Neuronal Mechanisms of Decision Making in the Prefrontal Cortex

Nishantha Malalasekera

Sobell Department of Motor Neuroscience and Movement Disorders

A dissertation submitted in partial fulfilment of the requirements for the degree of Doctor or Philosophy in University College London

I, Weeratunge Mudiyanselage Nishantha Malalasekera confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

This thesis examines several aspects of decision computations which are critical for understanding the processes by which decisions are made. It will show that subjects engaged covert attention to bias both saccadic and choice processes during simple decision making tasks even when these stimuli were novel. This saccadic behaviour was overridden when one presented stimulus is relatively more novel than the other implying the existence of separate value comparison circuits in the brain which deal with making value based decisions about attention and choice respectively. Even when the task was made more complex by introducing multiple decision variables this phenomenon of covert attention was maintained. This thesis will demonstrate that subjects controlled both the amount and manner of information gathering during decisions. This behaviour showed features of a confirmation bias.

Single cell neuronal recordings were performed while subjects executed a multi-attribute decision making task. ACC neurons represented action values and different populations of OFC neurons encoded attribute and attentional values. These neurons did not just reflect value (i.e. an input into a decision process) but instead evolved their coding to represent final choice thereby implying the existence of a parallel decision making circuit which compares value in different frames of reference. Information gathering strategy was also computed in the same frames of reference implying the existence of a common value comparison system which simultaneously drives both choice and information gathering. At the outcome of the decision ACC neurons encoded both categorical reward outcome and positive prediction errors. vmPFC neurons encoded prediction errors while OFC and ACC neurons encode fictive value when rewards were withheld. Finally frame of reference specific computations were observed in LPFC and OFC. The results in this thesis therefore provide novel insight into the role of valuation circuitry during value based decision making and outcome monitoring.

Acknowledgements

Dr. Steve Kennerley - As my primary supervisor it was Steve who gave me the opportunity to do this work that I have enjoyed so much. He is definitely a perfectionist and this attitude translates to the quality of the work presented in this thesis. Steve always pushed the project forwards by constantly coming up with new ideas for designing the task and analysing the data. He was always willing to forgive mistakes (and typos) I'm honoured to be his first student to graduate.

Dr. Simon Farmer – As my secondary supervisor Simon took a keen interest in my project throughout. He was also happy to provide insight and advice on any aspect of the project of my studies in general. He always had a keen eye on ensuring that my progress (and my research log) was on track. I also have to thank Simon for several enjoyable evenings at the Athenaeum Club and for paying my stipend for the final year of my PhD.

Dr. Bruno Miranda – He has helped me out in innumerable and unlistable ways. They say that primate neurophysiology is a team sport; if it were cricket then Bruno would be the wicketkeeper because he is always involved. He is also a great friend and somebody I continually look up to.

Dr. Laurence Hunt – Always happy to help me out whether it be with experiments or analysis, Laurence has been a constant source of ideas and criticism (the good kind). I have attained a mountain of knowledge from Laurence but I wish I had a few more years to mine his brain for more.

Prof. Tim Behrens – The genius that is Tim provided ideas and insight into analysis interpretation on countless occasions. His help is greatly appreciated for giving me a set of results that I am proud of.

Lianne McCombe – As lab manager Lianne has helped me in many ways from subject training, procedures to generally ensuring that the lab did not fall apart. She is also a very good friend who has helped make my time at the Sobell so enjoyable.

Chen Cheng – Despite only working with her for six months she soon became the standard by which I measured my ability to present data. She taught me to think clearer and write better and I am honoured to have worked with her and to call her my friend. She also regularly gave me food which ranks highly amongst my most prized qualities.

Sean Cavanagh – Performed the initial analysis of the data presented in Chapter 3. He was the first student I supervised and is all round good guy.

Prof. Roger Lemon and **Dr. Alexander Kraskov** – The two people I have to thank for introducing me to science. They introduced me warmly into their lab and their passion for neurophysiology rubbed off on me. Without their help and input over the course of my project I would never have completed it in time.

Prof. Gordon Stewart and **Susan Beesley** – The support and belief of Gordon and Sue made this experience one that I will never forget or regret.

Jonathon Henton and **Spencer Neil** – They built almost everything in the lab, without their tireless work and ingenuity I would have no data.

Deborah Hadley, **Kully Sunner**, **Hayley Mackenzie** and **Chris Seers** – Four people who have never been anything but helpful no matter how inane my administrative/computing query.

Anoma and Weeratunge Malalasekera – I would not be where I am today without the support of my parents. I hope that I have done them proud by matching their academic qualifications.

Shivan Kanani – One of my oldest and closest friends who had the misfortune to live with me for a year. I have shared some great moments with him and he has always there for moral support when needed.

Vishal Rawji – The person who I currently have the misfortune of living with, his own passion for neuroscience and clinical medicine pushes me to do better. He is also one of my closest friends.

Elisabeth Engl – Somebody who helped train my subjects even when it had no benefit to her. We became close friends and despite her moving on to pastures anew we have never forgotten the great time we had in the lab.

Maria Botcharova and **Hedi Young** – A pair of friends with highly charged opinions who would never fail to ignite an interesting topic of conversation at lunch.

Stephan Waldert, **Roland Philipp** and **Ganesh Vignesawan** – The 'three musketeers' from the Lemon/Kraskov lab who I have always enjoyed a joke and a scientific debate with. I have special memories of our time in San Diego together.

Rishi Banerjee and **Akash Maniam** – Leaving my medical peer group and setting off in a new direction was daunting but with friends like these two by my side I never felt alone. I have travelled to far flung shores with them in the pursuit of merriment and cricket.

RUMS Cricket Club – It continues to be an honour to play next to all the men at RUMS CC. After my last acknowledgee playing cricket has been my greatest passion outside the lab. They have given me lessons and memories to last a lifetime.

Thushara Kaluarachchi – Last but definitely not least the light of my life who has always been there to support me even when she had her own challenges to deal with. She made me want to do better and achieve more than I have ever dreamt of doing. The project never felt like work when I knew that I would come home to Thushara.

Table of Contents

Long Abstract

The motivation for the work presented in this thesis comes from a need to understand both the anatomical specificity and computational generalisability of value related computations throughout prefrontal cortex (PFC) during decision making. In this thesis inferences about the neuronal circuitry will be made using two tools: 1) behavioural dynamics such as reaction times, choice and information gathering behaviour, 2) the firing patterns of individual neurons in various parts of the PFC. This thesis will concentrate on four subregions of PFC which have all been implicated in various types of cognitive function; anterior cingulate cortex (ACC), lateral prefrontal cortex (LPFC), orbitofrontal cortex (OFC) and ventromedial prefrontal cortex (vmPFC).

Chapter 1 will aim to outline the current state of the neurophysiological evidence pertaining to decision making and the PFC. Although neurons throughout the brain respond to valuable stimuli, lesion studies in humans and animals imply that the PFC is critical for optimal value based decision making. Many neurons in PFC encode stimuli in a 'common value currency' (particularly in ACC) which may serve to feed into a general value comparison system in line with both psychological and some empirical models of decision making. However, evidence suggests that many OFC neurons do not indiscriminately encode value but instead do so only for specific value properties (i.e. attributes) such as reward size, risk or delay (to name but a few). These neurons may be critical when decisions are made on the basis of individual attributes rather than the overall values of each option.

The link between value and action is critical for converting the abstract decision to a definite movement which obtains the chosen goal. A subpopulation of neurons in ACC may support this link by multiplexing both value and action computations. Furthermore, ACC has been heavily implicated in the integration of physical effort into the decision process. These

findings lead to the possibility that decisions may also be resolved through competition between available valuable actions.

Human studies have revealed that the manner in which people make decisions is influenced by attention. Furthermore, vmPFC has been observed to modulate its value coding based on the location of current attention. These signals may also have a relevance to the decision making process.

The existence of value signals that have different 'frames of reference' (i.e. attribute in OFC, action in ACC and attention in vmPFC) leads to an important juncture in the field. Do these signals reflect a mere input into a serial decision system which compares the integrated values of various options? Or conversely is there a decision circuit based on parallel computations which allows value to be compared in various frames of reference? Although evidence from perceptual decision making implies that the latter is more likely, it remains unclear because the anatomical specificity and computational characteristics of these frame of reference signals remains under-explored. This question is therefore one that this thesis will attempt to answer.

ACC has been shown to be involved in many other cognitive processes including outcome mapping, behavioural flexibility, volatility, error monitoring and conflict. ACC may also play critical role in controlling and updating complex behaviour through its coding of reward prediction errors during the outcome phases of various tasks. The aforementioned findings could be reconciled by considering ACC to be an action-outcome predictor which can therefore exert behavioural control.

An often ignored aspect of optimal decision making is information gathering. Basic human information gathering behaviour appears to depend on the complexity of the decision both in terms of its difficulty (i.e. how easy it is to compute the various option values) and the number of parameters that need to be considered. As decisions become more complicated humans tend to both gather less information and switch from a strategy that gathers

information about individual options to one that gathers information about individual attributes. This strategy shift seems to be impaired in vmPFC lesion patients implying that this region may be important controlling information gathering in decision making.

Chapter 2 will lay out the basic anatomical connectivity of the PFC. The pattern of connectivity constrains the computation function of each area and individual PFC regions show a relative diversity in their connections. Basic sensory input from all sensory modalities is concentrated predominantly on OFC (and partially on vmPFC). In contrast ACC and LPFC are the major source of efferent connections to the motor system. ACC sends projections mainly to motor structures including presupplementary motor area, premotor area and primary motor cortex whereas LPFC connects strongly to supplementary eye fields. Limbic (and complex visual) inputs from the temporal cortex span the PFC, however some connections are region specific. OFC receive preferential connections from the hippocampus whereas the parraphippocampal cortex preferentially projects to ACC. Other less region specific projections exist from the entorhinal cortex, amygdala and temporal pole.

Dopaminergic projections to PFC also show topography with OFC connecting more strongly to ventral tegmental area (VTA) and ACC projecting more to Substantia Nigra pars compact (SNpc). Furthermore, there is a clear gradient in the number of dopamine receptors across the frontal cortex as a whole. The consequence of this is that ACC receives a much stronger dopaminergic input than other PFC regions.

The connectivity between PFC regions implies a flow of information from basic sensory areas to OFC which then relays subsequent computations on to LPFC and ACC (potentially though vmPFC in the latter case) which then communicate with motor and saccadic regions. Although this framework suggests a serial decision making pathway, little is known about circuits *within* PFC regions which may allow for parallel computations within each region.

Chapters 3 and 4 concentrate on characterising various aspects of primate behaviour during simple value based decision making. The first of these chapters looks at the link between attention and choice by examining the eye movements and the choices of subjects as they performed a simple decision making task using stimuli that are either well learnt or relatively novel to them. This chapter will show that both the direction of subjects' initial saccades and their final choices are guided by a value comparisons performed using covert attention. Both effects show characteristics of learning however there is a clear dissociation between initial saccadic and final choice when novel and overlearned stimuli are presented simultaneous. Covert attention constantly bias saccades *towards* novel stimuli but in contrast bias final choice *away* from novel stimuli early on in the session. This implies the existence of two at least partially separable valuation systems in the brain; one which guides overt attention and another which guides motor (i.e. limb) responses.

Subjects were then presented with a more complicated decision making task in the experiment discussed in Chapter 4. This involved them making decisions based on two separate attributes (probability and magnitude). In order for subjects to be optimal they were required to integrate the value of the two attributes to compute the option value of the two choices. Value was still seen to influence initial saccade direction however subjects rarely overtly gathered all available information before making a choice. However, choice behaviour was better explained by assuming that subjects considered all available information (covertly) before making decisions. This finding along with the fact that both the amount and the location of information gathering were guided by covert attention implies that in this task subjects principally perform value comparison through covert attention.

The second part of this chapter introduces the behavioural results of a variant of the multi-attribute decision making task where an element of information gathering was introduced. This was done by covering up all four stimuli (two for each option) and sequentially presenting one stimulus at a time. The location of the first two observed stimuli was manipulated by the experimenter so that subjects could initially be completely informed

about one option ('Option' trials) or completely informed about one attribute ('Attribute' trials). After viewing two stimuli subjects were allowed to either immediately make a decision based on the current state of evidence or gather as much more information as they wanted before responding. Subjects were still value driven in their choices but they were less optimal and the amount of information they chose to gather was influenced by the value of previous stimuli they had seen. In 'Option' trials they exhibited behaviour consistent with the presence of a confirmation bias where they would choose to gather information about a particular option even though it was clear that this option would be better than the alternative. In 'Attribute' trials their information gathering strategy also portrayed a confirmation bias because the location at which they chose to gather the third stimulus was driven by the value of the initial stimuli in a way that they would saccade towards the currently more valuable side. These results indicate that primate subjects are capable of both gathering information based on the current value of the known evidence in a manner biased by a need for confirmation and that they are also able to make decisions based on inferences about unknown information.

Chapters 5 and 6 will examine the neuronal correlates of the decision variables relevant to the multi-attribute decision making task described in Chapter 4. Specifically the former chapter examines neural firing during each sequential cue presentation and around the moment of response. At the first cue presentation neurons across all PFC regions encoded the generic value of the stimuli (i.e. without discriminating attribute or action properties). However, a specific and significant subpopulation of ACC neurons were seen to encode action value whereas a subpopulation of OFC neurons were seen to encode attribute value.

The 'Option' and 'Attribute' trial types provided a means for testing whether value computation in PFC regions were modulated by the location of attention. This is because at the second cue attention was oriented to the *same* side at the first cue in 'Option' trials whereas in 'Attribute' trials the second was on the *opposite* side of the first. It was observed

that value coding in OFC was indeed modulated by attention and that this modulation stemmed from a change in the sign of value coding of the *memory trace of previous cues*.

A projection analysis was used to compare how attentional, action and attribute value coding evolved over the course of the trial. OFC attentional value coding neurons were observed to reflect attentional choice with the relationship peaking immediately prior to response. ACC action value neurons were observed to evolve their coding over subsequent cues to no longer reflect action value but instead to reflect the final response. Finally the OFC attribute value neurons were observed to reflect attribute value coding at subsequent cues but also to reflect attribute specific choice immediately prior to the choice.

A further important element of task design was the fact that subjects were free to decide the location of the third piece of information that they gathered. Therefore it was possible to look for neuronal signals that predicted the location of information gathering (at the third cue) in each trial type. OFC attentional value neurons were observed to predict information gathering in both trial types in the frame of reference of whether information would be gathered in the same dimension as the currently attended cue or not. In 'Attribute' trials ACC action value neurons predicted the location (left or right) of the third saccade and in 'Option' trials OFC attribute value neurons weakly predicted the attribute to be viewed at the third cue. Finally, LPFC spatial position coding neurons were observed to predict the spatial location of future information gathering.

From the results presented in Chapter 5 it can be concluded that frames of reference in the PFC are not only region specific but also reflect generalised computations in PFC rather than those that are specific to certain situations. Furthermore, the fact that all three frames of reference are simultaneously represented and evolve to reflect choice implies that these signals do not act as mere inputs into a serial decision making pathway but instead take part in a parallel value comparison system. The fact that information gathering signals

are also frames by attention, action and attributes implies the existence of a common system of valuation which can drive both choice and information gathering.

Chapter 6 examines the neuronal activity at the moment of feedback onset. Many neurons across PFC distinguished whether the outcome was rewarded or not but the strength of outcome coding in ACC was proportional to the strength value coding in these neurons. In contrast probability preferring neurons in OFC were those that most strongly encoded reward/no reward (R/NR) while magnitude preference in LPFC reflected R/NR coding. Also LPFC action value coding neurons discriminated outcomes contingent on the actions that led to them. ACC value neurons were observed to encode both positive probabilistic prediction errors whereas both positive and negative prediction error coding was observed in vmPFC. In contrast both ACC and OFC neurons were observed to encode both probabilistic and magnitude fictive values when rewards were withheld. These results imply dissociable roles of PFC regions in learning where ACC and vmPFC reflect violations of expectations which may be critical for overall behavioural control In contrast, OFC neurons compute specific learning signals relevant for updating the value of more uncertain stimuli while LPFC neurons may be important for monitoring the value of actions.

The final chapter briefly summarises the salient results of this thesis and attempts to fit them into a broader understanding of the field of value based decision making. This thesis concludes that functional subdivisions of PFC perform vital, specific and unique value based computations during decision making in order to resolve choice, guide information seeking and learn about the environment. Through analysis of frames of reference I have shown that many of the abovementioned functions of PFC are actually subserved by the same valuation circuitry.

Chapter 1: The Functional Neurophysiology of the Prefrontal Cortex

In this chapter I will attempt to outline the current state of the decision making field with a particular emphasis on the findings of primate studies into both value and decision making computations. I will also briefly discuss two prevalent models of the decision making process discussing their similarities, differences and the predictions of their behaviour in example decisions.

PFC Lesions Suggest Functional Subdivisions

Our modern understanding of the role of the prefrontal cortex in complex behaviours and cognitive function can be traced back to the unfortunate case of Phineus Gage. At 4.30pm on 13th September 1848, Phineus Gage suffered an explosive accident which propelled an iron rod upwards through his left maxilla, causing an appalling insult to the medial and orbital parts of his frontal lobe. Miraculously, he survived this trauma and his recovery over the following days, months and years was carefully documented by Dr John M Harlow. Although many of the behavioural changes that Gage is best known for (such as his change in personality) were noticed months after his injury (Harlow, 1993), perhaps the most telling sign comes from a conversation between patient and doctor just 28 days after the accident (Harlow, 1848):

"Does not estimate size or money accurately, though he has memory as perfect as ever. He would not take $1000 for a few pebbles which he took from an ancient river bed where he was at work."

In this one sentence Harlow unknowingly provided vital insight into the critical function of PFC to estimate value and make decisions. Since then there have been multiple descriptions of patients with ventral and medial prefrontal cortex damage, who tend to make impulsive, risky and inappropriate decisions (Bechara et al., 1994, Manes et al., 2002, Camille et al., 2011a). Such observations have led to a wealth of interest in PFC in an attempt to understand the signals and computations that occur during the decision making process.

Importantly, value-based decision making deficits are typically present in human patients only when PFC areas are damaged (Kennerley and Walton, 2011). Among patients with PFC damage, those with damage to OFC, ACC and vmPFC typically show much more severe value-based decision making deficits than those with LPFC damage, though LPFC damage is arguably more often associated with cognitive and executive function deficits (Baxter et al., 2008, Bechara et al., 1998, Fellows, 2006, Fellows and Farah, 2005, Fellows and Farah, 2007, Petrides, 2005, Stuss et al., 2001). It should be noted however, that damage to PFC is associated with behavioural deficits in value-based, but not perceptual decision-making (Fellows and Farah, 2007, Manes et al., 2002). These observations suggest that there is functional segregation of decision making related computations across the brain, with PFC being particularly important for value-based decisions.

Indeed there has been a recent move to dissociate the functional subdivisions of PFC by the use of focal lesions. There has been a clear dissociation between observed deficits after ACC and OFC lesions. Both primate and human subjects with ACC lesions show deficits in action based decision making while those with OFC lesions show deficits in

stimulus based decision making (**Figure 1.1**, Rudebeck et al., 2008, Camille et al., 2011b, Kennerley et al., 2006). These findings will be discussed in greater detail in later sections. ACC lesions also specifically disrupt physical effort based decisions while OFC lesions affect decisions made based on delay effort (Rudebeck et al., 2006b). This further supports the idea that ACC is specialised to compute value in 'action space' whereas OFC does so in the abstract 'stimulus space'.

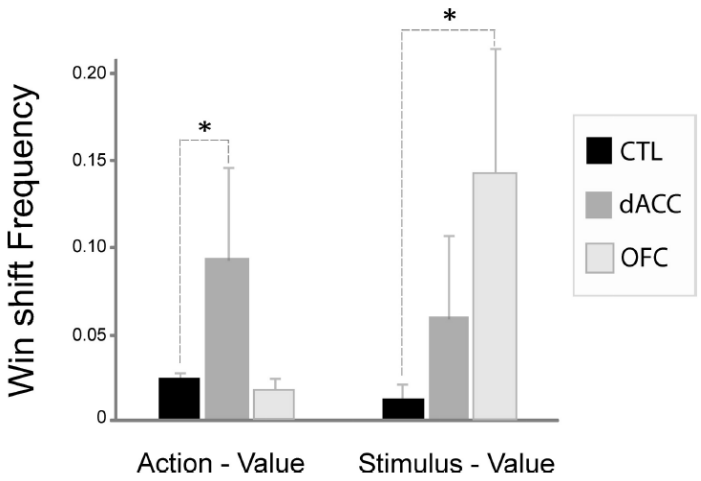

Figure 1.1: The double dissociation of action and stimulus based decision making between ACC and OFC. The frequency of shifting away from a rewarded response in an action and stimulus based reversal learning task for controls (CTL), ACC lesion patients (dACC) and OFC lesion patients (OFC). From Camille et al. (2011b).

Subjects with OFC lesions show deficits in contingent learning where they are unable to correctly assign the value of an outcome to the stimulus which preceded it (Walton et al., 2010). These subjects also exhibit altered fear responses (Izquierdo et al., 2005, Rudebeck et al., 2006a). Even within OFC there appears to be function specialisation. Lesions to medial portions affect optimal value comparison but damage to more lateral aspects cause deficits in credit-assignment (Noonan et al., 2010). Finally, damage to anterior cingulate gyrus (ACCg) causes noticeable deficits in social interest in other individuals (Rudebeck et al., 2006a).

A Framework for Decision Making Computations

In order to understand which regions perform critical computations during decision making we must first understand what unique signatures (either within a specific area or distributed across many) a decision process would exhibit. In their review Rangel et al. (2008) set out a framework for the types of decision computations would expect to see. They postulate that that the decision process involves the individual representation of various components of the decision (for example the hunger or specific decision variables), the integration of these variables in order to compute the value of each option, selecting an option/action based on competition between various valuations, evaluating the outcome and finally updating the representation, valuation and choice processes in order to optimise future decision making behaviour. In his 'goods based' model Padoa-Schioppa (2011) takes this basic framework and applies anatomical components in order to provide a plausible description of where in the brain each decision computation may occur. Although questions remain over whether this framework occurs in a serial or parallel manner (to be discussed in further detail in later sections) evidence for neuronal correlates of these signals will be considered in this chapter.

Encoding of Predictive Reward Signals within PFC

In the context of the wealth of lesion data showing decision making deficits, it is no surprise that neural correlates of reward, value and decision signals have been recorded throughout primate and rodent PFC and basal ganglia. Single neuronal recordings performed in 1970s described neurons in the anterior cingulate cortex (ACC) and lateral prefrontal cortex (LPFC) which modulated their firing based on whether or not the animal received a reward (Niki and Watanabe, 1976, Niki and Watanabe, 1979). Such value based signals have also been described in neurophysiological recordings in orbitofrontal cortex

(OFC) and ventromedial PFC (vmPFC) (Kennerley and Wallis, 2009a, Monosov and Hikosaka, 2012).

However, value based signals are not solely confined to PFC but have also been reported in areas such as the premotor cortex (Roesch and Olson, 2003), temporal areas (Baxter and Murray, 2002, Paton et al., 2006), visual cortices (Serences, 2008), parietal cortex (Platt and Glimcher, 1999, Foley et al., 2014), subcortical regions (Lau and Glimcher, 2008, Hikosaka et al., 2006, Kim and Hikosaka, 2013, Yasuda et al., 2012, Cai et al., 2011, Schultz et al., 1992) and even in electromyography signals when an animal is given a cue indicating the size of an upcoming reward (Roesch and Olson, 2003). This ubiquitous value signal may therefore reflect an aspect of motivation, arousal or attention, which can be difficult to distinguish from value (Kennerley and Walton, 2011, Roesch and Olson, 2004).

As described in the Rangel et al (2008) framework above, in order to make a decision, one has to represent the characteristics of options on offer. OFC seems to perform this function as it is capable of encoding the sensory properties of reward (Rolls and Baylis, 1994, Bouret and Richmond, 2010). It has been shown that OFC can linearly encode the value of rewards for individual decision attributes, including reward size (Kennerley et al., 2009, O'Neill and Schultz, 2010, Tremblay and Schultz, 1999, Morrison and Salzman, 2011, Rolls, 2000), reward probability (Kennerley et al., 2009), risk (O'Neill and Schultz, 2010), delay (Roesch et al., 2006), reward type (Hikosaka and Watanabe, 2000, Padoa-Schioppa and Assad, 2006) and effort (Kennerley et al., 2009, Kennerley and Wallis, 2009a). Several studies have shown that some OFC neurons orthogonally encode these decision attributes (O'Neill and Schultz, 2010, Morrison and Salzman, 2009, Roesch et al., 2006, Kennerley et al., 2009, Kennerley et al., 2011) although it has been shown that the representation of uncertainty and reward is integrated in OFC neurons (Raghuraman and Padoa-Schioppa, 2014).

Though OFC may tend to encode particular decision attributes as described above, owning to its strong sensory connections including gustatory input (Carmichael and Price, 1995b), it may be have an important role in encoding the integrated (i.e. subjective) value of options about juice rewards (Padoa-Schioppa and Assad, 2006, Raghuraman and Padoa-Schioppa, 2014). Although these findings may appear contradictory it is possible that these two populations are separate and have different computation roles in decision making. Neurons in OFC have been reported to code both pre-decision value computations such as 'offer value' (i.e. the representation of the options available) as well as post-decision signals such as 'chosen value' (i.e. the chosen option) (see **Figure 1.2A** and **B** for examples and further description) (Padoa-Schioppa and Assad, 2006). Such offer value coding is mostly absent from ACC neurons (Cai and Padoa-Schioppa, 2012). This implies that OFC performs critical decision computations, although it remains unclear whether the aforementioned finding reflects generalised decision computations or whether it is specific for the paradigm.

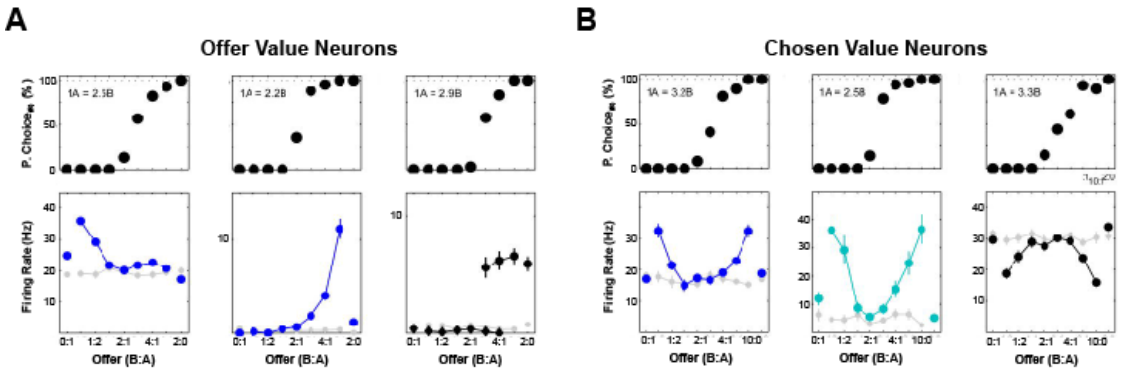

Figure 1.2: Pre-decision and post-decision computations in OFC. **(A)** The psychometric functions (top) and firing rates of three example offer value neurons (bottom). Each of these neurons reflects only the value of one of the juices on offer (i.e. either the value of juice A or B) through modulations in their mean firing rate. **(B)** The psychometric functions (top) and mean firing rates of three neurons which encode chosen value at various time points after cue presentation (coloured lines). Each neuron does not discriminate which juice was chosen but only the value of what was chosen. Adapted from Padoa-Schioppa and Assad (2006).

Like OFC, LFPC and ACC also encode different decision attributes and interestingly there does not seen to be any specialisation of decision attribute encoding within the three regions (Kennerley et al., 2009). However, a greater proportion of neurons in ACC tend to encode each decision attribute and also more often encode multiple decision attributes compared to both OFC and LPFC (Kennerley et al., 2009, Kennerley et al., 2011). This suggests that ACC might be integrating information about individual decision attributes into a single common value signal (Hosokawa et al., 2013). This is further backed up by the finding that ACC neurons can also encode integrated chosen value (Cai and Padoa-Schioppa, 2012). There is evidence that ACC neurons can also encode integrated information about value and other task relevant variables (Hayden and Platt, 2010, Amemori and Graybiel, 2012). Finally some evidence suggests that, in contrast to OFC, ACC only computes post-decision variables (Cai and Padoa-Schioppa, 2012), which might suggest that ACC sits further downstream in the decision pathway than OFC allowing it to integrate information multiple input areas.

As well as being invariant to the type of value information on offer, it has also been shown that neuronal firing in OFC is invariant to the menu on offer (Padoa-Schioppa and Assad, 2008, Tremblay and Schultz, 1999). This is to say that OFC neurons will encode value when choosing what to drink, but many of the same neurons will also encode value when deciding where to go on holiday. Central to this finding is the observation that some OFC neurons perform range adaptation based on the context of the choice (Padoa-Schioppa, 2009, Kobayashi et al., 2010). For example, OFC neurons with use the same range of firing to represent £1 to £5 in one context as is used to represent £100 to £1 million pounds in another context. Range adaption is important because it allows for maximal discrimination between different values on offer (Kennerley and Walton, 2011). However, other neurons in OFC do not show range adaption which may be vital in understanding the absolute value of options (Kennerley and Walton, 2011, Kobayashi et al., 2010). Range adaptation has also been described in ACC, ventral striatum and the dopamine system

implying that this functional property of some reward sensitive neurons may be vital for effective decision making (Cromwell et al., 2005, Tobler et al., 2005, Sallet et al., 2007).

Similar to the idea of range adaption, it has been suggested that value sensitive neurons exhibit context dependent modulation based on the value of other stimuli presented in the environment (i.e., value normalization). This relative valuation mirrors the idea of receptive fields in the sensory system and can explain some irregularities in subject choice behaviour (Louie et al., 2013). Furthermore, neuronal responses in LIP have been shown to exhibit relative value coding in simple saccade tasks (Louie et al., 2011).

Assigning Values to Stimuli Versus Actions

Making optimal value based decisions depends on the ability of organisms to learn about the environment. Animals and humans with OFC and vmPFC lesions are more erratic when presented with subjective choices between foods, suggesting that they are incapable of assigning a consistent value to these foods (Baylis and Gaffan, 1991, Camille et al., 2011a). The same animals also show deficits in a basic stimulus guided decision making task suggesting that they are also unable to form a stimulus-outcome association (Baylis and Gaffan, 1991).

Evidence also suggests that this stimulus-outcome learning may be specific to subregions of the orbital surface. Lesions to medial parts of OFC (including vmPFC) cause deficits in reward-guided decision making but not in learning and updating stimulus-outcome associations (Noonan et al., 2010). However, the same study showed primate subjects who received damage to more lateral parts of OFC appeared unable to correctly assign credit to rewarded outcomes in changeable environments. This implied that neurons in lateral OFC were critical for updating stimulus-outcome relationships. Furthermore, in three-arm bandit tasks, although OFC lesioned animals are initially capable of tracking the best stimulus

(irrespective of its value), when it changes value these subjects show an inability to track the new best stimulus (Walton et al., 2010). This effect came about because the OFC-lesioned animals were influenced less by recent stimulus-outcome associations than control animals.

OFC's function in forming stimulus-outcome relationships is not surprising in the context of its strong reciprocal connections to the amygdala, which has also been implicated in linking stimuli to outcomes in both lesion and neurophysiological experiments (Baxter and Murray, 2002, Carmichael and Price, 1995a, Paton et al., 2006, Hirai et al., 2009, Izquierdo et al., 2004, Ghashghaei et al., 2007). Furthermore, lesions to the amygdala cause observable changes to the functional properties of OFC neurons, such as attenuating OFC sensitivity to value as well as outcomes (Rudebeck et al., 2013a, Schoenbaum et al., 2003). Several studies have shown that OFC responses change rapidly to reflect new stimulus-outcome relationships (Tremblay and Schultz, 2000a, Schoenbaum et al., 1998, O'Doherty et al., 2003).

A natural follow-up question to these observations is; is it actually the amygdala that is critical for learning stimulus outcome relationships? The lesion data would suggest not, as lesions to OFC cause stimulus-outcome relationships in the amygdala to become inflexible, thereby causing reversal learning deficits (Saddoris et al., 2005, Stalnaker et al., 2007). Furthermore, in reversal learning tasks, OFC activity can be seen to encode the value of positively rewarded stimuli with a faster latency than in amygdala activity, and OFC exerts a disproportionate influence on the amygdala after learning compared to before learning (Morrison et al., 2011).

However, a critical finding by Rudebeck et al. (2013b) has shown that OFC grey matter tissue may not be critical for supporting reversal learning. These authors compared the effects of OFC aspiration versus excitotoxic lesions, where only the latter spares damaging the underlying white matter. They found that reversal learning deficits were only present in the aspiration lesion group, suggesting that rather than OFC itself, certain white

matter pathways which run close to OFC may be critical for updating stimulus-outcome relationships (Rudebeck and Murray, 2011, Rudebeck et al., 2013b). Therefore it remains ambiguous whether the encoding of stimulus-outcome relationships is an essential function of OFC neurons.

In contrast to OFC's role in stimulus-outcome associations, ACC has been shown to be important in action-outcome associations and damage to this area causes an inability to sustain rewarded actions (Ostlund and Balleine, 2005, Kennerley et al., 2006, Rudebeck et al., 2008, Hadland et al., 2003). Human lesion experiments have also shown similar results with OFC lesion patients relatively less able to form stimulus-outcome associations compared to ACC lesion patients whereas ACC lesion patients are slower to learn action-outcome associations when compared with OFC patients (Camille et al., 2011b). Furthermore, neurons both ACC and cingulate motor area have been observed to encode action based prediction error signals which are consistent with the learning and updating of action-outcome relationships (Shima and Tanji, 1998, Matsumoto et al., 2007).

Value and Actions

Action can be driven by decision making processes in two potential ways. Firstly, action plans can be generated post-decision once the chosen goal is known by the motor system. In this hypothesis the motor system does not actually play any active role in decision making but instead simply acts as a conduit for an abstract decision system to act on the environment (see **Figure 1.4A**) (Padoa-Schioppa, 2011). The other explanation is that decisions can be made in terms of actions through a competition between pools of neurons representing each action and the value associated with it (see **Figure 1.4B**). Such action value coding neurons have been observed in several brain regions including ACC (Hayden and Platt, 2010, Matsumoto et al., 2007, Luk and Wallis, 2009, Matsumoto et al., 2003), LPFC (Kim et al., 2008), dorsal premotor cortex (PMd) (Pastor-Bernier and Cisek, 2011), the

basal ganglia (Lau and Glimcher, 2008) and LIP (Platt and Glimcher, 1999). Of these regions, ACC and LFPC are of particular interest to value based decision making for several reasons. Firstly, as previously discussed, subjects with ACC lesions show specific deficits in action based decision making tasks (Camille et al., 2011b, Rudebeck et al., 2008, Kennerley et al., 2006, Hadland et al., 2003). Also neurons in LPFC have been observed to change their coding patterns over the course of a decision from action value to choice implying the existence of a decision process (Kim et al., 2008). It is possible that other areas take part in action selection computations because PMd neurons have been observed to code action values which are competitively modulated by other available values and actions thereby exhibiting many of the necessary properties of an action selection circuit therefore (Pastor-Bernier and Cisek, 2011). These finding may support the idea of parallel decision making competitions even within the domain of action selection. In contrast to ACC and LPFC, neurons in OFC very rarely encode the action value or the final response during the choice (Wallis and Miller, 2003, Padoa-Schioppa and Assad, 2006, Cai and Padoa-Schioppa, 2014).

An important aspect of selecting between actions is the computation of the physical effort involved for each action. ACC lesions lead to effort averse decision making in rodents (Rudebeck et al., 2006b). Furthermore, fMRI studies show that BOLD signal in ACC encodes effort discounted value (Prevost et al., 2010, Croxson et al., 2009).

However, not all decisions are made in terms of actions, and ACC and LPFC both may play a role in transforming value signals into response signals in these cases. During decision making, ACC neurons have been shown to code not just value but also response and the value x response interaction implying a transformation of value to action in this region (Kennerley et al., 2009). When subjects are forced to make abstract decisions before mapping the decision to an action this effect is also observable in LPFC (Cai and Padoa-Schioppa, 2014). Both ACC and LPFC are well positioned to facilitate value to action computations because of their extensive links to both critical PFC structures such as OFC,

and their link to saccadic and limb motor structures (Morecraft and Van Hoesen, 1992, Luppino et al., 2003, Matelli et al., 1986, Ongur and Price, 2000).

vmPFC, Internal State and Value Comparison

There is a relative paucity of electrophysiological data surrounding the function of vmPFC. One of the first single neuron reports of vmPFC implicated neurons here in the encoding of subjective value with reference to internal state. Its neurons specifically encode states such as satiety or fatigue (Bouret and Richmond, 2010) and furthermore the coding of value in vmPFC is modulated by current motivation (i.e. internal state) (Abitbol et al., 2015). This finding may be explained by vmPFC's connections to medial hypothalamus, which is known to be important for food intake and metabolic balance (Berthoud, 2003, Leibowitz, 1988, Ongur and Price, 2000).

Like many regions throughout the brain, vmPFC neurons encode the value of reward predictive cues (Monosov and Hikosaka, 2012). vmPFC neurons also appear to exhibit an anterior-posterior axis in functional coding, where more anterior neurons (Brodmann's area 14) code specifically for appetitive stimuli whereas neurons in posterior portions (Brodmann's area 25) specifically encode the value of aversive stimuli (Monosov and Hikosaka, 2012). Related medial-versus lateral anatomical claims about rewards and punishers, respectively, have been made about OFC, though the evidence is inconsistent (Rich and Wallis, 2014, Morrison and Salzman, 2011, O'Doherty et al., 2003, Hayes and Northoff, 2012, Liu et al., 2011, Fujiwara et al., 2008).

But perhaps most importantly, vmPFC, at least from human studies, has been implicated in value comparison processes. Several human studies have demonstrated value comparison signals in vmPFC during decision making (Basten et al., 2010, Philiastides et al., 2010), particularly signals reflecting the difference between the chosen and unchosen value

(Chau et al., 2014, Hunt et al., 2012). Such value comparison signals are key component of activity measured from plausible biophysical models of decision circuits (Hunt et al., 2012). In sequential decision making tasks, vmPFC neurons are observed to compute the difference in value between the two options as well as representing chosen value (Strait et al., 2014). However, it is unclear whether this value comparison signal occurs in a temporal frame of reference (i.e. the first option versus the second option) or in an attentional frame of reference (i.e. the option currently attended to versus the unattended option). Data from the human fMRI field would lead us to believe that the latter comparison is the more likely. In a heavily constrained economic decision making task, Lim et al. (2011) showed that the vmPFC BOLD signal correlated positively with the value of the attended item and negatively with that of the unattended item, with the inference that attention may be relevant for the comparison processes. The importance of attention in learning, valuation and value comparison processes will be further explored in Chapters 3-5.

Frames of Reference in PFC

At this point I will touch upon a concept that has been implied in several of the results that I have discussed above; that is the idea of frames of reference for valuation in the PFC. The three well established frames of reference in which value is often associated within PFC are that of attribute, actions and attention. The attribute frame of reference is based on the observation that some neurons in OFC specifically code for single decision attributes (Kennerley et al., 2009, Morrison and Salzman, 2011, O'Neill and Schultz, 2010) (see **Figure 1.3B** for an example neuron with description) and that damage to OFC disrupts decision making when a decision attribute (e.g., reward amount or probability) is associated with a stimulus (Fellows, 2006, Camille et al., 2011b, Rudebeck et al., 2008). The idea of the action reference frame comes from the observations that ACC (Hayden and Platt, 2010, Matsumoto et al., 2007) (see **Figure 1.3A** for an example), LPFC (Kim et al., 2008), basal

ganglia (Lau and Glimcher, 2008, Hunt et al., 2014) and parietal cortex (Dorris and Glimcher, 2004, Platt and Glimcher, 1999) all encode the predictive value of actions. Furthermore, as previously discussed, damage to ACC typically produces bigger deficits is action-based decision making (Rudebeck et al., 2008, Camille et al., 2011b), though subtle deficits have been reported with ACC dysfunction even in stimulus-based decision tasks (Kennerley and Walton, 2011, Amiez et al., 2006).

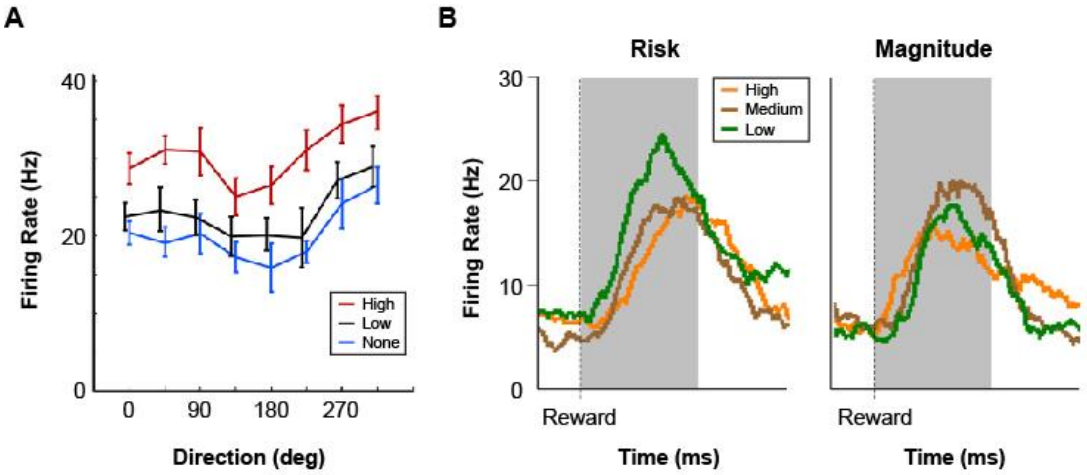

Figure 1.3: Action and attribute value coding neurons. **(A)** An example action value in ACC. Each point indicates the mean firing rate of the neuron for each combination of value and direction of saccade required to obtain reward. Adapted from Hayden and Platt (2010) **(B)** An example attribute value neuron in OFC. The left panel indicates the neuron's firing rate modulation to risk information, and the right panel the firing rate to magnitude information. This neuron responds only to risk and not to magnitude value. Adapted from O'Neill and Schultz (2010).

Finally, the attentional frame of reference has been described in the human vmPFC and ventral striatum BOLD signal (Lim et al., 2011) and suggested in vmPFC neuronal firing (Strait et al., 2014). This reference frame implies that value signals are tied to what the subject is directly attending to, and recent models of decision-making have integrated this idea in the competitive process toward decision formation (Krajbich et al., 2010, Krajbich and Rangel, 2011).

The idea of frames of reference coding is central to the concept of parallel decision making systems (which will be discussed at greater length at a later point in this chapter). By representing value in these frames of reference, it may allow decisions to be evaluated in parallel and choices potentially determined in individual frames of reference, rather than requiring single regions to integrate all relevant information/attributes into an integrated value of each option. Furthermore this could provide an explanation as to why very specific value based decision making deficits are reported in ACC and OFC lesioned subjects (Rudebeck et al., 2008, Camille et al., 2011b). Evidence for frame of reference specific value comparison remains relatively unidentified, however a study by Hunt et al. (2014) showed that action and stimulus based value comparison was observed in the dorsolateral striatum and OFC respectively. They also found changes in the functional coupling between these regions and parietal cortex which correlated with behavioural measures of which frame of reference was most influencing the choice on each trial.

Serial vs Parallel Value Comparison

A prevailing debate surrounding the issue of decision making in the brain revolves around whether decisions are made through a single serial pathway or parallel competitions across the brain. One of the most influential models of a serial decision process is the 'goods based model' of Padoa-Schioppa (2011) (**Figure 1.4A**). This model proposes that decisions are made through comparisons of abstract integrated values for various potential options.

These integrated values arise from the integration of many intrinsic and extrinsic factors such as thirst, satiety, effort, risk and commodity (to name but a few). The integration and comparison is said to occur in either OFC or vmPFC, based on the evidence that shows both abstract coding and integrated value signals in both regions (Padoa-Schioppa and Assad, 2006, Plassmann et al., 2007, Kennerley et al., 2009). An important aspect of the

model is that only after making a decision in abstract 'goods' space can the choice be converted into a relevant action (hypothesised to be in either LPFC or ACC) leading to action selection. Such a value to action signal has indeed been observed when decisions are made in abstract terms (Cai and Padoa-Schioppa, 2014). However this does not sit squarely with the evidence that when decisions and actions links are predefined, there is clear evidence of motor planning before decisions are completed (Selen et al., 2012).

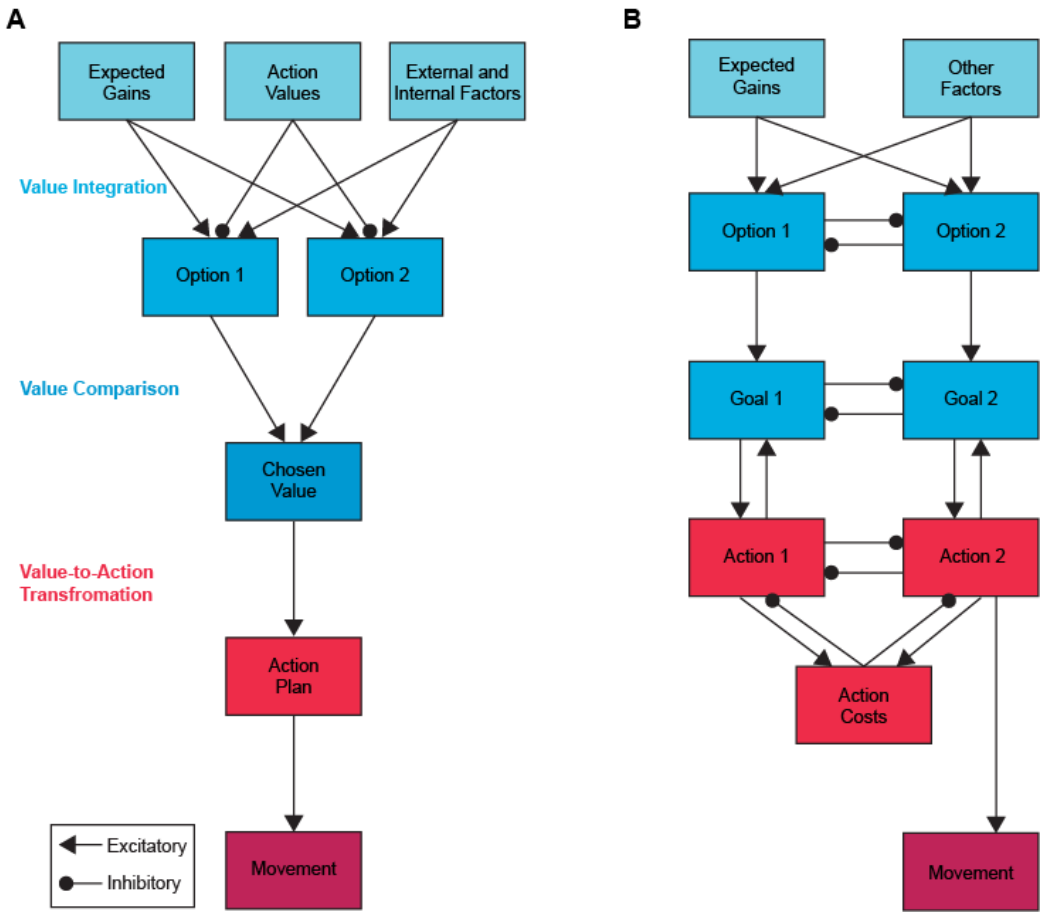

Figure 1.4: Serial versus parallel decision making pathways. **(A)** A schematic representation of the 'goods based model' of decision making (Padoa-Schioppa, 2011). Each box represents a computation performed and each line indicates the direction and the type of interaction (i.e. excitatory or inhibitory) between computations. **(B)** A schematic representation of the 'decision through consensus model'. The important feature of this model is that options are represented in several manners (in this case as options, goals and actions) and competitions occur at the level of all of these representations through mutual inhibition. Adapted from Cisek (2012).

In contrast one of the leading theories relating to the parallel decision making hypothesis is the idea of 'decision making through consensus' by Cisek (2012) (**Figure 1.4B**). This model, unlike the 'goods based model' does not dictate a single region of value comparison but instead states that this process can occur throughout cortex and subcortical regions. Furthermore, central to this hypothesis is the idea that value comparison could occur in different frames of references such as those of action, stimulus, internal/goals or subjective value (Cisek, 2007a). The representations at any level would compete in decision circuits through mutual inhibition. Such a model is theoretically appealing because it provides a plausible hypothesis as to how action costs might be integrated into a decision process without transforming them into abstract values. This model also allows for both habit- and goal-directed influences (commonly associated with the basal ganglia and frontal cortex, respectively) to operate simultaneously (Dolan and Dayan, 2013), and provides a flexible framework as to how decisions of various forms may be solved (Cisek, 2012). Furthermore, such a model of decision making reflects what is known about brain architecture, such as parallel information flow through the dorsal and ventral visual streams (Ungerleider, 1982).

Cisek's model also goes some way to explaining why decision variables are represented in classically sensorimotor regions of the brain (Hernández et al., 2010, Roitman and Shadlen, 2002, Thevarajah et al., 2009, Pastor-Bernier and Cisek, 2011). Furthermore, several studies have demonstrated that during perceptual decision making, although there is a flow of information from sensory areas to frontal regions and then on to motor regions, simultaneous comparison signals are observable in many regions of parietal, frontal and temporal lobes, implying that there are instead multiple levels of comparison (Siegel et al., 2015, Hernández et al., 2010).

Importantly, there are several different predictions that the goods based model and the decision through consensus model provide that are testable in the context of decision making. The first and most obvious of these is that the goods based model stipulates that

response and other post-decision signals should only be observable *after* the decision has been made. In contrast, the parallel processing hypothesis predicts that action selection (i.e. response selection) should occur *simultaneously* with abstract goods selection. Next, while the goods based model predicts that the only observable value comparison and choice signals should be in abstract 'goods' space, the decision by consensus hypothesis predicts comparisons in multiple frames of reference and corresponding choice signals for each frame. Finally the 'goods based model' hypothesises that value comparisons only occur in one brain region (OFC/vmPFC) (Padoa-Schioppa, 2011). Therefore, every region downstream of that (i.e. ACC, LPFC and motor structures) should only reflect post-decision computations. Chapter 5 will test, and arguably refute, many of these hypotheses of the goods-based model.

Representing Outcomes in PFC

In order to understand our world we must be able to learn the value of our environment. Central to achieving this is being able to represent outcomes that we experience in the context of what choices led to those outcomes. Neurons in ACC, LPFC and OFC have been shown to fire in the presence and absence of reward (Kennerley et al., 2009, Seo et al., 2007, Tremblay and Schultz, 1999, Kennerley et al., 2011). However, beyond the simple discrimination of outcomes, neurons in both ACC and OFC have been reported to encode reward prediction errors (i.e. when outcomes are better or worse than expected) (Kennerley et al., 2011, Sul et al., 2010, Matsumoto et al., 2007, Amiez et al., 2006, Seo and Lee, 2007). **Figure 1.5** shows an ACC positive prediction error coding neuron. This neuron responds positively with increasing chosen value during the choice epoch (**Figure 1.5A**), and responds negatively with chosen value at feedback on rewarded trials (**Figure 1.5B**), which reflects positive prediction error coding.

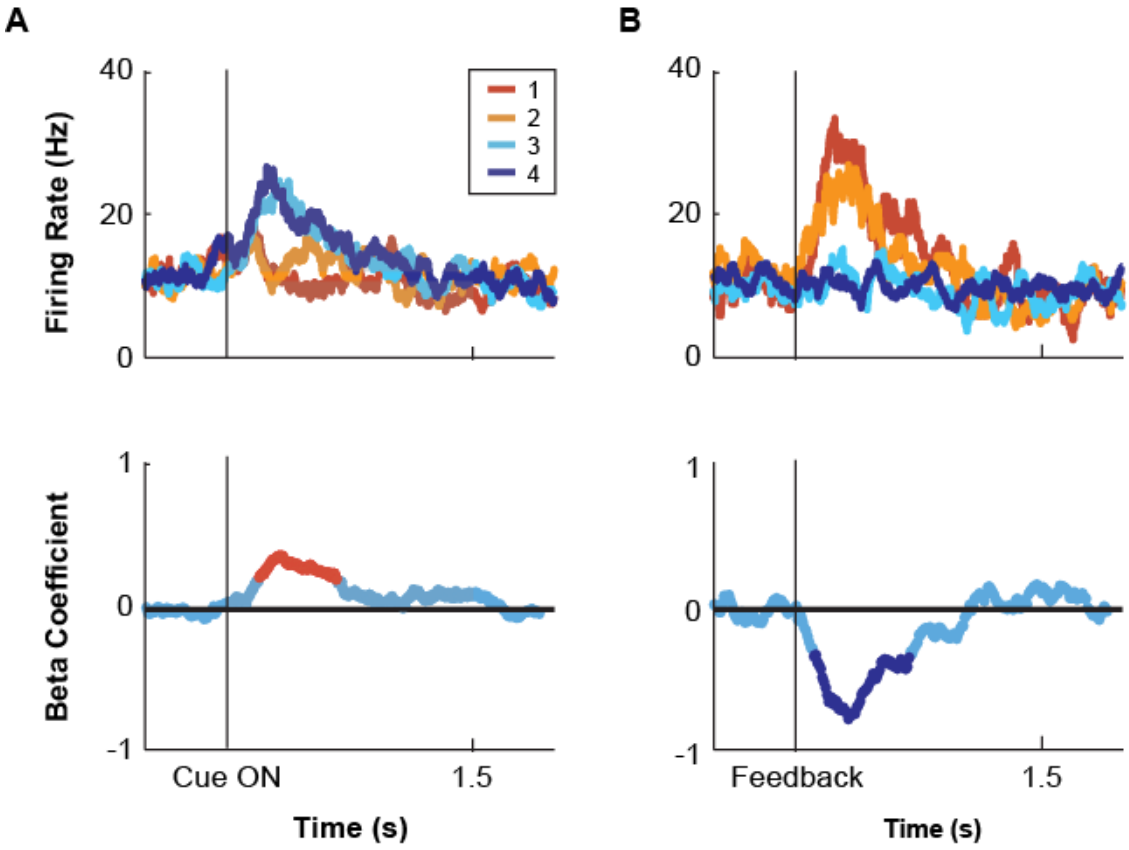

Figure 1.5: Positive prediction error coding in ACC. **(A)** The firing rate of an example ACC neuron at stimulus presentation (vertical line) for various chosen values (coloured line) (top panel) and the accompanying linear regression of firing rate and value (bottom panel). Red dots indicate significant bins and blue non-significant bins. **(B)** The firing rate of the same neuron at feedback onset (vertical line) on rewarded trials for each chosen value (top panel) and the accompanying linear regression (bottom). Blue dots indicate significant bins. Adapted from Kennerley et al. (2011).

Prediction error coding in ACC and OFC is not unexpected given both regions' have connections with the ventral tegmental area (VTA), a region containing dopamine neurons where classical reward prediction errors have been described (Schultz, 1986, Berger et al., 1988, Williams and Goldman-Rakic, 1993). Neurons in ACC, LPFC and OFC also encode the magnitude of a reward given (Rolls, 2000, Kennerley et al., 2009, Watanabe, 1996, Roesch et al., 2006). Such a ubiquitous encoding of outcomes in ACC conforms to results seen in human EEG studies where "feedback related negativity" over ACC-related electrodes has been recorded when subjects are given outcomes indicating monetary losses or wins (Wu and Zhou, 2009, Hajcak et al., 2005). Recent thinking has postulated that ACC

prediction error coding may reflect the computations of an action-outcome predictor, whose function is to control complex behaviour (Alexander and Brown, 2011).

As learning is predicated on using outcomes to inform future choices, a natural question is the relationship between neuronal activity on current and past choices. At the time of choice, OFC chosen value coding is influenced by the recent history of choice offers, whereas ACC neurons appear to encode chosen value of the current trial and the amount of reward received on the previous trial (Kennerley et al., 2011). This OFC result may be a variant of relative or adaptive coding as discussed previously (Tremblay and Schultz, 1999, Padoa-Schioppa, 2009). Neurons in ACC also encode previous choices and reward history (Seo and Lee, 2007). Reward history coding is seen to be heterogeneous across ACC neurons with neurons exhibiting different time constants (Bernacchia et al., 2011).

In a stimulus-guided strategy task, it has been observed that after the reward outcome, neurons in OFC encode the response made by the animal irrespective of whether it was rewarded (Tsujimoto et al., 2009). Such a signal may be useful in allowing the brain to understand the relationship been actions and outcomes which is critical for dynamic learning (Tsujimoto et al., 2012). However such a conclusion remains incongruent with the fact that lesions of OFC do not affect the learning of action-outcome associations and the fact that OFC itself rarely encode action values before the point of choice (Rudebeck et al., 2008, Wallis and Miller, 2003). One possibility is that OFC response selectivity in the Tsujimoto et al., 2009 study actually reflects a signal critical for planning the next trial's choice, since the experimental design required a win/stay, lose/switch response strategy dependent on whether the current response was rewarded.

Behavioural Flexibility and ACC

As well as its role in encoding upcoming reward, it has been suggested that ACC is critical to outcome mapping and behavioural flexibility (Hayden and Platt, 2010, Kerns et al., 2004). For example, in foraging tasks, neuronal firing in ACC correlates with the relative value of foraging compared to exploiting a resource (Hayden et al., 2011). Signals in ACC also correlate with volatility (Kolling et al., 2012), behavioural errors (Falkenstein et al., 2000), the speed-accuracy trade-off (Yeung and Nieuwenhuis, 2009), behavioural switching (Quilodran et al., 2008, Johnston et al., 2007) and conflict (van Veen et al., 2001). Furthermore, there is clear evidence that ACC neurons (among other regions) encode both positive and negative prediction errors in both stimulus based (Kennerley et al., 2011) and action based tasks (Matsumoto et al., 2007). Although these cognitive processes seem disparate, recent thinking has put forward the idea that the above described processes are reconcilable as phenomenon of a generalised action-outcome predictor model (Alexander and Brown, 2011). The basic concept of this model revolves around the idea that ACC neurons predict the expected outcome of a given action and respond when the expectation is not met. This prediction error is then used to update the prediction of the outcome for future behaviour. A potentially important implication of this model is that if true, ACC value coding during the choice phase of a decision may actually reflect a prediction about an outcome, rather than a computation immediately relevant to making choice.

Another influential unifying idea of ACC function is that of expected value of control. In this theory ACC computes the overall value of allocating ‘control’ resources to a given problem (Shenhav et al., 2013). This computation may be relevant for deciding which tasks or actions to engage in at any given time.

Executive Function in LPFC

As stated earlier, LPFC is known to encode value; however lesions to this area do not cause severe decision making deficits (Kennerley and Walton, 2011, Kobayashi et al., 2006). This has led to the suggestion that LPFC may be important in guiding attention towards behaviourally relevant information rather than the decision process itself (Rushworth et al., 1997, Buckley et al., 2009, Funahashi et al., 1989, Lebedev et al., 2004). For example in a spatially guided saccade task, it has been shown that ventral parts of LPFC (VLPFC) encode both the size of the reward on offer and the spatial location of the saccade required to obtain the reward (Kennerley and Wallis, 2009c). In contrast, OFC and ACC do not encode spatial information in the same task (Kennerley and Wallis, 2009a). LPFC neurons have been observed to encode locations of attention rather than a simple reflection of working memory (Lebedev et al., 2004). These two factors are often correlated in working memory tasks (for example. Funahashi et al. (1989)).

LPFC function has been linked to executive control (Mansouri et al., 2009, Miller and Cohen, 2001, Tanji and Hoshi, 2008, Tsujimoto et al., 2012). In support of this, neuronal firing within LPFC correlates with task relevant rules and task context (White and Wise, 1999, Hoshi et al., 2000, Asaad et al., 2000, Hoshi et al., 1998), planning (Collins et al., 1998, Gaffan et al., 2002) and goal selection (Saito et al., 2005). Finally LPFC activation is also noted when behavioural adjustments are implemented to maintain optimal behaviour (Egner and Hirsch, 2005a, Egner and Hirsch, 2005b, Kerns, 2006). These results imply that LPFC enjoys a more general function in cognition, potentially relating to allocating attentional resources and control, rather than a specific role in decision making per se. Such control functions may nonetheless be critical for decision-making processes, such as in prioritizing particular attributes or valuation processes in other brain regions (Hare et al., 2009).

Information Gathering Strategies

Gathering information is a critical feature of complex decision making. Without adequate information decisions are difficult to make, and outcomes are highly uncertain. Human information gathering behaviour has been characterised by psychologists and economists for several decades. Predominantly this has revolved around describing the manner in which humans choose to compare individual pieces of information when presented with multi-option and multi-attribute decisions. When the decision space is small (i.e. few options and few attributes to consider), subjects tend to look at most available information and compare information in terms of options (Payne, 1976, Cook and Swain, 1993, Lohse and Johnson, 1996). However, as the complexity of the task increases, humans shift towards comparing information across individual attributes and also begin to 'screen' options by using attributes of importance to exclude options (Sundström, 1987). The established thought is that this process of screening is a method by which subjects can reduce cognitive demand by reducing the total amount of information that needs to be considered (Weenig and Maarleveld, 2002, Kerstholt, 1992). This idea is also supported by the fact that subjects typically become less accurate and also gather a smaller amount of the total available information as task complexity increases (Kerstholt, 1992). Other manipulations of task complexity such as increasing decision difficulty have also been demonstrated to shift subject information comparison from that of 'option based' to 'attribute based' (Arieli, 2011).

Neurobiological studies of decision making typically bypass this point by simplifying decisions down to one attribute (e.g. Kennerley et al. (2009)) or conversely merge multiple attributes into single stimuli (e.g. (Hosokawa et al., 2013)). However, some inferences about the brain structures that may support information gathering can be drawn from human lesions studies. Patients with vmPFC damage exhibit a pattern of behavioural deficits which reflects a change in the way they gather information relative to controls. In a study by Leslie

Fellows (2006), three groups of subjects (vmPFC lesion patients, LPFC lesion patients and controls) were asked to choose an apartment in which they would like to live. Each subject was presented with several potential apartments. Each apartment was associated with a variety of decision attributes (e.g. its size, cost and location), however this information was initially covered up. Subjects were instructed to uncover any information they wanted, and that they were free to choose whenever they felt they had sufficient information. By allowing subjects to freely sample information, the authors could investigate whether there were any information gathering strategies present in the groups. LPFC lesion and control subjects tended to use the well described 'attribute based' information gathering strategy, in which they preferred to uncover the same attribute information for each apartment before moving on to a new attribute to uncover. However, patients with vmPFC damage tended to compare 'across options' by uncovering all attribute information for a single apartment before moving on to the next apartment. vmPFC patients also often picked different apartments at the final choice, even though vmPFC patients collected the same amount of information as the other groups before making their choice. An interpretation of these results could be that vmPFC is important in biasing behaviour towards attributes that are more relevant at that particular time. However, it should be noted that this result does not necessarily imply that vmPFC is critical for guiding information gathering; it is possible that the lesion subjects shifted their information gathering strategy in this task because they were unable to compare information within attributes due to destruction of attribute specific neurons in OFC.

Gambling paradigms have also shown that human vmPFC lesion patients are more likely to make risky choices irrespective of the odds of winning (Clark et al., 2008). In the context of Fellows' findings, this risk seeking behaviour could be the result of an inability of these patients to bias their behaviour towards the probability attribute of the gamble.

Open Questions

Despite the wealth of neurophysiological findings about value and the PFC, one of the great deficits in knowledge is the mechanism of choice. Although plausible biophysical models based on neuronal data exist which give a mechanistic account of perceptual decision making (Wang, 2008), analogous single neuronal data in support of value-based decision mechanisms remains elusive (but see Hunt et al. (2012), Chau et al. (2014), Strait et al. (2014)). Arguably, one reason why our understanding of value-based decision mechanisms is limited can be traced to classical experimental design approaches in many electrophysiological studies. For example, most decision making tasks are designed with a choice epoch of a fixed duration (e.g., 1000ms) such that even if subjects make a decision quickly, they must wait for the imperative 'Go' cue before making a decision (Hosokawa et al., 2013, Kennerley et al., 2009, Seo et al., 2007, Cai and Padoa-Schioppa, 2012, Cai and Padoa-Schioppa, 2014, Padoa-Schioppa and Assad, 2006). Many human fMRI studies have similar methodological constraints due to delays in the hemodynamic response. Such constrained tasks (as opposed to allowing subjects to choose freely as soon as they decide) make it impossible to determine exactly *when* the decision was made. Such imposed delays may also make it difficult to dissociate decision processes from working memory processes. Furthermore, most electrophysiological studies require animals to maintain central fixation throughout the choice epoch, which makes it difficult to understand the subject's information gathering strategy and the importance of attention of valuation processes. Therefore a task that allows an unconstrained reaction time and eye movements may yield an advance in our understanding of decision strategies and dynamics.

As discussed in this chapter, several studies have observed frame of reference-specific coding in PFC. However the role of these frames of reference in decision making remains unclear because it is unknown whether neuronal computations reflecting these reference frames are simply a by-product of the decision making paradigm used, or whether

neurons may indeed have preferred reference frames if the task incorporated multiple reference frames (Hunt et al., 2014). Furthermore, it is unclear whether these frames of reference are simply inputs into a serial decision system which could be congruent with the 'goods based model', or whether PFC circuits compare value in these different reference frames. This open question has implications for understanding the debate between serial and parallel processing in decision-making because the former model dictates that value comparison only occurs in abstract goods space whereas the parallel processing model predicts comparison and choice signals in all frames of reference.

An extremely under-sampled aspect of decision making is information gathering. Although human information gathering behaviour is well characterised, little is known about how other animals resolve this problem. Furthermore, despite the lesion evidence linking vmPFC in information gathering strategies, and the suggestion that vmPFC valuation processes may be linked to attention (Lim et al., 2011), it remains unknown how vmPFC and other PFC neurons support information gathering processes (Fellows, 2006). Understanding this may help us understand whether any link exists between high level cognitive control and general decision-making processes in PFC.

Finally, our understanding remains limited with respect to what information neurons compute when outcomes are delivered. Arguably, to learn optimally, one needs to know exactly what aspects of a choice outcome (i.e., attributes) deviate from the prediction. Although it is known that ACC and OFC neurons compute prediction errors at outcome when only reward probability is considered, it is unknown whether these prediction error computations are specific to certain attributes of value, or whether they instead reflect the difference between the expected and experienced *integrated* value of the outcome. Understanding the relationship between activity at choice and outcome may offer a better understanding of the role of PFC in value-guided choice and behavioural control mechanisms.

Conclusions

This review of the decision making literature has shown that there is direct lesion evidence that ACC, vmPFC and OFC are critical for the processes necessary for normal value-based decisions. The functional segregation between these three regions may reflect frame of reference specific decision computations which will be explored in this thesis. Furthermore, areas such as vmPFC and LPFC may be vital for supporting decision making behaviour through behavioural biasing, information gathering and attentional control.

Models of the decision making process are fundamentally centred around the concept of serial versus parallel decision making processes. Each model therefore makes several unique testable predictions about decision related computations in various parts of the brain. This question currently remains relatively open in the context of value guided decisions.

Chapter 2: The Anatomy of the Macaque Prefrontal Cortex

The anatomical connections of the PFC constrain and define the types of computations that it is capable of performing during decision making (Passingham et al., 2002). The PFC makes up the vast majority of the frontal cortex of both humans and primates. It consists of a multiple heterogeneous Brodmann's areas located on the lateral, orbital and medial surfaces of the brain and also shows large amounts of variation between primates and lower mammals (Paxinos and Mai, 2004, Wise, 2008, Fuster, 2001). Brodmann's areas within PFC are often themselves subdivided into functional areas (**Figure 2.1**). This section will provide an overview of the anatomy of the PFC of primates with particular attention paid to the intrinsic and extrinsic connections.

Although the PFC is made up of many Brodmann's area which all have varying patterns of inputs and outputs, these individual divisions are not generally considered when electrophysiological and lesion studies are undertaken. Instead, more general areas are commonly defined which span across Brodmann's areas. In keeping with this convention I will now define four regions that are relevant to the work in this thesis and describe the general pattern of efferent and afferent connections within these regions. The first region that I will define is Anterior Cingulate Cortex (ACC) This region encompasses areas within the ACC gyrus including Brodmann's areas 32, 24, 25 and area 9 within dorsal ACC. There is some disagreement over the anatomical definition of the dorsal bank of ACC sulcus anterior to the genu of the corpus callosum (Carmichael and Price, 1995a, Middleton and Strick, 2001, Luppino et al., 2003). Personal communication with Brent Vogt leads me to believe that this region shares cytoarchitectural properties with area 9. It should be noted that posterior portions of the ACC sulcus (24c' and 24a'b') may be defined as cingulate motor area (CMA). The electrophysiological recordings centre around area 9/24c but I shall consider the connections of the entire ACC. Lateral prefrontal cortex (LPFC) includes areas

8, 9, 46, 9/46 and 47/12. All electrophysiological recordings will come from area 9/46d and to a lesser extent 9/46v. Orbitofrontal cortex (OFC) spans almost the entire orbital surface including areas 11, 13, the orbital parts of area 12 and orbital parts of insula cortex. My electrophysiological recordings come from areas 11 and 13 between the medial and lateral orbital sulci. Finally, I define ventromedial prefrontal cortex (vmPFC) as solely area 14 with my recordings coming from the medial and orbital parts. For the purposes of this anatomical review I will state the general connections into the above defined regions rather than connections within individual Brodmann's areas. This description will therefore define the broad patterns of connectivity in PFC which are critical to decision making.

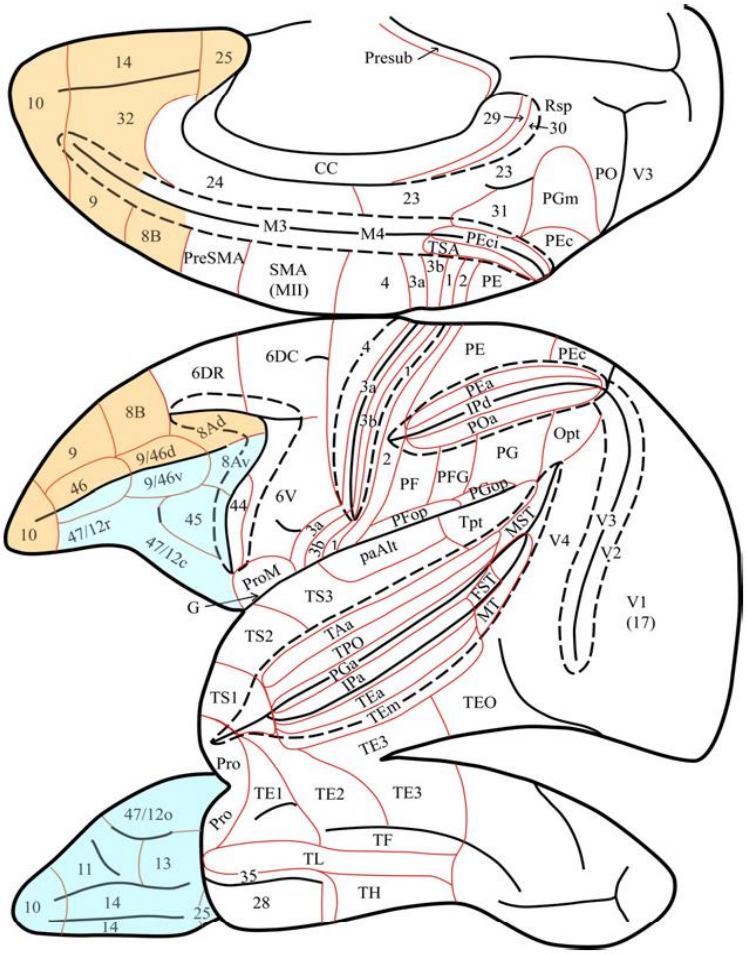

Figure 2.1: A diagrammatic depiction of the position of Brodmann's areas in the PFC (highlighted blue and orange). Taken from Yeterian et al. (2012).

Anatomical Basis of Sensory Input into PFC

In order to facilitate value based decision making the PFC must have access to various types of sensory modalities such as somatosensory, visual, olfactory, visceral, gustatory and auditory information. Tracer studies in primates have shown that these inputs enter the PFC in different regions (Carmichael and Price, 1995b). PFC has access to two sources of visual information. The inferior temporal cortex projects strongly to lateral OFC with lighter projections to the more medial parts. Also, lateral and posterior parts of OFC are strongly connected to superior temporal gyrus anterior and posterior (STGa and STGp) (Carmichael and Price, 1995b). Neurons in these temporal areas are thought to encode complex multimodal characteristics of visual and auditory stimuli implying that these connections are the source of the complex visual information to OFC (Bruce et al., 1981, Poremba and Mishkin, 2007).

Neurons in secondary somatosensory cortex (SII) terminate in the more central parts of OFC and originate from two separate clusters within SII (Carmichael and Price, 1995b). One of these clusters is known to correspond to the representation of the digits on the somatosensory body map within SII (Carmichael and Price, 1995b). Furthermore, primary somatosensory cortex (areas 1 and 2) has wide ranging projections within the central parts of OFC.

Visceral inputs to the PFC come from the parvocellular division of the ventroposterior medial nucleus of the thalamus (VPMpc), which in turn receives gustatory input from the solitary nucleus (Carmichael and Price, 1995b). These inputs terminate in the most posterior (insula) parts of OFC. This posterior part of OFC then projects on to many parts of PFC including all parts of OFC and even to more caudal parts of vmPFC (Carmichael and Price, 1995b). Also, projections from both gustatory cortex (in the insula) and olfactory cortex terminate in the central and posterior portions of OFC (Carmichael and Price, 1995b).

In general, ACC and LPFC receive fewer sensory inputs than OFC, however, the inputs they do receive tend to be more multisensory in nature from areas such as agranular temporal pole and superior temporal sulcus (Kondo et al., 2003, Saleem et al., 2014).

In conclusion, OFC receives every type of basic sensory input except auditory input. However, it is important to note that sensory inputs are distributed across a large extent of the OFC cortex suggesting that information from various sensory modalities needs to be integrated within OFC circuits in order to form complex sensory representations of the environment. These inputs are also completely unique to OFC implying that it enjoys a specific role in the process of value guided decision making although ACC and LPFC also receive an amount of multi-sensory information.

Premotor Connections of the PFC

As the PFC lies directly anterior to the premotor areas of the brain, it is no great surprise that connections exist between these regions. Tracer studies have shown that ventral parts of LPFC (area 47/12) have strong connections with supplementary eye fields (SEF) (Huerta and Kaas, 1990, Gerbella et al., 2010). SEF is then in turn strongly connected to frontal eye fields (FEF) which is the principle area concerned with voluntary eye movement (Huerta and Kaas, 1990). More spare connections also exist between SEF and area 9/46 (Huerta and Kaas, 1990). Weak connections also exist between central and posterior portions of ACC and SEF implying two potential PFC inputs into the saccadic system (Luppino et al., 2003). Somatotopic connections exist between posterior parts of ACC (i.e. CMA) and motor structures including presupplementary motor area (pre-SMA), premotor area (PM) and primary motor cortex (Morecraft and Van Hoesen, 1992, Arikuni et al., 1994, Hatanaka et al., 2003, Luppino et al., 2003). Direct projections also exist from LPFC to premotor cortex (Matelli et al., 1986).

All of the connections described above provide a direct pathway by which value information can be relayed from PFC to the motor system. However, there are relatively few paths by which this information can be relayed for motor output. The two main outputs are from ACC and LPFC, both of which can communicate with the saccadic and motor systems to one extent or another. Information transferred via LPFC is probably more likely to be relevant for eye movements whereas information given to PM is more likely to relate to body movement (Carmichael and Price, 1995b, Luppino et al., 2003). The main cortical pathway that OFC and vmPFC can influence motor action is therefore through communication with ACC or LPFC.

Limbic Connections with the PFC

Sources of limbic connections to PFC include the amygdala, hippocampus, temporal pole and the entorhinal, perirhinal and parahippocampal cortices. Although these regions are considered limbic in nature, many of them can also compute complex visual information about the environment. Therefore the connections I will describe in this connection may well outline possible pathways for complex sensory information as much as they could be for limbic information. Projections from the basal nucleus of the amygdala terminate across the entire extent of OFC and the more anterior portions of ACC (area 32 and 24a) (Carmichael and Price, 1995a). However, OFC connects preferentially to the medial portion of the basal nucleus whereas the anterior ACC connections lie more towards the lateral part of the nucleus (Carmichael and Price, 1995a). The divisions of the accessory basal nucleus also show differential projections to PFC. The magnocellular portion projects extensively throughout OFC and vmPFC with some minor connections to anterior ACC (Carmichael and Price, 1995a). In contrast the parvocellular portion is only seen to project specifically to posterior portions of vmPFC and very lateral and posterior parts of OFC (Carmichael and Price, 1995a). Finally, the lateral nucleus of the amygdala, periamygdaloid cortex and the

anterior cortical nucleus all project almost exclusively to posterior agranular parts of OFC (Carmichael and Price, 1995a). Although limbic input spans the entire PFC, many of the more selective connections described above terminate exclusively in OFC which suggests that OFC may enjoy unique access to certain limbic information from amygdala. This conclusion is backed up by simultaneous lesion and electrophysiology which shows a specific decrease in OFC value coding when amygdala lesions are administered (Rudebeck et al., 2013a). The amygdala has connections with almost every cortical visual area and therefore the amygdaloprefrontal connection may be another pathway by which complex visual information is relayed to PFC (Carmichael and Price, 1995a). It has also been suggested that the function of the connections between the lateral nucleus and the PFC may support the integration of ingestion related sensory information which may be critical for the valuation of food stuffs (Carmichael and Price, 1995a). The basolateral complex (which includes both the basal and lateral nuclei of the amygdala) is known to be important in linking and updating the association between stimuli and outcomes which may explain its extensive OFC projections (Baxter and Murray, 2002, Saddoris et al., 2005). The amygdala may play a role in providing emotional context to the PFC (Barbas et al., 2011).

