REPRESENTATIONS OF RELATIVE VALUE CODING IN THE ORBITOFRONTAL CORTEX AND AMYGDALA

Rebecca Allyson Saez

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2013

Abstract

Representations of relative value coding in the orbitofrontal cortex and amygdala

Rebecca Allyson Saez

In order to guide behavior, humans and animals must flexibly evaluate the motivational significance of stimuli in the environment. We sought to determine if, in different contexts, neurons in the amygdala and orbitofrontal cortex (OFC) indeed rescale their calculation of the motivational significance of stimuli that predict rewards. We used a "contrast revaluation" task in which the reward associated with one stimulus is held constant while other rewards within a particular context (or block of trials) change. This manipulation modulates the relative significance of the reward associated with one stimulus without changing its absolute amount.

We recorded the activity of individual neurons in the amygdala and OFC of two monkeys while they performed the contrast revaluation task. On every trial, a monkey viewed one of two conditioned stimuli (CSs; distinct fractal patterns), each predictive of a different reward amount. CSs were novel for every experiment. Unconditioned stimulus (US, liquid reward) delivery followed CS presentation and a brief temporal gap (trace interval). The task consisted of three trial blocks, with switches between blocks occurring without warning. The presentation of CS_2 predicted either a small (first and third blocks) or large US (second block). The presentation of CS_1 predicted delivery of a medium US in all blocks. Thus CS_1 corresponded to the "better" trial type in blocks 1 and 3, but not 2.

Anticipatory licking behavior indicated that the monkey adapted its behavior depending upon the relative amount of expected reward. Although the reward amount associated with CS_1 remained constant throughout the experiment, anticipatory licking decreased in block 2 and increased in block 3 – the blocks in which CS_1 trials had become relatively less (block 2) and more (block 3) valuable. Strikingly, many individual amygdala and OFC neurons also modulated their responses to CS_1 depending upon the block. Because this CS predicts the exact same reward in each block, these neurons cannot simply represent the sensory properties of a US associated with a CS. This finding demonstrates that amygdala and OFC neurons are often sensitive to the relative motivational significance of a CS, and not just to the sensory properties of its associated US or to the absolute value of the specific reward. Neurons in both the OFC and amygdala encode the relative value of CS_1 but OFC neurons significantly encode relative value earlier than amygdala neurons.

Cells in the amygdala and OFC code different properties during different time intervals during the trial and are consistent in valence when they code multiple properties. This implies that neurons are tracking state value: the overall motivational value of an organism's internal and external environment across time and sensory stimuli. Neurons that code relative value during the CS-trace interval and during reinforcement are also consistent in the valence that they code further supporting that these cells track state value. The neurons code with the same sign and strength whether the neuron is representing the relative value of the reward with no sensory input of the reward during CS or trace interval, or actually experiencing the reward during the US interval.

Further, amygdala and OFC neural activity was correlated with the animal's behavioral performance, suggesting that these neurons could form the basis for animal's behavioral adaptation during contrast revaluation. These neural representations could also support behavior in other situations requiring flexible and adaptive evaluation of the motivational significance of stimuli.

Table of contents

List of charts, graphs, and illustrations

- **Chapter II**
  -
  -
  -
  -
  -
  -
  -
  -
  -
  -
  -
  -
  -
  -
  -
  -

Acknowledgments

- I would like to thank my mentor Dr. C. Daniel Salzman for helping me through this entire process, for always catching all the bugs or errors in the analysis even before I could finish explaining how it was done, for all the jokes, for never letting me forget where I came from, and for all the good memories, and most of all for being a caring person and a great scientist.
- I am grateful to my thesis committee – Dr. Michael Goldberg, Dr. Stefano Fusi, Dr. Joshua Gordon, and Dr. Barry Richmond– for their thoughtful consideration of my work, and for their gift of time out of their incredibly busy schedules.
- I am grateful to Karen Marmon for her technical support. I would also like to thank all the current and former members of the Salzman lab for their support and help throughout the years. It has been a great place to spend these last six years. I might never work with such a talented and fun group again.
- Thank you to my family, especially my parents and my sister, for their advice, support over the years.
- Finally thank you to my husband Alex, I could not have done this without you. Thank you for all the emotional and scientific support. You have been integral to this work in every way.

Chapter I. Introduction

"Because it's all relative. You're pinned down in some filthy hellhole of a paddy, getting your ass delivered to kingdom come, but then for a few seconds everything goes quiet and you look up and see the sun and a few puffy white clouds, and the immense serenity flashes against your eyeballs - the whole world gets rearranged - and even though you're pinned down by a war you've never felt more at peace."

— Tim O'Brien, The Things They Carried (O'Brien 1990)

O'Brien is describing the mental process of relative valuation. This is a neural computation that computes an emotional/affective value of a recent event in comparison to a reference point, as frame of reference shifts our values of things shift along with it. In daily life this mental process happens so often and fluidly we might not even realize it. This mental process of relative valuation could underlie leaving your current spouse for someone else, or becoming a restaurant snob. This can happen because a stimulus is given a certain positive value and then later something better than that stimulus enters the environment and the original stimulus now diminishes in value in comparison.

To understand how the mental process of relative valuation might happen, we have to start with the steps the brain could take to reach that value judgment. First the brain must have the ability to recognize that a stimulus is behaviorally significant; this

requires a process through which the visual representation of a stimulus becomes associated with its affective value. The primary affective value of a stimulus varies depending on the current motivational state of the organism. The motivational state reflects the combination of internal needs and external possibilities and as the internal needs and external possibilities change or are reevaluated, the motivational state may change at any time. Cells which represent the current affective value are coding for the current reinforcement value of a specific stimulus given the current motivational state. A neuron encoding relative value cannot solely represent the sensory properties of the reward because those properties remain constant while the relative value of that reward compared to other rewards in the environment changes.

Types of value coding

The term “value” is used in neuroscience, economics and philosophy, and each field uses the term slightly differently; even within neuroscience, a different experimental task or parameters can adjust the meaning of the term. I am going to define below some different types of value and then later give examples of where and how those types of value are encoded in the brain.

Subjective value

Subjective value is an internal property that is defined by an individual. The subject gives a particular value to a stimulus based on its desirability given the circumstances. The value can be determined using many different factors such as usefulness, risk (or probability of delivery), effort needed to obtain it, amount of time

until receiving it. Subjective value integrates these factors listed above to assign value to a particular stimulus. The subjective value is the internal value a stimulus has to motivate behavior. Subjective value integrates information about internal state such as hunger or thirst with external factors such as reward amount to determine updates of value (Minamimoto, La Camera and Richmond 2009).

"Action value" vs. "stimulus value"

When neurons code "action value" of the available options they are encoding signals reflecting the predicted value of performing an action. When neurons code "stimulus value" they are encoding signals that a certain stimulus predicts a particular outcome. To differentiate between these two types of coding, the "value of action" and "value of the stimulus" have to be separated in the experimental paradigm. One example is that the monkey has a choice between two stimuli one with reward and the other with no reward, one presented on the left hand side of the screen and one on the right hand side respectively. The monkey must signal its choice with an eye movement. If a neuron increases firing when the monkey makes saccades to the right, it would be impossible to separate if it is coding rightward saccades or the value of the stimulus on the right. However, if the experiment is counterbalanced and the images alternate sides from trial-to-trial, then the two types of coding can be deciphered. A neuron that codes "action value" would differentiate between the two stimuli but only on one side. Alternatively, if the neuron increases its firing whenever the monkey chooses a specific stimulus, then it would be classified as coding "stimulus value", such as firing higher for the rewarded stimulus regardless of which side it is presented on.

However there can be interactions between value and action such as when the monkey has a choice between two stimuli A (low value) and B (high value). If the neuron increases firing when the monkey makes a rightward saccade to stimulus B, it can be coding a rightward saccade and increasing firing to the high value predicted outcome. However if it fires less when it makes a rightward saccade to stimulus A, and even less when it makes a leftward saccade to stimulus A. The neuron is coding rightward saccades and firing is increased for higher value outcomes.

An example of the distinction between "action value" and "stimulus value" is in a study of humans with focal frontal lobe damage (Camille, Tsuchida and Fellows 2011). Orbitofrontal damage disrupted the ability to sustain the correct choice of stimulus, but not of action, after positive feedback, while damage centered on dorsal anterior cingulate cortex led to the opposite deficit. These findings argue that there are distinct, domain-specific mechanisms by which outcome value is applied to guide subsequent decisions, depending on whether the choice is between stimuli or between actions.

This same result was also found in monkeys with OFC lesions and ACC lesions (Rudebeck et al. 2008). Both lesions disrupted decision making, but their effects were differentially modulated by the dependence on action- or stimulus-value contingencies. OFC lesions caused a deficit in stimulus but not action selection, whereas ACC lesions had the opposite effect, disrupting action but not stimulus selection.

Operationally defining value

We think it is important to define value operationally in the current work, as well as in previous studies from our group. We define positive value as those stimuli that elicit

approach or acquisitive behavior, and negative value as those stimuli that elicit defensive or withdrawal behavior. In the absence of choices which give an exact readout of stimulus value, we define positive stimuli as those that elicit approach behavior with a greater probability, intensity, duration or rapidity, and we define negative stimuli as those that elicit defensive/withdrawal behavior with a greater probability, intensity, duration or rapidity. In the work discussed in Chapter II I will only be using stimuli with positive value.

Relative value

In order to guide motivated behavior, humans and animals need to be able to flexibly evaluate the affective values of stimuli in the environment. To decide on an appropriate behavior, the nervous system could estimate the value of each potential action, convert it to a common scale, and use this scale to compare the different options and determine a course of action (Padoa-Schioppa 2011). Without an internal currency computed by the nervous system, an organism would be unable to assess the relative value of different events.

The relative value of a reward can change depending upon context. For example, an ice cream cone can be appealing in hot weather but if it is freezing outside it might not be, or if you just ate a bunch of ice cream the value could decrease as well. There are two different types of situations that can change relative value. One is the comparison of stimuli or actions with a fixed value to each other to facilitate choices (a choice between a vanilla or chocolate ice cream) and the other is the dynamic adjustment of the subjective value of a single stimulus or action in different contexts (ex. the presence of a larger

reward in the environment). In the work discussed in Chapter II I will use the second type of relative value modulation which we termed “revaluation” during a contrast revaluation task.

An economic framework of relative value

The branch of economics called prospect theory has shown that humans’ decisions are altered by contextual factors such as framing effects (*e.g.*, the presentation of a problem as losses or gains) and reference points - an internal zero point where everything above is considered a gain and everything below a loss (Tversky and Kahneman 1981). Human decision making does not generally follow a utility function, which uses multiplication between known probabilities and magnitude to calculate value. A perfectly rational decision-maker would make choices that obtain the maximum value as defined by the utility function. Humans and monkeys have access to the necessary parameters for rational decision-making – *e.g.*, probability and magnitude of reward – in various brain areas at the level of single neurons (Cromwell and Schultz 2003, Leon and Shadlen 1999, Platt and Glimcher 1999, Dorris and Glimcher 2004, Lau and Glimcher 2008). However, it is clear that they make value judgments using other factors, because humans and monkeys do not obtain maximum rewards when given the opportunity in a laboratory setting (La Camera and Richmond 2008).

Conditioning and reinforcement learning

Reinforcement learning is the process of learning about stimuli or actions solely on the basis of the rewards and punishments associated with them. To study reinforcement learning, researchers often use classical, or Pavlovian, conditioning

procedures (Pavlov 1927) or instrumental, also known as operant, conditioning procedures (Skinner 1938, Thorndike 1911). Classical conditioning differs from the operant conditioning in that the former does not contain an instrumental component; *i.e.*, the subject is not required to make a choice or take any action. In classical conditioning, the reinforcers (rewarding or aversive stimuli) are delivered independently of any actions taken by the animal. In instrumental conditioning, by contrast, the actions of the animal determine which reinforcement will be delivered.

During a classical conditioning procedure, subjects learn that an initially neutral conditioned stimulus (CS), such as a tone or abstract image, is followed by the presentation of an unconditioned stimulus (US), which is biologically relevant and has inherent positive or negative meaning, such as a food reward or an aversive shock. Through repeated presentations of a CS-US pair, subjects learn to predict a US based on a CS and generate a conditioned response (CR) to the CS, such as salivating or blinking. Several types of associations between a CS and US may be formed, such as those between the CS and the motor response elicited by the US, or between the CS and the sensory properties of the US.

Theories of reinforcement learning describe reward-based decision-making and adaptive choice of actions by the following three steps: (i) The organism estimates the value of a stimulus, defined as how much reward value (probability times volume) it is associated with. (ii) It selects an action by comparing the action values of multiple alternatives. (iii) It updates the new value of the action based on the expected or predicted value compared to the actual received value. If those two values are different the brain signals an error in those value representations and updates them to more accurately

predict future reinforcements. The magnitude of learning that occurs on each trial decreases as the reward becomes fully expected upon presentation of the stimulus.

The seminal Rescorla-Wagner (Rescorla and Wagner 1972, Sutton and Barto 1998, Schultz, Dayan and Montague 1997) model posits that the value representation should be updated on a trial-by-trial basis by "error signals" – *i.e.*, signals reflecting the difference between expected and received reinforcement. The function of the model is listed below.

$$\Delta V = \alpha (\lambda - V)$$

V is the associative strength of a CS on a given trial; ΔV is the change in strength of the association for a particular trial; λ is the maximum strength that the CS-US association can achieve (determined by the magnitude of the US); and α is a learning rate that takes on values between 0 and 1 (influenced by the intensity of the CS). At the start of learning, V is equal to 0; if learning is allowed to proceed towards its endpoint, V will asymptote at λ, and ΔV will trend towards 0.

Recent theories, such as temporal difference (TD) models, have extended this model so that reinforcement learning may be described quantitatively in real-time (Rescorla and Wagner 1972, Sutton and Barto 1998, Schultz et al. 1997). The function of the model is listed below.

$$V(t+1) = V(t) + \alpha * \delta(t)$$

V is the value function; α is a learning rate; and δ is the error signal, defined by $\delta(t) = R(t) – V(t)$, where R(t) is reinforcement as a function of time. Thus, in this type of model,

the value signal V(t) is reciprocally connected to the error signal, δ(t); *i.e.*, an error signal is generated when the derivative of the value signal is non-zero, and the error signal feeds back onto the value signal to render it more accurate.

The experiments discussed in Chapter II are of particular theoretical interest because of the neural calculations that must take place for the process of "revaluation" to occur. In this paradigm the contrast revaluation task, the reward associated with one stimulus is always the same size; therefore, expected reinforcement and received reinforcement remain the same upon block transitions, and this should not update the value function. Instead, the mechanism responsible for updating the representation of value must integrate information from other trial types to compute the "relative" value of the CS. Thus, the neural computation that takes place during this task needs to be modeled using a TD paradigm that incorporates an overall reinforcement rate term.

A paper by Hikosaka and colleagues revised the TD model, creating a contextual version of the model in which the value term is a function of the current sensory input and the context (Nakahara et al. 2004). In the task that these authors modeled, the probability of reward was conditional on the presence of reward in the previous trial. It was not a local context that could be maintained only within a trial, but a global context that had to be maintained in memory over many trials. A context term like the one used in their model incorporates reward information from previous trials, and might hold promise for modeling the type of learning mechanism needed to perform the contrast revaluation task.

