Medial and lateral orbitofrontal cortices represent unique components of cognitive maps of task space.

Laura A. Bradfield^{1,2} and Genevra Hart³.

- 1. Centre for Neuroscience and Regenerative Medicine, University of Technology Sydney, NSW, 2007, Australia.**
- 2. St. Vincent's Centre for Applied Medical Research, St. Vincent's Hospital Sydney Limited, Sydney, NSW 2010, Australia.**
- 3. Decision Neuroscience Laboratory, School of Psychology, University of New South Wales, NSW 2052, Australia.**

Corresponding Author:

Dr. Laura Bradfield at laura.bradfield@uts.edu.au

**Centre for Neuroscience and Regenerative Medicine,
School of Life Sciences,
University of Technology Sydney (St. Vincent's Campus)
405 Liverpool St
Darlinghurst, NSW 2010.
PH: 02 83824950**

Acknowledgements: We thank Bernard Balleine, Melissa Sharpe, Amir Dezfouli, and Marios Panayi for helpful discussions and comments regarding this manuscript.

This work was supported by the National Health and Medical Research Council of Australia Grant (GNT 1148244) awarded to L.A.B.

Abstract

The orbitofrontal cortex (OFC) has recently been proposed to function as a cognitive map of task space: a mental model of the various steps involved in a task. This idea has proven popular because it provides a cohesive explanation for a number of disparate findings regarding the OFC's role in a broad array of tasks. Concurrently, mounting evidence has begun to reveal the functional heterogeneity of OFC subregions, particularly the medial and lateral OFC. How these subregions might uniquely contribute to the OFC's role as a cognitive map of task space, however, has not been explored. Here we propose that the lateral OFC represents the agent's initial position within that task map (i.e. initial state), determining which actions are available as a consequence of that position, whereas the medial OFC represents the agent's desired future position within the task map (i.e. terminal state), influencing which actions are selected to achieve that position. We argue that these processes are achieved somewhat independently and somewhat interdependently, and are achieved through similar but non-identical circuitry.

Introduction

The notion that goal-directed decision-making relies on a ‘cognitive map’ or ‘mental model’ of the different steps involved in a task has gained significant traction in recent years. A number of studies have begun to identify the putative neural structures and circuits supporting this function, with several lines of evidence suggesting the orbitofrontal cortex (OFC) plays a key role¹⁻⁴. Recent studies have also demonstrated a high degree of functional heterogeneity within the OFC⁵⁻¹⁰, yet how these functionally distinct OFC subregions might differentially contribute to its function as a cognitive map of task space has never been considered.

In this perspective, we argue that current findings are consistent with the lateral OFC inferring the organism’s *initial* position within a task and thus what actions are available for selection as a consequence (i.e. initial state), whereas the medial OFC provides information about what *future* position within task space is likely if the animal selects a particular action (i.e. the terminal state, represented by what goal will be achieved). To put it another way, the function of lateral OFC has been likened to the dropping of a google maps ‘pin’ to determine one’s position within the task¹¹. As a consequence of the pin being in a particular position, the lateral OFC can calculate the actions available (routes), and which of those actions is most desirable (e.g. the fastest route). Using this same analogy, we suggest that the medial OFC can also influence actions (routes), but that it does so by providing information about the desired final destination. In this way, the functions of lateral OFC and medial OFC appear to be highly related and even interdependent (it is difficult to determine which route is most desirable without knowing the final destination), however they are not interchangeable. Finally, we argue that both initial and terminal state information appears to be particularly OFC-dependent when it relies on partially unobservable information, because several studies have shown that actions can be selected accurately in spite of OFC inactivation (lateral and/or medial) when all of the necessary state information is observable.

The Role of States in Decision-Making

Within the laboratory, decision-making tends to be operationalized as the selection of goal-directed actions. Such actions are effortful, flexibly deployed, and must be motivated by both the value of the outcome or goal (goal criterion), and the contingency between the action and outcome (instrumental criterion)¹². Such actions are held in contrast to habitual actions that are elicited automatically without regard for the outcome. Both goal-directed and habitual processes are integral to everyday functioning, yet are neurally separable as shown by experiments that have demonstrated deficits in one system but not the other as a result of damage to distinct brain structures¹³⁻¹⁵.

Computationally, reinforcement learning (RL) models have posited that goal-directed actions are akin to model-based models¹⁶. This means that in order to exert goal-directed control, animals build a mental model (i.e. a cognitive map of task space) of the various steps that are involved in a task. Animals can then simulate their prospective journey through the steps of the model to estimate the probability of obtaining their goal. Habits, on the other hand, are thought to be more akin to model-free models¹⁶ (although this position is not without controversy¹⁷), because they do not rely on cognitive maps in the same manner. Instead, they are elicited in response to various stimuli in the environment. Model-free RL posits that habits become valuable over the course of learning by being paired with reinforcers, causing value to be 'cached' at the level of the action itself. Unlike model-based RL, altered reinforcer value cannot immediately modify the cached value of actions until that reinforcer is experienced as a consequence of the action taken. Further, model-based but not model-free RL recognises the identity of the outcome; model-free RL recognises only its value.