In the hippocampus only the rostral subiculum has significant projections to PFC. Its projections terminate throughout vmPFC and more medial potions of OFC with no projection to ACC (Carmichael and Price, 1995a). These connections have been shown to come specifically from CA1 neurons in the hippocampus (Cavada et al., 2000). The hippocampus plays a role in memory and it may play a role in integrating spatial information (Thierry et al., 2000). Connections between PFC and hippocampus are critical for providing contextual information which facilitates hippocampal memory function (Preston and Eichenbaum, 2013). In contrast to the rostral subiculum, the parahippocampal cortex has been shown to project more medially to ACC and posterior vmPFC (Kondo et al., 2005). This part of the limbic system is known to deal with spatial memory, which may be important in selecting relevant actions which may explain its connections to ACC (Squire et al., 2004). Posterior

parahippocampal cortex also sends projections mainly to the medial surface with connections with ACC and vmPFC as well to medial OFC (Carmichael and Price, 1995a).

Strong reciprocal connections exist between entorhinal cortex and much of OFC, vmPFC and parts of ACC (i.e. areas 24a and 24b) (Carmichael and Price, 1995a, Arikuni et al., 1994). The entorhinal cortex is known to be an important relay for information from limbic areas throughout the brain into the hippocampus and is also thought to play a vital role in many aspects of cognition including attention, stimulus conditioning and working memory (Coutureau and Di Scala, 2009). In contrast to the entorhinal cortex, the perirhinal cortex has relatively few connections to PFC. It only connects to posterior parts of OFC (Kondo et al., 2005). The perirhinal cortex has been implicated in spatial recognition memory and memory consolidation which implies that OFC receives more limbic information pertaining to the visual or spatial properties of a stimulus compared to ACC which be relevant for its computations in more abstract frames of reference (Suzuki, 1996). Furthermore, the connections between OFC and the perirhinal cortex have been implicated in credit assignment (Clark et al., 2013, Walton et al., 2010).

In conclusion, what is striking about the limbic connectivity to the PFC is that almost every PFC area seems to receive input from at least one source. This suggests that the limbic system may exert huge influence on the function of the entire PFC although the specific influence will almost certainly depend on the origin of limbic input to each area.

Basal Ganglia Connections with PFC

Basal ganglia connections are known to be different between various parts of PFC. Anterior ACC (area 25) projects predominantly to the medial portion of the head, body and tail of the caudate, core of the nucleus accumbens and rostral ventral putamen (Ferry et al., 2000). Central OFC exhibits a similar pattern of connections. However, more anterior OFC

has restricted projections to the medial edge of the caudate and medial ventral striatum (Ferry et al., 2000, Haber et al., 1995). Projections from posterior vmPFC terminate exclusively in the core and shell of the nucleus accumbens (Nakano et al., 1999). However, more dorsal vmPFC/ventral ACC also has light projections to medial caudate as well as the strong projections to ventral striatum (Ferry et al., 2000, Kunishio and Haber, 1994, Eblen and Graybiel, 1995). More lateral and posterior regions of the lower bank of the cingulate sulcus (CMA) exhibit connections with dorsal (sensorimotor) striatum in both the caudate and putamen (Kunishio and Haber, 1994, Ferry et al., 2000). Unlike OFC, vmPFC and ACC, LPFC does not connect to ventral striatum but instead projects widely to lateral caudate and medial putamen (Yeterian and Pandya, 1991, Calzavara et al., 2007).

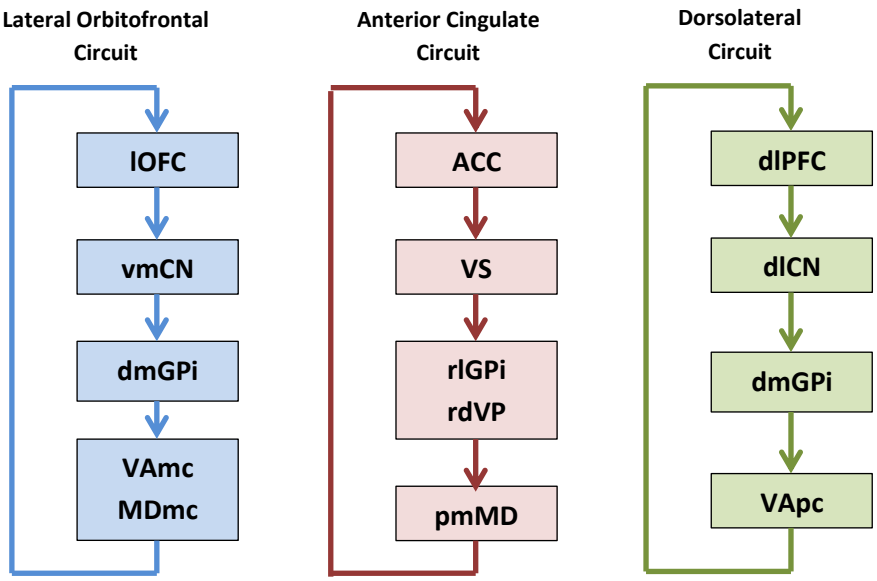

Figure 2.2: A schematic diagram depicting the three basal ganglia loops pertaining to PFC. Abbreviations: lateral orbitofrontal cortex (lOFC), caudate nucleus (CN), globus pallidus internal segment (GPi), ventral anterior thalamic nucleus (VA), medial dorsal nucleus of thalamus (MD), anterior cingulate cortex (ACC), ventral striatum (VS), ventral pallidum (VP), dorsolateral PFC (dlPFC), ventromedial (vm), dorsomedial (dm), dorsolateral (dl), rostrolateral (rl), rostrodorsal (rd), posteromedial (pm), magnocellular portion (mc), parvocellular portion (pc). Adapted from Alexander et al. (1986).

In addition to the focal PFC projections to ventral striatum, more diffuse PFC-striatum connections have also been described (Haber et al., 2006). These projections often cross

functional areas and enter the dorsal striatum (Haber et al., 2006). These projections may serve to modulate the strength of the focal topographical signal under certain conditions (Haber et al., 2006).

The topography of PFC inputs into the basal ganglia has particular importance considering that striatal efferants also show topography. Anterograde tracing studies have shown that projections from medial ventral striatum terminate at the medial edge of the ventral pallidum whereas projections from more central regions (such as the shell of the nucleus accumbens) terminate on the border between the ventral pallidum and the bed nucleus of the stria terminalis (Haber et al., 1995). This shows that the parallelism of cortical processing between the OFC, vmPFC, ACC and LPFC may be maintained throughout basal ganglia loops (Haber et al., 1995). PFC is known to have three separate loops that are known to pass through the basal ganglia (**Figure 2.2**). The 'lateral orbitofrontal' circuit projects to ventromedial caudate which in turn projects to the dorsomedial sector of the internal pallidal segment (Alexander et al., 1986). This part of the pallidum then projects to the magnocellular portions of both the ventral anterior thalamic nucleus (VAmc) and the medial dorsal nucleus (MDmc). Both of these nuclei then project back to lateral OFC to complete the loop (Alexander and Crutcher, 1990). The 'anterior cingulate' circuit has projections from ACC to ventral striatum and then on to rostrolateral internal globus pallidus (GPi) and rostrodorsal ventral pallidum (VP) (Alexander and Crutcher, 1990). These parts of GPi and VD then project on to posteromedial MD nucleus, which then sends projections back to ACC. Finally, the 'dorsolateral prefrontal' circuit originates from LPFC, which send input into dorsolateral head of the caudate and the rostrocaudal part of the tail of the caudate (Alexander and Crutcher, 1990). These parts of the caudate then project on to dorsomedial parts of the globus pallidus which then project on to parvocellular portion of the ventral anterior thalamic nucleus (VApc). The loop is then completed by projections from VApc back to LPFC. The function of loops may be to allow the continuous processing of

complex chains of events and may also imply specific functionality within these regions of PFC (Haber and Calzavara, 2009).

An important source of basal ganglia input to the PFC comes from the ventral tegmental area (VTA) and the substantia nigra pars compacta (SNpc). The major type of neuronal output from both these regions is dopaminergic in nature. In general, a gradient of dopaminergic innervation has been observed running from medial (higher innervation) to lateral (lower innervation) (**Figure 2.3B**) (Williams and Goldman-Rakic, 1998). Dopamine has two major receptors on which it acts. These D_1 and D_2 receptors are thought to have functional differences although these remain unclear (Glausier et al., 2009). Comparative differences have been noted between the density D_1 and D_2 receptors in certain parts of PFC. ACC contains significantly greater concentrations of D_1 receptors compared to D_2 although the functional significance of this is unclear (Richfield et al., 1989).

Retrograde tracer studies have found that both VTA and SNpc send projections to various parts of the PFC (**Figure 2.3A**) (Porrino and Goldman-Rakic, 1982). LPFC has connections with anterior parts of VTA and antero-medial and antero-dorsal SNpc (Porrino and Goldman-Rakic, 1982). Injections into OFC all find labelled neurons throughout the VTA but not with SNpc (Porrino and Goldman-Rakic, 1982). ACC receives the largest dopaminergic inputs (Porrino and Goldman-Rakic, 1982, Gaspar et al., 1989, Berger et al., 1988). Projections to anterior ACC originate from more medial regions of VTA whereas posterior ACC receives input from lateral parts of VTA (Williams and Goldman-Rakic, 1998, Raghanti et al., 2008). Even within ACC there appears to be a sharp transition from higher to lower dopamine innervation between areas 24b and 24c (Williams and Goldman-Rakic, 1998).

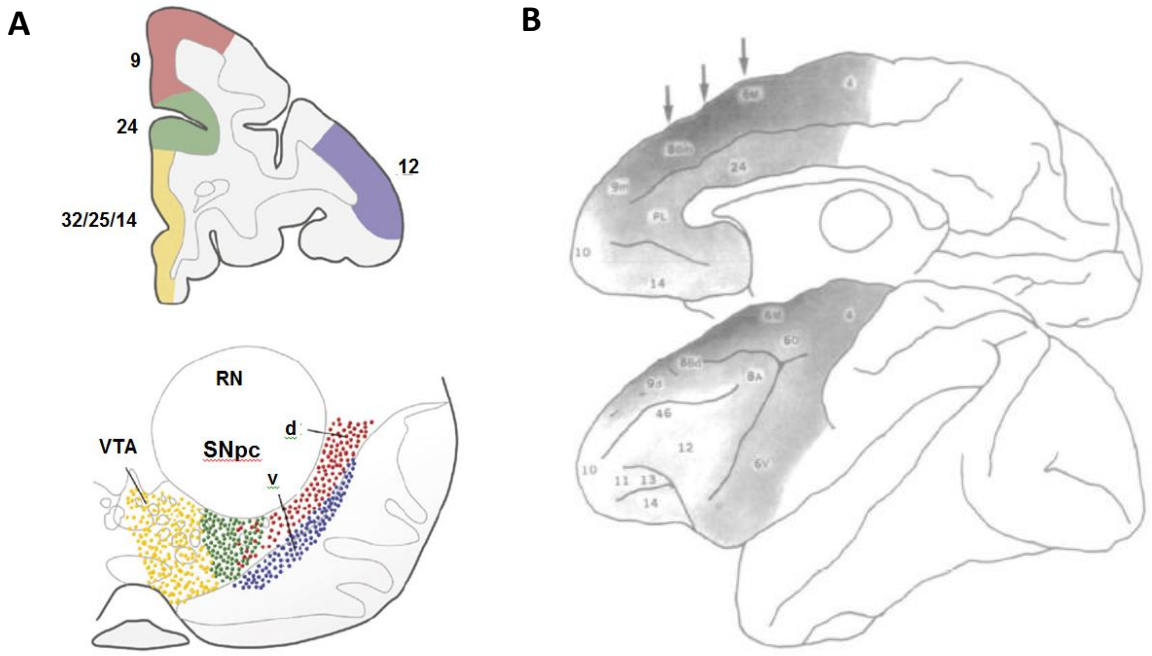

Figure 2.3: (A) A diagrammatic representation of the topographical arrangement of PFC inputs into midbrain. Abbreviations: ventral tegmental area (VTA), red nucleus (RN), substantia nigra pars compacta (SNpc), dorsal (d), ventral (v). Taken from Bjorklund and Dunnett, (2007). (B) The density of dopamine terminals in the frontal lobe. A density gradient can be seen in the rostral-caudal and the medial-lateral axis. ACC contains a significantly higher density of dopamine receptors than OFC or vmPFC. Taken from Williams and Goldman-Rakic (1993).

A study by Frankle et a.l (2006) in primates found a sparse projection to VTA and SNpc from ACC. Injections into OFC show connections to medial, dorsal SNpc and rostral VTA . In contrast, vmPFC is strongly connected to VTA . LPFC send projections mainly to medial and dorsal parts of SNpc with very few to VTA . The pattern of connectivity described by Frankle et al. appears to be quite similar to the pattern of efferent projections from midbrain dopamine to the PFC, suggesting that information transfer is likely to be reciprocal and topographical. However, the relative strength of the efferent dopamine-PFC connections is not equal. This may have some relevance to the computations performed through these reciprocal connections.

In conclusion, midbrain dopamine projections show strong topography with prefrontal organisation. Activity of dopaminergic neurons in VTA and SNpc is known to encode reward

prediction errors that are thought to be important for reward guided behaviour and reinforcement learning which are important functions of PFC (Schultz, 1998, Matsumoto and Hikosaka, 2009).

Networks in PFC

It is generally accepted that subregions of PFC can be divided into broad anatomical networks which partially transcend the definitions of ACC, OFC and vmPFC that I defined previously (Öngür D, 2000). Öngür and Price (2000) describe the existence of anatomically separate 'medial' and 'orbital' networks based on patterns of cortico-cortical and cortico-subcortical connections (**Figure 2.4**).

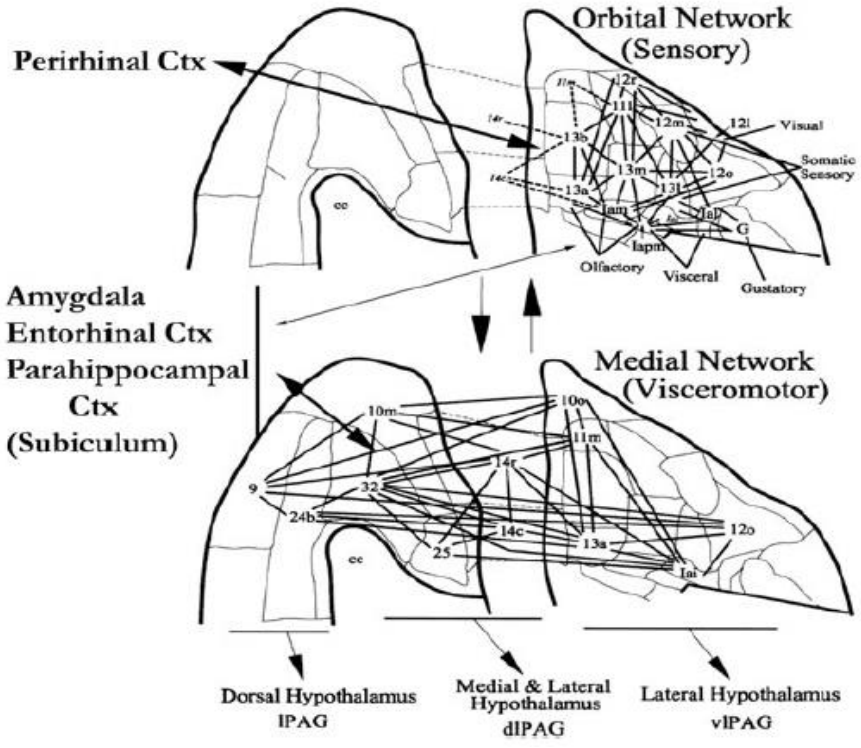

Figure 2.4: The structure and layout of the 'medial' and 'orbital' networks. Taken from Price (2007).

The so-called ‘orbital network’ consists solely of areas found on the orbital surface of the brain: Brodmann’s areas 11m, 11l, 12r, 12m, 12l, 12o, 13a, 13b, 13m, 13l, 14, Iam, Iapm and Ial which show a high degree of reciprocal interconnectivity (**Figure 2.4**) (Carmichael and Price, 1996). This network encompasses the entire OFC as previously defined but also includes vmPFC. As previously stated, areas within the ‘orbital network’ receive a huge amount of sensory input from multiple areas including those dealing with visual, gustatory, olfactory, somatosensory and visceral information (Öngür D, 2000). It also receives a large limbic input from enterhinal and perirhinal cortices and the basal, accessory basal and lateral nuclei of the amygdala (Carmichael and Price, 1995a, Öngür D, 2000).

The ‘medial network’ consists of reciprocal connections between brain areas located on both the medial and orbital surfaces: Brodmann’s areas 9, 10m, 10o, 11m, 12o, 13a, 14r, 14c, 24b, 25, 32 and Iai (**Figure 2.4**) (Carmichael and Price, 1996). This not only encompasses ACC but also vmPFC and parts of OFC. Unlike the ‘orbital network’, it receives very few sensory projections, although in general it does receive similar limbic input in comparison with the ‘orbital network’ with the exception that is more connected to the ventromedial part of the basal nucleus of the amygdala (Öngür D, 2000). Whereas the ‘orbital network’ seems to be a system that predominantly receives input from other areas, the ‘medial network’ tends to project outputs to other areas of the brain. It is known to send strong projections to the hypothalamus and peri-aqueductal grey (PAG), which are both associated with visceral and autonomic function (Öngür D, 2000, Keay et al., 1994, Price, 1999).

It should be noted that there are several areas of PFC which span both ‘medial’ and ‘orbital’ networks (Price, 1999). These include areas 11m, 12o and 13a. Areas 12o and 13a are both connected to area 24b which may allow the orbital network direct access to one of the major motor outputs of the PFC. Another interpretation is that these areas may be the main point of reciprocal communication between ACC and OFC and as these two areas

have extensive projections to all parts of ACC, OFC and vmPFC providing an ideal interface for communication (Carmichael and Price, 1996).

Not every region of PFC has access to information from sensory and limbic areas. Furthermore, not every region is able to send outputs to motor regions. In order for information to be relayed between different areas these areas need to be connected. The anatomical 'medial' and 'orbital' networks provide one potential solution to this problem but does not necessarily explain the broad connective properties of PFC regions. It is no surprise that the overall connectivity of PFC is extremely complex (**Figure 2.5**). Statistical analysis of the connections between regions of PFC has shown that regions have between 3 and 13 outputs with the mode number of 8 outputs per region (Averbeck and Seo, 2008). The same study also showed that every PFC area was able to access all types of extra-PFC information within two connections of its anatomical position suggesting that every PFC region has the potential to access information of almost all modalities (Averbeck and Seo, 2008).

Tracer studies indicate that area 13 (in central OFC) has widespread connections throughout PFC and most notably, these connections seem to span 'medial' and 'orbital' networks (i.e. physically link ACC and OFC) (Barbas and Pandya, 1989, Price, 2007, Carmichael and Price, 1996). Area 12o, which is the other area thought to be the point of connection between the two networks, also shows broad connectivity throughout PFC and importantly exhibits strong connections with area 24b (Price, 2007, Barbas and Pandya, 1989). vmPFC may be well located to act as an interface or conduit for communication between ACC and OFC given its anatomical position and connections between the two regions.

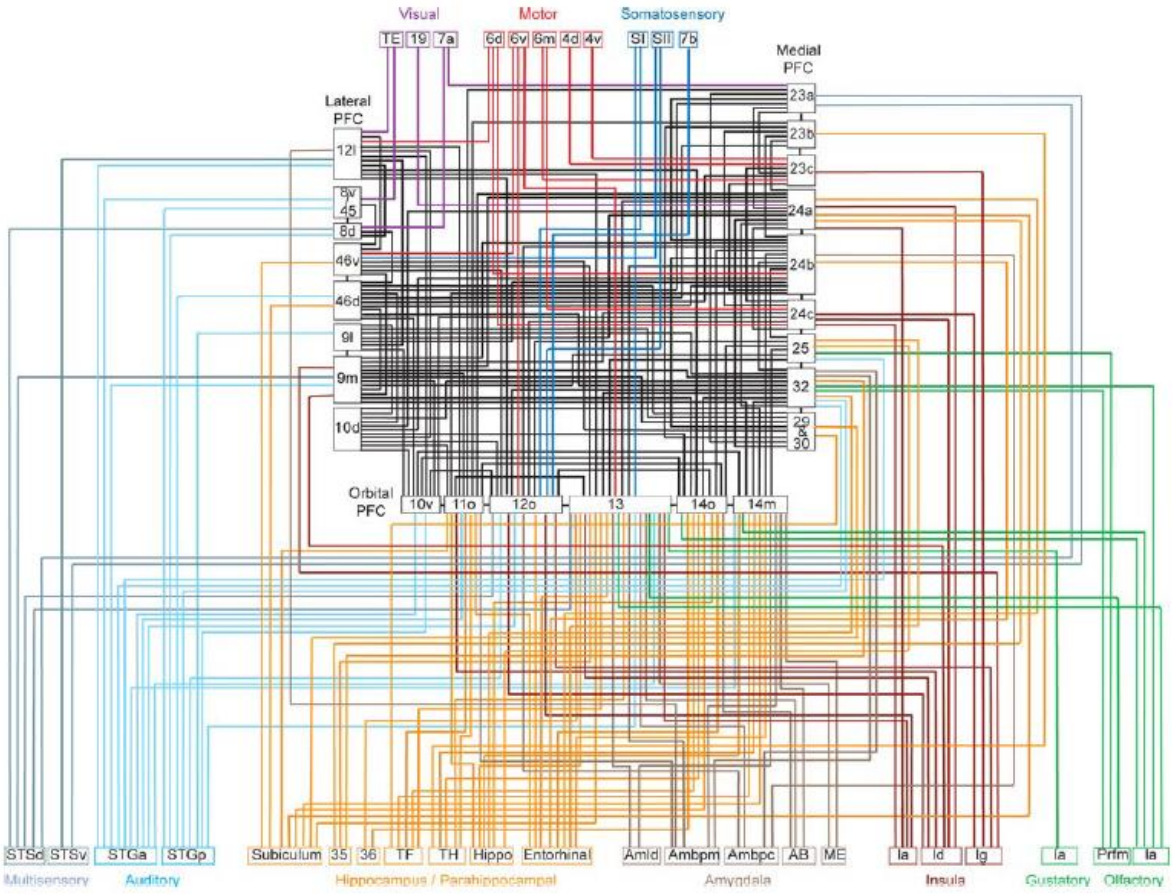

Figure 2.5: A schematic representation of the intrinsic and extrinsic connections of areas in the PFC. Common abbreviations are used. Taken from Averbeck and Seo (2008).

Primate-Human Homology

From an empirical point of view there is strong homology between human and macaque brain (Ongur and Price, 2000). Cytoarchitecturally, studies have shown great similarities between human and macaque OFC, although the relative sizes of the two structures differ (Wise, 2008). Both human and macaque OFC contains granular cytoarchitecture which rodents lack (Wise, 2008). These primate specific parts of OFC are found in the more anterior portions of OFC, which as stated above, are the areas which receive greater complex visual input and less gustatory and visceral input. These areas are also the same areas of OFC and vmPFC that are commonly recorded from in primate neurophysiology experiments (for example, Kennerley and Wallis (2009a), Kennerley et al.

(2009), O'Neill and Schultz (2010), Padoa-Schioppa and Assad (2006), Strait et al. (2014)). The agranular cortex, which is conserved between rats and humans is very posterior in humans and macaques (Wise, 2008).

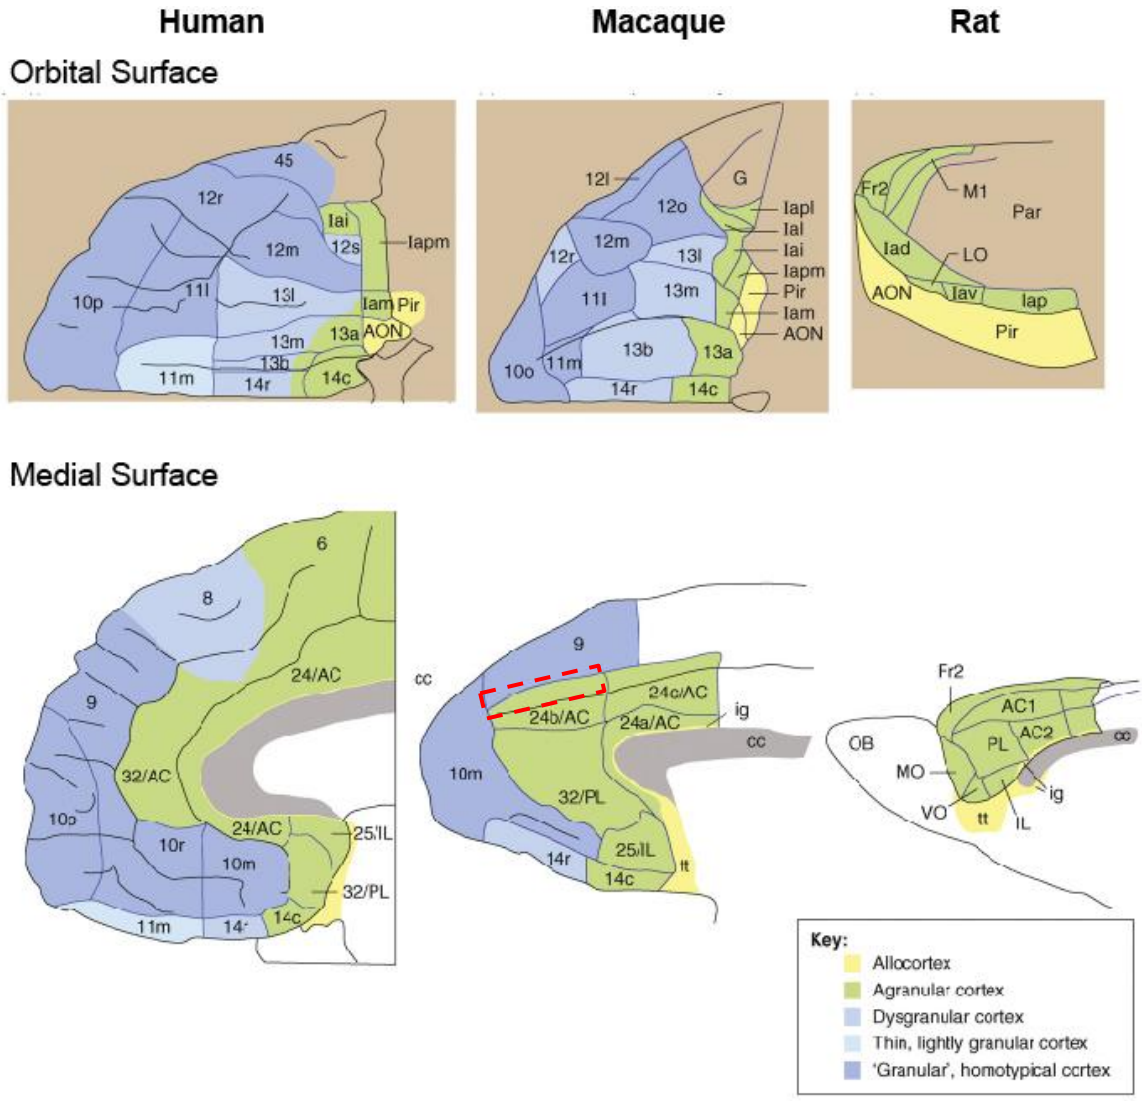

Figure 2.6: A diagrammatic representation of the cytoarchitecture of human, macaque and rat prefrontal cortices. The area within the red rectangle denotes the pre-genual part of ACC which is disputed in nomenclature. Abbreviations: a, agranular; AC, anterior cingular area; AON, anterior olfactory nucleus; cc, corpus callosum; Fr2, second frontal area; I, insula; MO, medial orbital area; LO, lateral orbital area; M1, primary motor cortex; Par, parietal cortex; Pir, piriform cortex; PL, prelimbic cortex; VO, ventral orbital area; l, lateral; m, medial; o, orbital; r, rostral; c, caudal; i, inferior; p, posterior; s, sulcus; v, ventral. Taken from Wise (2008).

Much like OFC, the cytoarchitectural properties of the ACC are very similar between humans and macaques (Wise, 2008). However, unlike OFC, macaque ACC consists almost solely of agranular cortex which is homologous to rodent ACC (Wise, 2008). A comparative connectivity study of human and macaque LPFC reported strong homology between the two species. The only exception was that the pattern of connections observed in human lateral frontal pole (area 10) was more similar to connections seen in macaque LPFC (around the principal sulcus) than that of macaque lateral frontal pole (Sallet et al., 2013, Neubert et al., 2014). A similar study examining the correspondence of areas such as ACC and vmPFC also found strong similarities between humans and macaques (Neubert et al., 2015). There are also several empirical difference between human and macaque PFC (Öngür D, 2000). These include disagreements over the similarity of Brodmann's area 32 (anterior ACC) and area 12 (lateral OFC) as classified in the human compared to primates (Öngür D, 2000). However, these differences have little relevance to the work presented in this thesis.

Overall, based on the cytoarchitectural and connection studies, it appears that strong homologies exist between the human and macaque PFC which in means that the neural dynamics of value based decision making in primates is likely to bare close resemblance to neural dynamics in humans.

Conclusions

The extensive connectivity of the PFC helps explain the flexibility in its function. The medial and orbital walls of the PFC (i.e. ACC and OFC) are anatomically separate which is no doubt vital in explaining functional differences between regions. In the context of a visually guided value based decision making tasks several areas of both networks may be of particular importance. **Figure 2.7** depicts the important intrinsic and extrinsic connections of PFC which may be relevant to decision making.

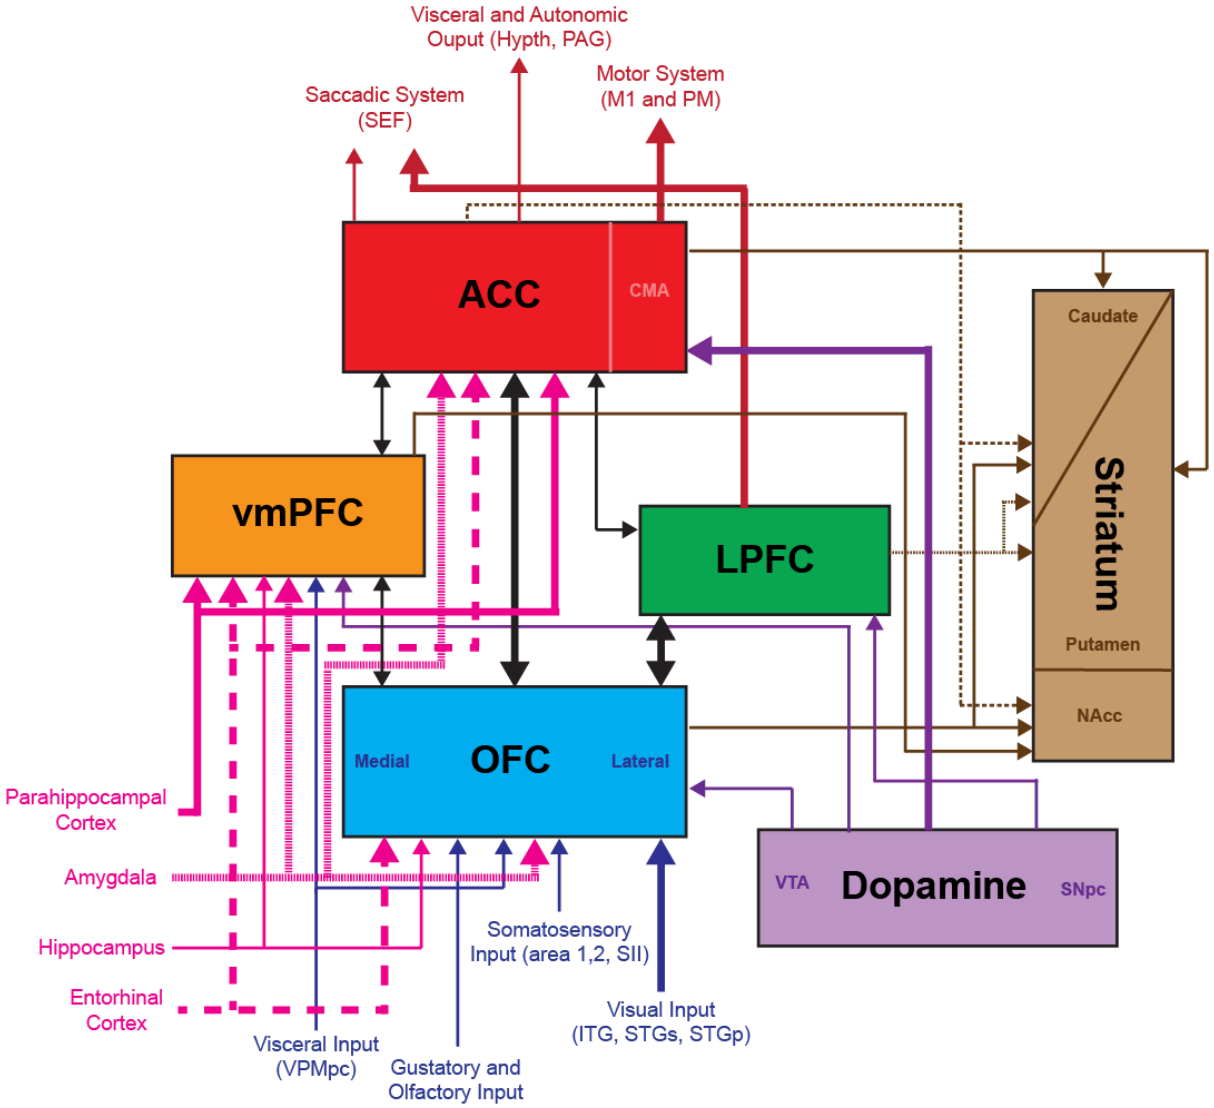

Figure 2.7: A diagrammatic representation of the main intrinsic and extrinsic connections of PFC as relevant to decision making. Dark blue arrows indicate sensory connections, magenta arrows show limbic connections, dark red arrows depict motor connections, purple arrows show dopaminergic input, brown arrows show projections to striatum and black arrows show connections between PFC regions. The thickness of each arrow represents the relative at size or importance of each input. General input and output topology is represented for OFC, striatum and dopamine. Reciprocal connectivity is depicted within PFC but not for regions outside PFC. Abbreviations: Hypth, hypothalamus; ITG, inferior temporal cortex. All other abbreviations are the same as previously stated.

Lateral and central parts of OFC receive visual information from both STG and inferior temporal cortex. It also has extensive connections with ventral striatum and receives dopaminergic input from VTA. Through its connections OFC is then capable of integrating

sensory and limbic information and passing this on either through its medial connections to ACC, its lateral connections to LPFC which both can then go on to manipulate the motor system in order to effect a decision. Conversely, vmPFC may be integrate information from OFC and pass it on to ACC through its reciprocal connections with both regions. vmPFC itself has several unique limbic and striatal connections that may convey specific computational properties.

Chapter 3: Covert and Overt Attention During Learning and Decision Making

Introduction

Every day our environment presents us with a rich assortment of visual information, some of which is highly relevant to potential future decisions. To make decisions efficiently we must quickly orient our eyes towards things that provide us with the most important data about upcoming decisions. How such attentional strategies are implemented in the brain is an intriguing question. A growing body of thought suggests that these strategies are a type of value based decision process which computes relevant visual locations to gather future information (Gottlieb, 2012, Jovancevic-Misic and Hayhoe, 2009, Sullivan et al., 2012, Brockmole and Henderson, 2005, Anderson, 2013). In this chapter we will explore the behaviour of primates performing a simple free gaze value decision task in order to characterise the roles of covert and overt attention in primate learning and decision making. We will then make inferences about the structure and connectivity between potential attention driving circuits and other decision making circuits.

It has long been known that basic physical visual properties of our environment (e.g., brightness) can have a powerful influence on saccadic behaviour (Theeuwes, 1992, Nothdurft, 2002, Itti and Koch, 2001, Folk et al., 1992). However, it is also well recognised that higher order non-physical properties of our environment (i.e., associations between value and stimuli) can also modulate saccadic attention through covert attention in both humans and monkeys (Yasuda et al., 2012, Anderson et al., 2011, Anderson and Yantis, 2012, Kim and Hikosaka, 2013, Milstein and Dorris, 2007). For example, both primate and human studies have found that with increasing exposure, subjects become significantly more

efficient at saccading towards valuable stimuli or locations when presented with multiple targets (Yasuda et al., 2012, Anderson et al., 2011). Both species are also significantly faster to saccade towards more rewarding stimuli during saccade contingent tasks (Milstein and Dorris, 2007, Kim and Hikosaka, 2013, Yasuda et al., 2012). Such findings are fascinatingly similar to the pattern of behavioural results commonly seen in value based decision making experiments where subjects choose more rewarding stimuli and make decisions faster when final choices are more valuable (Krajbich et al., 2012, Rudebeck et al., 2008, Philiastides and Ratcliff, 2013). From this, we could be tempted to conclude that neural systems that guide spatial attention may perform analogous value-based computations to those computations performed when choosing which option to eventually select.

If this is true, the question naturally follows: are the decision circuits that underlie value based choice and those underlying saccadic attention one and the same? Potential circuits for these dual roles might be frontal (Kennerley et al., 2009, Kennerley and Wallis, 2009c, Strait et al., 2014, Hayden and Platt, 2010, Padoa-Schioppa and Assad, 2006), parietal (Platt and Glimcher, 1999, Rorie et al., 2010, Sugrue et al., 2004) or subcortical (Lau and Glimcher, 2008, Kim and Hikosaka, 2013, Yasuda et al., 2012, Cai et al., 2011) areas, given neurons in these areas have all been implicated in value-based decision-making and exhibit responses correlating with attention and/or eye movements.

If the same circuit does drive both value guided saccadic attention and choice, we can make the straightforward prediction that output dynamics in both systems should be strongly correlated. From an anatomical and functional point of view, it seems unlikely that one circuit can subserve both behaviours for the following reasons. Firstly, although value signals are almost ubiquitous in the brain, lesion studies suggest the prefrontal cortex (PFC) may be the critical set of areas for value based decision making (Rudebeck et al., 2008, Camille et al., 2011a, Camille et al., 2011b, Fellows, 2006). Yet, although it is generally unreported in the literature, value signals in PFC are relatively slow, especially when compared to saccade latencies. Typical PFC value coding latencies are approximately 200-

300ms (Kennerley and Wallis, 2009a) but typical saccade latencies are 150-300ms (Enderle, 2002). This implies that PFC neurons might not be best suited for biasing saccades based on value information. On the other hand the subcortical saccadic network (i.e. Caudate, Superior Colliculus (SC), Substantia Nigra pars Reticulata (SNpr)) would be perfectly placed for value based computations with its extensive connections and overlap the reward network (Munoz, 2002). Yet despite the fact that neurons in areas like SNpr and Caudate encode value at fast latencies, these regions may not be critical in value based motor decision making (Hikosaka, 2007, Hikosaka et al., 2006, Kim and Hikosaka, 2013). Secondly, and more straightforwardly, it is normal for us to gather information with our eyes, but rare to directly act upon the world with saccades. The converse is true for our limbs. It therefore seems likely that *information value* will have a stronger bearing on saccades than the *reward value* of approaching or selecting an object, but the opposite may be true for action selection. Often information value and reward value will be strongly correlated, although they can be separated in certain circumstances (Gottlieb et al., 2014).

Whether covert attention (i.e. attending to stimuli in the environment without overtly saccading towards them) influences decision making is of particular relevance when considering contemporary models of value based choice, which argue that only information that is directly fixated or viewed (i.e., through overt saccades) correlates with eventual choice (Krajbich et al., 2010, Krajbich et al., 2012, Towal et al., 2013, Krajbich and Rangel, 2011). However, these attentional models frequently depend on the subject selecting a location to fixate at random, and only using overt attention to bias evidence accumulation towards making a choice. Yet, if value influenced where saccades were directed, as opposed to evidence accumulation proceeding only after an overt saccades, then we would be forced to reconsider whether evidence accumulation via covert attention biases decision making processes. Equally, because primate economic choice experiments typically require central fixation during the choice epoch, the importance of covert attention in evidence accumulation and decision-making processes remains largely unknown.

In order to better understand how ocular dynamics and covert attention correlate with evidence accumulation and value comparison processes, we recorded behavioural and eye position data during a binary free saccade decision making paradigm. On each trial, subjects chose between differently valued stimuli that were either well learnt or had been learned that day. We found that subjects invariably used covert attention to guide their saccades towards more valuable stimuli, and did so with increasing accuracy as the value of the stimuli became more learned. We will argue that this value guided saccadic decision is largely independent from that of the value guided final choice, because of clear differences between saccade and choice behaviours in terms of the speed of learning and novelty bonus. Thus, when covert attention can be used to evaluate information in the environment choice behaviour can become dissociable from overt saccadic behaviour.

Methods

Subjects

Two adult male rhesus monkeys (Macaca mulatta), M and F, were used as subjects in the study. All experimental procedures were approved by the Local Ethical Procedures Committee and carried out in accordance with the UK Animals (Scientific Procedures) Act.

Behavioural Protocol

Subjects sat in a behavioural testing chair facing a 19" computer monitor placed approximately 57cm away from the subjects' eyes. The height of the screen was adjusted so that the centre of the screen aligned with neutral eye level for the subject. An analog joystick (APEM Components, UK) was placed in front of the subject out of his line of sight and was used to make Left/Right manual responses during the task. Eye position and pupil tracking was achieved using an infrared camera (ISCAN ETL-200) sampled at 240Hz.

The behavioural paradigm was run using the MATLAB based toolbox MonkeyLogic (http://www.monkeylogic.net/, Brown University, USA). All joystick and eye position data was relayed to MonkeyLogic and used online during the task, and also interpolated and recorded by MonkeyLogic at 1000Hz.

Task

A representation of a single trial timeline can be found in **Figure 3.1B**. Subjects initiated each trial by returning the joystick to its centre position. At this point a white background appeared on the screen with a red centre fixation square (0.5 x 0.5 visual degrees in size). Subjects were required to fixate the red square for a continuous 500ms (fixation radius of 3 visual degrees) within a 10s time period. If this was not achieved, a short

'timeout' was given and the trial restarted. Once the fixation period was completed the fixation spot disappeared and two isoluminant stimuli (**Figure 3.1C**, 100 x 100 pixels) were presented 6.5 visual degrees to the left and right of the centre. Importantly, once the stimuli appeared, subjects were free to saccade to anywhere on (or off) the screen and to choose a cue using a left/right joystick response at any time. If subjects did not make a joystick response within 5s of stimuli onset, the trial was aborted. Once the response was made, a grey square was drawn around the chosen stimulus and a 500ms pre-feedback epoch was initiated, after which the unchosen stimulus was removed from the screen and the reward was initiated (feedback epoch). Subjects were rewarded with varying volumes of juice delivered to the mouth using a precise peristaltic pump (ISMATEC IPC). Subject M worked for a 50:50 water:apple juice mixture while Subject F worked for a 50:50 water:mango juice mixture.

Subjects were presented stimuli from one or more of four stimulus sets (see **Figure 3.1C** for an example set). Two sets consisted of cues associated with one of five probabilistic outcomes (10%, 30%, 50%, 70%, 90%) and the other two sets were associated with one of five magnitudes of reward size (0.14g, 0.33g, 0.51g, 0.71g, 0.90g). For each attribute (i.e. probability and magnitude), one stimulus set contained cues which subjects were highly 'Overtrained' on in previous training sessions (M: ~1500, F: ~3000 total exposures to the set). The other set contained 'Novel' cues which subjects had only experienced that day, where they were given limited exposure to each stimuli's value during the 'conditioning phase' (10 forced stimulus-outcome exposures with secondary conditioning per stimulus) which immediately preceded the 'choice phase' (see **Figure 3.1A**). Subjects could be given trials consisting of both overtrained stimuli ('Overtrained' trial), two novel stimuli ('Novel' trial) or one of each ('Mixed' trial) (see **Figure 3.1D**). Subjects were always asked to make choices within a certain attribute (e.g. choosing between probabilities) and never between attributes. All trial types were pseudorandomly interleaved. Optimality was measured by whether the subjects chose the highest value stimulus.

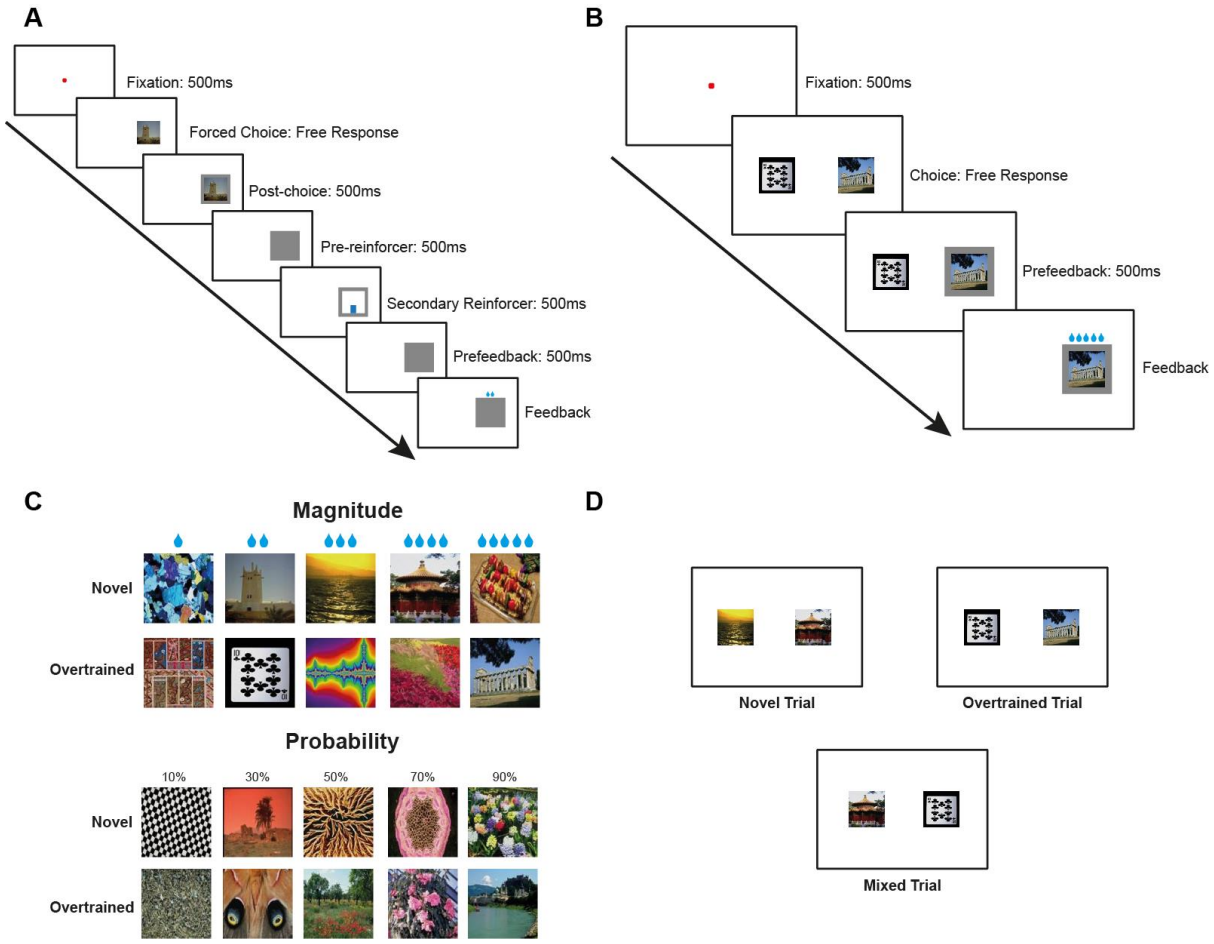

Figure 3.1: Overview of the task design. **(A)** Subjects began each session with a 'conditioning phase', where they were given 10 forced-choice trials of each of the 10 novel stimuli (100 trials in total) in order to learn the value assigned to each stimulus. Secondary conditioning (with a pre-learnt stimulus where bar height equalled value; blue bar: magnitude, black bar: probability attribute) was used to aid learning in these trials. 40 novel choice trials were periodically interleaved between forced-choice trials. **(B)** Once the 'conditioning phase' was complete, subjects moved to the 'choice phase' of the experiment where they were presented only with choice trials of various conditions. Subjects were free to saccade between the cues and make a manual joystick response to indicate choice at any time. **(C)** Example cue Set. Cues could either predict reward magnitude or probability and could either be well known to the subjects (Overtrained) or relatively new (Novel). **(D)** Example trials. On any given trial in the 'choice phase', subjects could be presented with two overtrained stimuli (Overtrained Trials), two novel stimuli (Novel trials) or one of each (Mixed Trials). Subjects were always asked to make choices within an attribute dimension and never had to choose between probability and magnitude stimuli.

Behavioural Analysis

The focus of this chapter is on the relationship between attention, learning and decision-making. As such, all analyses were performed on the choice trials (**Figure 3.1.B**).

All analyses were performed using MATLAB (MathWorks, USA). Reaction time was defined as the time between cue presentation and initial joystick movement. All eye position analysis was performed during this period. Eye position data were pre-processed by removing periods of time corresponding to either blinks or off screen gazes. For the purposes of our analysis, we collapsed across trials of different attributes for all results shown in this chapter because no meaningful differences were observed between attributes.

In order to decode whether the subject attended to information on the screen, a 16 square degree area (4 x 4 degrees) was defined around the centre of each cue. If the subject's gaze entered this area he was considered to have viewed or 'fixated' the stimulus. Time of the first saccade was defined as the time when the X position of the eye first left the 3 degree centre radius. First cue dwell time was defined as the period of time that the subject's eye position first entered the area around the cue until when it first left.

In order to avoid perfect separation in some regressions, all data analyses were performed using data collapsed across all sessions for a given subject. Logistic regressions were performed using **Equation 3.1** where YP is the probability of observing an event, b_0 is the inherent tendency to observe that event irrespective any other variables, b_n is the weighting coefficient and x_n is the regressor:

$$\boldsymbol{YP = \frac{1}{1 + e^{-(b_0 + b_1 x_1 + b_2 x_2 + \cdots + b_n x_{n)}}}} \qquad \textbf{Equation 3.1}$$

We characterised subject choice behaviour using logistic regression. We regressed the probability of left choice against the left-right value difference (collapsed across attribute types) for both Overtrained and Novel trials within the same regression. Doing this allowed us to compare the size of the beta coefficients for Overtrained and Novel trials using a linear hypothesis test. We also used the same regression model described above to test the effect of value on the probability of making a left first saccade.

In order to test whether subjects became more optimal at making a saccade toward the more valuable choice stimulus over the course of the session, we used a modified regression model which used saccade direction in the first and last 50 Overtrained and Novel trials as the dependent variable. The independent variables were as follows: left-right value difference in the first 50 Overtrained trials, left-right value difference in the first 50 Novel trials, left-right value difference in the last 50 Overtrained trials, left-right value difference in the last 50 Novel trials, and two binary terms which coded for Overtrained and Novel trials respectively.

To assess whether picture novelty influenced subject choice behaviour in Mixed trials, we used a regression model where the dependent variable was whether or not the novel choice stimulus was selected, and this was regressed against the value of the novel stimulus and the value of the overtrained stimulus, as well as a constant term which described the subjects' bias to choose novel stimuli irrespective of value. Finally, changing the dependent variable to the probability of saccading to a novel stimulus allowed us to test the effect of novelty on initial saccade behaviour.

We also used a logistic regression to test whether value influenced the probability of subjects making more than one saccade in a trial. The probability of making two saccades was the dependent variable and the independent variables were the value of the fixated stimulus and the non-fixated stimulus, as well as a constant term. It should be noted that for simplicity this regression was not split by trial type.

All linear regressions were performed using **Equation 3.2** where Y is the dependant variable, b_0 is the constant term and X_n are the regressor and are weighted by coefficients b_n:

$$\boldsymbol{Y = b_0 + b_1 X_1 + b_2 X_2 + \cdots + b_n X_n} \qquad \textbf{Equation 3.2}$$

We tested the effect of value on how long subjects fixated a stimulus (dwell time) using linear regression. The dependent variable was dwell time (normalised within each session) and the independent variables were the value of the fixated stimulus in Novel trials, the value of the fixated stimulus in Overtrained trials, the value of the non-fixated stimulus on Novel trials, the value of the non-fixated stimulus on Overtrained trials and a constant term.

In order to test whether our ‘Covert’ or ‘Overt’ models of left choice better explained subject choice behaviour, we computed model evidence for each model on a session by session basis. We achieved this by first estimating model parameters by performing a logistic regression of left choice using first the ‘Covert’ model which used the actual left and right stimulus values as input, then estimating model parameters of an ‘Overt’ model which contained identical regressors with the exception that whenever a stimulus was not overtly fixated, its value in the ‘Overt’ model for that trial was set to 3 (i.e. average value). Having estimated parameters for both models we then calculated the log likelihood (LL) for each model for each session. We then performed a binomial test in order to test whether one of the models consistently gave higher log likelihood estimates than the other model.

Results

Two macaque monkeys (subjects M and F) were trained to perform a free response value based decision making task in which they selected between two Overtrained or Novel stimuli, of differing value (see Methods, **Figure 3.1B-D**). Prior to performing the free response task, they completed 100 'conditioning trials' in which they learnt the reward probability and reward magnitude predicted by novel stimuli (see Methods). Subject M performed 14 sessions completing 9518 choice trials in total, while Subject F performed 20 sessions completing a total of 11593 choice trials. Trials were pseudorandomly selected from one of three conditions; Novel, Overtrained or Mixed (**Figure 3.1D**). All regression analyses were performed using data pooled across all sessions to avoid perfect separation. In this first section of results we will only consider behavioural differences between Novel and Overtrained trials. Here we collapsed across trials where stimuli reflected reward probability and reward magnitude, as both trial types showed similar results. A breakdown of subject choice optimality for each attribute and trial condition is shown in **Table 3.1**

Trial Condition	Subject M	Subject F
Novel Magnitude	91.6% (4.5)	86.1% (10.8)
Novel Probability	87.2% (3.9)	81.8% (8.5)
Novel Overall	89.4% (3.2)	83.9% (7.9)
Overtrained Magnitude	97.1% (1.5)	96.6% (2.9)
Overtrained Probability	96.2% (2.6)	93.7% (3.7)
Overtrained Overall	96.6% (1.7)	95.1% (2.4)
Mixed Magnitude	92.4% (3.1)	89.5% (4.9)
Mixed Probability	90.5% (4.4)	87.1% (4.7)
Mixed Overall	91.4% (2.8)	88.3% (3.9)

Table 3.1: Condition trial optimality. A table showing the subjects' choice optimality (i.e. choosing the most valuable stimulus) in different trial conditions. Numbers in brackets represent standard errors of the mean (SEM).

Choice Behaviour

Subjects were very good at selecting the more valuable option (**Table 3.1**, **Figure 3.2**). A logistic regression of choice left against left-right value difference found a strong value based effect (M: $T_{(Novel)}$=31.46, $T_{(Overtrained)}$=27.00, $p_{(Novel)}<3x10^{-217}$, $p_{(Overtrained)}<2x10^{-160}$, F: $T_{(Novel)}$=34.53, $T_{(Overtrained)}$=32.18, $p_{(Novel)}<3x10^{-261}$, $p_{(Overtrained)}<4x10^{-227}$). Notably, their choices were more sensitive to value in Overtrained trials than in Novel trials (linear hypothesis test comparing parameter estimates for Novel vs. Overtrained stimuli, M: $p<8x10^{-29}$, F: $p<9.2x10^{-57}$; difference in slopes in **Figure 3.2**). However, both subjects showed a significant increase in optimality across the session in Novel trials (Pearson's correlation coefficient, M: r = 0.178, p<0.04, F: r = 0.245, $p<6x10^{-4}$) but not Overtrained trials (Pearson's correlation coefficient, M: r = 0.039, p>0.5, F: r = -0.011, p>0.5). This implies that subjects continued to learn the value of Novel stimuli throughout the task.

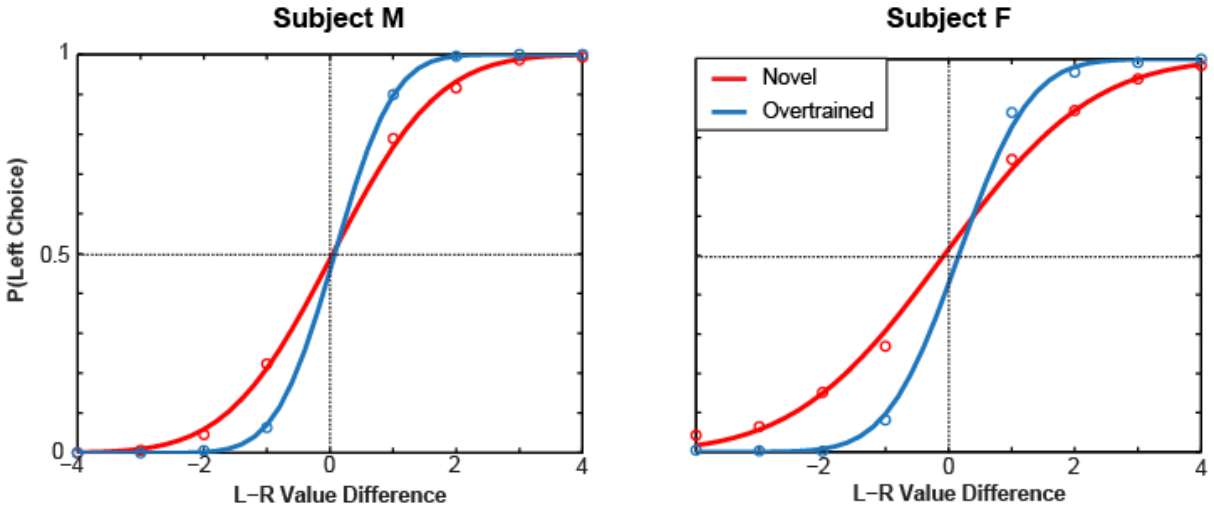

Figure 3.2: Psychometric function of final choice. The probability of choosing the left option as a function of the value difference between the left and right stimuli collapsed across all Novel (red) and Overtrained (blue) trials.

Value influences the first saccade direction, within 170ms of stimulus onset

An important feature of this task was that at cue presentation, subjects were free to saccade around the screen and make their choice (via joystick response) at any time within

5s of cue presentation. We therefore asked to what degree subjects' saccades were influenced by the value of pictures presented, and to what extent they were correlated with subjects' eventual choice. Importantly, subjects would almost always make a saccade to a stimulus prior to responding with the joystick (M and F: >99%) and their first saccades were very fast (**Figure 3.3**, median; M: 138ms, F: 170ms).

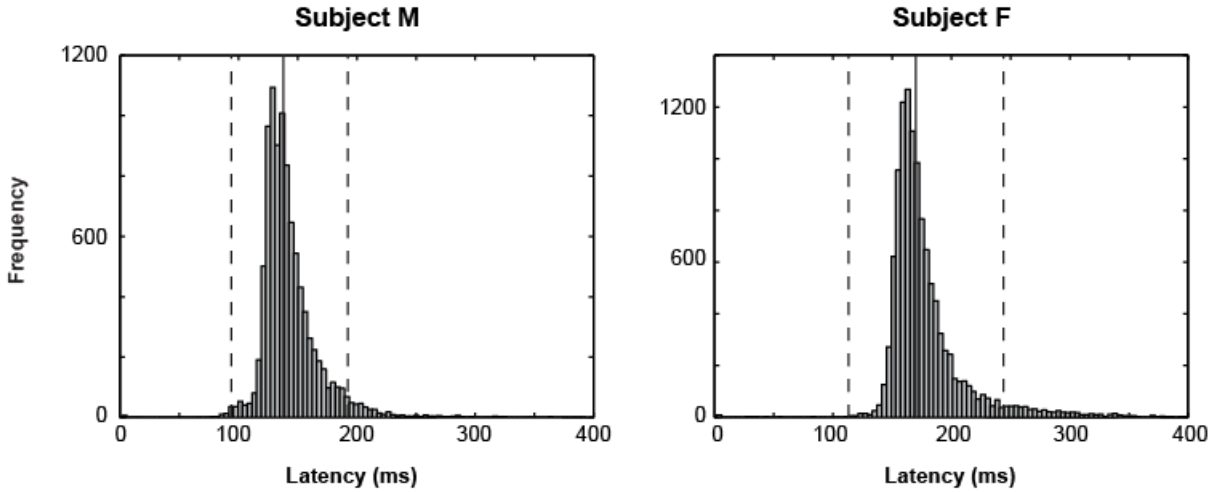

Figure 3.3: Distribution of first saccade latencies. A histogram of first saccade latencies for each subject showing the median (solid line, M: 138ms, F: 170ms) and the 95% confidence interval around the mean (dashed lines).

Surprisingly (given this latency), the direction of this fast first saccade was not random but instead was strongly influenced by the value of the presented pictures. A logistic regression of saccade direction against left-right value difference found that the direction of the first saccades in both trial types were significantly more likely to be towards the more valuable stimulus (M: $T_{(Novel)}=17.90$, $T_{(Overtrained)}=30.96$, $p_{(Novel)}<2x10^{-71}$, $p_{(Overtrained)}<2x10^{-210}$, F: $T_{(Novel)}=16.77$, $T_{(Overtrained)}=29.40$, $p_{(Novel)}<5x10^{-63}$, $p_{(Overtrained)}<6x10^{-190}$) (**Figure 3.4**). As with subjects' eventual joystick responses (choices), this effect was stronger in Overtrained trials than Novel trials (linear hypothesis test comparing parameter estimates for Novel vs. Overtrained stimuli, M: $p<2x10^{-61}$, F: $p<2x10^{-33}$).

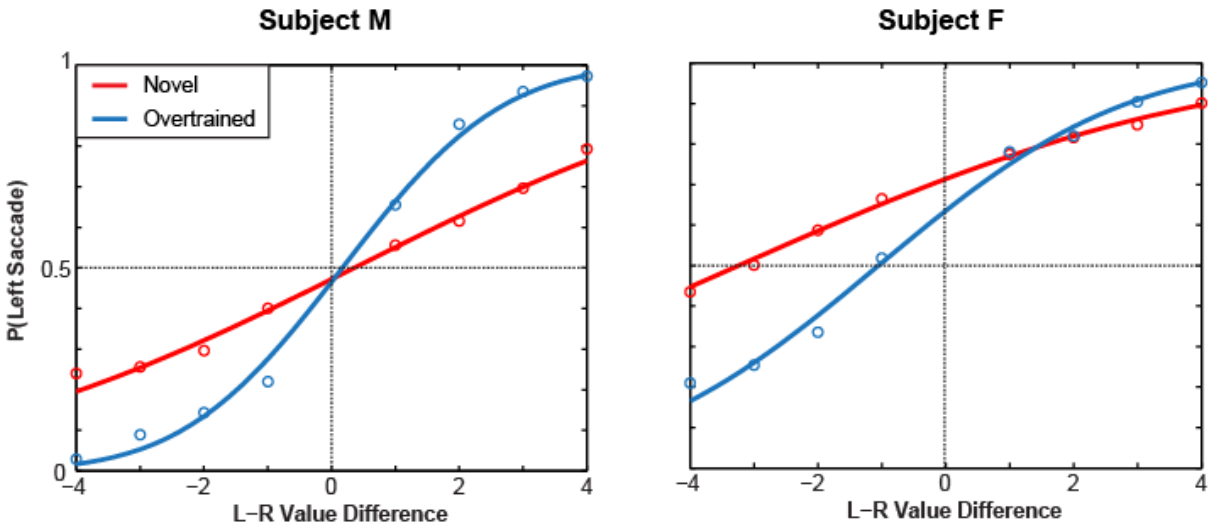

Figure 3.4: Psychometric function of first saccade direction. The probability of saccading to the left stimulus as a function of the value difference between the left and right stimuli collapsed across all Novel (red) and Overtrained (blue) trials. Note that Subject F showed a degree of left saccade bias which was particularly strong in Novel trials.

These results imply that subjects were carrying out reliable covert value comparisons within 170ms of stimulus onset that guided initial saccade direction. It also suggested that subjects' optimality in this covert value comparison varied depending on how familiar subjects were with the stimuli. To further explore this latter finding, we repeated the logistic regression using only the first 50 Novel trials of each session and compared the results to a regression using only the last 50 Novel trials of each session. This analysis showed that the influence of value on saccade direction was significantly greater in the last 50 Novel trials than the first 50 (linear hypothesis test of parameter estimates for value difference in first vs. last 50 Novel trials, M: $p<2x10^{-4}$, F: $p<3x10^{-4}$) (inset **Figure 3.5**). Furthermore, both subjects showed a positive correlation between the probability of saccading to the most valuable stimulus in Novel trials and session decile number (Person's correlation coefficient, M: $r = 0.291$, $p<6x10^{-4}$, F: $r = 0.257$, $p<3x10^{-4}$) (**Figure 3.5**). The same analysis on Overtrained trials found no significant differences between the start and end of the session (linear hypothesis test of parameter estimates of first and last 50 Overtrained trials and Pearson's correlation coefficient, $p>0.05$). From these results we can conclude that initial saccades can

be strongly value driven using covert attentional comparisons that span very short periods of time, and that the strength of this effect increases with learning the value of the stimuli.

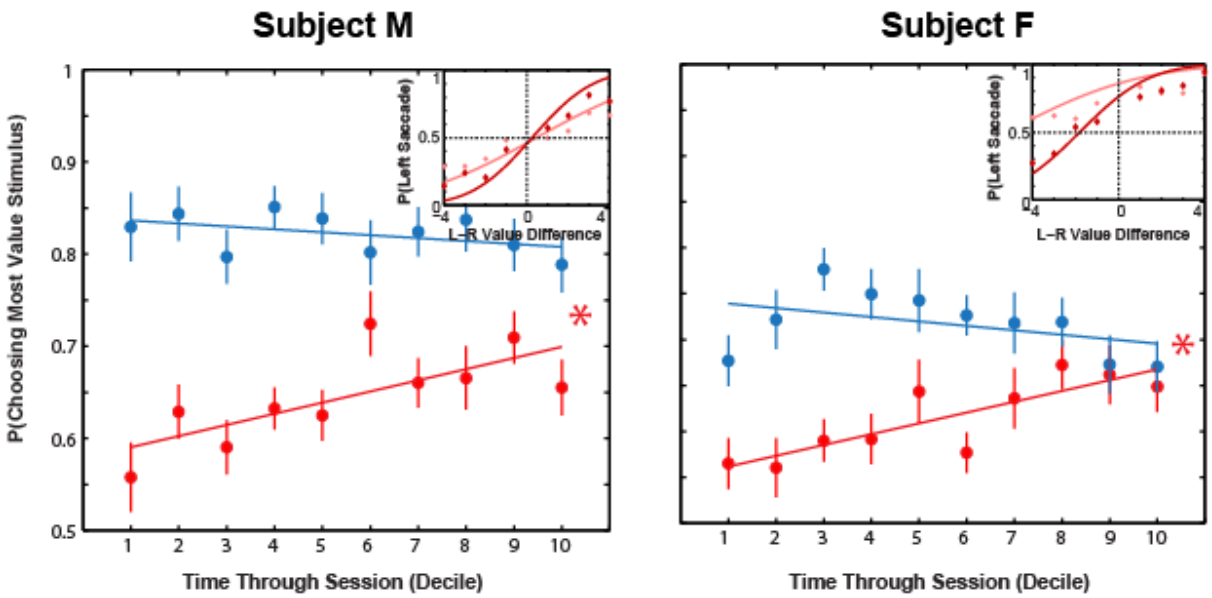

Figure 3.5: Probability of optimal first saccades across the session. Each dot is the mean probability of saccading to the most valuable stimulus across sessions for individual decile for Novel (red) and Overtrained (blue) trials. Vertical lines show the SEM for each decile. Relevantly coloured asterisks denote significant Pearson's correlation coefficients. Inset: Psychometric functions for left saccade probability the first 50 (light) and last 50 (dark) Novel trials collapsed across sessions.

With clear influences of value on both initial saccade direction and eventual choice behaviour, a natural question to ask is whether a single valuation and comparison system drives both of these behaviours or whether two separate systems are used. **Figure 3.6** depicts two possible evidence accumulation systems which consider information about both stimulus novelty and saccade direction in the final choice mechanism.

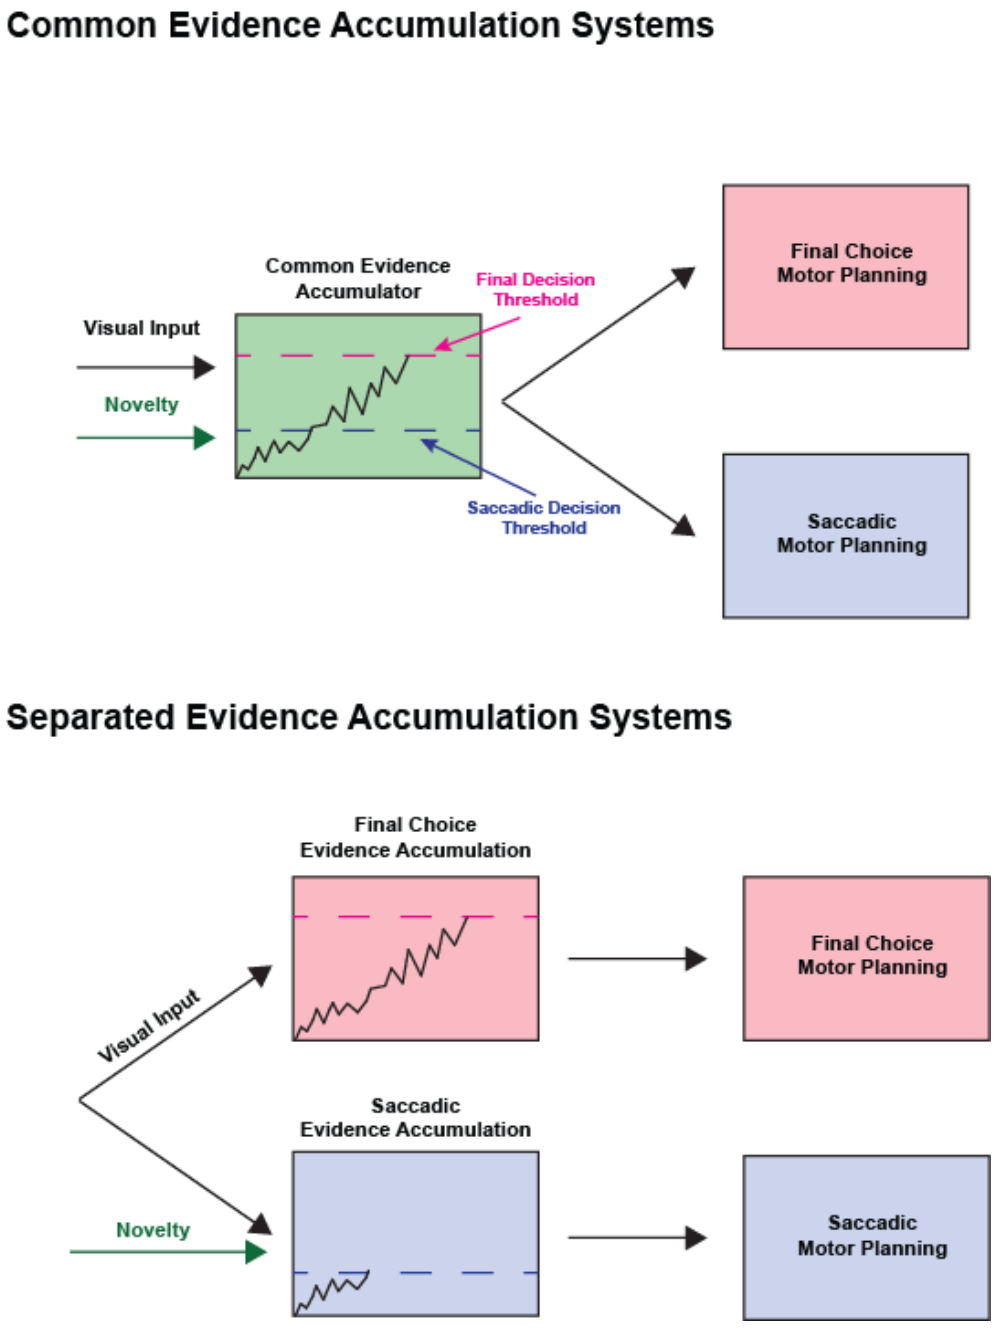

Figure 3.6: Empirical models of potential saccadic and motor evidence accumulation systems. The Common Evidence Accumulation Model (top) predicts that both visual input and novelty information enter a common evidence accumulator which has two evidence thresholds; a lower saccadic threshold (blue dashed line) and a higher final decision threshold (pink dashed line). Once each threshold is reached then relevant saccade and motor plans can be enacted respectively. The Separated Evidence Accumulation Model (bottom) predicts that there are separate saccadic and final motor choice evidence accumulation systems which both receive the same visual input but only the saccadic system receives novelty information.

To explore the relationship between saccadic and overt choice behaviour, we first calculated the proportion of optimal choices across the experiment (height of bars in **Figure 3.7**). We then asked within these optimal choices, what was the proportion of optimal vs non-optimal initial saccades (shaded area in **Figure 3.7** bars)? We reasoned that if a single valuation system was employed for both behaviours, the ratio between these two measures would stay constant across the session. However, in both subjects we observed the opposite result (**Figure 3.7**, black diamonds). The optimal saccade:optimal choice ratio *increased* significantly across the experiment (Pearson's correlation coefficient, M: $r = 0.32$, $p<0.0002$, F: $r = 0.16$, $p<0.02$).

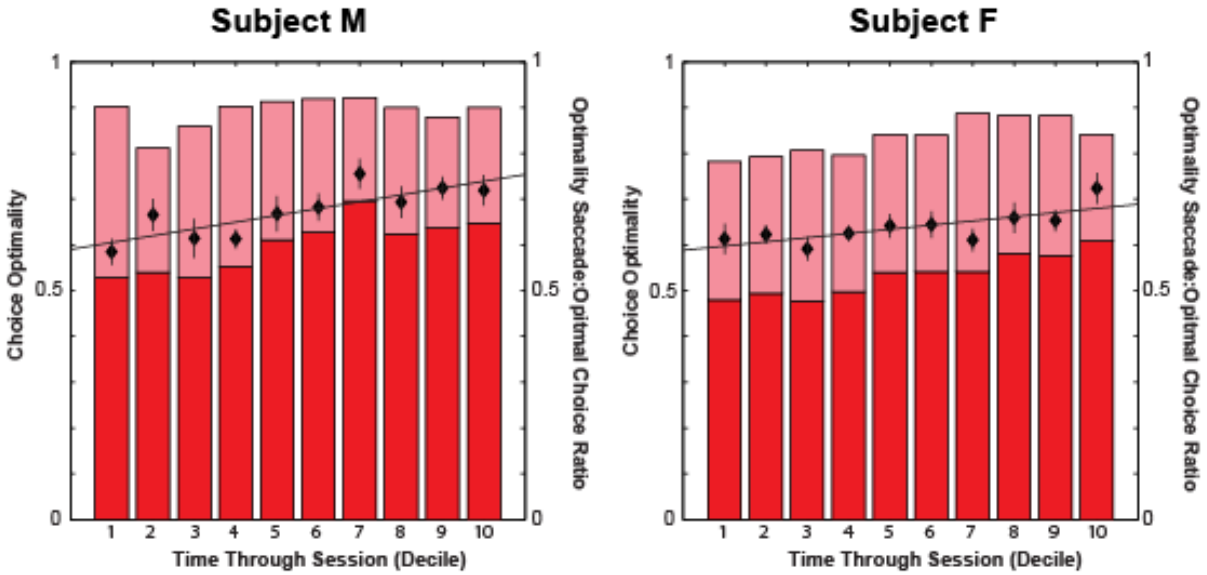

Figure 3.7: The ratio of optimal choices and optimal saccades. The total bar size indicates overall choice optimality per decile in Novel trials. The red sub-bar indicates the proportion of optimal choices that were associated with optimal saccades. Black diamonds indicate the ratio of optimal saccades to choice optimality. Asterisks indicate significant correlations between the ratio and decile number.

This result shows that the ability to direct saccades toward more valuable information takes longer to learn relative to overt choice behaviour. However, another way of viewing these effects is to focus on the first decile of a session; here subjects are effectively random in their direction of the first saccade, yet despite only 10+ exposures to each stimulus, both subjects choose the better stimulus close to 80% of the time, with little further improvement across the session. In other words, the very fast (~150ms) covert attentional system is slow

to learn to bias saccades toward the optimal stimulus and is still learning in this regard by the end of the session (Novel versus Overtrained performance in **Figure 3.5**). In contrast, subjects' overt choices exhibited little evidence of learning across deciles, implying our secondary conditioning protocol (**Figure 3.1A**) was effective in assigning value to stimuli. Given the large differences in optimal saccadic versus choice performance, we infer that valuation in these two systems is at least partially dissociable, meaning that optimal choice performance can recover from a suboptimal saccade via further information gathering mechanisms (see below). Notably, when the analysis was repeated on Overtrained trials (where no learning is occurring across the experiment) no significant correlation was found (Pearson's correlation coefficient, $p>0.05$).