Encoding of value

Encoding reward magnitude

Processing reward information is a function necessary for the survival of any animal foraging in a dynamic environment. In numerous areas of the brain, extending from the limbic system to the motor system, neuronal activity varies according to the size of the reward (usually liquid reward) for which a monkey is working (Cromwell and Schultz 2003, Hassani, Cromwell and Schultz 2001, Ikeda and Hikosaka 2003, Lauwereyns et al. 2002, Leon and Shadlen 1999, Platt and Glimcher 1999, Roesch and Olson 2003, Schultz 2002, Shidara and Richmond 2002a, Stuphorn, Taylor and Schall 2000, Thorpe, Rolls and Maddison 1983, Tremblay and Schultz 1999). The easiest quantifiable measure of reward for animals is the volume of juice, which animals can discriminate in sub-milliliter quantities (Tobler, Fiorillo and Schultz 2005). In a number of brain areas, neurons show increasing responses to conditioned stimuli (CSs) that predict higher volumes of reward than those that predict smaller volumes; these include the striatum (Cromwell and Schultz 2003) dorsolateral and orbital prefrontal cortex (Leon and Shadlen 1999, Roesch and Olson 2004, Wallis and Miller 2003a), parietal and posterior cingulate cortex (McCoy et al. 2003, Musallam et al. 2004, Platt and Glimcher 1999, Dorris and Glimcher 2004) amygdala (Paton et al. 2006b, Belova et al. 2007, Belova, Paton and Salzman 2008) and mid-brain dopamine neurons (Satoh et al. 2003, Tobler et al. 2005). Behaviorally, monkeys also show increased responses to large rewards, compared to small rewards, by decreased saccade latency (Takikawa et al. 2002), decreased error rates (Leon and Shadlen 1999), and increased anticipatory licking (Hassani et al. 2001).

Encoding of subjective value

A rewarding stimulus can decrease in subjective value when a better reward enters the environment; in psychology, this process is termed “contrast” (Flaherty, Turovsky and Krauss 1994). In rats, intake of a sweet solution is suppressed when the rat knows that an even sweeter solution will soon become available (Flaherty et al. 1994). Lesions of the amygdala eliminate this behavioral phenomenon (Gilbert and Kesner 2002). Most studies using this paradigm, however, only examined behavioral outcomes, and did not investigate the neurophysiological processes and circuitry involved.

Roesch and Olson previously found that neurons in the primate orbitofrontal cortex OFC are modulated both by the quantity of juice delivered to the animal and by the delay intervening before juice delivery (Roesch and Olson 2005). Wallis and colleagues found that the activity of neurons in the OFC depends not only on the quantity of juice delivered to a monkey, but also on the probability and on the physical effort exerted by the animal to obtain the juice (Kennerley, Behrens and Wallis 2011, Kennerley and Wallis 2009a). Another example of integration came from a study by Platt: in this experiment, monkeys chose between drinking a given amount of juice and drinking a different amount of juice while watching the image of a conspecific (Deaner, Khera and Platt 2005). Male macaques were willing to forgo some amount of juice for the opportunity to observe female perinea. Importantly, the activity of neurons in OFC encoded subjective value as defined by this behavioral trade-off.

Another form of subjective value is illustrated by two studies that use price to determine value. In the first study, they found that subjects reported analgesics relieved

pain more, if they were thought to be more expensive, despite the fact that the products was a placebo (Waber et al. 2008). In the second study, subjects were given several wines, and provided with information regarding the retail price of each (Plassmann et al. 2008). Subjects tasting wine they believed to be expensive found it significantly more pleasant than the same wine labeled as being cheap. Neural responses in medial orbitofrontal cortex correlated with the experienced pleasantness, rather than the identity of the wine.

Sometimes just framing and expectation can alter the subjective value of a stimulus. Rolls and colleagues gave subjects isovaleric acid, which has an odor similar to cheese, to subjects in an fMRI scanner, and accompanied it with the words 'cheddar cheese' or 'body odor' (de Araujo et al. 2005). They found that not only did subjects greatly prefer the scent when labeled 'cheddar cheese', but that activity in medial orbitofrontal cortex coded this subjective experience.

The amygdala and orbitofrontal cortex

Considerable evidence now indicates that individual neurons in the amygdala and OFC often encode information about the overall affective or motivational significance of USs associated with CSs (reviewed below in the section entitled "Role of OFC and amygdala in value and reward coding"), as well as the relative value of USs in OFC (reviewed in the section entitled "Relative value or absolute value representations in the OFC"). This is why we chose to focus our studies on these two brain areas during a task in which the relative value of a stimulus is altered but its sensory properties remain constant: to look for neural representations of relative value, as discussed in Chapter II.

Anatomy

The OFC and the amygdala are prime candidates for encoding relative value. These brain areas are involved in a major neural circuit that assigns values to stimuli and uses those values to guide emotional responses and decision-making (Nishijo, Ono and Nishino 1988b, Davidson 2002, Nakamura, Mikami and Kubota 1992, Nishijo, Ono and Nishino 1988a, Rolls et al. 1996, Hikosaka and Watanabe 2000, Rolls and Baylis 1994). Moreover, the OFC is reciprocally connected to the amygdala (Carmichael and Price 1995a, Ongur, Ferry and Price 2003, Ghashghaei 2002, Ghashghaei, Hilgetag and Barbas 2007, Barbas 2007) facilitating interactions that could play an important role in encoding relative value.

The amygdala is located in the anteromedial part of the temporal lobe. It lies ventromedial to the striatum and anterior to the ventral portion of the hippocampal formation. The primate amygdala may be divided into five major nuclei: the lateral (LA), basal, accessory basal, medial, and central nuclei. Cells in the lateral, basal, and accessory basal nuclei comprise the basolateral amygdala (BLA), and are made up of two types of neurons, projection neurons and interneurons (McDonald 1998).

The orbitofrontal cortex is defined as Walker's areas 11, 12, 13, and 14: the anterior, lateral, central, and medial orbital area respectively (Walker 1940). More recently, Price and colleagues refined the definitions of the orbitofrontal cortex to include further subdivisions, such as areas 11a, 13m, etc., based on cytoarchitectural features and anatomical connections (Carmichael and Price 1994, Ongur and Price 2000, Price 2007).

The amygdala sends projections most densely to the caudal regions of the OFC area 13. The OFC sends projections to numerous nuclei of the amygdala including the basal, accessory basal, and lateral nuclei, and the intercalated masses. The amygdala, meanwhile, receives projections from multiple structures that might modulate value representations depending upon context: *e.g.*, prefrontal cortex and hippocampus (Penick and Solomon 1991, Ghashghaei and Barbas 2002). The OFC receives innervation from polysensory cortices and higher sensory areas of every modality (Carmichael and Price 1995c, Cavada et al. 2000, Romanski, Bates and Goldman-Rakic 1999), with prominent inputs from adjacent gustatory and olfactory cortices. The OFC also has direct and indirect connections with limbic structures and the striatum (Haber et al. 1995, Carmichael and Price 1995b, Kondo, Saleem and Price 2005, Barbas and Blatt 1995).

Despite the tight anatomical interaction between the OFC and amygdala, the nature of the functional interaction between the regions remains unclear. In many respects, the OFC and amygdala resemble each other in their patterns of connections: both areas receive visual information from the rostral interior temporal cortex, and somatosensory information and gustatory information via the insula (Penick and Solomon 1991, Ghashghaei and Barbas 2002). Also, they both send outputs to several sites regulating autonomic and motor responses to stimuli, such as the lateral hypothalamus and nucleus accumbens. This could mean to a certain extent that the OFC and amygdala have potential for parallel processing, meaning they can simultaneously process the same stimulus either dissecting different elements of it or obtaining the same information from it and then combining this information for a value judgment of the stimulus.

Role in value and reward coding

Through lesion studies, the amygdala and OFC have been shown to be involved in incentive value: the positive motivational drive a stimulus has on behavior. Animals with lesions of either OFC (Baxter et al. 2000, Gallagher, McMahan and Schoenbaum 1999) or amygdala (Baxter et al. 2000, Gallagher et al. 1999, Hatfield et al. 1996) can no longer modify their behavior towards a reward-predicting cue when the outcome predicted by the cue changes in value. The similarities in the effects of OFC and amygdala lesions on behavior in certain tasks suggest that these two structures form a functional system involved in the acquisition and use of incentive information to guide goal-directed behavior.

The work that has contributed to this conclusion largely relies on one task in which the value of a stimulus is updated, known as reinforcer devaluation. In the Pavlovian version of the devaluation task, via conditioning, a cue comes to predict food. Then, the normally rewarding food is devalued (in the absence of the cue) by pairing the food with illness or satiation. After devaluation, normal animals spontaneously reduce responding in the presence of the cue that predicts availability of the "devalued" food.

There is evidence from both rats and monkeys that the amygdala and orbitofrontal cortex play a role in reinforcer devaluation. Rats with basolateral amygdala lesions are insensitive to post-conditioning changes in the value of the reinforcer, whereas rats with central nucleus lesions, like normal rats, are able to spontaneously adjust their conditioned responses to the current value of the reinforcer (Hatfield et al. 1996). Subsequently, Schoenbaum and colleagues found that the basolateral amygdala was critical for forming representations linking cues to the incentive properties of outcomes,

but not for maintaining these representations in memory (Pickens et al. 2003). Holland and colleagues clarified that when multiple reinforcers were used, but not under single-outcome conditions, post-training basolateral amygdala lesions disrupted the expression of devaluation performance in rats (Johnson, Gallagher and Holland 2009).

In monkeys the amygdala has also been shown to be necessary for reinforcer devaluation. Excitotoxic amygdala damage interferes with reinforcer devaluation effects (Malkova, Gaffan and Murray 1997). When the basolateral amygdala was inactivated using muscimol during selective satiation, devaluation was blocked; in contrast, muscimol infusion after satiation, so that it was present only during the testing period, did not impair devaluation (Wellman, Gale and Malkova 2005). Unilateral OFC- and amygdala-lesioned animals had attenuated reinforcer devaluation effects (Izquierdo and Murray 2004). Monkeys with excitotoxic amygdala lesions (Murray and Izquierdo 2007, Izquierdo and Murray 2007) are unable to refrain from choosing objects associated with the devalued reward. Thus, it seems that *updating* the value signal in the amygdala is required for reinforcer devaluation, but an active stimulus value representation in the amygdala is not necessary for expression of the updated preferences; these might instead be guided by expected outcome signals in the OFC or elsewhere in the absence of amygdala activity.

The OFC has also been shown to be essential for successful reinforcer devaluation. Monkeys with bilateral OFC lesions showed a significant attenuation of reinforcer devaluation (Izquierdo, Suda and Murray 2004). It was also shown in rats that OFC is critical for the maintenance of information about the current incentive value of

reinforcers, or the use of that information to guide behavior, in reinforcer devaluation tasks (Pickens et al. 2003, Pickens et al. 2005) .

Turning to the human brain, Dolan and colleagues presented two arbitrary visual stimuli, both before and after devaluation, in a paradigm of appetitive conditioning (Gottfried, O'Doherty and Dolan 2003). In the human amygdala and orbitofrontal cortex, responses evoked by the conditioned stimulus, as measured by fMRI, were decreased after devaluation, whereas responses to the non-devalued stimulus were maintained. The BOLD signal in amygdala and orbitofrontal cortex encoded the current value of reward representations. In another fMRI study, the authors found a significant correlation between the activation of a region of the orbitofrontal cortex and the decrease in subjective pleasantness when a liquid food is eaten to satiety (Kringelbach et al. 2003).

The data discussed above provide powerful evidence that the amygdala and OFC are both critical for reinforcer devaluation. The behavioral evidence from these studies indicates that the amygdala and OFC are both involved in a circuit that assigns value to stimuli and allows for updating of that value through experience.

The devaluation paradigm is a powerful technique to study the updating of reward value; however, in a set-up like the one used in the experiments discussed in Chapter II – in which only a single reward type (water) is used – we cannot use satiation to devalue the reward because the monkey simply will not perform. Instead we use the availability of other, larger rewards in the environment to contextually “devalue” the current reward. This form of devaluation is very different in the experimental design, but the principles and neural substrates could be similar based on the evidence presented in this section. The studies strongly suggest that when a CS-US association has a specific value and then

the value of the US is changed by something, in some cases satiation or illness the amygdala and OFC are involved in the updating of the association so that the CS, and the US, is devalued. These same mechanisms could be involved when a CS-US association is formed for a specific reward magnitude, and then a bigger reward enters the environment, requiring an integration of information about the new reward and the current one, as well as a updating of value, which, may take place in the OFC and the amygdala.

Further evidence that these signals exist in the amygdala and OFC come from previous work performed in the Salzman laboratory. We find neurons in the amygdala and OFC that respond more strongly to a CS that has been associated with reward (positive value-coding neurons), and other neurons that respond more strongly to a CS that has been associated with air-puff (negative value-coding neurons). Moreover, this value representation in the amygdala and OFC is not limited to recently learned CSs; rather, value-coding neurons respond in a consistent fashion to all events during a trial, including the rewards and air-puffs themselves (stimuli that are primary reinforcers) and a fixation point (a mildly positive stimulus conditioned over a long time period). These findings have led to the idea that neurons in the amygdala and OFC track "state value" – loosely defined as the overall value of an organism's situation – which may be modulated as different stimuli appear in the environment (Belova et al. 2007, Belova et al. 2008, Salzman et al. 2007, Paton et al. 2006a, Morrison and Salzman 2009, Morrison et al. 2011, Morrison and Salzman 2011). Finally, our group has shown that individual amygdala and OFC neurons often integrate information about impending reinforcement of both valences, and therefore do not simply represent the sensory properties of a US associated with a CS.

In all of these studies, value signals in the amygdala and OFC change as fast as or faster than monkeys learn new associations, as demonstrated by their behavior. The signals are available to the monkey before having to lick or blink, meaning the brain could be reading out these signals and using those representations to then drive behavior which most likely takes place in another brain region.

Another study performed in the lab investigated the dynamics of OFC and amygdala during learning (Morrison et al. 2011). Our group found that negative value coding neurons in amygdala update more quickly than negative value coding neurons in OFC, and positive value coding neurons in OFC update more quickly than positive value coding neurons in the amygdala. After learning the OFC always preceded the amygdala, and analysis of local field potentials (LFPs) also showed that the OFC was influencing the amygdala after learning.

This result implies that different populations of cells in the two brain areas signal the information on different time scales during learning. These results motivated the idea to record from OFC and amygdala during the contrast revaluation task to see if information from that task was represented by the two brain areas on different time scales. The contrast revaluation task has no aversive reinforcement, no reversals and might be more cognitively demanding, so it was an unknown what dynamics would occur between the OFC and amygdala.

Evidence of relative value is also found in a study by another group. In this experiment one block consists of two CSs one paired with juice and one with water. In another block there are another two CSs one paired with water and one paired with electric shock (Hosokawa et al. 2007) . Orbitofrontal neurons respond to the water in the

water, electric shock block, precisely as they do to juice reward in the juice, water block, tracking the relative value. The water is the same in both blocks but in one block it is the better of the two rewards and in the other block the worse of the two rewards. The neurons are not coding the sensory properties or a fixed value of water but instead are representing the relative value of water which changes based on the context and other rewards in the environment that it is compared to.

Relative value or absolute value representations in the OFC

It is currently a debated issue in the field whether neurons in OFC encode value on an absolute or relative scale. Two key papers addressing this issue, discussed below, have reached different conclusions that are difficult to reconcile.