Cognitive task maps consist of 'states', the concept of which has been useful for conceptualizing how action selection is influenced by internal and external cues. States are roughly akin to 'situations' determined by alterations in external stimuli, such as a light or a noise, or by

internal information garnered from memory, or some combination of both. Within each state there are certain actions available and taking a particular action will transition the agent into the next state. Figure 1 shows an example of a state representation for the task of making a cup of tea. In this example, an individual might initiate making a cup of tea whilst in State 1, which they enter into upon walking into the kitchen and observing the kettle. For this task map then, State 1 is the initial state. That individual then switches the kettle on, causing the person to transition into the ‘kettle on’ state (State 2) determined by the noise, sight, and feeling of the kettle switching on. The agent continues in this manner until the tea is finally being drunk, at which point the terminal ‘goal’ state is reached (Figure 1).

Figure 1: Simplified state space representation of making a cup of tea by the agent according to model-based RL. S1, the initial state is determined by being in the kitchen, taking ‘A1’ action of switching on the kettle transitions the agent into state S2, determined by the kettle being on. State-to-state transitions are achieved between each state from S1-S7 in this manner by taking actions A1-A6, until the terminal state ‘S7’ determined by the tea being drunk (goal), is reached. S – state, A – action.

State representations are complex, and often require the integration of several sensory inputs with internal information. In our tea example above, the ‘kettle on’ state is determined by combining information from (at least) three sensory modalities (the noise, sight, and feeling of the kettle switching on) with the internal knowledge that ‘tea’ is the final goal. As has been noted elsewhere^{1,18}, the anatomical circuitry of the OFC is particularly well placed to integrate external and

internal information in this way, as it receives inputs from many sensory areas (e.g. auditory cortex, olfactory tubercle), as well as areas central to learning and memory such as the hippocampus, amygdala, and ventral tegmental area¹⁸⁻²⁰. Moreover, the OFC has substantial outputs onto striatal action selection centres²⁰, providing a pathway via which OFC-dependent state representations could directly influence actions. Both medial and lateral subregions reflect this pattern of connectivity to some extent, with a significant degree of overlap, but there are also some differences. For example, although both medial and lateral OFC project to basolateral amygdala, only medial OFC appears to project to the central amygdala²⁰⁻²¹. Similar differences exist with regards to cortico-cortical, hippocampal, striatal and other thalamic connections with to medial versus lateral OFC¹⁹⁻²¹. These differences are sometimes subtle (i.e. projections to different portions of the same structure¹⁹⁻²¹), but are present nonetheless. Overall, these studies suggest the anatomical suitability of both medial and lateral OFC for representing state, but the differences in their connectivity suggest that their role in doing so is not identical.

Thus, there is anatomical support for our claim that the lateral OFC represents initial states for action selection (state S1 in Figure 1) whereas medial OFC represents terminal states (i.e. the goal, state S7 in Figure 1). The empirical evidence for these claims is presented below, with a summary of the relevant findings presented in Table 1.

Table 1. Summary of behavioral findings with regards to orbitofrontal cortex subregions.

<u>Task</u>	<u>Structure inactivated/lesioned</u>	<u>Finding</u>	<u>Reference</u>
Instrumental Devaluation	Anterior medial OFC	Impaired	Bradfield et al., (2015; 2018)
Instrumental Devaluation	Posterior medial OFC	Intact	Munster & Hauber (2017); Bradfield et al., (2018).
Instrumental Devaluation	Lateral OFC	Intact	Ostlund & Balleine, (2007); Lichtenberg et al., (2017); Panayi & Killcross, 2018a; Parkes et al., (2018).
Instrumental Devaluation	Lateral OFC	Impaired	Gremel & Costa (2013; 2016)
Contingency reversal (value-based)	Lateral OFC	Impaired (slower reversal)	Chudsama & Robbins (2003); Dalton et al., 2016; Schoenbaum et al. (2003); Wilson et al. (2014)
Contingency reversal (identity-based)	Lateral OFC	Impaired	Parkes et al. (2018)
Contingency reversal (value-based)	Medial OFC	Higher perseveration during reversal	Gourley et al. (2010)
Contingency reversal (probabilistic task)	lateral OFC	Impairment	Dalton et al., (2016)
Contingency reversal (delay discounting)	Medial OFC	Facilitation	Mar et al. (2011)
Contingency Degradation	Lateral OFC	Impaired	Zimmerman et al., (2017)
Contingency Degradation	Lateral OFC	Encoding intact, sustained performance impaired	Zimmerman et al, (2018)
Contingency Degradation	Medial OFC	Intact	Bradfield et al. (2015)
Pavlovian instrumental transfer	Lateral OFC	Impaired	Ostlund & Balleine, (2007)
Pavlovian instrumental transfer	Medial OFC	Impaired	Bradfield et al., (2015; 2018)

OFC – Orbitofrontal cortex.

Instrumental Outcome Devaluation

Instrumental outcome devaluation is considered the gold standard protocol to assess whether actions are under goal-directed (model-based) control. The typical two-outcome version of this task is summarised in Table 2: In Stage 1, animals are trained to press two levers (A1 and A2) for two outcomes (O1 and O2). They are then given the opportunity to eat one of these outcomes to satiety to reduce its value relative to the alternate outcome¹². Immediately following satiety treatment rats are placed back into the operant chambers and allowed to respond freely on each lever under extinction (i.e. with no outcomes delivered). According to model-based RL, if animals are selecting actions according to the value of prospective outcomes, then they should preferentially respond on the lever associated with the valued outcome relative to the lever associated with the devalued outcome (valued > devalued).

Table 2. Summary of relevant behavioural paradigms

A – action, O – outcome, S – stimulus.