Subsequent saccades reflect the value of non-fixated stimuli

Initial saccades tended to be towards more valuable pictures, implying a covert comparison process driving saccadic behaviour. In light of this putative process, we asked whether it was necessary to overtly attend to both stimuli in order to make an optimal choice. There were many trials in which subjects did not saccade to both stimuli: subject M fixated both stimuli on 56.2% of trials, and subject F did so on only 33.2% of trials. Both subjects tended to view more stimuli on Novel than Overtrained trials (Chi2 test, M: Chi = 108.1, $p<0.01$, F: Chi = 133.4, $p<0.01$).

We next asked whether the propensity to make a second saccade was influenced by the value of the non-fixated stimulus. We found the probability of making a second saccade (collapsed over trial types) was negatively influenced by the value of the fixated stimulus, but crucially, also *positively* influenced by the value of the non-fixated stimulus (logistic regression, M: $T_{(Fixated)}=-14.91$, $T_{(Non\text{-}fixated)}=23.57$, $p_{(Fixated)}<3\times10^{-50}$, $p_{(Non\text{-}fixated)}<7.5\times10^{-123}$,F: $T_{(Fixated)}=-35.95$, $T_{(Non\text{-}fixated)}=32.13$, $p_{(Fixated)}<5\times10^{-283}$, $p_{(Non\text{-}fixated)}<2\times10^{-226}$). In other words, both

stimuli had significant but opposing influences on whether a second saccade was generated (towards the non-fixated stimulus) before a choice was made.

Perhaps unsurprising, both subjects were significantly more likely to perform a second saccade when they initially fixated the less valuable of the two stimuli (binomial test, M: $p<1x10^{-15}$, F: $p<1x10^{-15}$). This could explain why subjects were more likely to make two saccades in Novel trials since they were also more likely to make suboptimal initial saccades. Subjects tended to make two saccades when they made errors in their initial covert valuation. But rather than execute a suboptimal choice following the suboptimal initial saccade, subjects could covertly attend the non-fixated stimulus peri- or post-first saccade, and then make a second saccade (if necessary) before making a choice. Analysis of final choice reaction time showed that subjects were also significantly slower to respond on two saccade trials (One-way ANOVA, M: F stat=750.92, $p<1x10^{-158}$, F: F stat=280.16, $p<1x10^{-61}$) which could be explained by subjects requiring more time to overcome erroneous initial evidence accumulation (reflected by suboptimal saccades) in order to maintain final choice optimality.

In addition, on two saccade trials, a linear regression onto first stimulus dwell time revealed a positive influence of fixated stimulus value and a negative influence of non-fixated stimulus value in both Overtrained and Novel trials (M: all unsigned T statistics>9.73, all p values$<2x10^{-22}$, F: all unsigned T statistics>7.00, all p values$<3x10^{-12}$) (**Figure 3.8**). Further, the value of the non-fixated stimulus had a significantly greater influence on dwell time than that of the fixated stimulus (linear hypothesis test comparing parameter estimates for fixated and non-fixated stimuli, M: $p_{(Novel)}<2x10^{-30}$, $p_{(Overtrained)}<2x10^{-62}$, F: $p_{(Novel)}<4x10^{-31}$, $p_{(Overtrained)}<2x10^{-34}$). It is only possible for non-fixated stimulus value to influence first fixation dwell time if subjects covertly processed the value of the non-fixated stimulus. This evidence coupled with the influence of value on first saccade optimality, is strongly supportive of the idea that direct fixation of stimuli is not necessary for evidence accumulation; covert attentional processes are fast and ongoing throughout the decision process.

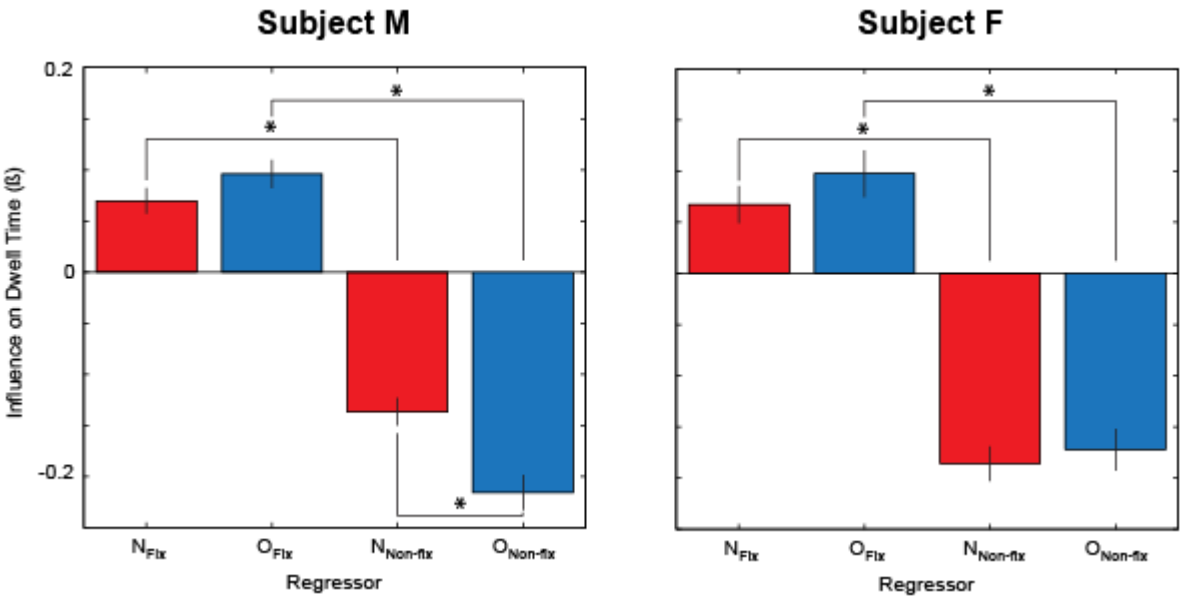

Figure 3.8: A regression of value on first stimulus dwell time on two saccade trials. A linear regression of the normalised first stimulus dwell time against fixated stimulus value in Novel trials (N_{Fix}), fixated stimulus value in Overtrained trials (O_{Fix}), non-fixated stimulus value in Novel trials ($N_{Non\text{-}fix}$) and non-fixated stimulus value in Overtrained trials ($O_{Non\text{-}fix}$). Asterisks indicate significantly different coefficients between trial types based on a linear hypothesis test (see text).

Another hypothesis for how subjects may solve this task is that their saccades may be driven by covert attentional processes, but that evidence accumulation for choices varies only after a given stimulus is fixated. To test whether this hypothesis better explains choice behaviour than covert attention, we directly compared the fits of two different logistic regression models to subjects' choices. In the first ('Covert' attention) model we used the actual left and right stimulus values for each trial to predict left choice irrespective of whether both stimuli were fixated. The second ('Overt' attention) model, was identical to the 'Covert' model with the exception that on any trial where the subject did not saccade to a given stimulus, that stimulus value was replaced with the *average* stimulus value (i.e. a rank of 3 out of a range of 1-5). We performed this analysis specifically on Overtrained trials in order to maximize effects of value on saccade behaviour and to avoid the potential confound of learning affecting choice behaviour. We then computed the model evidence for the two models for each session. Across sessions model evidence was greater for the 'Covert' model compared to the 'Overt' model (M: 14/14 sessions, binomial test, $p<0.007$, F: 15/20 sessions, binomial test, $p<0.03$). This result suggests that choice behaviour was better

explained by assuming that subjects were using covert information in order to make choices, rather than only accumulating evidence based on stimuli that had been overtly fixated. The covertly attended value of stimuli thus strongly influences both the probability of saccading towards those stimuli, and the probability of choosing them.

Saccades, but not choices, show a novelty bonus

Finally, we examined Mixed trials, where one Overtrained stimulus was compared against one Novel stimulus. Here, a key difference in behaviour was observed to Overtrained-only or Novel-only trials, namely that subjects preferred to make initial saccades toward the novel stimulus (**Figure 3.9A**). A logistic regression of the probability of the first saccade being directed toward the novel stimulus revealed a strong effect of both novel and overtrained stimulus value on saccade direction (M: $T_{(Novel)}$=25.49, $T_{(Overtrained)}$=-26.52, $p_{(Novel)}<3x10^{-143}$, $p_{(Overtrained)}<6x10^{-155}$, F: $T_{(Novel)}$=11.61, $T_{(Overtrained)}$=-16.00, $p_{(Novel)}< 4x10^{-31}$, $p_{(Overtrained)}< 2x10^{-57}$) (**Figure 3.9B**). The overtrained stimulus value had a significantly greater influence on initial saccades relative to novel stimulus value (linear hypothesis test on the parameter estimates for overtrained and novel value, M: $p<5x10^{-30}$, F: $p<2x10^{-4}$). Strikingly, however, this value based effect was diluted by a strong bias towards saccading to the novel stimulus first, irrespective of its value (logistic regression, M: T=17.35, $p<2x10^{-67}$, F: T=9.23, $p<3x10^{-20}$) (light red bar **Figure 3.9B**). This 'novelty bonus' shows that subjects *prefer to saccade to novel pictures*. If this is the case, do they also prefer to choose novel pictures? We found that they did not; a logistic regression of probability of choosing the novel stimulus showed that subjects' choices were both empirically more sensitive to value than subjects' initial saccades, and that the relative novelty bias was negative (i.e. they had a bias to *choose Overtrained stimuli*) (inset **Figure 3.9B**). Furthermore, examining subjects' choices throughout the session revealed a small bias to choose the *overtrained* stimulus in the first 50 Mixed trials of each session (binomial test, M: p<0.004, F: $p<2x10^{-4}$) and no bias at all by

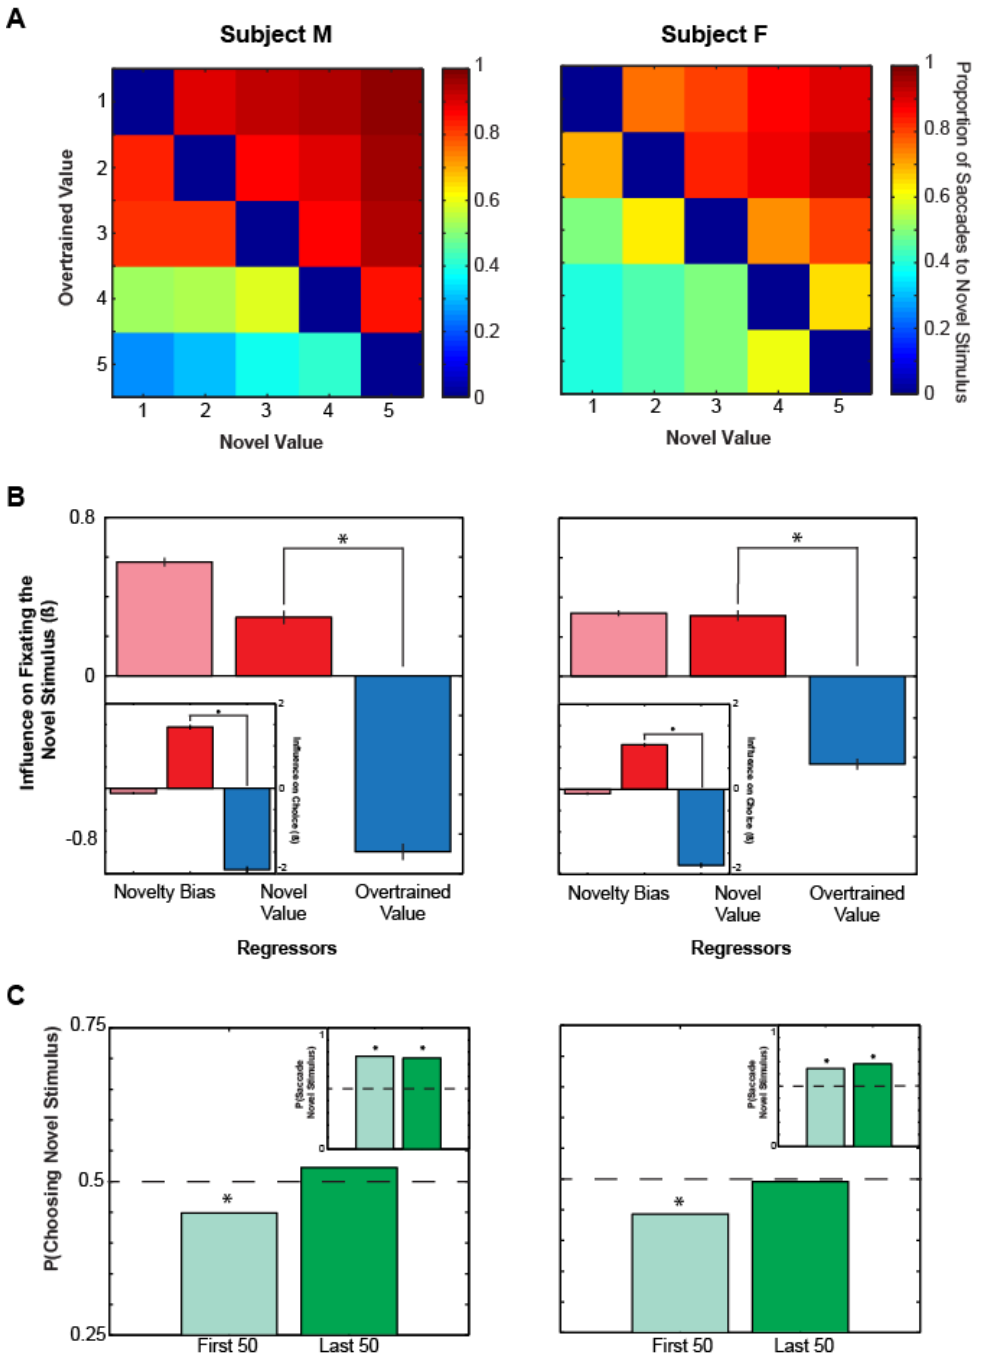

Figure 3.9: Influences of saccade and choice behaviour in mixed trials. **(A)** Heat maps of the probability of saccading to the novel stimulus in Mixed trials as a function of the value of the novel and overtrained stimulus. We never presented Mixed stimuli of equal value, hence the blue region along the diagonal. **(B)** Beta coefficients for a regression of novel and overtrained stimulus value predicting saccades towards the novel stimulus. Asterisks indicate coefficients that are significantly different in terms of unsigned magnitude. Inset: The same regression performed on the probability of choosing the novel stimulus at the final choice. **(C)** The mean probability of choosing the novel stimulus across the first and last 50 trials of a session, collapsed across sessions. Inset: The probability of the initial saccade being toward the novel stimulus across the first and last 50 trials of a session, collapsed across sessions. Asterisks indicate groups that are significantly different from chance.

the last 50 trials (binomial test, $p>0.05$) (**Figure 3.9C**). In contrast, an equivalent analysis of initial saccade direction revealed a significant bias towards the novel stimulus in both the first and last 50 Mixed trial of a session (**Figure 5.3C**, inset). This dissociation between saccadic and final choice behaviour provides further evidence that the two processes are separable from one another (see **Figure 3.6**).

Finally, both subjects were observed to improve choice optimality across the session in Mixed trials (Pearson's correlation coefficient, M: $r = 0.249$, $p<0.004$, F: $r = 0.205$, $p<0.004$). Yet surprisingly, subjects showed no decrease in the probability of initial saccades being toward the novel stimulus (novelty bias), despite the novel stimuli becoming increasingly more familiar as the session advanced (Pearson's correlation coefficient, M: $r = -0.106$, $p>0.2$, F: $r = 0.117$, $p>0.09$). This suggests that the subjects' saccadic novelty bonus reflects more of a categorical distinction between the novel and overtrained stimuli on Mixed trials, rather than the relative novelty between these stimuli.

.

Discussion

In this chapter we have shown that saccadic eye movements and motor responses of primate subjects during value guided decision making are heavily influenced by value and novelty through covert attentional processes.

Covert Attention Influences Saccades

This finding that our subjects' initial saccades tended to be directed towards move valuable stimuli agrees with the findings of Yasuda et al. (2012), who also found this effect as well as a strong learning effect which plateaued after approximately 300 trials spread across 5 different sessions. Importantly, their task did not require subjects to perform any action in order to obtain rewarded outcomes, but instead just allowed them to saccade freely around the screen. The current study suggests that given a smaller set of stimuli (but a larger 'value space') value-driven saccades can be deployed within as few as 15 initial exposures to the stimulus-outcome relationship. Results from this study and human saccade distractor studies all feed into the idea that non-physical stimulus properties modulate covert attention, which in turn drives initial saccade behaviour. The latency of this unconstrained value guided saccade in the current study was <170ms, the first report of such value-guided latencies that we are aware of in the primate literature.

Physical properties of stimuli can attract overt attention (Nothdurft, 2002, Theeuwes, 1992, Itti and Koch, 2001). In this study, we controlled for luminance but not for other properties that may influence visual salience such as contrast or colour. Nevertheless, the stimulus to value assignment was random and although it is possible that there may be some weak relationship between value and visual salience for particular stimulus sets, there

would be no systematic relationship between salience and value across the experiment, where subjects learned at least 14 different novel stimulus sets.

Covert Attention During Decision Making

A particularly appealing hypothesis for why saccadic behaviour is not random is that by using covert attention to guide our eyes we can orient our overt attention to relevant stimuli in our environment (Gottlieb, 2012, Gottlieb et al., 2014, Anderson, 2013). The results of the current study significantly extend this hypothesis by showing multiple influences of covert attention on saccadic and choice behaviour. Signatures of covert attention can be seen in the probability of optimal initial saccades, the first stimulus dwell time, the probability of subsequent saccades and at final choice. This suggests that subjects are likely to use covert attention to accumulate evidence during the entire decision making process. Indeed, the evidence in the current study suggests that in this task, subjects solve the decision largely (if not solely) through covert attention.

Given the above, if overt saccades to stimuli are not necessary for evidence accumulation, why are they present at all? We postulate that the use of covert attention to bias overt attention is an inherent property of the attentional system; covert attention may be optimal for distinguishing relatively simple features or categorizing stimuli very quickly, whereas in more complex naturalistic environments with complex visual features, direct fixation of stimuli may be critical for identifying stimulus features or contextualizing stimuli within the environment. The fact that monkeys exhibit markers of covert attention is unsurprising given neurons are modulated by the value of peripheral cues even when the subject is required to maintain central fixation, and thus well before an overt saccade to a stimulus is initiated (Kennerley et al., 2009, Padoa-Schioppa and Assad, 2006, Louie et al., 2011).

Such covert evidence accumulation processes are incompatible with the basic assumption of attention drift diffusion models (aDDMs), which is that subjects only use overt attention to bias the accumulation of evidence (Krajbich et al., 2010, Krajbich et al., 2012). In tasks which test these models, subjects show clear deliberative behaviour, where they make multiple saccades between stimuli (Krajbich et al., 2010, Krajbich et al., 2012, Krajbich and Rangel, 2011, Towal et al., 2013). The fact that our subjects do not regularly overtly attend to both stimuli indicates that there is very little deliberation about their choices. There are many reasons for this apparent dissociation between primate and human behaviour. Perhaps the most obvious reason is that humans in aDDM studies have very little exposure to a large set of stimuli (or the task) prior to data collection, whereas in the current study, even for novel stimuli, the subjects had approximately 15 exposures before data collection, and subjects knew the task design very well. As shown in this study and previous studies, subjects become better able to use covert attention with increased stimulus exposure (Yasuda et al., 2012), and therefore require diminishing amounts of overt attention to perform the task. Another reason why overt attention may be less important in our task relative to aDDM studies is that in aDDM studies, subjects are generally presented with multi-attribute stimuli (e.g., a picture of a candy bar) and with choices that are relatively similar in subjective value. In contrast, our stimuli had a clearly defined value based on a single attribute, making choices easier overall and potentially removing the need for long deliberation. Many other differences in human and primate task designs may also contribute to the behavioural differences we observe, though our results suggest future aDDMs models should incorporate covert attentional processes.

Circuits for Saccadic and Motor Value Comparison

Although further insight into whether saccadic and motor choice is driven by the same neural circuits may be achieved by neurophysiological investigation, our current

results allow us to make some inferences about this question. Firstly, the speed of saccadic behaviour implies the engagement of brain circuits which receive visual information quickly. This circuit seems unlikely to be the commonly described decision circuit involving prefrontal cortex because, although rarely explicitly described in the literature, many studies seem to show a relatively slow encoding of value in prefrontal neurons, on the order of 200-500ms post- stimulus onset (based upon single cell examples and population analysis from Kennerley et al. (2009), Hayden and Platt (2010), Strait et al. (2014), Roesch and Olson (2003)).

Instead, the subcortical saccadic system - including the Caudate, SNpr and SC - is a strong candidate circuit for these covert attentional processes. Firstly, this system lies considerably closer to visual system than frontal circuits (Munoz, 2002), and regions directly connected to this circuit such as visual cortex respond to visual information well within 100ms of stimulus onset (Schmolesky et al., 1998, Thorpe and Fabre-Thorpe, 2001). Secondly, Substantia Nigra pars Reticulata and the head of the caudate (both critical elements in saccadic control circuitry) have been shown to be capable to discriminate the value of stimuli with a relatively short latency (Kim and Hikosaka, 2013, Yasuda et al., 2012, Hikosaka and Wurtz, 1985, Kori et al., 1995). Given the SC, a main generator of saccades, receives visual input from both the retina and visual cortex, in addition to projections from other eye-related areas such as lateral intraparietal cortex (LIP), frontal eye field (FEF) and SNpr (Hikosaka et al., 2000), this region is ideally positioned to generate fast saccades based on simply visual features, as well as slower saccades based on more complex visual information including stimulus-reward associations.

However, saccade latency in itself is not final proof that saccadic and motor evidence accumulation are separate, because one could envisage a single drift process with two separate evidence bounds, one high (motor choice) and one low (saccadic choice, **Figure 3.6**). In this case one would expect to see fast, but potentially suboptimal saccades, yet because further evidence is required to reach the choice bound, final choices would be

slower and more accurate, in agreement with our results. However, although the concept of evidence bounds explain behavioural decision parameters well (Ratcliff, 1976, Ratcliff and McKoon, 2008), even in regions such as FEF and LIP which show evidence accumulation signals, the firing rate threshold for movement initiation appears static even when the speed-accuracy trade-off is manipulated (Hanks et al., 2014, Heitz and Schall, 2012). This suggests that the drift-diffusion process of evidence accumulation may not map directly to neuronal circuits, thereby making a single accumulator with two bounds less physiologically plausible.

Novelty Bonus

Our results show a strong saccadic bias towards overtly attending to novel stimuli in Mixed trials. Novel stimuli are known to capture both human and monkey attention (Johnston et al., 1990, Wilson and Goldman-Rakic, 1994, Goh et al., 2009, Foley et al., 2014). Studies have also found that midbrain dopaminergic neurons respond to novel stimuli (Horvitz et al., 1997) and that dopamine manipulation can perturb novelty responses (Costa et al., 2014). In attempting to reconcile this with dopamine's role in reinforcement learning, Kakade and Dayan put forward the idea of novelty bonuses, in which novelty itself can have its own inherent value in order to promote exploratory behaviour and reduce uncertainty. Although this idea is primarily oriented toward choices, it may be possible to similar inferences about saccades. Interestingly, our subjects only seem to assign an inherent value to saccading to novel stimuli, not to choosing them. Not only does this provide our clearest evidence for the functional separation between saccadic and motor evidence accumulation, it also begs the question what do subjects gain from this? One explanation could be that subjects are relatively uncertain about novel stimuli due to a lack of experience, therefore subjects could reduce uncertainty by directing overt saccades to novel stimuli (Dayan et al., 2000). However, if this is true, the subjects are only reducing identity uncertainty, because subjects rarely go on to subsequently choose the novel stimulus if it is the least valuable choice.

Similarly, it may be the case that subjects cannot easily discriminate the visual properties (and therefore their reward assignment) of novel stimuli in our task without direct foveation. If so, exhibiting a preference for fixating novel stimuli may be an optimal strategy for reward-based visual discrimination. As such, novel stimuli may be more valuable in terms of the information about the world that they yield. However, in contrast to the theory which hypothesises a rapid extinction of the novelty bonus with repeated exposure (Kakade and Dayan, 2002), we found the novelty bonus effects were consistent across the session in Mixed trials. This suggests that at least within our task which required value-based choices between novel versus overtrained stimuli, our subjects did not experience the novel stimuli sufficiently enough to allow covert attentional mechanisms to direct saccades to the optimal, rather than novel, stimulus.

A final potential explanation for the observed saccadic novelty bias is that the overt attentional system is biased towards salient objects in the environment (Gottlieb et al., 2014). If this concept holds true then things that are relatively novel may therefore be more salient to subjects when compared to relatively over exposed objects (even if these are of high value). In support of this idea, recent neurophysiological recordings from LIP (an area known to directly influence saccades) have implied that neuronal responses correlate with stimulus salience rather than just value (Leathers and Olson, 2012). Furthermore, LIP neurons respond to novelty and these responses do not diminish over time (Foley et al., 2014). It is unclear whether LIP could play such a role in the current study given the extremely fast latencies of the saccades, though LIP latencies are typically faster than PFC latencies. Finally, the fact that the novelty bonus was only present in saccade behaviour and was actually negative in choice behaviour is perhaps the strongest evidence put forward in this chapter that the decision circuit governing evidence accumulation to bias saccades is separate from the evidence accumulation circuit governing final choice.

In conclusion, in this chapter we have presented behavioural findings from a simple primate free saccade decision making study which demonstrates that subjects use covert

attention to guide both overt attention and choice. The use of covert attention increases with increasing familiarity with presented stimuli. However, relative novelty holds a strong influence on overt attention even though it does not bias choice. Taken as a whole, these results imply the employment of at least partially separable evidence accumulation systems for decision-making. The first is a fast, moderately accurate process which is biased by novelty and is responsible for orienting attention towards more informative stimuli. The second is a slower, more accurate process, which governs the final motor response of subjects based on the most valuable information at hand.

Chapter 4: Information Gathering Behaviour in Multi-attribute Decision Making

Introduction

In the previous chapter we showed that the saccadic system exhibits behavioural reflections of learning the value of stimuli in the environment and biasing visual attention towards more valuable stimuli within relatively few experiences of outcomes. At the same time, subjects showed the ability to make choices using information gathered by covert attention despite being free to use overt attention if they should so choose. One explanation of the latter finding is that the task was very simple and therefore there was very little cognitive demand associated with using covert attention. Therefore an obvious follow up question is: are these covert information gathering effects still present when the task or environment gets more complicated? This is one of several questions that we will try to examine in the following chapter. We will also examine how information gathering behaviour is influenced by task parameters and make inferences about the cognitive processes which govern information gathering using the behavioural results of an information gathering paradigm. Firstly we will show that covert attention still plays a large role in decision making even in more complex tasks. Secondly, we will demonstrate that by changing the way in which information is presented to subjects we can influence decision making behaviour. Finally, we will show that subjects choose both the amount of information to gather as well as the manner in which to gather it based upon the current state of evidence associated with a decision.

Gathering information is an important step to making optimal decisions (Furl and Averbeck, 2011). But what drives us to gather information? Some have argued that information gathering is driven by a need to reduce uncertainty about the environment, particularly when faced with novel stimuli (Feldman and Friston, 2010, Kakade and Dayan, 2002). Others have suggested that it is a process for maximising average reward while minimising average effort (Furl and Averbeck, 2011). These ideas are by no means mutually exclusive but outright evidence in support of either remains elusive.

A simple extension of the latter idea is that subjects bias their attention towards information (e.g., stimuli) in the environment which they believe will inform future choice; they attend to stimuli which have salience and value properties which are pertinent to the decision at hand (Gottlieb et al., 2014). Evidence for this idea comes from many levels. Firstly, neurons in the saccadic basal ganglia system (namely Caudate nucleus and Substantia Nigra pars Reticulata) encode and retain the value of stimuli over many days (Yasuda et al., 2012, Kim and Hikosaka, 2013). These circuits are known to be critical for the control of gaze and therefore by extension overt attention (Gottlieb et al., 2014, Yamamoto et al., 2012, Hikosaka et al., 2006). Furthermore, regions such as Lateral Intraparietal Area (LIP), which are known to directly influence saccades, have neuronal response patterns to stimuli consistent with the encoding of salience (Gottlieb et al., 2014, Leathers and Olson, 2012), and therefore may provide an information signal to saccadic initiators. The presence of these computations within brain regions provides a plausible pathway by which attention may be guided towards stimuli that are perceived to be important in the environment. Furthermore, information can have its own value which may be represented in OFC neurons (Blanchard et al., 2015). Importantly, this coding appears to be orthogonal to value coding in OFC implying (as expected) that information value is not integrated with decision value when choices are made (Blanchard et al., 2015). This is to be expected if one assumes that decisions about what information to gather are different from decisions about what to eventually choose, as our behavioural data in Chapter 3 might indicate.

One relatively unexplored area of behavioural neuroscience is the question of how value information is compared in complex environments. In most decision making studies the way in which information is gathered during decision making is irrelevant because decision are typically made within a single attribute or multiple attributes are merged into a single stimulus (for example Kennerley et al. (2009), Kim et al. (2008), Strait et al. (2014), Padoa-Schioppa and Assad (2006)). However, in a natural environment decisions are often made between many options which may all have multiple attributes associated with them, making the problem of information gathering a pertinent one. The domain of psychology has been studying this problem for many decades. When presented with complex multi-attribute and multi-alternative choices, human information seeking behaviour is well characterised. As the number of attributes and alternatives increase, subjects unsurprisingly tend to gather increasingly less of the total available information (Payne, 1976, Cook and Swain, 1993, Lohse and Johnson, 1996). Increasing the number of attributes also increases the variability in choices among subjects, decreases the optimality of choices and yet increases the decision makers confidence in their decision (Slovic and Lichtenstein, 1971, Payne, 1976). It has been hypothesised that this task complexity effect may arise from subjects switching from linear decision rules to choice heuristics in order to reduce cognitive demand during complex decisions (Payne, 1976).

Furthermore, as task complexity increases, subjects move away from performing ‘within option’ information comparison towards ‘within attribute’ comparison (Sundström, 1987). The purpose of these ‘within attribute’ comparisons is to ‘screen’ the large dimensional option space based on the attribute(s) of the highest importance, thereby reducing the number of options – and the total amount of information - needing consideration, thus avoiding the evaluation and integration of large volumes of information (Weenig and Maarleveld, 2002, Kerstholt, 1992). When time constraints are added to these multi-attribute, multi-alternative problems, this phenomenon of ‘screening’ is observed to

increase (Weenig and Maarleveld, 2002). Mathematical descriptions of choice processes governed by attribute screening have also been developed (Tversky, 1972).

However, shifting from 'within option' to 'within attribute' information comparison may not be limited to decision contexts having a large attribute or option space. Arieli (2011) presented human subjects with a binary multi-attribute decision making task where they had to integrate explicit probability and magnitude information in order to choose the option with the highest overall value. When the probabilities and magnitudes were relatively simple (e.g. $3000 at 15% against $4000 at 11%) subjects preferred to make 'within option comparisons'. However, when the values became very complex (e.g. $637 at 64.9% versus $549 at 73.2%) subjects shifted towards 'within attribute' comparisons. Therefore the body of the psychology research suggests that 'within attribute' information gathering arises from a need to reduce computational load when decisions become difficult, such as when option values are very similar, or with increasing option or attribute space of the decision context.

The neural processes that drive information gathering behaviour remain poorly understood. A study by Fellows (2006) provides a glimpse into how brain regions support information gathering behaviour. Fellows took a group of vmPFC lesion patients, a group of LPFC lesions patients and a healthy control group and observed their information gathering behaviour during multi-attribute, multi-option decision-making. It was found that although controls and LPFC lesion patients adopted a 'within attribute' comparison strategy, vmPFC lesion patients instead used a 'within option' comparison strategy. It was also noted that although the vmPFC lesion patients gathered as much information - and took as long to decide - as the other groups, they often made a different final choice to the other two groups. One interpretation of these findings is that vmPFC is either critical for facilitating information gathering processes, or this region plays a critical role in 'within attribute' comparisons.

Therefore given the state of the psychology literature pertaining to multi-attribute decision-making, we aimed to characterise the information gathering and choice behaviour

of our two monkeys as they performed a multi-attribute decision-making task where we manipulated how information was presented to them. The ultimate aim of this task is to better understand the kinds of neuronal computations which take place during information gathering behaviour (Chapter 5).

Methods

The same two male rhesus monkeys (*Macaca mulatta,* subjects F and M) used in the single attribute decision making experiment (Chapter 3) concurrently performed two variants of a multi-attribute decision making task, referred henceforth as the 'Simultaneous' and the 'Information Gathering' task. A detailed description of the task structure can be found below. Full details of the behavioural protocols (e.g., eye tracking, behavioural acquisition equipment, fluid control protocol) can be found in Chapter 3. Typically, subjects would perform a 3-4 day sequence of stimulus-outcome learning ('day 1', Chapter 3) followed by an information gathering experiment ('days 2-4, Chapter 4).

Task

The structure and timeline of the task are shown in **Figure 4.1**. Subjects initiated the trial by maintaining saccadic fixation on the centre of the screen and central fixation of the joystick for 500ms. Once this was achieved two options were presented on the screen (seven visual degrees left and right of centre). Each option consisted of two pre-learned picture cues assigned to two different value attributes; probability of reward (10%, 30%, 50%, 70%, 90%) and magnitude of juice reward (0.15AU, 0.35AU, 0.55AU, 0.75AU, 0.95AU). The actual reward magnitude was calculated by multiplying the arbitrary unit by the maximum available duration for each subject (M: 2500ms, F:2750ms). Each cue was 7 visual degrees above/below the horizontal centre of the screen. Reward magnitude was varied by manipulating the length of time a reward pump was driven. The overall (integrated) value of each option was defined as the product of the probability and magnitude cues. In 'Simultaneous' trials (**Figure 4.1C**), all of the cues were presented at once and subjects were free to saccade around the screen and make their choice (by L/R joystick movement) within 3000ms of cue onset. If subjects did not respond within this time then a short time out

was given and a new trial was initiated. Once a response was made all of the cues were uncovered (for 500ms for Subject F and 1000ms for Subject M) following which juice reward feedback was given based on the probability and reward magnitude chosen by the subject.

'Information Gathering' trials were identical to 'Simultaneous' trials except that stimuli were now presented sequentially and under experimenter control. In 'Information Gathering' trials, all four picture cues were covered up by grey squares with the exception of one which was covered by a blue square. The blue square informed the subject of the required location for a saccade. Once the subject fixated the blue square, the picture cue replaced the blue cue and the subject was required to continuously fixate this location for 300ms. If continuous fixation was not achieved within 1200ms the trial was aborted and subjects received a short timeout. Once this fixation period was finished, the cue was covered with a grey square, and a second blue square was presented at a different location. The position of this blue square indicated to the subject the type of trial being experienced. If the blue square was for the second cue of the same option, subjects were in an 'Option' trial; if the blue square was for the same attribute cue of the second option then this was an 'Attribute' trial. Selection of trial types was pseudorandom. The subject was again required to acquire and maintain fixation of the second cue for 300ms before it was also covered up by a grey square. After this point, the subjects were now free to either i) choose an option using a joystick movement (left/right) based on the value of the currently known information or ii) view one or both of the remaining cues (in any order) before making a choice, with the third cue requiring 300ms of uninterrupted fixation before the fourth cue could be viewed. Importantly, however, they were prevented from viewing any cue that they had already seen.

'Option' and 'Attribute' trials were pseudorandomly interleaved during blocks of 50 trials. Between each of these blocks subjects were presented with a block of 25 'Simultaneous' trials.

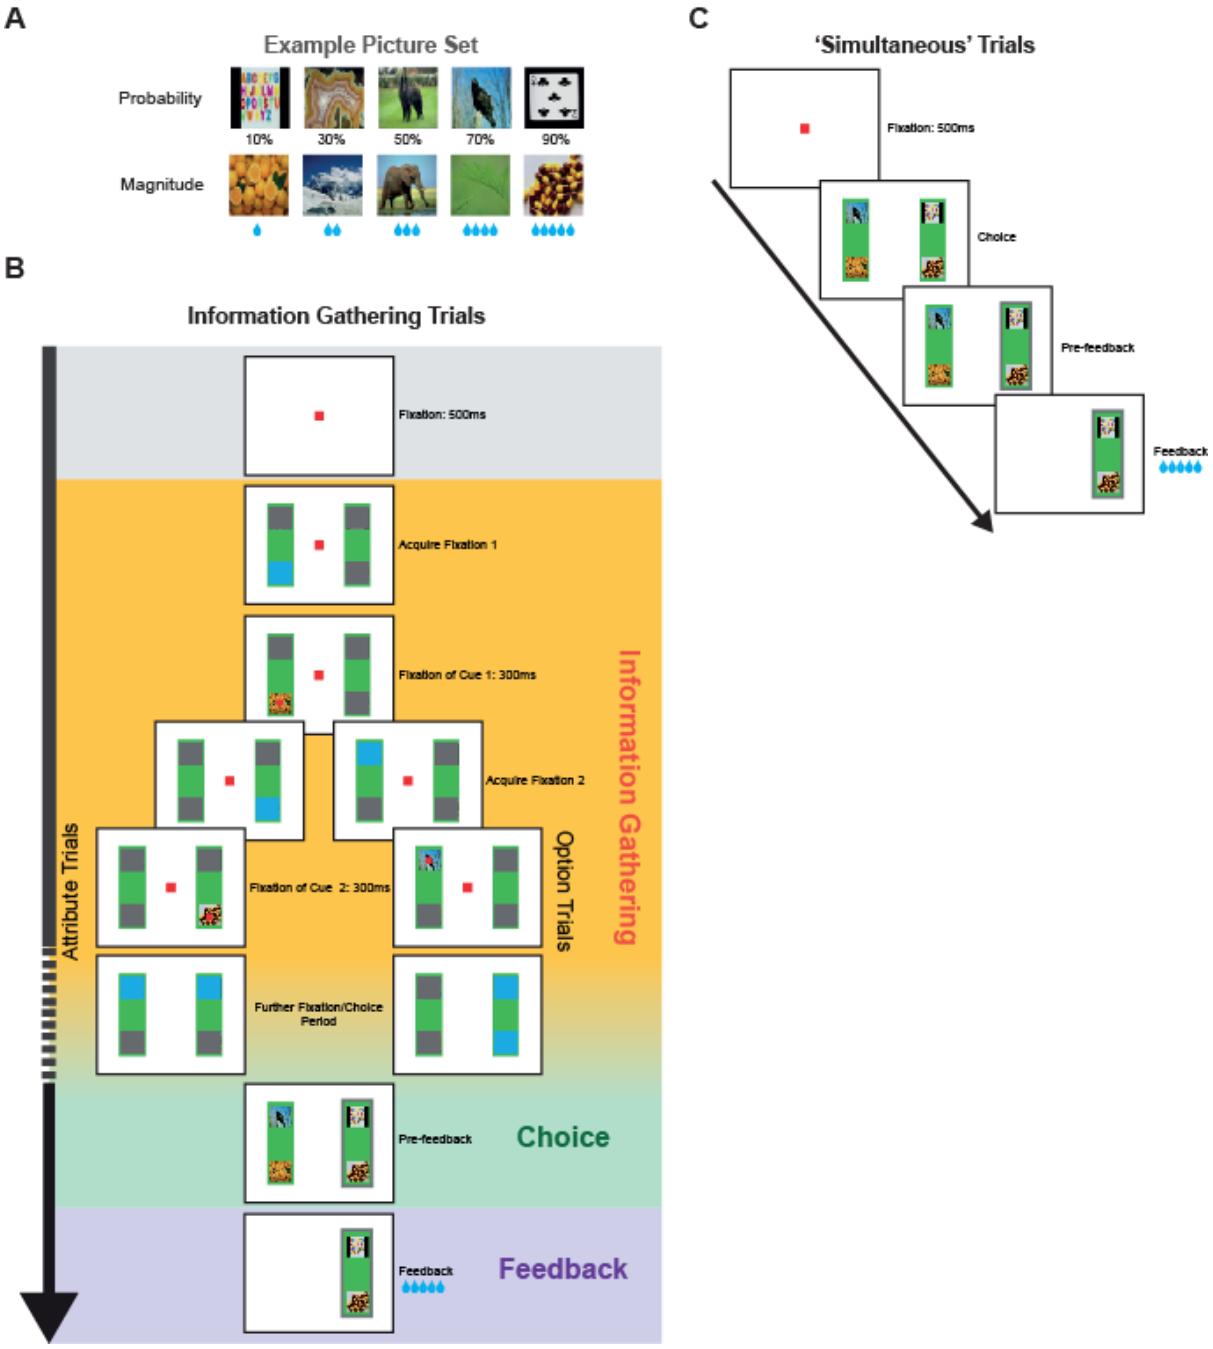

Figure 4.1: Task structure. **(A)** An example picture set. **(B)** A diagrammatic representation of the 'Information Gathering' task structure. Subjects saccade around the screen in order to gather information about two options. The initial two saccades are dictated to the subject but after this point subjects are free to gather as much information as they want and to choose at any time (see Methods for a more detailed description). **(C)** A representation of the timeline of 'Simultaneous' trials. Subjects were free to saccade around the screen and were given up to 3000ms to make a choice.

Behavioural Analysis

All behavioural analyses were performed using self-written scripts in MATLAB (MathWorks, USA). The most basic analysis performed was to characterise subject choice behaviour using logistic regression. This was achieved by fitting the data to **Equation 4.1** where YP is the probability of observing an event, b_0 was the inherent tendency to observe that event irrespective of any other variables, b_n was the weighting coefficient and x_n was the regressor:

$$\boldsymbol{YP = \frac{1}{1+e^{-(b_0+b_1x_1+b_2x_2+\cdots+b_nx_n)}}} \qquad \textbf{Equation 4.1}$$

We collapsed the data across all sessions in order to achieve the maximum possible power for the regression. **Table 4.1** contains a list regressors used to characterise the factors that predict subject left choice behaviour in 'Simultaneous' trials.

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Left-Right Probability Difference	Value	4	First Saccade Direction (Left-Right)	Choice Bias
2	Left-Right Magnitude Difference	Value	5	Constant	Left Choice Bias
3	Orthogonalised Left-Right Expected Value Difference	Value			

Table 4.1: A list of regressors and their interpretations used in the logistic regression of left choice probability.

In order to test whether the direction of the subjects' first saccade was driven by value, we used a regression model which tested the probability of left saccades based on a model which used the cue values at each of the four spatial locations, collapsed across

attribute (see **Table 4.2**). We also used the same regression model to test the influence of spatial position on final choice.

#	**Regressor**	**Interpretation**	#	**Regressor**	**Interpretation**
1	Top Left Value (TL)	Value	4	Bottom Right Value (BR)	Value
2	Top Right Value (TR)	Value	5	Constant	Left Saccade Bias
3	Bottom Left Value (BL)	Value			

Table 4.2: A list of regressors and their interpretations used in the logistic regression of left saccade and choice probability based on the value information at each spatial location.

We performed the following linear hypothesis test to work out whether certain specific locations had a greater influence on saccades (or choice) than others:

Top Coding > Bottom Coding = (TL +TR) – (BL + BR)

In 'Information Gathering' trials, we characterised subject choice behaviour using the regression model defined in **Table 4.3**. We then used linear hypothesis tests to examine whether there was a significant different between magnitude and probability regressors or trial type regressors.

#	Regressor	Interpretation	#	Regressor	Interpretation
1	'Option' Left-Right Probability Difference	Value	5	'Attribute' Left-Right Probability Difference	Value
2	'Option' Left-Right Magnitude Difference	Value	6	'Attribute' Left-Right Magnitude Difference	Value
3	'Option' First Saccade Direction (Left-Right)	Choice Bias	7	'Attribute' First Saccade Direction (Left-Right)	Choice Bias
4	'Option' Trials Constant	Left Choice Bias	8	'Attribute' Trials Constant	Left Choice Bias

Table 4.3: A list of regressors and their interpretations used in the logistic regression of left choice probability on 'Option' and 'Attribute' trials.

Next we aimed to characterise whether subjects made immediate choices after viewing three cues or whether they chose to gather the last piece of information. Therefore in 'Option' trials, we used a logistic regression model where the dependant variable was whether the choice happened after the third cue and the independent variables were the sum of the values of the first two cues and the value of the third cue. In 'Attribute' trials the independent variable was the absolute difference between the value of the first two cues and the value of the third cue.

After subjects viewed the first two cues, they were free to saccade to either of the remaining two locations. We were therefore interested in whether the subjects' information gathering strategy (i.e., third saccade) was driven by value. A preliminary analysis of third saccades revealed stereotypical behaviour on 'Option' trials (i.e., Subject M would saccade horizontally on the majority of trials: M = 67.7%, while subject F would always make a saccade towards the bottom stimulus: F = 100%) thus rendering insufficient variance to run

this model. We therefore focussed this analysis only on 'Attribute' trials, and performed a logistic regression of the probability of making a vertical saccade toward cue three (as opposed to a diagonal saccade) against the value of the first and second cues.

In 'Simultaneous' trials we examined whether the number of cues fixated was a function of cue value split by relative locations on the screen by using a regression model shown in Table 4.4:

#	**Regressor**	**Interpretation**	#	**Regressor**	**Interpretation**
1	Value Current Fixated Cue	Value	5	Option 1 Orthogonalised EV (fixated side)	Integrated Value
2	Value Cue Vertical to Fixation	Value	6	Option 2 Orthogonalised EV (non-fixated side)	Integrated Value
3	Value Cue Diagonal to Fixation	Value	7	Constant	Average Cues Viewed
4	Value Cue Horizontal to Fixation	Value			

Table 4.4: A list of regressors and their interpretations used in the linear regression of number of cues viewed and the logistic regression of the probability of option type saccades during the second saccade.

Finally, for several analyses (e.g., dwell times, reaction times) we used linear regression using **Equation 4.2** where Y was the dependant variable, b_0 was the constant term and X_n were the regressor and were weighted by coefficients b_n:

$$\boldsymbol{Y = b_0 + b_1 X_1 + b_2 X_2 + \cdots + b_n X_n} \qquad \textbf{Equation 4.2}$$

Eye position data was pre-processed by removing (and interpolating) data which

contained eye blinks. In ‘Simultaneous’ trials, we defined a stimulus as ‘viewed’ if the subject’s eye position entered an 8 x 8 visual degree area around the centre of each stimulus. The direction of saccades was defined using a saccade detection algorithm which detected changes in eye position that were faster than 7 degrees/s and lasted longer than 20ms (Engbert and Kliegl, 2003).

Results

Two subjects, M and F, were trained to perform two variants of a multi-attribute decision making task, the 'Simultaneous' task and the 'Information Gathering' task. The 'Information Gathering' task contained two types of trials; 'Option' and 'Attribute' which were pseudorandomly selected on every trial (**Figure 4.1**). Subject M performed 7632 'Simultaneous trials', 7187 'Option' trials and 7064 'Attribute' trials over 32 recording sessions. Subject F performed 5524 'Simultaneous trials', 4823 'Option' trials and 5040 'Attribute' trials over 25 recording sessions. Choice optimality was defined by whether the subject chose the option with the highest expected value (EV: probability value x magnitude value) irrespective of whether the subject viewed and fixated all available information. **Table 4.5** shows the choice optimality for each subject across trial types.

Trial Type	Subject M	Subject F
'Simultaneous'	85.9% (0.77)	80.7% (0.79)
'Option'	81.8% (0.45)	78.4% (0.62)
'Attribute'	80.2% (0.60)	78.4% (0.65)

Table 4.5: Choice optimality for each subject and trial type. Choices were considered optimal if subjects chose the option with the highest EV on a given trial.

The first part of this chapter will consider only behavioural data from 'Simultaneous' trials, the results of which will then inform the analytical approach used during the second part of the analysis with respect to 'Information Gathering' trials. Data were collapsed across all recording sessions to provide the maximum available power.

Part 1: Simultaneous Trials

Subjects Equally Consider both Magnitude and Probability Information When Making Decisions

An important aspect of the behavioural task was that in order to choose optimally (from the point of view of utility) subjects had to select the option with the highest Expected Value (EV), which in our task was defined as the product of reward probability and reward magnitude. Empirical analysis of the choice data suggested that both subjects were more likely to choice the left option as left EV increased and less likely to choose the left option when right EV increased in any given trial (**Figure 4.2**).

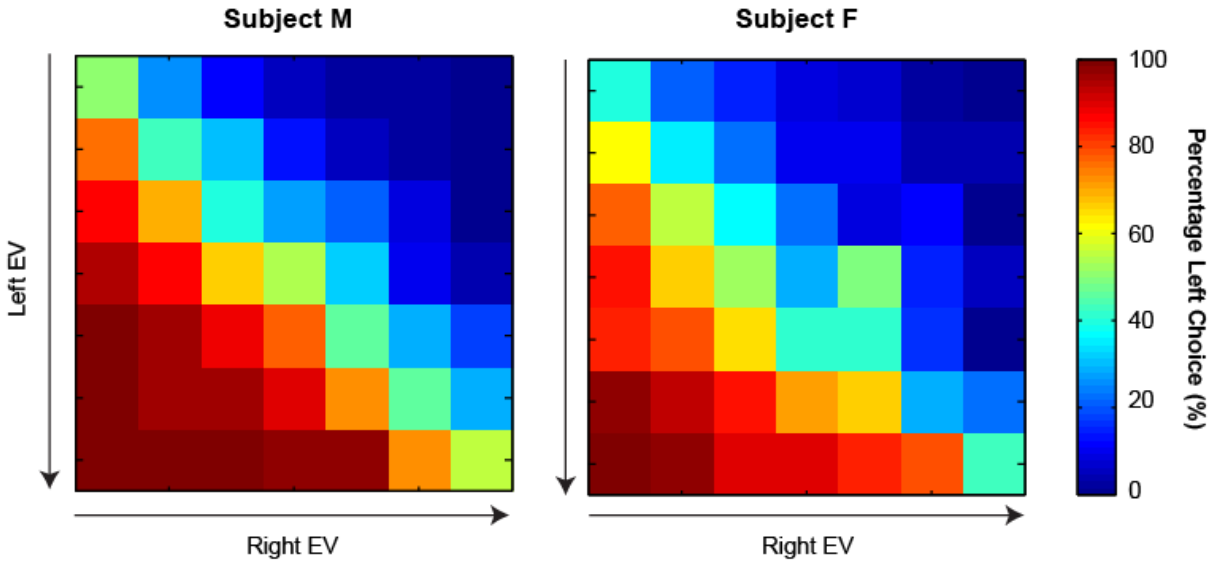

Figure 4.2: Expected value (EV) influences choice. Heat maps of left choice probability against left and right EV. Heat indicates the probability of choosing left as left (y-axis) and right EV (x-axis) increases.

However, in order to test whether subjects considered both attributes, a logistic regression was performed on left choice probability against left-right probability value difference and left-right magnitude difference (**Table 4.1**, **Figure 4.3**). This analysis indicated that both attributes significantly influenced subject choice (M: both T statistics>38.9, both

$p<2x10^{-308}$, F: both T statistics>32.00, both $p<2x10^{-224}$). Another important finding was the small but significant influence of the direction of first saccade on final choice direction implying that options that subjects viewed first had a bias on final choices (logistic regression, M: T=12.57, $p<3x10^{-36}$, F: T=15.43, $p<1x10^{-53}$).

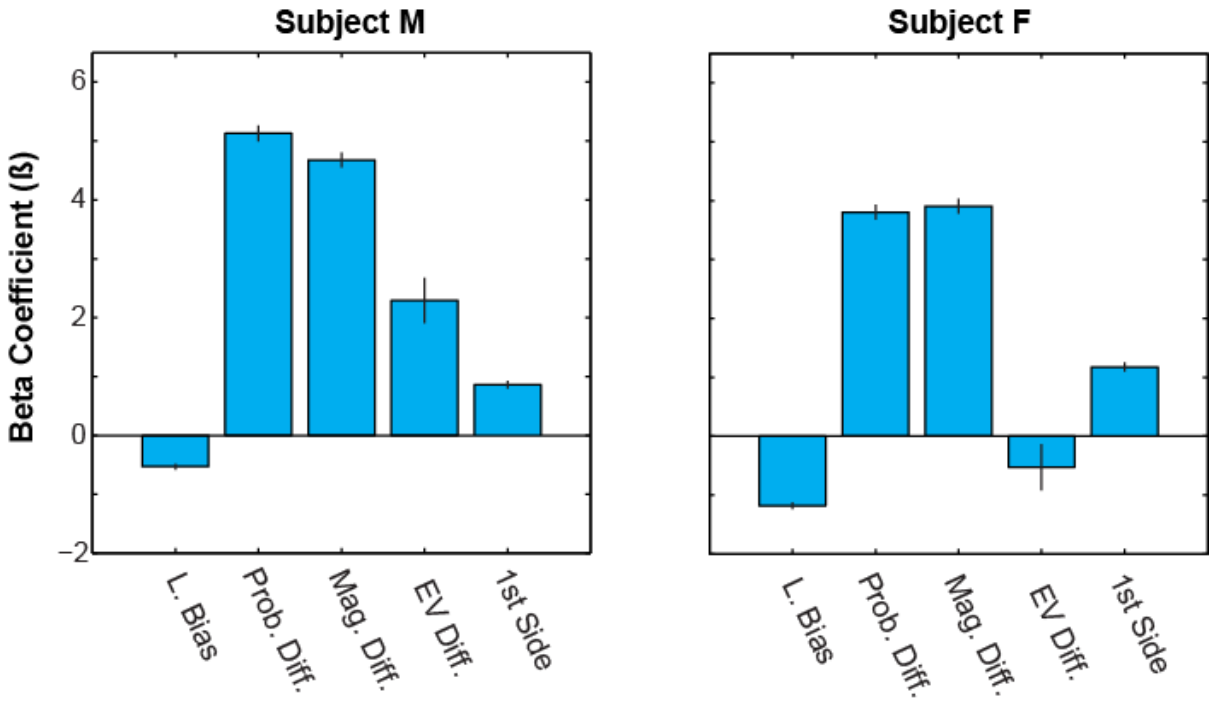

Figure 4.3: Logistic regression of left choice. The regressors used were a constant term for left choice bias (L. Bias), left-right probability value difference (Prob. Diff.) and left-right magnitude value difference (Mag. Diff.), orthogonalised left-right EV difference (E.V. Diff.) and a binary term for a first side choice bias (1^{st} Side). Error bars indicate the standard error of the mean (SEM).

Importantly, subjects weighted each attribute equally as would be expected if they calculated the utility of the option (linear hypothesis test of the beta coefficient for probability value difference compared against magnitude value difference, p>0.05). A linear regression of decision time against chosen-unchosen magnitude and probability value difference produced negative beta coefficients (i.e. as either value difference increased, choices became faster) but neither regressor had a greater influence (linear regression, M: both unsigned T statistics>13.81, both $p<2x10^{-35}$,F: both unsigned T statistics>5.79, both $p<8x10^{-9}$, linear hypothesis test of betas coefficients for chosen-unchosen probability value difference compared against chosen-unchosen magnitude value difference, p>0.05).

Subjects' Initial Saccades Are Value Driven

By allowing the subjects to freely saccade around the screen during the choice phase, it was possible to make some inferences about the processes that occur in the brain during decision making. The eye position data allowed us to determine the direction of saccades as well as identify the spatial position of each fixation; the latter was used to classify whether a particular picture had been overtly "attended".

When the cues were initially presented on the screen, we observed that initial saccades were of relatively short latency (median M: 146±20ms (1 S.D.), F: 153±30ms) and almost always (99%) towards cues on the bottom half of the screen. However, initial saccades were approximately evenly distributed between left and right sides of the screen (M: 50:50%, F: 47:53%). Were these left vs right saccades just random or was there a pattern to them? To test this, we performed a logistic regression of the probability of a left initial saccade against cue value at each position on the screen **(Figure 4.4A)**. This analysis found that although saccade direction was influenced by value cues presented at *all* positions on the screen (M: all unsigned T statistics>9.15, all $p<6x10^{-20}$, F: all unsigned T statistics>4.85, all $p<2x10^{-6}$), the influence was stronger for cues presented on the bottom of the screen (linear hypothesis test of beta coefficients for left saccade probability for value cues presented on the top versus bottom half of the screen, M: T=282.70, $p<3x10^{-62}$, F: T=163.54, $p<7x10^{-37}$). We repeated this analysis separated by attribute, but the type of attribute had no differential influence on saccade direction (linear hypothesis test of beta coefficients for left saccade probability for magnitude value cues presented anywhere on the screen against those for probability value, M: T=0.066, p=0.83, F: T=0.070, p=0.79). The influence of cues presented on the bottom of the screen was surprisingly also present at the level of choice, although the effect was weaker (**Figure 4.4B**) (linear hypothesis test of beta coefficients for left choice probability for magnitude and probability cues presented on the top half of the screen against those presented on the bottom half, M: T=43.14 $p<6x10^{-11}$, F: T=416.18, $p<3x10^{-308}$).

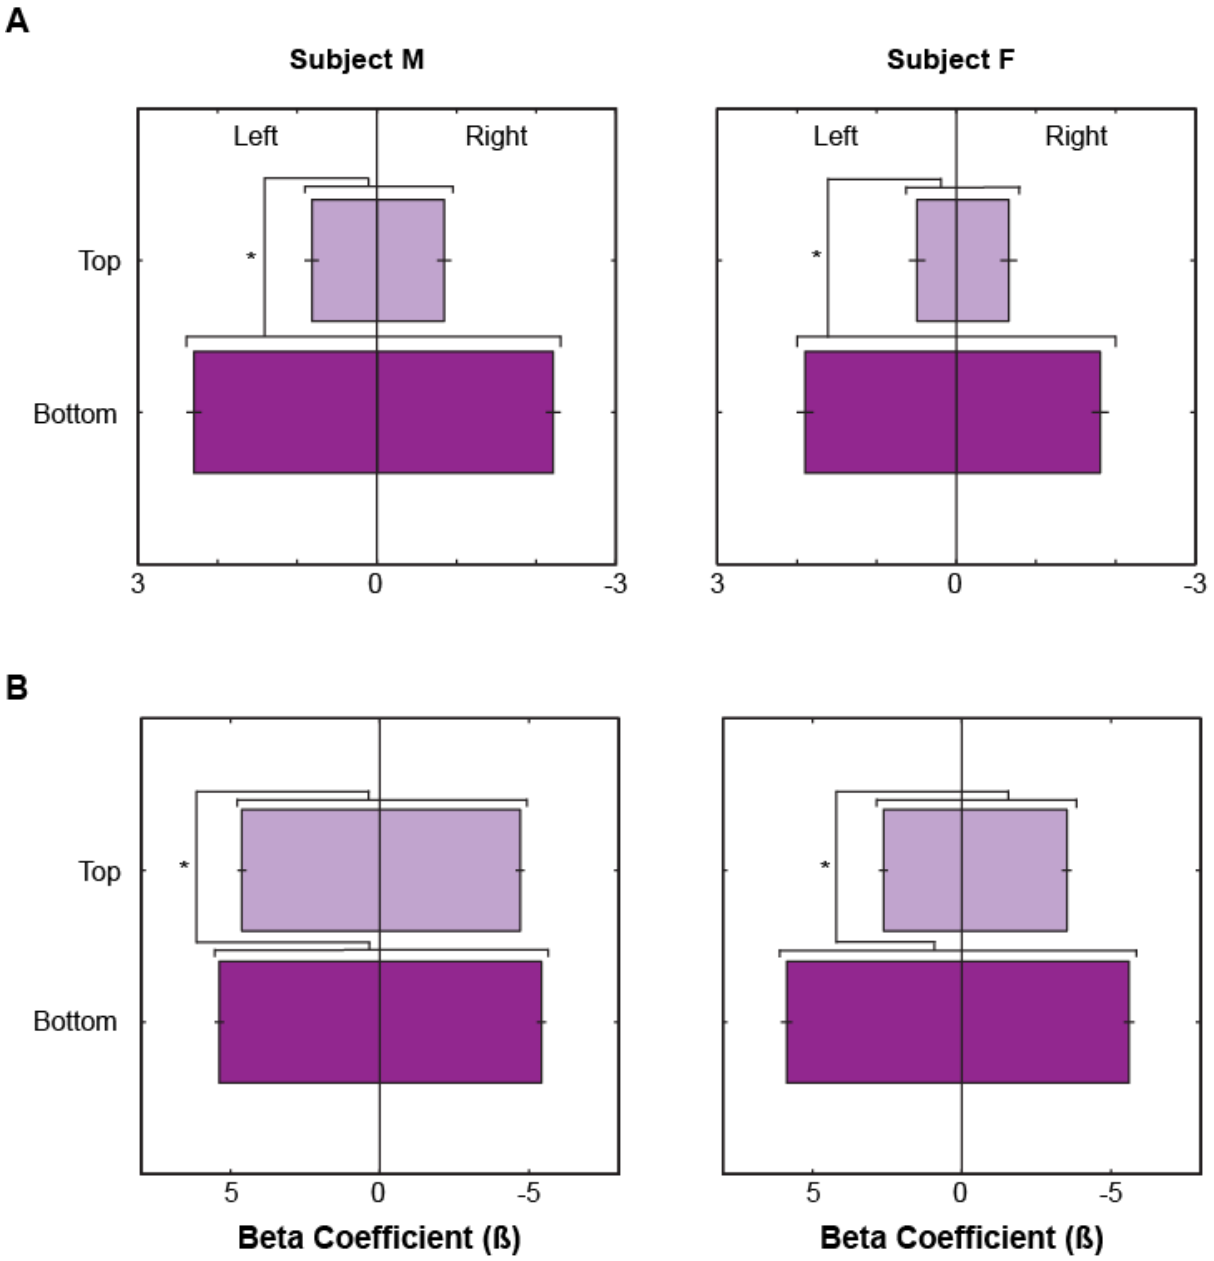

Figure 4.4: The influence of value and spatial position on saccades and choices. **(A)** A logistic regression of left saccade probability against cue values separated by spatial location. Asterisks indicate beta coefficients that are significantly different between value regressors on the top (light purple) and bottom (purple) of the screen (linear hypothesis test, $p<0.05$). **(B)** A logistic regression of left choice against cue values separated by spatial location.

We next wanted to explore whether this initial saccade was value driven. We identified the position of the first saccade and calculated the probability that this saccade was towards the more valuable picture when compared with the corresponding cue of the

other option (i.e. the value of the cue directly horizontal to it). **Figure 4.5** shows that the probability that the initial saccade would be directed towards the more valuable stimulus of the (typically bottom) horizontal pair ("saccade optimality") increased as a function of the best available value in that pair. Importantly, saccade optimality was also above chance for almost all given values (binomial test, M: all $p<0.03$, F: all $p<6x10^{-8}$ except value 2 where $p=0.33$), indicating that subjects weren't simply tracking only the most valuable stimuli, but instead tracking the value of all of the stimuli. The patterns of value-based initial saccades imply a covert value comparison process which directs overt attention toward more valuable information as a function of the value of that information (c.f., Chapter 3).

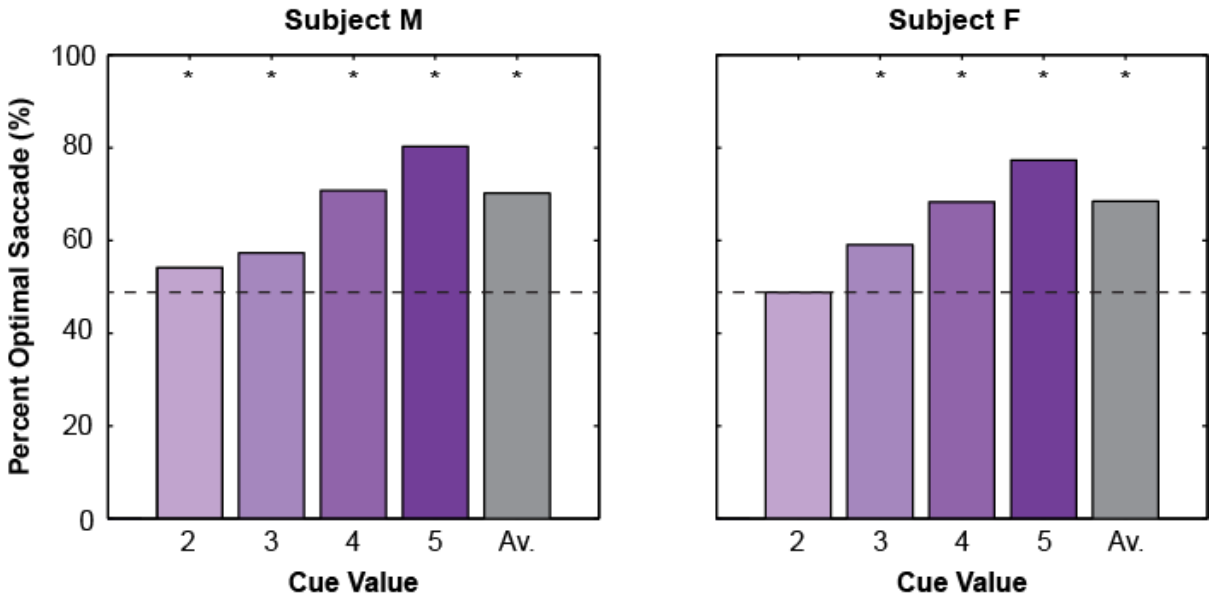

Figure 4.5: Probability of optimal saccades for various best available cues. Each bar indicates the probability that the subject's initial saccade was towards the most valuable cue of the horizontal cue pair (i.e., top or bottom pair) which were attended (i.e., if the subject's first saccade was towards the bottom left cue, was the value of this cue higher than then bottom right cue). The grey bar indicates the average optimality of saccades. Asterisks indicate saccade probabilities that are significantly above chance (binomial test, $p<0.05$).

Subsequent Saccades and Choice Utilise Covert Attention

Surprisingly, subjects rarely overtly saccaded to all available information on a given trial (M: 2%, F: <1%) and were mostly likely to fixate only two cues (M: 67%, F: 63%)

(**Figure 4.6**) before making a choice via joystick movement. This begs the question of whether the value of the cues presented influenced the number of cues subjects would view?.

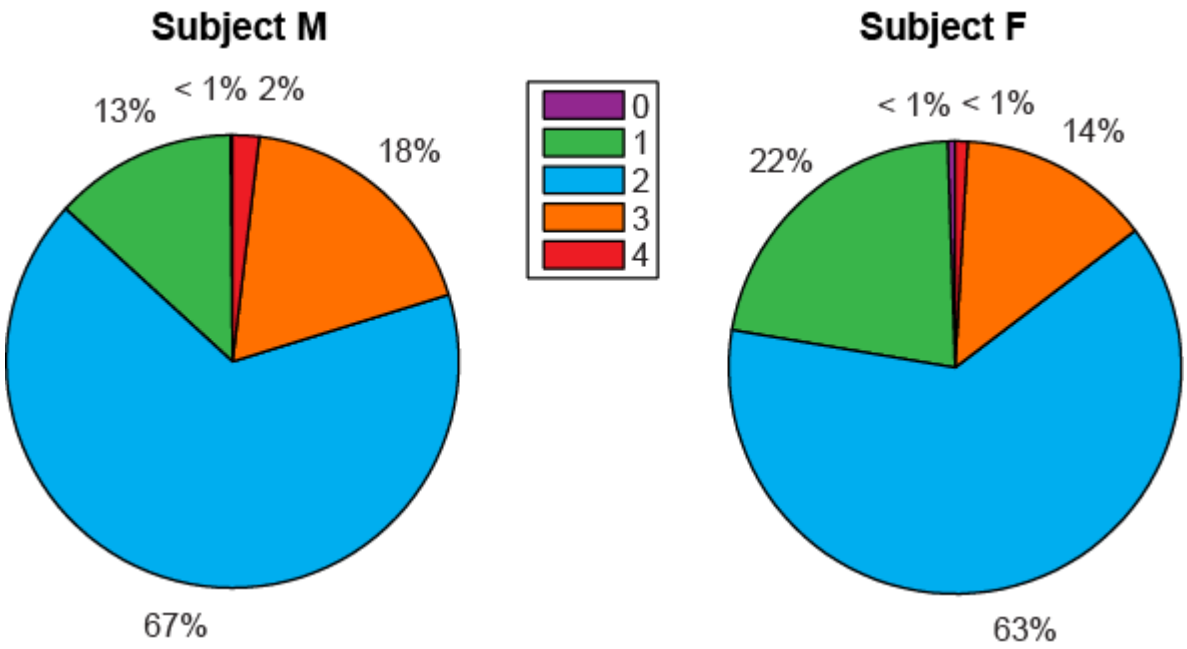

Figure 4.6: The number of cues viewed in 'Simultaneous' trials. The number of cues viewed was calculated by a saccade direction algorithm (see Methods).

Based on the finding that subjects tended to make initial saccades towards the better of the two cues on the bottom half of the screen, we hypothesised that the attended side of the screen may have a strong influence on future information gathering behaviour. We therefore examined the mean number of saccades observed for each possible EV of the first attended option (called Option 1) and the unattended option (called Option 2) (**Figure 4.7A**). We also performed a linear regression of the number of cues viewed against cue values at each screen position relative to the position of the first fixated cue, which included orthogonalised EV terms for Option 1 and Option 2 (**Figure 4.7B**).

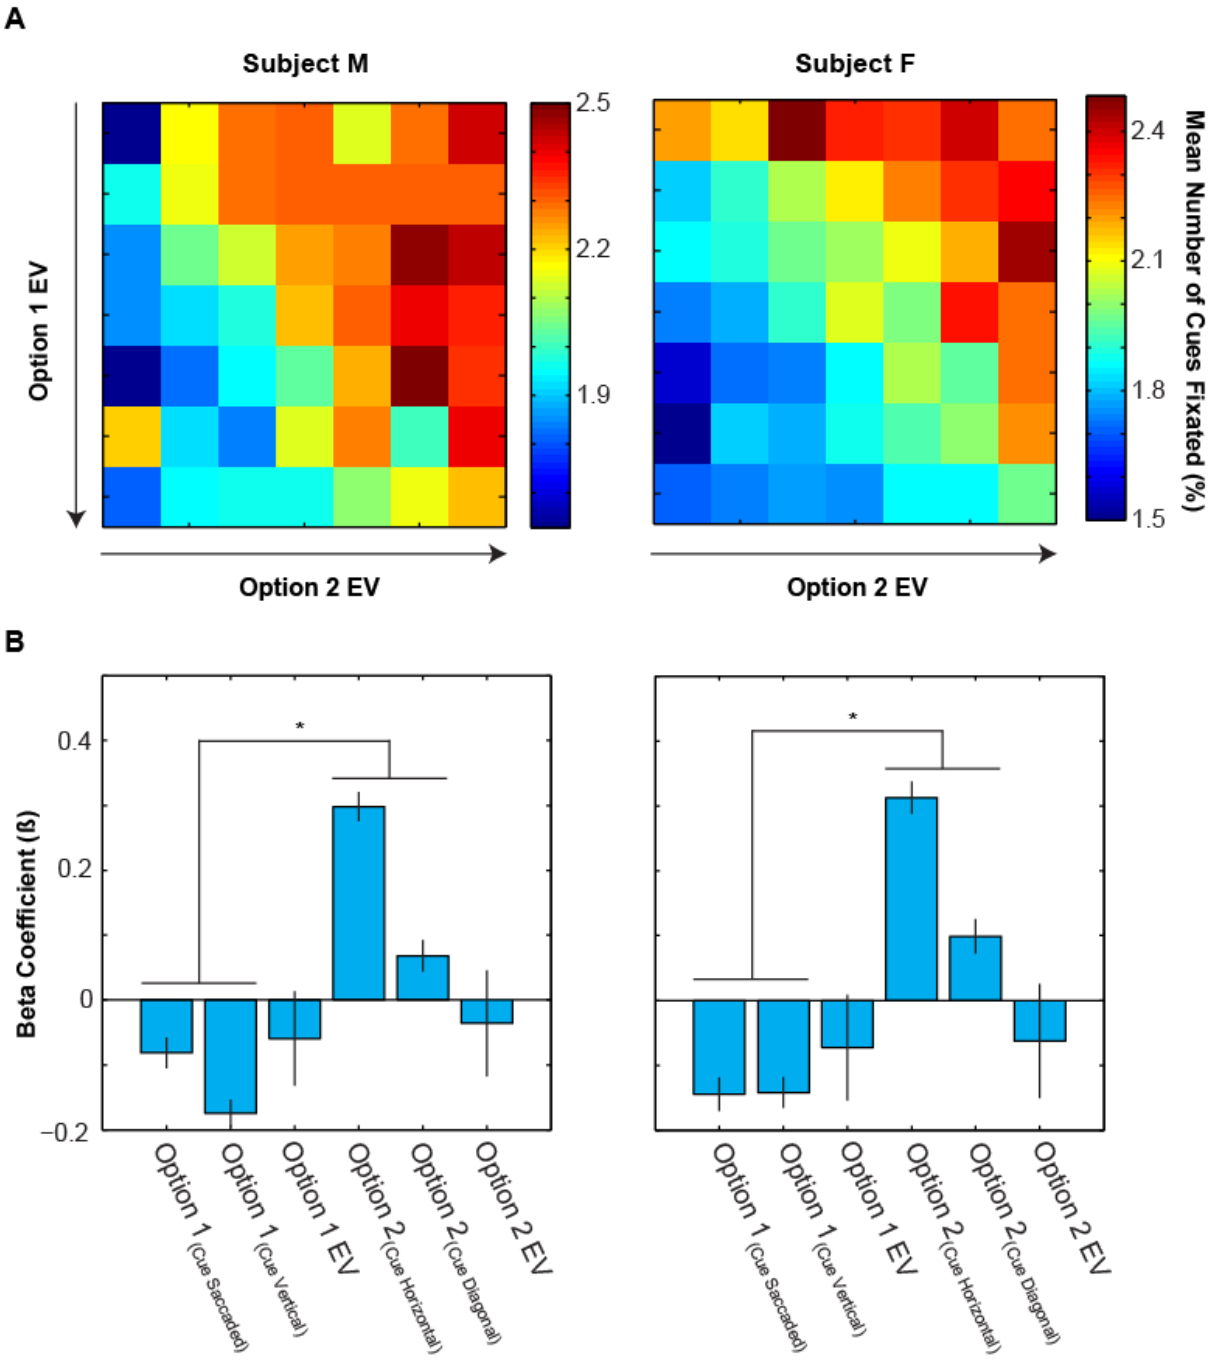

Figure 4.7: Subject saccade behaviour is driven by all cues. **(A)** A heat map of mean number of cues viewed as a function of left EV (y-axis) and right EV (x-axis). **(B)** Beta coefficients for the linear regression of mean number of cues viewed against cue values separated by relative positions with respect to the first viewed cue. Option $1_{(\text{Cue Saccaded})}$ was the initially fixated cue, Option $1_{(\text{Cue Vertical})}$ was the cue directly vertical to the initially fixated cue, Option $2_{(\text{Cue Horizontal})}$ was the cue horizontal to the initial cue and Option $2_{(\text{Cue Diagonal})}$ was the cue diagonal to the initial cue. All of figure properties are the same as **Figure 4.4.**

Both analyses found that there was a strong tendency to saccade to fewer cues as the value of first fixated option (i.e., “attended option”) increased, and conversely make more saccades when the value of cues of the unattended option increased (linear regression, M: all unsigned T-statistics>3.47, all $p<6x10^{-4}$, F: all unsigned T-statistics>3.73, all $p<2x10^{-4}$). However, on average the cues that made up Option 2 exerted a greater influence on the number of pictures viewed than those of Option 1 (linear hypothesis test of the beta coefficients for Option 1 cues against Option 2 cues, M: T=4.30, p<0.04, F: T=4.69, p<0.04). Neither of the orthogonalised EV regressors were significant (p>0.05) although this is unsurprising considering the strong correlation between integrated option value and the value of the cues. These results imply that the better the covert value comparison is prior to the initial saccade, the less information the subject will need to make a final choice.

We next investigated whether subjects exhibited any specific patterns or strategies in their information gathering behaviour after making their initial saccade to the first cue. We therefore classified three types of comparison; option (where subjects make a second saccade vertically within the two cues of the same option), attribute (where subjects saccade to the horizontally to the same attribute cue of the second option) or diagonal (where they saccade diagonally across both options and attributes). This analysis found that subjects had an overwhelming tendency to perform option and attribute saccades (M: 56% option, 36% attribute, F: 45% option, 52% attribute) (**Figure 4.8**).