Tremblay and Schultz (Tremblay and Schultz 1999) use an operant task in which the monkey has been over-trained on several visual cues. Each cue predicts one of several possible rewarding outcomes: one of either three different foods or three different liquids that are of known preference order (as determined by an interleaved choice task). Two cues are presented together and the monkey is instructed which cue to choose to obtain a reward. In this task, the activity of orbitofrontal neurons does not appear to encode the fixed physical properties of the predicted reward such as the sensory properties or a fixed reward value, but rather reflects the motivational value of one reward relative to the other available reward on that trial, as expressed by the behavioral preference. Just as each reward can have a higher or lower motivational value, behaviorally, relative to the reward with which it is compared, orbitofrontal neurons can be more or less activated by one reward, depending upon which alternative reward is available.

Meanwhile, Padoa-Schioppa and Assad (Padoa-Schioppa and Assad 2008a) take a different approach using a choice task: monkeys choose between two beverages offered in variable amounts. Monkeys' choices provided an operational measure of the values that the monkeys assigned to the two juices. They found three types of neuronal responses in OFC: "offer value" responses, which encoded the quantity or value of one of the two offered juices; "chosen value" responses, which encoded the value of the chosen juice independently of the juice type; and "taste" responses, which were binary responses reflecting which one of the two juices was chosen independent of the amount. Interestingly, they found that neuronal responses in the OFC are typically invariant for "changes of menu." For example, if a monkey chooses among juices A, B and C offered pair-wise, the activity of neurons encoding the value of juice B does not depend on whether juice B is offered against juice A or against juice C. Thus, Padoa-Schioppa and Assad argue that Juice B is represented in the same way in OFC regardless of the context. In contrast, the results of Tremblay and Schultz predict that juice B would be coded differently depending on whether it was offered against juice A or juice C, as the monkey had a different preference for the three juices and so the representation in OFC should reflect that relative value.

Using the theoretical approach of efficient statistics, it seems hard to imagine that having each neuron code one thing – for example, a fixed reward value associated with Juice B, regardless of the context – is biologically efficient. This is the case in terms of both 1) number of neurons available, a fixed number that is not large compared to the number of stimuli in the environment; and 2) the limited amount of information that can be encoded using spikes: only a certain amount of variability in firing rate can be

achieved, with fixed minimum and maximum points, producing a finite capacity for information that can be read out from changes in firing rate. One way the brain might deal with this limited capacity is to have cells code different information during different time intervals.

There were differences between the study of Tremblay and Schultz and that of Padoa-Schioppa and Assad that might account for some of the discrepancies in their results. First, the two groups recorded in partially overlapping areas of OFC, but Tremblay and Schultz were more lateral and anterior, recording in area 11 and 13L, whereas Padoa-Schioppa and Assad concentrated all of their recordings to the medial part of OFC, area 13m. Moreover, Padoa-Schioppa and Assad found a larger percentage of task-related neurons (35%) than Tremblay and Schultz (10%). It is unclear whether this would be the case if the two groups were recording from exactly the same area. Secondly, Tremblay and Schultz use a block design, compared to Padoa-Schioppa and Assad's interleaved choice task trials. In a block design, neurons might adapt to the context of each block, such that A: B blocks might be considered high-value blocks, whereas B: C blocks might be considered low-value blocks.

The work discussed in Chapter II is more similar in task design to that of Tremblay and Schultz than that of Padoa-Schioppa and Assad: like Tremblay and Schultz, I use a task, in which the monkey does not have a choice of reward outcome (although, unlike the prior study, it does not require any action on the part of the monkey). I also use a block design in which there are overall low-value and high-value blocks. However, the areas of OFC in which I recorded – mainly area 13m – are more similar to Padoa-Schioppa and Assad. Thus, this work might help reconcile some

differences between the two studies: specifically, because I find relative value coding cells in area 13m (see Chapter II), I can rule out differing sub-regions of OFC as a reason why Padoa-Schioppa and Assad found “absolute” value coding, but Tremblay and Schultz did not. However, like Tremblay and Schultz, I still have a block design, which would allow monkeys to learn new contexts adaptively.

Finally, note that a human analogue of the paradigm developed by Tremblay and Schulz was used with subjects during an fMRI scan (Elliott, Agnew and Deakin 2008). The analysis compared BOLD responses to two identical events, which differed only by the previous context. Medial OFC response to the same stimulus was greater when the stimulus predicted the more valuable of two rewards available in a given trial than when it predicted the less valuable. The central finding is consistent with the primate results and suggests that OFC neurons code relative rather than absolute reward value.

Neural mechanisms of encoding relative value

Normalization

One way the brain could encode relative value is through normalization. This process is found in the visual system. The theory states that there is normalization stage, where a given cell's response is divided by a quantity representing the pooled activity of a large number of other similar neurons (Heeger 1992). This was tested in the lateral intraparietal cortex (LIP) using three targets with one target in the receptive field (Louie, Grattan and Glimcher 2011). The results imply that neurons integrate information about

the surrounding regions of visual space and the data were best explained by a normalization model. This mechanism could be a general technique of cortical computation and could be a way for the brain to encode relative value.

Ratio vs. difference models

It is unknown how the brain compares two options and determines the relative value. It could be by taking the ratio or the difference between the two. Several models compute the probability of selecting a given option by comparing the expected value (EV) of each option. However, a subtle but important difference between two common rules used to compute the probability is often ignored. Specifically, one common rule type, the *exponential* rule, compares EVs via a *difference* operation, whereas another rule type, the *power* rule, uses a *ratio* operation. To test the validity of each rule type, multiple laboratories performed a choice tasks in which either the difference or the ratio between the reward values was altered relative to a control condition.

One study by Corrado and colleagues (Corrado et al. 2005) proposed a mechanistic model of choice that describes the behavior of primates in a dynamic foraging environment. They found that a model based on differential value describes the choice behavior of both animals in their study better than the fractional value. Another study suggests that human participants are able to compare expected rewards by either operation – ratio or difference – but that altering the ratio between rewards leads to the most drastic changes in behavior (Worthy, Maddox and Markman 2008). This notion is consistent with the work of Dorris and Glimcher in LIP (Dorris and Glimcher 2004), who found that when reward amounts were doubled, but the relative ratio relationship was

kept constant, neurons in LIP did not change their responses for all trial epochs. Meanwhile, Lau and Glimcher (Lau and Glimcher 2005) created a response-by-response model that treats as separable processes the effects due to past reinforcers and past choices. They found that decision rules use differences in values.

Range adaption

Values computed in different behavioral contexts can vary by orders of magnitude. For example, the same individual might choose sometimes between different sodas and other times between different houses for sale. If neurons encode values in a linear way, this could pose a serious challenge due to a limitation in firing rate range. When reward is likely to occur within a specific range, focusing the sensitivity on the predicted range would optimize the discrimination of small reward differences.

One way to overcome the trade-off between the ability to code largely different rewards and still be able to discriminate them optimally is to adapt reward sensitivity dynamically to the available rewards, through gain adaptation. Two studies have demonstrated the existence of such mechanisms (Padoa-Schioppa 2009, Kobayashi, Pinto de Carvalho and Schultz 2010). In the study by Schultz and colleagues, animals performed a task in which a fixation cue predicted the standard deviation of the probability distribution of juice volumes, while the expected mean volume was kept constant. A subsequent cue specified the exact juice volume obtained for a correct saccade response. Population responses of orbitofrontal neurons that reflected the predicted juice volume showed adaptation to the reward distribution. In the study by Padoa-Schioppa, monkeys chose between different juices and their choice patterns

provided a measure of subjective value. Value ranges were varied from session to session and, in each session, OFC neurons encoded values in a linear way. The neuronal activity range did not depend on the value range. Thus, Padoa-Schioppa argued, the activity of each neuron adapts to the range values it encodes, but does not depend on other available goods (Padoa-Schioppa and Assad 2008b). In both studies, the authors found that the encoding of value undergoes range adaptation such that a given range of firing rates represents different ranges of values in different behavioral contexts.

Chapter Summary

The current work hopes to clarify three issues that exist in the field. First, as discussed above, most previous studies using a "contrast" paradigm have examined only behavioral outcomes, and did not investigate the neurophysiological processes and circuitry involved. Second, it is a currently unresolved issue in the field whether neurons in OFC encode value on an absolute or relative scale, and it remains unknown whether the amygdala tracks relative value when the US associated with a CS changes in relative value, but absolute reward value remains constant. Third, the similarities in effects of OFC and amygdala lesions in certain tasks suggest that these two structures form a functional system involved in the acquisition and use of incentive information to guide goal-directed behavior; therefore, the proposed experiments may shed light on the workings of this neural circuit by obtaining neural recordings simultaneously from both brain areas during an incentive task.

In **Chapter II**, I present a body of data demonstrating that individual neurons in the primate OFC and amygdala track the relative value of a US associated with a CS

(CS1) that predicts the exact same reward in each block. I further show that the rates at which this representation occurs varies with brain area. OFC neurons encode the relative value of CS1 significantly faster than amygdala neurons. These neurons also fit with theories of state value defined as the value of the overall situation of an organism at a given moment, evidenced by neurons that are consistent in valence when they code multiple properties, and by neurons representing the relative value of the reward with no sensory input of the reward during CS or trace interval, or actually experiencing the reward during the US interval.

Further, amygdala and OFC neural activity was correlated with the animal's behavioral performance, suggesting that these neurons could form the basis for animal's behavioral adaptation during contrast revaluation. These neural representations could also support behavior in other situations requiring flexible and adaptive evaluation of the motivational significance of stimuli.

Chapter II. Neural representations during the contrast revaluation task

Introduction

Relative valuation allows the brain to flexibly adapt to changing environments where rewards can be more or less valuable depending on the situation. In a situation where one could win either $10 or $100, $10 might seem to be a disappointing outcome. On the other hand, if in a different situation the only two possible outcomes were $10 and $1, $10 might seem rewarding. The $10 did not change in actual fixed monetary value but its subjective value in comparison to other potential rewards changed. The plasticity of the subjective value of a particular reward due to relative valuation allows the current representation of the reward to be subjectively accurate and useful under the specific state of the environment.

The orbitofrontal cortex (OFC) and the amygdala are prime candidates for encoding the signals underlying relative value in the brain. These brain areas are involved in a major neural circuit that assigns values to stimuli and uses those values to guide emotional responses and decision-making (Nishijo et al. 1988b, Davidson 2002, Nakamura et al. 1992, Nishijo et al. 1988a, Rolls et al. 1996, Hikosaka and Watanabe 2000, Rolls and Baylis 1994, Padoa-Schioppa and Assad 2008b, Kennerley et al. 2011, Padoa-Schioppa and Assad 2006, Cardinal et al. 2002, Sugase-Miyamoto and Richmond 2005). Moreover, the OFC is reciprocally connected to the amygdala (Carmichael and Price 1995a, Ongur et al. 2003, Ghashghaei 2002, Ghashghaei et al. 2007, Barbas 2007)

facilitating interactions that could play an important role in encoding relative value. The amygdala sends projections most densely to the caudal regions of the OFC known as area 13. The OFC sends projections to numerous nuclei of the amygdala (including the basal, accessory basal and lateral nuclei and the intercalated masses). It remains unknown whether the amygdala tracks relative value when the relative value of a stimulus changes, but its absolute reward value remains constant. The dynamics of encoding relative value in the OFC and amygdala is an area that has not been explored.

To investigate these two questions, we recorded activity in the amygdala and OFC in two awake, behaving monkeys while they performed a contrast revaluation task that manipulates the relative value of a stimulus, while holding its absolute value constant (Figure 1A). During each trial, the monkey was presented with one of two visual stimuli (or conditioned stimuli, CS) consisting of novel, distinct fractals. Following the visual stimulus and a delay, a liquid reward was delivered. The task consists of three trial blocks. In the first block, CS2 is followed by a small reward and CS1 by a medium reward. In the second block, CS2 is followed by a large reward while CS1 continues to be followed by the same medium reward. In the first block, the medium reward was the larger of the two rewards whereas in the second block, it is the smaller reward. In the third block, the initial reinforcement contingencies from block 1 are reinstated. Thus, in this task, the relative value of CS1 (as computed by comparing its associated reward to the other available reward) is manipulated while its absolute value is held constant.

Materials and methods:

General Methods

For these experiments, two rhesus monkeys (*Macaca mulatta*) weighing 10-11 kg were used. During experiments, monkeys sat in a Plexiglas primate chair (Crist Instruments), facing a computer screen on which visual stimuli was displayed. The monkey's head was restrained so that eye position could be monitored using an infrared eye tracking system (ASL/Eyelink). Licking was measured by passing an infrared beam between the monkey's mouth and the reward delivery tube; during each lick, the monkey's tongue interrupts the beam. The TEMPO software package (Reflective Computing) was used for experimental control. The monkeys did not work for all of their daily water during the task (remaining water was given in the cage) to ensure that experiments took place in a water-deprived and motivated state and to minimize the effects of satiation.

Prior to the start of experiments, magnetic resonance imaging (MRI) was used to locate the amygdala and OFC. A surgery was performed to install an MRI-compatible head-holder and recording chamber. After recovery from surgery, MRI was again used to verify chamber placement, and measure distances from the bottom of the chamber to brain landmarks in order to guide electrode placement. Finally, a scan was performed with an electrode inserted in the brain in order to check distances and verify that we reach the amygdala and OFC.

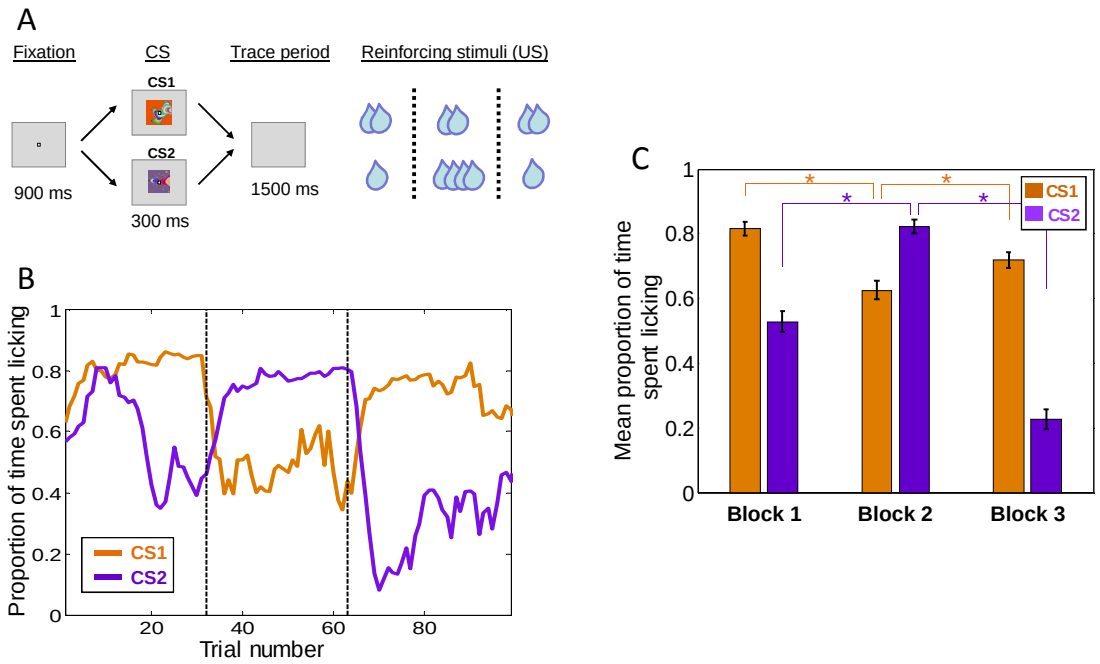

Figure 1: Contrast revaluation task and behavior.
(A) Sequence of events for the three blocks. Following fixation, one of the two stimuli is presented. After a trace period delay, reward reinforcement is delivered according to the block and stimulus presented. CS2 is associated with a small reward in blocks 1 and 3 and a large reward in block 2. CS1 is associated with a medium reward in all three blocks. Blocks are switched at a time unsignaled to the monkey after the monkey's behavior reflects learning of the new CS-reward associations in effect.
(B) Licking behavior from a representative session. The proportion of time spent licking during the CS and trace intervals is plotted as a function of trial number for CS1 and CS2 separately. Block switches are indicated by vertical dashed lines. The licking rate was smoothed using a 5-trial sliding window. The decrease in licking for CS1 in block 2 reflects the monkey's relative valuation of this stimulus when a more valuable one, CS2, enters the environment.
(C) Average licking rate across all experiments for the two stimuli in each block. The monkeys' learning of the reward associated with CS2 is indicated by a significant increase in licking for CS2 in block 2 relative to blocks 1 and 3. Relative valuation is shown by a significant decrease in licking for CS1 in block 2 relative to blocks 1 and 3 (stars : $p<0.01$, one-way ANOVA).