<u>Paradigm</u>	<u>Stage 1</u>	<u>Stage 2</u>	<u>Test</u>
Instrumental Devaluation	A1-O1 A2-O2	O1: 1hr	A1 < A2
Contingency Reversal (valued-based)	A1-O1 A2-	A1- A2-O1	
Contingency Reversal (identity-based)	A1-O1 A2-O2	A1-O2 A2-O1	O1: 1hr Original: A1 < A2 Reversed: A2 < A1
Contingency Degradation	A1-O1 A2-O2	A1-O1/O1 A2-O2	A1 < A2
Pavlovian-instrumental transfer	S1-O1 S2-O2	A1-O1 A2-O2	S1: A1 > A2 S2: A2 > A1

Critical to this task, animals must be able to accurately represent the outcomes of their actions (i.e. the terminal states). This is represented in Fig 2A. We have demonstrated⁴⁻⁵ that disruption to medial OFC functioning (anterior medial OFC in particular^{5,22}) via either permanent excitotoxic lesions, or temporary inactivation across test using chemogenetics (hM4Di DREADDs²³), abolished rats' capacity to show goal-directed choice on an outcome devaluation task. Instead, rats performed the two available actions (devalued and non-devalued) at equal rates, and this result is represented in Fig 2B. Importantly, when an almost-identical devaluation test was conducted, except that each outcome was actually delivered as a result of each lever press on test (i.e. outcomes were observable), medial OFC inactivated animals showed intact devaluation (valued > devalued). This pattern of responding matched predictions from our computational modelling in which we removed the assumption of the ability to infer terminal states defined by absent outcomes (the code for which can be found here: <https://github.com/adezfouli/OFCSim>, and the state space representation for which is in Figure 2A).

In contrast to this role for medial OFC, a number of experiments have demonstrated that inactivating the lateral OFC via either pre- or post-training lesions²⁴⁻²⁵, chemogenetic inactivation of lateral OFC cell bodies²⁶ or its connections with the basolateral amygdala²⁷ leave instrumental outcome devaluation intact, represented in Fig 2B. This is consistent with the argument that lateral OFC represents unobservable initial states, because both actions (left and right lever press) are available simultaneously in a single initial state (state 1 in Figure 2A). Thus, the levers are observable such that available actions need not be inferred from memory. In two exceptions to these findings, Gremel and Costa²⁸⁻²⁹ used a different kind of instrumental outcome devaluation task to infer that the lateral OFC might regulate instrumental outcome devaluation. The state space representation of this task is presented in Fig 2C. They trained animals to press a single lever in one context by using a random interval (RI) schedule to promote habits, and trained the same lever in another context using a random ratio (RR) schedule to promote goal-directed responding. On test, devaluation was demonstrated (Valued > Devalued) for lateral OFC-intact mice in the RR but not RI context, whereas

lateral OFC inactivated mice (by way of lesion, chemogenetics, or optogenetics) did not show devaluation in either context (Valued = Devalued, Fig 2D). As a result, the authors concluded that the lateral OFC is necessary for instrumental outcome devaluation.

Figure 2. State space and data representations for the role of medial versus lateral OFC in

instrumental outcome devaluation tasks. **A**, State space representation of a two-lever two-outcome instrumental outcome devaluation task. Rats in the initial state (S1) can choose to make one of two actions (A1 or A2) which will transition them into one of two terminal states (O1 and O2). **B**, When tested under extinction following outcome devaluation, intact goal-directed rats and rats with lateral OFC disruption preferentially choose the action that had given the still-valued outcome (A1, blue), whereas rats with mOFC disruption choose A1 and A2 equally^{4-5, 24-26}. **C**, State space representation of a single lever contextual discrimination task; rats need to use state information (S1 or S2) to determine the reinforcement schedule that delivers O1. **D**, When tested under extinction in the goal-directed context, intact rats show outcome devaluation but rats with IOFC disruption don't²⁸⁻²⁹. S – State, A – Action, O – Outcome, mOFC – medial orbitofrontal cortex, IOFC – lateral orbitofrontal cortex.

Although this remains a possibility, the involvement of IOFC in this task can equally be explained by its representation of initial state information. Specifically, it is likely that the alteration in reward schedules (RI vs. RR) between contexts required the formation of two different states to disambiguate these different contingencies³⁰⁻³¹ (e.g. State 1 in context 1, State 2 in context 2, Fig 2C). These latent states are internally generated, and function to retrieve the appropriate contingency and reduce interference by the competing contingency. As argued by Wilson et al.¹, animals with whole/lateral OFC lesions are unable to partition information based on latent states and so would attribute competing sets of contingencies to the same 'default' state such that they interfere with each other. Thus, we suggest that a similar process could be in effect here, such that intact animals partitioned information appropriately and demonstrated intact devaluation in the appropriate (goal-directed/RR) context, animals without a functional lateral OFC likely generalised between these states, or at the very least suffered interference between them.

Together, the results of these experiments neatly demonstrate an important difference between the role of initial and terminal state representation in action selection and the respective roles of lateral OFC and medial OFC in each. That is, as illustrated in Figure 2C, in the two action, two outcome devaluation task, initial state information is not necessary because both levers are observable, but animals must represent the different terminal states (O1 and O2, which are not presented on test) to select the appropriate action (valued > devalued). This capacity is profoundly disrupted by inactivations of medial but not lateral OFC^{4,5,9,24-27} (Figure 2 A-B). By contrast, an intact lateral OFC is only necessary for instrumental devaluation when animals must use initial state information to partition learning about two distinct action-outcome contingencies in order to apply the appropriate response strategy²⁸⁻²⁹ (Fig 2C-D).