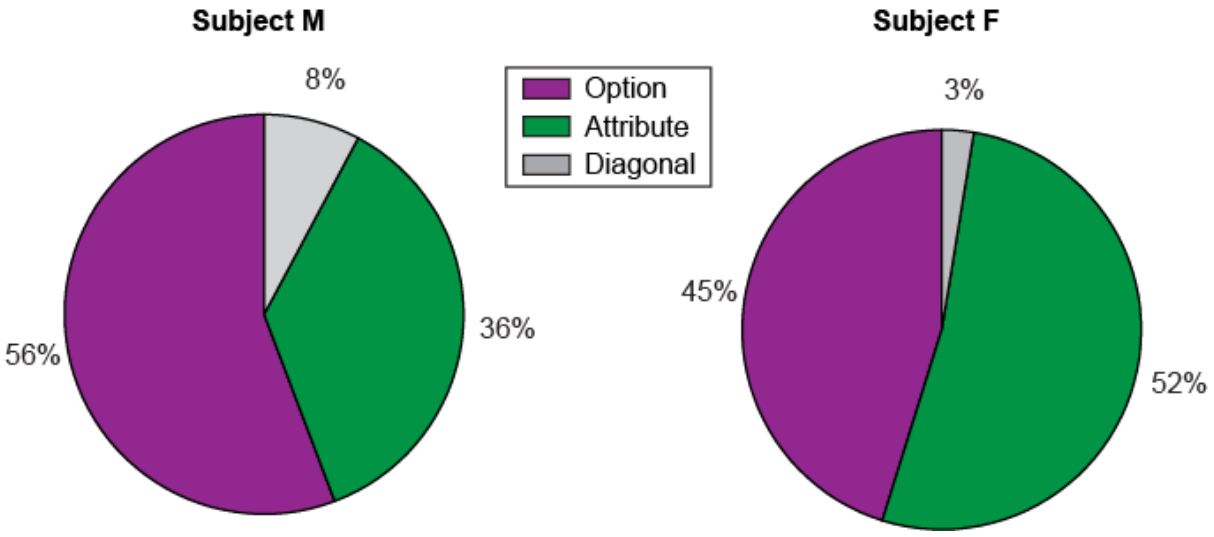

Figure 4.8: Direction of second saccades. A pie chart showing the proportion of second saccades that are made towards the other cue of the same option as the first saccade ('Option'), the same attribute cue of the other option ('Attribute') or to the other attribute cue of the other option ('Diagonal').

We then considered whether there was any pattern to these option and attribute saccades. We performed a logistic regression of option saccade probability against the value of cue at the three possible relative locations with respect to the first saccaded cue position, including orthogonalised EVs for each option (**Figure 4.9**). This analysis indicated that the value of all of the cues influenced whether subjects made 'option' or 'attribute' type saccades. (M: All unsigned T-statistics>4.58, all $p<5x10^{-6}$, F: All unsigned T-statistics>8.80, all $p<6x10^{-16}$). Again, the value of cues on the Option 2 side had a greater influences on future saccade direction than cues on the Option 1 side (linear hypothesis test of the beta coefficients for Option 1 cues against Option 2 cues, M T=65.81, $p<6x10^{-16}$, F: T=43.23, $p<6x10^{-11}$). Subjects tend to make saccades towards the more valuable side. This suggests that rather than just randomly saccading around the screen, subjects use covert attention in order to plan the direction of future saccades. Furthermore, there is evidence that subjects have covertly attended to all of the information on the screen before deployment of the second saccade.

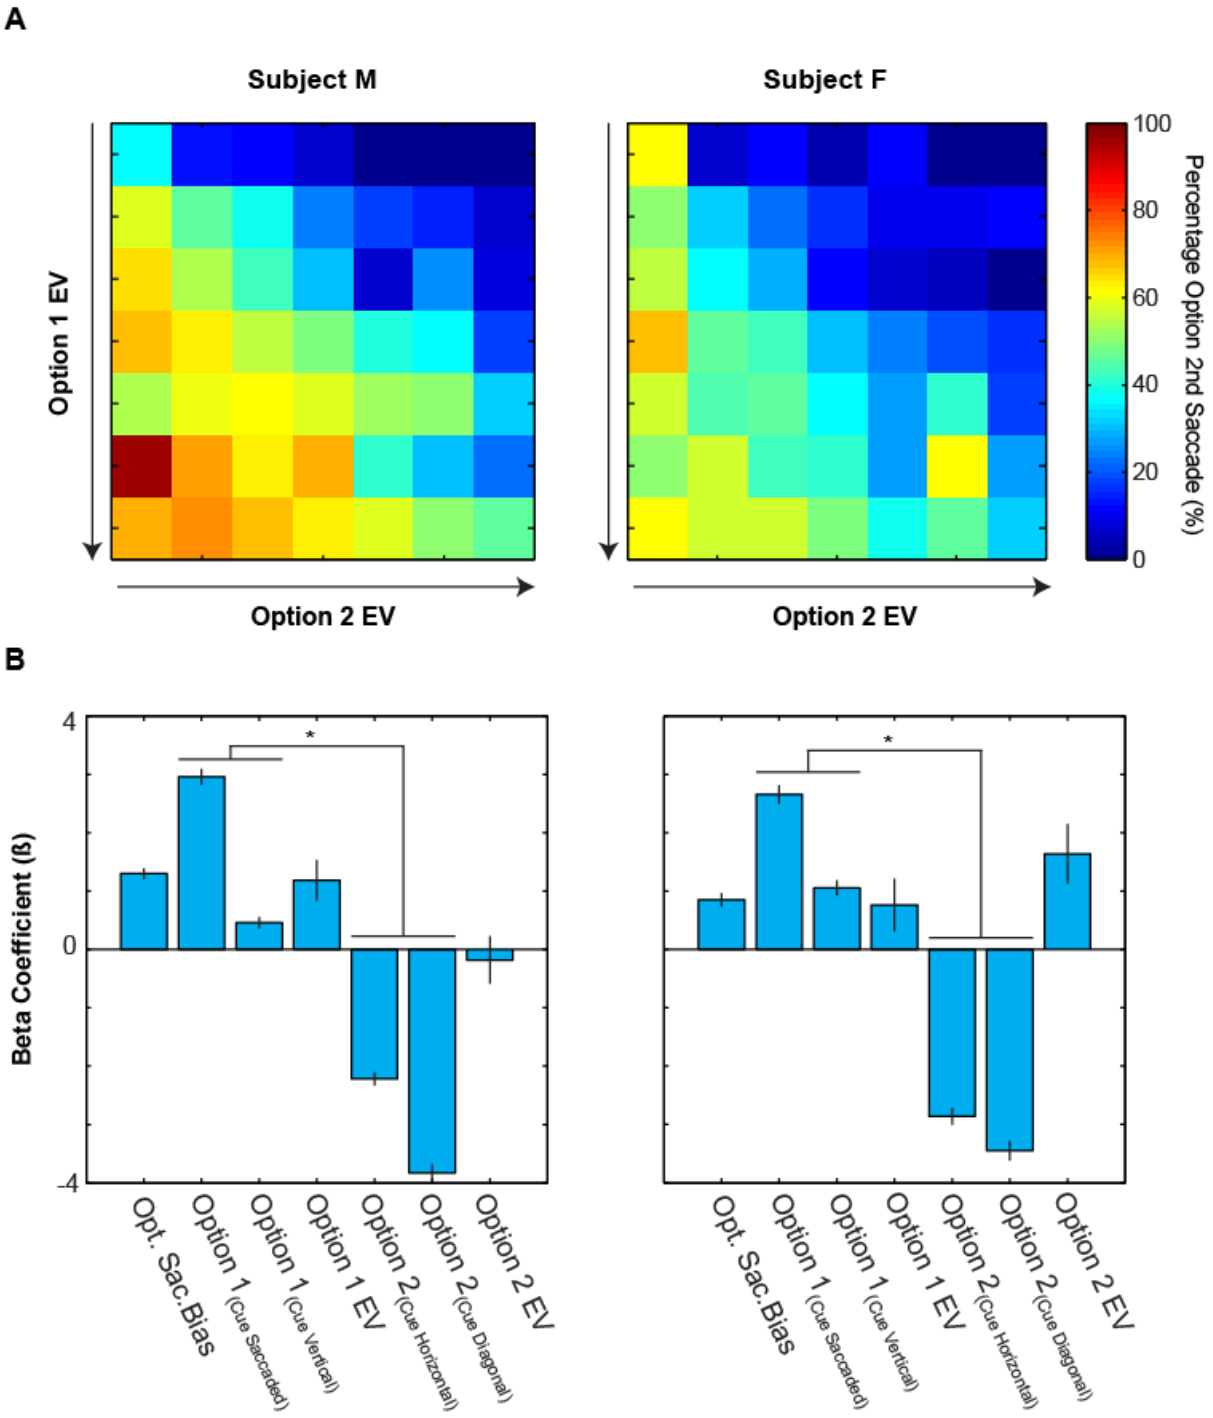

Figure 4.9: Subject saccade direction is influenced by cue value. **(A)** A heat map of probability of option saccade as a function of left EV (y-axis) and right EV (x-axis). **(B)** Beta coefficients for the linear regression of mean number of cues viewed against cue values separated by relative positions with respect to the first viewed cue. All of figure properties are the same as **Figure 4.7**.

These analyses have shown that subjects rarely overtly attend to all available information before making a choice, and that saccades are heavily driven by covert

attention. We therefore investigated whether the final choice was driven predominantly by overt attention (i.e. that subjects only consider the value of overtly attended cues when choosing) or by covert attention (i.e. that subjects consider the value of all cues when choosing even if they are not overtly fixated). To do this, we performed model comparison of two logistic regression models (dependent variable: probability of Left choice), called the 'covert' and 'overt' model. Each model contained only one regressor; left versus right EV difference. In the 'covert' model, the value difference regressor was constructed using the probability and magnitude values presented on each trial on the left and right. For the 'overt' model the same method was used to construct the value difference regressor with the exception that on each trial, if the subject did not saccade to one or more of the cues, the values of these cues were set to average value (i.e. 3 out of 5). The rationale behind this was that if subjects were only using overtly gathered information to make choices then they must make inferences about the value of unseen cues. If this was true then the intuitive inference they could make would be that the cues were of average value. We then obtained the residuals for each model, and compared how well each model predicted left choice. The median residuals for the 'covert' model was seen to be significantly smaller than that of the 'overt' model (Kruskal-Wallis test, M: $Chi^2=2.79x10^3$, $p<3x10^{-308}$, F: $Chi^2=1.26x10^3$, $p<3x10^{-276}$), indicating that behaviour was better explained by assuming that subjects were using covert attention to consider all available information rather than assuming that they were only using overt attention and inferring the value of unseen cues.

In conclusion, behaviour in the 'Simultaneous' task indicates that subjects integrate both probability and magnitude information when executing multi-attribute binary decisions. However, primate saccadic systems rapidly utilise covert attention to modulate both initial and later saccades based on the value of all cues presented in an attribute non-specific manner. Finally, subjects also use covert attention to guide final choice despite this not being a constraint of the task.

Part 2: Information Gathering Trials

Subjects Sensitivity to Value Varies With Trial Type

Because of the clear influence of covert attention during 'Simultaneous' trials, it is difficult to make conclusions about how subjects may be comparing information during the decision making process. In contrast, the sequential yet relatively unconstrained nature of the 'Information Gathering' task provides the ideal opportunity to assess the relationship between value, information comparison and choice.

In order to examine the extent to which choices are driven by value in 'Option' and 'Attribute' trials of the 'Information Gathering' Task, we performed a logistic regression of left choice probability against left-right probability and left-right magnitude value difference for each trial type and the side of the first cue presentation (left/right) (**Figure 4.10**, **Table 4.3**). Any cue that was not considered to have been seen was given average value (i.e., 3) before value difference was computed. This led to the surprising finding that unlike in 'Simultaneous' trials, the beta coefficients for probability value difference were significantly different compared to those for magnitude value difference in both 'Option' and 'Attribute' trials (linear hypothesis test of the beta coefficients for probability value difference against magnitude value difference, M: absolute $T=49.80$, $p<2x10^{-12}$, F: absolute $T=4.62$, $p<0.04$). Interestingly, this difference was of opposite signs for each subject (i.e. Subject M was more sensitive to probability while subject F was more sensitive to magnitude value). The size of the coefficients for both probability and magnitude value difference were significantly smaller for 'Attribute' trials compared to 'Option' trials indicating that subjects were slightly less sensitive to value in 'Attribute' trials (linear hypothesis test of the beta coefficients for probability and magnitude value difference in 'Option' trials against those in 'Attribute' trials, M: $T_{(Probability')}=8.38$, $T_{(Magnitude)}=9.79$, $p_{(Probability)}<0.004$), $p_{(Magnitude)}<0.002$, F: $T_{(Probability')}=67.30$, $T_{(Magnitude)}=43.36$, $p_{(Probability)}<3x10^{-308}$), $p_{(Magnitude)}<3x10^{-308}$). These results suggest that presenting information to subjects sequentially exposes attribute biases in subject choice,

and forcing subjects to perform initial option comparisons makes them more sensitive to value than when they perform initial attribute comparisons.

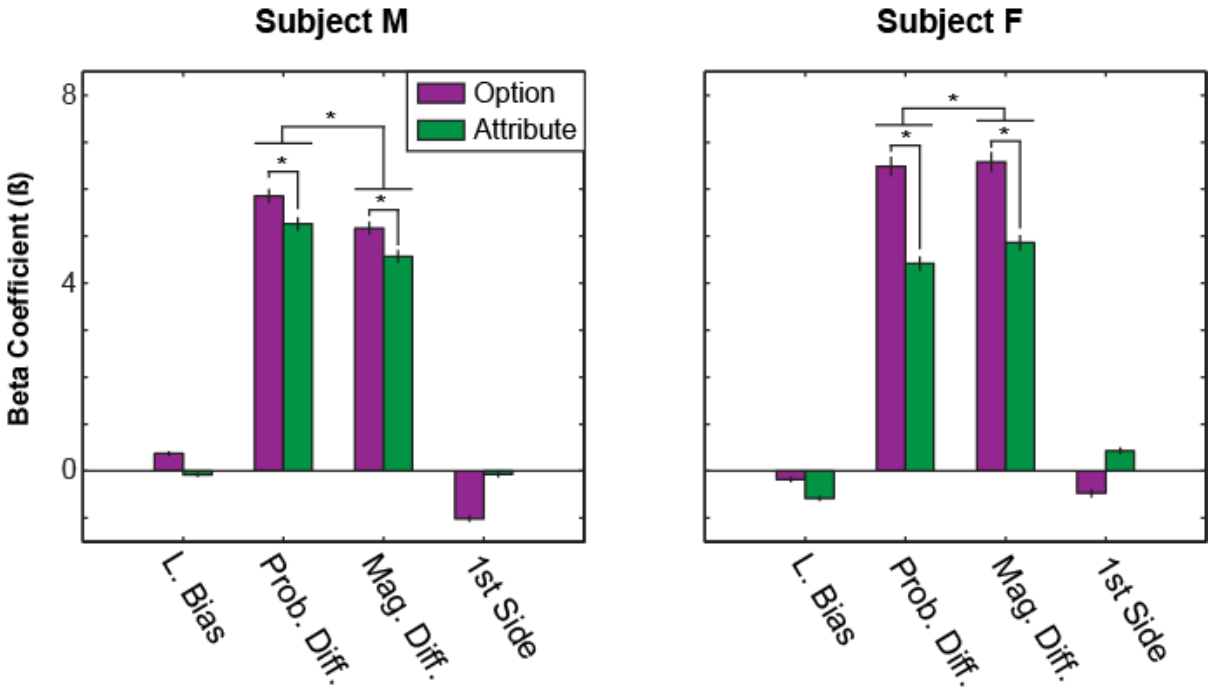

Figure 4.10: Trial type affects the influence of value on choice. Logistic regression of left choice against left-right probability and magnitude value difference. 1st Side is defined by left-right.

Subjects also had a tendency to choose away from the direction of the first saccade in 'Option' trials (**Figure 4.10**: logistic regression, M: T=-14.15, $p<2x10^{-45}$, F: T=-5.24, $p<2x10^{-7}$). Subject F also had the opposite bias in 'Attribute' trials, where he tended to choose the side of the first presented cue (logistic regression, T=5.70, $p<2x10^{-8}$).

Subjects Gather Incomplete Information During Information Gathering Trials

In both 'Option' and 'Attribute' trials, subjects rarely attended to all four cues before making a decision (M: 8%, F: 5%) (**Figure 4.11**). However, the number of cues that subjects viewed was similar between the two trial types.

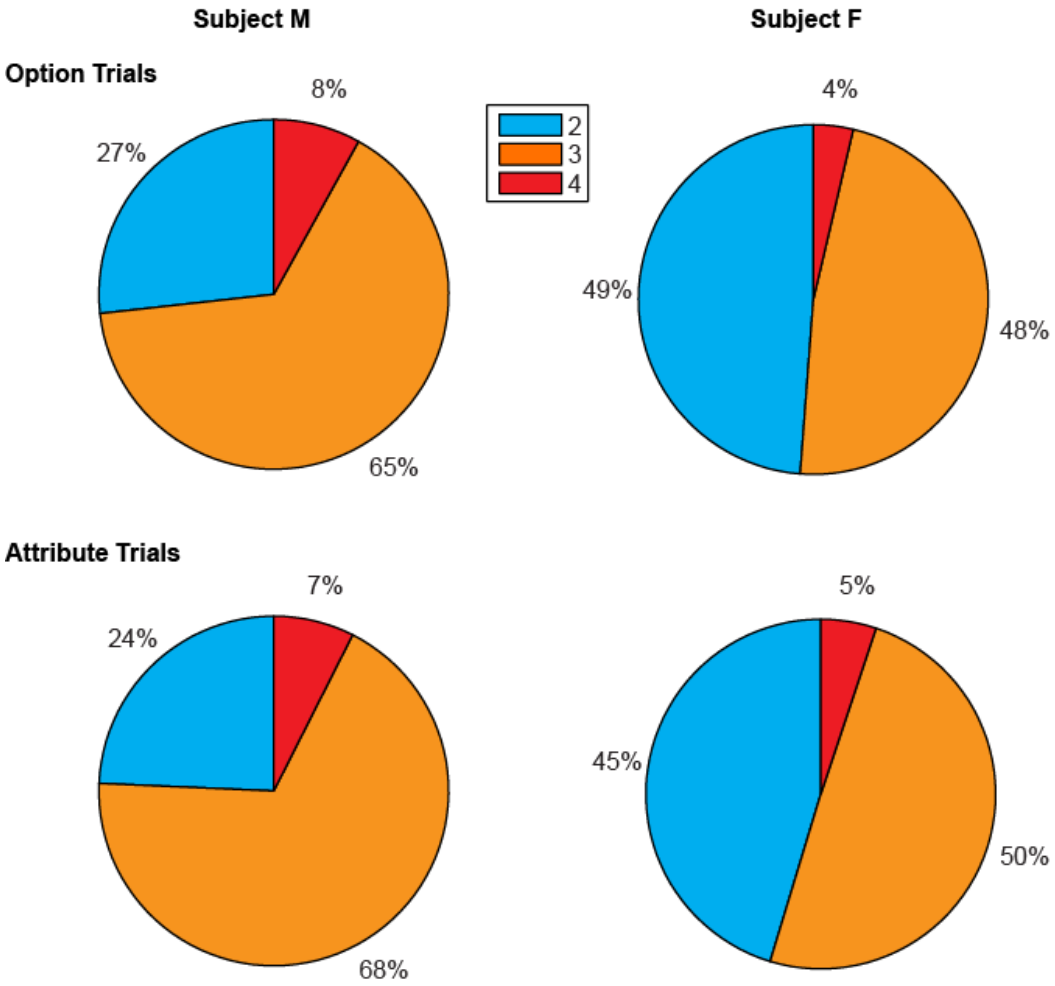

Figure 4.11: Subjects rarely gather all available information. The Number of Cues Viewed in 'Option' and 'Attribute' Trials.

To explore how cue values influenced information gathering behaviour, we examined the mean number of cues viewed for each combination of Cue 1 and Cue 2 values separated by trial type (**Figure 4.12A** and **Figure 4.13A**). We then performed a linear regression of the number of pictures viewed against the cue 1 and 2 values (**Figure 4.12B** and **Figure 4.13B**). Both analyses showed that in 'Option' trials (**Figure 4.12B**), as both cue values increased, subjects were likely to view fewer cues in total (M: both unsigned T-statistics>22.41, $p<2x10^{-107}$, F: both unsigned T-statistics>15.41, $p<3x10^{-52}$). Cue 1 also had a greater influence on the number of cues viewed than Cue 2 (linear hypothesis test of beta coefficients for Cue 1 against those for Cue 2, M: T=195.42, $p<8x10^{-44}$, F: T=14.87, $p<2x10^{-4}$). Interestingly, trials when subjects were initially presented with two low value cues were

those when they would view the most cues. This was a form of 'confirmation bias', where rather than just choosing the other unseen (almost certainly better) option, they preferred to first view the cues of that option before choosing it.

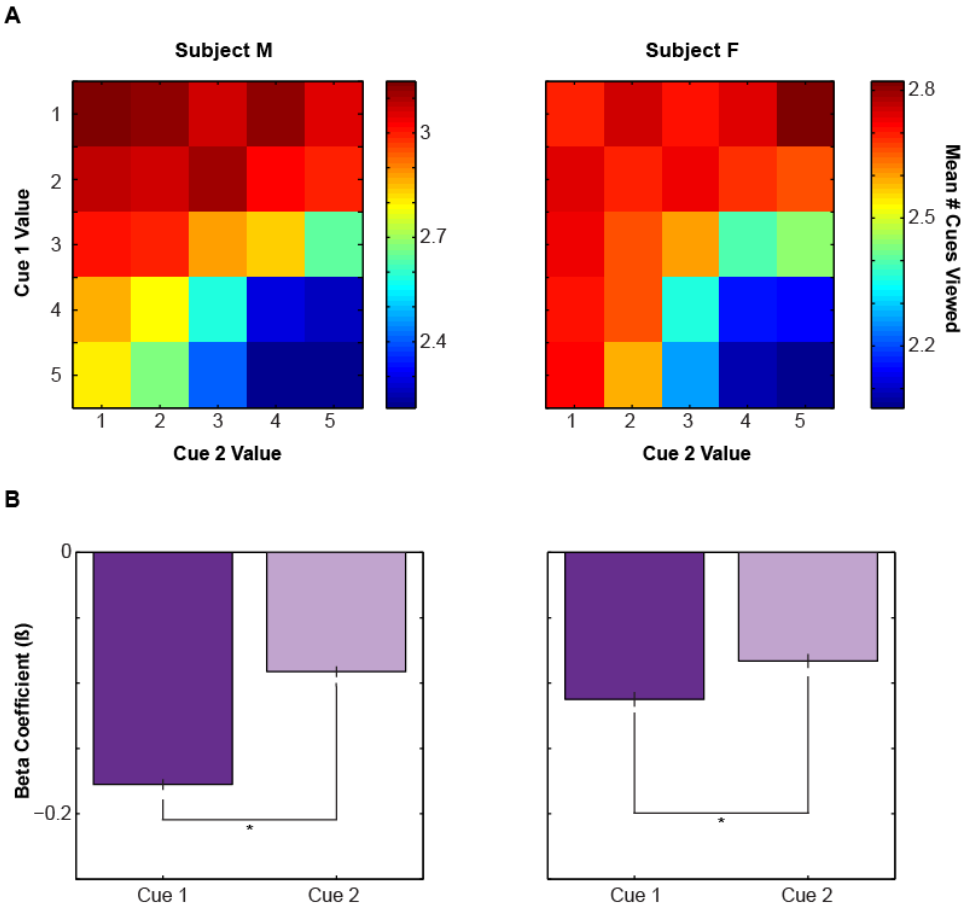

Figure 4.12: Information gathering in 'Option' trials. **(A)** A heat map of mean number of cues viewed as a function of the value of the first and second cues presented in an 'Option' trials. **(B)** A linear regression of cues viewed against the value of the first and second cues. All figures properties are the same **Figure 4.9**.

In contrast, in 'Attribute' trials, the two subjects exhibited different behavioural patterns. It was the case that as Cue 2 value increased, both subjects viewed more cues (linear regression, M: $T=12.25$, $p<4x10^{-34}$,F: $T=2.46$, $p<0.008$) (**Figure 4.13B**). However, Subject M had a positive influence of Cue 1 value on total viewed cues (linear regression, $T=18.41$, $p<6x10^{-74}$) whereas Subject F had a negative influence (linear regression, $T=-4.96$, $p<8x10^{-7}$). This suggests that the amount of information Subject M chose to gather was related to maximum expectation of reward; i.e. having seen two low value cues in 'Attribute' trials he

was less motivated to gather more information because the EV of either option would be much lower than if he had seen two high value cues initially. On the hand, Subject F seemed to implement a strategy where he looked for evidence to choose/reject Option 1; i.e. if Cue 1 was high in value he needed much less information to confirm a choice of that option than if Cue 2 was high in value. This could be interpreted as a confirmation bias towards Option 1.

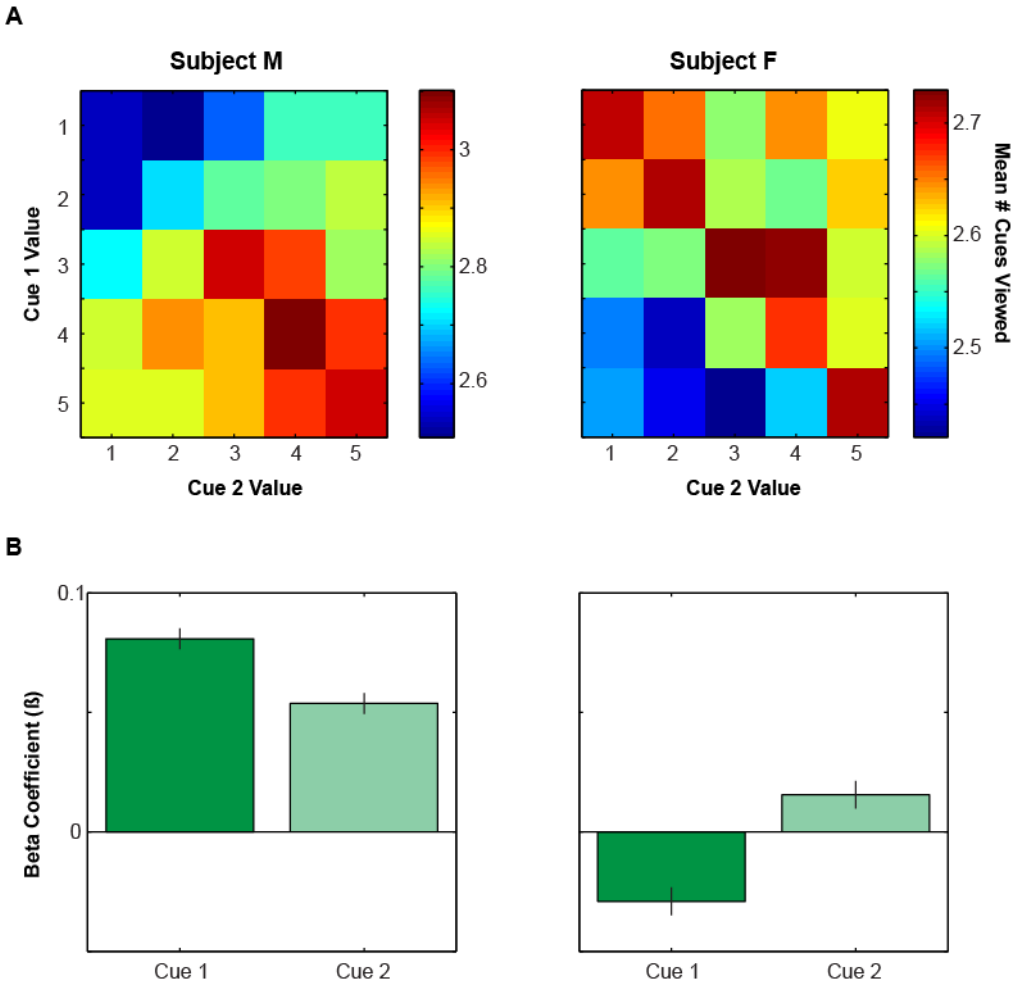

Figure 4.13: Information gathering in 'Attribute' trials. **(A)** A heat map of mean number of cues viewed as a function of the value of the first and second cues presented in 'Attribute' trials. **(B)** A linear regression of cues viewed against the value of the first and second cues. All figures properties are the same **Figure 4.12**.

Subjects Decide To Respond Based on Current Cue Comparisons

What drives subjects to terminate information gathering and initiate a response? At Cue 2 the examination of the number of cues viewed provided an indirect measure of this. However, to answer this question at a more advanced stage in the decision we examined

the probability of choosing immediately after viewing three cues based on the values of the seen cues. **Figure 4.14A** shows the probability of immediate choice in 'Option' trials as function of Option 1 EV (Cue 1 value x Cue 2 value) and Cue 3 value. For Subject M, when Option 1 increased in EV he became more likely to choose immediately, but when Cue 3 value increased, he was more likely to defer his choice until gathering the last piece of information. A logistic regression revealed both Option 1 and Cue 3 values significantly influenced the probability of immediate choice ($T_{(Option\ 1)}$=7.45, $T_{(Cue\ 3)}$=-10.94, $p_{(Option\ 1)}<9x10^{-86}$, $p_{(Cue\ 3)}<7x10^{-28}$) (**Figure 4.14B**). For Subject F, only the value of the Cue 3 influenced the immediacy of choice (logistic regression, $T_{(Option\ 1)}$= 1.19, $T_{(Cue\ 3)}$=-4.78, $p_{(Option\ 1)}>0.05$, $p_{(Cue\ 3)}<2x10^{-6}$).

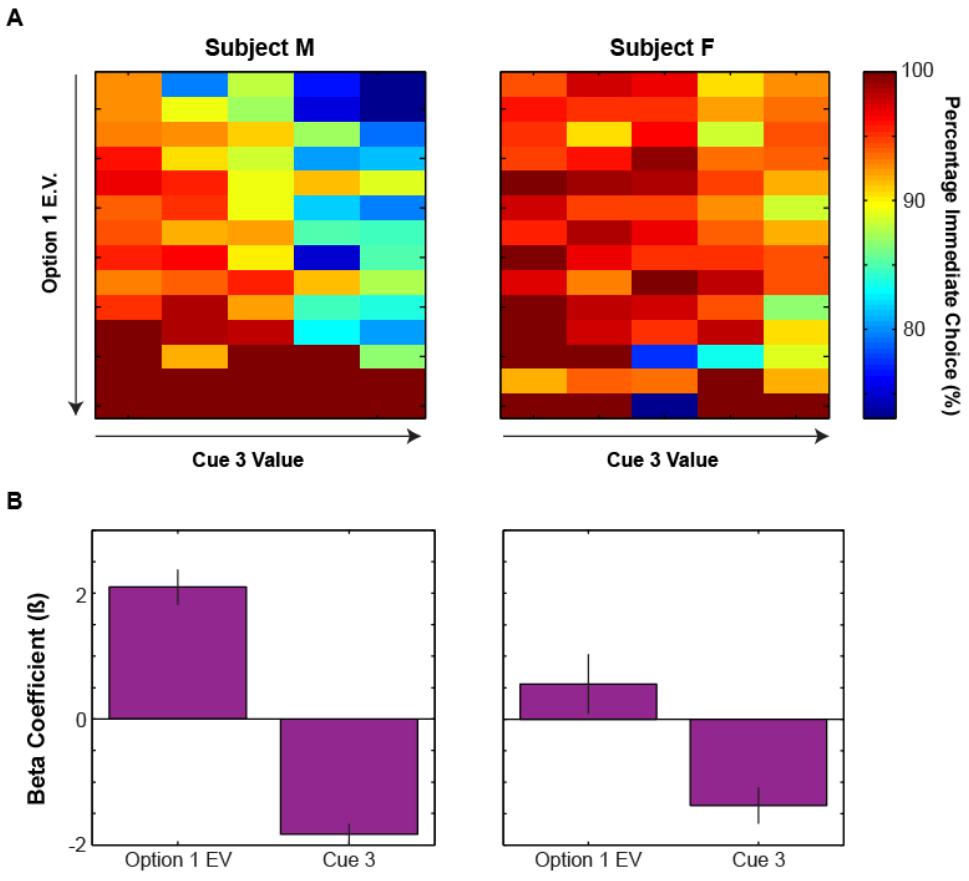

Figure 4.14: Probability of choice after Cue 3 in 'Option' trials. **(A)** A heat map of the probability of choosing immediately after Cue 3 presentation (as opposed to gathering information about a fourth cue) as a function of increasing Option 1 integrated value (Option 1 EV) and increasing Cue 3 value. **(B)** A logistic regression of the probability of immediate choice after the third cue against Option 1 EV and Cue 3 value. All figures properties are the same **Figure 4.9**.

For ‘Attribute’ trials, we examined the probability of immediate choice as a function of the absolute value difference between Cue 1 and Cue 2 (an indicator of decision difficulty) and the value of Cue 3 (**Figure 4.14A**). When the absolute difference between Cue 1 and 2 was high, Subject M was more likely to choose immediately. In contrast, only Cue 3 value had a positive influence on choice immediacy for Subject F (logistic regression, M: $T_{(Cue\ 1\ vs\ 2)}=8.74$, $T_{(Cue\ 3)}=0.07$, $p_{(Cue\ 1\ vs\ 2)}<3x10^{-18}$, $p_{(Cue\ 3)}>0.05$, F: $T_{(Cue\ 1\ vs\ 2)}=1.72$, $T_{(Cue\ 3)}=3.49$, $p_{(Cue\ 1\ vs\ 2)}>0.05$, $p_{(Cue\ 3)}<5x10^{-4}$) (**Figure 4.14B**).

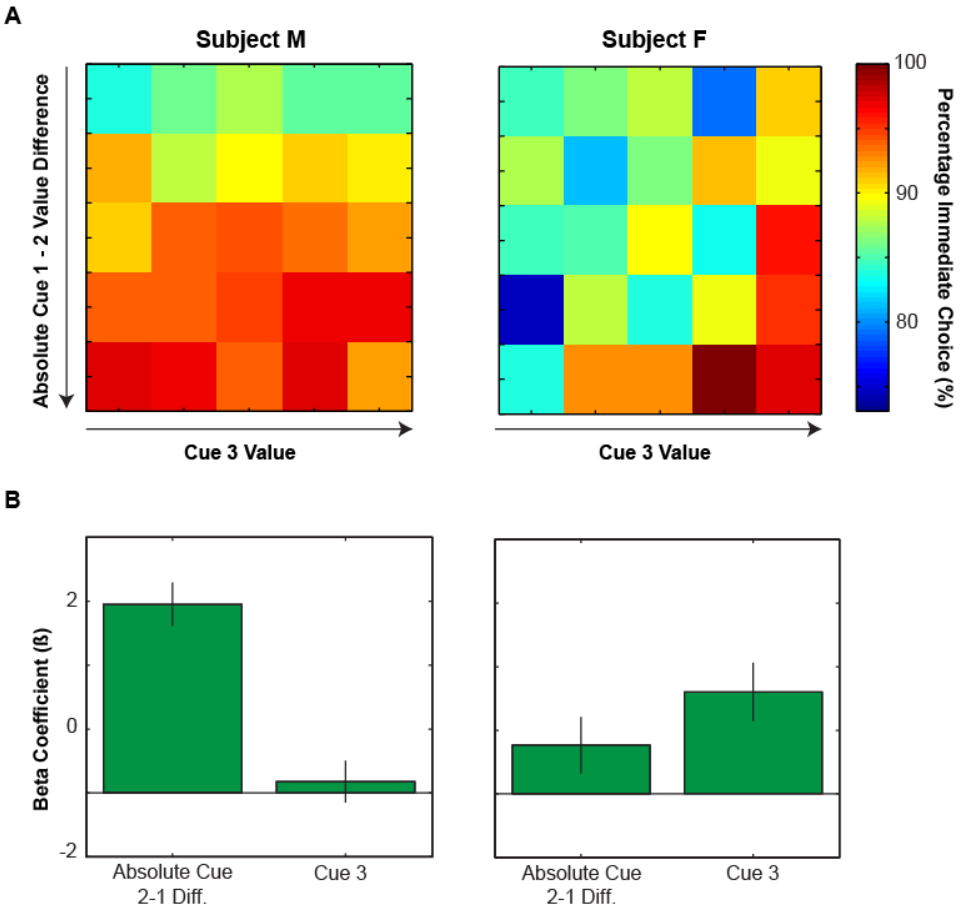

Figure 4.15: Probability of choice after Cue 3 in ‘Attribute’ trials. **(A)** A heat map of the probability of choosing immediately after Cue 3 presentation (as opposed to gathering information about a fourth cue) as a function of increasing absolute difference between the value of Cue 1 and 2 and increasing Cue 3 value. **(B)** A logistic regression of the probability of immediate choice after the third cue against absolute Cue 1- 2 difference and Cue 3 value. All figures properties are the same **Figure 4.9**.

Subjects Gather Future Information Based on the Current Value State of Options

Having established that both the amount information gathered and the timing of the response are related to the value of cues seen by subjects, we next explored whether subjects employed any strategy when gathering information. Recall that after seeing two cues, subjects were free to saccade to either one of the remaining cues to uncover the 3rd piece of information. Therefore, we were interested in whether, in 'Attribute' trials, subjects seek to gather more information about the first option viewed (i.e., the 3rd saccade being a diagonal saccade back to the first option), or more information about the currently attended option (i.e., the 3rd saccade being a vertical saccade toward the other piece of information of the second option). To answer this question, we looked at the saccade behaviour of subjects in 'Attribute' trials on occasions when they chose to fixate a third cue (**Figure 4.16**). We did not consider 'Option' trials for this analysis because Subject F exhibited no variance in this third saccade behaviour and because from a normative point of view there was no benefit to choosing the gather one location over another.

We examined the probability of making a vertical (as opposed to diagonal) 3rd saccade as a function of Cue 1 and Cue 2 values. Subject M made vertical saccades on 53.2% of occasions whereas Subject F did so on 60.1% of occasions. This revealed that as Cue 1 value increased, subjects were significantly more likely to saccade diagonally back towards Option 1; whereas when Cue 2 value increased, subjects were more likely to make vertical saccades (logistic regression, M: $T_{(Cue\ 1)}$=-17.00, $T_{(Cue\ 2)}$=23.00, $p_{(Cue\ 1)}<9x10^{-65}$, $p_{(Cue\ 2)}=5<10^{-117}$, F: $T_{(Cue\ 1)}$=-3.80, $T_{(Cue\ 2)}$=16.44, $p_{(Cue\ 1)}<2x10^{-4}$, $p_{(Cue\ 2)}< 1x10^{-60}$). Furthermore, it was observed that Cue 2 value had a stronger influence on saccades than Cue 1 value (linear hypothesis test of the beta coefficients for Cue 1 value against those for Cue 2 value, M: T=17.64, $p<3x10^{-}5$, F: T=75.24, $p<6x10^{-18}$). These results indicate that in 'Attribute' trials, rather than randomly gathering information, subjects gathered information about options that currently have the highest value. In other words, subjects gather information about options that they currently believe they will choose, indicative of a confirmation bias.

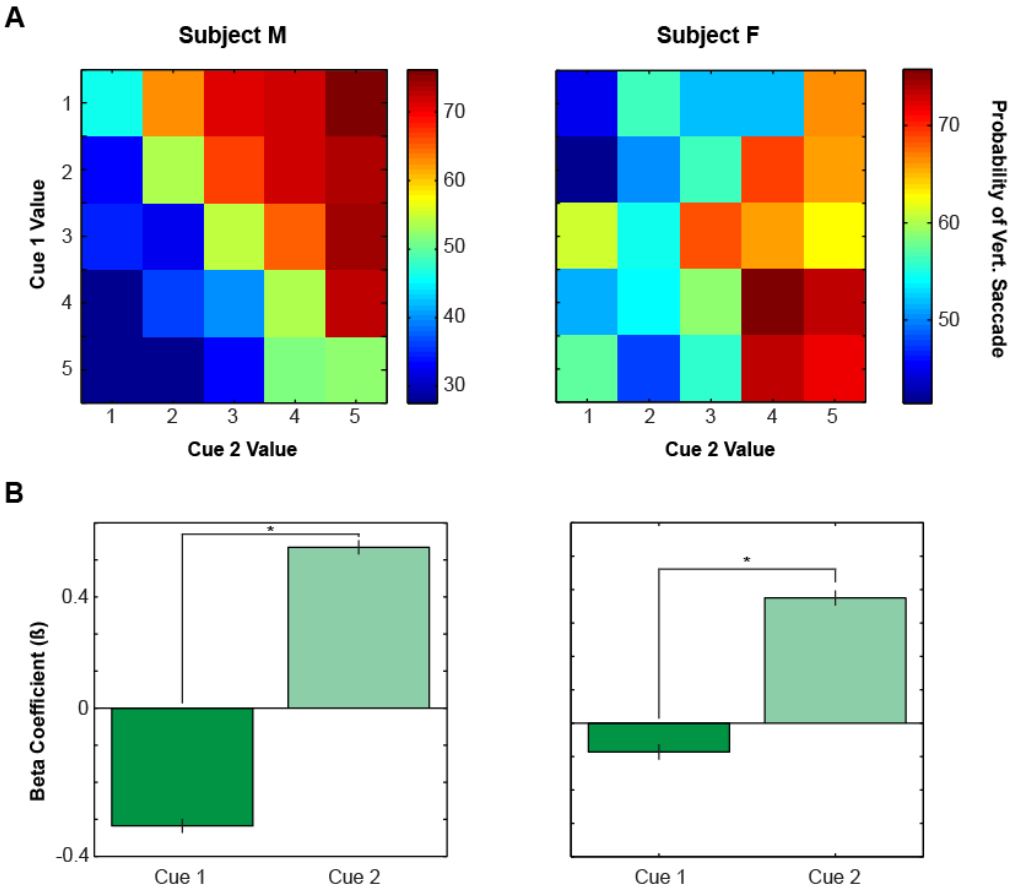

Figure 4.16: Information gathering Strategy is value driven. **(A)** A heat map of the probability of making a vertical saccade in 'Attribute' trials as a function of the value of the first two cues. **(B)** The beta coefficients of a logistic regression of cue value on saccade direction. All of figure properties are the same as **Figure 4.13**.

In conclusion, the behavioural results obtained during the 'Information Gathering' task indicate that subjects make value based decisions using incomplete information and change their information gathering strategies based on the state of current value information in a trial. Their choices are also biased by both attribute and first saccade location. Furthermore, manipulating the way in which information is initially present to subjects has a profound influence on sensitivity to value.

Discussion

In the first part of this chapter we presented evidence for the covert value comparison during multi-attribute decision making. Unsurprisingly primate subjects are known to be able to perform multi-attribute decisions using both covert and overt attention (e.g. Hosokawa et al. (2013), Padoa-Schioppa and Assad (2006), Strait et al. (2014)), though rarely are primate subjects allowed to freely view information during choice. When they were permitted to freely view information their saccades did not seems to reflect information comparison but instead provided insight into the extent to which they were using covert attention. One interesting and surprising result in our dataset is that when all information is presented to subjects simultaneously, they show no biases towards particular attributes but when the information is presented sequentially (i.e. 'information gathering' trials), both subjects show a slight tendency to weight one attribute over the other. This effect may reflect a biasing or heuristic behaviour that takes over when task complexity (through adding working memory components) increases similar to the ideas put forward by Payne (1976). Subjects also appears to show a change in sensitivity to value when information is presented sequentially which is borne out by the lower optimality exhibited in 'Option' and 'Attribute' trials compared to 'Simultaneous' trials (see **Table 4.5**). This is no doubt partially down to the use of heuristics in solving decisions but could also come about from the fact that the 'Information Gathering' trials required subjects to retain previously viewed information in working memory for several seconds. The implication of these simple observations is that the by manipulating the information presentation/gathering process we have caused changes in the way that the decision process unfolds.

Covert Attention in Multi-Attribute Decision Making

We have shown that subjects utilize covert attention to drive overt fixation at multiple points during the choice process, even when the task parameters get more complex (e.g.,

compared to the task in Chapter 3). When subjects are provided which an unconstrained environment in which to gather information towards making a decision we still observed strong deployment of covert attention. On 'Simultaneous' trials, our subjects seemed to decrease their cognitive effort by concentrating their initial covert evaluation and subsequent initial value-based saccade towards stimuli presented on the bottom of the screen. As a rule this would have still allowed subjects to saccade to the more valuable option more times than not. In 'Simultaneous' trials, there was also clear evidence of covert attentional processes in both the direction of the second saccade, and the total amount of information subjects chose to gather.

Our results therefore provide further suggests that covert attention is a factor that needs to be considered when considering the mechanisms by which subjects make decisions in primate decision making paradigms. Although we did not perform a formal test of this, the fact that saccade behaviour appears to be more strongly influenced by the value of unattended compared to attended stimuli implies that covert attention may play a stronger role than overt attention. This is in direct conflict with the well described behaviour in the human literature and suggests that the mathematical approximations of evidence accumulation and decision making described by attentional drift diffusion models may not apply to how primates make simple economic decisions (Krajbich et al., 2010, Krajbich et al., 2012, Krajbich and Rangel, 2011).

Again, just as in Chapter 3, this indicates saccadic behaviour does not reflect the way in which subjects are comparing information. As previously discussed, the neural circuits that likely support these covert valuation and attention processes are likely to lie in subcortical structures known to perform value based computations such as Caudate or Substantia Nigra (Kim and Hikosaka, 2013, Yasuda et al., 2012). These regions are heavily implicated in the control of eye movements (Hikosaka et al., 2014, Beckstead and Frankfurter, 1982, Parent et al., 1983). Furthermore, it is unlikely that PFC regions are the ones that drive this fast initial saccade because the because the latencies of these initial saccades (approximately

150ms) are often faster than the latency of value coding in PFC (Kennerley et al., 2009, Strait et al., 2014, Hayden and Platt, 2010, Roesch and Olson, 2003, Hikosaka et al., 2014). One finding in the human fMRI literature is that BOLD signal in vmPFC correlates positively with the value of what is being attended and negatively with that of what is unattended (Lim et al., 2011). Although (as described above) it is unlikely that in 'Simultaneous' trials vmPFC value comparisons are driving this overt attention, it is possible that this overt value comparison is being biased through the control of overt attention by subcortical systems or possibly through indirect anatomical connections to vmPFC from subcortical nuclei.

Primate Multi-Attribute Information Gathering Strategies

Despite the influence of covert attention on saccade in 'Simultaneous' trials, these and 'Information Gathering' trials may provide some insight into information gathering behaviour in primate subjects. Unlike many of the classical multi-attribute decision making paradigms, the paradigm in this task was relatively simple with only two options and two attributes for subjects to consider. The data of Sundström (1987) suggest that given such a small number of attributes to consider, subjects would tend to prefer an option based comparison strategy. Furthermore, the data of Arieli (2011) suggests that subjects may be more likely to make attribute saccades when the trial was more difficult. However, we found that in 'Simultaneous' trials there was a similar number of option and attribute based saccades overall, and most importantly, the direction of this second saccade was heavily dependent on the EVs of the fixated option and the unfixated option. This implies that our subjects chose to make their second saccade toward the more valuable side and therefore by extension, the side that they would most likely go on to choose. Another important implication is that when primate subjects are allowed to gather information in any way they choose, they prefer to use covert attention rather than overt saccades. However, the discrepancies between our data and these previous studies may be attributable to task

differences, perhaps the most important being task difficulty. In the tasks described in the psychology literature, human subjects often have to make decisions based on many attributes and these decisions often have no 'right' answer. In contrast, our experiment only required subjects to evaluate and integrate information about two attributes and two options, such that there was always one clear objectively optimal choice. Therefore subjects do not need to deliberate like human subjects do on multi-attribute decision making paradigms. Furthermore, the influence of covert attention on the decision process means that saccades themselves have less relevance to the information gathering process.

We have demonstrated that by manipulating the manner in which information was presented to subjects, we could successfully change their choice behaviour. The fact that optimality changes between 'Option' and 'Attribute' trials suggests that different computational processes, and therefore potentially different brain areas, may be used for each trial type. Fellows (2006) found that patients with vmPFC lesions gather the same amount of information as control subjects, but are less likely to perform attribute based comparisons, implying that they may be unable to compare information in this manner. They also tended to make different choices compared to controls. We found that subjects gathered similar amounts of information having performed initial option and attribute comparisons in 'Information Gathering' trials. Surprisingly however, unlike the data of Fellows (2006), we found that there was very little difference in choices (defined by overall optimality) between 'Option' and 'Attribute' trials. The most likely explanation for this discrepancy is that in this paradigm decisions can be solved equally easily through either type of information comparison because there is relatively little information to consider.

One finding of particular interest in the current study is that in 'Attribute' trials, subjects choose the direction of their third saccade (and therefore their information gathering strategy) based on the value of the previous information. They tended to look towards the side of the option that has the highest current value and therefore by extension, the one that they are more likely to eventually choose. This reflects features of a confirmation bias which

(although is broadly used throughout psychology) is commonly considered to pertain to the seeking or interpretation of evidence that supports a current hypothesis, expectation or belief (Nickerson, 1998). This result bears a resemblance to the simple rule based reasoning tasks used by Wason (1968). In this task subjects were presented with 4 cards (showing a vowel, a consonant, an even number and an odd number respectively) and given a rule (in this case if a card has a vowel (*P*) on one side then is has an even number (*Q*) on the other side). They were told to indicate which cards they would turn over to confirm/deny the rule. Choosing the vowel (*P*) and the odd number card (*not Q*) would allow for the optimal resolution of this problem. However, subjects tended choose the vowel card (i.e. the card which would confirm the rule) and rarely chose the odd number card (i.e. the card that disproves the rule) thereby exhibiting information gathering behaviour which was biased towards confirmation of the current hypothesis. Further evidence of a primate confirmation bias comes from the fact that in 'Option' trials, when subjects are presented with two low value initial stimuli, they tend to gather the most information (i.e., all 4 cues) instead of immediately choosing the unattended (but most likely more valuable) option. Some evidence suggests that confirmation biases are exaggerated when information is presented sequentially rather than simultaneously (Jonas et al., 2001). It is difficult to test whether this is also true of the current study because of the influence of covert attention which renders saccades uninterpretable from an information gathering point of view. However, in all three types of trial we observed a tendency for subjects to make saccades towards the more valuable options which by extension were more likely to be chosen.

Our results suggest that during multi-attribute value-based decision making, primate subjects can use covert attention to gather complex multi-attribute information and compute EVs which inform not only choices but also saccades during the choice phase. However when covert attention cannot be utilised, subjects make decisions about how much information to gather and where to gather it based on the currently known information. They also exhibit features of confirmation bias in both the amount and the position of future

information gathering. Finally, although subjects appear to consider all available information in ‘Simultaneous’ trials (presumably using covert attentional valuation processes for unattended cues), when the extra cost of time and effort is introduced in sequential ‘Information Gathering’ trials, subjects prefer to make inferences about some unknown information.

Chapter 5: Frames of Reference in the Prefrontal Cortex

What processes do organisms use to make decisions? Scientists have tackled this question in many different ways: 1) mathematical descriptions of choice behaviour (Kahneman D, 1979), 2) mathematical approximations of decision processes (Ratcliff, 1976, Krajbich et al., 2010, Louie et al., 2013), 3) plausible biophysical models of decision circuits (Wang, 2002), 4) empirical models of decision pathways (Padoa-Schioppa, 2011, Cisek, 2012). This chapter will concentrate on examining the evidence supporting various empirical models of decision making by asking whether inferences can be drawn about the arrangement of decision circuits within the brain based on the representation of value in different so called *frames of reference*. We will first briefly outline the current state of evidence pertaining to frames of reference specific computations in prefrontal cortex (PFC). We will then present a set of neurophysiological findings that support the idea of parallel reference frame specific value computations which unify findings from several other neurophysiological and lesion studies. Finally we will present evidence that information gathering processes are also subject to the same frames of reference implying the existence of a common value comparison system for choice and information gathering.

Introduction

Understanding the value of objects in our environment is a critical prerequisite to making optimal value based decisions. Many neurophysiological studies have found value representation across the PFC (Kennerley et al., 2009, Roesch and Olson, 2003, Strait et al., 2014, Padoa-Schioppa, 2009, O'Neill and Schultz, 2010, Lim et al., 2011, Blanchard et al., 2015). Furthermore, in decision making tasks correlates of 'chosen value' (i.e. a post-

decision signal) have been reported across PFC regions (Kennerley et al., 2009, Padoa-Schioppa and Assad, 2006, Cai and Padoa-Schioppa, 2012, Cai and Padoa-Schioppa, 2014). However this abundance of value correlates does not improve our understanding of the functional specialisations that exist within PFC based on the lesion evidence. (Bechara et al., 1994, Clark et al., 2008, Camille et al., 2011a, Camille et al., 2011b, Rudebeck et al., 2008, Rudebeck and Murray, 2011, Noonan et al., 2010, Walton et al., 2002). Consequently, the natural question is what critical decision making computations do specific regions of PFC perform?

Recently two divergent ideas surrounding the potential mechanisms and pathways involved in value based decision making have attained prominence (Cisek, 2012, Padoa-Schioppa, 2011). One of these models puts forward the concept of 'decision making through consensus': the notion that multiple cortical and subcortical areas are able to compute and/or compare value representations of various options through mutual inhibition and that decisions are reached through a consensus between regions (Cisek, 2012). Central to this hypothesis is the idea that different regions may compare options in different so called 'frames of reference' (see **Figure 5.1**), and that these computations can take place in a simultaneous and parallel manner during decision making (Cisek, 2006, Cisek, 2007a, Selen et al., 2012). The extensive anatomical connectivity pattern of the PFC suggests that such parallel processing could well be possible (see Chapter 2). The other model of prominence is called the goods based model. In this model, decisions are made solely through a serial pathway where internal and external properties of options are represented in various parts of the brain but are integrated to form 'offer values' only in OFC/vmPFC (Padoa-Schioppa, 2011). Comparison is then performed in the abstract 'goods space' before the chosen value is converted into an action plan in other PFC and premotor areas. The critical difference in experimental predictions between the two models are that the 'goods space model' predicts that comparison signals will only be observed in goods space and in no other frame of

reference. In contrast, the ‘decision making through consensus model’ predicts that value comparisons may in fact occur in many different frames of reference.

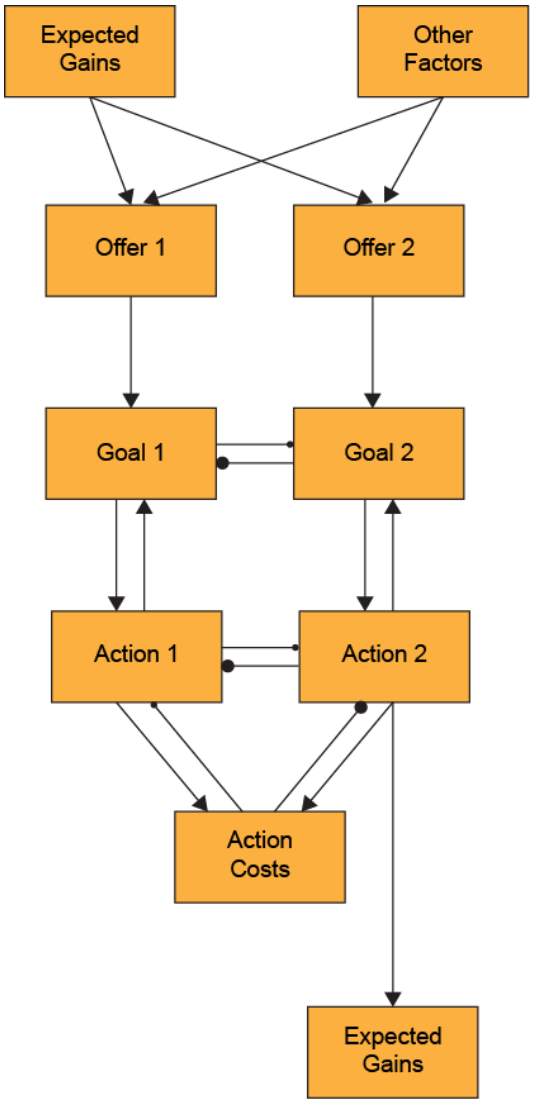

Figure 5.1: Decision making through consensus. A description of the potential processes related to making decisions through consensus. Different options can be compared at multiple levels (in this case at both the level of the goal and of the action) through mutual inhibition between pools of neurons representing the different options. Arrows indicate the direction of excitatory connections and line width the activity of these connections in an example given trial. Blunt arrows indicate inhibitory connections. Adapted from Cisek (2012).

Evidence for value representation in different frames of reference has been presented from various human and animal studies. Firstly, primate lesion experiments have demonstrated that damage to dorsal anterior cingulate cortex (ACC) causes selective

deficits in action guided decision making whereas damage to orbitofrontal cortex (OFC) specifically impairs stimulus guided decisions (Rudebeck et al., 2008, Kennerley et al., 2006). Similar studies examining human lesion patients have also replicated these findings (Camille et al., 2011b). The presence of these deficits heavily implies that ACC and OFC must engage in frame of reference specific computations. This is further backed up by the fact that the two regions show disparate connectivity patterns with ACC connecting strongly to the motor system and OFC receiving more basic sensory input (see Chapter 2 for more details). In further support of this inference, neurons in primate ACC have been observed to encode action value in various tasks (Hayden and Platt, 2010, Matsumoto et al., 2003, Matsumoto et al., 2007).

If OFC encodes value in a stimulus-based reference frame, this may in fact reflect a specific role in coding particular decision attributes. Subpopulations of neurons in OFC have been observed to specifically encode some decision relevant attributes such as risk and reward size (O'Neill and Schultz, 2010), reward size and delay (Roesch et al., 2006) and rewarding and aversive outcomes (Morrison and Salzman, 2009) in largely separate populations. Although not all OFC neurons discriminate attribute type (Kennerley et al., 2009, Padoa-Schioppa and Assad, 2006), this subset of neurons are potentially capable of making attribute specific comparisons when subjects are presented with multi-attribute decisions, a hypothesis that is supported by some human lesion data (Fellows, 2006, Fellows and Farah, 2005).

There is also strong evidence for frame of reference specific computations from the human imaging literature. Several studies have demonstrated a physical separation of the representation of some decision attributes such as risk and value (Wright et al., 2013), variance and skewness (Symmonds et al., 2011) or food and money (Levy and Glimcher, 2011) although there are also many examples of common valuation signals particularly in vmPFC (Lim et al., 2013, Lin et al., 2012, Chib et al., 2009). BOLD signal in ACC correlates with both the effort and reward for a given action (Croxson et al., 2009, Prevost et al., 2010)

which could be interpreted as an action value signal. Perhaps most convincingly, in a multi-attribute trinary decision making study, Hunt et al. (2014) showed that fMRI BOLD signal in intraparietal sulcus (IPS) correlated with the value comparison of the frame of reference (action value or stimulus value) that subjects used to make decisions on a trial-by-trial basis. IPS also showed specific functional connectivity with OFC and Putamen on trials where the stimulus value and action value drove the final choice, respectively. Interestingly, the dorsal medial frontal cortex encompassing ACC encoded an integrated value difference signal of both stimulus and action value. These results suggest that value is calculated in different frames of reference, and that value comparisons in different frames of reference can have differential influences on choice in different scenarios (Hunt et al., 2014).

Another frame of reference that has been suggested in the human imaging literature is that of attention. It has already been demonstrated that attention plays a large role in simple human decision making, with attention now incorporated into drift diffusion models of decision making (Krajbich et al., 2010, Krajbich et al., 2012, Krajbich and Rangel, 2011, Shimojo et al., 2003). An fMRI study by Lim et al. (2011) demonstrated that the BOLD signal in vmPFC and ventral striatum correlates positively with value of attended items and negatively with unattended items during simple binary choices, suggesting that value computations in OFC/vmPFC may be framed by current attention. While it remains to be seen whether the level of attention toward an item/option modulates value-coding neurons, it is noteworthy that subjects tend to attend more to the item they will choose (Krajbich et al., 2010). Given 'chosen value' coding is a ubiquitous signal in single neurons across the brain (Kennerley and Walton, 2011), taken together, these results imply that attention may well contribute, or even confound, the interpretation of 'chosen value' responses.

Despite the suggestive evidence that there may be at least three different valuation reference frames used by the brain (attentional, action and stimulus/attribute), three important questions remain unanswered: 1) how do these signals evolve throughout the decision making process, 2) are they relevant to decisions themselves, 3) are there

reference frames simultaneously represented during decision making? These questions are important to answer in order to understand whether attribute and action value signals simply provide inputs into a general value comparison process or whether frame of reference specific competition truly occurs in line with the 'decision making by consensus' hypothesis. The primate behavioural paradigm presented in this chapter provides a novel insight into these questions by providing the subject with three potential frames of reference with which to make decisions; general (attentional) value, action value and attribute specific value. Of particular importance is the fact that the task employs sequential and discreet information gathering events, allowing the representation of all three reference frames to be isolated and simultaneously tracked across the decision process. Furthermore, recording neuronal data from ACC, OFC and vmPFC (three regions heavily implicated in the aforementioned three value reference frames) as well as LPFC – an area implicated in attention and eye movements (Kennerley and Wallis, 2009c, Lebedev et al., 2004) - allows us to examine whether these frames of reference are specific to different regions.

Another vital process in decision making is information gathering. However, very little is known about how information gathering behaviour is controlled by the brain. As discussed and shown in Chapter 3, covert attention and the brain systems which support it may allow subjects to quickly extract salient features of our environment and then bias overt attention (saccades) towards the most relevant information (Gottlieb, 2012, Gottlieb et al., 2014, Anderson, 2013). Such a mechanism would be perfectly suited to solve a problem such how to find a preferred chocolate bar on a supermarket shelf. However, when our environments contain stimuli that are too complex or abstract to attend to covertly, for example the dozens of different properties displayed in an estate agent's window front, we may require a higher level strategy to optimise our information search.

In one of the few studies to examine the neural basis of information gathering, Fellows (2006) found that when presented with multiple options, each consisting of multiple attributes, patients with vmPFC lesions failed to follow the within-attribute comparison

strategy that control patients used, and instead used a within option comparison. This result implied that damage to vmPFC somehow prevented these patients either from employing the normal information gathering strategy, or disrupted some aspect of the valuation or value comparison process for decisions of this complexity.

In decision making, information itself has its own inherent value which can either be correlated or dissociated from the value of the potential options (Loewenstein, 1994). Neurons in OFC are known to encode the informativeness of stimuli, and given subjects prefer informative over uninformative cues, such OFC activity might reflect a value signal (Blanchard et al., 2015). Interestingly, these neurons are orthogonal to value coding neurons in OFC (Blanchard et al., 2015), in contrast to a common value signal for reward and information by dopamine neurons (Bromberg-Martin and Hikosaka, 2009). Importantly however, in this paradigm the informative cue has no bearing on the choice or information gathering behaviour; it simply randomly or fully predicts the outcome (Bromberg-Martin and Hikosaka, 2009, Blanchard et al., 2015). Thus, despite some evidence that neurons in the brain encode information that may have value to predicting outcomes, neuronal signals which reflect the value of information, or reflect future information gathering strategies which could influence choice, remain unknown.

The results reported in this chapter will test the ideas of frames of reference in a unifying paradigm. We will show that value signals are evident throughout PFC, yet they are also functionally dissociable. We will demonstrate a subpopulation of OFC neurons uniquely compare information in the frame of reference of attention, which is converted over the course of the trial to an attentional choice signal as the subjects near a response. A subpopulation of ACC neurons encode value in the frame of reference of actions during early information gathering, but this action value coding becomes weaker over the trial, and instead evolves into a signal encoding the chosen action as the response nears. In contrast, another subpopulation of OFC neurons encode value in the frame of reference of attributes throughout the entire trial, and also encodes an attribute specific correlate of choice which

peaks around the time of the response. Finally, we will present the completely novel finding that future information gathering strategies are also coded in the three described frames of reference prior to information gathering saccades. These results suggest unique contributions of different PFC areas in how determining not only what choice is most valuable, but also the process of deciding how to decide.

Methods

Subjects

Two adult male rhesus monkeys (*Macaca mulatta*), M and F, aged 5 and 6 years respectively were used as subjects in the study. All experimental procedures were approved by the Local Ethical Procedures Committee and carried out in accordance with the UK Animals (Scientific Procedures) Act. Fluid was controlled to ensure that subjects received their daily allotment of fluid during the course of the testing session.

Behavioural Protocol

A representation of the task structure is shown in **Figure 5.2B**. Subjects initiated the trial by maintaining saccadic fixation on the centre of the screen and central fixation of the joystick for 500ms. Once this was achieved two options were presented on the screen (seven visual degrees left and right of centre). Each option consisted of two pre-learned picture cues assigned to two different value attributes, probability of reward (10%, 30%, 50%, 70%, 90%) and magnitude of juice reward (0.15AU, 0.35AU, 0.55AU, 0.75AU, 0.95AU). An example set can be seen in **Figure 5.2A**. The actual reward magnitude was calculated by multiplying the arbitrary unit by the maximum available duration for each subject (M: 2500ms, F:2750ms). Reward magnitude was varied by manipulating the length of time a reward pump was driven and the absolute values (i.e. reward time) of each stimulus was different between subjects. Each cue was 7 visual degrees above/below the horizontal centre and 7 degree left/right of the vertical centre of the screen.

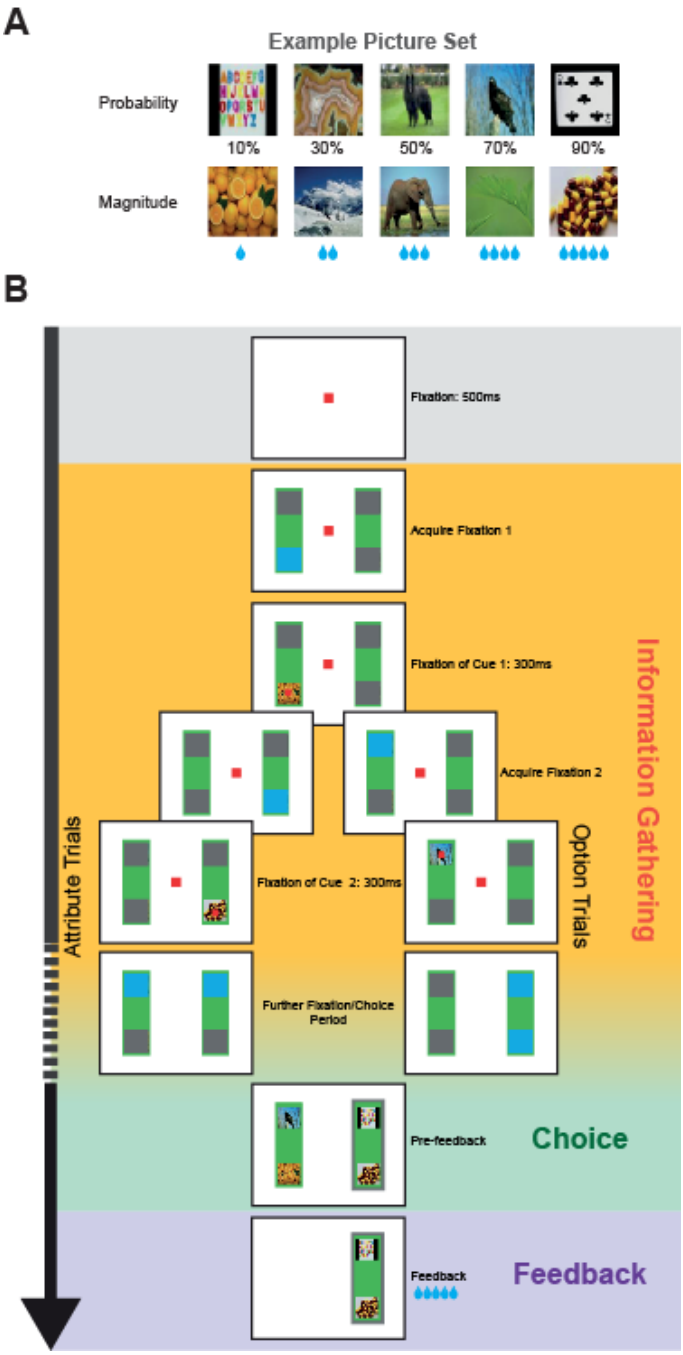

Figure 5.2: Task structure. **(A)** An example picture set. **(B)** Task design and structure. Subjects saccade around the screen to gather information about the two options. The initial two saccades are experimentally controlled, but then subjects are free to either gather more information (1 or 2 cues), or make a choice at any time (see Methods for further details).

At the start of a trial, all four picture cues were covered up by grey squares with the exception of one which was covered by a blue square (**Figure 5.2B**). The blue square informed the subject of the required location for a saccade. Once the subject fixated the blue

square, the picture cue replaced the blue cue and the subject was required to continuously fixate this location for 300ms. If continuous fixation was not achieved within 1200ms the trial was aborted and subjects received a short timeout. Once this fixation period was finished, the cue was covered with a grey square, and a second blue square was presented at a different location. The position of this blue square indicated to the subject the type of trial being experienced. If the blue square was for the second cue of the same option, subjects were in an 'Option' trial; if the blue square was for the same attribute cue of the second option then this was an 'Attribute' trial. Selection of trial types was pseudorandom. The subject was again required to acquire and maintain fixation of the second cue for 300ms before it was also covered up by a grey square. After this point, the subjects were now free to either i) choose an option using a joystick movement (left/right) based on the value of the currently known information or ii) view one or both of the remaining cues (in any order) before making a choice, with the third cue requiring 300ms of uninterrupted fixation before the fourth cue could be viewed. Importantly, however, they were prevented from viewing any cue that they had already seen. Once a response was made, all of the cues were uncovered (for 500ms for Subject F and 1000ms for Subject M) following which juice reward feedback was given with the probability and reward magnitude chosen by the subject.

'Option' and 'Attribute' trials were pseudorandomly interleaved during blocks of 50 trials. Between each of these blocks subjects were presented with a block of 25 trials where all of the picture cues were presented immediately (so called 'Simultaneous' trials). Data from these trials will not be discussed in this chapter.

Neuronal Recordings

Subjects were initially implanted with a titanium headpost in order to achieve head restraint (for eye tracking and electrophysiological recordings) before undergoing the behavioural protocol. Subjects were then subsequently implanted with bilateral circular

recording chambers (19mm internal diameter) which were located using pre-operative MRI and peri-operative stereotactic measurements. Post-operatively, gadolinium attenuated MRI imaging and electrophysiological mapping of gyri and sulci was used to confirm accurate chamber placement. The centre of each chamber was as follows; Subject M: left: AP 30.5, right: AP 33, Subject F: left: AP 34, right: AP 32.5. Craniotomies were then performed inside each chamber.

During each recording session, neuronal activity was measured using tungsten microelectrodes (FHC Instruments, Bowdoin, USA) which were driven through the brain using custom-built manual microdrives mounted to a grid. During a typical recording session, 8-20 electrodes were lowered bilaterally into multiple target regions until well isolated neurons were found. Neuronal data was recorded at 40kHz using a Plexon Omniplex system (Dallas, USA). Neuronal isolation was done through manual spike sorting using Plexon Offline Sorter (Dallas, USA).

Neuronal data was recorded from four target regions; ACC, LPFC, OFC, vmPFC. We considered ACC to be the entire dorsal bank of the anterior cingulate sulcus from AP 27-37. LPFC recordings spanned both dorsal and ventral banks of the principal sulcus but were concentrated towards the former. All neurons recorded lateral to the medial orbital sulcus and medial to the lateral orbital sulcus were considered OFC. Finally, vmPFC was considered to be a continuous region which was ventral of the genu of ACC and medial to the medial orbital sulcus. Electrophysiological and depth observations (i.e., gyral and sulcal landmarks, white matter zones) obtained from each electrode during the electrode lowering process were used to estimate the location of each recorded neuron with reference to previously obtained MRI images. The full reconstruction of all recorded neurons is shown in **Figure 5.3**.

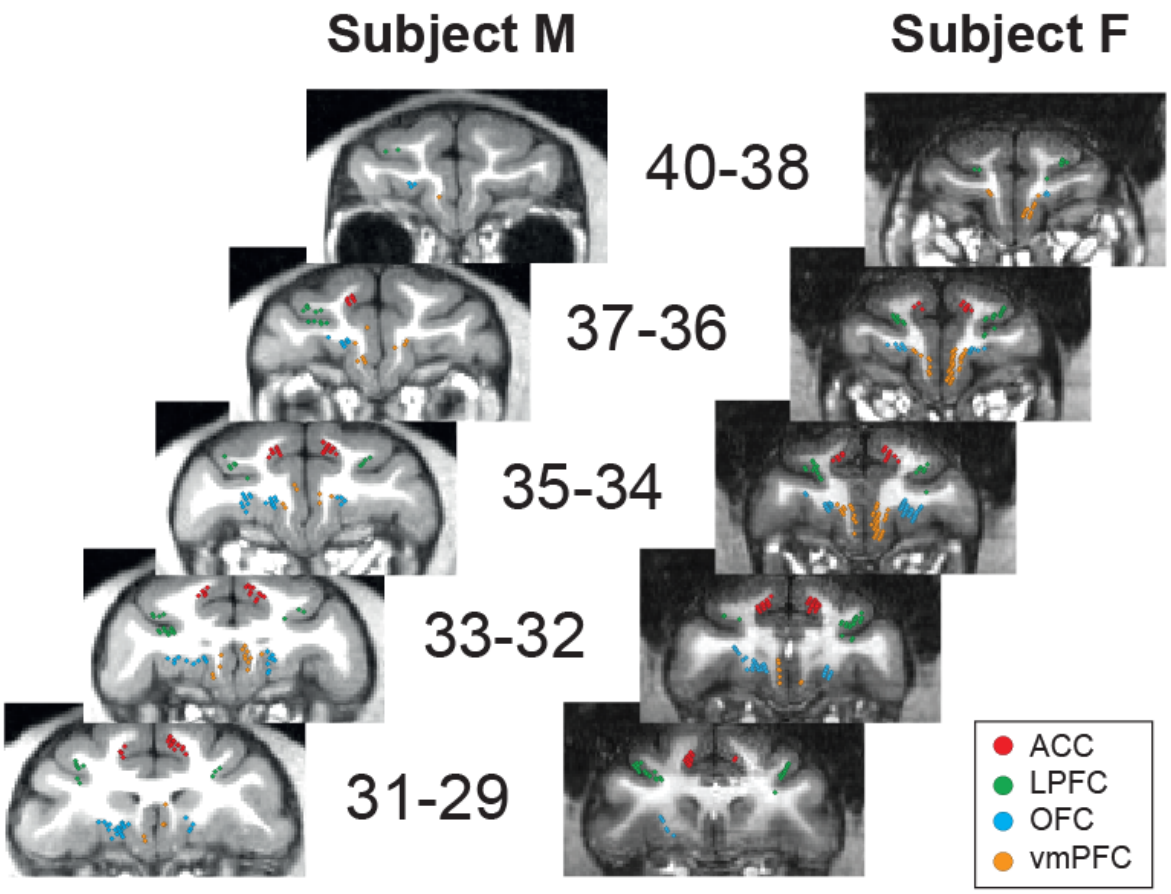

Figure 5.3: Approximate locations of neurons recorded from each subject. Each dot represents the location of one neuron. Each location was estimated based on depth of penetration, electrophysiological observations during recordings and registration of the recording grid to post-operative MRI scans. Numbers correspond to the distance in millimetres anterior of the inter-aural line (AP).

Data Analysis

All data analysis was performed using MATLAB (Mathworks, USA). In order to analyse the data the first step was to separate neuronal firing into four different epochs of relevance. These epochs were named Cue 1, Cue 2, Cue 3 and Response. Each 'Cue' epoch included data 200ms prior to cue onset and 600ms post-onset. The response epoch spanned from 900ms prior to response to 300ms after the response. These large windows were used in order to maximise the chances of observing task relevant computation although when analyses were repeated with shorter windows similar results were seen.

To test whether neurons encoded experimentally relevant variables, multiple linear regression was used (**Equation 5.1**) where Y was the dependent variable and X_n were the regressor and were weighted by coefficients β_n. Before performing regression, neuronal firing rate (FR) was normalised by subtracting the mean trial FR across all trials and dividing by the standard deviation in FR across all trials. For each of the cue epochs normalised FR was then averaged into 61 200ms time bins with a 10ms slide between adjacent bins. The same procedure was implemented for the response epoch with the exception that, due to its larger size, FR was averaged into 101 200ms bins. Regression was then performed at each time bin.

$$Y = \beta_0 + \beta_1 X_1 + \cdots + \beta_n X_n \qquad \textbf{Equation 5.1}$$

In order to correct for multiple comparisons across time, we performed a permutation test for every regressor and neuron. In order to do this we permuted the trial order of the normalised FR matrix (before averaging into time bins) and then performed the same regression on this permuted data. This was repeated for 1000 permutations. For each permuted regression and each regressor, the maximum and minimum T-statistic observed over all time bins was drawn to create a distribution of maximum and minimum T-statistics observed from random data. We then took the 97.5th percentile of the maximum T-statistic distribution and the 2.5th percentile of the minimum T-statistic distribution as the upper and lower significance thresholds respectively. This protocol meant that all regressors and neurons had individual threshold T-statistics for significance. Using this permutation method, a typical computed upper threshold T-statistic (computed from approximately 400-450 trials) was between 2.7 and 3, which was considerably stricter than the typical upper threshold computed from a standard T distribution.

In order to test our hypotheses pertaining to neuronal coding, we used several regression models at various epochs during the task and then used a ‘population code projection’ to examine the relationship between computations at each stage of the decision.

The regression model we used at Cue 1 can be seen in **Table 5.1**. We performed linear hypothesis tests (i.e. contrasts) for the following combinations of regressors in order to compute other task variables (abbreviations to be found in **Table 5.1**):

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Magnitude Value Top Left (MTL)	Value	7	Probability Value Bottom Left (PBL)	Value
2	Magnitude Value Top Right (MTR)	Value	8	Probability Value Bottom Right (PBR)	Value
3	Magnitude Value Bottom Left (MBL)	Value	9	Top Left Cue Presentation	Position Response
4	Magnitude Value Bottom Right (MBR)	Value	10	Top Right Cue Presentation	Position Response
5	Probability Value Top Left (PTL)	Value	11	Bottom Left Cue Presentation	Position Response
6	Probability Value Top Right (PTR)	Value	12	Bottom Right Cue Presentation	Position Response

Table 5.1: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to Cue 1 onset. All regressors refer to the properties of Cue 1.

