of Psychology, Program in Neuroscience and Cognitive Science, University of Maryland, College Park, Maryland 20742, USA.

Geoffrey Schoenbaum is also at the Department of Psychiatry, University of Maryland School of Medicine, Baltimore, Maryland 21201, USA, and at the Department of Psychology, University of Maryland Baltimore County, Baltimore, Maryland 21228, USA.

Correspondence to G.S.
e-mail: schoenbg@schoenbaumlab.org

doi:10.1038/nrn2753

Published online 11 November 2009

- Harlow, J. M. Recovery after passage of an iron bar through the head. *Publ. Massachusetts Med. Soc.* **2**, 329–346 (1868).
- Damasio, H., Grabowski, T., Frank, R., Galaburda, A. M. & Damasio, A. R. The return of Phineas Gage: clues about the brain from the skull of a famous patient. *Science* **264**, 1102–1105 (1994).
- Schoenbaum, G. & Roesch, M. R. Orbitofrontal cortex, associative learning, and expectancies. *Neuron* **47**, 633–636 (2005).
- Preuss, T. M. Do rats have prefrontal cortex? The Rose-Woolsey-Akert program reconsidered. *J. Comp. Neurol.* **7**, 1–24 (1995).
- Jones, B. & Mishkin, M. Limbic lesions and the problem of stimulus–reinforcement associations. *Exp. Neurol.* **36**, 362–377 (1972).
- Butter, C. M. Perseveration and extinction in discrimination reversal tasks following selective frontal ablations in *Macaca mulatta*. *Physiol. Behav.* **4**, 163–171 (1969).
- Rolls, E. T., Hornak, J., Wade, D. & McGrath, J. Emotion-related learning in patients with social and emotional changes associated with frontal lobe damage. *J. Neurol. Neurosurg. Psychiatry* **57**, 1518–1524 (1994).
- Meunier, M., Bachevalier, J. & Mishkin, M. Effects of orbital frontal and anterior cingulate lesions on object and spatial memory in rhesus monkeys. *Neuropsychologia* **35**, 999–1015 (1997).
- McAlonan, K. & Brown, V. J. Orbital prefrontal cortex mediates reversal learning and not attentional set shifting in the rat. *Behav. Brain Res.* **146**, 97–130 (2003).
- Bissonette, G. B. *et al.* Double dissociation of the effects of medial and orbital prefrontal cortical lesions on attentional and affective shifts in mice. *J. Neurosci.* **28**, 11124–11130 (2008).
- Bechara, A., Damasio, H., Tranel, D. & Damasio, A. R. Deciding advantageously before knowing the advantageous strategy. *Science* **275**, 1293–1294 (1997).
- Izquierdo, A. D., Suda, R. K. & Murray, E. A. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *J. Neurosci.* **24**, 7540–7548 (2004).
- Chudasama, Y. & Robbins, T. W. Dissociable contributions of the orbitofrontal and infralimbic cortex to pavlovian autoshaping and discrimination reversal learning: further evidence for the functional heterogeneity of the rodent frontal cortex. *J. Neurosci.* **23**, 8771–8780 (2003).
- Fellows, L. K. & Farah, M. J. Ventromedial frontal cortex mediates affective shifting in humans: evidence from a reversal learning paradigm. *Brain* **126**, 1830–1837 (2003).
- Hornak, J. *et al.* Reward-related reversal learning after surgical excisions in orbito-frontal or dorsolateral prefrontal cortex in humans. *J. Cogn. Neurosci.* **16**, 463–478 (2004).
- Schoenbaum, G., Setlow, B., Nugent, S. L., Saddoris, M. P. & Gallagher, M. Lesions of orbitofrontal cortex and basolateral amygdala complex disrupt acquisition of odor-guided discriminations and reversals. *Learn. Mem.* **10**, 129–140 (2003).
- Ferrier, D. *The Functions of the Brain* (G. P. Putnam's Sons, New York, 1876).
- Fuster, J. M. *The Prefrontal Cortex* (Lippincott-Raven, New York, 1997).
- Damasio, A. R. *Descartes' Error* (Putnam, New York, 1994).
- Eagle, D. M. *et al.* Stop-signal reaction-time task performance: role of prefrontal cortex and subthalamic nucleus. *Cereb. Cortex* **18**, 178–188 (2008).
- Wallis, J. D., Dias, R., Robbins, T. W. & Roberts, A. C. Dissociable contributions of the orbitofrontal and lateral prefrontal cortex of the marmoset to performance on a detour reaching task. *Eur. J. Neurosci.* **13**, 1797–1808 (2001).
- Gallagher, M., McMahan, R. W. & Schoenbaum, G. Orbitofrontal cortex and representation of incentive value in associative learning. *J. Neurosci.* **19**, 6610–6614 (1999).
- Baxter, M. G., Parker, A., Lindner, C. C. C., Izquierdo, A. D. & Murray, E. A. Control of response selection by reinforcer value requires interaction of amygdala and orbitofrontal cortex. *J. Neurosci.* **20**, 4311–4319 (2000).
- Pickens, C. L., Saddoris, M. P., Gallagher, M. & Holland, P. C. Orbitofrontal lesions impair use of cue–outcome associations in a devaluation task. *Behav. Neurosci.* **119**, 317–322 (2005).
- Pickens, C. L. *et al.* Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. *J. Neurosci.* **23**, 11078–11084 (2003).
- Burke, K. A., Franz, T. M., Miller, D. N. & Schoenbaum, G. The role of orbitofrontal cortex in the pursuit of happiness and more specific rewards. *Nature* **454**, 340–344 (2008).
- Kazama, A. & Bachevalier, J. Selective aspiration or neurotoxic lesions of orbital frontal areas 11 and 13 spared monkeys' performance on the object discrimination reversal task. *J. Neurosci.* **29**, 2794–2804 (2009).
- Chudasama, Y., Kralik, J. D. & Murray, E. A. Rhesus monkeys with orbital prefrontal cortex lesions can learn to inhibit prepotent responses in the reversed reward contingency task. *Cereb. Cortex* **17**, 1154–1159 (2007).
- Rolls, E. T. The orbitofrontal cortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **351**, 1433–1443 (1996).
- Thorpe, S. J., Rolls, E. T. & Maddison, S. The orbitofrontal cortex: neuronal activity in the behaving monkey. *Exp. Brain Res.* **49**, 93–115 (1983).
- Rolls, E. T., Critchley, H. D., Mason, R. & Wakeman, E. A. Orbitofrontal cortex neurons: role in olfactory and visual association learning. *J. Neurophysiol.* **75**, 1970–1981 (1996).