Behavioral task

Two monkeys were trained on a contrast revaluation conditioning paradigm; a schematic diagram of the proposed task is shown in Figure 1A. In trace conditioning, a short interval – the trace period – is inserted between the presentation of the CS and the delivery of the US (reward). The trace period is important in this experimental design for the following reasons: a) it allows time for the development of anticipatory licking behavior indicative of learning and b) it facilitates separate analysis of neural signals related to US expectation/prediction, and signals related to delivery of the reward itself.

During each experimental session, the monkey learned the values of two novel visual stimuli. For visual stimuli, we used abstract fractal patterns to ensure that the stimuli did not initially hold any emotional meaning for the monkey. At the beginning of the session, each stimulus in a pair is assigned a value – one a small reward and one a medium reward. A reward consists of drops of water: the bigger the reward, the longer the reward system takes to deliver it (100 ms ≈ 50 µl). The size of the rewards is as follows for monkey T, small reward (50 µl of reward), medium reward (250 µl of reward), and large reward (1250 µl of reward). The size of the rewards for monkey P, was small reward (20 µl of reward), medium reward (100 µl of reward), and large reward (500 µl of reward). On a typical trial, the monkey centers its gaze at a fixation point in the center of the screen, and is required to maintain fixation for 900 ms before presentation of the image. The monkey must also fixate during stimulus presentation, which is kept short (300 ms) so that the monkey does not have a chance to break fixation

for a less desirable trial. After presentation of the CS, a trace interval lasting 1500 ms ensues (no fixation required), followed by delivery of the US: a liquid reward of small, medium, or large size (defined above) depending upon the association with the stimulus.

During each experimental session, we monitored licking behavior in order to gauge the monkey's learning of the stimulus values. After the monkey had learned the small and medium values of the image (usually at ~5-10 trials of each type) and stably maintained behavioral learning (~20-30 trials of each type), we changed the assigned values of CS2 from small to large reward and the medium reward kept constant. After the monkey learned these new values of the image, we changed the assigned values of each stimulus back to small and the medium. The block switches occurred in an unsignaled, unpredictable manner, such that the monkey must learn the new task contingencies through experience.

Data collection

We recorded neural activity from 96 neurons in the right OFC and 166 neurons in the right amygdala of two rhesus monkeys (*Macaca mulatta*): 62 OFC cells and 100 amygdala cells from a 11 kg male (monkey T); 34 OFC cells and 66 amygdala cells from a 10 kg male (monkey P). In each recording session, we individually advanced up to four tungsten microelectrodes (impedance: approximately 2 MΩ; FHC Instruments) into each brain area using a motorized multielectrode drive (NAN). We used the Plexon system for signal amplification, filtering, digitizing of spike waveforms, and spike sorting using a principal component analysis platform (online with offline verification). We analyzed all

well isolated neurons; monkeys performed a fixation task or no task during the search for well-isolated neurons.

Data Analysis

Construction of behavioral curves

To construct licking curves we calculated the amount of time spent licking during the CS-trace interval and divided by the interval duration (1800ms). The licking curves were smoothed using a 5 trial sliding window. To obtain population licking, we calculated the mean licking for each CS and block during each experiment and averaged across experiments.

Classifying Cell Coding Properties

To characterize the coding properties of each cell recorded, we performed a regression analysis with the following factors: image identity, block, an interaction between image identity and block, and a constant term, with the corresponding coefficients β_{IM}, β_{BL}, β_{INT} and α, respectively. The image factor vector consisted of 1s for CS1 trials and 0s for CS2 trials. The block factor vector consisted of 0s for blocks 1 and 3 and 1s for block 2 trials. The interaction between image identity and block factor was the product of the image and block factors. The equation for the model is as follows:

$$FR = \alpha + \beta_{IM}*Im + \beta_{BL}*Block + \beta_{INT}*ImxBlock + Err$$

"Reward magnitude coding" is defined as: $\beta_{BL} = -\beta_{INT}$ ($\neq 0$), coefficients β_{BL} and β_{INT} are both non-zero (significant block and interaction factors), equal and opposite in

sign. When β_{BL} and β_{INT} are equal and opposite, there is no change for CS1 from block 1 to 2 (see Fig. 4).

"Block coding" is defined as: $\beta_{BL} \neq 0$, $\beta_{INT} = 0$ β_{BL} is non-zero to account for a change from block 1 to block 2 for both CSs and β_{INT} is zero so that the change from block 1 to 2 is equal for CS1 and CS2 (see Fig. 4).

"Relative value coding" is defined as: $\beta_{INT} \neq 0$, and $\beta_{BL} \neq -\beta_{INT}$. β_{INT} is non-zero to account for a change in firing for CS1 in block 2 that is different from the change in firing to CS2, and therefore not merely a block effect, while $\beta_{BL} \neq -\beta_{INT}$ ensures that the difference in firing between CS1 and CS2 in block 2 is not due to an unchanged firing rate to CS1 across the experiment. An additional condition, $\beta_{BL}*(\beta_{BL}+\beta_{INT}) \leq 0$, requires that the change in firing to CS1 in block 2 occurs in the opposite direction from the one for CS2 (as one stimulus becomes more valuable, the other one becomes less so), which simply translates consistency in the coding of relative value (see Fig. 4).

Sliding regression and optimal window of coding

Unless stated otherwise, we characterized the firing patterns of all recorded cells by using a sliding regression with the factors described above. We used a sliding window of 200ms stepped every 10ms throughout the trial. We identified an optimal window of coding (relative-value, reward magnitude etc…) as the longest stretch within the trial where 70% of 30 consecutive bins satisfy the criteria (defined above) for that specific type of coding. Figure 2 is an illustration of which signal the optimal window selects. Figure 2 shows the optimal window outlined in red on the firing rate plots and a plot of the significance of the criteria for that specific cell property. Of note, it is apparent in the

PSTHs of Figure 2C that this cell also encodes relative value at a later time during the trial (optimal window of coding not shown here). This cell was then classified as relative value coding in addition to block coding. This type of dual encoding within the same cells will be addressed later in this work.

Dynamics of the OFC and amygdala

For each cell identified as “relative value coding”, we collected the firing rate within the optimal window of coding for CS1 trials only. The firing rates were z-scored and aligned to the transition from block 1 to 2, then sign-flipped, according to the cell’s valence, to go from low to high. All the cells in each brain area were averaged together and normalized by subtracting the average firing rate of the last 10 trials before the transition and dividing that by the average firing rate of trials 15-25 after the transition minus the average firing rate on the last 10 trials before the transition. This procedure, analogous to a max-normalization, ensured that the mean firing rate values started at 0 at the time of the transition and asymptoted at 1 once learning has taken place, usually after 15 trials of each CS. CS1 trials after the transition were ranked according to the number of CS2 trials experienced so far (rather than the total number of trials since the transition), since only CS2 trials can be used to infer that block 2 is in effect.

In order to compare how quickly cells in the OFC and in the amygdala update their coding after the transition between blocks 1 and 2, we fit the mean, normalized

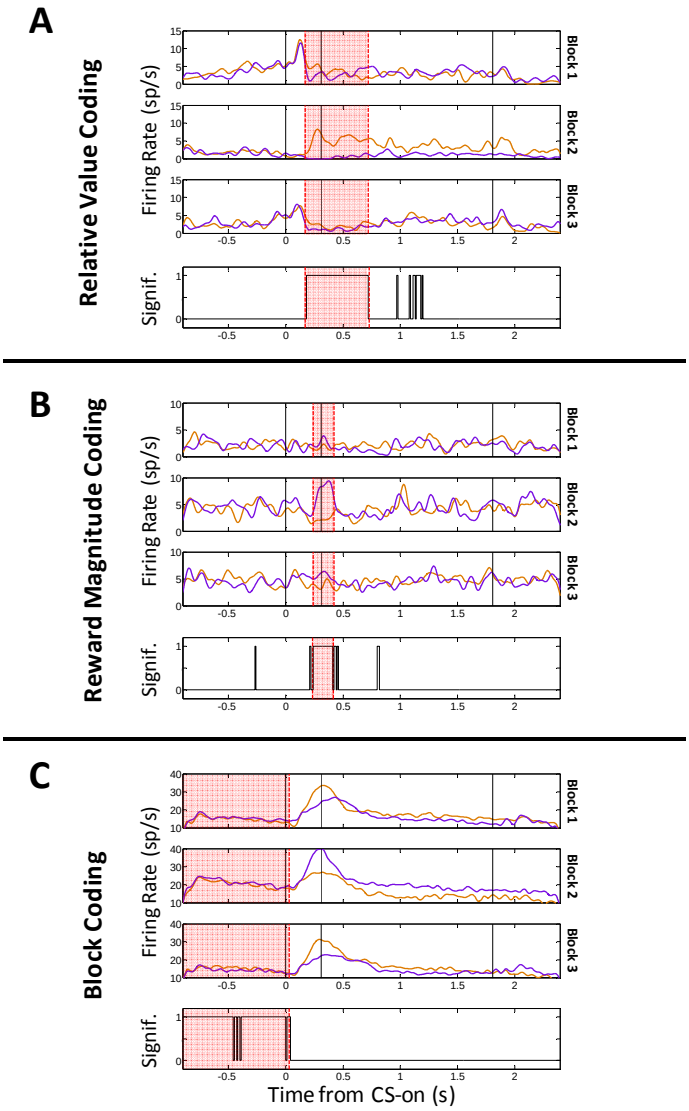

Figure 2: Optimal window of coding

An optimal window of coding was defined as the longest stretch within the trial where 70% of 30 consecutive bins satisfy the criteria for a specific type of coding.

(A) Example relative value coding cell and its identified optimal window of coding. Top three panels: PSTHs for each of the three blocks for CS1 trials (orange) and for CS2 trials (purple). The optimal window of coding is represented by the colored area between red dashed lines. Bottom panel: schematic plot representing when the criteria for relative value coding reached significance during the trial: 1 significant, 0 not significant.

(B) Example reward magnitude coding cell and its identified optimal window of coding. Same conventions as in A.

(C) Example block coding cell and its identified optimal window of coding. Same conventions as in A.

firing rate with a cumulative Weibull distribution function given by:

$$f(x) = 1 - e^{-(x/a)^b},$$

where the scale parameter *a* is informative of how quick the transition happens and *b* is a shape parameter. The derived scale parameters were compared by permutation test.

The permutation test used in this analysis randomly assigned each cell to either the OFC or amygdala population and then performed the same operations as just described. This was done 1000 times and the two scale parameters found at each iteration were compared. The p-value returned by the test corresponds to the proportion of times that a difference in the scale parameters at least as large as the true one was found by randomly shuffling cells between the two populations.

Fisher Method

We combined the p-values associated to the correlation coefficients between firing rates and licking rates using Fisher's combined probability test. This method consists in adding the logarithms of the k independent p-values p_i according to the formula:

$$\chi^2 = -2\sum_{i=1}^{k} log_e(p_i)$$

When all the null hypotheses are true, *i.e.*, when the p_i are not significant, χ^2 has a chi-squared distribution with $2k$ degrees of freedom. This is used to test the significance of the population of p-values as a whole. Indeed, when the p_i tend to be small (*i.e.*, the

correlation coefficients are significant), χ^2 will be large and deviate from a chi-squared distribution.

Results:

Behavior

In order to assess the monkeys' subjective valuation of the stimuli, we measured their anticipatory licking – an approach behavior that indicates expectation of reward – in response to each visual stimulus, during the CS presentation period (300ms) and the trace period (1500ms). We had a lick tube in front of the monkey's mouth to deliver rewards and between the monkey's mouth and lick tube we placed an infrared beam such that when the monkey would stick its tongue out in anticipation of a reward the beam is interrupted allowing us to calculate the proportion of time spent licking for the CS and trace periods combined (1800ms). Figure 1B shows the anticipatory licking behavior during a representative experiment for each CS independently. In order to determine whether the licking behavior reflects learning of the CS-reward associations in each block, we averaged the licking rate for each CS and each block and combined these rates across experiments (Figure 1C). The first 15 trials of each block were excluded to restrict analysis to post-learning trials. The monkeys systematically licked more in response to the CS that predicted the larger of the two rewards in each block (Wilcoxon rank-sum test, $p<0.01$). We then performed, for each image, a one-way ANOVA with a main effect of block and post-hoc tests (Tukey) to test for adaptive licking to a given CS according to the situation, or block. We found a significant effect of block, thus showing a modulation of the subjective valuation of each image in each context ($p< 0.01$) Indeed, the monkeys'

learning of the association of CS2 with changing reward size is indicated by lower licking in block 1 and 3 and higher licking in block 2. Meanwhile, learning of the relative value of CS1 is indicated by higher licking in block 1 and 3 and lower licking in block 2. Licking decreases for CS1 in block 2 despite the fact that the absolute reward amount associated with CS1 remains constant (Figure1C, stars). This result implies that the monkey's anticipatory licking is not just a response to the sensory properties or the volume amount of reward, but instead a response to a relative value modulated by other rewards in the environment (*i.e.*, by the context within the block of trials).