Cue 1 Attentional Value = MTL + MTR + MBL + MBR + PTL + PTR + PBL + PBR

A neuron which has a positive coefficient for this linear hypothesis test fires more when the attended cue increases in value.

Cue 1 Action Value = (MTL + MBL + PTL + PBL) – (MTR + MBR + PTR + PBR)

A neuron which has a significant positive coefficient for this hypothesis test fires more as the value of stimuli presented on the left increase compared to those on the right.

Cue 1 Attribute Value = (MTL + MTR + MBL + MBR) – (PTL + PTR + PBL + PBR)

A neuron which has a significant positive coefficient for this hypothesis test fires more as the value of the magnitude stimulus increases compared to that of probability stimulus.

Cue 1 Top-Bottom Value = (MTL + MTR + PTL + PTR) – (MBL + MBR+ PBL + PBR)

A neuron which has a significant positive coefficient for this hypothesis test fires more as the value of stimuli presented on the top half of the screen increase compared to those on the bottom half.

In order to examine attention dependent value computations within PFC we constructed regressions models for the Cue 2 and 3 epochs which took into account value (of current and previous cues), trial type and third saccade direction (see **Table 5.2** and **5.3**). For the Cue 2 epoch we then performed linear hypothesis tests to compute the following variables (all abbreviations can be found in **Table 5.2**):

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Option Trial Cue 1 Value Horizontal 3rd Saccade (OptC1Horz)	Attentional Value	6	Attribute Trial Cue 1 Value Vertical 3rd Saccade (AttC1Vert)	Attentional Value
2	Option Trial Cue 2 Value Horizontal 3rd Saccade (OptC2Horz)	Attentional Value	7	Attribute Trial Cue 2 Value Vertical 3rd Saccade (AttC2Vert)	Attentional Value
3	Option Trial Cue 1 Value Diagonal 3rd Saccade (OptC1Diag)	Attentional Value	8	Attribute Trial Cue 1 Value Diagonal 3rd Saccade (AttC1Diag)	Attentional Value
4	Option Trial Cue 2 Value Horizontal 3rd Saccade (OptC2Diag)	Attentional Value	9	Attribute Trial Cue 2 Value Horizontal 3rd Saccade (AttC2Diag)	Attentional Value
5	Left Response	Constant	10	Right Response	Constant

Table 5.2: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to Cue 2 onset.

Cue 1 Option Trials = OptC1Horz + OptC1Diag

A neuron that encodes a positive coefficient for this linear hypothesis test fires more (at Cue 2) as the value of the Cue 1 increases on 'Option' trials.

Cue 2 Option Trials = OptC2Horz + OptC2Diag

A neuron that encodes a positive coefficient for this linear hypothesis test fires more as the value of Cue 2 increases on 'Option' trials.

Cue 1 Attribute Trials = AttC1Horz + AttC1Diag

A neuron that encodes a positive coefficient for this linear hypothesis test fires more as the value of Cue 1 increases in 'Attribute' trials.

Cue 2 Attribute Trials = AttC2Horz + AttC2Diag

A neuron that encodes a positive coefficient for this linear hypothesis test fires more as the value of Cue 2 increases in 'Attribute' trials.

We also constructed a similar regression model at Cue 3 which again included all current and previous value information, trial type and third saccade direction (see **Table 5.3**). Linear hypothesis tests were performed on the beta coefficients of the following regressors to compute other task variables (all abbreviations can be found in **Table 5.3**):

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Option Trial Cue 1 Value Horizontal 3rd Saccade (OptC1Horz)	Value	7	Attribute Trial Cue 1 Value Vertical 3rd Saccade (AttC1Vert)	Value
2	Option Trial Cue 2 Value Horizontal 3rd Saccade (OptC2Horz)	Value	8	Attribute Trial Cue 2 Value Vertical 3rd Saccade (AttC2Vert)	Value
3	Option Trial Cue 1 Value Diagonal 3rd Saccade (OptC1Diag)	Value	9	Attribute Trial Cue 1 Value Diagonal 3rd Saccade (AttC1Diag)	Value
4	Option Trial Cue 2 Value Horizontal 3rd Saccade (OptC2Diag)	Value	10	Attribute Trial Cue 2 Value Horizontal 3rd Saccade (AttC2Diag)	Value
5	Option Trial Cue 3 Value Horizontal 3rd Saccade (OptC3Diag)	Value	11	Attribute Trial Cue 3 Value Horizontal 3rd Saccade (AttC3Diag)	Value
6	Left Response	Constant	12	Right Response	Constant

Table 5.3: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to Cue 3 onset.

Cue 1 Option Trials = OptC1Horz + OptC1Diag

Cue 2 Option Trials = OptC2Horz + OptC2Diag

Cue 3 Option Trials = OptC3Horz + OptC3Diag

Cue 1 Attribute Trials = AttC1Horz + AttC1Diag

Cue 2 Attribute Trials = AttC2Horz + AttC2Diag

Cue 3 Attribute Trials = AttC3Horz + AttC3Diag

Previous Cue Saccaded Side = AttC1Diag + AttC2Horz

A neuron which encodes this coefficient positively fires more (at Cue 3) as the value of the previous cue located on the side of current attention increases in 'Attribute' trials.

Previous Cue Unsaccaded Side = AttC2Diag + AttC1Horz

A neuron which encodes this coefficient positively fires more (at Cue 3) as the value of the previous cue located on the side *away from* current attention increases in 'Attribute' trials.

In order to examine whether neuronal firing also correlated with various task parameters at Cue 2, we used a simple regression model which encoded only the value of the current cue separated by presentation side (L/R) and attribute type (P/M). It also contained binary regressors for whether the Cue 2 side was chosen differentially based on the current attribute type. A similar binary regressor was included for when the cue was unchosen. Other binary regressors that we included were for final response direction (left/right), whether the first or second option was chosen and whether the attended stimulus was eventually chosen. An accurate description of all regressors entered into this model can be found in **Table 5.4**. We performed linear hypothesis tests on the following regressors to examine current action and attribute value (all abbreviations can be found in **Table 5.4)**:

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Probability Value Left (PL)	Value	6	Left-Right Response	Response
2	Probability Value Right (PR)	Value	7	Current Cue Side Chosen (Probability Value – Magnitude Value)	Attribute Specific Choice
3	Magnitude Value (ML)	Value	8	Current Cue Side Unchosen (Probability Value – Magnitude Value)	Attribute Specific Choice
4	Magnitude Value (MR)	Value	9	Constant	Constant
5	Choose Option 1 - 2	Temporal Choice	10	Choose Attended Side	Attentional Choice

Table 5.4: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to Cue 2 onset. All regressors refer to the properties of Cue 2.

Current Cue Attentional Value = (PL + PR + ML + MR)

Current Cue Action Value = (PL + ML) – (PR + MR)

Current Cue Attribute Value = (PL + PR) – (ML + MR)

At Cue 3 we used a similar regression model to Cue 2 which encompassed only current cue value split by attribute and presentation and again binary regressors for response, current cue choice differentially by attribute and previous cue choice (i.e. Cue 2) differentially by attribute (see **Table 5.5** for a full description of the model). The same linear hypothesis tests were performed as for the Cue 2 epoch regression to compute *Current Cue Action Value* and *Current Cue Attribute Value* at Cue 3.

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Probability Value Left (PL)	Value	7	Left-Right Response	Response
2	Probability Value Right (PR)	Value	8	Cue 2 Chosen (Probability – Magnitude)	Attribute Specific Choice
3	Magnitude Value (ML)	Value	9	Cue 2 Unchosen (Probability – Magnitude)	Attribute Specific Choice
4	Magnitude Value (MR)	Value	10	Constant	Constant
5	Cue 3 Chosen (Probability – Magnitude)	Attribute Specific Choice	11	Choose Attended Side	Attentional Choice
6	Choose Option 1 -2	Temporal Choice			

Table 5.5: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to Cue 3 onset. All regressors refer to the properties of Cue 3.

Finally, we also performed a regression analysis of neuronal data aligned to the subjects' joystick response. In order to perform this analysis, several assumptions were made. Because subjects were not forced to view every cue before making a choice, we made the assumption that if a cue was not viewed by the time of response, the subject inferred the cue's value as the average value (i.e. a rank of 3 out of 5) and made decisions based on this inference. Also, subjects had a tendency to fixate a new picture while already making a joystick response, therefore in these cases the last cue that subjects viewed could not possibly have influenced choice; therefore, any occasion where final cue acquisition occurred <100ms before the joystick response we considered the final cue to not have been viewed for the purpose of the regression, thereby assuming it was of average value. The regression model included the probability and magnitude values presented in the trial sorted by side of presentation and whether each was chosen or unchosen. There were also binary

terms for response direction, for whether the side of the final stimulus was chosen or unchosen, for whether subjects were choosing the first or second presented option, for when subjects were choosing the side of the final attended cue differentially based on its attribute type and finally for when subject did not choose the side of the final attended cue differentially based on attribute type. A full description of this model can be found in **Table 5.6**. Note that we define (un)chosen value as the sum of the value of the two cues rather than the product in order to simplify the analysis. We then performed the following linear hypothesis tests to examine other task variables (all abbreviations can be found in **Table 5.6**):

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Chosen Probability Value Left (ChPL)	Value	8	Unchosen Magnitude Right (UnMR)	Value
2	Chosen Magnitude Value Left (ChML)	Value	9	Left Response (Lresp)	Left Response
3	Chosen Probability Value Right (ChPR)	Value	10	Right Response (Rresp)	Right Response
4	Chosen Magnitude Value Right (ChMR)	Value	11	Last Attended Side Chosen Differentially by Attribute	Attribute Specific Choice Correlate
5	Unchosen Probability Value Left (UnPL)	Value	12	Last Attended Side Unchosen Differentially by Attribute	Attribute Specific Choice Correlate
6	Unchosen Magnitude Left (UnML)	Value	13	Last Attended Side Chosen - Unchosen	Attentional Choice
7	Unchosen Probability Right (UnPR)	Value	14	Choose Option 1 – Choose Option 2	Stimulus Specific Choice Correlate

Table 5.6: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to response.

Action Specific Chosen Value = (ChPL +ChML) – (ChPR + ChMR)

Attribute Specific Chosen Value = (ChPL + ChPR) – (ChML +ChMR)

Response = Lresp – Rresp

To determine how much variance in firing rate was accounted for by value, action value and attribute value at Cue 1, we calculated the coefficient of partial determination (CPD) and then averaged this value across all neurons in each brain area to obtain an estimate of the amount of variance explained for each regressor at the population level. The CPD for regressor Xi is defined as:

$$CPD(Xi) = \{SSE(X - i) - SSE(X - i, Xi)\}/SSE(X - i) \qquad \textbf{Equation 5.2}$$

where SSE(X) refers to the sum of squared errors in a regression model that includes a set of regressors X, and X_{-i} a set of all the regressors included in the full model except X_i. Binary terms for stimulus position were also included and full list of regressors can be found in **Table 5.7**.

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Value	Value	5	Top Left Cue Presentation	Position Response
2	Action Value (Left – Right)	Action Value	6	Top Right Cue Presentation	Position Response
3	Attribute Value (Probability – Magnitude)	Attribute Value	7	Bottom Left Cue Presentation	Position Response
4	Left – Right Side Presentation	Side Response	8	Bottom Right Cue Presentation	Position Response

Table 5.7: A list of regressors and their interpretations used in the CDP of neuronal firing rate aligned to Cue 1 presentation.

One of the main objectives of this chapter is to look at the evolution of early value signals in various frames of reference across a trial in order to understand how these neurons perform frame of reference relevant computations. In order to achieve this goal, we used the mean beta coefficients (from 100-500ms post-cue onset) for the 'attentional value', 'attribute value' and 'action value' regressors at Cue 1 as 'population value codes'. We then regressed each of these fixed population codes against sliding population codes for other variables and regressors obtained at time-points during the Cue2, Cue 3 and response epochs. This analysis tested whether the strength and direction of the encoding of a regressor at Cue 1 was related to the strength and direction of coding of task-related variables at other time points. We compared the strength of these 'sliding projection analyses' (i.e., correlations) between regions every 100ms using linear hypothesis tests (with Bonferroni correction for repeated tests across areas), in order to observe whether one or more regions had a significantly stronger correlation than others.

To test whether any neurons in PFC encoded future information gathering behaviour, we used a regression model at Cue 1 which predicted third saccade behaviour in three separate frames of reference (**Table 5.8**). In 'Option' trials we used a binary regressor which described whether subjects made a horizontal or diagonal third saccade (i.e. in the frame of reference of the *currently attended attribute* but not in an attribute frame of reference). We also used a binary co-regressor which described whether the third cue was of probability or magnitude type in 'Option' trials. In 'Attribute' trials, we included a binary regressor which described whether subjects made vertical or diagonal saccades (i.e. whether they chose to saccade towards or away from the currently attended side) and a second binary co-regressor which described whether the saccade finished on the left or right side of the screen (i.e. in the action frame of reference). We also accounted for attentional, action and attribute value coding in the same model. We also used the same regression model at Cue 2 and Cue 3 to investigate whether neuronal firing predicted third saccade behaviour (i.e. post-

saccade signals) with the exception that at each of these cues, we also added value regressors for previously seen cues along with the currently attended cues.

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Probability Value Left (PL)	Value	6	Left-Right Response in 'Option' trials	Response
2	Magnitude Value Left (ML)	Value	7	Magnitude vs Probability Saccades in 'Option' trials	Attribute Saccade
3	Probability Value Right (PR)	Value	8	Horizontal vs Diagonal Saccades in 'Option' trials	Attentional Saccade
4	Magnitude Value Right (MR)	Value	9	Left vs Right Saccades in 'Attribute' trials	Action Saccade
5	Left-Right Response in 'Option' trials	Response	10	Vertical vs Diagonal Saccades	Attentional Saccade

Table 5.8: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to Cue 1 for predicting third saccade behaviour.

Results

Two adult rhesus macaque monkeys (Subjects M and F) were trained to perform a multi-attribute sequential information gathering and decision task (**Figure 5.2**), in which they made manual responses to choose one of two options based on information they acquired through saccading to cues on the screen (see Methods). Subject M performed 32 recording sessions completing an average of 445 trials per session. Subject F performed an average of 394 trials per session over 25 recording sessions. A more comprehensive breakdown of behaviour can be found in Chapter 4. During the task we recorded single neuronal activity from four regions of PFC; ACC, LPFC, OFC and vmPFC. The total number of neurons recorded from each region for each subject can be found in **Table 5.9** and the estimated locations of each neuron can be seen in **Figure 5.3**.

	ACC	**LPFC**	**OFC**	**vmPFC**
Subject M	101	49	87	35
Subject F	97	107	108	125
Total	198	156	195	160

Table 5.9: The numbers of neurons recorded in each brain area split by subject.

Neurons in PFC Encode Value in Different Frames of Reference at Cue 1

A critical feature of the task that subjects performed was that each presented cue consisted of three important properties: it's attentional value (invariant of any action or

attribute type), it's action (i.e. the joystick response required to choose the option pertaining to the cue) and it's attribute (i.e. whether it was of magnitude or probability type).

In order to understand whether neurons in PFC encoded value in these three reference frames, we performed a sliding multiple regression of mean neuronal firing rate at Cue 1 presentation against several regressors (see Methods, **Table 5.1**). We then separated neuronal selectivity by area to investigate region specific computations within PFC (**Figure 5.4A**). Significant proportions of neurons across all four regions encoded the attentional value of Cue 1 (binomial test, all $p<8x10^{-16}$). However, a significantly smaller proportion of vmPFC neurons encoded attentional value compared to all other regions (pairwise Chi^2 test, all $p<0.007$). Furthermore, neurons in both ACC and OFC were significantly more likely to encode attentional value than LPFC (pairwise Chi^2 test, all $p<0.003$). From these findings we concluded that although value coding is particularly strong in ACC and OFC, value representations are a ubiquitous signal throughout PFC.

We then investigated whether any neurons discriminated their value code based on the action associated with the cue (i.e. whether they encoded action value). A significant subset of ACC and LPFC neurons (approximately 20%) encoded action value (binomial test, all $p<3x10^{-7}$). Both of these populations were significantly greater than OFC and vmPFC, which themselves only encoded action value at chance level (pairwise Chi^2 test, all $p<0.006$, binomial test for OFC and vmPFC, both $p>0.05$). **Figure 5.4B** shows an example ACC action value neuron which encodes the value of a cue presented on the left with a positive relationship, yet reverses this relationship when the presentation is on the right. It is possible that rather than encoding action value, the ACC and LPFC neurons are in fact encoding value with reference to various parts of space. If this were true, then this would be indistinguishable from action value in the left/right domain. However one might expect to also find neurons which differentiated value when cues were presented on the top part of the screen compared to the bottom part. Indeed, in LPFC an equally prevalent population of top-

bottom value neurons were observed (Chi^2 test, p>0.05), whereas in ACC this population

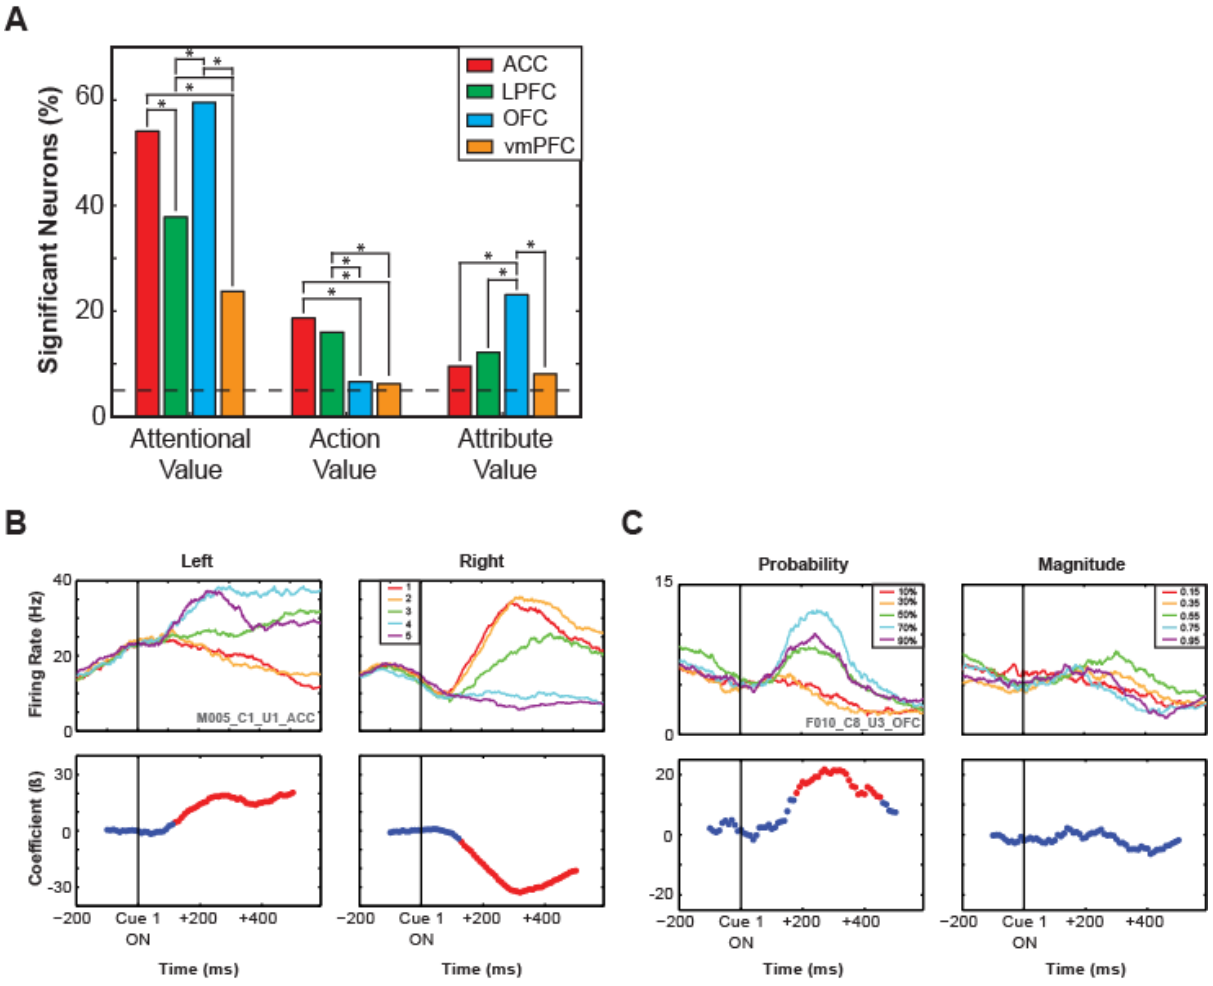

Figure 5.4: Frames of reference specific value coding in PFC **(A)** Percentage of neurons encoding attentional value, action value and attribute value at Cue 1. Asterisks indicate populations of neurons that are significantly different (Chi^2 test). The dashed line indicates the 5% chance level. **(B)** The firing rate of an example action value ACC neuron aligned to Cue 1 onset (solid black line) separated by side of presentation and by cue value (coloured lines). The plot below represents the beta coefficients for a sliding linear regression of firing rate against left and right cue value respectively. Red dots indicate significant coefficients (defined by a permutation test) and blue dots coefficients that are non-significant. **(C)** The firing rate of an example attribute value OFC neuron separated by attribute type and cue value. All other features are the same as (B).

was significantly smaller than the left-right value population (Chi^2 test, p<5.4x10^{-4}). From this we conclude that LPFC is likely to encode value differentially based on spatial position whereas ACC is more likely to encode value in the frame of reference of actions necessary for choice.

Next we considered whether neurons encode attribute value (i.e. neurons which discriminate value more for one attribute than another). In this case, OFC was found to be significantly more likely to do this (**Figure 5.4A**) (pairwise Chi^2, all $p<0.009$). **Figure 5.4C** is an example of such a neuron. This OFC neuron encoded that value of probability cues irrespective of which side of the screen the cue was presented, yet this neuron did not encode value of magnitude cues. However, it should be noted that ACC and LPFC also contained a small population of attribute value coding neurons which exceeded chance level (binomial test, both $p<0.02$).

To further examine the action value codes we made a scatter plot of the maximum left value encoding (i.e. T-statistic) against maximum right value encoding for neurons in each region (**Figure 5.5A**). If neurons were insensitive to the presentation side of the cue, the expected observation would be that all neurons would fall approximately on the equality line (dashed line). This was seen to be the case for OFC which also had a significantly stronger Person's correlation coefficient when compared to all other regions (Fisher's r-to-z transformation, z test, all $p<0.005$). From this we concluded that OFC neurons are significantly less likely than other regions to discriminate action value. In contrast many ACC neurons did not lie on the equality line and coupled with the fact that many neurons were seen to encode value, this suggests a strong influence of action on value computation at Cue 1.

We then repeated this analysis but instead used the maximum probability and magnitude value coding in order to probe attribute specific value coding (**Figure 5.5B**). In this analysis, ACC was seen to have a significantly stronger correlation coefficient compared to all other regions (Fisher's r-to-z transformation, z test, all $p<2x10^{-10}$), indicating that ACC is extremely insensitive to attribute specific value, whereas OFC had many neurons which fell well off the equality line, indicative of attribute-specific value coding.

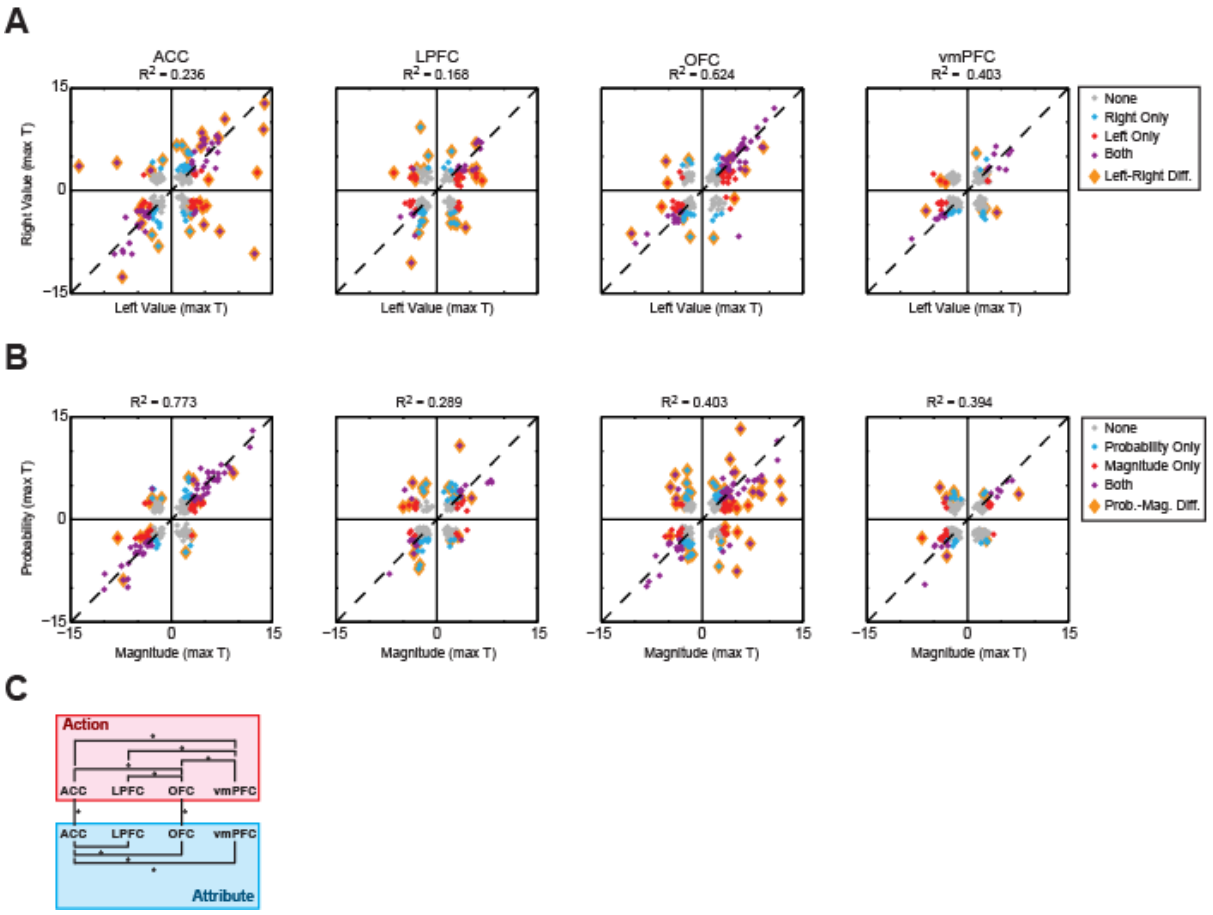

Figure 5.5: Maximum action and attribute value encoding across regions. **(A)** Maximum T-statistics for right and left value encoding after cue presentation for individual neurons. The colour of the dot indicates whether the regression model considered the neuron to encode either, both and neither left and right value (see figure legend). Diamonds indicate neurons that were considered to differentially encode left and right value (i.e. encode action value). **(B)** Maximum T-statistics for probability and magnitude value encoding after cue presentation for individual neurons. Colour legends are the same as (A) except for probability and magnitude value encoding. **(C)** A representation of the Person's correlation coefficients that are significantly different using Fisher's r-to-z transformation ($p<0.05$). Each area was compared with all other areas within each frame of reference (i.e. action or attribute) and with itself across frames of reference.

Having observed the differences in correlation coefficients between ACC and OFC within a frame of reference, we wanted to test whether the observed decrease in correlation was actually specific for the action frame of reference for ACC and the attribute frame of reference for OFC. In order to achieve this, we performed a statistical comparison of the correlation coefficients observed for each frame of reference for individual regions (**Figure**

5.5C). This test showed that the correlation coefficients were significantly different across frames in both ACC and OFC (Fisher's r-to-z transformation, z test, all $p<0.005$). One weakness of this correlation analysis is that areas that do not strongly encode value may also give low correlation coefficients in this analysis. Therefore this test may not be sensitive for areas such as LPFC and vmPFC.

In order to examine how much neuronal variance was explain by attentional value, action value and attribute value at the entire population level in each brain area, we calculated the coefficient of partial determination (CPD) during the Cue 1 epoch. This analysis included the aforementioned variables amongst its regressors (see **Table 5.7** for a full list of regressors). The CPD analysis found that at its peak, attentional value accounted for approximately 3% of the variance in neuronal firing in ACC and OFC (**Figure 5.6**). On the other hand, action and attribute value only accounted for 0.5-1% of variance in ACC and OFC, respectively (**Figure 5.6**). This analysis also revealed the time course and latency of the value signals in the three different reference frames. The attentional value signal was seen to be relatively fast, occurring as 100ms post-cue 1 onset, while the action and attribute signals appeared to be slightly slower, only occurring shortly before 200ms post-cue onset. These initial observations imply that early on during the trial, many neurons throughout PFC encode attentional value. Furthermore, significant subsets of ACC neurons encode action value and a significant subset of OFC neurons encode attribute value. Also, both action and attribute value coding occurs later than attentional value coding.

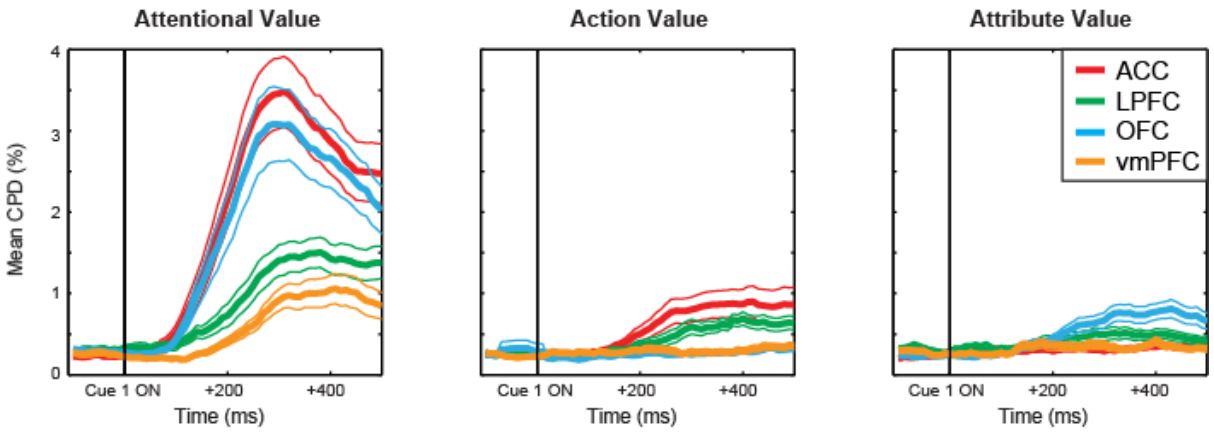

Figure 5.6: The time course of reference frame specific value coding. Mean coefficient of partial determination (CPD) across all neurons for value, action value and attribute value encoding over time split by region. Thick lines indicate the mean CPD for the population and the shaded line the SEM. All data is aligned to the Cue 1 onset (black line).

OFC Compares Value in an Attentional Frame of Reference

Having identified populations of neurons throughout PFC that encode the general (attentional) value of the first cue, we aimed to track the computations of these neurons over the course of the trial. By doing this we hoped to gain insight into whether these neurons simply encode the value of attended stimuli, or whether they perform task relevant computations (value summations or subtractions) framed by current attention. The first course of action was to examine attentional value coding during Cue 2 presentation. This period has a critical relevance to any attentional signal because at this point subjects could be diverted into either an 'Option' or 'Attribute' trial. The importance of this is that in 'Option' trials, the previously seen cue remained on the attended side of the screen, whereas in 'Attribute' trials the first cue was now on the unattended side of the screen (see **Figure 5.3** for example trials). Therefore, if value representation in any region was dependent on the attended side (given the task requires a left or right choice), then differential effects should be observed between trial types.

To test this we regressed the mean population code for attentional value at Cue 2 between 200-500ms post-cue onset against the mean population memory trace code for

Cue 1 value at Cue 2 during the same time period for both 'Option' and 'Attribute' trials. Uniquely, for the OFC population, this analysis showed a positive correlation between Cue 1 and 2 value coding specifically in 'Option' trials (indicating the population was encoding the Option sum), and a negative correlation between the two Cues in 'Attribute' trials (indicating the population was encoding the difference in value between Cue 1 and Cue 2), with a significant difference between the trial types (**Figure 5.7A**) (linear regression; $T_{('Option')}$=2.60, $T_{('Attribute')}$=-3.53, $p_{('Option')}$=0.01, $p_{('Attribute')}$<0.0005, linear hypothesis test of coefficients for 'Option' trials compared to 'Attribute' trials, T=19.21, $p<2x10^{-5}$). OFC therefore sums information on the attended side but subtracts value information from unattended side. This result implies that OFC computes value in the frame of reference of the attended versus the unattended side/option. Although this attentional value code is represented at the population level, individual examples of attentional coding can be observed in single neurons in OFC (**Figure 5.7B)**. This example neuron encodes the option sum value negatively in 'Option' trials, and encodes value difference negatively (i.e. with a preference for Cue 1) in 'Attribute' trials. Surprisingly, LPFC also showed a significant negative correlation reflecting a difference computation specific to 'Option' trials (**Figure 5.7A**) (linear regression; T=-2.27, p<0.03). This LPFC signal on 'Option' trials may reflect a signal for future information gathering, as large differences between Cue 1 and Cue 2 value on 'Option' trials would be associated with options of average value; hence further information may be necessary to decide which option to choose. Relatedly, such a value difference on 'Option' trials may be important for signalling which attribute to saccade to at Cue 3 for subsequent value comparisons and decision-making.

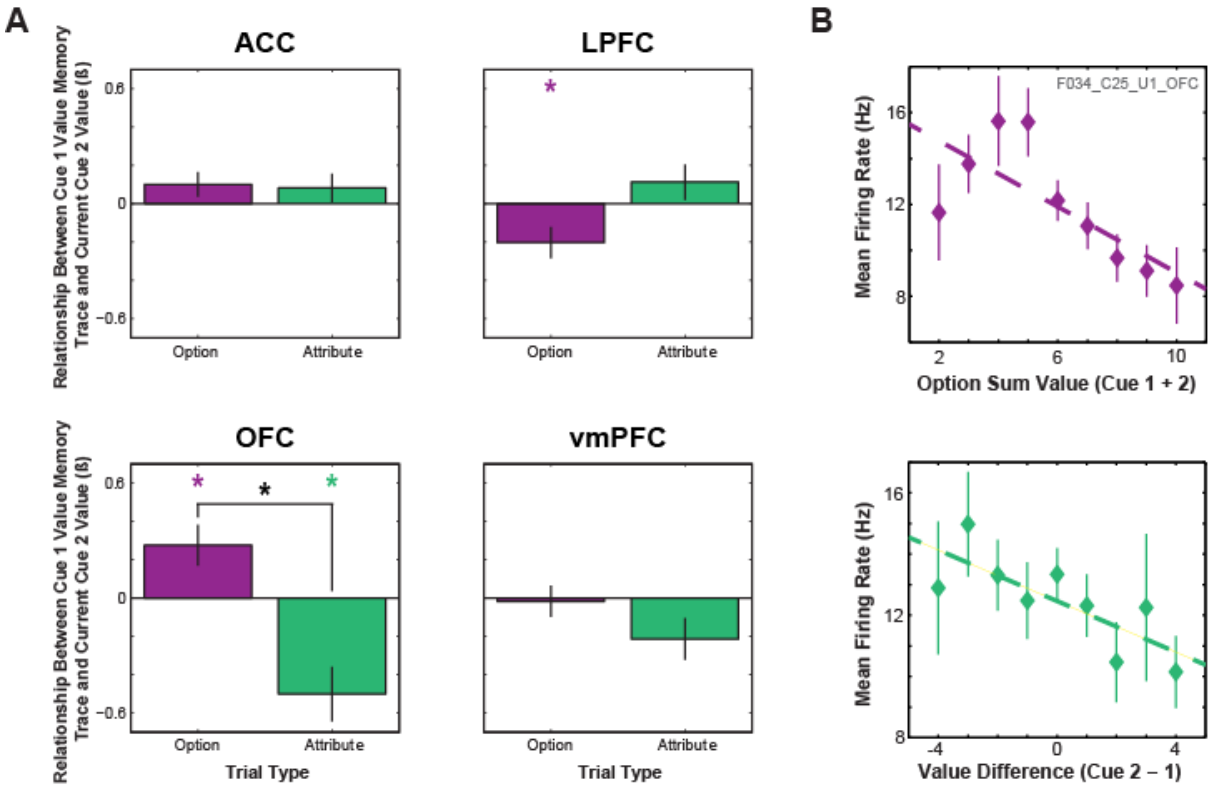

Figure 5.7: Population value computations at Cue 2. **(A)** The linear regression of the population Cue 1 attentional value memory trace for each trial type against Cue 2 attentional value population code. Coloured asterisks indicate beta coefficients that are significant ($p<0.05$) and black asterisks indicate coefficients that are considered significantly different from one another by a linear hypothesis test ($p<0.05$). **(B)** An example OFC neuron in 'Option' trials (top) and 'Attribute' trials (bottom). Each diamond indicates the mean firing rate at Cue 2 for a given option sum (top) or value difference (bottom). Vertical lines show the SEM and the dashed line the fit of the data using least-squares regression.

In order for OFC to compute attentional value difference in 'Attribute' trials, the value code for Cue 2 must be of opposite sign to the memory code of Cue 1. This begs the question: was the Cue 2 value code opposite to the original Cue 1 code or did the Cue 1 memory trace invert its code when Cue 2 was presented? To test this we performed a regression of the mean Cue 1 attentional value code at Cue 1 against the mean Cue 1 attentional value memory trace at Cue 2 for each trial type (**Figure 5.8A** and **5.8B**). When computing the mean code using 200-400ms cue onset, there was a significant negative correlation specific to 'Attribute' trials, implying a sign flip of the memory trace when attention was diverted away from the side of the Cue 1 (**Figure 5.8B**) (linear regression; $T=-2.11$, $p<0.04$). When the window was moved forward (100-300ms post cue onset) a significant

positive correlation was observed in OFC specifically in 'Option' trials, which would be expected if these neurons computed attended option value (**Figure 5.8A**) (linear regression; T=2.02, p<0.05). The fact that the two above results present themselves at slightly different time windows may imply that computation dynamics in the two trial types occur with different temporal profiles.

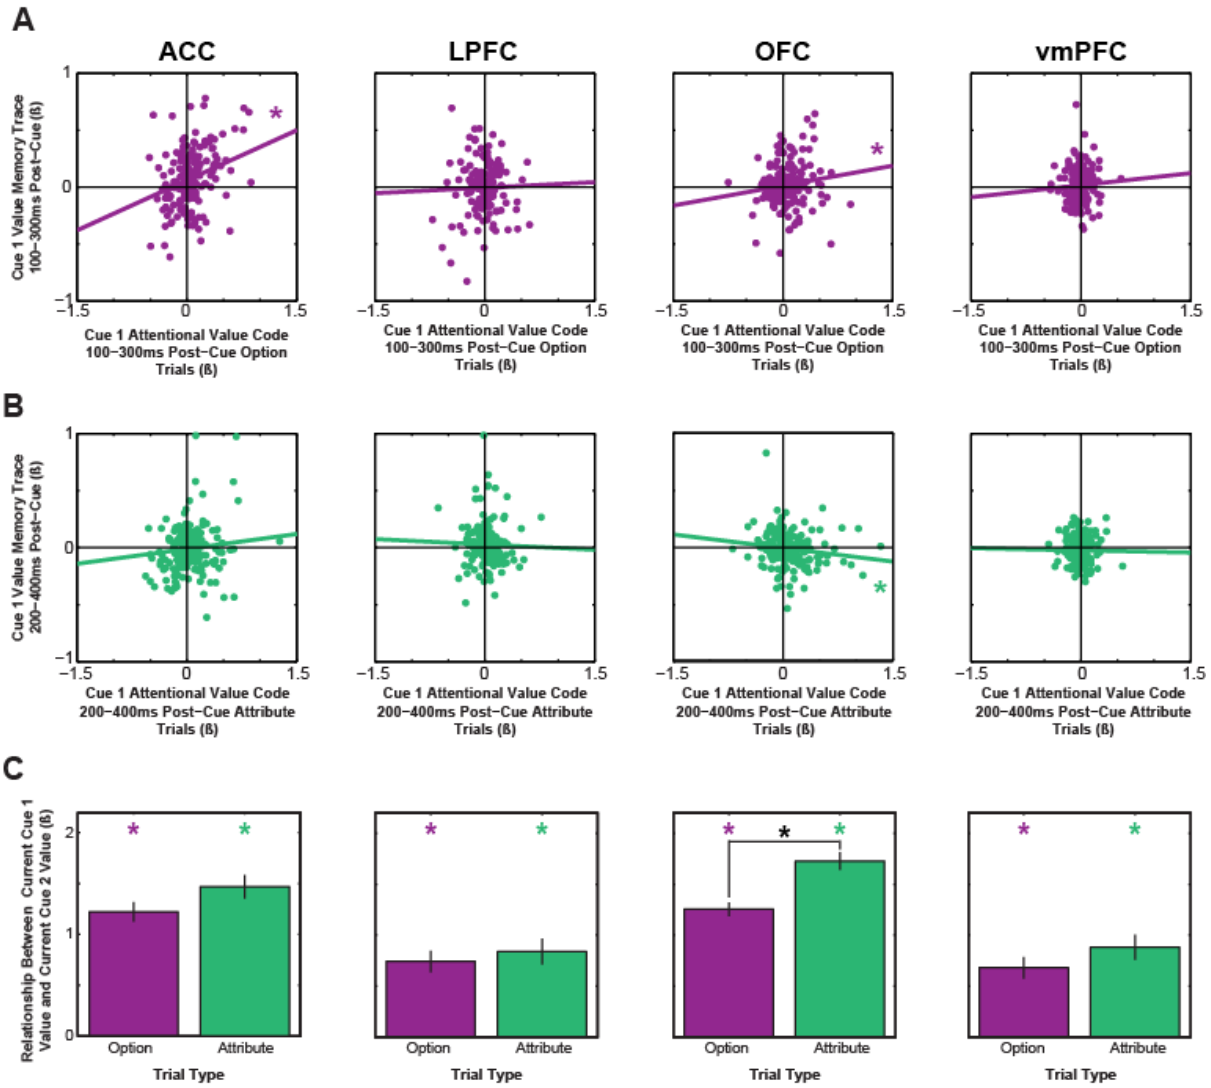

Figure 5.8: The relationship between Cue 1 attentional coding across cues. **(A)** Scatter plots of attentional value coding of Cue 1 at Cue 1 against the memory value code for Cue 1 at Cue 2 on 'Option' trials. Dashed lines indicate the least-squares regression slope and asterisks indicate significant relationships (p<0.05). **(B)** The same scatter plots for 'Attribute' trials. **(C)** The beta coefficients of a regression of Cue 1 attentional coding against Cue 2 attentional coding split by trial type.

We also performed an identical regression on the Cue 1 attentional value code (i.e. at Cue 1) against the Cue 2 attentional value code (at Cue 2) in each type of trial and found strong positive correlations in all areas (**Figure 5.8C**) (linear regression; all $T_{('Option')}>6.37$, all $T_{('Attribute')}>7.15$, all $p_{('Option')}<7x10^{-10}$, all $p_{('Attribute')}<7x10^{-12}$). However, there was a significant difference in the Cue1-Cue 2 relationship in OFC when compared between trial types (linear hypothesis test of the coefficient for 'Option' trials against the coefficient for 'Attribute' trials, $T=18.25$, $p<3x10^{-5}$), implying that the population representation of Cue 2 attentional value was weaker in OFC in 'Option' trials, although the functional relevance of this remains unclear. Therefore, we can conclude that the value of currently attended information is always represented in the same manner (if not always as strongly), but the memory representations of previous value information is flipped specifically in OFC when attention shifts to a different choice option.

We then went on to test attentional value computations at Cue 3. In 'Option' trials, the third cue was invariably presented on the opposite side of the screen to the first and second cues, and hence attention was involuntarily forced away from the side of Cues 1 and 2. On performing a linear regression of the mean value code at attended Cue 3 against the mean unattended code for the memory trace Option 1 value (i.e. Cue 1 + Cue 2) at Cue 3, we again found a negative correlation between the attended and unattended value codes in OFC (**Figure 5.9A**) ($T=-2.56$, $p<0.02$). Surprisingly, a borderline negative correlation was also observed in LPFC ($T=-1.86$, $p=0.067$). We then performed a sliding regression of the Cue 3 value code against the Option 1 memory trace code at Cue 3 to examine the time course of this computation. This analysis showed that in OFC the attentional comparison was strongest 100-400ms post-Cue 3 onset (**Figure 5.9B**). The LPFC negative correlation also began at approximately 100ms but was sustained past 500ms post-Cue 3 onset. We then repeated this analysis but used Cue 1 and Cue 2 value instead of Option 1 value. The negative correlation was borderline significantly stronger for Cue 1 than Cue 2 in OFC (**Figure 5.9C**) (linear hypothesis test of beta coefficients for Cue 1 against Cue 2; $T=3.22$,

p=0.073). For LPFC, only the Cue 2 value code correlated with the Cue 3 code (linear regression; T=-2.99, p<0.003).

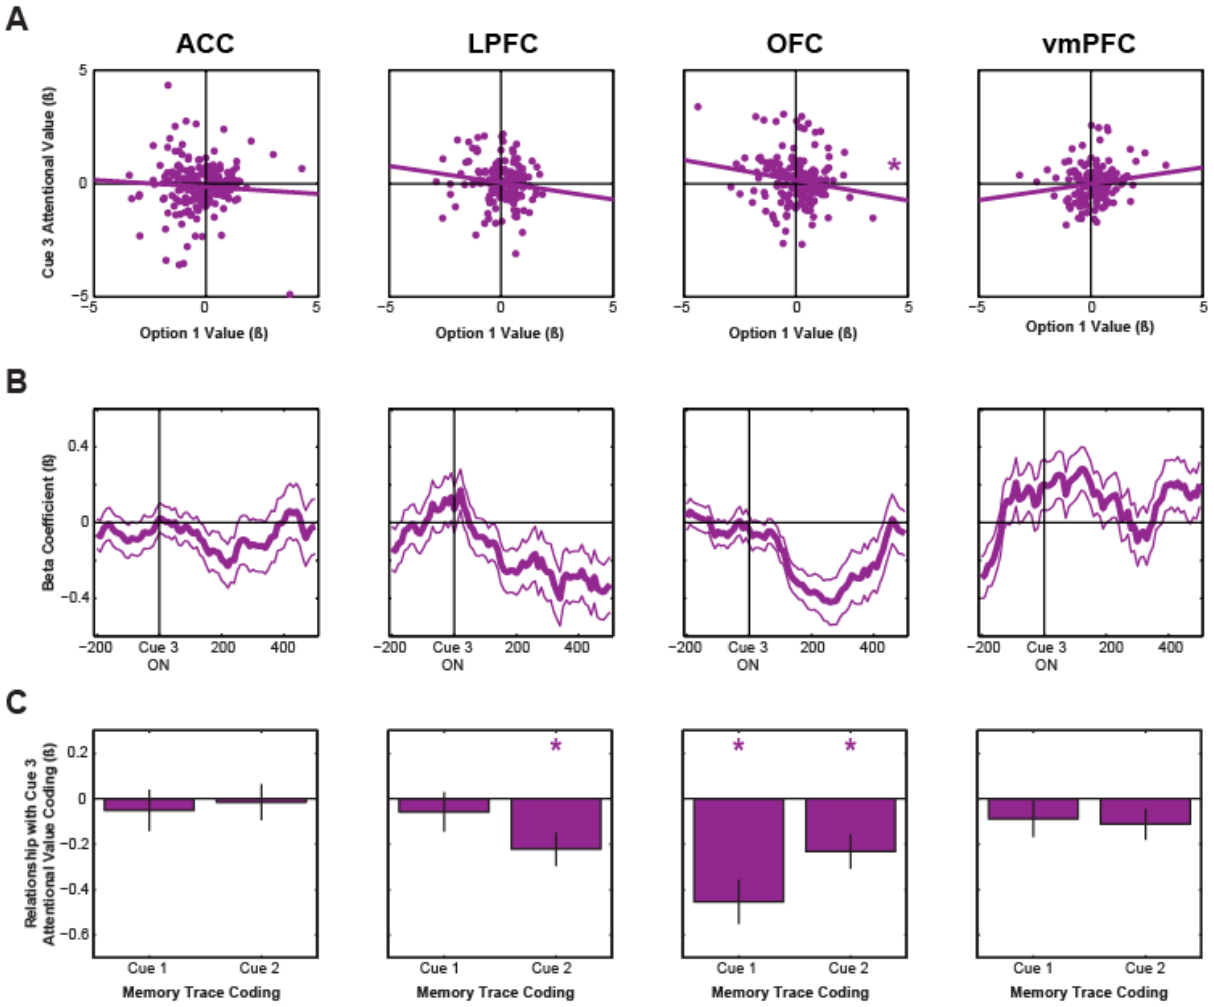

Figure 5.9: Attentional value computations in 'Option' trials at Cue 3. **(A)** A scatter plot of the beta coefficients for Option 1 (i.e. the unattended side) value memory coding at Cue 3 against the attentional value code of Cue 3. All other figure properties are the same as **Figure 5.8**. **(B)** A sliding linear regression of the Option 1 value memory coding at Cue 3 against Cue 3 attentional value coding. Thin lines indicate the SEM. **(C)** The linear regression of Cue 3 attention value coding against the Cues 1 and 2 value memory codes. All other figure properties are the same as **Figure 5.8**.

As expected, a sign flip was also observed in the Cue 1 value memory trace at Cue 3 compared to the original value code at Cue 1 in OFC (**Figure 5.10**) (linear regression; T=-2.02, p<0.05). However, surprisingly this flip was not the case with Cue 2 value memory trace in OFC, which remained strongly positively correlated with the original Cue 2 value

code (**Figure 5.10**) (linear regression; T=3.22, p<0.002). This positive Cue 2 effect was also observed in vmPFC (linear regression; T=2.71, p<0.008).

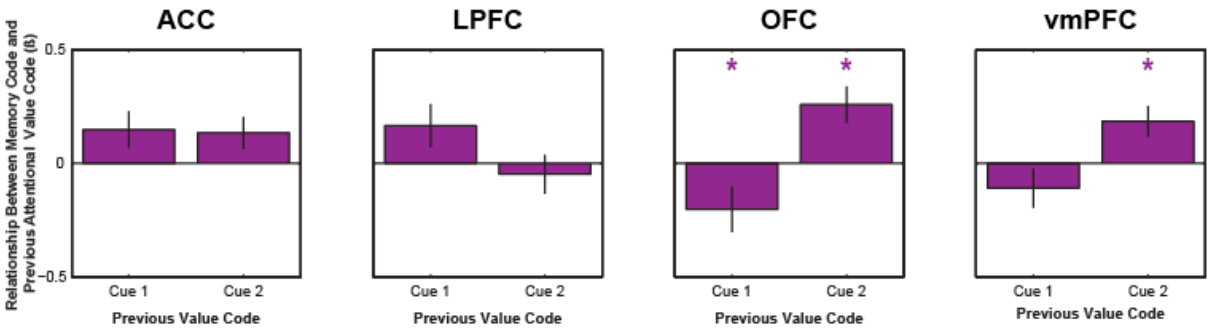

Figure 5.10: The relationship between the value memory trace and previous attentional value codes. The beta coefficients of two separate linear regressions, firstly of Cue 1 attentional value at Cue 1 against Cue 1 value memory code at Cue 3, and secondly of Cue 2 attentional value at Cue 2 against Cue 2 value memory code at Cue 3. All other figure properties are the same as **Figure 5.8**.

The third cue in 'Attribute' trials presents a complication to the analysis of attentional value because subjects were free to choose whether to saccade back to Option 1 (the side of Cue 1) or towards Option 2 (the side of Cue 2). This meant that the memory trace codes for Cues 1 and 2 at Cue 3 sometimes corresponded to the saccaded option and other times to the unsaccaded option. To mitigate this, we added a contrast term into the linear regression of firing rate at Cue 3 (see Methods) which accounted for the memory coding of the value of the cue (i.e. Cue 1 or 2) on the saccaded side of the screen and another term for the memory code for value on the unsaccaded side. We then performed a linear regression of the population Cue 3 value code in 'Attribute' trials against the value codes for the saccaded and unsaccaded Cue and 2 memory traces. For consistency between analyses we again took the mean value code from the 200-400ms post-cue onset time period. As predicted by the attention hypothesis of OFC, there was a significant difference between coefficients for the saccaded and unsaccaded memory traces, with the saccaded trace being positive (although non-significant) and the unsaccaded trace significantly

negative with a significant difference between the two coefficients (**Figure 5.11**) (linear regression; $T_{(Saccaded)}$=1.58, $T_{(Unsaccaded)}$=-3.03, $p_{(Saccaded)}$=0.11, $p_{(Unsaccaded)}$=0.003, linear hypothesis test of coefficients for the memory trace for the saccaded side cue compared to the unsaccaded side , T=10.72, p<0.002).

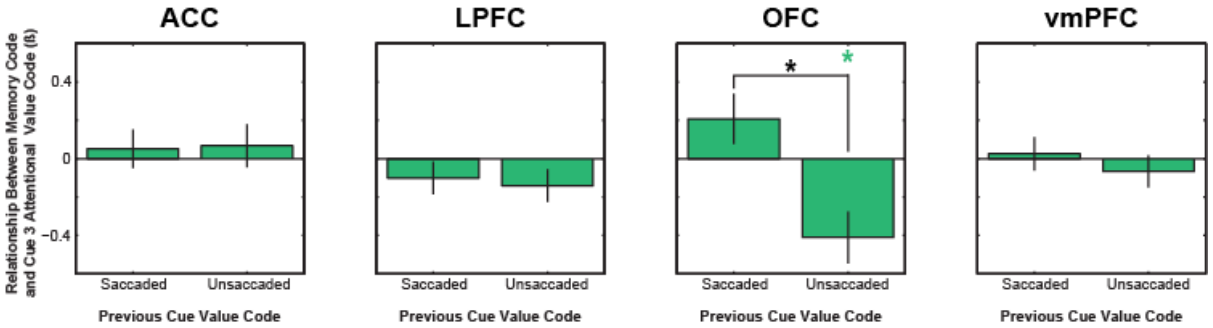

Figure 5.11: Attentional value computations in 'Attribute' trials at Cue 3. The beta coefficients for the regression of the Cue 3 attentional value code against the value memory traces for previous cues on the currently saccaded and unsaccaded sides. All other figure properties are the same as **Figure 5.8**.

To summarise these results, at a population level, the value computations between current and past information in OFC are clearly framed by the location of current attention irrespective of whether the location of the attention is voluntary (as is 'Attribute' trials) or involuntary (as is 'Option' trials).

OFC Value Neurons Represent Choice in an Attentional and Temporal Frame of Reference

Having found that value computations in OFC are framed by attention, one might expect to find OFC neurons that discriminate whether currently attended cues will go on to be chosen or not. We therefore explored how the representation of value and attentional choice (i.e. whether the currently attended side will be chosen) evolved in value coding neurons (**Figure 5.12**). Having found that only OFC neurons represented value with respect to attention, we restricted this element of the analysis solely to the OFC value population. This analysis revealed that there was a clear decrease in attentional value coding over cue

presentations, but at the same time there was a small but clear increase in the representation of attentional choice in the attentional value population (i.e., whether the attentional value neurons discriminate whether subjects will choose the Option of the currently attended Cue), implying an evolution of the coding from value to choice in the frame of reference of attention (**Figure 5.12**, purple square).

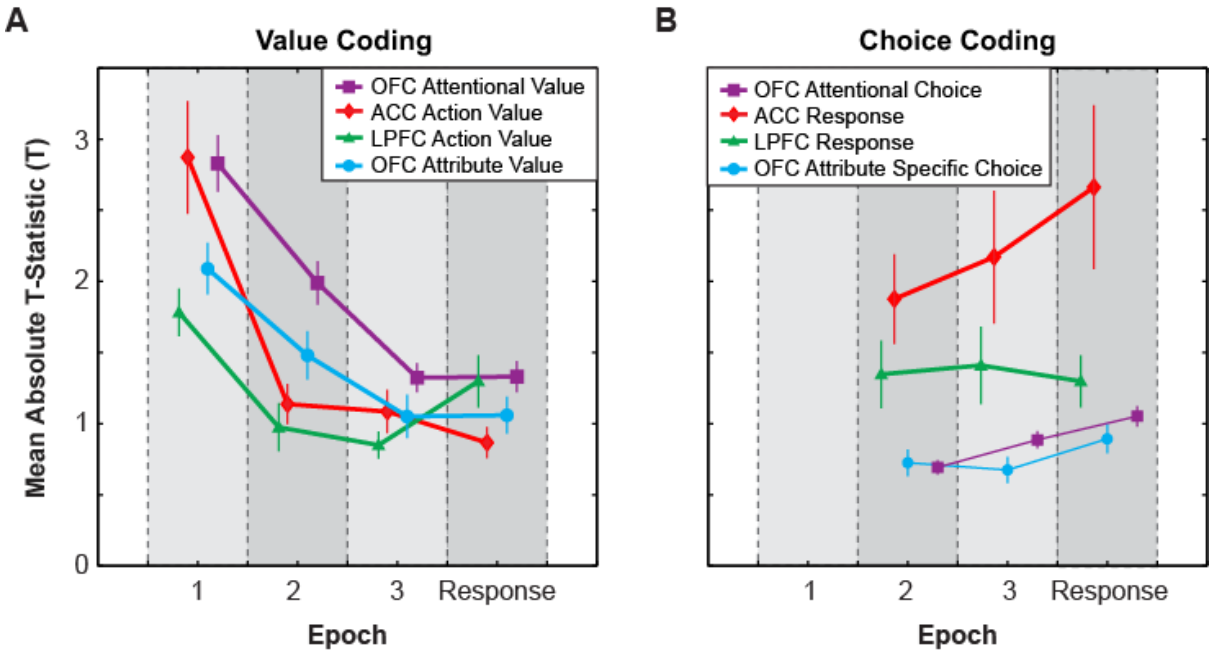

Figure 5.12: The evolution of value and choice coding in PFC. **(A)** The mean absolute T-statistic for OFC attentional value (purple squares), ACC and LPFC action value (red squares, green triangles respectively) and OFC attribute value (blue circles) neurons across the Cue 1, Cue 2, Cue 3 and response epochs. At response the regressors shown are chosen value (because subjects typically choose the currently attended Option), chosen action value, and chosen attribute value respectively **(B)** The mean absolute T-statistics for attentional choice (choose attended vs unattended side) encoding in OFC attentional value neurons (purple squares), response (left-right movement) encoding in ACC and LPFC action value neurons (red diamonds, green triangles respectively) and attribute specific choice (attribute coding when the attended side is chosen) in attribute specific OFC neurons. Vertical lines indicate SEM.

Next we asked whether there was a direct relationship between value coding and attentional choice coding and if so, was this more prevalent in OFC compared to other regions? To test this we performed a projection analysis in which we first identified a population code for value coding in each region by computing the average beta coefficient

for value for each neuron from Cue 1 onset until 500ms post-onset. We then used linear regression to probe the relationship between this attentional value code and the attention choice code at Cue 2, Cue 3 and response across time. This analysis found that although there were few significant differences between regions at Cue 2, at Cue 3 and response the projection correlation was significantly stronger in the OFC than all other areas, and peaked 200ms prior to the response (**Figure 5.13**) (linear hypothesis test of the beta coefficients for OFC against other regions with Bonferroni correction, $p<0.008$). This provides further evidence for the conclusion that value neurons in OFC have a strong tendency to go on to represent attentional choice (i.e., whether the currently attended stimulus of an Option will be chosen).

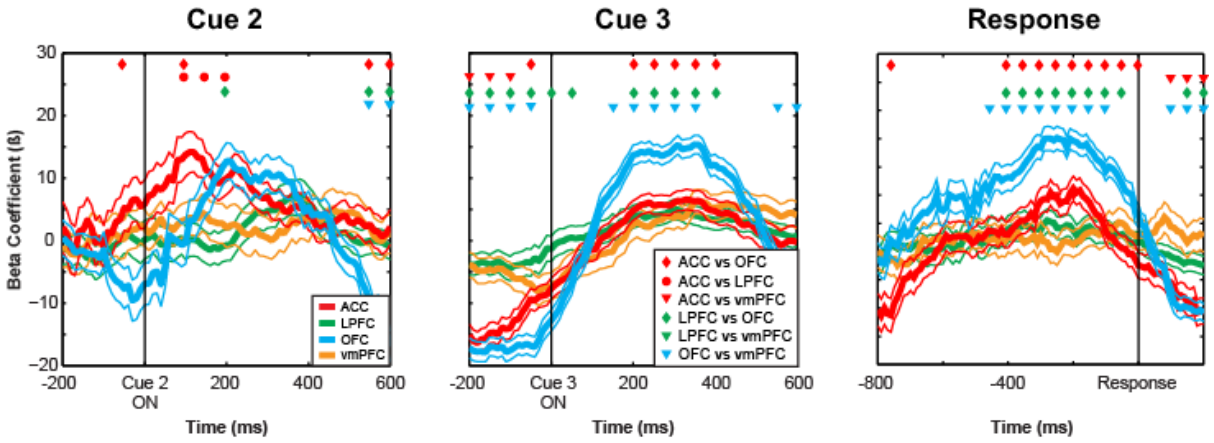

Figure 5.13: The evolution of attentional choice coding in attentional value neurons. The sliding regression of the Cue 1 attentional value code in each region against the attentional choice code (i.e. whether the currently attended stimulus of an option will be chosen or unchosen) at Cue 2 (left), Cue 3 (middle) and Response (right). Symbols indicate coefficients between that are significantly different (linear hypothesis test with Bonferroni correction). All other figure properties are the same as **Figure 5.6**.

Do computations in OFC value neurons truly evolve over time? In order to answer this question we performed a multiple linear regression of the Cue 1 OFC population attentional value code against both the value and choice codes estimated at Cues 2 and 3. Firstly, the value code at Cue 3 was found to be significantly smaller than that of Cue 2 (linear hypothesis test of coefficients for Cue 2 value against Cue 3 value, $T=4.50$, $p<0.04$).

Furthermore, the Cue 3 choice code coefficient was significantly greater than that of Cue 2 (linear hypothesis test of coefficients for Cue 2 choice against Cue 3 choice, $T=6.26$, $p<0.02$). This result demonstrates that early on in the trial the OFC value neurons tended to compute attentional value but as the trial progressed their computations shifted towards attentional choice and that this evolution was proportional to the extent to which neurons initially encoded value. The implication of the latter is that strong value coding neurons went on to become strong attentional choice coders.

Different Action and Attribute Value Codes are Equally Represented in PFC

Having surmised that ACC and OFC represent value in action and attribute frames of reference we investigated whether specific action or attribute values were comparatively over-represented by neurons within these areas. To do this we could not use the contrast terms because due to the fact neurons can (and do) encode value either positively or negatively, the sign of the contrast terms did not specify which action or attribute a neuron preferred. Therefore we computed attribute preference by subtracting the mean absolute T-statistic for left value coding against the same for right value coding. Neurons with a positive difference in this measure preferred left and those with negative differences preferred right. We also repeated this for probability versus magnitude. There was no significant difference between the number of neurons that preferred left or right value in ACC (binomial test, $p>0.05$). There was also no significant difference when considering the proportion of magnitude and probability value preferring neurons in OFC (**Figure 5.5B**) (binomial test, $p>0.05$). There was also no significant skew in the strength (i.e. the mean absolute T-statistic difference) of action or attribute value encoding across all ACC or OFC neurons respectively (one sample T-test, $p>0.05$). When the above tests were repeated specifically for neurons that were considered significant for action value coding in ACC and attribute value coding in OFC all results were still non-significant (all $p>0.05$).

Finally, it is possible that ACC neurons encode action value ipsilaterally or contralaterally to the recorded hemisphere. To test this we examine the number of ipsi/contralateral action value preferring neurons. This analysis revealed that there was a weak tendency to encode contralateral action value *specifically in action value coding neurons* in ACC (binomial test, $p<0.05$).

ACC Neurons Encode Action Related Computations

Are neurons that encode action value at Cue 1 more likely to perform action related computations later on in the trial? We attempted to answer this question on two levels. First, we looked at how strongly Cue 1 action value neurons encoded action value in other epochs (**Figure 5.12)**. We used our first level analysis of action value coding at Cue 1 to restrict this analysis solely to ACC and LPFC neurons. ACC action value neurons did represent action value at Cue 2 but the representation diminished substantially over the trial to the point where the chosen action value was essentially not represented at response (**Figure 5.12A**). In contrast the encoding of the response in the ACC Cue 1 action value population increased substantially over the trial (**Figure 5.12B**). This result suggested that ACC action value neurons had a propensity to evolve their neuronal coding to represent final response as the trial progressed.

In LPFC a different pattern was observed. Although action value coding diminished progressively across Cues 1-3, at response there was a noticeable increase in coding of chosen action value, exceeding that of the ACC action value population (**Figure 5.12A**). Furthermore, unlike ACC, the LPFC action value neurons did not increase their representation of response as the trial progressed (**Figure 5.12B**). This may indicate that the LPFC action value population serves a separate function to the ACC population, and may represent value in a spatial reference frame (left/right, up/down) and the chosen spatial

value at response, rather than computing the value of the action/Option and the final action like ACC neurons.

Next we performed a more advanced analysis which asked whether the strength of action value coding was related to the strength of later action computations. This method had the added advantage of allowing us to visualise the temporal dynamics of further computations of action value which neurons performed. As with the OFC attentional value population, we used linear regression to probe the relationship between the mean Cue 1 action value code and the action value and response codes at Cues 2, 3 and response (**Figure 5.14**). This analysis found a positive correlation between Cue 1 action value and Cue 2 action value coding only in ACC and LPFC (**Figure 5.14A**) (linear hypothesis test of coefficients for ACC/LPFC against OFC or vmPFC with Bonferroni correction, $p<0.008$). Surprisingly, however, when the same analysis was applied to Cue 3, the Cue 1 action value code in both ACC and LPFC *did not* correlate with Cue 3 action value (**Figure 5.14A**), showing that action value coding diminished in both regions over time (linear hypothesis test of coefficients for ACC/LPFC against OFC or vmPFC with Bonferroni correction, $p>0.008$).

However, as early as Cue 2, ACC Cue 1 action value coding was seen to significantly predict final response coding more than any other region (**Figure 5.14B**) (linear hypothesis test of coefficients for ACC against all other regions with Bonferroni correction, $p<0.008$). This correlation was maintained across the Cue 3 and response epochs and increased with proximity to the response (**Figure 5.14B**) (linear hypothesis test of coefficients for ACC against all other regions with Bonferroni correction, $p<0.008$). In contrast to ACC, LPFC action value neurons did not correlate with response during other cues with the exception of immediately prior to the response, although this effect may be extremely weak because it was only significantly different from vmPFC (**Figure 5.14B**) (linear hypothesis test of coefficients for LPFC against vmPFC with Bonferroni correction, $p<0.008$).

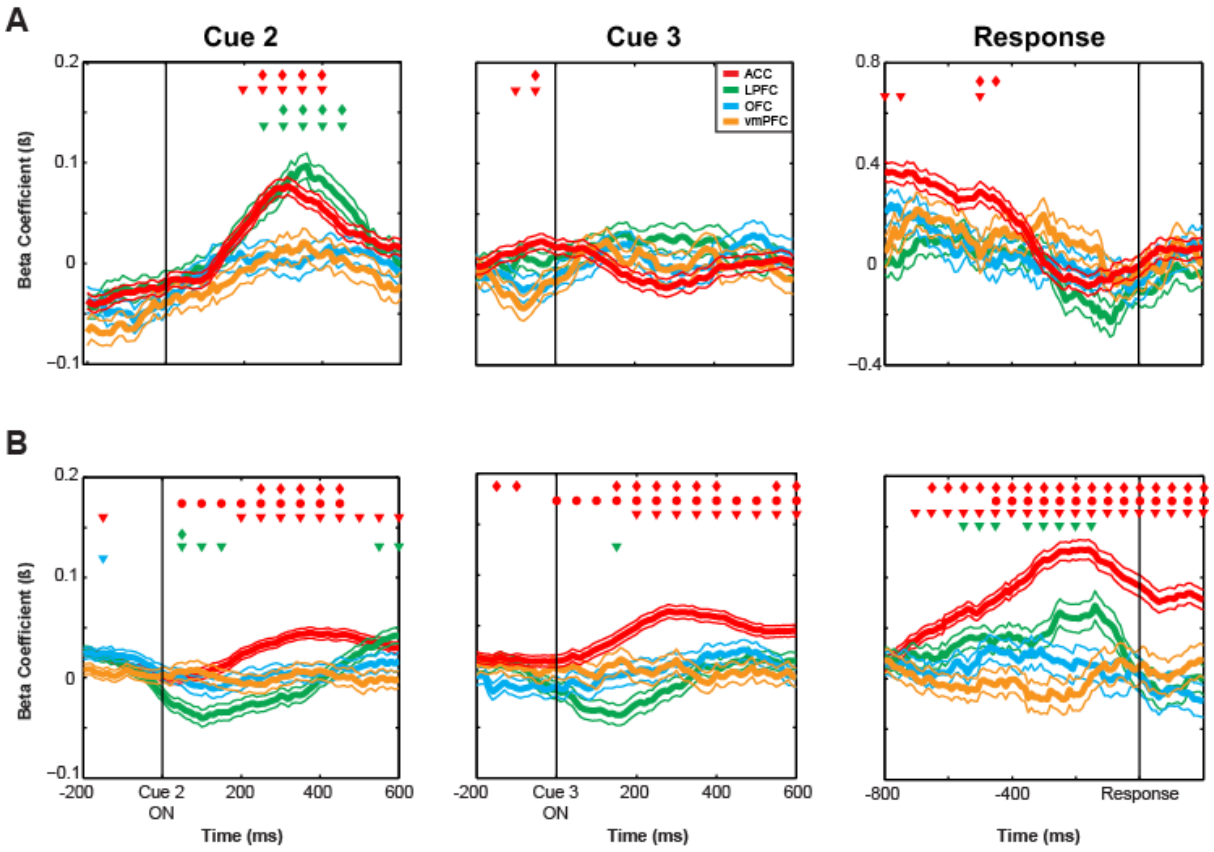

Figure 5.14: The evolution of action value coding to response. **(A)** The linear regression of the Cue 1 action value against action value coding at Cue 2 and Cue 3 and chosen action value coding at response for each area. **(B)** The linear regression of Cue 1 action value against response coding at Cue 2, 3 and response. All other figure properties are the same as **Figure 5.13**.

To summarise, these results show that the firing patterns of certain populations of ACC neurons evolve over the course of the trial from the representation of action value to that of the final response of the subject. In contrast LPFC neurons cease coding action value as the trial progresses but also do not go on to represent final chosen action.

OFC Neurons Encode Attribute Computations

Having identified the evolution of action value coding into final response in ACC the natural follow up question was to ask whether a similar process can be identified in attribute value coding OFC neurons. Again, by looking at the strength of attribute value coding across

cues in this predefined OFC population, we were able to show that the strength of the attribute representation decreased markedly over the course of a trial (**Figure 5.12A**).

One of the constraints of the behavioural task is that although there is a clear index of subject choice in the frame of reference of actions (i.e. the eventual joystick movement); there is no similar index for choice in the frame of reference of attribute. Therefore in order to study an 'attribute framed' choice signal similar (but not analogous) to the ACC response signal, we attempted to find a neural correlate of attentional choice that also discriminated the attribute type of the cue (i.e. a signal that discriminated attribute type of the currently attended cue when it was on the side of the chosen option). We refer to this signal as an 'attribute specific choice' signal. For Cue 2 Subjects M and F chose the attended side on 48.2% and 48.1% of occasions respectively. At Cue 3 it was 58.3% and 55.6% of occasions respectively. On examining the representation of the 'attribute specific choice' regressor, there appeared to be a small increase in its representation within OFC as the response neared (**Figure 5.12B**).

We used the same projection analysis described above to test the evolution of the OFC the attribute value code. We found that OFC neurons which encoded attribute value at Cue 1 also encoded attribute value at Cue 2 significantly more than all other regions (**Figure 5.15A**) (linear hypothesis test of the beta coefficients for the OFC attribute code against the same coefficients for other regions with Bonferroni correction, $p<0.008$). However, unlike the ACC action value code, the OFC attribute value code was also clearly maintained at Cue 3 (**Figure 5.15A**) (linear hypothesis test of the beta coefficients for the OFC attribute code against the same coefficients for other regions with Bonferroni correction, $p<0.008$). Finally, when we analysed the Cue 1 attribute value code with respect to the response epoch, we found that the Cue 1 attribute value code was strongly positively correlated with the attribute specific encoding of the chosen option, implying a maintenance of the attribute value code over the course of the trial (**Figure 5.15A**) (linear hypothesis test of the beta coefficients for

the OFC attribute code against the same coefficients for other regions with Bonferroni correction, p<0.008).

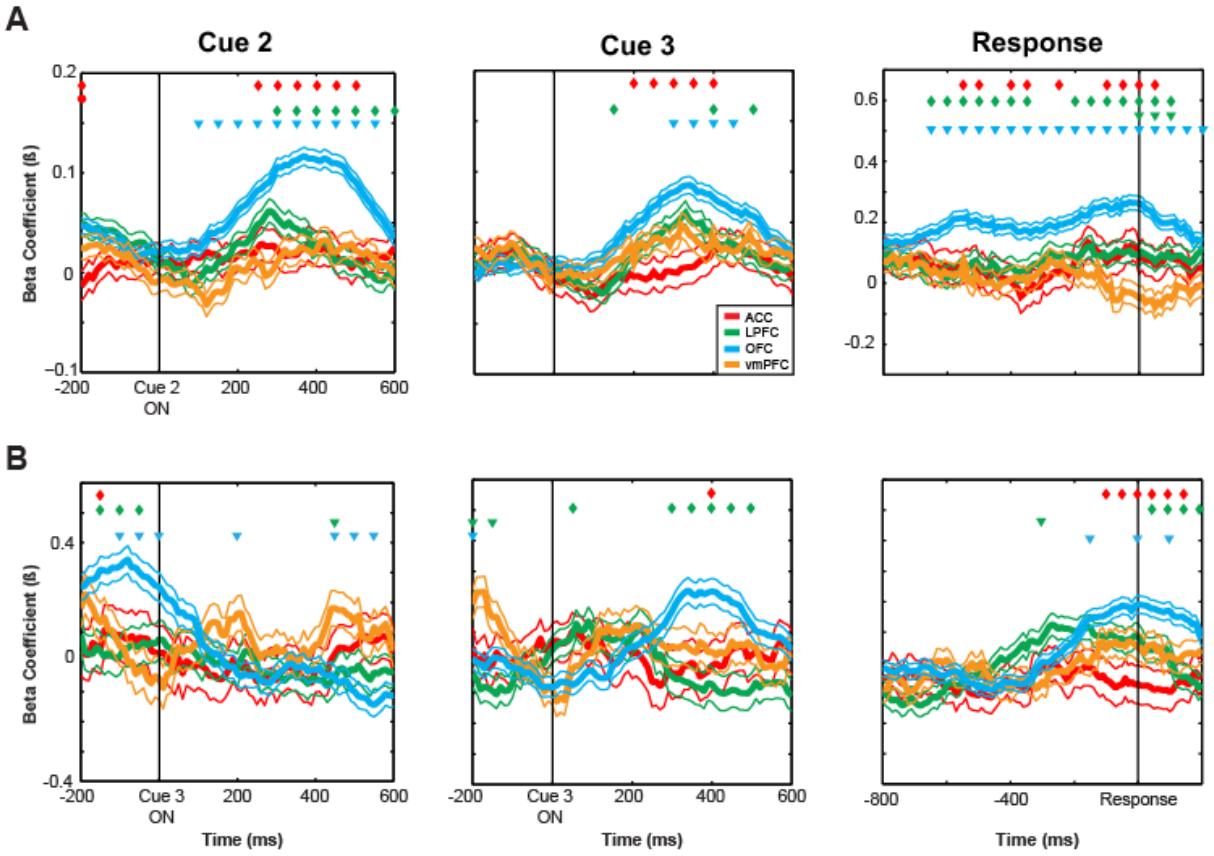

Figure 5.15: The evolution of attribute value coding to 'attribute specific choice'. **(A)** The linear regression of the mean Cue 1 attribute value code against attribute value coding at Cue 2 and Cue 3 and chosen attribute value coding at response for each area. **(B)** The linear regression of Cue 1 attribute value coding against 'attribute specific choice' coding (i.e. discriminating attribute type of the currently attended cue when it was on the side of the chosen Option) at Cue 2, 3 and response. Note that the left most figure comes from data aligned to Cue 3 onset but pertains to attribute specific choice coding *at Cue 2*. All other figure properties are the same as **Figure 5.13**.

When we correlated the Cue 1 attribute code against the 'attribute specific choice' code at Cue 2, there was no significant difference between areas (linear hypothesis test of the beta coefficients for the OFC 'attribute specific choose Cue 2 side' code against the same coefficients for other regions with Bonferroni correction, p>0.008). However, when this same 'attribute specific choose Cue 2' regressor was time locked to Cue 3, the correlation was seen to be much stronger peaking just prior to the onset of Cue 3 (linear hypothesis test

of the beta coefficients for the OFC ‘attribute specific choose Cue 2 side’ code against the same coefficients for other regions with Bonferroni correction, $p<0.008$) (left plot of **Figure 5.15B**). In the same regression, an ‘attribute specific choose Cue 3’ signal was also seen from the OFC Cue 1 attribute value neurons approximately 400ms post-Cue 3 onset (linear hypothesis test of the beta coefficients for the OFC ‘attribute specific choose Cue 3 side’ code against the same coefficients for other regions with Bonferroni correction, $p<0.008$) (**Figure 5.15B**). Next, we correlated the OFC Cue 1 attribute value code against the ‘attribute specific choose Cue 3’ signal now referenced to the attribute identity of the last cue viewed before the response if that attended cue was of the side that was chosen. This revealed that the ‘attribute specific choice’ signal in the OFC attribute neurons peaked around the response time, which was both positive and significantly different from the other regions (**Figure 5.15B**).