32. Schoenbaum, G., Chiba, A. A. & Gallagher, M. Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J. Neurosci.* **19**, 1876–1884 (1999).
33. Stalnaker, T. A., Roesch, M. R., Franz, T. M., Burke, K. A. & Schoenbaum, G. Abnormal associative encoding in orbitofrontal neurons in cocaine-experienced rats during decision-making. *Eur. J. Neurosci.* **24**, 2643–2653 (2006).
34. Patton, J. J., Belova, M. A., Morrison, S. E. & Salzman, C. D. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* **439**, 865–870 (2006).
35. Stalnaker, T. A., Franz, T. M., Singh, T. & Schoenbaum, G. Basolateral amygdala lesions abolish orbitofrontal-dependent reversal impairments. *Neuron* **54**, 51–58 (2007).
36. Schoenbaum, G., Chiba, A. A. & Gallagher, M. Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature Neurosci.* **1**, 155–159 (1998).
37. Hollerman, J. R. & Schultz, W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neurosci.* **1**, 304–309 (1998).
38. Roesch, M. R., Calu, D. J. & Schoenbaum, G. Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nature Neurosci.* **10**, 1615–1624 (2007).
39. Pan, W.-X., Schmidt, R., Wickens, J. R. & Hyland, B. I. Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network. *J. Neurosci.* **25**, 6235–6242 (2005).
40. Waelti, P., Dickinson, A. & Schultz, W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature* **412**, 43–48 (2001).
41. Bayer, H. M. & Glimcher, P. W. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* **47**, 129–141 (2005).
42. Schoenbaum, G., Setlow, B., Saddoris, M. P. & Gallagher, M. Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. *Neuron* **39**, 855–867 (2003).
43. Takahashi, Y. *et al.* The orbitofrontal cortex and ventral tegmental area are necessary for learning from unexpected outcomes. *Neuron* **62**, 269–280 (2009).
44. Wallis, J. D. & Miller, E. K. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur. J. Neurosci.* **18**, 2069–2081 (2003).
45. Gottfried, J. A., O'Doherty, J. & Dolan, R. J. Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science* **301**, 1104–1107 (2003).
46. O'Doherty, J., Deichmann, R., Critchley, H. D. & Dolan, R. J. Neural responses during anticipation of a primary taste reward. *Neuron* **33**, 815–826 (2002).
47. Hikosaka, K. & Watanabe, M. Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. *Cereb. Cortex* **10**, 263–271 (2000).
48. Hikosaka, K. & Watanabe, M. Long- and short-range reward expectancy in the primate orbitofrontal cortex. *Eur. J. Neurosci.* **19**, 1046–1054 (2004).
49. Padoa-Schioppa, C. & Assad, J. A. Neurons in orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).
50. Tremblay, L. & Schultz, W. Relative reward preference in primate orbitofrontal cortex. *Nature* **398**, 704–708 (1999).
51. Sugase-Miyamoto, Y. & Richmond, B. J. Neuronal signals in the monkey basolateral amygdala during reward schedules. *J. Neurosci.* **25**, 11071–11083 (2005).
52. Watanabe, M. Reward expectancy in primate prefrontal neurons. *Nature* **382**, 629–632 (1996).
53. Shidara, M. & Richmond, B. J. Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* **296**, 1709–1711 (2002).
54. Saddoris, M. P., Gallagher, M. & Schoenbaum, G. Rapid associative encoding in basolateral amygdala depends on connections with orbitofrontal cortex. *Neuron* **46**, 321–331 (2005).
55. Schoenbaum, G. & Eichenbaum, H. Information coding in the rodent prefrontal cortex. I. Single-neuron activity in orbitofrontal cortex compared with that in pyriform cortex. *J. Neurophysiol.* **74**, 733–750 (1995).