Contingency Reversal

In a typical contingency reversal protocol, the animal is initially trained such that one action-outcome (A1-O1) contingency is rewarded but another is not (A2-), and this relationship is later reversed (A1-, A2-O1, see Table 2). Over time, intact animals will alter their response preference according to the currently rewarded contingency. Figure 3A-B shows a cognitive map of task space for reversal learning. According to this representation, when animals enter the reversal training phase in Fig 3B, they partition the new contingency into a separate state (S2) from the original contingency (S1) to prevent interference. Consistent with the presumed role of lateral OFC in initial state representation, there are several demonstrations that lateral/ventrolateral OFC inactivations (mainly lesions) impair contingency reversal performance, generally being slower to acquire the reversed response and making more perseverative errors³²⁻³⁵, shown in Figure 3C. Animals with lateral OFC lesions are presumed unable to partition information in this way, instead attributing both contingencies to the same 'default' state such that they suffer interference leading to impaired task performance. Likewise, lateral OFC lesions disrupt reversal learning when outcome identities are reversed²⁷. Rats were trained on two response-outcome pairings, A1-O1 and A2-O2, and then these contingencies were reversed; A1-O2, A2-O1, under chemogenetic inactivation of the lateral OFC. In between each phase of training, animals were tested for outcome devaluation to assess which response-outcome association was dominant. Whereas instrumental devaluation was intact following initial instrumental training (as mentioned above), when the lever press-outcome contingencies were reversed, lateral OFC inactivation did impair devaluation performance.

Figure 3. State space and data representation for single-outcome contingency reversal.

A, Rats in the initial state (S1) are trained such that one action is reinforced (A1-O) and the other isn't (A2-no O). **B**, When these contingencies are reversed (A1-no O, A2-O) rats form a new state (S2) within which to encode the new learning. **C**, Following reversal, intact rats rapidly bias responding to the newly reinforced action (accurate responses) and cease responding on the non-reinforced action (perseverative errors), whereas rats with mOFC disruption show increased perseverative errors, and rats with IOFC disruption show both increased perseverative errors and decreased accurate responses^{1,32,35,37}. S – State, A – Action, O – Outcome.

There are relatively fewer investigations of the role of the medial OFC in reversal learning, and several of the studies that have investigated it have done so in probabilistic tasks^{33,36}, confounding their interpretation. As we have argued previously⁴ in tasks such as probabilistic learning, delay discounting, and even progressive ratio tasks, where medial OFC lesions are often found to impair performance^{6,22,33,36}, this is likely because each of these tasks require the estimation of likely outcomes when those outcomes are not directly observable. Thus it is straightforward to understand how removing the medial OFC would impair the ability to properly make such estimations and ultimately impair performance.

Gourley et al.³⁷ did, however, examine the effect of medial OFC lesions on a more straightforward reversal task in mice, and found them to impair reversal learning. The manner of this impairment was somewhat different, however, to that resulting from whole or ventrolateral OFC lesions in the same study. Specifically, whereas ventrolateral OFC-lesioned mice were slower to learn

the new (reversed) contingencies, medial OFC-lesioned mice were able to learn the new contingencies as readily as controls (represented in Figure 3C) but additionally performed more perseverative responses on the previously rewarded nose-poke. The fact that lesioned mice learned the new contingencies as quickly as controls suggests they were able to partition conflicting contingencies according to initial state (see¹ and Figure 3A-B). Instead, in line with our framework, the deficit may lie in their inability to appropriately extinguish the previously rewarded response, due to an inability to retrieve a representation of the expected outcome and appropriately assign an error-based learning mechanism to drive model-based extinction. Therefore, mice with medial OFC lesions would have had to rely on model-free processes alone to extinguish the initial contingency, causing slower extinction and more perseverative responses relative to controls.

Contingency Degradation

Like outcome devaluation, sensitivity to contingency degradation is considered characteristic of goal directed instrumental actions. A typical protocol is presented in Table 2. In Stage 1, animals are trained on two action-outcome contingencies, and then in Stage 2, animals receive additional, unsignalled deliveries of one of the outcomes (A1-O1/O1). This effectively reduces the contingency between that action and outcome^{4,38}. An animal using goal-directed control will reduce its responding on the degraded lever, whilst maintaining responding on the non-degraded lever (in accordance with the contingency criterion). We have previously argued that contingency degradation relies on the ability to separate the different contingencies learned in Stage 1 (original contingencies) and Stage 2 (degraded/nondegraded contingencies) into separate internal states³⁸. Importantly, however, because outcomes are present and observable throughout this task, an inability to infer terminal outcome states should not be necessary for intact degradation performance. Thus, our framework predicts that the lateral OFC but not the medial OFC should be critical for performance on this task.

There is some support for this prediction. Specifically, we have found that an intact medial OFC was not necessary for contingency degradation performance using this design⁴. Although the role of the lateral OFC in this exact procedure has not been tested, the selective knockdown of brain-derived neurotrophic factor (BDNF) in the lateral OFC in mice has been found to abolish nominally instrumental (nose-poke) contingency degradation performance³⁹. It should be noted however, that even more recently, Zimmerman et al.⁴⁰ used the same procedure and found that chemogenetic inactivation of ventrolateral OFC had no effect on degradation learning, but it did have an effect on sustained degradation performance. Although our framework predicts an effect of lateral OFC inactivation on both learning and performance of degradation, it is possible that the ventrolateral target of these inactivations has a slightly different function to the more lateral portion. Alternately, the reason for this discrepancy may be methodological, as these studies used a somewhat different procedure to that we have described (here and in Table 2), in which responding no longer produced any contingent outcomes (contingency of -1 rather than 0). Thus, it is possible that this more acute version of the task encouraged learning according to model-free processes that could proceed in the absence of state inference.