Recording Sites

We used magnetic resonance imaging (MRI) to guide our recording electrodes to the targeted structures: the amygdala and OFC. Amygdala recording sites were in the central, basal, accessory basal, or lateral nuclei. OFC recording sites were mainly in area 13m. We obtained data during the monkey's performance of the contrast revaluation task for 166 amygdala neurons. (Monkey T, 100 neurons, Monkey P, 66 neurons), and 96 OFC neurons (Monkey T, 62 neurons, Monkey P 34 neurons) (Figure 3)

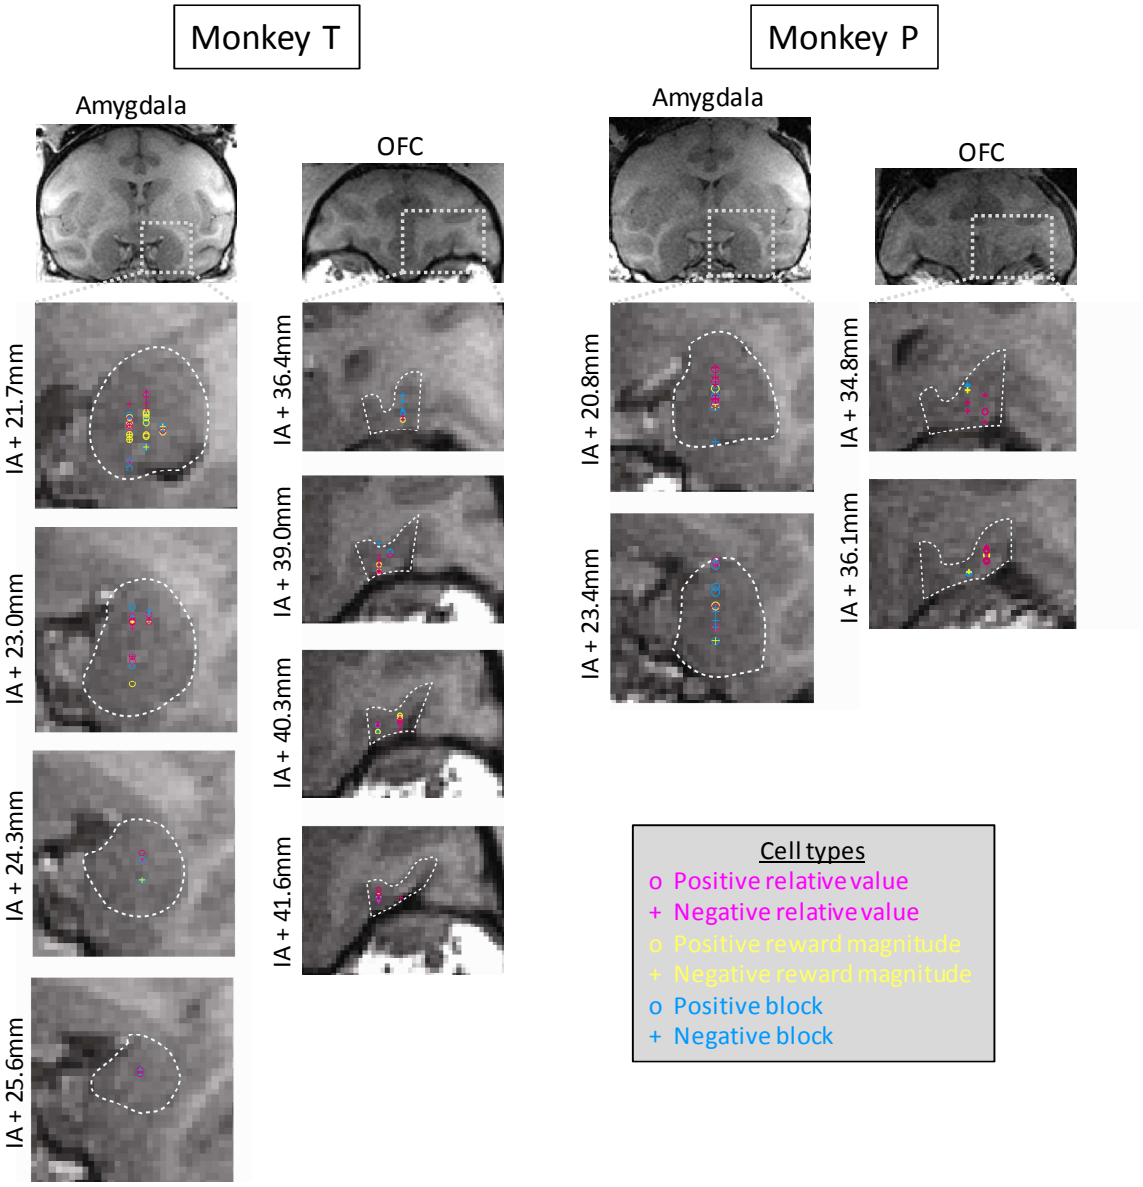

Figure 3: Reconstruction of the recording sites
The recording location for each cell of this study is represented on the corresponding MRI slice, with its category indicated by different colors and symbols. The amygdala and area 13 of OFC were identified according to standard conventions and are outlined in white for clarity. The coordinate of each plane along the anterior-posterior direction is indicated relative to the inter-aural plane (IA).

Characterizing cell properties

There are multiple aspects of the task that neurons can be selective to. In order to characterize the neural responses we recorded, we performed a regression analysis with the following factors: image identity, block, an interaction between image identity and block and a constant term, with the corresponding coefficients β_{IM}, β_{BL}, β_{INT} and α, respectively. Because of how the factor vectors of the regression were designed (see methods), and if we ignore the constant term, the firing rate in response to CS1 in blocks 1 and 3 is modeled by β_{IM} (Figure 4). Similarly, the firing rate for CS1 in block 2 is modeled by $\beta_{IM}+\beta_{BL}+\beta_{INT}$ and that of CS2 in block 2 simply by β_{BL}. Of particular interest is the difference in firing rate between blocks 1/3 and block 2 for CS1, since it can be seen as a measure of relative value coding in a cell. In our regression model, this quantity is modeled by $\beta_{BL}+\beta_{INT}$ (Figure 4). In a similar way, β_{BL} can be used as a measure of the difference in firing rate for CS2 between block 2 and blocks 1/3, when the absolute amount of associated reward with the CS changes (Figure 4). We define positive and negative cells as cells that increase or decrease, respectively, their firing rate in response to a larger reward.

Cells were classified according to the trial properties that where encoded in their firing rate. This was done by defining, for each type of coding, a set of criteria of the three parameters (β_{IM}, β_{BL} and β_{INT}) of the regression model (see methods for details). Because the regression was performed on a sliding time window stepped across the trial, the different types of coding could be specified on a time-bin basis. This allowed us to

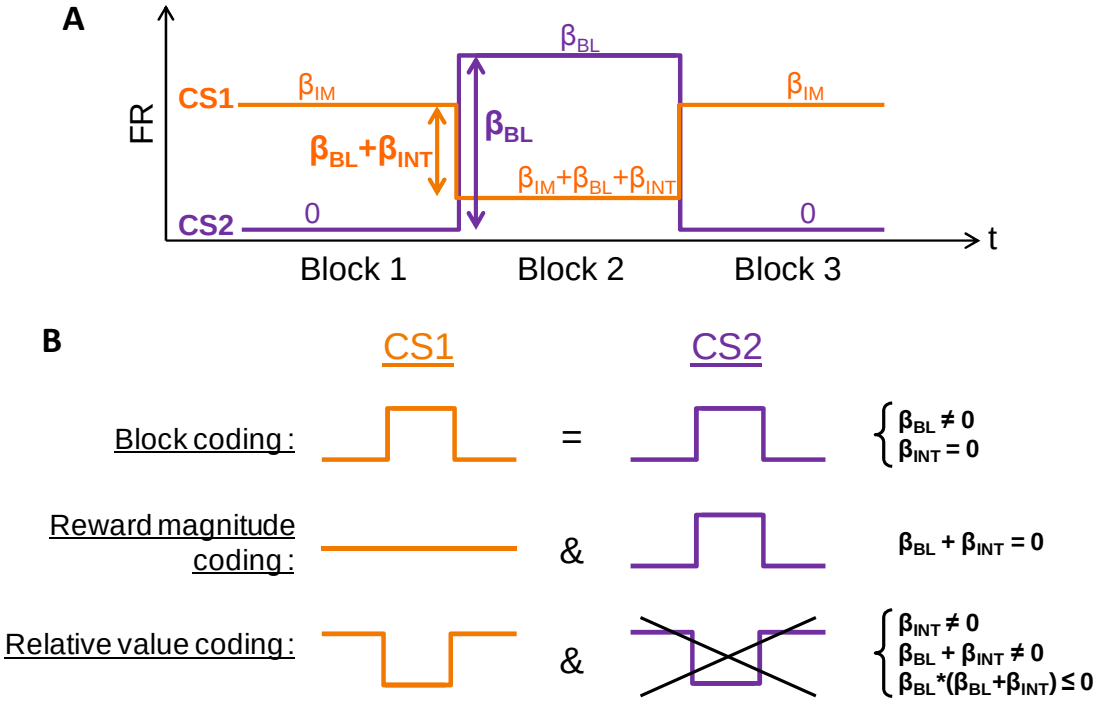

Figure 4: Schematic diagram of regression model and criteria for cell classification
(A) Schematic firing rate from a fictive relative value coding cell for CS1 trials and CS2 trials separately. We show how the different parameters of the regression model fit the different portions of the firing rate.
(B) Schematic firing rate for CS1 and CS2 for each type of coding as well as the corresponding criteria for the parameters of the regression model.

define an optimal window of coding as the longest stretch during the trial where the criteria for a given type of coding were satisfied (see methods and Figure 2).

Neurons were categorized as “relative value coding” if their firing rate tracked the relative value of CS1, that is, if the firing rate on CS1 trials changed in block 2 compared to blocks 1 and 3. even though the reward associated with CS1 has remained constant.

Here we present two examples of relative value coding cells in the amygdala and in OFC (Figure 5). The amygdala cell (Fig. 5A and 5B) is a positive relative value coding neuron, increasing its firing for larger rewards compared to smaller rewards. If we focus on the optimal window of coding, this is demonstrated by the increased response to CS2 in block 2, when it is associated with a large reward, compared to blocks 1 and 3, when it is associated with a small reward. Relative value coding is illustrated by the response to CS1, which decreases in block 2 when it is the smaller of the two rewards available in that block, in accordance with its relative value and not its reward magnitude, which remains constant in all three blocks. Conversely, the OFC cell presented in Figure 5C and 5D is a negative relative value coding neuron: it increases its firing rate for smaller rewards compared to larger rewards. This cell responds in the opposite direction of the positive cell. The cell decreases its firing for CS2 when it is associated with a large reward in block 2 and increases its firing for CS1 in block 2 when it is the smaller of the two rewards available. During the CS and/or trace intervals, we found 49 (29%) of such cells in the amygdala (24 positive, 25 negative), and 35 (36%) in OFC (13 positive, 22 negative). The fact that responses change in block 2 for CS1 even though the reward value has not changed suggests that knowledge of the total reward available within this block modulates estimation of value. This finding demonstrates that the amygdala and

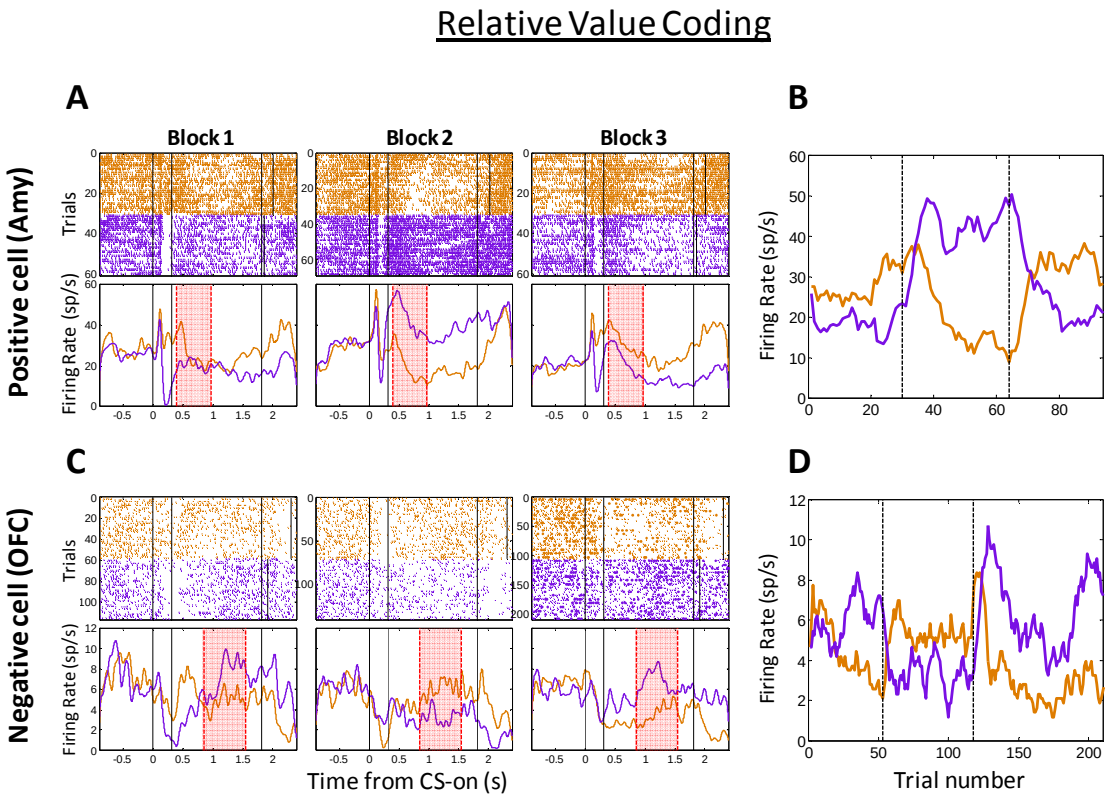

Figure 5: Example relative value coding cells
(A) Positive relative value coding cell from the amygdala. Top row: raster plots for CS1 trials (orange) and CS2 trials (purple) for each of the three different blocks of the experiment. Bottom row: PSTHs with optimal window of coding indicated by red shading between red dashed lines.
(B) Firing rate during the optimal window of coding as a function of trial number for the same amygdala cell. Orange line, CS1 trials; purple line: CS2 trials. Block switches are indicated by vertical dashed lines. Firing rate was smoothed using a 5-trial sliding window.
(C,D) Negative relative value coding cell from OFC. Same conventions as in A and B.

OFC respond to relative value (an abstract, subjective value) and not just to the sensory properties of its associated US or the absolute amount of the specific reward. Information about rewards in the task is integrated across trial-types and updated to modify the value of the medium reward in block 2.

Neurons were categorized as “reward magnitude coding” if their firing rate tracked the absolute amount of reward associated with the CS and not the changes in relative value. This means that the firing rate was affected by the change in reward associated with CS2 in block 2 (going from small to large) but remained unchanged for CS1, whose associated US stayed constant (medium to medium).

Figure 6 shows two example reward magnitude coding cells from the OFC and the amygdala. The OFC cell (Fig. 6A and 6B) is a positive reward magnitude coding cell: it increases its firing rate in response to CS2 in block 2, relative to blocks 1 and 3,. while its response to CS1 is not modulated. The amygdala neuron (Fig. 6C and 6D), on the other hand, is a negative reward magnitude coding neuron and shows the opposite response to CS2 . We identified a number of positive and negative reward magnitude coding cells in each brain area during the CS and/or trace intervals: 47 cells (28%) in amygdala (27 positive cells, 20 negative cells), and 15 cells (16%) in OFC (11 positive, and 4 negative).