Finally, in order to further examine the evolution of the attribute value code over the trial, we performed a linear regression of the Cue 1 OFC attribute value code against the OFC attribute value and ‘attribute specific choice’ codes at Cues 2 and 3. Unsurprisingly, there was no significant change in the attribute value codes between the two cues but there was a small but significant increase in the ‘attribute specific choice’ coefficients (**Figure 5.12**; linear hypothesis test of coefficients for Cue 2 attribute value against Cue 3 attribute value, $T=3.07$, $p>0.05$; linear hypothesis test of coefficients for ‘attribute specific choose Cue 2 side’ against ‘attribute specific choose Cue 3 side’, $T=4.04$, $p<0.04$).

These analyses suggest that attribute specific OFC neurons identified at Cue 1 continue to compute attribute specific value at subsequent cues, and also compute an ‘attribute specific choice signal’ at each cue which peaking immediately prior to response. Although we cannot conclude that the latter is an index for subjects choosing based upon particular attributes, these results reveal that attribute value neurons in OFC can track choice in an attribute frame of reference.

Frame of Reference Neurons Also Drive Information Gathering Behaviour

One of the unique aspects of the behavioural paradigm is that after viewing two cues, subjects were free to *choose where* to gather a third piece of information, or they could make their final choice via a joystick movement. One might therefore expect that PFC regions may guide subjects' information gathering behaviour through computations based on value. Might these putative computations also be in different frames of reference? If these signals do indeed drive information gathering behaviour, then it is necessary for these computations to occur *before* the third saccade. To explore this idea, we used linear regression of firing rates at Cue 1 to ask whether any neurons could decode future information gathering strategy at the third cue. We defined the information gathering strategies by attention, action and attribute. Therefore, in 'Option' trials, we examined whether neurons discriminated whether subjects made a third saccade to the *same* (attended) attribute as Cue 2 (i.e. horizontal in direction) or towards the other (unattended) attribute as Cue 2 (i.e. diagonal direction), reflecting coding in the attentional frame of reference. It is vital to note that this can be considered as a saccade contingent on the *currently attended attribute* but not attribute specific. Also in 'Option' trials, we considered whether the direction of the third saccade was towards a probability stimulus or a magnitude stimulus (i.e. in the attribute frame of reference). In 'Attribute' trials, we examined whether neurons encoded third saccades towards the same side as Cue 2 or away from the Cue 2 side (i.e. vertical versus diagonal saccades respectively) in order to test the attentional frame of reference. Finally, on 'Attribute' trials, we also measured whether the third saccade ended on the left or right hand side of the screen (irrespective of where it started) to examine the action frame of reference. Importantly, because saccade direction is known to be influenced by value (see Chapter 4), we accounted for this by adding co-regressors which could (through contrasts) account for attentional, action and attribute value coding. We repeated this analysis at Cue 2 and Cue 3 (i.e. after third saccade deployment).

On examining the proportion of neurons in each PFC region that encoded future saccades in an attentional frame of reference, we found that even at Cue 1, a significant proportion of LPFC neurons did so in both trial types (**Figure 5.16,** left panel) (binomial test, $p<0.05$). In 'Attribute' trials, this proportion in LPFC was significantly greater than in OFC and vmPFC neurons (Chi^2 test, $p<0.05$). Furthermore, at Cue 2, the proportion of LPFC neurons grew as high as 26% in 'Attribute' trials and 24% in 'Option' trials, although significant populations of OFC and ACC neurons were also observed to encode the future saccade in the attentional reference frame in both trial types (**Figure 5.16,** middle panel) (binomial test, all $p<0.05$). With respect to the attentional saccade in 'Attribute' trials, LPFC coding was much greater than all other areas (pairwise Chi^2 test, $p<0.05$). Even after the third saccade deployment, a significant proportion of ACC, LPFC and OFC still encoded the direction of the third saccade in the attentional reference frame in both trials types (**Figure 5.16,** right panel) (binomial test, all $p<0.05$).

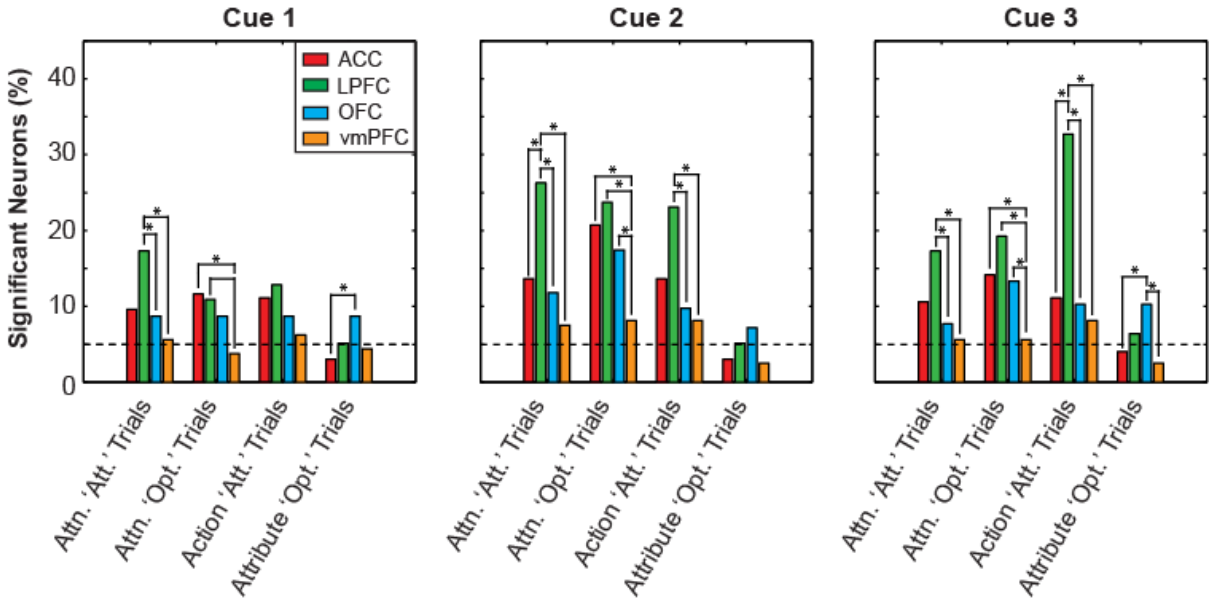

Figure 5.16: The encoding of third saccade behaviour across cues. The proportion of neurons that encode attentional saccades in 'Attribute' trials (Attn. 'Att' Trials), attentional saccades in 'Option' trials (Attn. 'Opt.' Trials), action saccades in 'Attribute' trials, attribute saccades in 'Option' trials. All of figure properties are identical to **Figure 5.4A**.

Having found that neurons in PFC predict information gathering strategy framed by current attention, the obvious follow up question was to ask whether these were the same

neurons that computed attentional value during the trial (as described earlier). To do this we performed a population projection of mean Cue 1 attentional value coding against the mean attentional saccade population code (for each trial type) at each cue epoch. In 'Attribute' trials, we observed a significant and specific positive relationship between value coding and attentional saccade coding in OFC at the second cue (**Figure 5.17A**) (linear regression, T= 3.97, $p<0.8x10^{-5}$, linear hypothesis test of the beta coefficient for OFC against each other regions with Bonferroni correction, all $p<0.008$). This implies that OFC attentional value neurons fire more when future information gathering is towards the side of current attention compared to when information is gathered away from the currently attended option. This correlation was not present at either Cue 1 or Cue 3, implying that this signal was a phasic computation in OFC attentional value neurons immediately prior to the volitional information gathering saccade.

We next performed a linear regression of the Cue 1 attentional value coding against attentional saccade coding in 'Option' trials (**Figure 5.17B**). Again, we found a singularly significant relationship in OFC which was not significantly different from any other region (linear regression, T= -2.57, $p<0.02$, linear hypothesis test of the beta coefficient for OFC against other regions with Bonferroni correction, $p>0.008$). Again, like in 'Attribute' trials, this relationship was positive. Therefore the positive attentional value coding neurons at Cue 1 respond most when subjects make horizontal saccades, making comparisons within the currently attended attribute. Again, repetition of this analysis at Cue 1 and 3 found no significant relationship (linear regression, $p>0.05$).

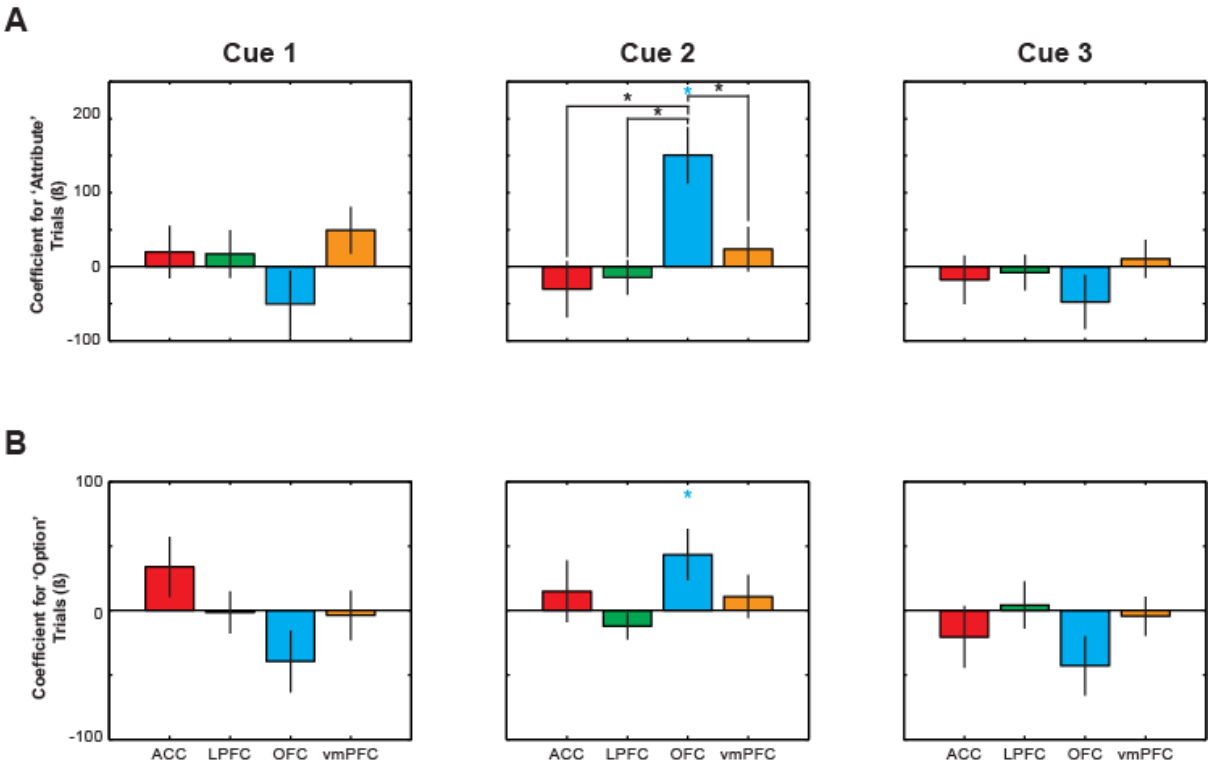

Figure 5.17: OFC value neurons predict future information gathering in the attention frame of reference. **(A)** The linear regression of mean Cue 1 value coding against the coding of vertical versus diagonal Cue 3 saccades at each cue in 'Attribute' trials. **(B)** The linear regression of mean Cue 1 value coding against the coding of horizontal versus diagonal Cue 3 saccades at each cue in 'Option' trials. Coloured asterisks indicate regions that have significant beta coefficients. Black asterisks indicate areas that are significantly different from each other (linear hypothesis test with Bonferroni correction). Vertical lines indicate the SEM.

Taken together, these results suggest that LPFC neurons strongly represent information gathering strategies framed by attention immediately prior to saccade deployment. However, OFC attentional value neurons multiplex this information gathering signal with value and choice computations peaking immediately prior to the saccade.

We next considered whether future saccades were also encoded in an action frame of reference (i.e. action saccades). We performed this analysis exclusively in 'Attribute' trials where subjects were free to choose which option they wanted to obtain information about. Also, unlike in 'Option' trials, the position of the first and second cues in 'Attribute' trials did not predict the direction of the third saccade. Upon examining the proportion of significant neurons that encoded the direction of these forthcoming action saccades at Cue 1, we found

that only ACC and LPFC contained a significant proportion of neurons which encoded action saccades (**Figure 5.16**, left panel) (binomial test, both $p<0.05$). At Cue 2, the proportion of ACC and LPFC neurons coding for the action saccade at Cue 3 increased to approximately 15% and 23% of neurons in these areas, respectively (**Figure 5.16**, middle panel). LPFC neurons encoded action saccades significantly more than all other areas (pairwise Chi^2 test, all $p<0.05$). The proportion of LPFC neurons encoding action saccades increased further at Cue 3 but diminished in ACC, such that LPFC encoded the direction of the completed action saccade significantly more than all other regions (**Figure 5.16**, right panel) (pairwise Chi^2 test, $p<0.05$).

We next considered the pattern of this action saccade coding by examining the number of significant neurons (for all neurons) across time at all three cue epochs (**Figure 5.18**). This analysis showed that at Cue 1, most of the ACC saccade direction coding occurred around the presentation of the cue, whereas the LPFC coding peaked approximately 300-400ms after cue presentation (**Figure 5.18**, left panel). After Cue 2 presentation, the saccade direction coding in LPFC rapidly increased and peaked around 100ms post-cue presentation (**Figure 5.18**, middle panel), whereas the coding in the ACC population increased slowly and generally plateaued soon after the second cue. Both ACC and LPFC saccade direction coding was maintained in the time period immediately prior to the third saccade (**Figure 5.18**, right panel). While the ACC representation slowly decreased after the saccade, the LPFC representation increased, peaking for a second time approximately 200ms post-cue onset (**Figure 5.18**, right panel).

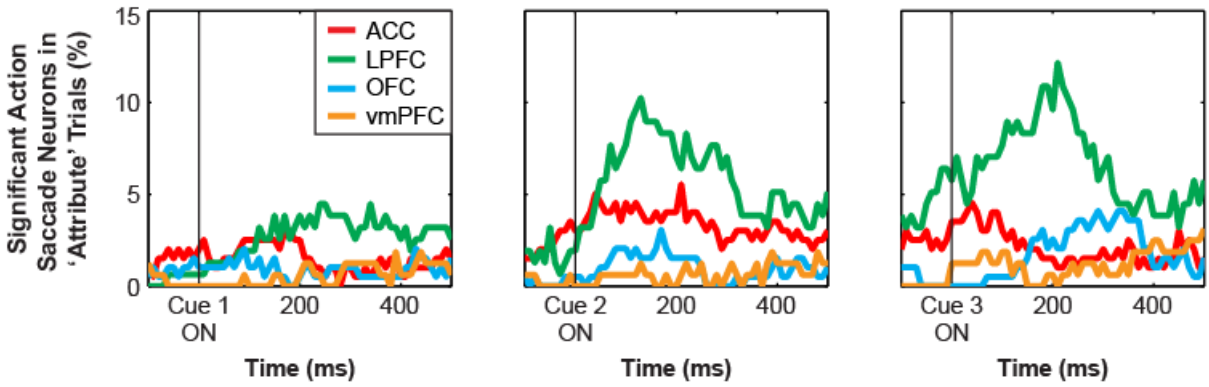

Figure 5.18: The proportion of action saccade selective neurons at each cue epoch.

Are the neurons that encode action saccades the same neurons that encode action value in ACC and LPFC? To test this we performed a population projection of the Cue 1 action value code in each area against action saccade code at each cue on 'Attribute' trials (**Figure 5.19A**). Upon doing this analysis we found that ACC Cue 1 action value coding was significantly correlated with action saccade coding at Cue 2 (linear regression, $p<5x10^{-8}$). Even though there was no significant relationship in any other region, the ACC effect was only significantly greater than vmPFC (linear hypothesis test of the beta coefficient for ACC against that of vmPFC with Bonferroni correction, $p<0.008$). Furthermore, this relationship between action value coding and action saccades in ACC was not present either at Cue 1 or Cue 3 (linear regression, $p>0.05$).

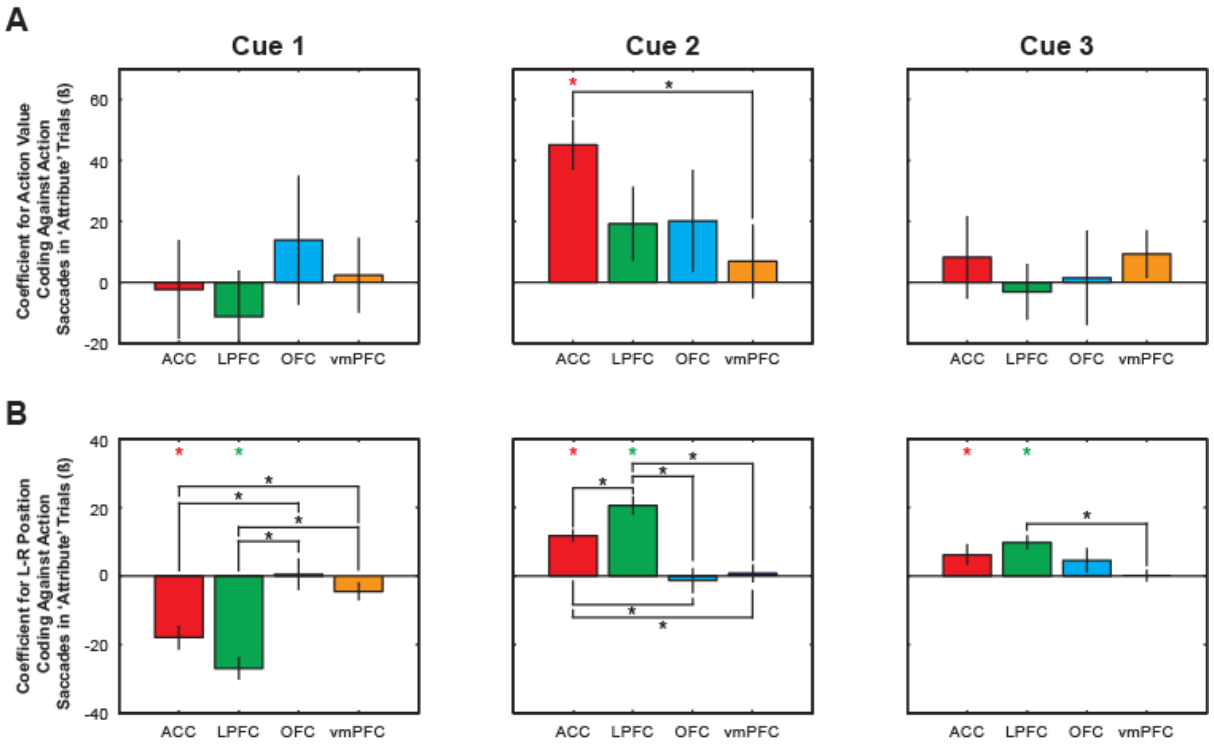

Figure 5.19: Action and position neurons predict future information gathering strategies. **(A)** The beta coefficients of a linear regression of Cue 1 action value coding against the third saccade direction (left-right) in 'Attribute' trials as estimated at each cue. **(B)** The beta coefficients of a linear regression of Cue 1 left-right cue position coding against the third saccade direction (left-right) in 'Attribute' trials as estimated at each cue. All other figure properties are the same as **Figure 5.17**.

Figure 5.20 shows an example of an ACC action value neuron which also encodes action saccades. This neuron responds only to the value of cues presented on the left side of the screen at Cue 1 within approximately 200ms and not to those on the right side (**Figure 5.20A**). However, when the neuron's firing rates are reorganised by the combinations of possible third saccades and final choice, the neuron encodes the direction of the future saccade at approximately 450ms post-Cue 1 onset (**Figure 5.20B**). This is then maintained during Cue 2 presentation, but stops encoding the future saccade well before Cue 3 presentation (and hence before the volitional third saccade).

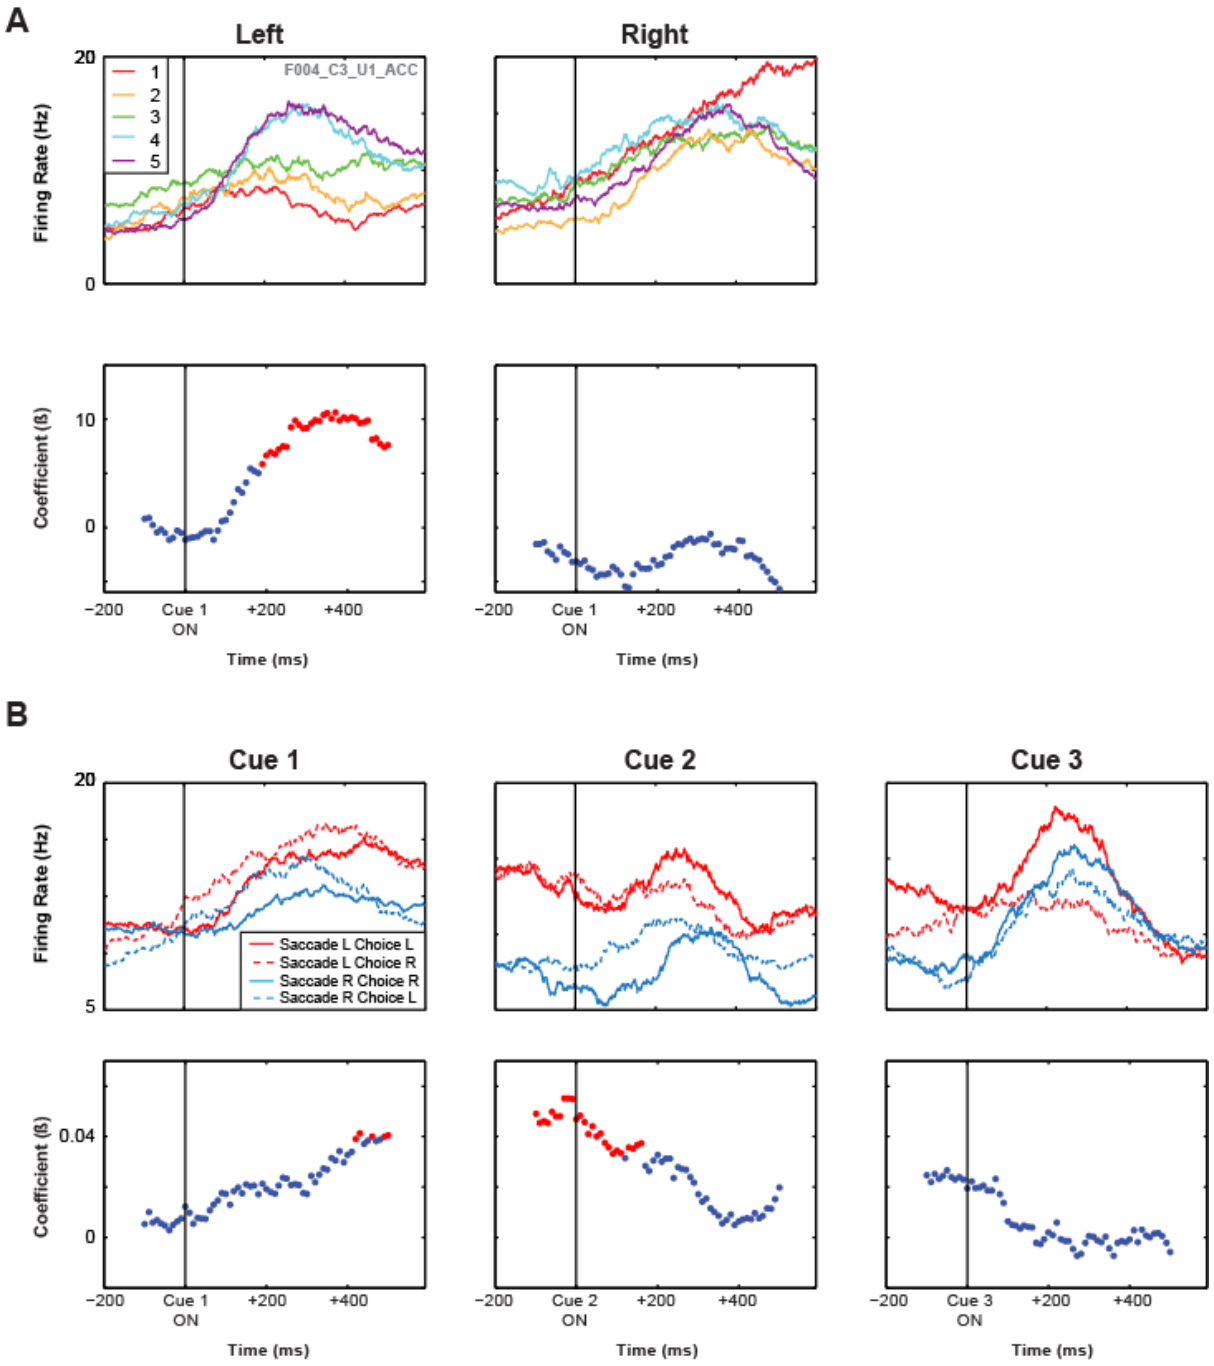

Figure 5.20: An example of an ACC neuron which encodes both action value and action Saccades. **(A)** The neuronal response of this neuron to value split by action at Cue 1 (upper panel) and the corresponding beta coefficients for left and right value coding (lower panel). **(B)** The same neurons firing rate at Cue 1, 2 and 3 separated by the combination of third cue saccades and final responses (upper panel) and the sliding beta coefficients for action saccade (left-right) coding across time (lower panel). All other figure properties the same as **Figure 5.4B**.

In order to better understand the temporal profile of this action saccade computation, we performed a sliding projection analysis of the mean ACC Cue 1 action value code against the action saccade code across time at all three cues (**Figure 5.21**, solid line). The positive relationship between action value and action saccade strategy began approximately 500ms after Cue 1 presentation and was then sustained throughout Cue 2 presentation. After the third saccade this relationship slowly decreased and was around zero after approximately 300ms post-Cue 3.

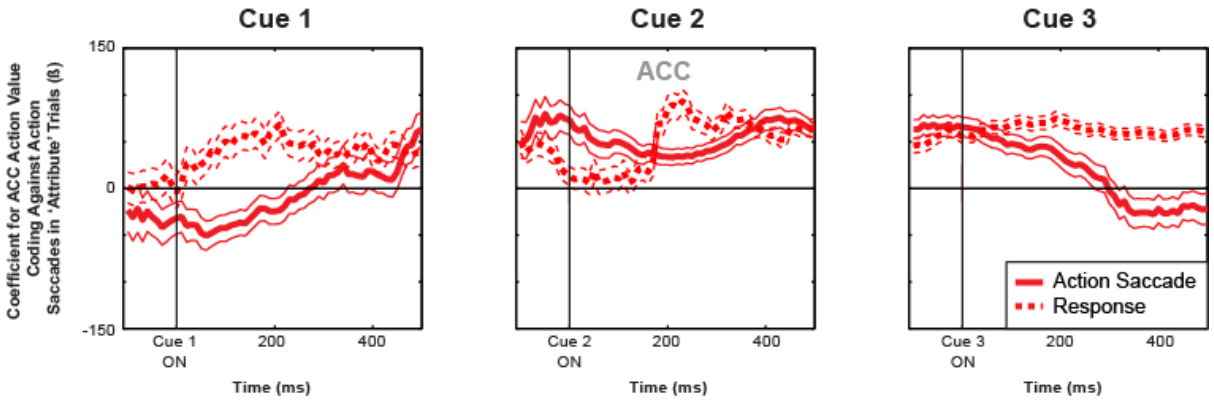

Figure 5.21: Action saccades are mutually represented with final response in ACC action value neurons. A sliding linear regression of Cue 1 ACC action value coding against 'Attribute' trials action saccade (solid line) and final response (dashed line) coding at each cue. Thin lines indicate the SEM.

Subjects have a tendency to saccade towards what they eventually choose in 'Attribute' trials; therefore were the action saccade codes and final response codes simultaneously represented during the trial or was there an evolution of coding from that of the saccade to the response after Cue 3? To qualitatively examine this we also performed a projection of Cue 1 action value coding in ACC against final response coding specifically in 'Attribute' trials as estimated from the regression models set out in **Table 5.8** (**Figure 5.21**, dashed line). At Cue 2 it was clear that both response and action saccades were represented by the action value population and after saccade deployment at Cue 3, the saccade code died away while the response code was maintained.

Surprisingly, the Cue 1 LPFC action value code did not correlate with action saccade coding (**Figure 5.19A**, linear regression, $p>0.05$). Is it therefore possible that this LPFC action saccade population (**Figure 5.16, 5.18**) constitutes a population of neurons that perform a different computation early on in the trial? An obvious candidate computation is spatial position coding. Therefore we performed a projection analysis of Cue 1 left-right spatial position coding against the action saccade code at each cue **(Figure 5.19B)**. At Cue 1 both ACC and LPFC populations exhibited a significant negative relationship with the action saccade direction (linear regression, both $T<-5.21$, both $p<4x10^{-7}$). Both were also significantly different from OFC and vmPFC (linear hypothesis tests of beta coefficients for ACC and LPFC individually tested against OFC and vmPFC with Bonferroni correction, $p<0.008$). By Cue 2 this correlation became strongly positive in ACC and LPFC, indicating that neurons that preferred cues located on the left also fired more when subjects would make future saccades towards the left side of the screen. This correlation in LPFC was significantly stronger than all other regions (linear regression, both $T>6.57$, both $p<1x10^{-10}$, linear hypothesis test of beta coefficients for LPFC against those for each other region with Bonferroni correction, $p<0.008$) while the correlation in ACC was significantly stronger than OFC and vmPFC (linear hypothesis test of beta coefficients for ACC against those for OFC and vmPFC individually with Bonferroni correction, $p<0.008$). The significant positive relationships in ACC and LPFC were then maintained throughout the Cue 3 period, but by this time only LPFC was significantly stronger than any region (linear regression, both $T>2.02$, both $p<0.05$ linear hypothesis test of beta coefficients for LPFC against those for vmPFC with Bonferroni correction, $p<0.008$).

We also performed a sliding projection analysis using the left-right position selective code to examine the time course of the ACC and LPFC action saccade computations (**Figure 5.22**). This revealed that at Cue 1, the negative correlation between these two regressors in the ACC population was present even before cue presentation (although the subjects already knew the position of the first cue at this point) (**Figure 5.22A**), whereas in

LPFC the negative correlation began about 50-100ms post-cue presentation (**Figure 5.22B**). It is unclear, however, why this signal would be negative. Both areas showed a flip in the sign of the projection coefficient around Cue 2 onset, but again the ACC effect was earlier than in LPFC. The positive relationship was then clearly maintained until approximately 300ms post-Cue 3 onset in both areas. Interestingly, unlike the action value code (**Figure 5.21**), neurons that encoded the spatial L-R code at Cue 1 did not represent the final response over each cue, implying that these neurons did not represent final choice (**Figure 5.21**). These results indicate that non-value neurons in PFC are also able to drive information gathering behaviour, in this case, in a spatial frame of reference.

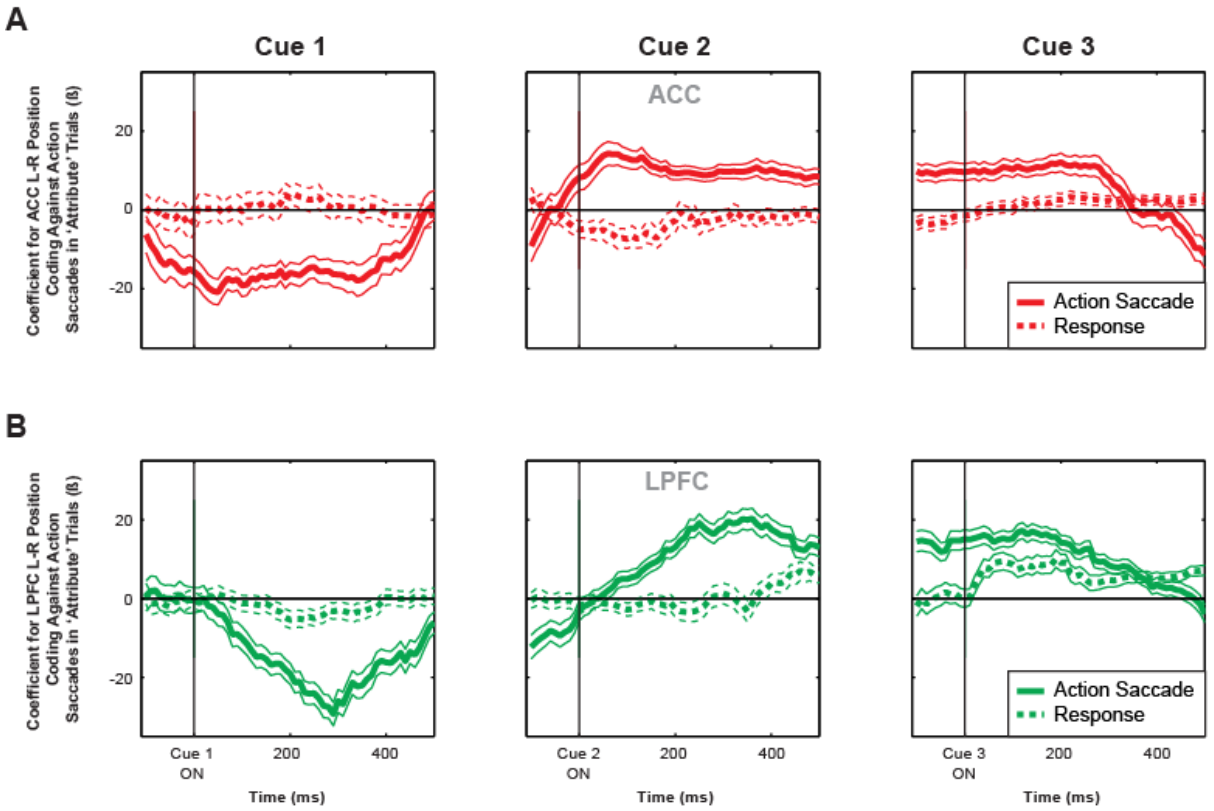

Figure 5.22: Position selective neurons predict future saccade behaviour in 'Attribute' trials. **(A)** A linear regression of the Cue 1 left-right spatial selectivity code in ACC against third saccade action (left-right) saccade (solid line) and final response (dashed line) coding at each cue. **(B)** The same analysis as (A) using LPFC. All other figure properties are the same as **Figure 5.21**.

The final question we asked with respect to information gathering was whether attribute specific neurons in OFC could also predict attribute specific saccades in 'Option'

trials (i.e. predict which attribute the third saccade would be directed towards)? At both Cue 1 and Cue 2, no area significantly encoded future attribute saccades (**Figure 5.16**, left and middle panels) (binomial test, all p>0.05). However, after the saccade 10% of OFC neurons encoded the currently attended attribute type (**Figure 5.16**, right panel).

We next performed a projection analysis of Cue 1 attribute value coding against the attribute saccade code at each cue. At Cue 1, LPFC showed a significant negative relationship (linear regression, p<0.05). At Cue 2, there was a small but significant positive relationship in OFC Cue 1 attribute value coding neurons, which was significantly greater than ACC (linear regression, T=1.98, p<0.05, linear hypothesis test of the beta coefficient for OFC against ACC, p<0.02). This implies that neurons that fire more for probability value in OFC also fire more when a probability stimulus will be viewed at Cue 3 in 'Option' trials. At Cue 3, this relationship in OFC had attenuated, but a significant positive relationship emerged in LPFC (**Figure 5.25**) (linear regression, T=3.01, p<0.003).

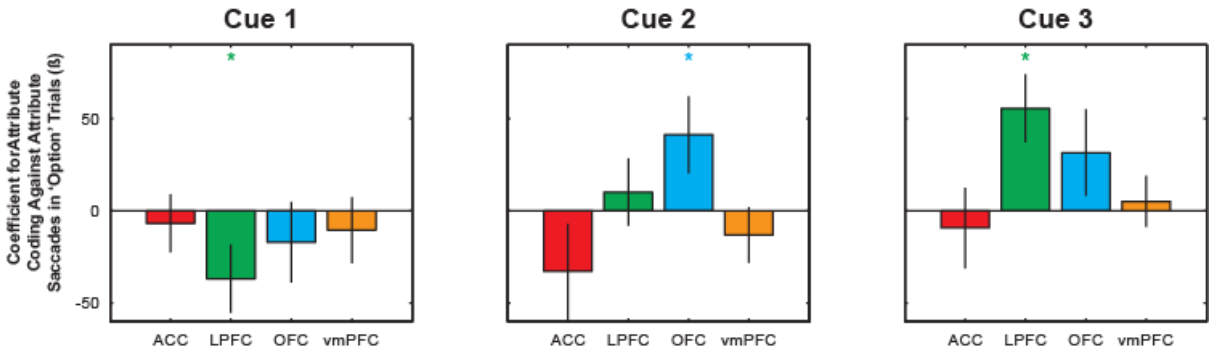

Figure 5.23: OFC attribute value neurons predict attribute saccades in 'Option' trials. A linear regression of mean Cue 1 attribute value coding against attribute saccade coding (i.e. saccades towards probability versus magnitude information) in 'Option' trials at each cue. All other figure properties are the same as **Figure 5.19**.

A sliding projection analysis allowed us to empirically inspect the time course of both the OFC and LPFC signal (**Figure 5.24)**. The LPFC signal was observed to be weak and phasic. It was strongest 200ms after Cue 2 onset and after the third was presented (i.e. post-saccade) (**Figure 5.24A**). OFC showed phasic coding of the attribute saccade. The peak strength was around 400ms post-cue onset at both the first and second cues (**Figure**

5.24B). In comparison to the action and attention saccade signals, this attribute signal was noticeably weaker and was also phasic in nature rather than being maintained across cues.

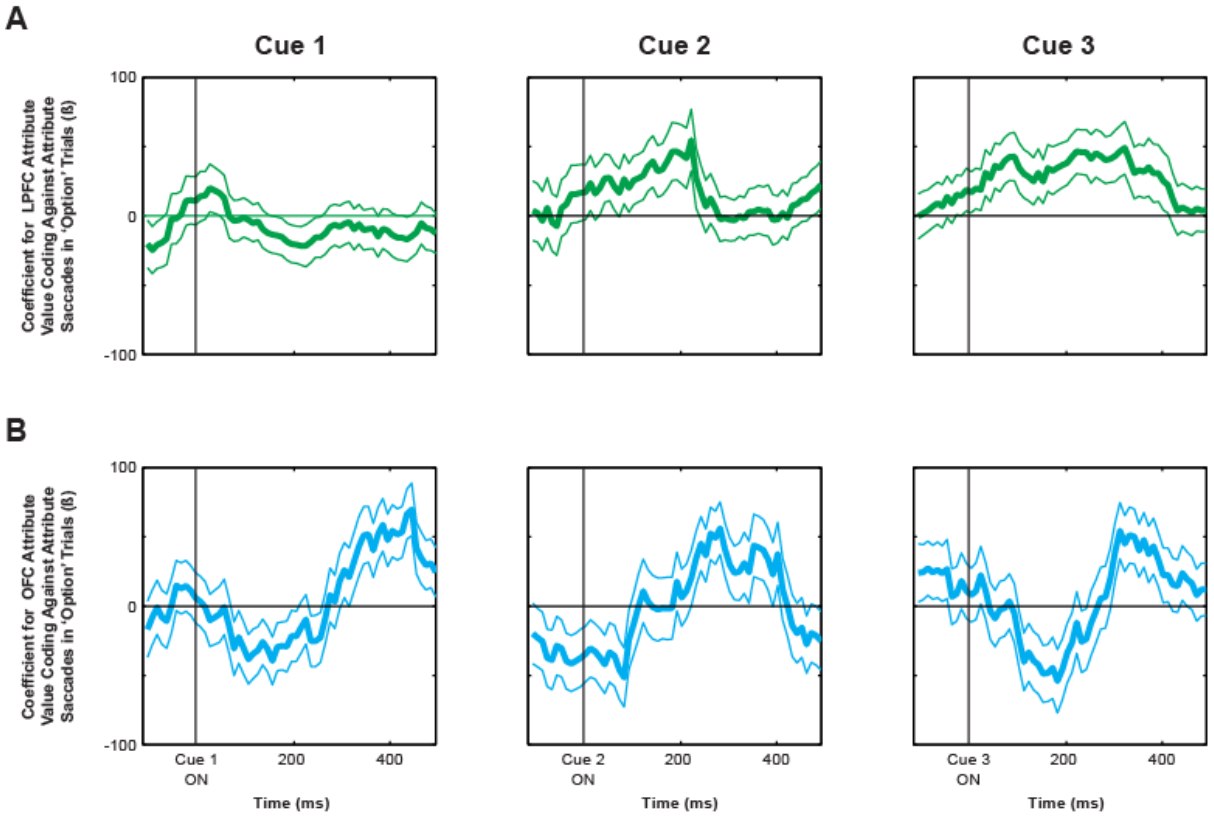

Figure 5.24: LPFC and OFC attribute value neurons encode future attribute saccades. **(A)** A linear regression of the Cue 1 LPFC attribute value code against Cue 3 attribute (probability-magnitude) saccade (solid line). **(B)** The same analysis as (A) using OFC neurons. All other figure properties are the same as **Figure 5.21**.

In conclusion, we have presented evidence that future saccade behaviour in this task can be predicted by neuronal firing patterns in ACC, LPFC and OFC. Exactly like value computations, these future saccade signals are differentially framed: by action in ACC, by space in LPFC and by attribute and attention in OFC. These results imply that the respective valuation reference frames in each region are simultaneously utilized for both valuing information and valuing choices.

Discussion

In this chapter we have reported experimental data which demonstrates that value and information gathering signals in primate PFC are *simultaneously* represented in three separate frames of reference; attention, action and attribute. Furthermore, we have shown that computations within each of these frames of reference reflect value at the earliest stages of information gathering, but also evolve to encoding elements of the final choice by the end of the decision.

Attentional Value and Choice Coding in OFC

In this experiment we observed ubiquitous coding of value (in terms of the value of the currently attended stimulus) throughout all four PFC regions much in line with many previous neurophysiological studies (for example, Kennerley et al. (2009), Roesch and Olson (2003), Strait et al. (2014)). This general coding of attentional value may well reflect a representation of subject motivation (Roesch and Olson, 2003, Roesch and Olson, 2004). If this were true then one might expect to see a positive correlation between the value memory trace of any previous cues and the value code for the current cue in all trial types. Instead it was only observed in 'Option' trials specifically in OFC. However, it could represent a simple and generic value signal that propagates throughout the brain when the environment changes.

It should also be noted that in our experiment, we were able to explicitly isolate both overt attention (by requiring fixation of a cue) and covert attention (by covering up all other information) at Cues 1-2. This meant we had better control of experimental variables that could confound a value signal (chosen vs unchosen offers, value difference calculations,

etc.). One obvious explanation for such ubiquitous value coding across the brain is that neurons in different brain regions are performing different value computations, but the experimental paradigm may not have afforded the control to isolate the specific computation or frame of reference in which these neurons were coding value. Whilst attentional value coding was nonetheless present across all regions in our dataset, the fact that these value signals evolved as the decision process evolved toward a choice emphasises the potential difficulty in interpreting previous and future studies where multiple stimuli are presented simultaneously and/or when the focus of attention cannot be confidently determined.

It is clear from this data set that the OFC value neurons have a unique roll. They are the only neurons that consistently compute the value difference of information with respect to the location of current attention. This is an almost analogous signal to the vmPFC and ventral striatum BOLD signal reported by Lim et al. (2011). The reason for the discrepancy between regions in the present study remains unclear. One possible explanation is that the BOLD signal is difficult to measure in the human OFC (Weiskopf et al., 2007) and therefore it is unknown whether attention related computations also occur in OFC. Another potential explanation for the absence of this signal in vmPFC in our study could be a combination of the fact that its neurons were relatively under-sampled compared to OFC and the fact that value coding power may be weakened by the fact that vmPFC neurons generally showed lower firing rates during the task. As the linear regression is very sensitive to both degrees of freedom and variance in the dependant variable, 160 neurons which only had small firing rate modulations may not be enough to bring out a small effect in vmPFC. Finally, given that resting state MRI has shown that connectivity patterns of human and primate vmPFC show strong homologies (Neubert et al., 2015), it is also possible that neuronal firing in vmPFC is not directly proportional to the BOLD signal in fMRI.

A completely novel finding in the current study is the representation of attentional choice in OFC at the population level. In their sequential option vmPFC study, Strait et al. (2014) report that residual firing in approximately 12% of vmPFC neurons accounted for

whether the first or second presented option was to be chosen. This may reflect an attentional choice signal, but because the temporal order of stimulus presentation and attention are correlated in their task this may actually represent a temporal choice signal. In the current behavioural task, subjects have a tendency to gather information about options that they go on to choose, but attention and cue order are overall much less correlated than in the Strait et al. 2014 study. Based on the work of Lim et al. (2011) one might hypothesise that attentional choice signal would be specific to vmPFC/OFC, but surprisingly this was not the case, instead ACC and OFC also encoded attentional choice but in the OFC the effect was stronger.

Action Value Coding in ACC and LPFC

It is of no surprise that neurons in both ACC and LPFC were observed to encode action value during the early stages of the decision, as might be predicted from the findings of Matsumoto et al. (2007) and Hayden and Platt (2010) who both found neurons in ACC which were modulated by value and action. Furthermore, the evolution of the action value signal to become a response signal is similar to the results observed by Kim et al. (2008) who found that during intertemporal choice, LPFC neurons encoded action value during the early phase of the decision and then went on to encode response. ACC neurons have also been shown to multiplex value and response coding with the value coding occurring before the response (Kennerley et al., 2009). One advance of the current study over these studies is that we have shown that the ACC action value neurons are the *same* neurons that encode the response during the choice phase. This result bears a resemblance to that of Cai and Padoa-Schioppa (2014), who reported that LPFC neurons encoded the value of the chosen item in good space before transforming to represent the response before the subjects respond. The major difference between the aforementioned study and the current study is that the task of Cai and Padoa-Schioppa study demanded that subjects solve the decision in

'goods space' because they were not informed of the relevant action until after the stimuli were presented. Therefore it is unsurprising that the LPFC neurons encoded the chosen value in goods space before evolving to encoding the response. In contrast, in the current study, subjects were always aware of the actions contingent to the choice, allowing them to solve the task in action space if they chose to. Unlike the two aforementioned studies, the current study finds action value and response coding in ACC rather than LPFC. However, these tasks required saccade-contingent choices, whereas the current study required a joystick movement to indicate choice. Therefore this discrepancy might be expected given LPFC's strong connections to supplementary eye fields (Huerta and Kaas, 1990, Gerbella et al., 2010) and cingulate cortex's connections to limb motor areas (Morecraft and Van Hoesen, 1992, Kennerley et al., 2009). Given the clear action value to response transformation within ACC, it is no surprise that ACC lesions cause deficits in action based decision making as described in humans and primates (Rudebeck et al., 2008, Camille et al., 2011b, Kennerley et al., 2006, Hadland et al., 2003).

Several strands of evidence lead us to the conclusion that LPFC neurons that were classified as action value coders at Cue 1 were in fact encoding spatial value. Firstly, unlike ACC, a roughly equal proportion of neurons encoded top-bottom value as encoded left-right value. Such an encoding scheme has no relevance to solving the task. Secondly, unlike ACC neurons, the LPFC action value coders did not go on to encode response or action value by the end of the trial. Finally, it is well known that LPFC represents spatial aspects of the environment as well as the interaction between space and value (Kennerley and Wallis, 2009c, Rao et al., 1997).

Attribute Specific Coding in OFC

The current study demonstrates the existence of specific neurons in OFC which discriminate the value of stimuli based on their attribute type. However, this is a

subpopulation of neurons and in fact many neurons in OFC simply encode general attentional value. These findings are in some ways very similar to those of Kennerley et al. (2009) with the exception that we found that attribute specific neurons are specific to OFC whereas Kennerley et al. found no region more likely to demonstrate attribute selectivity. In contrast to studies which report that attributes are represented by different populations of neurons (O'Neill and Schultz, 2010, Morrison and Salzman, 2011, Roesch et al., 2006), our data indicates that there are some OFC neurons which differentially represent probability and magnitude values, but at the same time encode both variables (purple dots with orange diamonds in **Figure 5.5B**). In other words, these are neurons which either encode the two attributes with opposite firing rate relationships with value, or neurons which encode the value of both attribute, but one attribute significantly more than another.

The existence of OFC attribute value neurons may provide a neurophysiological explanation as to why primate and human subjects with OFC lesions struggle to perform stimulus based learning and decision making tasks (Rudebeck et al., 2008, Camille et al., 2011b). These stimulus based decision making tasks usually involve subjects making decisions within one attribute and therefore if subjects do not have attribute specific neurons it may impair their ability to do this. However, a recent study by Rudebeck et al. (2013b) showed that the classically reported OFC lesion finding of impaired stimulus based decision making may actually arise from damage to white matter tracts around OFC rather the loss of OFC itself (Schmahmann et al., 2007). However, the previously observed deficits in reinforcer devaluation were still present when white matter sparing OFC lesions were given (Rudebeck et al., 2013b). Although these results might imply that OFC is not necessary for stimulus (or attribute) based decision making, Rudebeck and colleagues suggest that there may be redundancy within these stimulus/attribute computations where the loss of OFC may allow the problem to be solved in areas such as amygdala or ventral striatum. Damage to the white matter tract around OFC may therefore give rise to the observed deficits by disrupting communication of stimulus specific learning signals to PFC areas.

As stated previously, the current task had no measure of attribute based choice, and therefore there is no way to know when subjects were guiding their behaviour based on information from one attribute or another. Therefore the 'attribute specific choice' signal that we defined only represents a correlate of choice and not necessarily causative to choice. Furthermore, the signal is quite weak and this could correlate with the fact that behaviourally, subjects only a weak attribute bias (see Chapter 4).

Parallel vs Sequential Value Processing in Value Based Decision Making

The fact that three separate reference frames are represented in PFC during decision making lends strong evidence to the idea that decision making circuitry may be organized in a parallel manner which may be consistent with the decision through consensus hypothesis. This hypothesis states that values can be represented and compared in different frames of reference in multiple regions across the brain and that final decisions come about through consensus between areas (Cisek, 2012). In contrast the goods based model of decision-making proposes that options are represented and compared as 'goods' in OFC/vmPFC which are invariant of attribute type (Padoa-Schioppa, 2011). It postulates that only after this abstract comparison can actions be selected in order to attain a chosen option. This model makes several predictions that are not only clearly different from the decision making through consensus hypothesis, but are also directly tested and refuted by our study. First of all, as stated earlier, there is clear evidence for attribute and action value coding subpopulations in PFC, and these populations also go on to represent choice and response respectively. This implies that these subpopulations are not merely providing an input into a 'goods' valuation process, but are also exhibiting the features of a competitive process in biasing the actual choice. Secondly, the 'goods based' model stipulates that actions are only selected when the decision is already made as has been demonstrated in some tasks (Cai and Padoa-Schioppa, 2014). In contrast, in the current study, response

codes are seen in ACC as early as 700ms before the response, implying that the coding of response ramps over a long timescale rather than just simply after the comparison of goods. In contrast, these findings are in agreement with what would be expected of a parallel processing framework (**Figure 5.1**) (Cisek, 2012).

Information Gathering and PFC

Neurons in ACC, LPFC and OFC all predict future information gathering behaviour. Perhaps most surprisingly, many of these signals are specific to particular frames of reference. It is important to note that the regressions used to define these *future* information gathering signals accounted for – and therefore should be independent of – value. Therefore, there is no reason to expect that value sensitive neurons should be more likely to also be information gathering neurons. However, it is intriguing that in fact, value neurons in our study often also encoded information about future saccades. Although these results are novel, they share some similarities with results from previous studies. The fact that ACC action value neurons also predict the third saccade direction is similar to the finding of Hayden and Platt (2010) who described neurons that multiplexed both value and action in terms of saccade. The current study suggests that these neurons are capable of multiplexing action of different modalities (i.e. saccades as well as limb movements) during the decision process. The idea that LPFC spatial neurons can predict the spatial location (i.e. the left or right side of the screen) is also in line with the finding that LPFC neurons can code for both value (Kennerley et al., 2009), spatial position (Rao et al., 1997) and can also multiplex both (Kennerley and Wallis, 2009c). The role of ACC and LPFC in saccade behaviour is also plausible from an anatomical point of view given that both areas send direct connections to supplementary eye fields (Gerbella et al., 2010, Luppino et al., 2003).

Although we defined third saccade behaviour in the frame of reference of attention in both ‘Attribute’ and ‘Option’ trials, these signals could also be considered to be abstract

strategy signals. It has been demonstrated that LPFC neurons encode strategies and rules in various types of tasks (Asaad et al., 2000, White and Wise, 1999, Genovesio et al., 2005). This may explain why a large proportion of LPFC neurons in the current study were observed to encode the 'attentional saccade' but (unlike OFC) did not correlate with value coding.

Fellows (2006) found that patients with vmPFC lesions avoided gathering information in an attribute based manner in favour of option based comparisons. In the current study, subjects were not given the freedom to choose the initial information gathering strategy but we found no evidence that vmPFC encoded information gathering strategies in any of the tested frames of reference. One potential explanation for Fellows' findings is that many of the vmPFC lesion patients also had subtotal damage to OFC. These patients would therefore have lost the attribute value neurons described in the current study, which we observed also drive attribute based information gathering behaviour. Therefore, these patients would not be able to perform both present and future attribute comparisons, and instead would be forced to compare value and gather information in a different manner, such as between options or actions.

Very few neurophysiological studies have considered the how information gathering is represented in the brain. One of the few studies that have is that of Blanchard et al. (2015) who found that OFC neurons encode value of choices orthogonally to the value of information. This is in contrast to dopamine neurons which appear to encode reward and information value with a common signal (Bromberg-Martin and Hikosaka, 2009). The findings of our study would suggest that information gathering signals are more similar to the behaviour of dopamine neurons although there are several critically difference between the abovementioned studies and our study. Firstly, previous studies examine the *value* of information rather than *how* information is gathered. It was not possible for us to examine information value because this was perfectly predicted by the value of previous stimuli in our task. By examining information gathering strategies we examined the output of the

information gathering decision rather than the input which is information value. Furthermore, in their task this information value signal could not be used to drive behaviour other than to reduce uncertainty about future outcomes. In contrast in our study the act of information gathering was a critical part of the choice process and our information gathering signals also had a clear output (i.e. the saccade). Therefore, it is possible that their finding in OFC and our finding may reflect different computation signals entirely.

In conclusion, the neuronal data presented in this chapter shows the simultaneous representation of value computations and choice signals in three separate frames of reference, which implies the existence of parallel decision circuitry in PFC. These findings unify several ideas about value computations and competition within PFC. We have also shown that these frames of reference also extend to the domain of information gathering, which suggests that the concept of frames of reference may extend across cognitive domains, or that both information and choices/outcomes are valued in a similar way by the PFC valuation circuitry.

Chapter 6: Outcome and Prediction Error in the Prefrontal Cortex

When we make a decision how do we know whether it is right one or not? One method by which one might do this is to consider one's outcome compared to one's expectation. If the outcome exceeds the expectation then logic dictates that given the same situation in the future, one should not only make the same choice but also assign more value to the chosen option. However, if the expectation was greater than the actual outcome then one must assign lower value to this option when making choices in the future. This chapter will examine whether PFC neurons perform such outcome related computations and whether such computations are specific to decision attributes (such as probability or reward size) or to functional subpopulations of neurons.

Introduction

Tracking the values of our outcomes is an important process in allowing us to adapt our behaviour to be as optimal as possible in uncertain environments. It has been established in many studies that neurons in ACC, LPFC, OFC and vmPFC can encode information about unexpected outcomes as well as the presence or absence of a reward (Kennerley and Walton, 2011, Ito et al., 2003, Kennerley and Wallis, 2009b, Strait et al., 2014, Monosov and Hikosaka, 2012, Quilodran et al., 2008, Seo et al., 2007, Seo and Lee, 2009, Tremblay and Schultz, 2000b, Roesch and Olson, 2003, Sallet et al., 2007). Neurons in ACC, LPFC and OFC have also been shown to encode the amount or type of reward given (Hikosaka and Watanabe, 2000, Kennerley and Wallis, 2009b, Rolls et al., 1990). Interestingly, although many of these neurons encode outcomes during the feedback phase

of the experiment, neurons in ACC, LPFC and vmPFC have been shown to maintain this coding into subsequent trials (Seo et al., 2007, Seo and Lee, 2009, Strait et al., 2014). One possible function of these sustained outcome signals is to modulate the representation of value with respect to the recent history of outcomes or the internal state of the subject. The former representations have been reported in OFC where encoding of value on the current trial is modulated by the value of the choice on the previous trial (Kennerley et al., 2011), whereas vmPFC encodes value in the context of internal states (Abitbol et al., 2015, Bouret and Richmond, 2010).

Another important function of representing one's outcomes is to learn about the current environment. Over and above tracking the outcome of a decision, it is important to compute whether outcomes are expected or not. Several studies have shown that ACC neurons are responsive to behavioural errors (Totah et al., 2009, Quilodran et al., 2008, Ito et al., 2003). Such a signal in its self could be critical for selecting future actions that avoid such errors. Furthermore, the encoding of prediction errors (PEs) is central to the idea of learning and updating the value of one's stimuli in the environment. Classically neurons midbrain dopamine neurons have been shown to encode these PEs, but extensive neuronal data show that ACC and OFC can also encode PEs (Schultz et al., 1997, Matsumoto et al., 2007, Kennerley et al., 2011, Sul et al., 2010). Furthermore, some have suggested that the coding of prediction errors may be a general feature of adaptive neuronal circuits throughout the brain (Friston, 2009). Importantly, similar (but not identical) to dopamine neurons, ACC neurons also strongly code for the value of stimuli of various modalities before outcomes are presented (i.e. a prediction of their value) (Kennerley et al., 2009, Amiez et al., 2006). Interestingly, in ACC PEs have been reported in both the value domain (Kennerley et al., 2011) and the action domain (Matsumoto et al., 2007) although it is unclear whether this is a generalised computation that ACC can perform in multiple frames of reference or whether specific populations of neurons encode the two different PEs. The role of ACC in reinforcement learning has also been demonstrated in several studies (Nee et al., 2011,

Jessup et al., 2010, Silvetti et al., 2013, Brown and Braver, 2005). This may provide a functional basis for why PEs are observed within ACC.

Many human studies have reported error related negativity (ERN) detected on scalp electrodes positioned over ACC in a variety of tasks (Silvetti et al., 2014, Van 't Ent and Apkarian, 1999, Holroyd et al., 1998, Holroyd and Coles, 2002, Walsh and Anderson, 2012). Some evidence suggests that these so called error signals may also reflect uncommon results in trials (such as errors in most tasks) therefore reflect a surprise signal (Jessup et al., 2010, Brown and Braver, 2005). This negativity has also been shown to correlate with PEs in several simple tasks (Silvetti et al., 2014, Nunez Castellar et al., 2010, Talmi et al., 2013). The ERN has also been shown to come about from both positive PEs (i.e. unexpected rewards) and negative PEs (unexpected losses) (Silvetti et al., 2014). The idea that ACC is capable of computing PEs is given credence by the fact that it receives strong dopaminergic input from the ventral tegmental area (Williams and Goldman-Rakic, 1998, Berger et al., 1988, Gaspar et al., 1989).

Others have also suggested that ACC computes outcomes signals that represent both conflict (van Veen et al., 2001), volatility in the environment (Behrens et al., 2007) and error likelihood (Brown and Braver, 2005) although it has been postulated that all of these effects could rise from ACC acting as an actor-critic reinforcement learning model in which PEs are computed (Silvetti et al., 2011). Another unifying idea aimed at explaining this diversity of findings revolves around the concept of the ‘expected value of control’ where ACC integrates the payoff expected from the exertion of cognitive control and integrates this with the cost and amount of control required (Shenhav et al., 2013).

The results in this chapter will outline the neuronal correlates of outcome related computations in the PFC in our decision making paradigm. We aim to show that neurons in all four area of PFC that we record from encode the final outcome of a decision. However, we will show that specifically in ACC there is a clear positive relationship between value

coding at the choice phase of the trial and reward/no reward (R/NR) coding at the feedback phase. An evaluation of R/NR in the action/spatial domain will also be demonstrated in the LPFC. We will also show the presence of a double dissociation in R/NR coding in LPFC and OFC where the former is driven by neurons that prefer to encode reward magnitude whereas the latter is dominated by probability preferring neurons. We will present evidence that ACC neurons compute only positive probabilistic PEs whereas neurons in vmPFC encode both positive and negative probabilistic PEs. Furthermore, when rewards were omitted both ACC and OFC neurons represented the fictive value of these rewards. These results speak to the idea that ACC and other PFC regions play an active and complementary role in updating and monitoring the values of stimuli in order to optimise reward seeking behaviour.

Methods

The data presented in this chapter comes from a further analysis of the data set shown in Chapter 5. Therefore, all information pertaining to the subjects, behavioural protocol, task, neuronal recordings and the basic analysis techniques are discussed in detail in the Methods section of Chapter 5. In addition to the regression analyses defined in Chapter 5, additional regression models relevant to the current chapter will be described below.

During the feedback phase of the task, the grey box around the option chosen by the subject reliably changed colour (see **Figure 5.3**). If the subject received reward the grey box changed to a darker shade of grey for the entire time that the reward pump was in operation; if the subjects were unrewarded the box changed to a lighter shade of grey for the same amount of time that the subjects would have received reward. By having this visual cue informing the subjects about the presence or absence of reward we hoped to minimise any variance in the neuronal signal that may have come about from variation in the perception of reward timing. For analysis we used an epoch that began 200ms before feedback onset and ended 600ms post-onset.

In examining the response of neurons at the point of feedback, we wanted to ask whether neurons in PFC encoded outcome related information which might be relevant for guiding future decisions. Therefore at the response epoch we constructed a regression model which accounted for the chosen probability value and chosen magnitude value both separated by whether the trial was rewarded or unrewarded as well as a binary rewarded-unrewarded regressor. A detailed outline of this model can be found in **Table 6.1**.

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Chosen Probability Rewarded Trials	Value	4	Chosen Magnitude Unrewarded Trials	Value
2	Chosen Probability Unrewarded Trials	Value	5	R/NR (Rewarded-Unrewarded)	Categorical Reward
3	Chosen Magnitude Rewarded Trials	Value	6	Constant	Mean Firing Rate

Table 6.1: A list of regressors and their interpretations used in the multiple linear regression of neuronal firing rate aligned to Feedback onset.

In order to examine whether neurons multiplexed outcome and action signals at feedback we used a similar regression model to the model described in **Table 6.1** with the exception that we split each regressor down by the response that the subject made on each trial (see **Table 6.2**). We then used a linear hypothesis test to contrast R/NR coding on left response trials against outcome coding on right response trials.

#	Regressor	Interpretation	#	Regressor	Interpretation
1	Chosen Probability Rewarded Trials Left Choice	Value	7	Chosen Magnitude Rewarded Trials Right Choice	Value
2	Chosen Probability Unrewarded Trials Left Choice	Value	8	Chosen Magnitude Unrewarded Trials Right Choice	Value
3	Chosen Magnitude Rewarded Trials Left Choice	Value	9	Left Constant	Constant
4	Chosen Magnitude Unrewarded Trials Left Choice	Value	10	Right Constant	Constant
5	Chosen Probability Rewarded Trials Right Choice	Value	11	R/NR Left Choice	Categorical Reward
6	Chosen Probability Unrewarded Trials Right Choice	Value	12	R/NR Right Choice	Categorical Reward

Table 6.2: A list of action specific regressors and interpretations used in the multiple linear regression of neuronal firing aligned to Feedback onset.

Results

Two adult rhesus macaque monkeys (Subjects M and F) were trained to perform the same multi-attribute sequential information gathering and decision task as described in the previous chapter. As stated in Chapter 5, Subject M performed 32 recording sessions completing an average of 445 trials per session. Subject F performed an average of 394 trials per session over 25 recording sessions.

Neurons in PFC Encode Different Types of Outcome Information

If one wants to reflect on the astuteness of a decision the one must be able to represent the choice that was made as well as the outcome of the choice. Do PFC neurons in this task represent such signals? To test this we performed a linear regression of neuronal firing rate aligned to feedback onset against various relevant task parameters. We first examined the proportion of neurons that encoded whether reward was given or not at feedback. Neurons in all PFC regions significantly encoded rewarded trials (**Figure 6.1**). Encoding in ACC neurons was as high as 80% and was significantly greater than all areas (pairwise Chi^2 test, $p<0.05$). Furthermore, approximately 60% of OFC and LPFC neurons encoded the rewarded vs unrewarded (R/NR) regressor and both did so significantly more than vmPFC (pairwise Chi^2 test, $p<0.05$). We also noted that there was a significant skew towards negative coding of R/NR in both the ACC and LPFC populations (binomial test, both $p<8x10^{-4}$). This meant that most neurons in these regions responded when rewards were omitted rather than when they were present. These results led us to conclude that many ACC, LPFC and OFC neurons encode whether rewards are delivered or not.

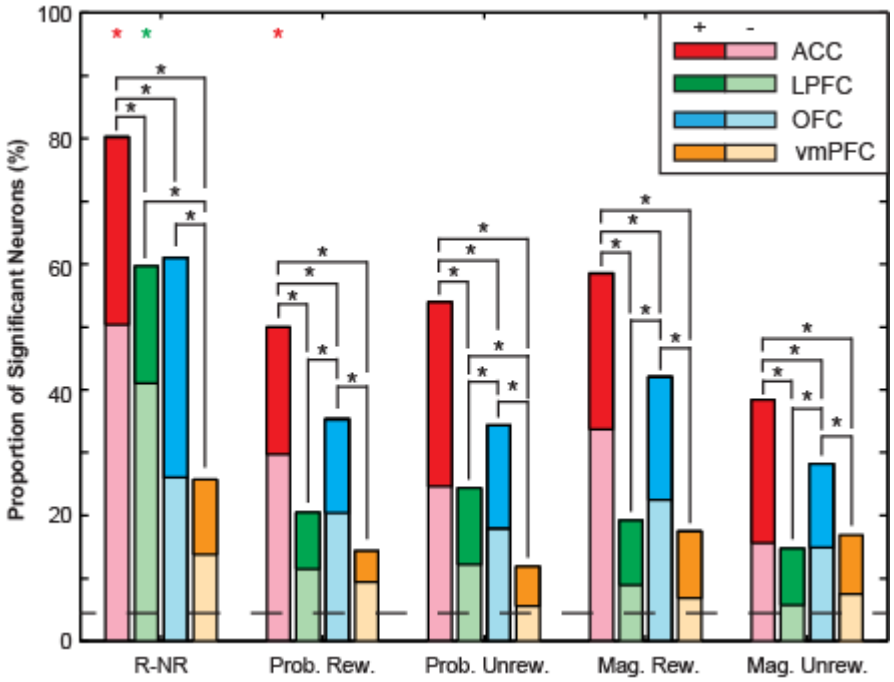

Figure 6.1: The neuronal coding of outcome parameters during reward feedback. The number of neurons in PFC which encoded whether a rewarded was delivered (R/NR), chosen probability value on rewarded trials (Prob. Rew.), chosen probability on unrewarded trials (Prob. Unrew.), chosen magnitude value on rewarded trials (Mag. Rew.) and chosen magnitude value on rewarded trials (Mag. Unrew.). Areas are split by whether neurons code these variables with a positive or negative sign. Black asterisks indicate regions that are significantly different from each other (Chi^2 test, $p<0.05$) and coloured asterisks indicate areas that show a skew in the coding sign of a regressor (binomial test, $p<0.05$). The dashed line indicates the 5% chance level.

Next we examined the profile of the encoding of the reward empirically by examining the pattern of coding across time for all neurons (**Figure 6.2A**). This demonstrated that coding in all areas was empirically sustained across the epoch but that even within single neurons; the coding in ACC was empirically stronger than other areas. It also appeared that ACC neurons tended to become significant earlier than other neurons (comparing the green lines and blue dashed lines between different areas). To test this quantitatively we performed a one-way ANOVA and multiple comparisons test of the latency of the first significant encoding of the R/NR variable for each neuron across regions which indicated that ACC coding was indeed significantly faster than all other regions (median latency; ACC: 200ms, LPFC: 300ms, OFC: 280ms, vmPFC: 350ms, One-way ANOVA, $p<4x10^{-9}$).

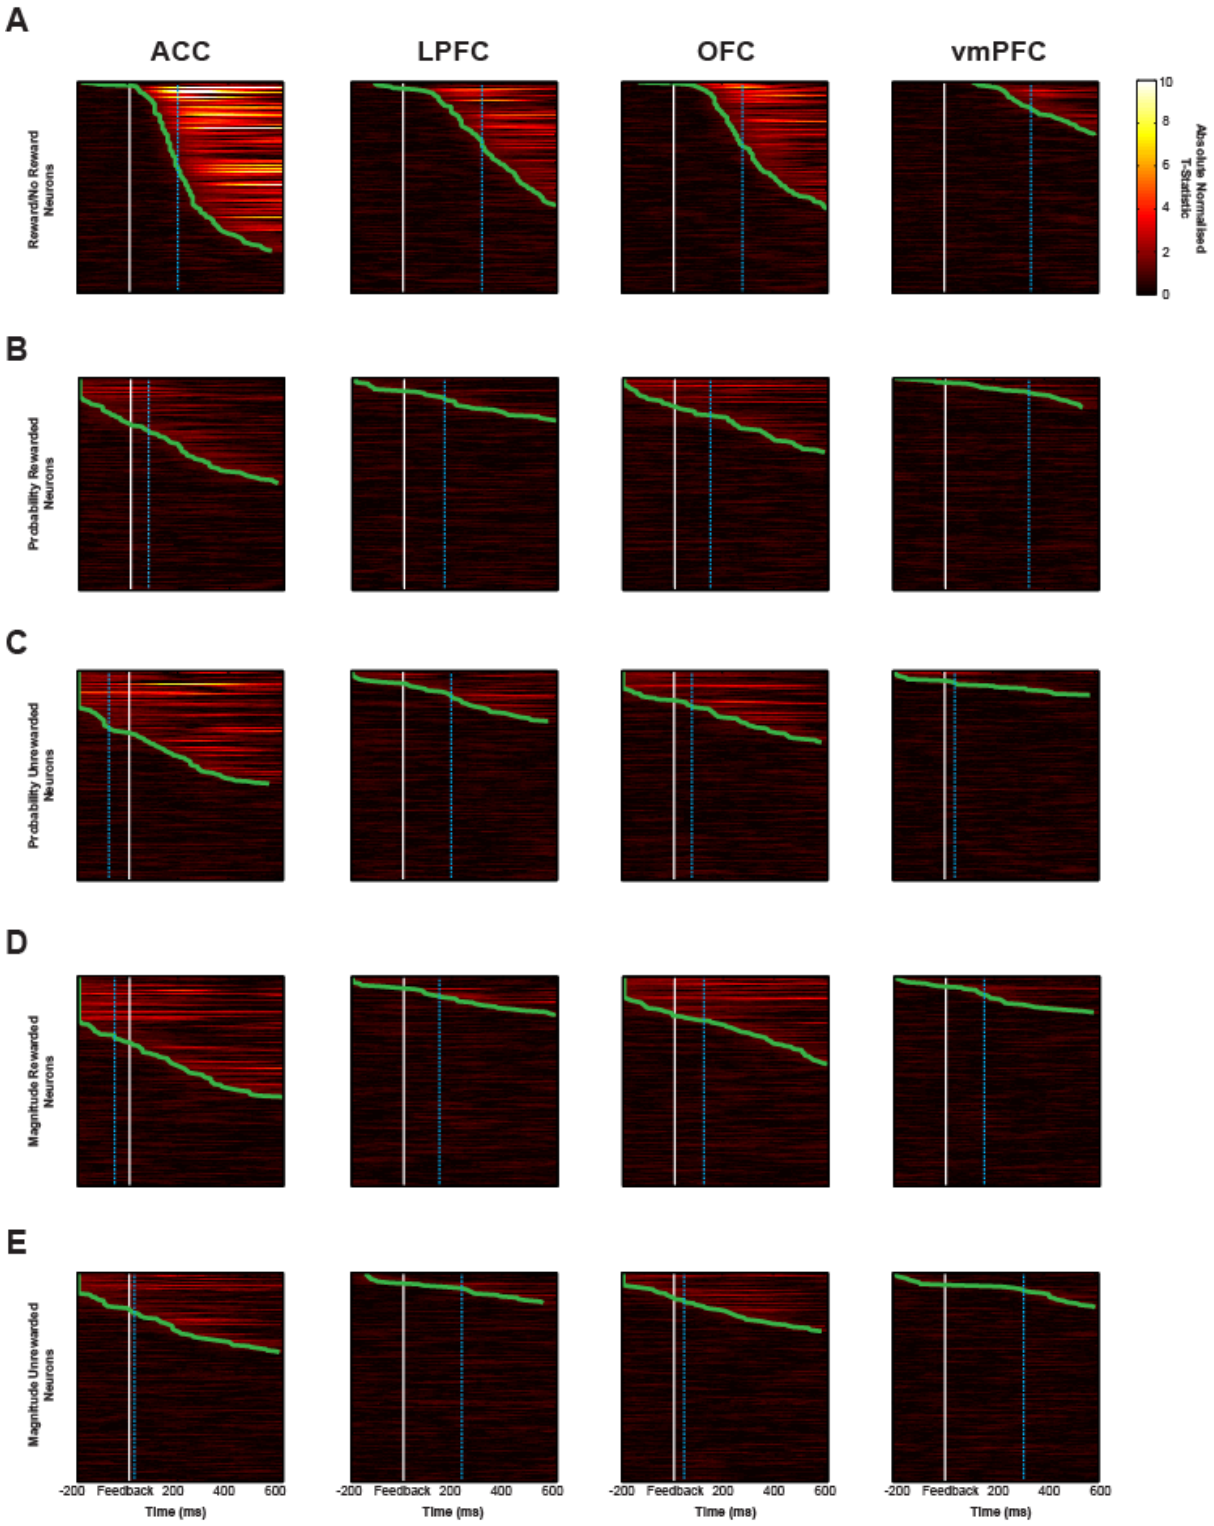

Figure 6.2: The temporal profile of value coding at feedback. **(A)** Heat maps of the absolute standardised T-statistics for outcome coding for every neuron in a region sorted by the latency of initial coding (green line). The white vertical line indicates the time of feedback onset. The median latency of encoding is shown for each region (blue dashed line). The same maps are drawn for chosen probability coding on rewarded trials **(B)**, chosen probability coding on unrewarded trials **(C)**, chosen magnitude coding on rewarded trials **(D)**, chosen magnitude coding on unrewarded trials **(E)**.

Next we asked whether PFC neurons coded for chosen value at outcome. As value was separated into probability and magnitude value we examined the coding of these separately. We found that neurons in all regions coded for chosen probability value on both rewarded and unrewarded trials (**Figure 6.1**) (binomial test, all $p<0.05$). However ACC encoded both significantly more than all other regions while vmPFC coding was significantly smaller than all other regions (pairwise Chi^2 test, all $p<0.05$). No region was observed to code probability value on rewarded trials any more than on unrewarded trials (Chi^2 test, $p>0.05$).

We next examined the coding of the chosen magnitude coding on rewarded and unrewarded trials. Chosen magnitude coding was significantly more prevalent in ACC and significantly less prevalent in vmPFC than all other regions (pairwise Chi^2 test, all $p<0.05$). Surprisingly, a significantly greater population of neurons in ACC, LPFC and OFC preferred to code the chosen magnitude on rewarded trials compared to unrewarded trials (Chi^2 test, all $p<0.05$). This latter result implies that the representation of the chosen magnitude in ACC, LPFC and OFC diminishes when this reward magnitude is not delivered.

We next examined the temporal profile of the representations of the four abovementioned variables around the feedback time (**Figure 6.2B-E**). As expected many neurons represented all of the variables before feedback onset. However, a larger subset of neurons also showed their first encoding of the chosen values (i.e. one of probability or magnitude) *after* feedback set implying a re-emergence of the value signal at feedback.

Whilst the above result suggested that chosen magnitude coding could be modulated by whether a reward was received, this did not specifically explore the effect of R/NR on each neuron. In other words, we wondered whether the same neurons encoded chosen probability (or magnitude) on both rewarded and unrewarded trials. To test this, we performed a contrast of the chosen probability on rewarded vs unrewarded trials, which revealed that all regions did significantly encode them but that the coding was significantly

greater in ACC and significantly lower in vmPFC (**Figure 6.3A**) (binomial test, p<0.05, pairwise Chi^2 test, <0.05). We performed a similar analysis for chosen magnitude differential on rewarded versus unrewarded trials and found similar results (**Figure 6.3A**) (binomial test, p<0.05, pairwise Chi^2 test, <0.05). However, 30% of OFC neurons were observed to perform this differential computation which was greater than both LPFC and vmPFC (pairwise Chi^2 test, <0.05). Finally, we examined the temporal profile of this differential coding scheme in each region (**Figure 6.3B** and **C**). This analysis showed that for both attributes the median latency of the differential code was relatively slow (more than 200ms in all regions) and that there were very few empirical differences between attributes.