Taken together, therefore, these studies provide partial support for the notion that lateral OFC regulates contingency degradation, but more studies are needed using a purely instrumental (e.g. lever press) version of the task that does not involve extinction. For this task, our framework makes the clear prediction that lateral but not medial OFC inactivation should impair performance.

Pavlovian-instrumental Transfer

Although the results of OFC dysfunction in each of the tasks reviewed above provides clear evidence of the functional heterogeneity of medial and lateral OFC, this is not always the case. In specific Pavlovian-instrumental transfer (PIT; Table 2), for example, lesions of both medial and lateral OFC have been found to impair performance^{4-5,24}. In this task, rats are first trained to

associate two stimuli with two unique outcomes (S1-O1, S2-O2). In the second, instrumental, phase of training, they are trained to perform two distinct actions for those same outcomes (A1-O1, A2-O2). On test, when presented with the stimuli and levers together for the first time, but no food outcomes, animals will typically respond on the lever associated with the same outcome as the stimulus that is being presented (e.g. S1: A1 > A2 due to the association of both S1 and A1 with O1). We suggest that specific PIT requires the inference of both initial and terminal state representations to guide performance, which is why lesions of both have been found to affect its performance. Specifically, initial state information is necessary on test because this is the first time that the Pavlovian and instrumental contingency information learned during training must be inferred from memory, and integrated to drive responding on the correct lever. In addition, unobservable terminal state information about which lever earned which outcome must be inferred to drive responding on the correct lever. Thus, consistent with the aforementioned findings, our framework predicts a role for both lateral and medial OFC in this task.

Do lateral and medial OFC represent state information in the same way for Pavlovian tasks?

All of the studies we have reviewed so far have been either entirely, or predominantly, instrumental in nature. However, a number of studies have shown that the lateral OFC also appears to regulate learning and/or expression in a number of Pavlovian tasks such as (Pavlovian) devaluation^{25,41-43}, overexpectation⁴⁴, identity unblocking⁴⁵, and preconditioning⁴⁶. Furthermore, if the same framework outlined above were applied to these tasks, then it would be difficult to argue that lateral OFC is representing initial but not terminal states. In Pavlovian devaluation, for example, this task is usually employed in the absence of any conflicting contingencies. Thus, if we apply the same logic outlined above for the instrumental version of this task, lateral OFC inactivations should leave it intact, but this is clearly not the case^{25,41-43}.

We believe that the reason for this might lie in the fundamental difference in how the cognitive map of a Pavlovian versus an instrumental devaluation task is mentally represented. Specifically, in an instrumental task an action must be performed in order to transition through states (referred to as the ‘transition function’⁴⁷). In our tea example above (Figure 1), the individual switches on the kettle to transition into the ‘kettle on’ state and the individual will not transition to that state if that action is not performed. This is not the case in Pavlovian studies in which a cue (e.g. a tone) is typically presented, and food is delivered to the food receptacle regardless of the animal’s actions. Thus, the animal transitions into a ‘food delivered’ state without performing any action (note that entry into the food receptacle could be seen as potentially instrumental, but could also be a Pavlovian conditioned response to the sight of the food). As such, it is not clear exactly what would lead to the state-to-state transition in a ‘model-based’ cognitive map of a Pavlovian task (time? Sequence of events?), and as has been noted in articles regarding the potential model-based nature of some Pavlovian responses⁴⁸⁻⁴⁹ the algorithm – and hence the transition function- for model-based Pavlovian learning is still unknown.

This is relevant to the current argument, because the separation of initial and terminal states in instrumental studies is reliant on the action-based transitions between them. In the absence of such transitions, it is difficult to say what the potential role of medial versus lateral OFC in Pavlovian studies might be. Further, it has been argued^{48,50} that in Pavlovian conditioning, the stimuli may take on properties of the outcome in a way that does not occur in instrumental conditioning (referred to as ‘stimulus substitution’), and indeed phenomena such as conditioned reinforcement⁵¹ and second-order conditioning⁵² support this claim. If this is indeed the case, then this would suggest that it is difficult (if not impossible) to separate initial versus terminal state information in such studies, because the presentation of a cue itself has taken on properties of both the cue and outcome – i.e. the initial and terminal states. As such, there may be no role for the medial OFC in Pavlovian learning, as this is fulfilled by the Pavlovian cue. Instead, such learning relies on the lateral OFC to retrieve a representation of task structure when Pavlovian cues are

ambiguous¹¹. Alternatively, it is possible that the framework we have outlined here applies solely to instrumental studies, and that medial versus lateral OFC might have very different roles in Pavlovian studies that may or may not be related to state representation. More research is clearly needed to tease apart these many possibilities.

The Orbitofrontal Cortex as an Interface between Memory and Action Selection

Although the studies reviewed here have been primarily conducted in rodents, we believe that these relative roles for medial and lateral OFC could apply across species. In support of this idea, recent evidence⁵³ has suggested that OFC circuitry is highly homologous between rodents and primates (despite differences in cytoarchitecture⁵⁴). Evidence of a degree of functional homology can be drawn from studies in humans that have also directly implicated medial OFC as representing a cognitive map of task space². Moreover, several primate studies^{2,9,55} have found differences in the function of lateral and medial OFC in decision-making tasks, although these tasks almost always involve the presence of both Pavlovian and instrumental contingencies (as is often necessary for practical reasons) as well as several different alterations in contingencies throughout the task, making it difficult to tease apart the relative representations by lateral and medial OFC with regards to state. Superficially, then, it is possible that these relative functions of medial and lateral OFC are preserved across species, although as noted by Wilson et al.,¹ the complexity of state and task representations in humans and primates is likely much richer than that which exists in the rodent brain.