Finally, we categorized neurons as “block coding” if they did not differentiate between CSs but still modified their firing rate according to which block was in effect, increasing or decreasing their activity in block 2 compared to blocks 1 and 3. These cells

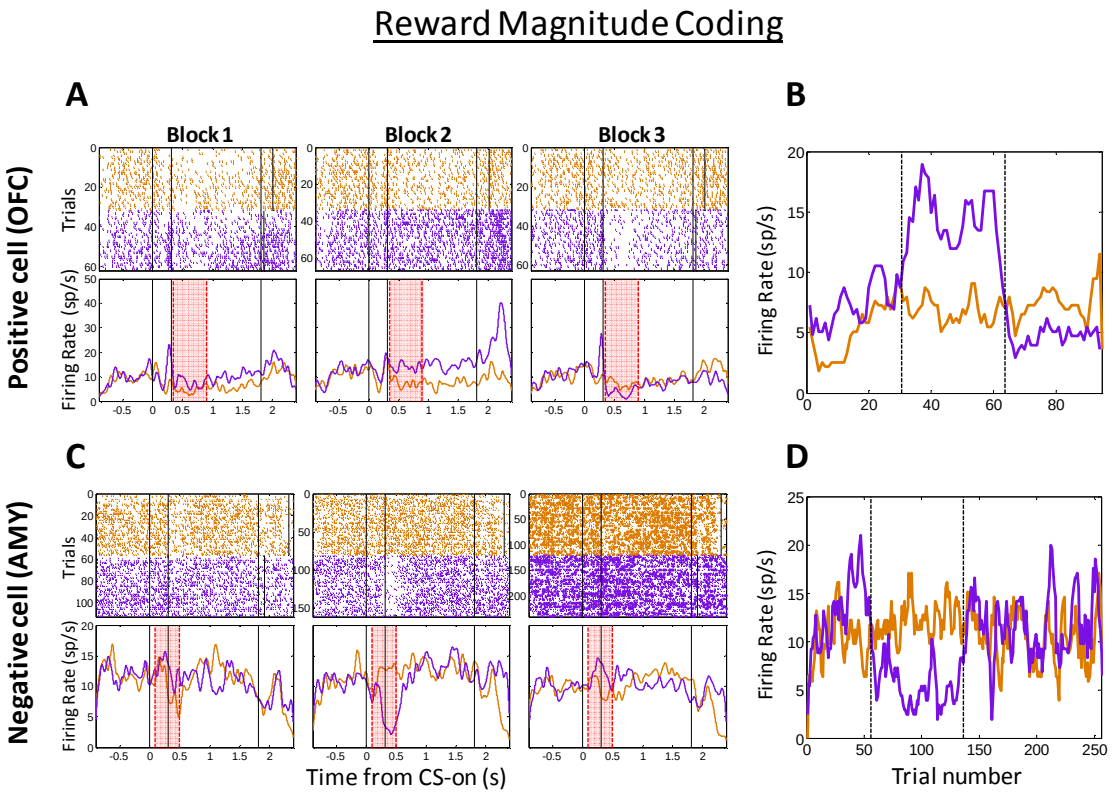

Figure 6: Example reward magnitude coding cells
(A) Positive reward magnitude coding cell from OFC. Top row: raster plots for CS1 trials (orange) and CS2 trials (purple) for each of the three different blocks of the experiment. Bottom row: PSTHs with optimal window of coding indicated by red shading between red dashed lines.
(B) Firing rate during the optimal window of coding as a function of trial number for the same amygdala cell. Orange line, CS1 trials; purple line: CS2 trials. Block switches are indicated by vertical dashed lines. Firing rate was smoothed using a 5-trial sliding window.
(C,D) Negative reward magnitude coding cell from the amygdala. Same conventions as in A and B.

can be seen as responding the overall environment, or to the expected value (as defined in statistics) of a trial in a given context.

We show in Figure 7 two example cells from the OFC (positive block coding, 7A and 7B) and amygdala (negative block coding, 7C and 7D). Both cells modify their firing rate during block 2, relative to blocks 1 and 3, regardless of which CS was presented.

Because, unlike the two other types of coding, "block" can be encoded as soon as the trial starts (before CS presentation), we extended the search for block coding cells to the fixation interval. We identified substantial populations of positive and negative block coding cells in each brain area during the fixation, CS and/or trace intervals: 80 cells (48%) in the amygdala (34 positive, 46 negative), and 35 cells (36%) in OFC (12 positive, 23 negative).

Figure 8 summarizes, for each brain area, the numbers of cells that we categorized according to each type of coding as well as their valence (positive or negative). We did not find any evidence as to whether cells tended to be anatomically clustered in the brain according to their type of coding and/or valence, as can be seen in Figure 3.

For some of the analyses presented below, it is useful to group the three identified cell populations into two larger groups: 1) cells sensitive to relative changes ("relative value coding") and 2) cells sensitive to absolute changes only, whether between reward amounts from trial to trial ("reward magnitude coding"), or between average reward amounts from block to block ("block coding").

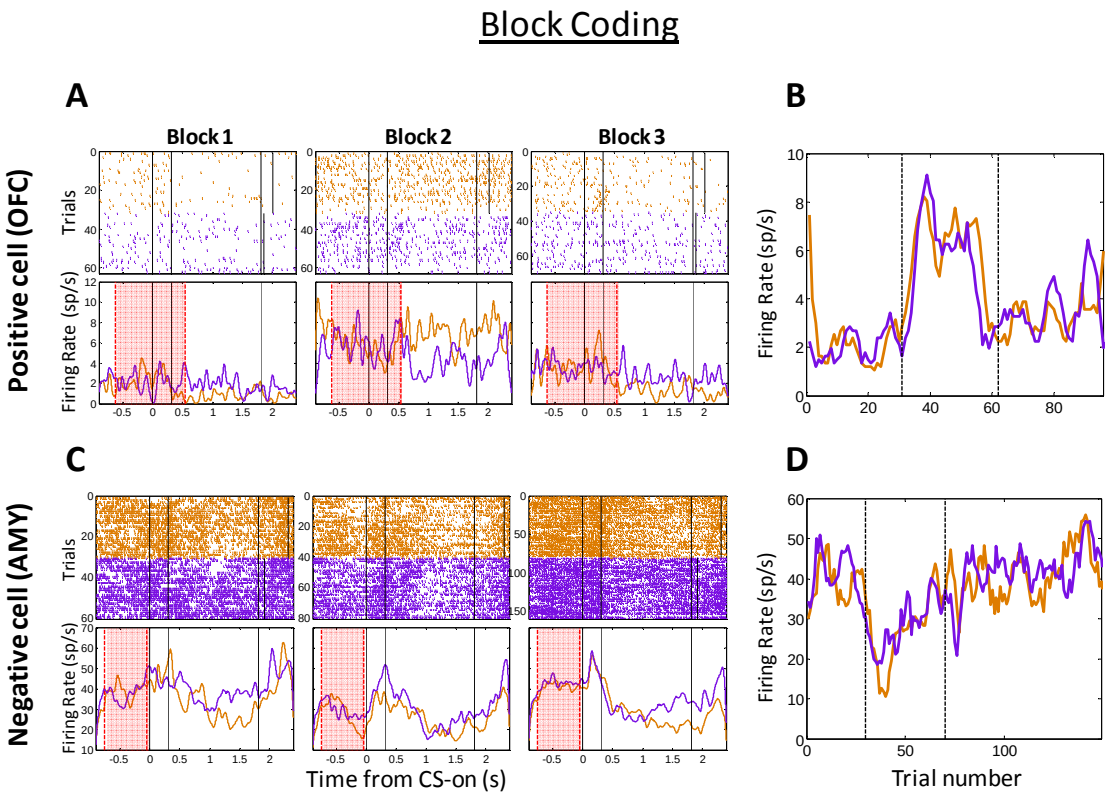

Figure 7: Example block coding cells

(A) Positive block coding cell from OFC. Top row: raster plots for CS1 trials (orange) and CS2 trials (purple) for each of the three different blocks of the experiment. Bottom row: PSTHs with optimal window of coding indicated by red shading between red dashed lines.

(B) Firing rate during the optimal window of coding as a function of trial number for the same amygdala cell. Orange line, CS1 trials; purple line: CS2 trials. Block switches are indicated by vertical dashed lines. Firing rate was smoothed using a 5-trial sliding window.

(C,D) Negative block coding cell from the amygdala. Same conventions as in A and B.

	OFC		Amygdala	
	Positive	**Negative**	**Positive**	**Negative**
Relative Value Coding cells:	13	22	24	25
	35 (36%)		**49 (29%)**	
Reward Magnitude Coding cells:	11	4	27	20
	15 (16%)		**47 (28%)**	
Block Coding cells	12	23	34	46
	35 (36%)		**80(48%)**	

Figure 8: Cell counts for each category
Table summarizing the number and percentage of cells identified in each of the three categories and for each valence

We next asked when during the trial each type of coding was more likely to be represented in the firing rate. Figures 9 and 10 show, for the OFC and the amygdala respectively, the location within the trial of the optimal window of coding for each cell that was classified as relative value-, reward magnitude- or block-coding. Here, we searched for these three types of coding anywhere from trial start to US delivery. These plots show that several cells that encode relative value and reward magnitude seem to reach significance before the onset of the CS, which might seem contradictory. In fact, in most instances, this is merely an artifact of the way the sliding regression analysis is performed (200ms sliding window, stepped every 10ms) as well as the criterion for the optimal window of coding (70% of 30 consecutive bins must be significant). This leads to effects being displayed up to 280ms before their actual occurrence. However, two reward magnitude coding cells, one in OFC and one in the amygdala, truly reach significance before CS-onset. This represents 0.008% of all cells and could be simply due to chance. Most block coding cells tend to do so during the fixation interval, with a clear decrease in the number of cells encoding block at the time of CS presentation (Figs. 9 and 10, bottom panels). One possible explanation is that a number of cells encode the overall "value" of the trial at trial start and then move on to encoding the actual reward amount or relative value of that specific trial once the information about the trial type becomes available. This would result in such a cell being categorized as block coding during fixation an as something else later in the trial.

Cells cannot exhibit more than one type of coding within the same time window because the criteria are mutually exclusive. However, we found a considerable number of cells that encode different properties in different, non-overlapping time periods. Figure

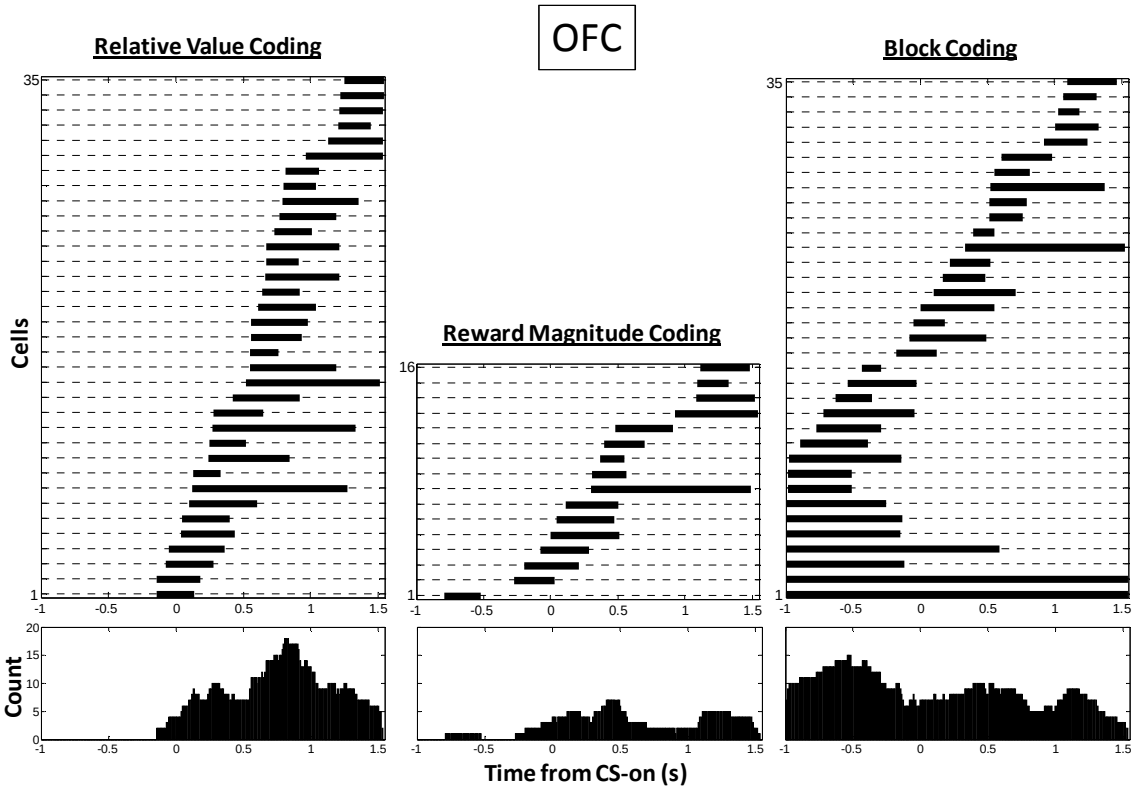

Figure 9: Time of coding of relative value, reward magnitude, and block for each OFC cell
Top: for each relative value, reward magnitude and block coding cell in OFC, we plotted the extent of the optimal window of coding within the trial. All cells for which a window of coding was found anywhere between the fixation interval and the end of the trace interval were plotted. Bottom: histogram with the number of cells encoding each category at any given time during the trial.

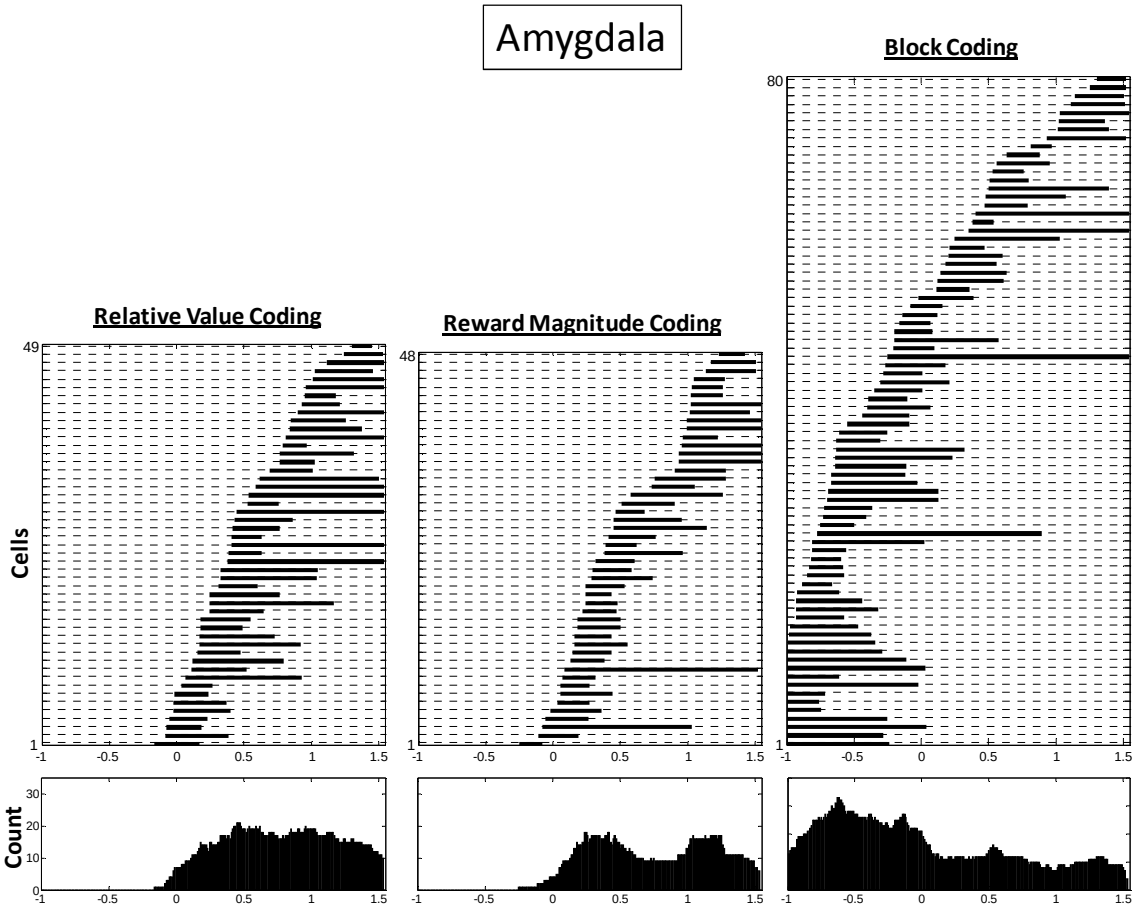

Figure 10: Time of coding of relative value, reward magnitude, and block for each amygdala cell

Top: for each relative value, reward magnitude and block coding cell in the amygdala, we plotted the extent of the optimal window of coding within the trial. All cells for which a window of coding was found anywhere between the fixation interval and the end of the trace interval were plotted.