56. Niv, Y. & Schoenbaum, G. Dialogues on prediction errors. *Trends Cogn. Sci.* **12**, 265–272 (2008).
57. Machado, C. J. & Bachevalier, J. The effects of selective amygdala, orbital frontal cortex or hippocampal formation lesions on reward assessment in nonhuman primates. *Eur. J. Neurosci.* **25**, 2885–2904 (2007).
58. O'Doherty, J. *et al.* Sensory-specific satiety-related olfactory activation of the human orbitofrontal cortex. *Neuroreport* **11**, 893–897 (2000).
59. Critchley, H. D. & Rolls, E. T. Hunger and satiety modify the responses of olfactory and visual neurons in the primate orbitofrontal cortex. *J. Neurophysiol.* **75**, 1675–1686 (1996).
60. de Araujo, I. E. *et al.* Neural ensemble coding of satiety states. *Neuron* **51**, 483–494 (2006).
61. Hatfield, T., Han, J. S., Conley, M., Gallagher, M. & Holland, P. Neurotoxic lesions of basolateral, but not central, amygdala interfere with Pavlovian second-order conditioning and reinforcer devaluation effects. *J. Neurosci.* **16**, 5256–5265 (1996).
62. Malkova, L., Gaffan, D. & Murray, E. A. Excitotoxic lesions of the amygdala fail to produce impairment in visual learning for auditory secondary reinforcement but interfere with reinforcer devaluation effects in rhesus monkeys. *J. Neurosci.* **17**, 6011–6020 (1997).
63. Wellmann, L. L., Gale, K. & Malkova, L. GABA_A-mediated inhibition of basolateral amygdala blocks reward devaluation in macaques. *J. Neurosci.* **25**, 4577–4586 (2005).
64. Mitchell, A. S., Browning, P. G. & Baxter, M. G. Neurotoxic lesions of the medial mediodorsal nucleus of the thalamus disrupt reinforcer devaluation effects in rhesus monkeys. *J. Neurosci.* **27**, 11289–11295 (2007).
65. Ostlund, S. B. & Balleine, B. W. Orbitofrontal cortex mediates outcome encoding in Pavlovian but not instrumental learning. *J. Neurosci.* **27**, 4819–4825 (2007).
66. Kheramin, S. *et al.* Role of the orbital prefrontal cortex in choice between delayed and uncertain reinforcers: a quantitative analysis. *Behav. Process.* **64**, 239–250 (2003).
67. Winstanley, C. A., Theobald, D. E. H., Cardinal, R. N. & Robbins, T. W. Contrasting roles of basolateral amygdala and orbitofrontal cortex in impulsive choice. *J. Neurosci.* **24**, 4718–4722 (2004).
68. Mobini, S. *et al.* Effects of lesions of the orbitofrontal cortex on sensitivity to delayed and probabilistic reinforcement. *Psychopharmacology* **160**, 290–298 (2002).
69. Cousins, G. A. & Otto, T. Neural substrates of olfactory discrimination learning with auditory secondary reinforcement. I. Contributions of the basolateral amygdaloid complex and orbitofrontal cortex. *Integr. Physiol. Behav. Sci.* **38**, 272–294 (2003).
70. Hutcheson, D. M. & Everitt, B. J. The effects of selective orbitofrontal cortex lesions on the acquisition and performance of cue-controlled cocaine seeking in rats. *Ann. NY Acad. Sci.* **1003**, 410–411 (2003).
71. Pears, A., Parkinson, J. A., Hopewell, L., Everitt, B. J. & Roberts, A. C. Lesions of the orbitofrontal but not medial prefrontal cortex disrupt conditioned reinforcement in primates. *J. Neurosci.* **23**, 11189–11201 (2003).
72. McDannald, M. A., Saddoris, M. P., Gallagher, M. & Holland, P. C. Lesions of orbitofrontal cortex impair rats' differential outcome expectancy learning but not conditioned stimulus-potentiated feeding. *J. Neurosci.* **25**, 4626–4632 (2005).