It is also worth noting that even within different subregions of OFC, subpopulations of neurons have shown a remarkable degree of heterogeneity with regards to the types of functions they subserve. Electrophysiological studies in rats and monkeys have discovered, for example, that a subpopulation of neurons located within the lateral OFC preferentially codes for value⁵⁶⁻⁵⁹, but others have different lateral OFC populations that respond to outcome identities, outcome location,

or even responses, independently of value⁵⁹⁻⁶⁰. Moreover, these differing responses are not always static across the changing conditions of the task, with some OFC neurons observed to respond to value prior to reversal and then stop responding to value after reversal, in a seemingly context (or state) dependent manner⁶¹. An even more recent study has managed to optogenetically manipulate distinct populations of cells within lateral OFC that either selectively responded to feeding or social interaction experiences, and found differing effects on both types of behaviour⁶². Studies such as these complement the notion that the OFC integrates the rich and varied range of information that is necessary to form representations of state, and further suggest that intercellular processes may contribute to such representations.

Finally, we wish to conclude by re-iterating that although we have here proposed different roles for medial and lateral OFC in representing terminal versus initial states, respectively, these roles are likely achieved interdependently rather than independently. For example, although we have argued that medial OFC does not infer initial states directly, medial OFC-dependent retrieval of particular outcomes (terminal states) may sometimes be necessary for the accurate inference of the initial state. A potential example can be drawn from Pavlovian-instrumental transfer, in which the presentation of a stimulus may retrieve a representation of its associated outcome (terminal state), and infer the availability of that outcome in the initial state, promoting selection of the response that will deliver it. An alternative real-world example might be if you visit your regular café under the understanding it is Saturday (initial state) and you therefore infer from your knowledge that the daily special is steak (outcome state). The converse can also be true; when you see that the daily special offered is steak (outcome state), you may infer that today must be Saturday (initial state). In this way, it is possible to imagine many scenarios in which the medial and lateral OFC regions might work in concert to achieve a cohesive cognitive map of task space.

References

1. Wilson, R.C., Takahashi, Y.K., Schoenbaum, G. & Niv, Y. Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267-279. (2014). <https://doi.org/10.1016/j.neuron.2013.11.005>.
2. Schuck, N.W., Cai, M.B., Wilson, R.C. & Niv, Y. Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* **91**, 1402-1412. (2016). <https://doi.org/10.1016/j.neuron.2016.08.019>
3. Sharpe, M.J., Wikenheiser, A.W., Niv, Y. & Schoenbaum, G. The state of the orbitofrontal cortex. *Neuron* **88**, 1075-1077. (2015). <https://doi.org/10.1016/j.neuron.2015.12.004>
4. Bradfield, L.A., Dezfouli, A., van Holstein, M., Chieng, B. & Balleine, B.W. Medial orbitofrontal cortex mediates outcome retrieval in partially observable task situations. *Neuron* **88**, 1268-1280. (2015). <https://doi.org/10.1016/j.neuron.2015.10.044>
5. Bradfield, L.B., Hart, G. & Balleine, B.W. Inferring action-dependent outcome representations depends on anterior but not posterior medial orbitofrontal cortex. *Neurobiol Learn Mem* **155**, 463-473. (2018). <https://doi.org/10.1016/j.nlm.2018.09.008>
6. Gardner et al. Medial orbitofrontal inactivation does not affect economic choice. *eLife* **7**, e38963. (2018). <https://doi.org/10.7554/eLife.38963>
7. Boudet, S. & Richmond, B. J. Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J Neurosci* **30**, 8591-8601. <https://doi.org/10.1523/JNEUROSCI.0049-10.2010>
8. Noonan et al. Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc Natl Acad Sci U S A* **107**, 20547-20552. (2010). <https://doi.org/10.1073/pnas.1012246107>
9. Rudebeck, P.H. & Murray, E.A. Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behaviour. *J Neurosci* **31**, 10569-10578. <https://doi.org/10.1523/JNEUROSCI.0091-11.2011>
10. Malvarez, M., Shieh, C., Murphy, M.D., Greenfield, V.Y. & Wassum, K.M. (2019). Distinct cortical-amygdala projections drive reward value encoding and retrieval. *Nat Neurosci* **22**, 762-769. (2019). <https://doi.org/10.1038/s-41593-019-0374-7>
11. Sharpe, M.J., Stalnaker, T., Schuck, N.W., Killcross, S., Schoenbaum, G. & Niv, Y. An integrated model of action selection: distinct modes of cortical control of striatal decision-making. *Ann Rev Psychol* **70**, 53-76. (2019). <https://doi.org/10.1146/annurev-psych-010418-102824>
12. Dickinson, A. & Balleine, B.W. Motivational control of goal-directed action. *Anim Learn Behav* **22**, 1-18. (1994).
13. 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). <https://doi.org/10.1111/j.1460-9568.2004.03095.x>
14. 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. (2005a). <https://doi.org/10.1111/j.1460-9568.2005.04219.x>
15. Yin, H.H., Ostlund, S.B., Knowlton, B.J. & Balleine, B.W. The role of the dorsomedial striatum in instrumental conditioning. *Eur J Neurosci* **22**, 513-523. (2005b). <https://doi.org/10.1111/j.1460-9568.2005.04218.x>