Bottom: histogram with the number of cells encoding each category at any given time during the trial.

11 shows, for OFC and the amygdala separately, the Venn diagrams for cells that present one, two or three types of coding.

Correlation between firing rate and behavior

As we saw earlier in this work, the licking behavior of monkeys performing this task shows relative valuation of CS1 according to the value associated with the other CS (Figure 1B and 1C). In addition, we identified a population of cells, the relative value coding cells, that reflect the relative value of CS1in their firing rate. The firing pattern of these cells resembles the licking pattern in the monkey's behavior in that the response to CS1 is modulated during block 2 although the associated US remains constant. This resemblance can be seen clearly by comparing Figure 5B to Figure 1B. We therefore wondered whether relative value coding neurons could underlie the animal's behavior during the performance of this task. In order to address this question, we performed a trial-by-trial correlation analysis between the firing rate of the identified relative value coding cells and the licking rate. For each relative value coding cell, we collected the firing rates on CS1 trials in blocks 1 and 3. In doing so, our goal was to collect the firing rate from as many identical trials as possible from each experiment. Although, technically, all CS1 trials in a given experiment are identical, including CS1 trials from block 2 would inevitably lead to a significant correlation between the firing rate and the behavior since, they both present similar patterns across the experiment. Firing rates were taken during the optimal window of relative value coding defined previously (see methods). They were then normalized by z-score and sign-flipped according to the cell's

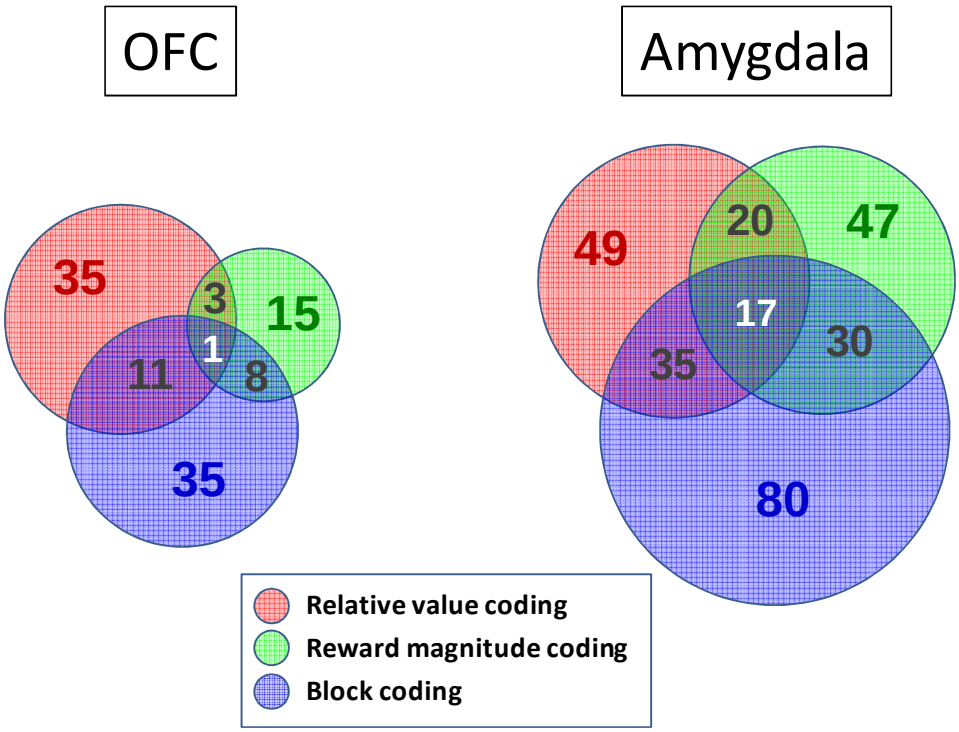

Figure 11: Overlap between cell populations for the three categories of coding
Venn diagrams showing the numbers of cells presenting one, two or three types of coding in their firing rate.

valence so that high firing rates would always correspond to more valuable trials. Licking rates were taken during the whole CS + trace interval. We combined the firing rates of all relative value coding cells into one single vector, where each element is a CS1 trial from blocks 1 or 3, and we correlated it with a vector containing the licking rates from the very same trials. For the populations of relative value coding cells in both the OFC and the amygdala, we obtained positive, significant correlation coefficients: OFC: $r = 0.065$, $p<10^{-3}$; Amygdala: $r = 0.053$, $p<10^{-3}$. This result did not hold true for the other two populations of cells identified in this study: Reward magnitude coding cells, OFC: $r = 0.007$, $p=0.79$; Amygdala: $r = 0.022$, $p=0.12$. Block coding cells, OFC: $r = 0.012$, $p=0.52$; Amygdala: $r = -0.010$, $p=0.45$. This finding suggest that the identified populations of relative value coding cells in the OFC and the amygdala might underlie the flexible representation of value needed to produce an appropriate behavioral response on a trial-by-trial basis.

The trial-by-trial correlation we found between the activity of relative value coding cells and the monkeys' licking behavior, although significant, might seem weak. Reports of significant correlations between trial-by-trial single neuron activity and behavior are relatively few in the literature. Correlations between single neuron activity and saccadic reaction time in LIP (Janssen and Shadlen 2005), and caudate nucleus (Itoh et al. 2003) report effects that are of the same order of magnitude as the effects we see here. Nevertheless, in other to test the robustness of the correlation between firing rate and behavior, we carried out a separate correlation analysis. We computed a correlation coefficient for each relative value coding cell independently and we analyzed the distributions of correlation coefficients. Figure 12 shows these distributions for the OFC

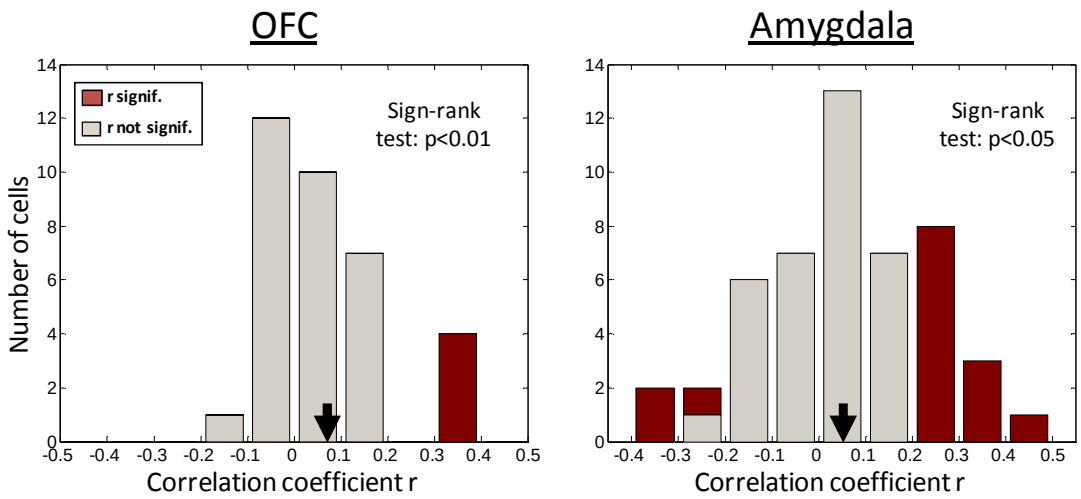

Figure 12: Distributions of correlation coefficients between firing rate and licking behavior for relative value coding cells

For each relative value coding cell, we calculated the trial-by-trial correlation coefficient *r* between its firing rate and the licking rate on CS1 trials during blocks 1 and 3. The distribution of these correlations coefficient is shown for OFC and amygdala cells separately. The mean of each distribution is indicated on the x-axis by a black arrow. Both distributions were significantly shifted towards the positive values of *r* (sign-rank test).

and the amygdala populations separately. The two distributions appear significantly shifted towards the positive values (Sign-rank test: OFC $p<0.01$; Amygdala $p<0.05$), thus showing that the firing rate of each relative value coding cell tends to be, after correcting for the cell's valence, positively correlated with the licking rate. Moreover, when we combined the p-values associated to each positive correlation coefficient using the Fisher method (see methods), we obtained a significant Fisher p-value for both populations: OFC $p_{Fisher}(r>0)<10^{-3}$; Amygdala $p_{Fisher}(r>0)<10^{-6}$. This means that the positive correlation coefficients have a significant trend to be associated with low p-values. This was not the case for cells whose correlation coefficients fell in the negative side, presumably by chance: OFC $p_{Fisher}(r<0)=0.70$; Amygdala $p_{Fisher}(r<0)=0.17$.

One possible explanation for the observed correlation between firing rate and behavior is that relative value coding cells, instead of underlying the representation of value used to drive behavior, are simply involved in the motor planning of the licking action. In order to rule out this possibility, we computed lick-triggered averages of the neural activity. If relative value coding cells were merely signaling motor plans, there should be a pattern of activity time-locked to each lick. Such a pattern was not observed in the present data set.

Relative dynamics of coding between the OFC and the amygdala

A previous study in the Salzman lab investigated the dynamics of value coding between the OFC and the amygdala during learning (Morrison et al. 2011). The study, using a trace conditioning task, found that, following a reversal of the CS-reward

associations, negative value coding neurons in the amygdala update their activity to reflect the new contingencies more quickly than negative value coding neurons in the OFC, while positive value coding neurons in the OFC do so more quickly than positive value coding neurons in the amygdala. This result implies that populations of cells in the two brain areas signal the information on different time scales during learning. The contrast revaluation task used in this work has neither aversive reinforcement nor reversals, and might be more cognitively demanding because it requires the integration of information across trial types in order to compute the relative value of CS1. Due to the higher level of processing and integration that is required for this task, we hypothesized that a prefrontal structure such as OFC might encode the information faster than, and possibly transmit it to, a limbic structure such as the amygdala.

We were interested, in particular, in how fast the cells from each brain area encode the relative value associated with CS1 immediately after the transition between block 1 and block 2 (Figure 13). Because only the reward associated with CS2 is informative as to whether this transition has happened, we used the number of CS2 trials that had been experienced as a metric for time since the transition, rather than the total number of trials since the transition. The normalized firing rates (see methods) of each relative value coding cell were aligned to the transition in this way and averaged together across OFC cells and amygdala cells separately. As previously, firing rates were taken during the optimal window of relative value coding. We assessed the relative dynamics of the two populations of neurons by fitting a Weibull distribution to these averages and comparing the scale parameters of the fitted curves (see methods). We found a significant difference between OFC and amygdala populations ($p<0.01$, permutation test), revealing

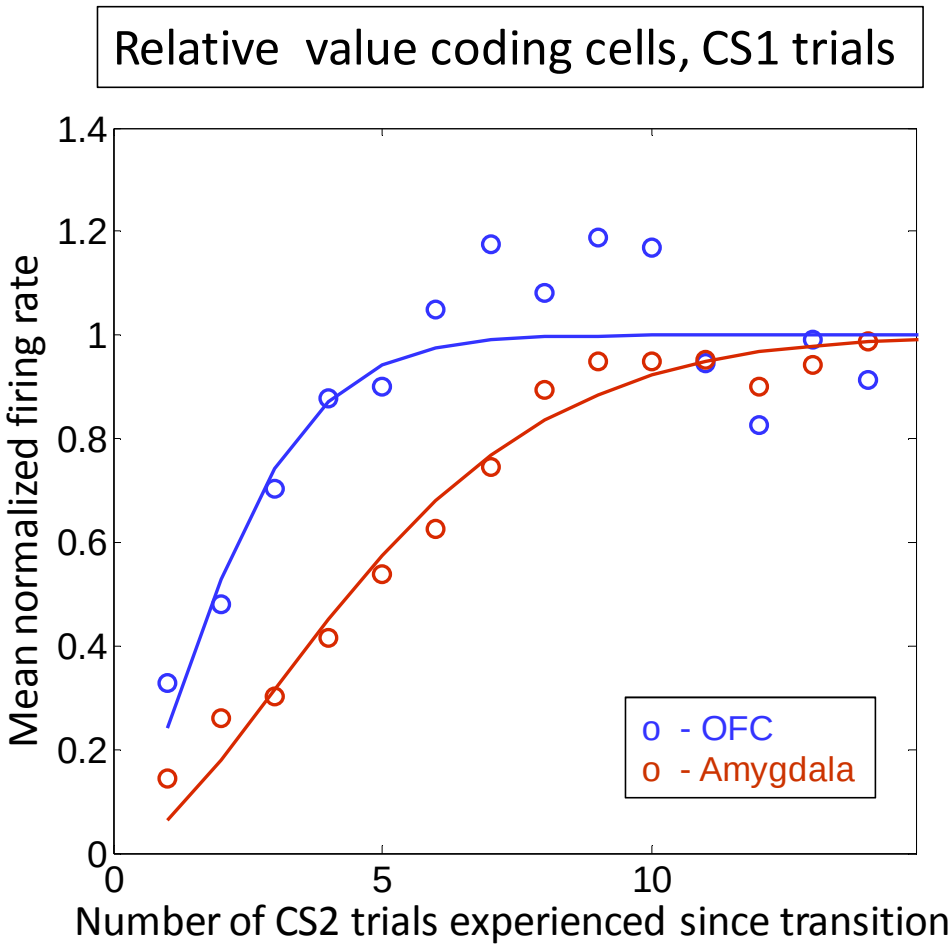

Figure 13: Dynamics of relative value coding in the amygdala and OFC
The firing rates during the optimal window of coding during CS1 trials from all relative value coding cells were normalized, sign-flipped according to the cell's valence to go from low to high, then aligned to the transition from block 1 to block 2. CS1 trials were ranked according to how many CS2 had been experienced previously since the transition. All the cells in each brain area were then averaged together and normalized. The average normalized firing rate was fitted with a Weibull distribution (solid line). OFC cells (blue) encode the relative value of CS1 sooner after the transition than amygdala cells (red)

that relative value coding cells in OFC encode the relative value of CS1 faster than the ones in the amygdala. This result is consistent with the idea that prefrontal signals, in particular from OFC, flexibly integrate information from a changing environment and update the representation of value in the amygdala in a top-down manner (Rolls and Grabenhorst 2008, Miller and Cohen 2001, Kennerley and Wallis 2009b, Sugrue, Corrado and Newsome 2005, Simmons and Richmond 2008).

We then asked whether such a difference in the dynamics of coding between brain areas was specific to relative value signals and the population of cells that encode them. We performed an analogous analysis on the two populations of cells that are only sensitive to absolute changes in reward amounts, namely, reward magnitude coding cells and block coding cells (Figure 14). These cells, by definition, do not modulate their activity on CS1 trials throughout the whole experiment. Reward magnitude coding cells only track the reward associated with CS2, while block coding cells track the mean expected reward of a trial across blocks regardless of the CS. We therefore carried out this analysis only on CS2 trials for the reward magnitude coding population and on all trials for the block coding one. We did not find a significant difference in the dynamics between brain areas in either population of cells: reward magnitude coding cells, $p=0.26$; block coding cells, $p=0.39$, permutation test (Figure 14). One interpretation is that cells that are merely tracking direct changes in reward amount do not need to integrate information across trial types; they are performing a low level computation that might well be executed just as quickly in the amygdala as in the prefrontal cortex.