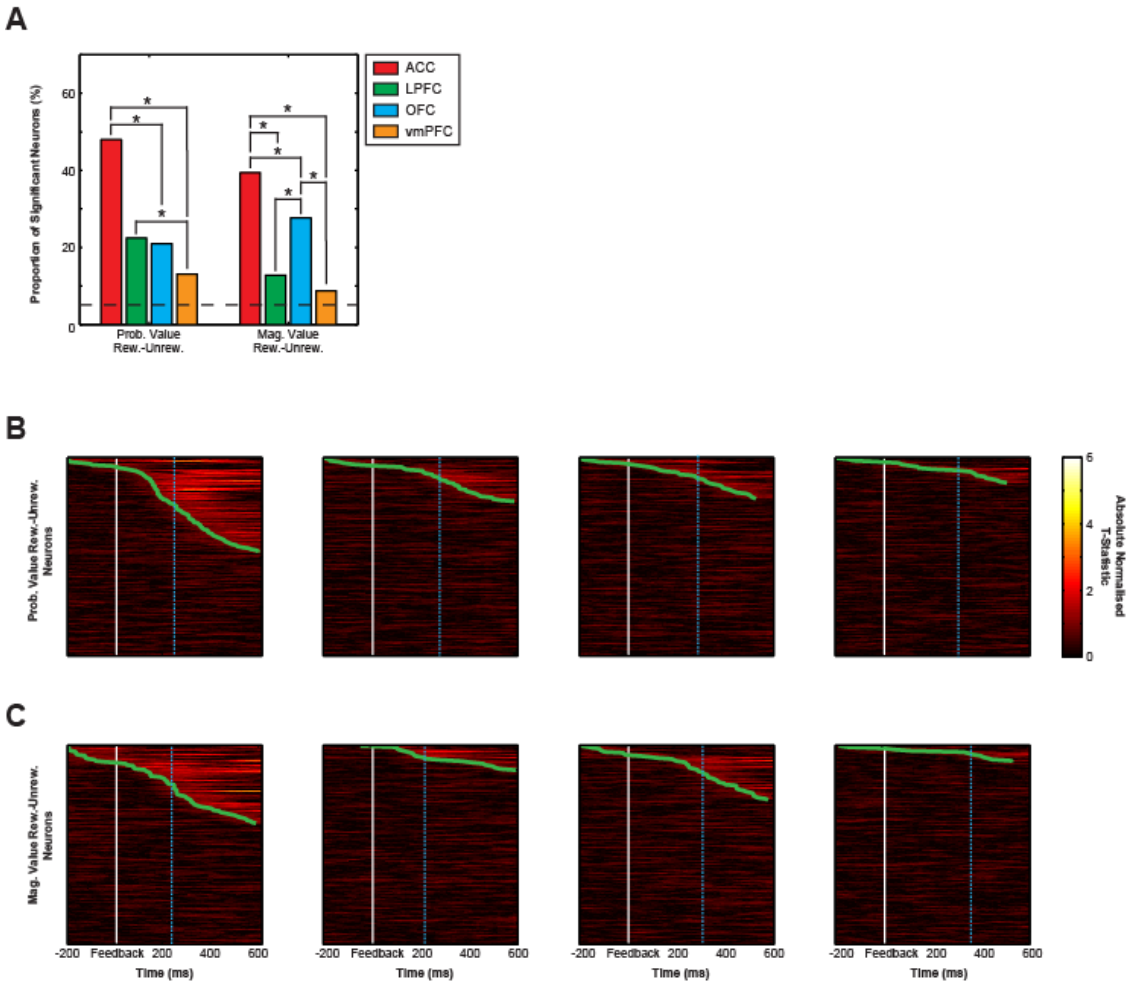

Figure 6.3: Differential coding of value based on categorical outcomes. **(A)** The number of neurons which significantly encode value probability value and magnitude value differentially based on whether the trial is rewarded or not. **(B)** Heat maps showing the strength of the differential coding of probability value. **(C)** The same heat maps for magnitude value. All other figure properties are the same as **Figures 6.1** and **6.2**.

Relationship Between Cue 1 Value Coding and R/NR Outcome Coding

Having identified a number of unique outcome-related signals across PFC neurons, we next examined whether neurons that encoded these outcome representations were also the same neurons that encoded value information when the subjects were forming predictions about the value of the upcoming outcome. For these projection analyses, we used information about Cue 1 value centred in a 500ms epoch time-locked to Cue 1 onset. We performed a projection analysis of the Cue 1 value code (defined in Chapter 5, Methods, **Table 5.1**) against the mean R/NR outcome code at feedback (**Figure 6.4A**). This indicated that there was a significant positive correlation between ACC value coding and outcome coding. This effect was significantly stronger than the effect measured in OFC but no other region (linear hypothesis test of the beta coefficient for ACC against OFC with Bonferroni correction, $p<0.008$). The interpretation of this is that ACC neurons that respond most strongly to value also best discriminate whether the subject received reward or not. We then examined the temporal profile of this signal by performing a sliding projection of the Cue 1 value code against the outcome code (**Figure 6.4B**). The relationship between these two codes in ACC was empirically observable within approximately 100ms after the feedback started but it reached its peak at approximately 300ms post onset. Furthermore, this signal was at no point significantly greater than LPFC or vmPFC and also only became significantly different from OFC at the peak of its coding (i.e. at 300ms post-onset) (linear hypothesis test of the beta coefficient ACC against the coefficient for OFC with Bonferroni correction, $p<0.008$). This suggests that neurons which encode value during decision making go on to monitor the outcomes of the same decisions.

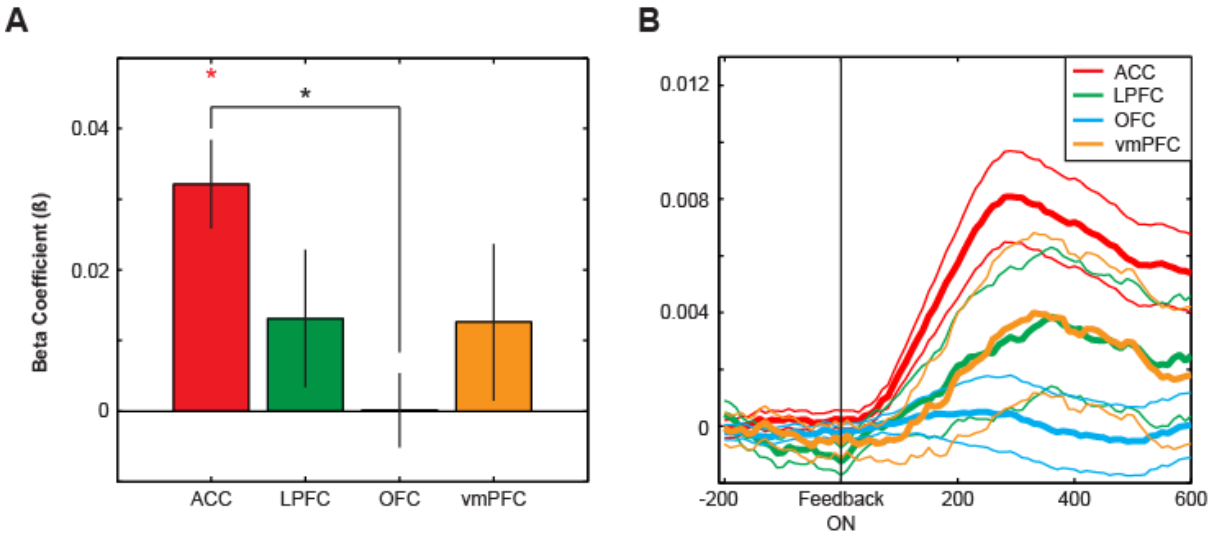

Figure 6.4: The relationship between value coding and outcome coding. **(A)** The results of a linear regression of the mean beta coefficients for value coding against the mean coefficients for R/NR coding between the start of the feedback onset and 500ms post-onset. Coloured asterisks indicate significant areas (linear regression, $p<0.05$) and black asterisks indicate areas that are significantly different from each other. **(B)** The sliding regression of mean Cue 1 value coding against outcome coding at feedback (vertical black line). Thin lines indicate SEM.

Although Cue 1 value neurons outside of ACC do not encode R/NR outcomes, are attribute specific neurons more likely to respond to whether trials are rewarded or not? In order to answer this question, we had to define a measure of attribute preference. In previous analyses (**Figure 5.4**), we used a contrast of the betas between the two attributes at Cue 1 to define differential attribute selectivity. However, the signed betas of these contrasts cannot reveal which attribute is more strongly encoded because value information itself can be encoded with either a positive or negative relationship with firing rate. Therefore, we computed the relative difference in the mean *absolute* coding strength for Cue 1 probability and magnitude value across all neurons. Then we computed the 75th percentile of the t-stat distribution for the relative coding of probability across all neurons ($\Delta T_{(Prob.-Mag.)}$=0.57) and all neurons that showed coding above this threshold were those defined as probability preferring neurons (see **Figure 6.5A**, green shaded area). We then correlated the relative probability preference for each area against the R/NR code at feedback (**Figure 6.5B**). This analysis revealed a significant correlation between the degree of probability

preference at Cue 1 and R/NR specifically in OFC neurons (linear regression, $T_{(OFC)}$=2.86, $p_{(OFC)}$<0.006). This suggests that the more an OFC neuron prefers to code probability value, the more it discriminates whether or not a reward is delivered.

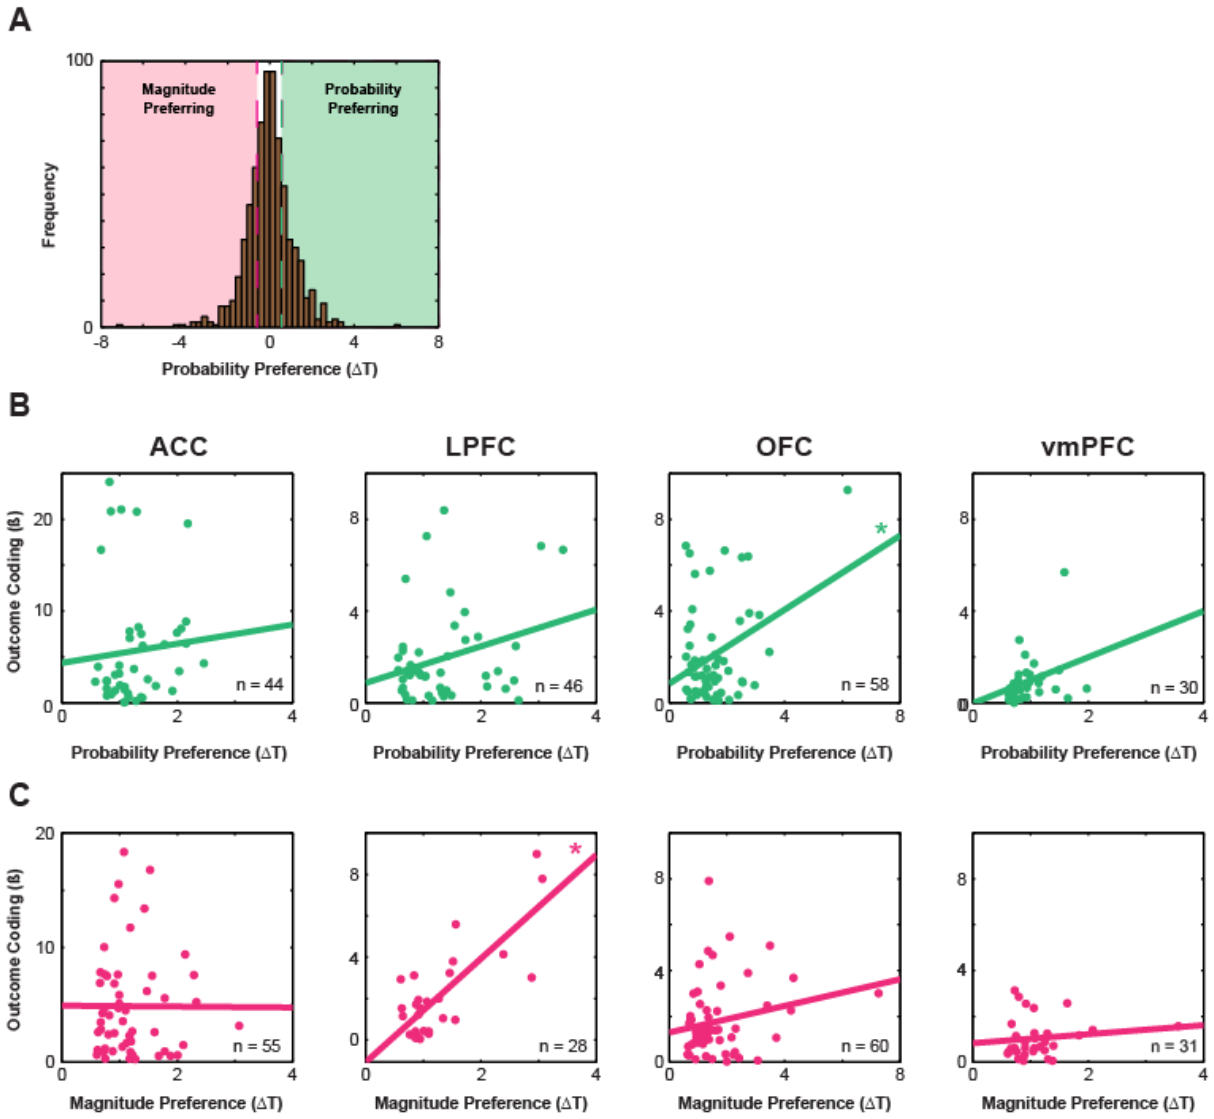

Figure 6.5: Attribute preferences predict outcome coding. **(A)** A histogram showing the distribution of absolute attribute preference across all neurons. The pink area indicates neurons that were considered to be magnitude preferring and the green area those that were considered to be probability preferring. **(B)** Scatter plots of mean absolute outcome coding against mean absolute Cue 1 probability preference for neurons within each area that pass preference criteria (as defined from the distribution shown in (A)). Asterisks indicate significant relationships (linear regression, p<0.05). **(C)** Scatter plots of mean absolute outcome coding against mean Cue 1 absolute magnitude coding for each region.

We next repeated the above analysis by defining a population of magnitude preferring neurons ($\Delta T_{(Mag.-Prob.)}$=0.60) (see **Figure 6.5A**, pink shaded area). Here, we found a significant correlation in LPFC (**Figure 6.5C**) (linear regression, $T_{(LPFC)}$=6.25, $p_{(LPFC)}<2x10^{-6}$). Therefore, the more these LPFC magnitude preferring neurons preferred magnitude value then the more they would discriminate rewarded and unrewarded outcomes. These results point to functional subdivision of attribute specific coding in LPFC and OFC in tracking whether or not outcomes are obtained.

LPFC Action Value Neurons Multiplex Chosen Action and Outcome

Do any neurons in the PFC reflect outcomes contingent on the actions that led to them? We first examined whether neurons in any region encoded the decision variables discussed above (i.e. R/NR, probability value rewarded and unrewarded or magnitude value rewarded and unrewarded) in the domain of the chosen action (**Figure 6.6**). We found that approximately 15% of ACC neuron did encode left-right action R/NR although this was only significantly different vmPFC (Chi^2 test, $p<0.05$). Significant proportions of neurons in ACC, LPFC and OFC were observed to encode probability and magnitude values on both chosen and unchosen trials (binomial test, $p<0.05$) but no clear pattern was observed.

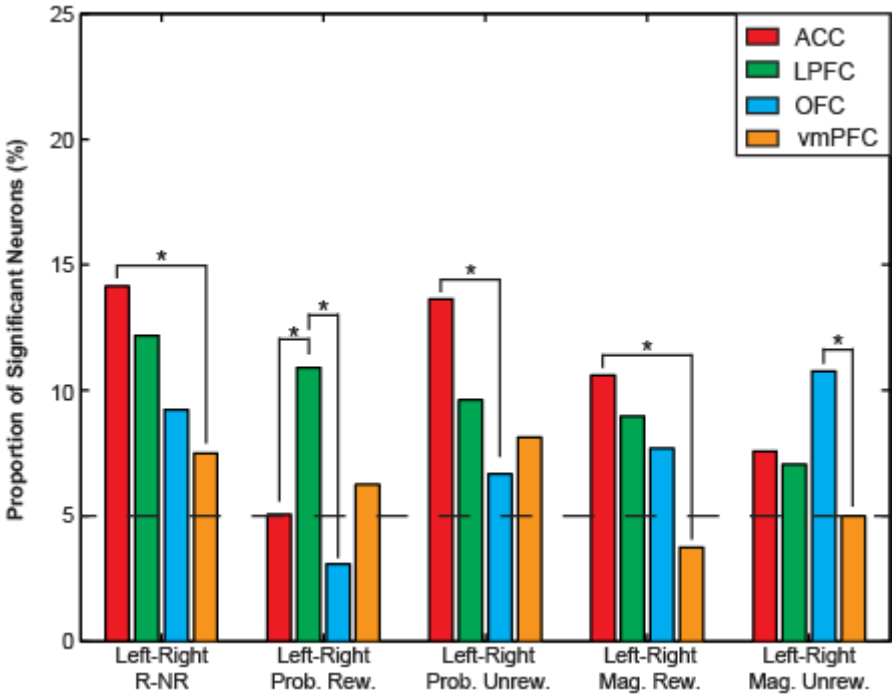

Figure 6.6: The neuronal coding of the action related outcome parameters during reward feedback. All other figure properties are the same as **Figure 6.1**.

We then performed a population correlation of Cue 1 action value coding against left-right action R/NR coding at feedback (i.e. a contrast of regressors 11 against 12 in **Table 6.2**). We found that there was a significant *positive* relationship between action value coding at Cue 1 and action dependent R/NR coding at outcome specifically in LPFC (**Figure 6.7A**) (linear regression, $T=2.81$, $p<0.006$). This means that a neuron that encoded left value preferentially over right value at Cue 1 is likely to discriminate the R/NR at feedback more when the subject chose left to obtain the outcome compared to occasions when the subject chose right. It should however be noted that this LPFC population may reflect spatial value computations rather than action value computations (see Chapter 5, Results). Though this effect was only present in LPFC, it was not significantly greater than in any other area (linear hypothesis test of the beta coefficient for LPFC against other areas with Bonferroni correction, $p>0.008$).

We then examined the temporal profile of this signal by performing a sliding projection analysis of mean Cue 1 action value coding against the action specific R/NR code

across the cue epoch (**Figure 6.7B**). This analysis showed that in comparison to the strong action-independent Cue 1 – R/NR projection found in ACC neurons (**Figure 6.4**), the LPFC action-specific Cue 1 – R/NR effect was empirically much smaller and slower. The relationship was only seen to come about after approximately 400ms of feedback onset but was maintained beyond the end of the epoch.

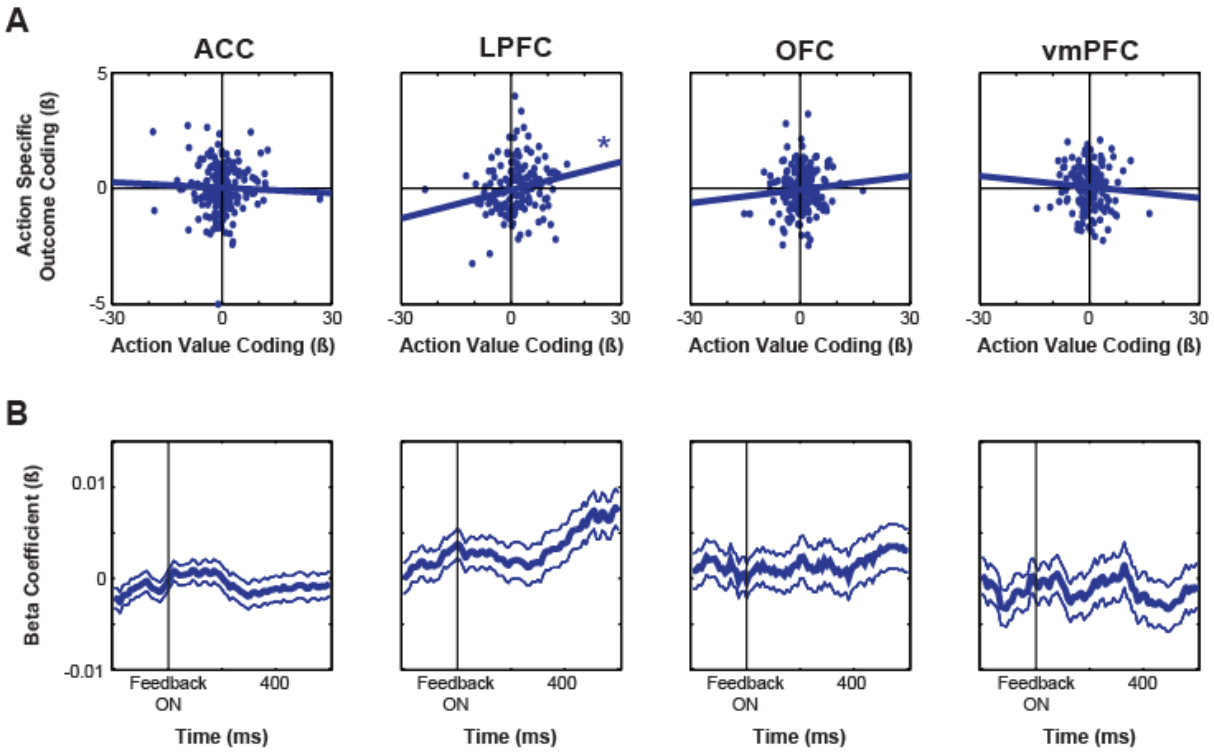

Figure 6.7: LPFC action/spatial value neurons code for action specific outcomes. **(A)** Scatter plots of the mean Cue 1 action value code against the mean action specific R/NR code at Feedback. Asterisks indicate significant correlations (linear regression, $p<0.05$) **(B)** A sliding projection analysis of the mean Cue 1 action value code against the action specific R/NR code across the Feedback epoch. Thin lines indicate SEM.

Figure 6.8 depicts an example LPFC neuron. Although this neuron has a very low firing rate at cue presentation and does not significantly encode the value of stimuli presented either on the left or the right sides of the screen, the left-right contrast term is considered significant (most likely because the neuron has a tendency to code left value negatively and right value positively) (**Figure 6.8A**). During the feedback epoch this neuron clearly differentiates the chosen action even before feedback onset (**Figure 6.8B**). However

after feedback onset the neuron increases its firing rate if the subject was not rewarded. This is particularly pronounced on occasions when the subject made a left choice and was unrewarded.

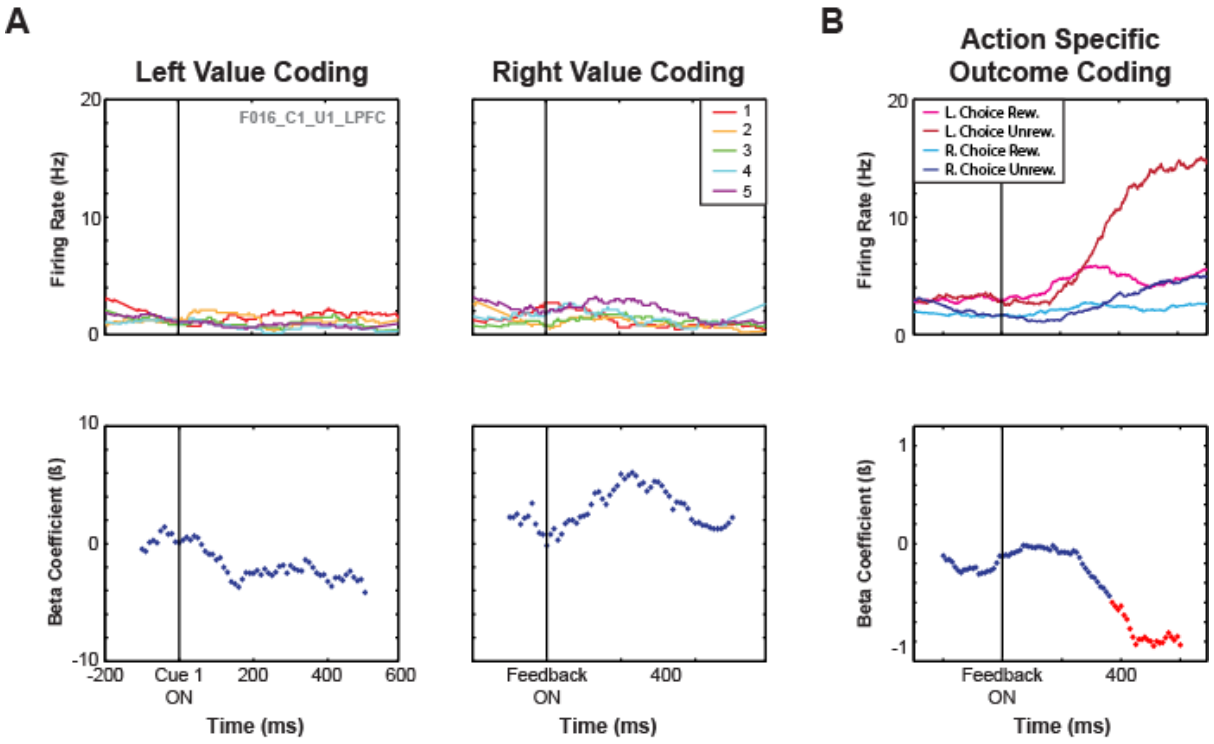

Figure 6.8: An example LPFC action specific outcome coding neuron. **(A)** The firing rate of the neuron at Cue 1 separated by value and cue presentation side. The corresponding beta coefficients are shown in the lower panels. **(B)** The firing rate of the same neuron at Feedback based on the chosen action and the outcome (top) and its corresponding beta coefficients (bottom).

Encoding of Probabilistic Reward Prediction Errors and Value Information at Outcome

The feedback period is an opportunity for the subject to update his estimation of value. Prediction errors are the product of such a computation. Although prediction errors (PE) have been previously reported in PFC neurons (Kennerley et al., 2011), one advantage of our current dataset is the multi-attribute aspect of the choices. Indeed, a critical element of learning is to update specifically what information, in this case, decision attributes, violated expectations (i.e. PE). Because neurons can encode value information both during prediction and during outcomes, it can be difficult to assess whether an outcome signal

reflects value or PE. However, since positive/negative PEs are defined as 1/0-value, respectively, if a neuron codes value at choice (e.g., fires the highest for the high probability stimulus), for its outcome activity to reflect a PE signal it would have to fire the most for the low probability stimulus, a signature evident in both dopamine and some PFC neurons (Kennerley et al., 2011, Tobler et al., 2005). Therefore, we classified PE signals as neurons that encoded value at both choice (i.e., Cue 1) and outcome, but with opposing signed regression coefficients. We define neurons with the same signed regression coefficients at choice and outcome as value coding outcome neurons.

We first sought characterise the number of PE coding neurons in each region. We found that approximately 20% of ACC and 10% of OFC neurons encoded positive PEs (**Figure 6.9**). However, ACC was observed to be significantly greater than all other regions (pairwise Chi^2 test, <0.05). Negative PE encode was empirically smaller in both ACC and OFC although ACC was still significantly greater than LPFC and vmPFC (pairwise Chi^2 test, <0.05). Finally, only 5% of ACC neurons encoded both positive and negative PEs. This was only significantly different from the LPFC population (Chi^2 test, $p<0.01$). From this simple analysis was conclude that ACC (and to some extent OFC) neurons appear to preferentially compute various types of PE at outcome.

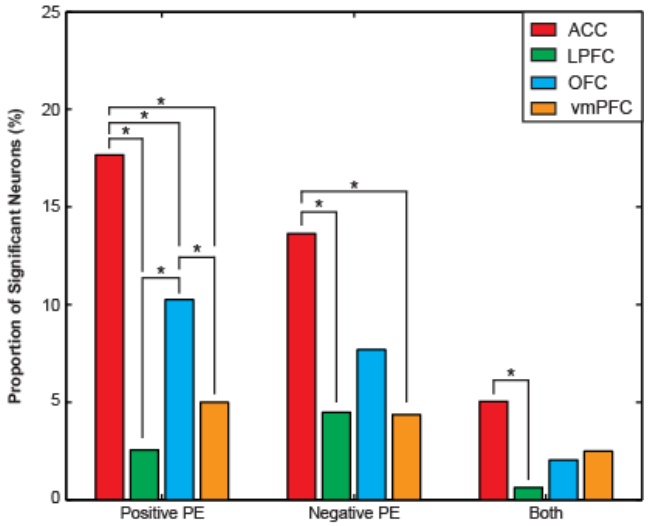

Figure 6.9: The neuronal coding of prediction errors at feedback. All other figure properties are the same as **Figure 6.1**.

We examined the relationship between Cue 1 value coding and chosen probability coding on rewarded trials using a projection analysis. There was a significant *negative* relationship between Cue 1 value coding in ACC and chosen probability in rewarded trials (**Figure 6.10A**) (linear regression, T=-2.01, p<0.03). This implies that neurons ACC neurons

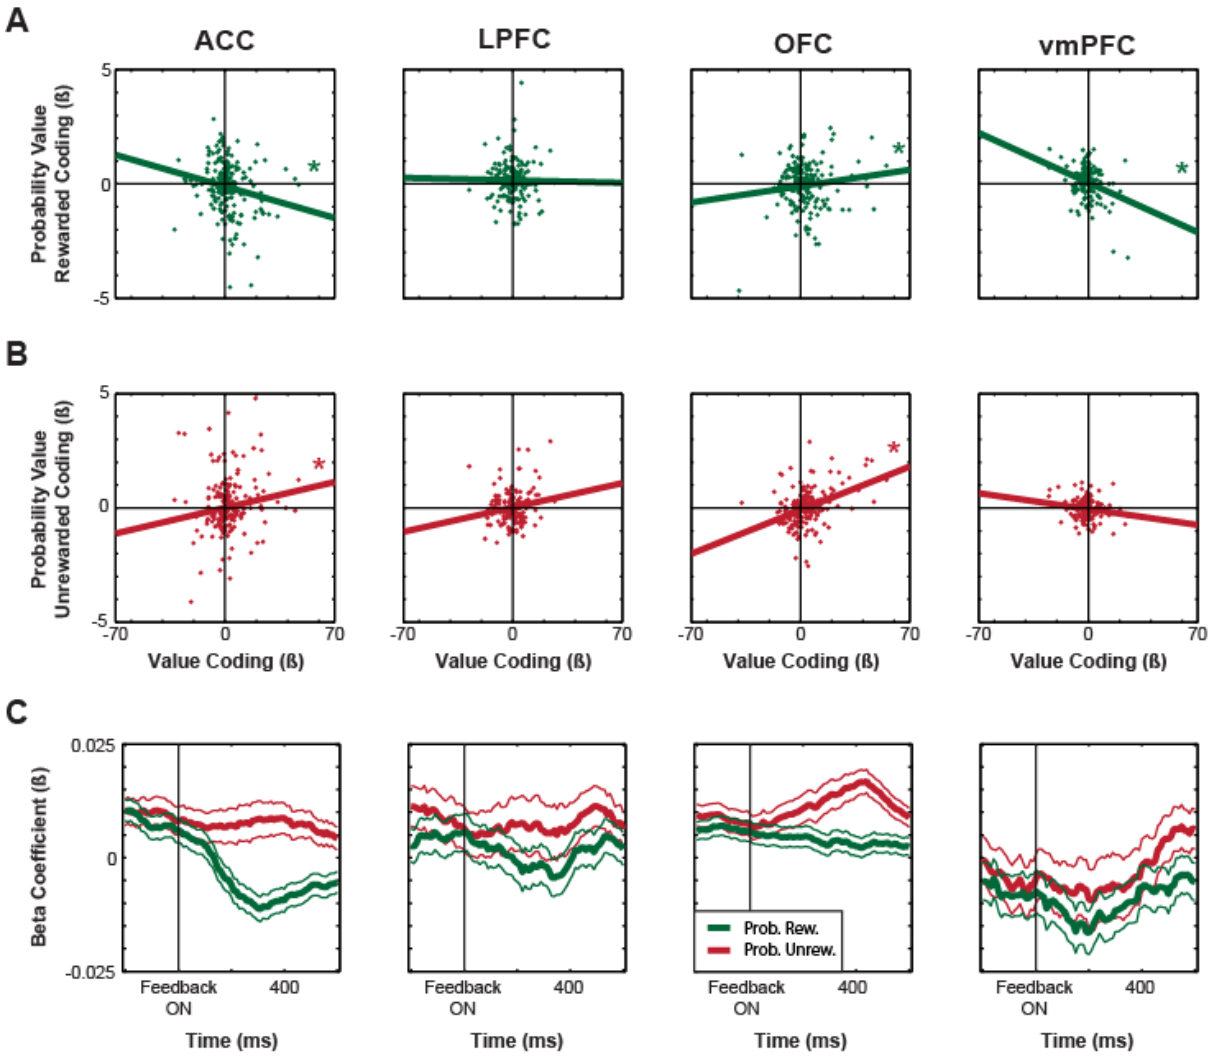

Figure 6.10: The coding of probabilistic prediction errors at feedback. **(A)** Scatter plots showing the relationship between mean Cue 1 value coding and mean probability value coding at feedback on rewarded trials. **(B)** The same scatter plots for unrewarded trials. **(C)** The sliding projection of mean value coding and probability coding on rewarded trials (maroon) and probability coding on unrewarded trials (dark green) over time. All other figure properties are the same as **Figure 6.7**.

that respond more (i.e. positively) to value at the choice then respond more when a low probability choice leads to a reward than a high probability choice (i.e. an unexpected reward over an expected reward). This is consistent with the encoding of a *positive PE* signal. We also noted a significant negative relationship in vmPFC neurons, but a positive relationship in OFC neurons (**Figure 6.10A**) ($T_{(OFC)}$=2.45, $T_{(vmPFC)}$=-2.57, $p_{(OFC)}$<0.02, $p_{(vmPFC)}$<0.02). This suggests that vmPFC may also encode positive PE, while OFC encodes only value information at outcome.

Figure 6.11 shows an example ACC neuron which encodes positive PEs. This neuron responds to value at Cue 1 with a firing rate that increases with increasing value (**Figure 6.11A**). This value coding was significant from approximately 100ms post-cue onset and was maintained over the entire epoch. In the feedback epoch, the same neuron encodes the value of the chosen probability stimulus with a positive relationship up until feedback is initiated; at this point the coefficient for value coding can be seen to flip in sign so that it now encodes the chosen probability with a negative relationship to value approximately 300ms after feedback begins (**Figure 6.11B**, right column), the key signature of a positive PE. However, on rewarded trials, the neuron does not show the same sign flip of regression coefficients, and instead maintains the positively signed representation of chosen probability value until 450ms post-feedback onset (**Figure 6.11B**, right column). This indicates that on unrewarded trials, this neuron encodes that value of the outcome not received, or foregone value.

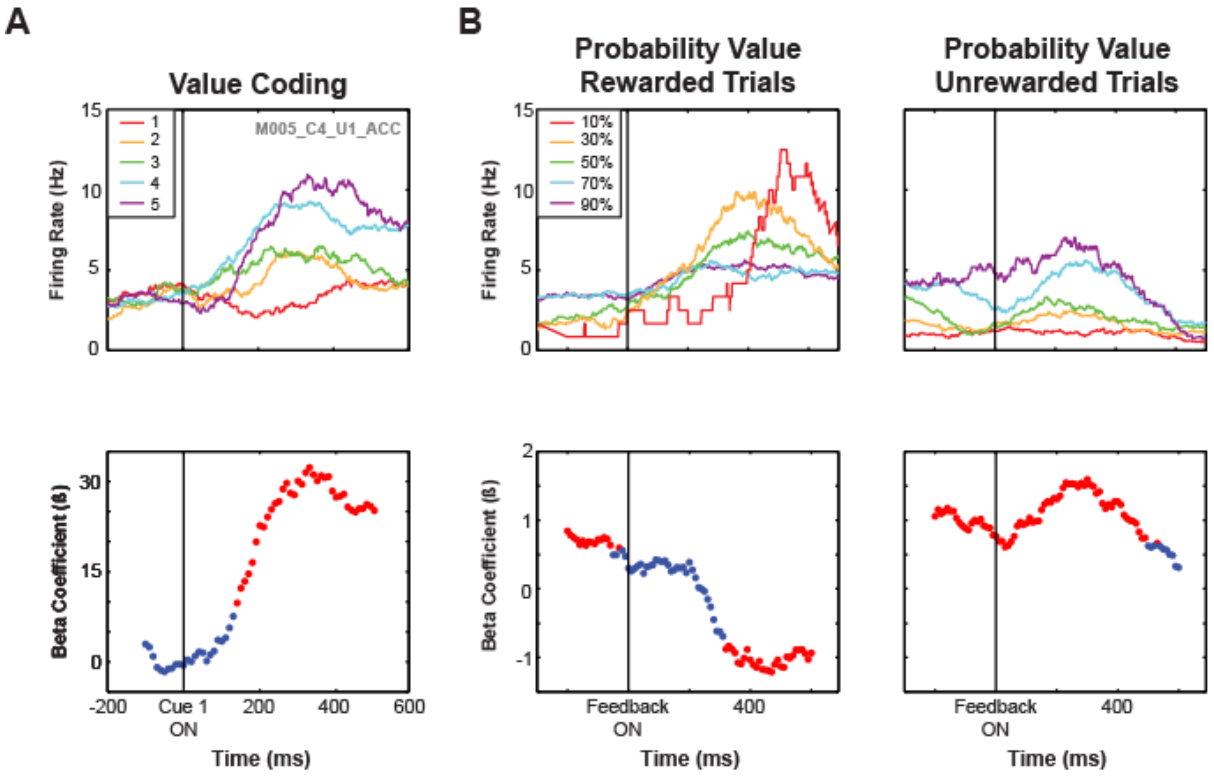

Figure 6.11: An example ACC positive prediction error neuron. **(A)** The average firing rate of the neuron split by the value of Cue 1 (top panel) and the corresponding coefficients for the linear regression of value onto firing rate (bottom panel). Red dots indicate significant bins in the regression and blue dots non-significant bins. **(B)** The firing rate of the same neuron at feedback split by the value of the chosen probability stimulus and the whether the trial was rewarded (left) or unrewarded (right) along corresponding beta coefficients (bottom panels).

Figure 6.12 shows an example vmPFC positive PE coding neuron. At Cue 1 presentation, this neuron codes for value positively at an approximate latency of 400ms (which is very slow when compared to the above ACC example neuron). At feedback the neuron does not maintain a representation of the chosen probability before feedback onset. However, specifically on rewarded trials this neuron negatively encodes chosen probability at a latency of 450ms, in other words, a positive PE signal. In contrast, on unrewarded trials there is no reactivation of the value representation implying a specific positive PE computation.

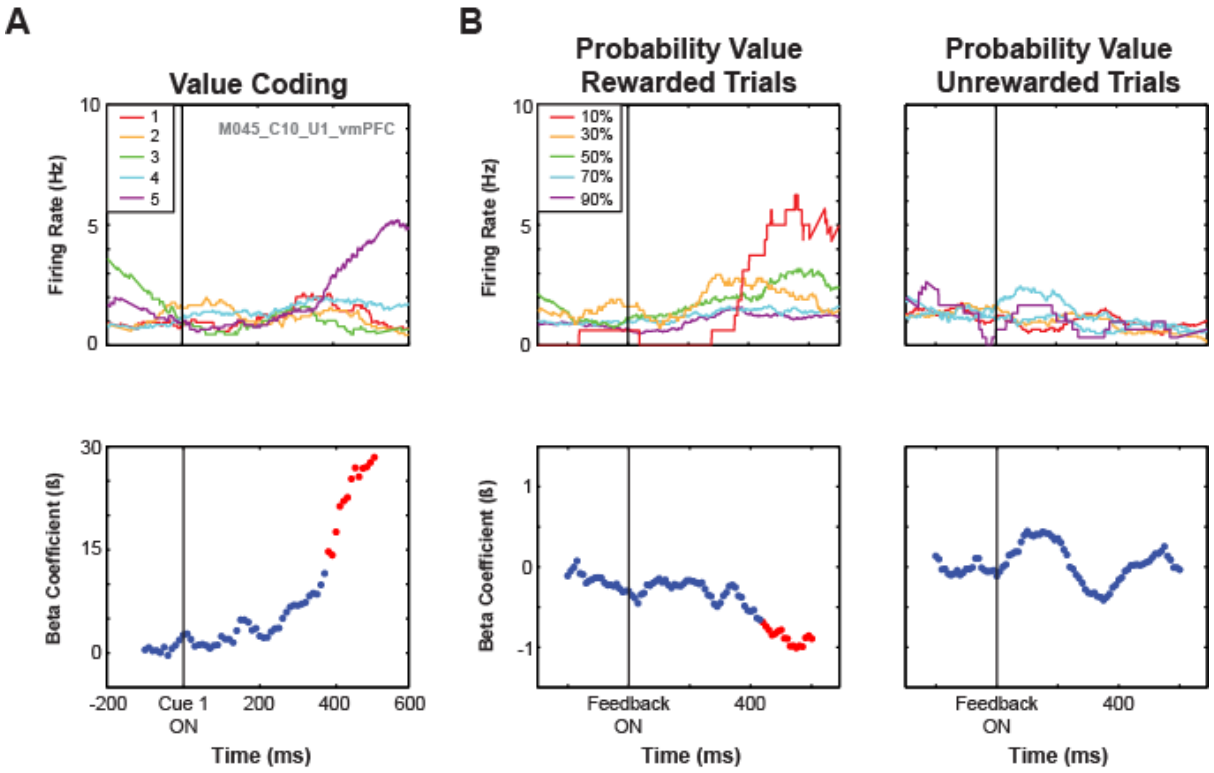

Figure 6.12: An example vmPFC positive prediction error neuron. **(A)** The average firing rate of the neuron split by the value of Cue 1 (top panel) and the corresponding coefficients for the linear regression of value onto firing rate (bottom panel). **(B)** The firing rate of the same neuron at feedback split by the value of the chosen probability stimulus and the whether the trial was rewarded (left) or unrewarded (right) along corresponding beta coefficients (bottom panels). All other figure properties are the same as **Figure 6.11**.

To examine whether PFC neurons had neural signatures resembling negative prediction errors, we performed a projection analysis of Cue 1 value onto chosen probability coding at outcome on unrewarded trials. Unlike on rewarded trials, ACC showed a significant *positive* relationship between Cue 1 value coding and probability value on unrewarded trials (**Fig. 6.10B**; linear regression, $p<0.05$). In other words, if an ACC neuron encodes value positively at choice then it will respond more to unrewarded high probability choices compared to unrewarded low probability choices. This is the opposite, and significantly different, to the neuronal response to rewarded trials (linear hypothesis test of the beta coefficient for ACC in rewarded trials against the corresponding coefficient in unrewarded trials, $T=9.16$, $p<0.003$), and indicates that on unrewarded trials ACC neurons encode how valuable the outcome would have been had it been rewarded, a type of *fictive* or *foregone* value signal. Thus, ACC only encodes positive PEs. OFC also exhibited a *strong* positive

correlation between value coding and unrewarded probability coding (linear regression, $p<2x10^{13}$). This positive correlation in OFC on unrewarded trials was much stronger than its positive relationship in rewarded trials (linear hypothesis test of the OFC beta coefficient for probability value on unrewarded trials against the same coefficient on rewarded trials, $T=16.96$, $p<5x10^{-4}$), and this positive correlation in OFC on unrewarded trials was significantly stronger than all other regions (linear hypothesis test of the beta coefficients for OFC individually tested against the coefficients for all other regions with Bonferroni correction, $p<0.008$). OFC therefore also encodes fictive value when the trial is unrewarded.

LPFC showed a similar positive correlation in this projection, though it did not reach significance (linear regression, $p>0.05$). vmPFC on the other hand, was the only region to exhibit a trend towards a negative correlation between the Cue 1 value coding and probability value on unrewarded trials, the classical signature of a negative PE. While this correlation did not quite reach significance (**Figure 6.10B**; linear regression, $p<0.15$), it also did not significantly differ from the correlation on rewarded trials (linear hypothesis test of the beta coefficient for vmPFC in rewarded trials against the corresponding coefficient in unrewarded trials, $T=0.66$, $p>0.05$), implying the vmPFC population encoded both positive PE and negative PE similarly. Taken together, these results suggest both OFC and ACC may encode similar information about foregone outcomes while vmPFC may encode general PEs.

Figure 6.10C shows how the PE projection evolved across the outcome epoch. We used a sliding projection of mean value coding against probability coding in rewarded and unrewarded trials. The positive PE signal (i.e., a negative correlation in the projection) in both ACC and vmPFC on rewarded trials was very clearly visible as early as 150ms post-feedback onset and peaked at approximately 300ms post-onset indicating a strong signal with a fast latency (**Figure 6.10C**). On unrewarded trials, only vmPFC exhibited a negative correlation in the projection indicative of negative PE coding. In contrast, the strong positive correlation on unrewarded trials indicative of encoding foregone value was particularly robust

in OFC, with the correlation increasing dramatically beginning 100ms post-feedback onset and peaking approximately 450ms post-onset (**Figure 6.10C**).

Figure 6.13 depicts two classes of PE neurons, both in ACC. **Figure 6.13A** shows a neuron which responds negatively to value at Cue 1. When the subject is given a reward, this neuron does not encode the value of the chosen probability stimulus. However, when a reward is omitted, the neuron strongly represents probability value *positively* if the trial is unrewarded. Therefore, this neuron specifically encodes negative PE. **Figure 6.13C** also shows a neuron which responds negatively to value at Cue 1. However, at response this same neuron positively encodes the value of the chosen probability stimulus (irrespective of whether the trial is rewarded or not) (**Figure 6.13B**). This neuron therefore shows a sign flip in coding at response which is indicative of the coding of both positive *and* negative PEs at feedback.

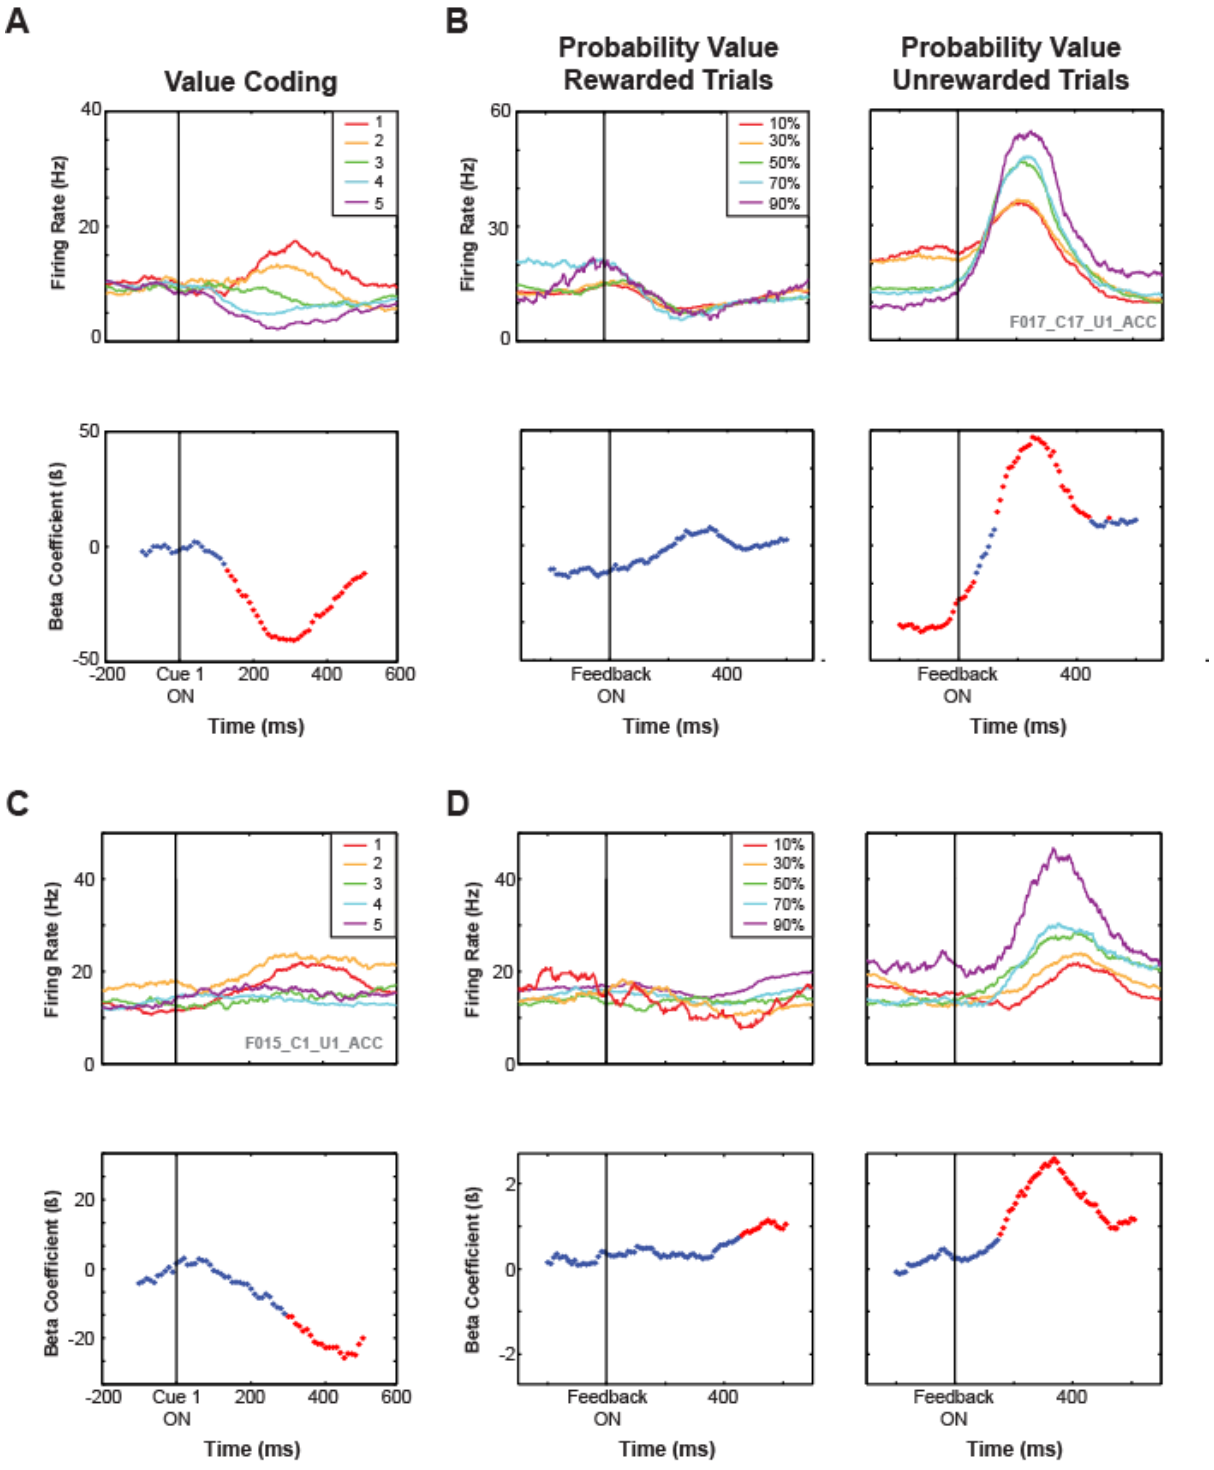

Figure 6.13: Example ACC prediction error coding neurons. The firing rates of each neuron at Cue 1 split by value (irrespective of attribute) (**A** and **C**). The firing rates of the same neurons at feedback split by the value of the chosen probability stimulus and whether the trial was rewarded (left column) or unrewarded (right column) (**B** and **D**). All other figure properties are the same as **Figure 6.11**.

The relationship between reward prediction and reward magnitude at outcome

Having demonstrated that ACC and vmPFC neurons encode PEs in the domain of probabilistic information we next considered whether these regions also computed PEs with respect to magnitude information. Our first step was to probe the relationship between initial value coding (i.e. at Cue 1) and magnitude coding at feedback during rewarded and unrewarded trials. Using the same projection analysis as outlined previously, we found a significant positive correlation between value coding and magnitude coding on rewarded trials only in OFC (**Figure 6.14A**) (linear regression, both $T_{(OFC)}=3.78$, $p_{(OFC)}<2x10^{-4}$). There was also a borderline significantly negative relationship on rewarded trials in vmPFC ($T_{(vmPFC)}=-1.94$, $p=0.0504$), which was similar to the negative relationship observed for probability coding on rewarded trials. On unrewarded trials, both ACC and OFC had a significant positive correlation (**Figure 6.14B**) (linear regression, both $T>2.98$, both $p<0.003$). None of the reward magnitude projections in any of the areas differed between rewarded and unrewarded trials (linear hypothesis test of the beta coefficient for magnitude coding in rewarded trials against that for unrewarded trials, $p>0.05$).

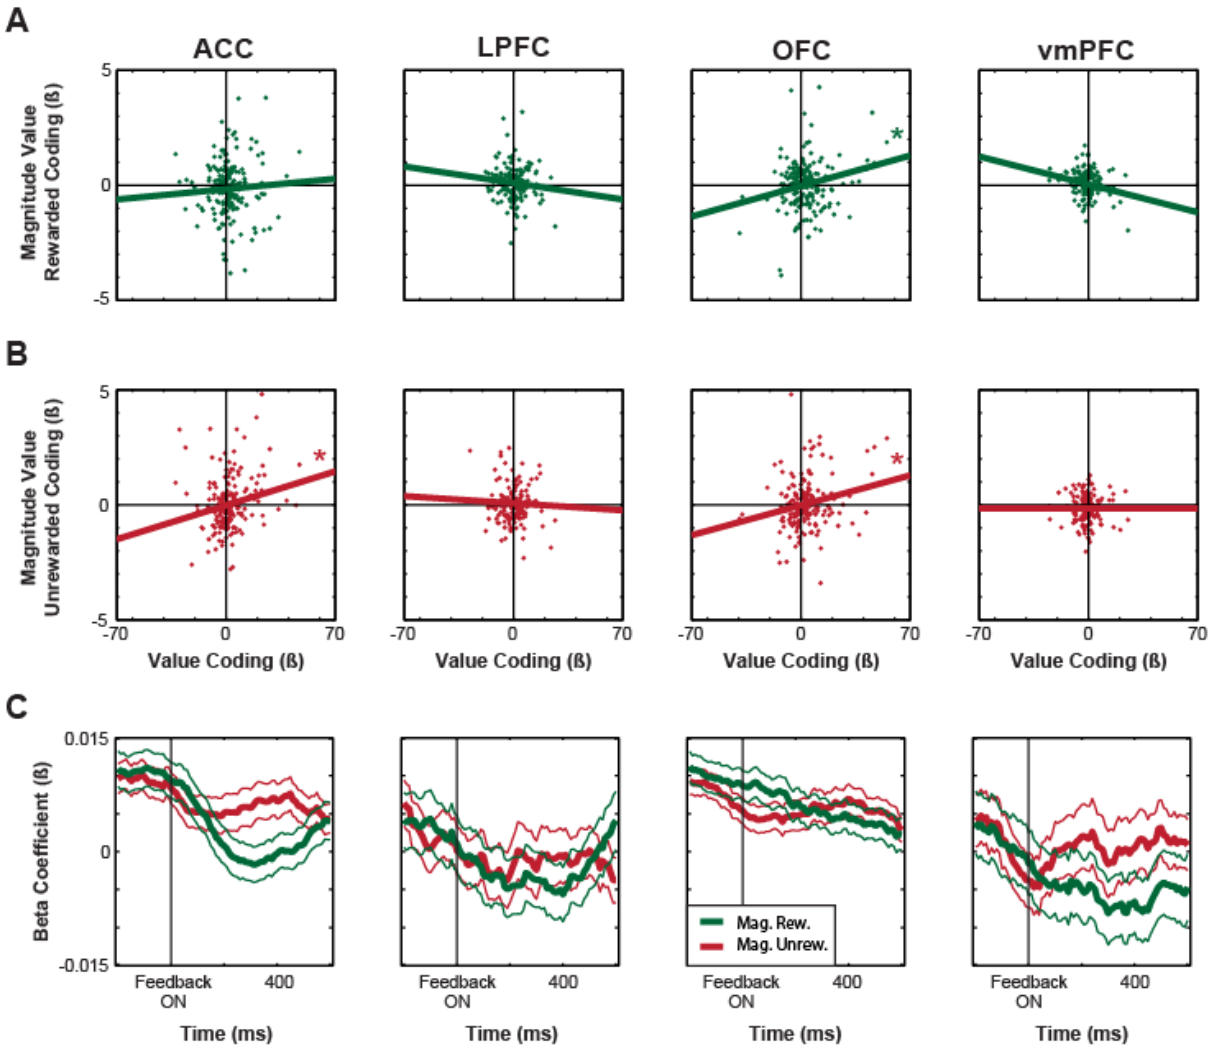

Figure 6.14: The coding of reward magnitude expectation at feedback. **(A)** Scatter plots showing the relationship between mean Cue 1 value coding and mean chosen magnitude value coding at feedback on rewarded trials. **(B)** The same scatter plots for unrewarded trials. **(C)** The sliding projection of mean value coding and magnitude coding on rewarded trials (maroon) and probability coding on unrewarded trials (dark green) over time. All other figure properties are the same as **Figure 6.10**.

Figure 6.14C shows the temporal evolution of the reward magnitude projection across the outcome epoch, using the same sliding projection analysis as outlined earlier for probability coding. In ACC this analysis revealed that approximately 200-500ms after feedback onset there was a qualitative difference between the two trial types, attributed to a quenching of the magnitude projection on rewarded trials (**Figure 6.14C**). In fact when the mean signal projection was constrained only to this time period there was a significant difference between reward and no reward (linear hypothesis test of the beta coefficient for

ACC in rewarded trials against unrewarded trials, $p<0.04$). This implies that when the subject is aware that an outcome will be rewarded, the ACC population stops encoding reward information. In contrast, when the subject is made aware that they will not receive a reward, the ACC population continues to encode the *foregone* reward magnitude. The OFC population exhibited positive projections on both rewarded and unrewarded trials which indicates that this region encodes the chosen reward magnitude at feedback irrespective of whether this is an experienced or foregone reward.

Having found that outcome coding in LPFC is modulated by action, one might expect that positive and negative prediction error codes are also frames by action. We therefore repeated the probability and magnitude (rewarded and unrewarded) value projection analyses outline above instead using the action specific contrasts defined from the regression model laid out in **Table 6.2**. However, we found that there was no significant correlation between Cue 1 action value coding and the coding of chosen probability contingent on the chosen action on either rewarded or unrewarded trials (linear regression, all $p>0.05$). This was also true in the magnitude domain (linear regression, $p>0.05$).

Discussion

R/NR Coding in PFC

We have shown that many neurons in ACC, LPFC, OFC and vmPFC discriminate whether or not a reward is delivered following a choice. This finding is in agreement with many studies that have also reported outcome related coding in these regions (Kennerley et al., 2011, Monosov and Hikosaka, 2012, Ito et al., 2003, Quilodran et al., 2008, Seo et al., 2007, Seo and Lee, 2007). In agreement with Kennerley et al. (2011) we found that ACC neurons were the single largest outcome sensitive population. However, unlike Kennerley et al. (2011) who found no skew in the coding pattern of R/NR, we found that this population was significantly skewed towards negative coding.

ACC has been implicated in monitoring behavioural errors in several human tasks (Critchley et al., 2005, Paulus et al., 2002, Holroyd et al., 2004). The ACC effect in the current study does not necessarily reflect errors in optimal choice because there only a weak correlation between selecting the best option and getting a rewarded outcome. Furthermore, the feedback period is not the point at which subjects would become aware of whether they made an error in selecting the best option, instead this would have been immediately after the response was made when all of the stimuli were uncovered allowing the subject to compare what he chosen against the unchosen option. Some studies define errors as failures to inhibit unwanted actions (Stemmer et al., 2004, Critchley et al., 2005). Although it is possible that there may be occasions when subjects choose suboptimally because they fail to inhibit a response bias, analysis of the choice behaviour (Chapter 4) shows that behaviour is strongly driven by value and that these types of errors are relatively uncommon. Based on these factors, it is therefore unlikely that the ACC neurons in this study were encoding behavioural errors, and instead likely coding information about the outcome received based on the choice that was made.

In this study we defined outcome coding as a differential response to a rewarded and unrewarded outcomes (R/NR). A study by Seo et al. (2007) where they analysed coding of wins and losses (which are slightly different to the definitions used in this study) the authors found that different populations of neurons in LPFC and ACC coded for wins as compared to losses. The method of analysis in this study does not necessarily conflict with this finding because a neuron that only responds to a reward (or no reward) can also appear to be significant for the differential (contrast) term if the strength of this coding is strong enough. Seo and Lee (2007) also found that the coding of a gain in ACC correlated with the mean value of the trial (called assets in their paradigm) but the same was not true for losses. However, this analysis was performed at feedback whereas our study shows that ACC neurons that encode value during the decision phase are specifically modulated later in the trial by outcomes.

The fact that OFC probability preferring neurons are more likely to represent the outcome is unsurprising as the outcome is critical for updating the value of the probability stimuli and should have no bearing on the value of magnitude stimuli. Lesion studies have shown that damage to parts of OFC can lead to deficits in credit assignment (Noonan et al., 2010). One possible explanation for this could be that lesions to OFC cause a loss in probability specific value neurons, or neurons which differentiate whether choices are rewarded, both of which are capable of tracking outcomes. Interestingly, in a previous study that compared outcome coding across PFC neurons for different decision attributes, OFC neurons were most sensitive to whether or not a choice was rewarded compared to the magnitude of the reward or the physical cost incurred to obtain the reward (Kennerley and Wallis, 2009b). We also found a positive correlation between LFPC magnitude value preference at Cue 1 and R/NR coding. This implies that the more an LPFC neuron prefers encoding magnitude over probability information at choice, the more they will differentiate whether or not that choice is rewarded at feedback.

Finally, the current study found that action value LPFC neurons at Cue 1 also encode R/NR contingent on the action performed to achieve the outcome. Neurons is SEF show a similar pattern of action specific outcome coding (Uchida et al., 2007). However, this finding is novel in the literature for LPFC, although it has been previously shown that LPFC neurons represent previous actions (several trials into the past) and current outcomes simultaneously during feedback (Barraclough et al., 2004). Such a signal may therefore be critical for specifically tracking the value of actions. It should be noted that the action outcome signal may still reflect a spatial value signal. This is because when feedback is provided, visual feedback is presented informing the subject as to whether he is being given reward or not. It is therefore possible that these LPFC neurons are actually responding to whether or not reward is presented on their preferred side of the screen.

Prediction Error Coding in PFC

This study has shown that two separate PFC regions compute PEs at feedback. The ACC population strongly encode positive PEs but show no coding of negative PEs. In contrast, vmPFC neurons encoded both positive PE and negative PE with similar negative relationships between Cue 1 value coding and the probability coding at outcome, though the negative PE coding on unrewarded trials did not quite reach significance. Several studies have shown that ACC neurons can encode either positive or negative PE (Matsumoto et al., 2007, Silvetti et al., 2014, Kennerley et al., 2011). However, Kennerley et al. (2011) found that ACC had a predilection for positive PE coding, consistent with our current ACC findings. Matsumoto et al. (2007) found that ACC neurons encoded prediction errors during motor learning, however, the computations themselves were not action specific which is in agreement with the current study. It should be noted that all regions contained some neurons that computed both types of PE. However, at the population level it was only ACC and vmPFC that encoded various forms of PE.

The finding that ACC neurons encode both positive PEs is consistent with the contemporary ERN data from the human EEG field (Silvetti et al., 2014). The ERN is seen to occur 100-200ms post-feedback onset which is the approximately timing of the PE signal in this study. This result is therefore a potential link between the small scale neuronal computations measured in primate studies and the large scale signals seen in EEG studies. However, one discrepancy is that in the current study the PE signal was observed to last beyond 700ms whereas the ERN seems to have a much shorter latency (approximately 300ms). The ERN may therefore reflect a shorter process such as synaptic activity rather than directly measuring neuronal firing. Dopamine neurons are known to encode predictions about the value of stimuli as well as both positive and negative PEs in a common value currency (Tobler et al., 2005). Like ACC, dopamine neurons have also been shown to prefer to code for positive PEs (Bayer and Glimcher, 2005). Given the strong anatomical connections between ACC and VTA it is no surprise that ACC and dopamine computations share some similarities (Williams and Goldman-Rakic, 1998).

Prediction error signals have not been reported in vmPFC in previous studies. In fact Monosov and Hikosaka (2012) found no evidence of PE coding in vmPFC neurons in their study. However, there are two important considerations to bear in mind. First, this study used a simple pavlovian task which lacked any choice component, thus the outcome information could only be used to update expectations rather than adapt future choice behaviour. Second, the authors did find that anterior parts of vmPFC were more responsive to value and outcome compared to more posterior portions. In the current study we recorded from this anterior portion (Brodmann's area 14) and also found value and outcome selective neurons. Again the use of a projection analysis may explain why we were able to detect a PE signal. It should however be noted that the positive PE signal in vmPFC was the weakest of the significant PE projection signals reported here. However, it was also not significantly different to the signal measured on unrewarded trials, suggesting vmPFC may code deviations from expectations on both rewarded and unrewarded outcomes. Therefore,

vmPFC may share more similarities with the dopamine signal than ACC. As with most vmPFC signals reported throughout this thesis, the vmPFC PE signals were much weaker than other value-related computations in ACC and OFC.

Studies have shown that some OFC neurons encode both positive and negative PEs (Sul et al., 2010, Kennerley et al., 2011, Noonan et al., 2011), though this encoding appears to quantitatively less prevalent in OFC compared to ACC (Kennerley et al., 2011). We found that few OFC neurons compute PEs and that these signals are not represented at all on the population level. Instead, of all the PFC areas sampled, arguably OFC exhibited the strongest positive correlation between prediction (i.e., Cue 1) and outcome coding, implying one of the key representations in OFC is to encode the value of outcomes, rather than their deviations from expectations.

We found in OFC that the Cue 1-outcome probability projection was significantly more positive on unrewarded than rewarded trials. ACC also exhibited a positive Cue 1-outcome probability projection on unrewarded trials. This implies that when outcomes are not rewarded, rather than coding negative PEs, neurons in both ACC and OFC encode information about the reward probability in a similar way as those neurons encoded Cue 1 value. Neural signatures of fictive or hypothetical outcomes (i.e., outcomes for choice not made) have been observed in ACC, OFC and LPFC (Abe et al., 2011, Hayden et al., 2009). Our results reflect value coding for choices made, but not experienced, or *fictive* value. Such fictive value signals are known to influence decision making behaviour (Chiu et al., 2008, Montague et al., 2006, Kim et al., 2015, McClure et al., 2003). Data from the current study provides evidence of a subtly different form of fictive value to that which is commonly reported in the neurophysiology literature. Whereas most studies show that PFC regions encode the value of options that should have been chosen at the moment of feedback (Abe and Lee, 2011, Hayden et al., 2011) we have shown that ACC and OFC neurons compute the value of the chosen option when it is not rewarded. Fictive reward encoding is difficult to separate from motivational coding. Such signals are observed across the frontal cortex

(Roesch and Olson, 2003). However the fact that value coding is seen to remerge at outcome and the fact that value coding changes based on whether the outcome is rewarded or not implies that this is unlikely to be a mere motivational signal.

Functional Subdivision of Prediction Error Coding

One of the major advantages of considering the outcome period in this behavioural paradigm is that outcome signals can be decomposed into ones that relate to the probabilistic and the reward size aspects of the choice. In this study we have shown that OFC neurons do not encode reward magnitude PEs but instead represent the chosen magnitude value at feedback onset on both rewarded and unrewarded trials. ACC neurons represent the chosen magnitude only when it is not delivered. Both these results on unrewarded trials are consistent with ACC and OFC neuronal populations encoding information about the reward magnitude in a similar way as those neurons encoded Cue 1 value; such foregone value coding about reward magnitude is remarkably similar to the foregone probability value discussed above.

The fact that we did not find negative projections between Cue 1 value and reward magnitude at outcomes may have been expected. From the point of view of learning it seems more optimal to code prediction errors solely for probability stimuli because unlike magnitude stimuli they are harder to learn the value of and have high variance in outcomes. Furthermore, there is very little variance associated with the outcomes of the magnitude stimuli because they were clearly separated in size and also have visual feedback associated with them in order to make their exact timing as clear as possible for the subjects. Such signals may be critical for ongoing credit assignment for probability stimuli during task (Noonan et al., 2011, Noonan et al., 2010).

In conclusion this chapter has presented evidence that specific subsets of ACC, LPFC and OFC neurons encode reward outcomes. ACC and vmPFC compute various types

of PE while OFC and ACC also represent fictive value. These complementary signals may be critical facilitating adaptive behaviour during decision making.

Chapter 7: Discussion

This thesis used both behaviour and single neuron physiology to make inferences about the layout and computations of decision making circuits which guide both overt attention and final choice. The main findings were as follows: 1) covert attentional systems use value guided decision making processes to bias overt attention, 2) neurons in PFC represent and compare values in different references frames consistent with a parallel decision making system, 3) future information gathering is also represented in the same reference frames in PFC, 4) there is functional specialisation within PFC regions reflecting different roles they play in optimising behaviour.

Covert Attention and Decision Making

This thesis examined the influence of covert attention in two separate but similar behavioural paradigms presented in Chapters 3 and 4. In both experiments subjects used covert attention to both decide where to make their first saccade and to decide what to eventually choose. The important difference between the two tasks was the fact the multi-attribute task was more complex to solve because subjects had to consider more information. However, despite this added complexity covert attentional behaviour was mostly unchanged. The use of covert attention may have several empirical benefits to subjects in solving the types of task presented in this thesis. Firstly, it drives overt attention towards more valuable options which may help overt attentional to bias subjects to choose these options (Krajbich et al., 2012) thereby increasing the probability of making optimal decisions. Secondly, covert attention almost certainly reduces the time required to sample available information overtly since overt saccades most likely require more planning and execution

time than shifting covert attention. Therefore the use of covert attention may increase subjects' reward rates.

Several behavioural differences between covert control of overt attention and final choice were observed in both experiments. Chief among these is the fact that novelty differentially influences saccades and choice. One potential inference of this finding is that these two value guided systems are separable. If this is true then this implies that neuronal circuitry in several regions of the brain is configured to resolve value comparison problems. This is not surprising given the fact that recordings from multiple regions during perceptual decision making have shown comparison signals (Siegel et al., 2015, Hernández et al., 2010) and also fits with the idea of parallel value comparison (Cisek, 2012).

Because the experiments presented in Chapter 3 and 4 are entirely behavioural it is impossible to state for certain that basal ganglia circuits do indeed perform this fast latency value comparison driven covert attention. However, as stated in Chapter 3, given the anatomical and neurophysiological information known about caudate and Substantia Nigra par reticulata this inference appears to be safe. The obvious follow up question to ask is whether the basal ganglia circuit interacts with the prefrontal circuit that most likely drives final choice? From an anatomical point of view is it plausible because loops exist between caudate PFC regions including OFC and LPFC (Alexander et al., 1986). However, without neurophysiological recording this remains an unsupported hypothesis.

Frames of Reference in the PFC

In Chapter 5, the finding was reported that different subpopulations of neurons encode value in specific frames of reference is not surprising and is in fact a confirmation of several previous studies which have individually each individually reported single frames of reference in single areas in different tasks (Hayden and Platt, 2010, Roesch et al., 2006,

O'Neill and Schultz, 2010, Matsumoto et al., 2007). However the major advance in this study is that we have shown that these frames of references are *simultaneously* represented in *specific* populations of PFC neurons. This finding along with the fact that future information gathering signals are also represented in the same reference frame implies that frames of reference represent either a more general feature of cognitive functions that can be driven by value or a common value comparison system for choices and information gathering in PFC.

fMRI BOLD activity in human vmPFC has been seen to correlate strongly with value in decision related computations in a huge number of studies (Bartra et al., 2013) and has also been named as part of the 'default mode network' of the brain (Laird et al., 2011). Therefore given the relatively strong signal measured in humans it is surprising that neuronal recordings from this region have shown relatively few computations occurring in vmPFC and also that any computations vmPFC does perform do not appear to be region specific and are also far weaker than other PFC regions. Furthermore, given the demonstration of attentional value computations in the vmPFC BOLD signal in the study of Lim et al. (2011) it is surprising that we find attentional value computations specifically in primate OFC. There are several possible explanations for this finding. Firstly, the BOLD signal in fMRI only a proxy for neuronal activity and can be modulated by several physiological factors including large veins several centimetres downstream of the actual activity (Kim and Ogawa, 2012, Arthurs and Boniface, 2002). Although it is possible that this vmPFC BOLD signal is therefore not accurately reflecting the exact location or amount of activity related to decisions this is unlikely to account for the mismatch between fMRI and the current study. Another potential answer to this question comes from the fact that BOLD signal from OFC is difficult to obtain because of signal dropout caused by the air sinuses of the cranium (Weiskopf et al., 2007). Therefore it is possible that the OFC activity may in actual fact be more modulated by attention than it appears. This mismatch in vmPFC may come about from fundamental differences in the task designs between humans and primates. Whereas humans are often

given little training and exposures to stimuli before tasks are initiated, primates are often heavily overtrained on both the task structure and the stimuli. Despite the fact that subject in this study were relatively less trained on the stimuli compared to other decision making task (for example Kennerley et al. (2009)) they still received more exposure than humans. This may lead to changes in the circuitry which deal with the resolving the decision thereby causing vmPFC to become inactive during the decision phase. Finally it is possible that human and primate vmPFC are fundamentally different in the computations that they perform. However, the fact that resting state MRI has shown connective homology between the two species (Neubert et al., 2015) indicates that this interpretation is unlikely.

What is the point of these frames of reference? The only frame of reference that could be considered suitable for solving all decisions is that of 'good'. However, as outlined in Cisek (2012) it is clear that some decisions (such as those intimately tied to actions) are more intuitively made in frames of reference other than 'goods'. Therefore, by representing many frames of reference it may allow PFC to be flexible enough to support all possible types of value based decisions an animal might face.

Parallel Value Comparison in the PFC

In the current study value coding in all three frames of a reference evolve over time to reflect that of choice (each in different ways). This implies that these reference frame specific computations do not simply reflect inputs into a common value comparison but instead reflect the *simultaneous* value comparison in line with a parallel computations predicted by the 'decision making through consensus' idea (Cisek, 2007b). It should be noted that in the current study we only present explicit evidence of value comparison in the frame of reference of attention, the evolution of all of the signals from value to choice heavily suggests some form of competitive process (Machens et al., 2005, Wang, 2008). The idea of

parallel value comparison is an appealing concept when considering how flexible behaviour might be achieved. This is because this model allows for the value of actions to constantly be compared (even while value comparison occurs in other frames of reference) until a decision is reached by consensus (Cisek, 2012). In the real world this would allow deciders to delay decisions for a long as possible when flexible behaviour is required. In contrast in the 'goods based model' the decider would have to make their abstract decision slightly earlier in order to have enough time to convert the decision to an action.

Real world decisions can have several outputs such as limb movements, saccades or even speech. Although it is conceivable that a single abstract decision making circuit could make decisions and then map to all of these output systems it is tempting to imagine that if the cortical areas associate with these outputs contained decision making circuitry then many different types (in terms of the required outputs) of decision could be solved more efficiently.

Furthermore, the task presented in the current study has several plausible ways in which the decision could be resolved. These include performing an abstract comparison of the available option values, comparing the values of available actions and making decisions based on flexible biases towards certain attributes. Although we have no way of measuring how subjects resolve each decision the parallel computation hypothesis provides a clear framework for how all methods may be flexibly employed by subjects if they choose to.

Learning From Outcomes in PFC

As discussed in Chapter 1 many studies have demonstrated that PFC neurons respond to the presence or absence of a reward and this has been essentially replicated in the current study. However, this study has taken this finding one step further by showing that specifically in ACC, the neurons that discriminate outcomes are neurons that code value

during the choice phase of a decision. The design of the task used in Chapter 6 has the added advantage that choices are made based on *two* attributes. Therefore, it is possible to make inferences about whether information outcome related computations are specific to certain attributes or based on a 'common currency'. ACC neurons encode only positive prediction errors (at the population level). This conclusion is further strengthened by the comparison to OFC and vmPFC which are both only encode probabilistic prediction errors. One completely novel finding is that vmPFC neuron may encode both positive and negative prediction errors (although our data only shows a trend towards the latter). Therefore, vmPFC neurons may show more similarities to the responses of midbrain dopamine neurons than ACC neurons. It has been suggested that vmPFC is less involved with stimulus learning than more lateral parts of OFC (Noonan et al., 2010) therefore it is more likely that both ACC and vmPFC signals may reflect an aspect of behavioural control where predictions about the environment and outcomes are continually monitored (Alexander and Brown, 2011).

The task used in the current study should in theory involve little learning because the subjects are exposed to the same stimuli in the paradigm described in Chapter 3 on the day preceding data acquisition. Optimality in choices between stimuli of a specific attribute is observed to be high in the results described in Chapter 3. As might be expected, subjects are slightly more optimal at magnitude choices than probabilistic choices. Because the probabilistic outcomes are associated with a relatively higher degree of uncertainty it is perhaps unsurprising that PFC neurons encode prediction errors about these stimuli. Importantly, we found that probability value preferring neurons in OFC were more likely to encode whether reward was given or not. OFC is thought to be critical for learning and updating stimulus-outcome relationships and therefore these neurons pay be particularly relevant for updating the value of probability stimuli over the course of the session. (Noonan et al., 2010, O'Doherty et al., 2003, Tremblay and Schultz, 2000a, Schoenbaum et al., 1998, Camille et al., 2011b, Rudebeck et al., 2008).