16. Daw, N.D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioural control. *Nat Neurosci* **8**, 1704-1711. <https://doi.org/10.1038/nn1560>
17. Dezfouli, A. & Balleine, B.W. Habits, action sequences and reinforcement learning. *Eur J Neurosci* **35**, 1036-1051. (2012). <https://doi.org/10.1111/j.1460-9568.2012.08050.x>
18. Schuck, N.W., Wilson, R. & Niv, Y. A state representation for reinforcement learning and decision-making in the orbitofrontal cortex. In Goal-directed decision-making: computations and neural circuits. Morris, R., Bornstein, A., & Shenhav, A. (Eds). (Elsevier: Academic Press, 2018). <https://doi.org/10.1016/B978-0-12-812098-9.00012-7>
19. Reep, R.L., Corwin, J.V. & King, V. Neuronal connections of orbital cortex in rats: topography of cortical and thalamic afferents. *Exp Brain Res* **111**, 215-232. (1996).
20. McDonald, A.J., Mascagni, F. & Guo, L. Projections of the medial and lateral prefrontal cortices to the amygdala: a Phaseolus vulgaris leucoagglutinin study in the rat. *Neuroscience* **71**, 55-75. (1996).
21. Hoover, W.B. & Vertes, R.P. Projections of the medial orbital and ventral orbital cortex in the rat. *J Comp Neurol* **519**, 3766-3801. (2011). <https://doi.org/10.1002/cne.22733>.
22. Munster, A. & Hauber, W. Medial orbitofrontal cortex mediates effort-related responding in rats. *Cerebral Cortex* **28**, 1-11. (2017). <https://doi.org/10.1093/cercor/bhx293>
23. Armbruster, B.N., Li, X., Pausch, M.H., Herlitze, S. & Roth, R.L. Evolving the lock to fit the key to create a family of G protein-coupled receptors potently activated by an inert ligand. *Proc Natl Acad Sci U S A* **104**, 5163-5168. (2007). <https://doi.org/10.1073/pnas.0700293104>
24. Ostlund, S.B., & Balleine, B.W. (2007). Orbitofrontal cortex mediates outcome encoding in Pavlovian but not instrumental conditioning. *Journal of Neuroscience*, **27**, 43-52. <https://doi.org/10.1016/j.neuron.2016.03.006> pmid:27100198
25. Panayi, M.C. & Killcross, S. Functional heterogeneity within the rodent orbitofrontal cortex dissociates outcome devaluation and reversal learning deficits. *eLife* **7**, e37357. (2018). <https://doi.org/10.7554/eLife.37357>
26. Parkes et al. Insular and ventrolateral orbitofrontal cortices differentially contribute to goal-directed behaviour in rodents. *Cereb Cortex* **28**, 2313-2325. (2018). <https://doi.org/10.1093/cercor/bhx132>
27. Lichentenberg et al. Basolateral amygdala to orbitofrontal cortex projections enable cue-triggered reward expectations. *J Neurosci* **25**, 486-417. (2017). <https://doi.org/10.1523/JNEUROSCI.0486-17.2017>
28. Gremel, C.M. & Costa, R.M. Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nat Comms* **4**, (2013). <https://doi.org/10.1038/ncomms3264>
29. Gremel et al. Endocannabinoid modulation of orbitostriatal circuits gates habit formation. *Neuron* **90**, 1312-1324. (2016). <https://doi.org/10.1016/j.neuron.2016.04.043>
30. Gershman, S.J., & Niv, Y. Context, learning, and extinction. (2010). *Psych Rev* **117**, 197-209. (2010). <https://doi.org/10.1037/a0017808>
31. Gershman, S.J., & Niv, Y. Exploring a latent cause theory of classical conditioning. *Learn Behav* **40**. (2012). <https://doi.org/10.3758/s13420-012-0080-8> pmid:22927000
32. 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**, 8771-80. (2003). <https://doi.org/10.1523/JNEUROSCI.23-25-08771.2003>