73. Camille, N. *et al.* The involvement of the orbitofrontal cortex in the experience of regret. *Science* **304**, 1168–1170 (2004).
74. Dias, R., Robbins, T. W. & Roberts, A. C. Dissociable forms of inhibitory control within prefrontal cortex with an analog of the Wisconsin card sort test: restriction to novel situations and independence from "on-line" processing. *J. Neurosci.* **17**, 9285–9297 (1997).
75. Schoenbaum, G., Nugent, S., Saddoris, M. P. & Setlow, B. Orbitofrontal lesions in rats impair reversal but not acquisition of go, no-go odor discriminations. *Neuroreport* **13**, 885–890 (2002).
76. Sage, J. R. & Knowlton, B. J. Effects of US devaluation on win-stay and win-shift radial maze performance in rats. *Behav. Neurosci.* **114**, 295–306 (2000).
77. Matsumoto, M. & Hikosaka, K. Lateral habenula as a source of negative reward signals in dopamine neurons. *Nature* **447**, 1111–1115 (2007).
78. Daw, N. D., Kakade, S. & Dayan, P. Opponent interactions between serotonin and dopamine. *Neural Netw.* **15**, 603–616 (2002).
79. Belova, M. A., Paton, J. J., Morrison, S. E. & Salzman, C. D. Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron* **55**, 970–984 (2007).
80. Matsumoto, M., Matsumoto, K., Abe, H. & Tanaka, K. Medial prefrontal cell activity signaling prediction errors of action values. *Nature Neurosci.* **10**, 647–656 (2007).
81. Izquierdo, A. D. & Murray, E. A. Selective bilateral amygdala lesions in rhesus monkeys fail to disrupt object reversal learning. *J. Neurosci.* **27**, 1054–1062 (2007).
82. O'Doherty, J. *et al.* Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* **304**, 452–454 (2004).
83. Yin, H. H., Knowlton, B. J. & Balleine, B. W. Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *Eur. J. Neurosci.* **19**, 181–189 (2004).
84. Rescorla, R. A. Pavlovian conditioning: it's not what you think it is. *Am. Psychol.* **43**, 151–160 (1988).
85. Dickinson, A. & Balleine, B. W. Motivational control of goal-directed action. *Anim. Learn. Behav.* **22**, 1–18 (1994).
86. Balleine, B. W. & Corbit, L. H. Double dissociation of nucleus accumbens core and shell on the general and outcome-specific forms of pavlovian-instrumental transfer. *Soc. Neurosci. Abstr.* **71**, 16 (Washington DC, 12–16 November 2005).
87. Valentin, V. V., Dickinson, A. & O'Doherty, J. P. Determining the neural substrates of goal-directed learning in the human brain. *J. Neurosci.* **27**, 4019–4026 (2007).
88. Roesch, M. R., Taylor, A. R. & Schoenbaum, G. Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron* **51**, 509–520 (2006).
89. Feierstein, C. E., Quirk, M. C., Uchida, N., Sosulski, D. L. & Mainen, Z. F. Representation of spatial goals in rat orbitofrontal cortex. *Neuron* **51**, 495–507 (2006).
90. Corbit, L. H. & Balleine, B. W. The role of prefrontal cortex in instrumental conditioning. *Behav. Brain Res.* **146**, 145–157 (2003).
91. Yin, H. H., Knowlton, B. J. & Balleine, B. W. Blockade of NMDA receptors in the dorsomedial striatum prevents action-outcome learning in instrumental conditioning. *Eur. J. Neurosci.* **22**, 505–512 (2005).

Acknowledgements

The authors would like to acknowledge support from the National Institute on Drug Abuse, the National Institute of Mental Health and the National Institute on Aging.

FURTHER INFORMATION

Geoffrey Schoenbaum's homepage:
<http://www.schoenbaumlab.org>

ALL LINKS ARE ACTIVE IN THE ONLINE PDF