Another particularly novel aspect of the results presented in this thesis is that many of the outcome monitoring and potential learning signals are specific to frames of reference. This might be expected because outcomes can be monitored in different frames of reference; the abstract outcome it's (irrespective of what occurs to bring it about), the action made to achieve the outcome and the specific properties of the chosen outcome. It therefore follows that frame of reference specific value neurons should monitor outcomes in the same frames of reference in order to optimise their future computations.

Final Thoughts: An Integrated View of PFC Computations During Value Based Decision Making

Given the large breath of both neuronal and behavioural findings reported in this thesis it is prudent to give a brief summary of the both signals and how they integrate into the wider field of decision making. **Figure 7.1** provides a basic overview of computations that each region performs in the tasks presented in this thesis. It also sets these results out in the context of a simplified anatomical framework. What is striking about this set of results is the clear implication that there are various decision systems in the brain and that even through their anatomical connections they exhibit a clear degree of parallelism. At the level of coarse anatomical regions this parallelism stems from the fact that many regions can simultaneously receive varying degrees of basic sensory, complex sensory or limbic inputs which are all useful for making value based computations. Furthermore, many of these regions widely interconnect with each other and also have connections to various motor and saccadic systems.

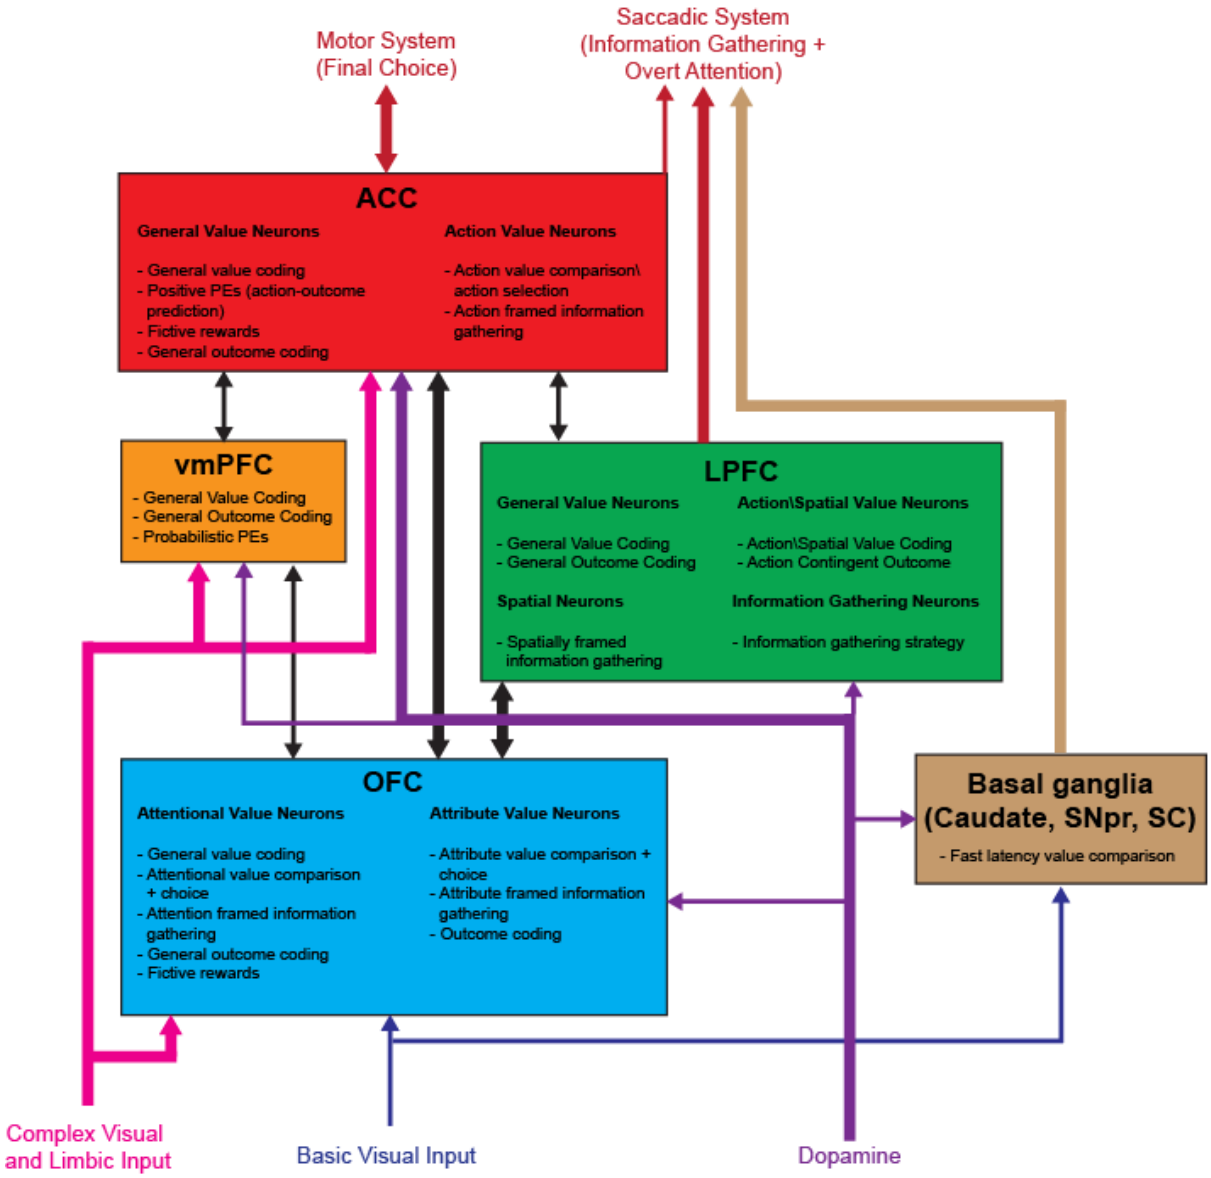

Figure 7.1: A summary of the important computations inferred in this thesis. A simplified outline of the main results of this thesis in the context of the simplified relevant anatomical connections. The thickness of the lines indicates both the strength and importance of the connections. For areas that have specific value subpopulations computations have been attributed to the relevant population. The extra-frontal systems and inputs shown are each composed of multiple regions and inputs and this figure does not attempt to make reference to connective differences between these regions and PFC. For example, the 'saccadic system' refers to all regions that perform saccadic control and therefore information from ACC, LPFC and Basal ganglia may well be passed on to different structures. Abbreviations: SNpr, Substantia Nigra pars reticulata; SC, Superior Colliculus. All other abbreviations are the same as previously stated.

It is also possible to use the data presented in this thesis in the context of the wider decision making literature to form a broad summary of the computations that take place

during simple multi-attribute value based decision making. When decision information is presented it is represented in an attribute specific manner in OFC neurons. These attribute specific neurons then feed into other neurons which integrate attributes to encode option values (in the frame of reference of attention) but importantly they also undergo a competitive value comparison process (most likely through mutual inhibition). The option value population pools also undergo competitive processes and also feed information into the action system which compares the value of the two available actions (taking into account effort). Attribute and attention level comparisons can interact and furthermore attention and action level comparisons can interact. Decisions are resolved through a consensus between competitive processes which results in a movement. These frames of reference specific computations are used for information gathering. The chosen value from the decision process is then represented for the purposes of prediction. When outcomes are presented some neurons may represent action-contingent outcome computations for the monitoring of actions. ACC and vmPFC neurons represent various forms of prediction errors which are used for behavioural monitoring and probability preferring OFC neurons represent whether reward is delivered which can be used for updating the value of the internal representations of probabilistic value stimuli. Therefore, to conclude, this thesis provides evidence of multiple types of decision related computations which add to our understanding of the cognitive and mechanistic processes that drive value based decision making.

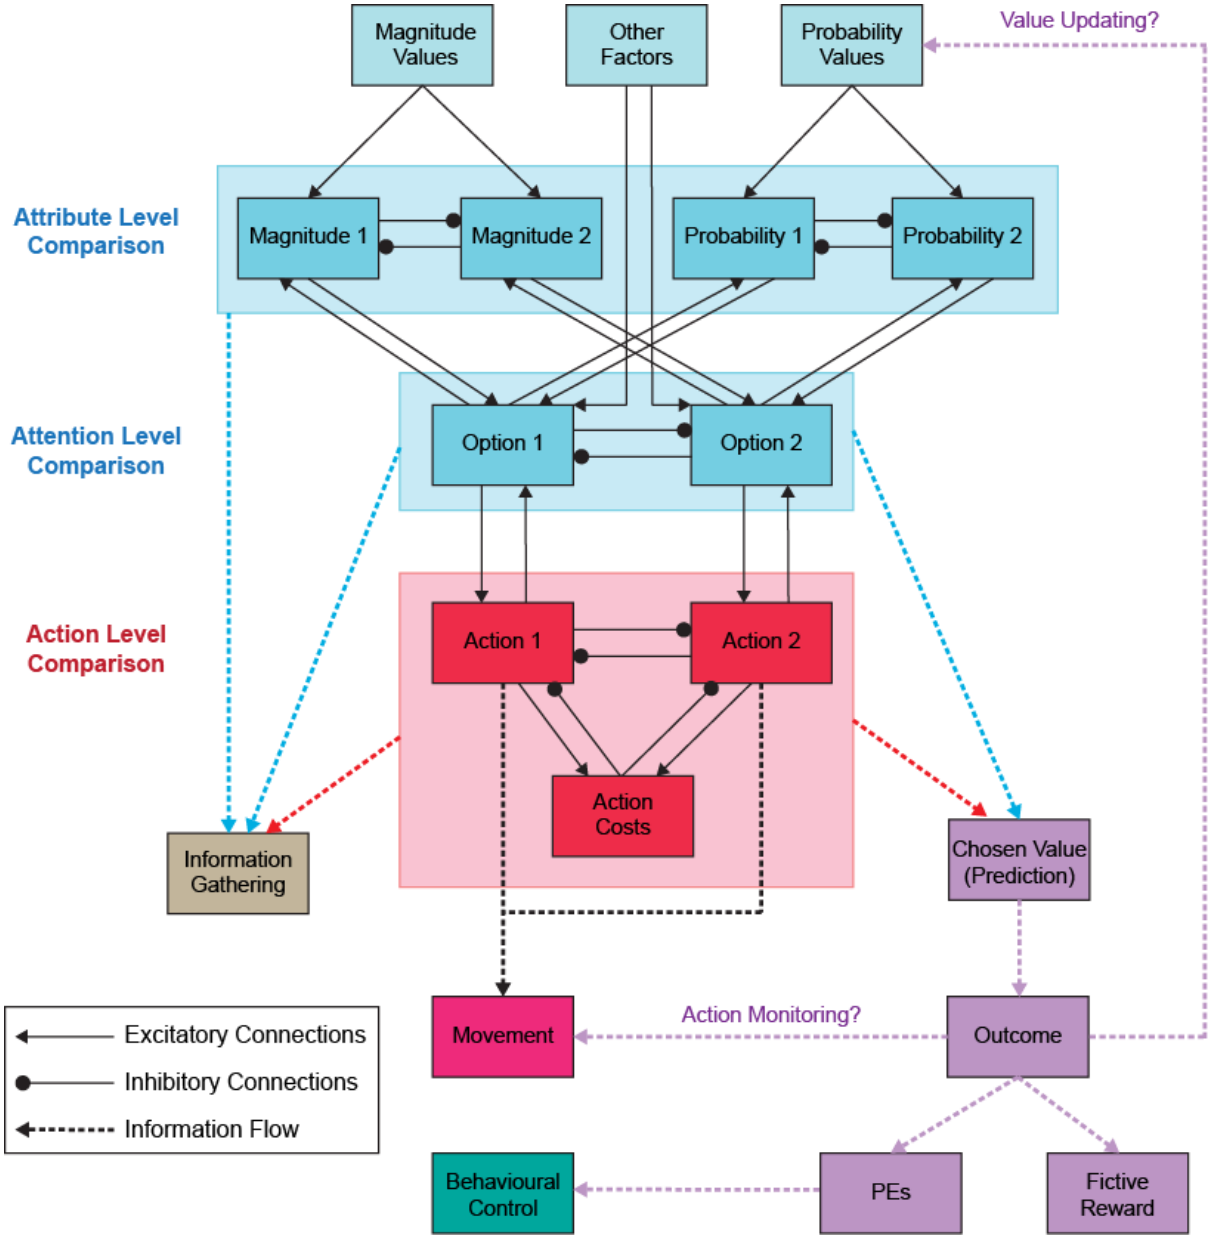

Figure 7.2: The computations and processes involved in multi-attribute decision making. Solid lines indicate mechanistic interactions between pools of neurons representing various decision properties (boxes) during decision making based on the findings of this thesis and the literature as a whole. Dashed lines indicate the non-mechanistic evolution of computations for other decision related processes.

Appendix

This section contains a supplementary analysis performed on the behavioural data for ‘Information Gathering’ trials presented in Chapter 4. The results of this analysis were used to set parameters for exclusion of viewed information for all other analyses presented in *Part 2* of Chapter 4.

Information Requires Time in Order to Influence Choice in Information Gathering Trials

Because of the clear influence of covert attention during ‘Simultaneous’ trials, it is difficult to make conclusions about how subjects may be comparing information during the decision making process. In contrast, the sequential yet relatively unconstrained nature of the ‘Information Gathering’ task provides the ideal opportunity to assess the relationship between value, information comparison and choice. However, one disadvantage of this freedom in subject information gathering behaviour is that subjects may habitually continue to gather information while moving the joystick to respond. If this were to be the case then the last cue that was viewed by the subject would have no influence on his choice.

In order to scrutinise whether subjects tended to gather their final piece of information while moving the joystick we computed the difference in time between when the joystick was initially moved and when the final cue was fixated. This revealed that on a large minority of occasions (M: 31.8%, F: 42.3%), subjects actually fixated to a new cue having already started to move the joystick. Furthermore, when this is considered alongside the finding that the overwhelming majority of joystick movements lasted less than 150ms (M: 96.0%, F: 85.5%), there is strong evidence to suggest that information obtained peri-movement is unlikely to influence choice.

In order to test this hypothesis quantitatively, we used model comparison of six logistic regression models. The dependent variable used in the regression was probability of left choice and the independent variable was always left-right EV difference which was computed using actual values of the cues fixated and the average value (i.e. 3 out of 5) for any cue not fixated. The way in which each model differed was as follows: before computing the EV difference for each trial, the last fixation to joystick movement time was considered. In one model, if the last fixation occurred after joystick movement, this cue was considered as unseen cue, and the average value was assigned to it. Four of the models used the same correction but changed the threshold to also include any cue viewed for <100ms, <200ms, <250ms and <300ms after movement initiation respectively. The final model had no correction. We found that the model with the 250ms and 300ms thresholds gave empirically larger parameter estimates and a significantly better explanation of trial-by-trial choice behaviour compared to the other four models (Kruskal-Wallis test with multiple comparison test, M: Chi^2=403.7, $p<5x10^{-85}$, F: Chi^2=305.7, $p<6x10^{-66}$), however neither was significantly better than the other in the multiple comparison test. This suggests that information presented to subjects requires approximately 250ms to exert an influence on decision processes that are already in motion. Based on this result, we have also disregarded (i.e., replaced its value with an average value of 3) any 'cue four' fixated for <250ms in any analysis of this task in Chapter 4. It should be noted that this analysis cannot account for whether subjects actually represent the value of the last cue but instead heavily underweight it in their final decisions.

Bibliography

ABE, H. & LEE, D. 2011. Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron,* 70**,** 731-41.

ABE, M., SCHAMBRA, H., WASSERMANN, E. M., LUCKENBAUGH, D., SCHWEIGHOFER, N. & COHEN, L. G. 2011. Reward improves long-term retention of a motor memory through induction of offline memory gains. *Curr Biol,* 21**,** 557-62.

ABITBOL, R., LEBRETON, M., HOLLARD, G., RICHMOND, B. J., BOURET, S. & PESSIGLIONE, M. 2015. Neural mechanisms underlying contextual dependency of subjective values: converging evidence from monkeys and humans. *J Neurosci,* 35**,** 2308-20.

ALEXANDER, G. E. & CRUTCHER, M. D. 1990. Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends Neurosci,* 13**,** 266-271.

ALEXANDER, G. E., DELONG, M. R. & STRICK, P. L. 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci,* 9**,** 357-81.

ALEXANDER, W. H. & BROWN, J. W. 2011. Medial prefrontal cortex as an action-outcome predictor. *Nat Neurosci,* 14**,** 1338-44.

AMEMORI, K. & GRAYBIEL, A. M. 2012. Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nat Neurosci,* 15**,** 776-85.

AMIEZ, C., JOSEPH, J. P. & PROCYK, E. 2006. Reward encoding in the monkey anterior cingulate cortex. *Cereb Cortex,* 16**,** 1040-55.

ANDERSON, B. A. 2013. A value-driven mechanism of attentional selection. *Journal of Vision,* 13**,** 7-7.

ANDERSON, B. A., LAURENT, P. A. & YANTIS, S. 2011. Value-driven attentional capture. *Proc Natl Acad Sci U S A,* 108**,** 10367-71.

ANDERSON, B. A. & YANTIS, S. 2012. Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Atten Percept Psychophys,* 74**,** 1644-53.

ARIELI, A. B.-A., Y & RUBINSTEIN, A. 2011. Tracking Decision Makers under Uncertainty. *American Economic Journal: Microeconomics,* 3**,** 68-76.

ARIKUNI, T., SAKO, H. & MURATA, A. 1994. Ipsilateral connections of the anterior cingulate cortex with the frontal and medial temporal cortices in the macaque monkey. *Neurosci Res,* 21**,** 19-39.

ARTHURS, O. J. & BONIFACE, S. 2002. How well do we understand the neural origins of the fMRI BOLD signal? *Trends Neurosci,* 25**,** 27-31.

ASAAD, W. F., RAINER, G. & MILLER, E. K. 2000. Task-specific neural activity in the primate prefrontal cortex. *J Neurophysiol,* 84**,** 451-9.

AVERBECK, B. B. & SEO, M. 2008. The statistical neuroanatomy of frontal networks in the macaque. *PLoS Comput Biol,* 4**,** e1000050.

BARBAS, H. & PANDYA, D. N. 1989. Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *J Comp Neurol,* 286**,** 353-75.

BARBAS, H., ZIKOPOULOS, B. & TIMBIE, C. 2011. Sensory pathways and emotional context for action in primate prefrontal cortex. *Biol Psychiatry,* 69**,** 1133-9.

BARRACLOUGH, D. J., CONROY, M. L. & LEE, D. 2004. Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci,* 7**,** 404-10.

BARTRA, O., MCGUIRE, J. T. & KABLE, J. W. 2013. The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *Neuroimage,* 76**,** 412-427.

BASTEN, U., BIELE, G., HEEKEREN, H. R. & FIEBACH, C. J. 2010. How the brain integrates costs and benefits during decision making. *Proc Natl Acad Sci U S A,* 107**,** 21767-72.

BAXTER, M. G., GAFFAN, D., KYRIAZIS, D. A. & MITCHELL, A. S. 2008. Dorsolateral prefrontal lesions do not impair tests of scene learning and decision-making that require frontal-temporal interaction. *Eur J Neurosci,* 28**,** 491-9.
BAXTER, M. G. & MURRAY, E. A. 2002. The amygdala and reward. *Nat Rev Neurosci,* 3**,** 563-73.
BAYER, H. M. & GLIMCHER, P. W. 2005. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron,* 47**,** 129-41.
BAYLIS, L. L. & GAFFAN, D. 1991. Amygdalectomy and ventromedial prefrontal ablation produce similar deficits in food choice and in simple object discrimination learning for an unseen reward. *Exp Brain Res,* 86**,** 617-22.
BECHARA, A., DAMASIO, A. R., DAMASIO, H. & ANDERSON, S. W. 1994. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition,* 50**,** 7-15.
BECHARA, A., DAMASIO, H., TRANEL, D. & ANDERSON, S. W. 1998. Dissociation Of working memory from decision making within the human prefrontal cortex. *J Neurosci,* 18**,** 428-37.
BECKSTEAD, R. M. & FRANKFURTER, A. 1982. The distribution and some morphological features of substantia nigra neurons that project to the thalamus, superior colliculus and pedunculopontine nucleus in the monkey. *Neuroscience,* 7**,** 2377-88.
BEHRENS, T. E. J., WOOLRICH, M. W., WALTON, M. E. & RUSHWORTH, M. F. S. 2007. Learning the value of information in an uncertain world. *Nat Neurosci,* 10**,** 1214-1221.
BERGER, B., TROTTIER, S., VERNEY, C., GASPAR, P. & ALVAREZ, C. 1988. Regional and laminar distribution of the dopamine and serotonin innervation in the macaque cerebral cortex: a radioautographic study. *J Comp Neurol,* 273**,** 99-119.
BERNACCHIA, A., SEO, H., LEE, D. & WANG, X. J. 2011. A reservoir of time constants for memory traces in cortical neurons. *Nat Neurosci,* 14**,** 366-72.
BERTHOUD, H. R. 2003. Neural systems controlling food intake and energy balance in the modern world. *Curr Opin Clin Nutr Metab Care,* 6**,** 615-20.
BLANCHARD, T. C., HAYDEN, B. Y. & BROMBERG-MARTIN, E. S. 2015. Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron,* 85**,** 602-14.
BOURET, S. & RICHMOND, B. J. 2010. Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J Neurosci,* 30**,** 8591-601.
BROCKMOLE, J. R. & HENDERSON, J. M. 2005. Prioritization of new objects in real-world scenes: evidence from eye movements. *J Exp Psychol Hum Percept Perform,* 31**,** 857-68.
BROMBERG-MARTIN, E. S. & HIKOSAKA, O. 2009. Midbrain Dopamine Neurons Signal Preference for Advance Information about Upcoming Rewards. *Neuron,* 63**,** 119-126.
BROWN, J. W. & BRAVER, T. S. 2005. Learned predictions of error likelihood in the anterior cingulate cortex. *Science,* 307**,** 1118-21.
BRUCE, C., DESIMONE, R. & GROSS, C. G. 1981. Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque. *J Neurophysiol,* 46**,** 369-84.
BUCKLEY, M. J., MANSOURI, F. A., HODA, H., MAHBOUBI, M., BROWNING, P. G., KWOK, S. C., PHILLIPS, A. & TANAKA, K. 2009. Dissociable components of rule-guided behavior depend on distinct medial and prefrontal regions. *Science,* 325**,** 52-8.
CAI, X., KIM, S. & LEE, D. 2011. Heterogeneous coding of temporally discounted values in the dorsal and ventral striatum during intertemporal choice. *Neuron,* 69**,** 170-82.
CAI, X. & PADOA-SCHIOPPA, C. 2012. Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *J Neurosci,* 32**,** 3791-808.
CAI, X. & PADOA-SCHIOPPA, C. 2014. Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. *Neuron,* 81**,** 1140-51.

CALZAVARA, R., MAILLY, P. & HABER, S. N. 2007. Relationship between the corticostriatal terminals from areas 9 and 46, and those from area 8A, dorsal and rostral premotor cortex and area 24c: an anatomical substrate for cognition to action. *Eur J Neurosci,* 26**,** 2005-24.
CAMILLE, N., GRIFFITHS, C. A., VO, K., FELLOWS, L. K. & KABLE, J. W. 2011a. Ventromedial frontal lobe damage disrupts value maximization in humans. *J Neurosci,* 31**,** 7527-32.
CAMILLE, N., TSUCHIDA, A. & FELLOWS, L. K. 2011b. Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *J Neurosci,* 31**,** 15048-52.
CARMICHAEL, S. T. & PRICE, J. L. 1995a. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J Comp Neurol,* 363**,** 615-641.
CARMICHAEL, S. T. & PRICE, J. L. 1995b. Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol,* 363**,** 642-664.
CARMICHAEL, S. T. & PRICE, J. L. 1996. Connectional networks within the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol,* 371**,** 179-207.
CAVADA, C., COMPANY, T., TEJEDOR, J., CRUZ-RIZZOLO, R. J. & REINOSO-SUAREZ, F. 2000. The anatomical connections of the macaque monkey orbitofrontal cortex. A review. *Cereb Cortex,* 10**,** 220-42.
CHAU, B. K., KOLLING, N., HUNT, L. T., WALTON, M. E. & RUSHWORTH, M. F. 2014. A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nat Neurosci,* 17**,** 463-70.
CHIB, V. S., RANGEL, A., SHIMOJO, S. & O'DOHERTY, J. P. 2009. Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. *J Neurosci,* 29**,** 12315-20.
CHIU, P. H., LOHRENZ, T. M. & MONTAGUE, P. R. 2008. Smokers' brains compute, but ignore, a fictive error signal in a sequential investment task. *Nat Neurosci,* 11**,** 514-20.
CISEK, P. 2006. Integrated neural processes for defining potential actions and deciding between them: a computational model. *J Neurosci,* 26**,** 9761-70.
CISEK, P. 2007a. Cortical mechanisms of action selection: the affordance competition hypothesis. *Philos Trans R Soc Lond B Biol Sci,* 362**,** 1585-99.
CISEK, P. 2007b. A parallel framework for interactive behavior. *Prog Brain Res,* 165**,** 475-92.
CISEK, P. 2012. Making decisions through a distributed consensus. *Curr Opin Neurobiol,* 22**,** 927-36.
CLARK, A. M., BOURET, S., YOUNG, A. M., MURRAY, E. A. & RICHMOND, B. J. 2013. Interaction between orbital prefrontal and rhinal cortex is required for normal estimates of expected value. *J Neurosci,* 33**,** 1833-45.
CLARK, L., BECHARA, A., DAMASIO, H., AITKEN, M. R., SAHAKIAN, B. J. & ROBBINS, T. W. 2008. Differential effects of insular and ventromedial prefrontal cortex lesions on risky decision-making. *Brain,* 131**,** 1311-22.
COLLINS, P., ROBERTS, A. C., DIAS, R., EVERITT, B. J. & ROBBINS, T. W. 1998. Perseveration and strategy in a novel spatial self-ordered sequencing task for nonhuman primates: effects of excitotoxic lesions and dopamine depletions of the prefrontal cortex. *J Cogn Neurosci,* 10**,** 332-54.
COOK, G. J. & SWAIN, M. R. 1993. A Computerized Approach to Decision Process Tracing for Decision Support System Design*. *Decision Sciences,* 24**,** 931-952.
COSTA, V. D., TRAN, V. L., TURCHI, J. & AVERBECK, B. B. 2014. Dopamine modulates novelty seeking behavior during decision making. *Behav Neurosci,* 128**,** 556-66.
COUTUREAU, E. & DI SCALA, G. 2009. Entorhinal cortex and cognition. *Prog Neuropsychopharmacol Biol Psychiatry,* 33**,** 753-61.
CRITCHLEY, H. D., TANG, J., GLASER, D., BUTTERWORTH, B. & DOLAN, R. J. 2005. Anterior cingulate activity during error and autonomic response. *Neuroimage,* 27**,** 885-95.

CROMWELL, H. C., HASSANI, O. K. & SCHULTZ, W. 2005. Relative reward processing in primate striatum. *Exp Brain Res,* 162**,** 520-5.
CROXSON, P. L., WALTON, M. E., O'REILLY, J. X., BEHRENS, T. E. & RUSHWORTH, M. F. 2009. Effort-based cost-benefit valuation and the human brain. *J Neurosci,* 29**,** 4531-41.
DAYAN, P., KAKADE, S. & MONTAGUE, P. R. 2000. Learning and selective attention. *Nat Neurosci,* 3 Suppl**,** 1218-23.
DOLAN, R. J. & DAYAN, P. 2013. Goals and habits in the brain. *Neuron,* 80**,** 312-25.
DORRIS, M. C. & GLIMCHER, P. W. 2004. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron,* 44**,** 365-78.
EBLEN, F. & GRAYBIEL, A. M. 1995. Highly restricted origin of prefrontal cortical inputs to striosomes in the macaque monkey. *J Neurosci,* 15**,** 5999-6013.
EGNER, T. & HIRSCH, J. 2005a. Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nat Neurosci,* 8**,** 1784-90.
EGNER, T. & HIRSCH, J. 2005b. The neural correlates and functional integration of cognitive control in a Stroop task. *Neuroimage,* 24**,** 539-47.
ENDERLE, J. D. 2002. Neural control of saccades. *Prog Brain Res,* 140**,** 21-49.
ENGBERT, R. & KLIEGL, R. 2003. Microsaccades uncover the orientation of covert attention. *Vision Research,* 43**,** 1035-1045.
FALKENSTEIN, M., HOORMANN, J., CHRIST, S. & HOHNSBEIN, J. 2000. ERP components on reaction errors and their functional significance: a tutorial. *Biol Psychol,* 51**,** 87-107.
FELDMAN, H. & FRISTON, K. J. 2010. Attention, uncertainty, and free-energy. *Front Hum Neurosci,* 4**,** 215.
FELLOWS, L. K. 2006. Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. *Brain,* 129**,** 944-52.
FELLOWS, L. K. & FARAH, M. J. 2005. Different underlying impairments in decision-making following ventromedial and dorsolateral frontal lobe damage in humans. *Cereb Cortex,* 15**,** 58-63.
FELLOWS, L. K. & FARAH, M. J. 2007. The role of ventromedial prefrontal cortex in decision making: judgment under uncertainty or judgment per se? *Cereb Cortex,* 17**,** 2669-74.
FERRY, A. T., ONGUR, D., AN, X. & PRICE, J. L. 2000. Prefrontal cortical projections to the striatum in macaque monkeys: evidence for an organization related to prefrontal networks. *J Comp Neurol,* 425**,** 447-70.
FOLEY, N. C., JANGRAW, D. C., PECK, C. & GOTTLIEB, J. 2014. Novelty enhances visual salience independently of reward in the parietal lobe. *J Neurosci,* 34**,** 7947-57.
FOLK, C. L., REMINGTON, R. W. & JOHNSTON, J. C. 1992. Involuntary covert orienting is contingent on attentional control settings. *J Exp Psychol Hum Percept Perform,* 18**,** 1030-44.
FRANKLE, W. G., LARUELLE, M. & HABER, S. N. 2006. Prefrontal cortical projections to the midbrain in primates: evidence for a sparse connection. *Neuropsychopharmacology,* 31**,** 1627-36.
FRISTON, K. 2009. The free-energy principle: a rough guide to the brain? *Trends Cogn Sci,* 13**,** 293-301.
FUJIWARA, J., TOBLER, P. N., TAIRA, M., IIJIMA, T. & TSUTSUI, K. 2008. Personality-dependent dissociation of absolute and relative loss processing in orbitofrontal cortex. *Eur J Neurosci,* 27**,** 1547-52.
FUNAHASHI, S., BRUCE, C. J. & GOLDMAN-RAKIC, P. S. 1989. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol,* 61**,** 331-49.
FURL, N. & AVERBECK, B. B. 2011. Parietal cortex and insula relate to evidence seeking relevant to reward-related decisions. *J Neurosci,* 31**,** 17572-82.
FUSTER, J. M. 2001. The prefrontal cortex--an update: time is of the essence. *Neuron,* 30**,** 319-33.

GAFFAN, D., EASTON, A. & PARKER, A. 2002. Interaction of inferior temporal cortex with frontal cortex and basal forebrain: double dissociation in strategy implementation and associative learning. *J Neurosci,* 22**,** 7288-96.
GASPAR, P., BERGER, B., FEBVRET, A., VIGNY, A. & HENRY, J. P. 1989. Catecholamine innervation of the human cerebral cortex as revealed by comparative immunohistochemistry of tyrosine hydroxylase and dopamine-beta-hydroxylase. *J Comp Neurol,* 279**,** 249-71.
GENOVESIO, A., BRASTED, P. J., MITZ, A. R. & WISE, S. P. 2005. Prefrontal cortex activity related to abstract response strategies. *Neuron,* 47**,** 307-20.
GERBELLA, M., BELMALIH, A., BORRA, E., ROZZI, S. & LUPPINO, G. 2010. Cortical connections of the macaque caudal ventrolateral prefrontal areas 45A and 45B. *Cereb Cortex,* 20**,** 141-68.
GHASHGHAEI, H. T., HILGETAG, C. C. & BARBAS, H. 2007. Sequence of information processing for emotions based on the anatomic dialogue between prefrontal cortex and amygdala. *Neuroimage,* 34**,** 905-23.
GLAUSIER, J. R., KHAN, Z. U. & MULY, E. C. 2009. Dopamine D1 and D5 receptors are localized to discrete populations of interneurons in primate prefrontal cortex. *Cereb Cortex,* 19**,** 1820-34.
GOH, J. O., TAN, J. C. & PARK, D. C. 2009. Culture modulates eye-movements to visual novelty. *PLoS One,* 4**,** e8238.
GOTTLIEB, J. 2012. Attention, learning, and the value of information. *Neuron,* 76**,** 281-95.
GOTTLIEB, J., HAYHOE, M., HIKOSAKA, O. & RANGEL, A. 2014. Attention, reward, and information seeking. *J Neurosci,* 34**,** 15497-504.
HABER, S. N. & CALZAVARA, R. 2009. The cortico-basal ganglia integrative network: the role of the thalamus. *Brain Res Bull,* 78**,** 69-74.
HABER, S. N., KIM, K. S., MAILLY, P. & CALZAVARA, R. 2006. Reward-related cortical inputs define a large striatal region in primates that interface with associative cortical connections, providing a substrate for incentive-based learning. *J Neurosci,* 26**,** 8368-76.
HABER, S. N., KUNISHIO, K., MIZOBUCHI, M. & LYND-BALTA, E. 1995. The orbital and medial prefrontal circuit through the primate basal ganglia. *J Neurosci,* 15**,** 4851-67.
HADLAND, K. A., RUSHWORTH, M. F., GAFFAN, D. & PASSINGHAM, R. E. 2003. The anterior cingulate and reward-guided selection of actions. *J Neurophysiol,* 89**,** 1161-4.
HAJCAK, G., HOLROYD, C. B., MOSER, J. S. & SIMONS, R. F. 2005. Brain potentials associated with expected and unexpected good and bad outcomes. *Psychophysiology,* 42**,** 161-70.
HANKS, T., KIANI, R. & SHADLEN, M. N. 2014. A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife,* 3.
HARE, T. A., CAMERER, C. F. & RANGEL, A. 2009. Self-control in decision-making involves modulation of the vmPFC valuation system. *Science,* 324**,** 646-8.
HARLOW, J. M. 1848. Passage of an Iron Rod through the Head. *The Boston Medical and Surgical Journal,* 39**,** 389-393.
HARLOW, J. M. 1993. Recovery from the passage of an iron bar through the head. *History of Psychiatry,* 4**,** 274-281.
HATANAKA, N., TOKUNO, H., HAMADA, I., INASE, M., ITO, Y., IMANISHI, M., HASEGAWA, N., AKAZAWA, T., NAMBU, A. & TAKADA, M. 2003. Thalamocortical and intracortical connections of monkey cingulate motor areas. *J Comp Neurol,* 462**,** 121-138.
HAYDEN, B. Y., PEARSON, J. M. & PLATT, M. L. 2009. Fictive reward signals in the anterior cingulate cortex. *Science,* 324**,** 948-50.
HAYDEN, B. Y., PEARSON, J. M. & PLATT, M. L. 2011. Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci,* 14**,** 933-9.
HAYDEN, B. Y. & PLATT, M. L. 2010. Neurons in anterior cingulate cortex multiplex information about reward and action. *J Neurosci,* 30**,** 3339-46.

HAYES, D. J. & NORTHOFF, G. 2012. Common brain activations for painful and non-painful aversive stimuli. *BMC Neurosci,* 13**,** 60.
HEITZ, RICHARD P. & SCHALL, JEFFREY D. 2012. Neural Mechanisms of Speed-Accuracy Tradeoff. *Neuron,* 76**,** 616-628.
HERNÁNDEZ, A., NÁCHER, V., LUNA, R., ZAINOS, A., LEMUS, L., ALVAREZ, M., VÁZQUEZ, Y., CAMARILLO, L. & ROMO, R. 2010. Decoding a Perceptual Decision Process across Cortex. *Neuron,* 66**,** 300-314.
HIKOSAKA, K. & WATANABE, M. 2000. Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. *Cereb Cortex,* 10**,** 263-71.
HIKOSAKA, O. 2007. Basal ganglia mechanisms of reward-oriented eye movement. *Ann N Y Acad Sci,* 1104**,** 229-49.
HIKOSAKA, O., KIM, H. F., YASUDA, M. & YAMAMOTO, S. 2014. Basal ganglia circuits for reward value-guided behavior. *Annu Rev Neurosci,* 37**,** 289-306.
HIKOSAKA, O., NAKAMURA, K. & NAKAHARA, H. 2006. Basal ganglia orient eyes to reward. *J Neurophysiol,* 95**,** 567-84.
HIKOSAKA, O., TAKIKAWA, Y. & KAWAGOE, R. 2000. Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiol Rev,* 80**,** 953-78.
HIKOSAKA, O. & WURTZ, R. H. 1985. Modification of saccadic eye movements by GABA-related substances. II. Effects of muscimol in monkey substantia nigra pars reticulata. *J Neurophysiol,* 53**,** 292-308.
HIRAI, D., HOSOKAWA, T., INOUE, M., MIYACHI, S. & MIKAMI, A. 2009. Context-dependent representation of reinforcement in monkey amygdala. *Neuroreport,* 20**,** 558-62.
HOLROYD, C. B. & COLES, M. G. 2002. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev,* 109**,** 679-709.
HOLROYD, C. B., DIEN, J. & COLES, M. G. 1998. Error-related scalp potentials elicited by hand and foot movements: evidence for an output-independent error-processing system in humans. *Neurosci Lett,* 242**,** 65-8.
HOLROYD, C. B., NIEUWENHUIS, S., YEUNG, N., NYSTROM, L., MARS, R. B., COLES, M. G. H. & COHEN, J. D. 2004. Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nat Neurosci,* 7**,** 497-498.
HORVITZ, J. C., STEWART, T. & JACOBS, B. L. 1997. Burst activity of ventral tegmental dopamine neurons is elicited by sensory stimuli in the awake cat. *Brain Res,* 759**,** 251-258.
HOSHI, E., SHIMA, K. & TANJI, J. 1998. Task-dependent selectivity of movement-related neuronal activity in the primate prefrontal cortex. *J Neurophysiol,* 80**,** 3392-7.
HOSHI, E., SHIMA, K. & TANJI, J. 2000. Neuronal activity in the primate prefrontal cortex in the process of motor selection based on two behavioral rules. *J Neurophysiol,* 83**,** 2355-73.
HOSOKAWA, T., KENNERLEY, S. W., SLOAN, J. & WALLIS, J. D. 2013. Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *J Neurosci,* 33**,** 17385-97.
HUERTA, M. F. & KAAS, J. H. 1990. Supplementary eye field as defined by intracortical microstimulation: connections in macaques. *J Comp Neurol,* 293**,** 299-330.
HUNT, L. T., DOLAN, R. J. & BEHRENS, T. E. 2014. Hierarchical competitions subserving multi-attribute choice. *Nat Neurosci,* 17**,** 1613-22.
HUNT, L. T., KOLLING, N., SOLTANI, A., WOOLRICH, M. W., RUSHWORTH, M. F. & BEHRENS, T. E. 2012. Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci*.
ITO, S., STUPHORN, V., BROWN, J. W. & SCHALL, J. D. 2003. Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science,* 302**,** 120-2.
ITTI, L. & KOCH, C. 2001. Computational modelling of visual attention. *Nat Rev Neurosci,* 2**,** 194-203.

IZQUIERDO, A., SUDA, R. K. & MURRAY, E. A. 2004. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *J Neurosci,* 24**,** 7540-8.
IZQUIERDO, A., SUDA, R. K. & MURRAY, E. A. 2005. Comparison of the effects of bilateral orbital prefrontal cortex lesions and amygdala lesions on emotional responses in rhesus monkeys. *J Neurosci,* 25**,** 8534-42.
JESSUP, R. K., BUSEMEYER, J. R. & BROWN, J. W. 2010. Error effects in anterior cingulate cortex reverse when error likelihood is high. *J Neurosci,* 30**,** 3467-72.
JOHNSTON, K., LEVIN, H. M., KOVAL, M. J. & EVERLING, S. 2007. Top-down control-signal dynamics in anterior cingulate and prefrontal cortex neurons following task switching. *Neuron,* 53**,** 453-62.
JOHNSTON, W. A., HAWLEY, K. J., PLEWE, S. H., ELLIOTT, J. M. & DEWITT, M. J. 1990. Attention capture by novel stimuli. *J Exp Psychol Gen,* 119**,** 397-411.
JONAS, E., SCHULZ-HARDT, S., FREY, D. & THELEN, N. 2001. Confirmation bias in sequential information search after preliminary decisions: an expansion of dissonance theoretical research on selective exposure to information. *J Pers Soc Psychol,* 80**,** 557-71.
JOVANCEVIC-MISIC, J. & HAYHOE, M. 2009. Adaptive gaze control in natural environments. *J Neurosci,* 29**,** 6234-8.
KAHNEMAN D, T. A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica,* 47**,** 263–291.
KAKADE, S. & DAYAN, P. 2002. Dopamine: generalization and bonuses. *Neural Networks,* 15**,** 549-559.
KEAY, K. A., CLEMENT, C. I., OWLER, B., DEPAULIS, A. & BANDLER, R. 1994. Convergence of deep somatic and visceral nociceptive information onto a discrete ventrolateral midbrain periaqueductal gray region. *Neuroscience,* 61**,** 727-32.
KENNERLEY, S. W., BEHRENS, T. E. & WALLIS, J. D. 2011. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci,* 14**,** 1581-9.
KENNERLEY, S. W., DAHMUBED, A. F., LARA, A. H. & WALLIS, J. D. 2009. Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci,* 21**,** 1162-78.
KENNERLEY, S. W. & WALLIS, J. D. 2009a. Encoding of reward and space during a working memory task in the orbitofrontal cortex and anterior cingulate sulcus. *J Neurophysiol,* 102**,** 3352-64.
KENNERLEY, S. W. & WALLIS, J. D. 2009b. Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci,* 29**,** 2061-73.
KENNERLEY, S. W. & WALLIS, J. D. 2009c. Reward-dependent modulation of working memory in lateral prefrontal cortex. *J Neurosci,* 29**,** 3259-70.
KENNERLEY, S. W. & WALTON, M. E. 2011. Decision making and reward in frontal cortex: complementary evidence from neurophysiological and neuropsychological studies. *Behav Neurosci,* 125**,** 297-317.
KENNERLEY, S. W., WALTON, M. E., BEHRENS, T. E., BUCKLEY, M. J. & RUSHWORTH, M. F. 2006. Optimal decision making and the anterior cingulate cortex. *Nat Neurosci,* 9**,** 940-7.
KERNS, J. G. 2006. Anterior cingulate and prefrontal cortex activity in an FMRI study of trial-to-trial adjustments on the Simon task. *Neuroimage,* 33**,** 399-405.
KERNS, J. G., COHEN, J. D., MACDONALD, A. W., 3RD, CHO, R. Y., STENGER, V. A. & CARTER, C. S. 2004. Anterior cingulate conflict monitoring and adjustments in control. *Science,* 303**,** 1023-6.
KERSTHOLT, J. 1992. Information search and choice accuracy as a function of task complexity and task structure. *Acta Psychol (Amst),* 80**,** 185-197.
KIM, H. F. & HIKOSAKA, O. 2013. Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. *Neuron,* 79**,** 1001-10.

KIM, K. U., HUH, N., JANG, Y., LEE, D. & JUNG, M. W. 2015. Effects of fictive reward on rat's choice behavior. *Sci Rep,* 5**,** 8040.
KIM, S., HWANG, J. & LEE, D. 2008. Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron,* 59**,** 161-72.
KIM, S. G. & OGAWA, S. 2012. Biophysical and physiological origins of blood oxygenation level-dependent fMRI signals. *J Cereb Blood Flow Metab,* 32**,** 1188-206.
KOBAYASHI, S., NOMOTO, K., WATANABE, M., HIKOSAKA, O., SCHULTZ, W. & SAKAGAMI, M. 2006. Influences of rewarding and aversive outcomes on activity in macaque lateral prefrontal cortex. *Neuron,* 51**,** 861-70.
KOBAYASHI, S., PINTO DE CARVALHO, O. & SCHULTZ, W. 2010. Adaptation of reward sensitivity in orbitofrontal neurons. *J Neurosci,* 30**,** 534-44.
KOLLING, N., BEHRENS, T. E., MARS, R. B. & RUSHWORTH, M. F. 2012. Neural mechanisms of foraging. *Science,* 336**,** 95-8.
KONDO, H., SALEEM, K. S. & PRICE, J. L. 2003. Differential connections of the temporal pole with the orbital and medial prefrontal networks in macaque monkeys. *J Comp Neurol,* 465**,** 499-523.
KONDO, H., SALEEM, K. S. & PRICE, J. L. 2005. Differential connections of the perirhinal and parahippocampal cortex with the orbital and medial prefrontal networks in macaque monkeys. *J Comp Neurol,* 493**,** 479-509.
KORI, A., MIYASHITA, N., KATO, M., HIKOSAKA, O., USUI, S. & MATSUMURA, M. 1995. Eye movements in monkeys with local dopamine depletion in the caudate nucleus. II. Deficits in voluntary saccades. *J Neurosci,* 15**,** 928-41.
KRAJBICH, I., ARMEL, C. & RANGEL, A. 2010. Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci,* 13**,** 1292-8.
KRAJBICH, I., LU, D., CAMERER, C. & RANGEL, A. 2012. The attentional drift-diffusion model extends to simple purchasing decisions. *Front Psychol,* 3**,** 193.
KRAJBICH, I. & RANGEL, A. 2011. Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc Natl Acad Sci U S A,* 108**,** 13852-7.
KUNISHIO, K. & HABER, S. N. 1994. Primate cingulostriatal projection: Limbic striatal versus sensorimotor striatal input. *J Comp Neurol,* 350**,** 337-356.
LAIRD, A. R., FOX, P. M., EICKHOFF, S. B., TURNER, J. A., RAY, K. L., MCKAY, D. R., GLAHN, D. C., BECKMANN, C. F., SMITH, S. M. & FOX, P. T. 2011. Behavioral Interpretations of Intrinsic Connectivity Networks. *J Cogn Neurosci,* 23**,** 4022-4037.
LAU, B. & GLIMCHER, P. W. 2008. Value representations in the primate striatum during matching behavior. *Neuron,* 58**,** 451-63.
LEATHERS, M. L. & OLSON, C. R. 2012. In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. *Science,* 338**,** 132-5.
LEBEDEV, M. A., MESSINGER, A., KRALIK, J. D. & WISE, S. P. 2004. Representation of attended versus remembered locations in prefrontal cortex. *PLoS Biol,* 2**,** e365.
LEIBOWITZ, S. F. 1988. Hypothalamic paraventricular nucleus: interaction between alpha 2-noradrenergic system and circulating hormones and nutrients in relation to energy balance. *Neurosci Biobehav Rev,* 12**,** 101-9.
LEVY, D. J. & GLIMCHER, P. W. 2011. Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. *J Neurosci,* 31**,** 14693-707.
LIM, S. L., O'DOHERTY, J. P. & RANGEL, A. 2011. The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. *J Neurosci,* 31**,** 13214-23.
LIM, S. L., O'DOHERTY, J. P. & RANGEL, A. 2013. Stimulus value signals in ventromedial PFC reflect the integration of attribute value signals computed in fusiform gyrus and posterior superior temporal gyrus. *J Neurosci,* 33**,** 8729-41.
LIN, A., ADOLPHS, R. & RANGEL, A. 2012. Social and monetary reward learning engage overlapping neural substrates. *Soc Cogn Affect Neurosci,* 7**,** 274-81.

LIU, X., HAIRSTON, J., SCHRIER, M. & FAN, J. 2011. Common and distinct networks underlying reward valence and processing stages: a meta-analysis of functional neuroimaging studies. *Neurosci Biobehav Rev,* 35**,** 1219-36.
LOEWENSTEIN, G. 1994. The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin,* 116**,** 75-98.
LOHSE, G. L. & JOHNSON, E. J. 1996. A Comparison of Two Process Tracing Methods for Choice Tasks. *Organizational Behavior and Human Decision Processes,* 68**,** 28-43.
LOUIE, K., GRATTAN, L. E. & GLIMCHER, P. W. 2011. Reward value-based gain control: divisive normalization in parietal cortex. *J Neurosci,* 31**,** 10627-39.
LOUIE, K., KHAW, M. W. & GLIMCHER, P. W. 2013. Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci U S A,* 110**,** 6139-44.
LUK, C. H. & WALLIS, J. D. 2009. Dynamic encoding of responses and outcomes by neurons in medial prefrontal cortex. *J Neurosci,* 29**,** 7526-39.
LUPPINO, G., ROZZI, S., CALZAVARA, R. & MATELLI, M. 2003. Prefrontal and agranular cingulate projections to the dorsal premotor areas F2 and F7 in the macaque monkey. *Eur J Neurosci,* 17**,** 559-78.
MACHENS, C. K., ROMO, R. & BRODY, C. D. 2005. Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science,* 307**,** 1121-4.
MANES, F., SAHAKIAN, B., CLARK, L., ROGERS, R., ANTOUN, N., AITKEN, M. & ROBBINS, T. 2002. Decision-making processes following damage to the prefrontal cortex. *Brain,* 125**,** 624-39.
MANSOURI, F. A., TANAKA, K. & BUCKLEY, M. J. 2009. Conflict-induced behavioural adjustment: a clue to the executive functions of the prefrontal cortex. *Nat Rev Neurosci,* 10**,** 141-52.
MATELLI, M., CAMARDA, R., GLICKSTEIN, M. & RIZZOLATTI, G. 1986. Afferent and efferent projections of the inferior area 6 in the macaque monkey. *J Comp Neurol,* 251**,** 281-98.
MATSUMOTO, K., SUZUKI, W. & TANAKA, K. 2003. Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science,* 301**,** 229-32.
MATSUMOTO, M. & HIKOSAKA, O. 2009. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature,* 459**,** 837-41.
MATSUMOTO, M., MATSUMOTO, K., ABE, H. & TANAKA, K. 2007. Medial prefrontal cell activity signaling prediction errors of action values. *Nat Neurosci,* 10**,** 647-56.
MCCLURE, S. M., BERNS, G. S. & MONTAGUE, P. R. 2003. Temporal prediction errors in a passive learning task activate human striatum. *Neuron,* 38**,** 339-46.
MIDDLETON, F. A. & STRICK, P. L. 2001. Cerebellar projections to the prefrontal cortex of the primate. *J Neurosci,* 21**,** 700-12.
MILLER, E. K. & COHEN, J. D. 2001. AN INTEGRATIVE THEORY OF PREFRONTAL CORTEX FUNCTION. *Annu Rev Neurosci,* 24**,** 167-202.
MILSTEIN, D. M. & DORRIS, M. C. 2007. The influence of expected value on saccadic preparation. *J Neurosci,* 27**,** 4810-8.
MONOSOV, I. E. & HIKOSAKA, O. 2012. Regionally distinct processing of rewards and punishments by the primate ventromedial prefrontal cortex. *J Neurosci,* 32**,** 10318-30.
MONTAGUE, P. R., KING-CASAS, B. & COHEN, J. D. 2006. Imaging valuation models in human choice. *Annu Rev Neurosci,* 29**,** 417-48.
MORECRAFT, R. J. & VAN HOESEN, G. W. 1992. Cingulate input to the primary and supplementary motor cortices in the rhesus monkey: evidence for somatotopy in areas 24c and 23c. *J Comp Neurol,* 322**,** 471-89.
MORRISON, S. E., SAEZ, A., LAU, B. & SALZMAN, C. D. 2011. Different time courses for learning-related changes in amygdala and orbitofrontal cortex. *Neuron,* 71**,** 1127-40.
MORRISON, S. E. & SALZMAN, C. D. 2009. The convergence of information about rewarding and aversive stimuli in single neurons. *J Neurosci,* 29**,** 11471-83.

MORRISON, S. E. & SALZMAN, C. D. 2011. Representations of appetitive and aversive information in the primate orbitofrontal cortex. *Ann N Y Acad Sci,* 1239**,** 59-70.
MUNOZ, D. P. 2002. Commentary: saccadic eye movements: overview of neural circuitry. *Prog Brain Res,* 140**,** 89-96.
NAKANO, K., KAYAHARA, T. & CHIBA, T. 1999. Afferent connections to the ventral striatum from the medial prefrontal cortex (area 25) and the thalamic nuclei in the macaque monkey. *Ann N Y Acad Sci,* 877**,** 667-70.
NEE, D. E., KASTNER, S. & BROWN, J. W. 2011. Functional heterogeneity of conflict, error, task-switching, and unexpectedness effects within medial prefrontal cortex. *Neuroimage,* 54**,** 528-540.
NEUBERT, F. X., MARS, R. B., SALLET, J. & RUSHWORTH, M. F. 2015. Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. *Proc Natl Acad Sci U S A,* 112**,** E2695-704.
NEUBERT, F. X., MARS, R. B., THOMAS, A. G., SALLET, J. & RUSHWORTH, M. F. 2014. Comparison of human ventral frontal cortex areas for cognitive control and language with areas in monkey frontal cortex. *Neuron,* 81**,** 700-13.
NICKERSON, R. S. 1998. Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology,* 2**,** 175–220.
NIKI, H. & WATANABE, M. 1976. Prefrontal unit activity and delayed response: relation to cue location versus direction of response. *Brain Res,* 105**,** 79-88.
NIKI, H. & WATANABE, M. 1979. Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res,* 171**,** 213-24.
NOONAN, M. P., MARS, R. B. & RUSHWORTH, M. F. 2011. Distinct roles of three frontal cortical areas in reward-guided behavior. *J Neurosci,* 31**,** 14399-412.
NOONAN, M. P., WALTON, M. E., BEHRENS, T. E., SALLET, J., BUCKLEY, M. J. & RUSHWORTH, M. F. 2010. Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc Natl Acad Sci U S A,* 107**,** 20547-52.
NOTHDURFT, H. C. 2002. Attention shifts to salient targets. *Vision Res,* 42**,** 1287-306.
NUNEZ CASTELLAR, E., KUHN, S., FIAS, W. & NOTEBAERT, W. 2010. Outcome expectancy and not accuracy determines posterror slowing: ERP support. *Cogn Affect Behav Neurosci,* 10**,** 270-8.
O'DOHERTY, J., CRITCHLEY, H., DEICHMANN, R. & DOLAN, R. J. 2003. Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *J Neurosci,* 23**,** 7931-9.
O'NEILL, M. & SCHULTZ, W. 2010. Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron,* 68**,** 789-800.
ONGUR, D. & PRICE, J. L. 2000. The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cereb Cortex,* 10**,** 206-19.
OSTLUND, S. B. & BALLEINE, B. W. 2005. Lesions of medial prefrontal cortex disrupt the acquisition but not the expression of goal-directed learning. *J Neurosci,* 25**,** 7763-70.
PADOA-SCHIOPPA, C. 2009. Range-adapting representation of economic value in the orbitofrontal cortex. *J Neurosci,* 29**,** 14004-14.
PADOA-SCHIOPPA, C. 2011. Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci,* 34**,** 333-59.
PADOA-SCHIOPPA, C. & ASSAD, J. A. 2006. Neurons in the orbitofrontal cortex encode economic value. *Nature,* 441**,** 223-6.
PADOA-SCHIOPPA, C. & ASSAD, J. A. 2008. The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat Neurosci,* 11**,** 95-102.
PARENT, A., MACKEY, A. & DE BELLEFEUILLE, L. 1983. The subcortical afferents to caudate nucleus and putamen in primate: a fluorescence retrograde double labeling study. *Neuroscience,* 10**,** 1137-50.
PASSINGHAM, R. E., STEPHAN, K. E. & KOTTER, R. 2002. The anatomical basis of functional localization in the cortex. *Nat Rev Neurosci,* 3**,** 606-16.

PASTOR-BERNIER, A. & CISEK, P. 2011. Neural correlates of biased competition in premotor cortex. *J Neurosci,* 31**,** 7083-8.
PATON, J. J., BELOVA, M. A., MORRISON, S. E. & SALZMAN, C. D. 2006. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature,* 439**,** 865-70.
PAULUS, M. P., HOZACK, N., FRANK, L. & BROWN, G. G. 2002. Error rate and outcome predictability affect neural activation in prefrontal cortex and anterior cingulate during decision-making. *Neuroimage,* 15**,** 836-846.
PAXINOS, G. & MAI, J. R. K. 2004. The human nervous system. 2nd / edited by George Paxinos & Jurgen K. Mai. ed. Oxford: Academic.
PAYNE, J. W. 1976. Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance,* 16**,** 366-387.
PETRIDES, M. 2005. Lateral prefrontal cortex: architectonic and functional organization. *Philos Trans R Soc Lond B Biol Sci,* 360**,** 781-95.
PHILIASTIDES, M. G., BIELE, G. & HEEKEREN, H. R. 2010. A mechanistic account of value computation in the human brain. *Proc Natl Acad Sci U S A,* 107**,** 9430-5.
PHILIASTIDES, M. G. & RATCLIFF, R. 2013. Influence of branding on preference-based decision making. *Psychol Sci,* 24**,** 1208-15.
PLASSMANN, H., O'DOHERTY, J. & RANGEL, A. 2007. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J Neurosci,* 27**,** 9984-8.
PLATT, M. L. & GLIMCHER, P. W. 1999. Neural correlates of decision variables in parietal cortex. *Nature,* 400**,** 233-8.
POREMBA, A. & MISHKIN, M. 2007. Exploring the extent and function of higher-order auditory cortex in rhesus monkeys. *Hear Res,* 229**,** 14-23.
PORRINO, L. J. & GOLDMAN-RAKIC, P. S. 1982. Brainstem innervation of prefrontal and anterior cingulate cortex in the rhesus monkey revealed by retrograde transport of HRP. *J Comp Neurol,* 205**,** 63-76.
PRESTON, ALISON R. & EICHENBAUM, H. 2013. Interplay of Hippocampus and Prefrontal Cortex in Memory. *Current Biology,* 23**,** R764-R773.
PREVOST, C., PESSIGLIONE, M., METEREAU, E., CLERY-MELIN, M. L. & DREHER, J. C. 2010. Separate valuation subsystems for delay and effort decision costs. *J Neurosci,* 30**,** 14080-90.
PRICE, J. L. 1999. Prefrontal cortical networks related to visceral function and mood. *Ann N Y Acad Sci,* 877**,** 383-96.
PRICE, J. L. 2007. Definition of the orbital cortex in relation to specific connections with limbic and visceral structures and other cortical regions. *Ann N Y Acad Sci,* 1121**,** 54-71.
QUILODRAN, R., ROTHE, M. & PROCYK, E. 2008. Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron,* 57**,** 314-25.
RAGHANTI, M. A., STIMPSON, C. D., MARCINKIEWICZ, J. L., ERWIN, J. M., HOF, P. R. & SHERWOOD, C. C. 2008. Cortical dopaminergic innervation among humans, chimpanzees, and macaque monkeys: a comparative study. *Neuroscience,* 155**,** 203-20.
RAGHURAMAN, A. P. & PADOA-SCHIOPPA, C. 2014. Integration of multiple determinants in the neuronal computation of economic values. *J Neurosci,* 34**,** 11583-603.
RANGEL, A., CAMERER, C. & MONTAGUE, P. R. 2008. A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci,* 9**,** 545-56.
RAO, S. C., RAINER, G. & MILLER, E. K. 1997. Integration of what and where in the primate prefrontal cortex. *Science,* 276**,** 821-4.
RATCLIFF, R. 1976. A mathematical model for paced serial addition. *Journal of Mathematical Psychology,* 13**,** 46-78.
RATCLIFF, R. & MCKOON, G. 2008. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput,* 20**,** 873-922.

RICH, E. L. & WALLIS, J. D. 2014. Medial-lateral organization of the orbitofrontal cortex. *J Cogn Neurosci,* 26**,** 1347-62.
RICHFIELD, E. K., YOUNG, A. B. & PENNEY, J. B. 1989. Comparative distributions of dopamine D-1 and D-2 receptors in the cerebral cortex of rats, cats, and monkeys. *J Comp Neurol,* 286**,** 409-26.
ROESCH, M. R. & OLSON, C. R. 2003. Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. *J Neurophysiol,* 90**,** 1766-89.
ROESCH, M. R. & OLSON, C. R. 2004. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science,* 304**,** 307-10.
ROESCH, M. R., TAYLOR, A. R. & SCHOENBAUM, G. 2006. Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron,* 51**,** 509-20.
ROITMAN, J. D. & SHADLEN, M. N. 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci,* 22**,** 9475-89.
ROLLS, E. T. 2000. The orbitofrontal cortex and reward. *Cereb Cortex,* 10**,** 284-94.
ROLLS, E. T. & BAYLIS, L. L. 1994. Gustatory, olfactory, and visual convergence within the primate orbitofrontal cortex. *J Neurosci,* 14**,** 5437-52.
ROLLS, E. T., YAXLEY, S. & SIENKIEWICZ, Z. J. 1990. Gustatory responses of single neurons in the caudolateral orbitofrontal cortex of the macaque monkey. *J Neurophysiol,* 64**,** 1055-66.
RORIE, A. E., GAO, J., MCCLELLAND, J. L. & NEWSOME, W. T. 2010. Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS One,* 5**,** e9308.
RUDEBECK, P. H., BEHRENS, T. E., KENNERLEY, S. W., BAXTER, M. G., BUCKLEY, M. J., WALTON, M. E. & RUSHWORTH, M. F. 2008. Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J Neurosci,* 28**,** 13775-85.
RUDEBECK, P. H., BUCKLEY, M. J., WALTON, M. E. & RUSHWORTH, M. F. 2006a. A role for the macaque anterior cingulate gyrus in social valuation. *Science,* 313**,** 1310-2.
RUDEBECK, P. H., MITZ, A. R., CHACKO, R. V. & MURRAY, E. A. 2013a. Effects of amygdala lesions on reward-value coding in orbital and medial prefrontal cortex. *Neuron,* 80**,** 1519-31.
RUDEBECK, P. H. & MURRAY, E. A. 2011. Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *J Neurosci,* 31**,** 10569-78.
RUDEBECK, P. H., SAUNDERS, R. C., PRESCOTT, A. T., CHAU, L. S. & MURRAY, E. A. 2013b. Prefrontal mechanisms of behavioral flexibility, emotion regulation and value updating. *Nat Neurosci,* 16**,** 1140-5.
RUDEBECK, P. H., WALTON, M. E., SMYTH, A. N., BANNERMAN, D. M. & RUSHWORTH, M. F. 2006b. Separate neural pathways process different decision costs. *Nat Neurosci,* 9**,** 1161-8.
RUSHWORTH, M. F., NIXON, P. D., EACOTT, M. J. & PASSINGHAM, R. E. 1997. Ventral prefrontal cortex is not essential for working memory. *J Neurosci,* 17**,** 4829-38.
SADDORIS, M. P., GALLAGHER, M. & SCHOENBAUM, G. 2005. Rapid associative encoding in basolateral amygdala depends on connections with orbitofrontal cortex. *Neuron,* 46**,** 321-31.
SAITO, N., MUSHIAKE, H., SAKAMOTO, K., ITOYAMA, Y. & TANJI, J. 2005. Representation of immediate and final behavioral goals in the monkey prefrontal cortex during an instructed delay period. *Cereb Cortex,* 15**,** 1535-46.
SALEEM, K. S., MILLER, B. & PRICE, J. L. 2014. Subdivisions and connectional networks of the lateral prefrontal cortex in the macaque monkey. *J Comp Neurol,* 522**,** 1641-90.

SALLET, J., MARS, R. B., NOONAN, M. P., NEUBERT, F. X., JBABDI, S., O'REILLY, J. X., FILIPPINI, N., THOMAS, A. G. & RUSHWORTH, M. F. 2013. The organization of dorsal frontal cortex in humans and macaques. *J Neurosci,* 33**,** 12255-74.
SALLET, J., QUILODRAN, R., ROTHE, M., VEZOLI, J., JOSEPH, J. P. & PROCYK, E. 2007. Expectations, gains, and losses in the anterior cingulate cortex. *Cogn Affect Behav Neurosci,* 7**,** 327-36.
SCHMAHMANN, J. D., PANDYA, D. N., WANG, R., DAI, G., D'ARCEUIL, H. E., DE CRESPIGNY, A. J. & WEDEEN, V. J. 2007. *Association fibre pathways of the brain: parallel observations from diffusion spectrum imaging and autoradiography*.
SCHMOLESKY, M. T., WANG, Y., HANES, D. P., THOMPSON, K. G., LEUTGEB, S., SCHALL, J. D. & LEVENTHAL, A. G. 1998. Signal timing across the macaque visual system. *J Neurophysiol,* 79**,** 3272-8.
SCHOENBAUM, G., CHIBA, A. A. & GALLAGHER, M. 1998. Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat Neurosci,* 1**,** 155-9.
SCHOENBAUM, G., SETLOW, B., SADDORIS, M. P. & GALLAGHER, M. 2003. Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. *Neuron,* 39**,** 855-67.
SCHULTZ, W. 1986. Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey. *J Neurophysiol,* 56**,** 1439-61.
SCHULTZ, W. 1998. Predictive reward signal of dopamine neurons. *J Neurophysiol,* 80**,** 1-27.
SCHULTZ, W., APICELLA, P., SCARNATI, E. & LJUNGBERG, T. 1992. Neuronal activity in monkey ventral striatum related to the expectation of reward. *J Neurosci,* 12**,** 4595-610.
SCHULTZ, W., DAYAN, P. & MONTAGUE, P. R. 1997. A neural substrate of prediction and reward. *Science,* 275**,** 1593-9.
SELEN, L. P., SHADLEN, M. N. & WOLPERT, D. M. 2012. Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J Neurosci,* 32**,** 2276-86.
SEO, H., BARRACLOUGH, D. J. & LEE, D. 2007. Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cereb Cortex,* 17 Suppl 1**,** i110-7.
SEO, H. & LEE, D. 2007. Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci,* 27**,** 8366-77.
SEO, H. & LEE, D. 2009. Behavioral and neural changes after gains and losses of conditioned reinforcers. *J Neurosci,* 29**,** 3627-41.
SERENCES, J. T. 2008. Value-based modulations in human visual cortex. *Neuron,* 60**,** 1169-81.
SHENHAV, A., BOTVINICK, M. M. & COHEN, J. D. 2013. The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron,* 79**,** 217-40.
SHIMA, K. & TANJI, J. 1998. Role for cingulate motor area cells in voluntary movement selection based on reward. *Science,* 282**,** 1335-8.
SHIMOJO, S., SIMION, C., SHIMOJO, E. & SCHEIER, C. 2003. Gaze bias both reflects and influences preference. *Nat Neurosci,* 6**,** 1317-22.
SIEGEL, M., BUSCHMAN, T. J. & MILLER, E. K. 2015. BRAIN PROCESSING. Cortical information flow during flexible sensorimotor decisions. *Science,* 348**,** 1352-5.
SILVETTI, M., NUNEZ CASTELLAR, E., ROGER, C. & VERGUTS, T. 2014. Reward expectation and prediction error in human medial frontal cortex: an EEG study. *Neuroimage,* 84**,** 376-82.
SILVETTI, M., SEURINCK, R. & VERGUTS, T. 2011. Value and prediction error in medial frontal cortex: integrating the single-unit and systems levels of analysis. *Front Hum Neurosci,* 5**,** 75.
SILVETTI, M., SEURINCK, R. & VERGUTS, T. 2013. Value and prediction error estimation account for volatility effects in ACC: A model-based fMRI study. *Cortex,* 49**,** 1627-1635.

SLOVIC, P. & LICHTENSTEIN, S. 1971. Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance,* 6**,** 649-744.
SQUIRE, L. R., STARK, C. E. & CLARK, R. E. 2004. The medial temporal lobe. *Annu Rev Neurosci,* 27**,** 279-306.
STALNAKER, T. A., FRANZ, T. M., SINGH, T. & SCHOENBAUM, G. 2007. Basolateral amygdala lesions abolish orbitofrontal-dependent reversal impairments. *Neuron,* 54**,** 51-8.
STEMMER, B., SEGALOWITZ, S. J., WITZKE, W. & SCHÖNLE, P. W. 2004. Error detection in patients with lesions to the medial prefrontal cortex: an ERP study. *Neuropsychologia,* 42**,** 118-130.
STRAIT, C. E., BLANCHARD, T. C. & HAYDEN, B. Y. 2014. Reward value comparison via mutual inhibition in ventromedial prefrontal cortex. *Neuron,* 82**,** 1357-66.
STUSS, D. T., FLODEN, D., ALEXANDER, M. P., LEVINE, B. & KATZ, D. 2001. Stroop performance in focal lesion patients: dissociation of processes and frontal lobe lesion location. *Neuropsychologia,* 39**,** 771-86.
SUGRUE, L. P., CORRADO, G. S. & NEWSOME, W. T. 2004. Matching behavior and the representation of value in the parietal cortex. *Science,* 304**,** 1782-7.
SUL, J. H., KIM, H., HUH, N., LEE, D. & JUNG, M. W. 2010. Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron,* 66**,** 449-60.
SULLIVAN, B. T., JOHNSON, L., ROTHKOPF, C. A., BALLARD, D. & HAYHOE, M. 2012. The role of uncertainty and reward on eye movements in a virtual driving task. *J Vis,* 12**,** 19.
SUNDSTRÖM, G. A. 1987. Information search and decision making: The effects of information displays. *Acta Psychol (Amst),* 65**,** 165-179.
SUZUKI, W. A. 1996. The anatomy, physiology and functions of the perirhinal cortex. *Curr Opin Neurobiol,* 6**,** 179-86.
SYMMONDS, M., WRIGHT, N. D., BACH, D. R. & DOLAN, R. J. 2011. Deconstructing risk: separable encoding of variance and skewness in the brain. *Neuroimage,* 58**,** 1139-49.
TALMI, D., ATKINSON, R. & EL-DEREDY, W. 2013. The feedback-related negativity signals salience prediction errors, not reward prediction errors. *J Neurosci,* 33**,** 8264-9.
TANJI, J. & HOSHI, E. 2008. Role of the lateral prefrontal cortex in executive behavioral control. *Physiol Rev,* 88**,** 37-57.
THEEUWES, J. 1992. Perceptual selectivity for color and form. *Percept Psychophys,* 51**,** 599-606.
THEVARAJAH, D., MIKULIC, A. & DORRIS, M. C. 2009. Role of the superior colliculus in choosing mixed-strategy saccades. *J Neurosci,* 29**,** 1998-2008.
THIERRY, A. M., GIOANNI, Y., DEGENETAIS, E. & GLOWINSKI, J. 2000. Hippocampo-prefrontal cortex pathway: anatomical and electrophysiological characteristics. *Hippocampus,* 10**,** 411-9.
THORPE, S. J. & FABRE-THORPE, M. 2001. Neuroscience. Seeking categories in the brain. *Science,* 291**,** 260-3.
TOBLER, P. N., FIORILLO, C. D. & SCHULTZ, W. 2005. Adaptive coding of reward value by dopamine neurons. *Science,* 307**,** 1642-5.
TOTAH, N. K., KIM, Y. B., HOMAYOUN, H. & MOGHADDAM, B. 2009. Anterior cingulate neurons represent errors and preparatory attention within the same behavioral sequence. *J Neurosci,* 29**,** 6418-26.
TOWAL, R. B., MORMANN, M. & KOCH, C. 2013. Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proc Natl Acad Sci U S A,* 110**,** E3858-67.
TREMBLAY, L. & SCHULTZ, W. 1999. Relative reward preference in primate orbitofrontal cortex. *Nature,* 398**,** 704-8.

TREMBLAY, L. & SCHULTZ, W. 2000a. Modifications of reward expectation-related neuronal activity during learning in primate orbitofrontal cortex. *J Neurophysiol,* 83**,** 1877-85.
TREMBLAY, L. & SCHULTZ, W. 2000b. Reward-related neuronal activity during go-nogo task performance in primate orbitofrontal cortex. *J Neurophysiol,* 83**,** 1864-76.
TSUJIMOTO, S., GENOVESIO, A. & WISE, S. P. 2009. Monkey orbitofrontal cortex encodes response choices near feedback time. *J Neurosci,* 29**,** 2569-74.
TSUJIMOTO, S., GENOVESIO, A. & WISE, S. P. 2012. Neuronal activity during a cued strategy task: comparison of dorsolateral, orbital, and polar prefrontal cortex. *J Neurosci,* 32**,** 11017-31.
TVERSKY, A. 1972. Choice by elimination. *Journal of Mathematical Psychology,* 9**,** 341-367.
UCHIDA, Y., LU, X., OHMAE, S., TAKAHASHI, T. & KITAZAWA, S. 2007. Neuronal activity related to reward size and rewarded target position in primate supplementary eye field. *J Neurosci,* 27**,** 13750-5.
UNGERLEIDER, L. M., M. 1982. Two cortical visual systems. *Analysis of visual behavior***,** 549-586.
VAN 'T ENT, D. & APKARIAN, P. 1999. Motoric response inhibition in finger movement and saccadic eye movement: a comparative study. *Clin Neurophysiol,* 110**,** 1058-72.
VAN VEEN, V., COHEN, J. D., BOTVINICK, M. M., STENGER, V. A. & CARTER, C. S. 2001. Anterior cingulate cortex, conflict monitoring, and levels of processing. *Neuroimage,* 14**,** 1302-8.
WALLIS, J. D. & MILLER, E. K. 2003. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci,* 18**,** 2069-81.
WALSH, M. M. & ANDERSON, J. R. 2012. Learning from experience: Event-related potential correlates of reward processing, neural adaptation, and behavioral choice. *Neuroscience & Biobehavioral Reviews,* 36**,** 1870-1884.
WALTON, M. E., BANNERMAN, D. M. & RUSHWORTH, M. F. 2002. The role of rat medial frontal cortex in effort-based decision making. *J Neurosci,* 22**,** 10996-1003.
WALTON, M. E., BEHRENS, T. E., BUCKLEY, M. J., RUDEBECK, P. H. & RUSHWORTH, M. F. 2010. Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron,* 65**,** 927-39.
WANG, X. J. 2002. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron,* 36**,** 955-68.
WANG, X. J. 2008. Decision making in recurrent neuronal circuits. *Neuron,* 60**,** 215-34.
WASON, P. C. 1968. Reasoning about a rule. *Q J Exp Psychol,* 20**,** 273-81.
WATANABE, M. 1996. Reward expectancy in primate prefrontal neurons. *Nature,* 382**,** 629-32.
WEENIG, M. W. H. & MAARLEVELD, M. 2002. The impact of time constraint on information search strategies in complex choice tasks. *Journal of Economic Psychology,* 23**,** 689-702.
WEISKOPF, N., HUTTON, C., JOSEPHS, O., TURNER, R. & DEICHMANN, R. 2007. Optimized EPI for fMRI studies of the orbitofrontal cortex: compensation of susceptibility-induced gradients in the readout direction. *MAGMA,* 20**,** 39-49.
WHITE, I. M. & WISE, S. P. 1999. Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res,* 126**,** 315-35.
WILLIAMS, S. M. & GOLDMAN-RAKIC, P. S. 1993. Characterization of the dopaminergic innervation of the primate frontal cortex using a dopamine-specific antibody. *Cereb Cortex,* 3**,** 199-222.
WILLIAMS, S. M. & GOLDMAN-RAKIC, P. S. 1998. Widespread origin of the primate mesofrontal dopamine system. *Cereb Cortex,* 8**,** 321-45.
WILSON, F. A. & GOLDMAN-RAKIC, P. S. 1994. Viewing preferences of rhesus monkeys related to memory for complex pictures, colours and faces. *Behav Brain Res,* 60**,** 79-89.

WISE, S. P. 2008. Forward frontal fields: phylogeny and fundamental function. *Trends Neurosci,* 31**,** 599-608.
WRIGHT, N. D., SYMMONDS, M. & DOLAN, R. J. 2013. Distinct encoding of risk and value in economic choice between multiple risky options. *Neuroimage,* 81**,** 431-40.
WU, Y. & ZHOU, X. 2009. The P300 and reward valence, magnitude, and expectancy in outcome evaluation. *Brain Res,* 1286**,** 114-22.
YAMAMOTO, S., MONOSOV, I. E., YASUDA, M. & HIKOSAKA, O. 2012. What and where information in the caudate tail guides saccades to visual objects. *J Neurosci,* 32**,** 11005-16.
YASUDA, M., YAMAMOTO, S. & HIKOSAKA, O. 2012. Robust representation of stable object values in the oculomotor Basal Ganglia. *J Neurosci,* 32**,** 16917-32.
YETERIAN, E. H. & PANDYA, D. N. 1991. Prefrontostriatal connections in relation to cortical architectonic organization in rhesus monkeys. *J Comp Neurol,* 312**,** 43-67.
YETERIAN, E. H., PANDYA, D. N., TOMAIUOLO, F. & PETRIDES, M. 2012. The cortical connectivity of the prefrontal cortex in the monkey brain. *Cortex,* 48**,** 58-81.
YEUNG, N. & NIEUWENHUIS, S. 2009. Dissociating response conflict and error likelihood in anterior cingulate cortex. *J Neurosci,* 29**,** 14506-10.
ÖNGÜR D, P. J. 2000. The Organsiation of Networks within the Orbital and Medial Prefrontal Cortex of Rats, Monkeys and Humans. *Cerebral Cortex,* 10**,** 206-19.