33. Dalton, G.L., Wang, N.Y., Phillips, A.G. & Floresco, S.B. Multifaceted contributions by different regions of the orbitofrontal prefrontal cortex to probabilistic reversal learning. *J Neurosci* **36**, 1996-2006. (2016). <https://doi.org/10.1523/JNEUROSCI.3366-15.2016>.
34. Schoenbaum, G., Nugent, S.L., Saddoris, M.P. & Setlow, B. Orbitofrontal lesions in rats impair reversal but not acquisition of go no-go odour discriminations. *Neuroreport* **13**, 885-890. (2002). <https://doi.org/10.1097/00001756-200205070-00030>
35. 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).
36. Mar, A.C., Walker, A.L., Theobald, D.E., Eagle, D.M. & Robbins, T.W. Dissociable effects of lesions to orbitofrontal cortex subregions on impulsive choice in the rat. *J Neurosci* **31**, 6398-6404. (2011). <https://doi.org/10.1523/JNEUROSCI.6620-10.2011>
37. Gourley, S.L., Lee, A.S., Howell, J.L., Pittenger, C. & Taylor, J.R. Dissociable regulation of instrumental action within mouse prefrontal cortex. *Eur J Neurosci* **32**, 1726-1734. (2010). <https://doi.org/10.1111/j.1460-9568.2010.07438>
38. Bradfield, L.A., Bertran-Gonzalez, J., Chieng, B. & Balleine, B.W. The thalamostriatal pathway and cholinergic control of goal-directed action: Interlacing new with existing learning in the striatum. *Neuron* **79**, 153-166. (2013). <https://doi.org/10.1016/j.neuron.2013.04.039>
39. Zimmerman, K.S., Yamin, J.A., Rainnie, D.G., Ressler, K.J. & Gourley, S.L. Connections of the mouse orbitofrontal cortex and regulation of goal-directed action selection by brain-derived neurotrophic factor. *Biol Psychiatry* **81**, 366-377. (2017). <https://doi.org/10.1016/j.biopsych.2015.10.026>.
40. Zimmerman, K.S., Li, C., Rainnie, D.G., Ressler, K.J. & Gourley, S.L. Memory retention involves the ventrolateral orbitofrontal cortex: comparison with the basolateral amygdala. *Neuropsychopharm* **43**, 373-383. (2018). <https://doi.org/10.1038/npp.2017.139>
41. Pickens et al. Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. *J Neurosci* **23**, 11078-11084. (2003).
42. Gallagher, M., McMahan, R.W. & Schoenbaum, G. Orbitofrontal cortex and representation of incentive value in associative learning. *J Neurosci* **19**, 6610-6614. (1999).
43. Izquierdo, A. & Murray, E.A. Combined unilateral lesions of the amygdala and orbital prefrontal cortex impair affective processing in rhesus monkeys. *J Neurophysiol* **91**, 2023-2039. <https://doi.org/10.1152/jn.00968.2003>
44. Takahashi et al. The orbitofrontal cortex and ventral tegmental area are necessary for learning from unexpected outcomes. *Neuron* **62**, 269-280. (2009) <https://doi.org/10.1016/j.neuron.2009.03.005>.
45. McDannald et al. Orbitofrontal neurons acquire responses to 'valueless' Pavlovian cues during unblocking. *eLife* **3**, e02653. (2014). <https://doi.org/10.7554/eLife.02653>
46. Jones et al. Orbitofrontal cortex supports behaviour and learning using inferred but not cached values. *Science* **338**, 953-956. (2012). <https://doi.org/10.1126/science.1227489>
47. Daw, N.D. Model-based reinforcement learning as cognitive search: neurocomputational theories. In *Cognitive Search: Evolution, algorithms and the brain*. Todd, P.M, Hilles, T.T., & Robbins, T. W. (Eds). 195-208. (Cambridge, MA: MIT Press, 2012).
48. Dayan, P. & Berridge, K.C. Model-based and model-free Pavlovian reward learning: revaluation, revision, and revelation. *Cogn Affect Behav Neurosci* **14**, 473-492. (2013). <https://doi.org/10.3758/s13415-014-0277-8>.
49. Pauli, W.M., Gentile, G., Collette, S., Tyska, J.M. & O'Doherty, J.P. Evidence for model-based encoding of Pavlovian contingencies in the human brain. *Nat Comms* **10**, 1-11. | <https://doi.org/10.1038/s41467-019-08922-7>

50. Pavlov, I.P. Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. (G.V. Anrep, Trans.). (London: Oxford University Press, 1927).
51. Robbins, T.W. The acquisition of responding with conditioned reinforcement: effects of pipradrol, methylphenidate, d-amphetamine, and nomifensine. *Psychopharm* **58**, 79-87. (1978).
52. Holland, P.C. & Rescorla, R.A. (1975). The effect of two ways of devaluing the unconditioned stimulus after first- and second-order appetitive conditioning. *J Exp Psychol: Anim Behav Proc* **1**, 355-363. <https://doi.org/10.1037/0097-7403.1.4.355>
53. Heilbronner, S.R., Rodriguez-Romaguera, J., Quirk, G.J., Groenewegen, H.J. & Haber, S.N. (2016). Circuit-based corticostriatal homologies between rat and primate. *Biol Psychiatry* **80**, 509-521. (2016). <https://doi.org/10.1016/j.biopsych.2016.05.012>
54. Wise, S.P. (2008). Forward frontal fields: phylogeny and fundamental function. *Trends Neurosci* **31**, 599-608. (2008). <https://doi.org/10.1016/j.tins.2008.08.008>
55. Rudebeck, P.H. & Murray, E.A. Balkanizing the primate orbitofrontal cortex: Distinct subregions for comparing and contrasting values. *Ann N Y Acad Sci* **1239**, 1-13. (2014). <https://doi.org/10.1111/j.1749-6632.2011.06267.x>
56. Padoa-Schioppa, C. & Assad, J.A. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223-226. (2006). <https://doi.org/10.1038/nature04676>
57. Stalnaker et al. Orbitofrontal neurons infer the value and identity of predicted outcomes. *Nat Comms* **5**, 3926. (2014). <https://doi.org/10.1038/ncomms4926>
58. 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 Neurophys* **74**, 733-750. (1995). <https://doi.org/10.1152/jn.1995.74.2.733>
59. Tremblay, L. & Schultz, W. Relative reward preference in primate orbitofrontal cortex. *Nature* **398**, 704-749-8. (1999).
60. 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). <https://doi.org/10.1016/j.neuron.2006.06.027>
61. 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). <https://doi.org/10.1523/JNEUROSCI.19-05-01876.1999>
62. Jennings et al. Interacting neural ensembles in orbitofrontal cortex for social and feeding behaviour. *Nature* **565**, 645-649. (2019). <https://doi.org/10.1038/s41586-018-0866-8>