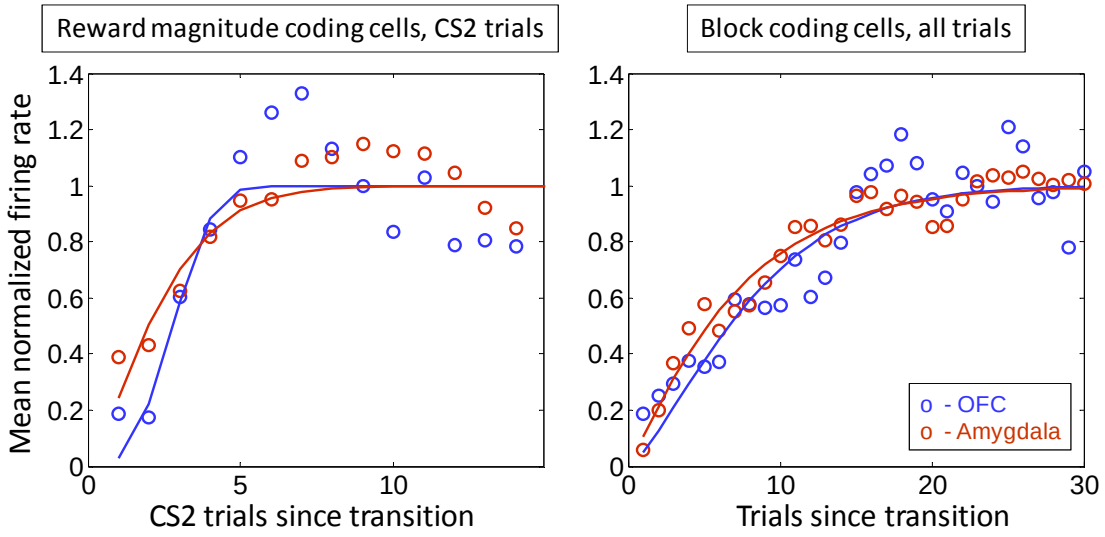

Figure 14: Dynamics of reward magnitude coding and block coding in the amygdala and OFC

Same analysis as in Figure 13 for reward magnitude coding cells (left) and block coding cells (right). Reward magnitude coding was assessed during CS2 trials whereas block coding was assessed on all trials. In both cases, the coding is not updated faster after the transition in one brain area than in another.

State Value

Our group has previously identified, in both the amygdala and OFC, neurons that respond more strongly to a CS that has been associated with reward (classified as positive value coding neurons) as well neurons that respond more strongly to a CS that has been associated with an aversive air-puff (negative value coding neurons). Using these classifications, our laboratory had previously studied the responses to the rewards and air-puffs themselves (stimuli that are primary reinforcers) and a fixation point (a mildly positive stimulus conditioned over a long time period). The neurons coded these other events in a way consistent with their classification. Positive value coding cells increased their response from baseline to the fixation point and reward, and negative value coding cells decreased their response from baseline. These findings have led to the idea that neurons in the amygdala and OFC track the "state value" – loosely defined as the overall value of an organism's situation – which may be modulated as different stimuli appear in the environment (Belova et al. 2007, Belova et al. 2008, Salzman et al. 2007, Paton et al. 2006a, Morrison and Salzman 2009, Morrison et al. 2011, Morrison and Salzman 2011, Bromberg-Martin, Hikosaka and Nakamura 2010).

In the present work, we have shown that cells in the amygdala and in OFC can encode value in different ways during different time intervals throughout the trial. For instance, a cell can encode block, or overall value, during the fixation point interval and reward magnitude during the trace interval. To address whether such cells encode value in a consistent manner throughout the trial, we identified all cells that encode block during the fixation interval (block coding cells) and are, in addition, sensitive to changes in reward amounts during CS or trace intervals (reward magnitude coding cells or block

coding cells). Because, for this analysis, we needed to compute unique regression coefficients for a given time interval (fixation, CS or trace), we fitted our regression model on the firing rate taken from the whole interval rather than using a sliding regression. We found that, overall, the valence of cells that encode block during fixation and cells that encode reward amount during CS or trace, as reported by the sign of β_{BL}, is preserved (Figure 15). This is true in both amygdala and OFC cells. This means that cells that "prefer" block 2 trials (where the overall value is higher) during the fixation interval also tend to prefer trials with a larger reward during the CS or trace intervals. The strength of coding also remains consistent: if the neuron encodes by block strongly during fixation, it is likely to encode value strongly during CS or trace. This implies that neurons are tracking state value across time and sensory stimuli.

In a similar way, relative value can also be coded in different ways throughout the trial. A cell can encode the relative value of a CS based on the size of the reward it predicts (this type of coding has been the main focus of the current work), but also the relative value of a US based on its actual sensory properties (such as water amount). Our task gives us access to this latter type of coding as well: we simply have to apply the same criteria for relative value coding to the firing rate recorded during the US interval, defined as a 500-ms period starting at the end of delivery of the medium reward. We found that neurons which code relative value during the CS-trace interval (CS and trace intervals combined) are significantly more likely to code relative value in the US interval as well. This is true for both OFC and amygdala cells and is shown in the contingency tables in Figure 16A and 16C, respectively, with their associated chi squared tests (OFC $p=0.014$; Amygdala $p=4.6x10^{-3}$). Moreover, cells that encode relative value in both

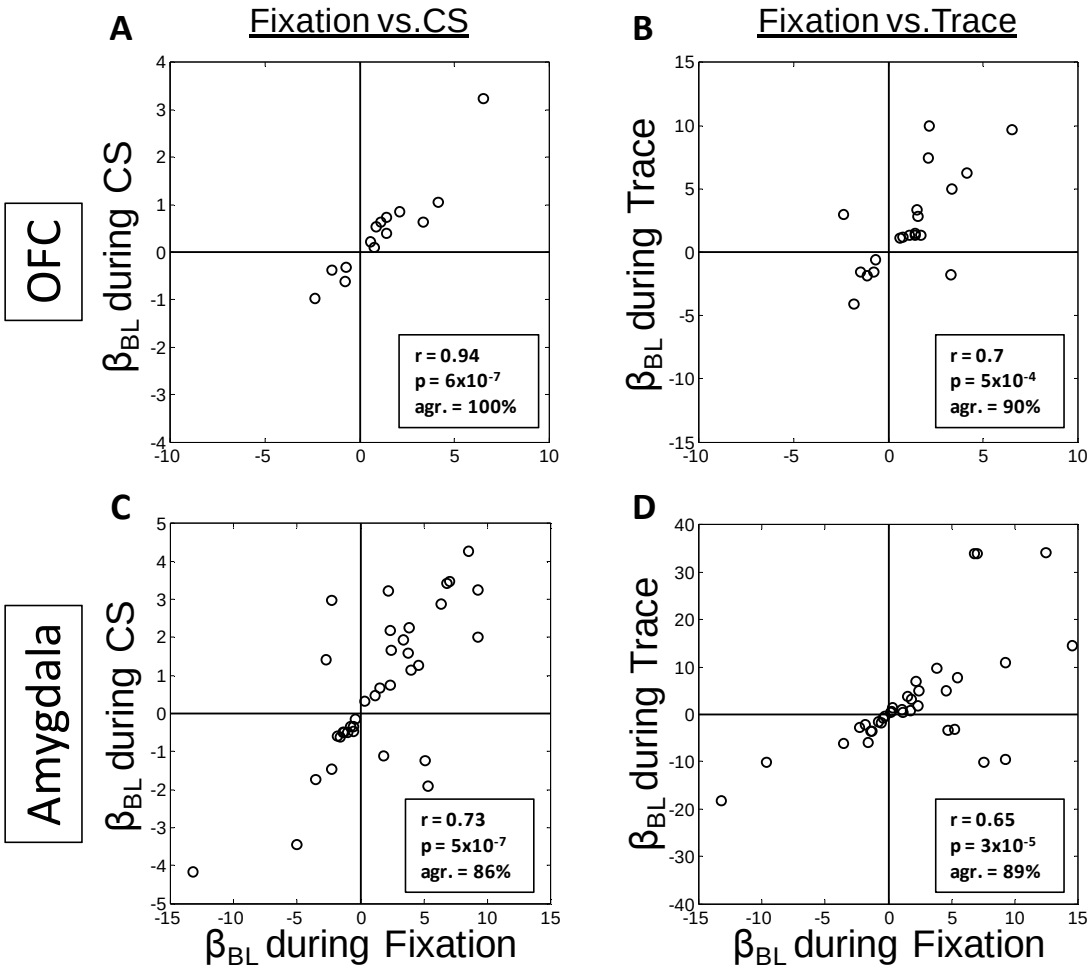

Figure 15: Neurons encode value properties in a consistent manner throughout the trial.
We identified all cells that encode block during the fixation interval and block or reward magnitude during CS or trace intervals. Both these properties can be quantified by the β_{BL} parameter of the regression model. We plotted, for each cell, the β_{BL} found during the CS (A,C) or the trace (B,D) interval against β_{BL} during the fixation interval for OFC cells (A,B) and amygdala cells (C,D) separtaely. We found, over all, that the valence of these cells is preserved, as indicated by the percentage of sign agreement ('agr.') reflecting the proportion of cells that fall in the bottom-left and top-right quadrants. The strength of coding remains consistent as well: if a neuron encodes block strongly during fixation, it is likely to encode block and reward magnitude more strongly during CS or trace. This is shown by the correlation coefficients ('r') found for each scatter plot and their associated p-values ('p').

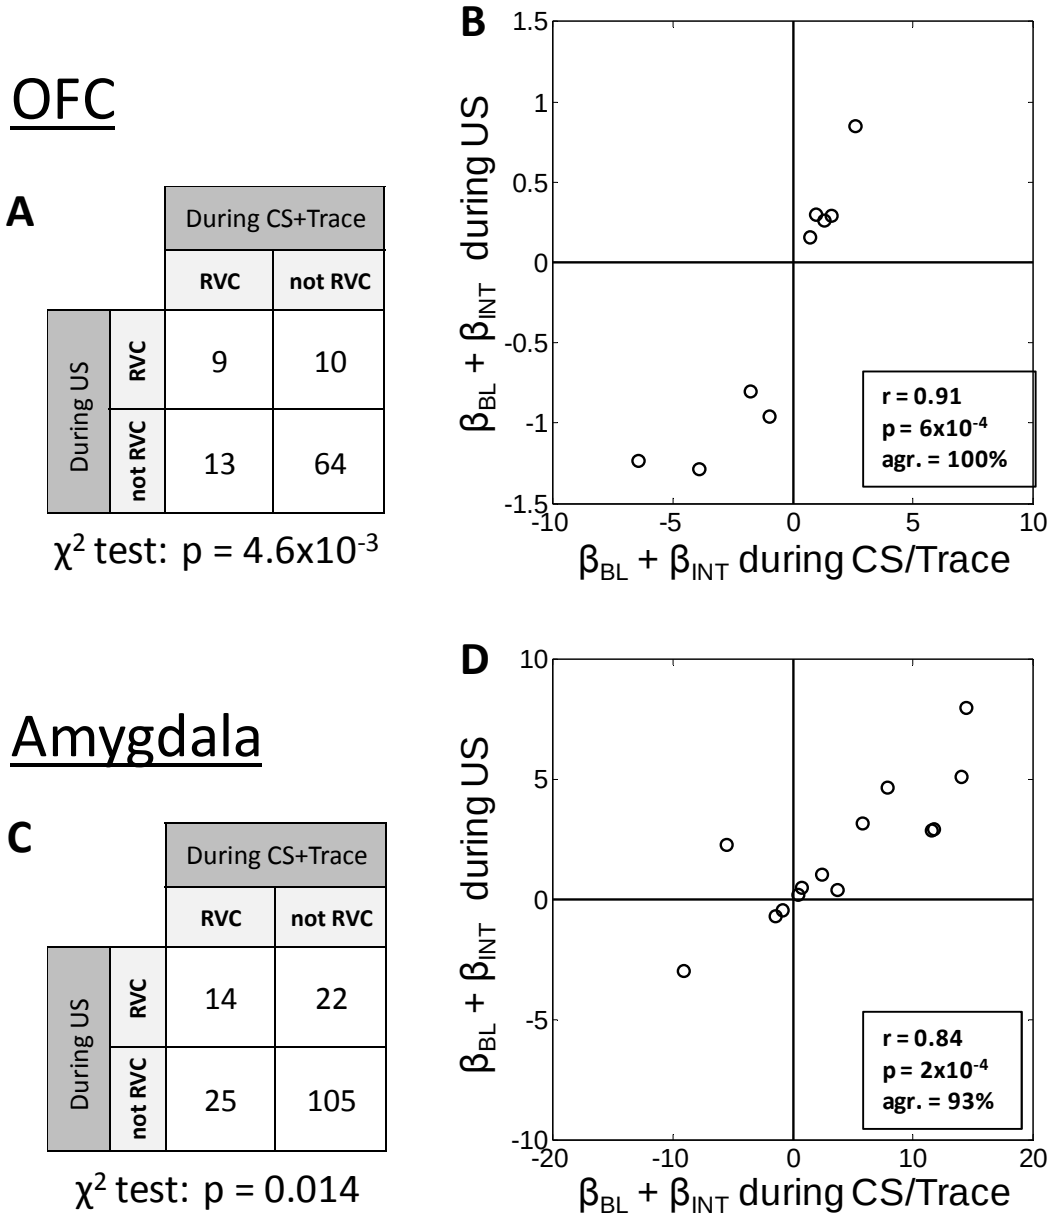

Figure 16: Neurons encode the relative value of the CS and of the US in a consistent manner
We identified all cells that encode relative value during the CS and trace intervals combined and during the US interval and examined the consistency between the two.
(A,C) Contingency tables with the number of cells that encode relative value during CS-trace and during US in OFC (A) and the amygdala (C). When a cell encodes relative value during one interval, it is significantly more likely to do so in the other interval, as indicated by the Pearson's chi-squared test associated with each table.
(B,D) We plotted $\beta_{BL} + \beta_{INT}$ (a measure of relative value coding) during the US interval versus $\beta_{BL} + \beta_{INT}$ during the CS-trace interval for OFC cells (B) and amygdala cells (D). The correlation coefficients ('r') and their associated p-values ('p') indicate that the sign and strength of the relative value coding remain consistent whether the neuron is representing the relative value of a CS with no sensory input during the CS-trace interval or actually experiencing the reward during the US interval.

intervals (14 in the amygdala and 9 in the OFC) tend to maintain a consistent coding, in both their valence and magnitude, as quantified by $\beta_{BL}+\beta_{INT}$ from the regression model (Figures 16B, 16D). These neurons code relative value with the same sign and strength whether they are representing the meaning of a CS learned through experience or actual reward amounts detected by direct sensory input (OFC: r= 0.84, $p<10^{-3}$; Amygdala: r=0.91, $p<10^{-3}$).

Discussion:

We wanted to investigate the type of reward information coded by the amygdala and the OFC, and found a cell population that is sensitive to relative changes ("relative value coding") as well as cell populations sensitive to two kinds of absolute changes: between US amounts ("reward magnitude coding") and between blocks ("block coding"). We were also interested in how this information was updated during learning and whether it differed based on the whether an integration of information across trial types was required. Comparing the dynamics of simultaneous recorded neural signals in the OFC and amygdala, we found that cells in the OFC encoded the relative value of a CS significantly more quickly than cells in the amygdala (Figure 12). There was no significant difference in dynamics between OFC and amygdala when considering reward magnitude coding cells or block coding cells, (as opposed to limiting our sample to relative value coding cells; Figure 13). Relative value requires the integration of information across trial types to adjust value. This might be a more complicated type of value coding than reward magnitude or block coding, which does not need information

from other trial types to determine value. This higher order processing might originate in the OFC before being transmitted to the amygdala.

What type of information is coded in the OFC and amygdala during the contrast revaluation task?

Several types of associations can be formed between a CS and a US (Everitt et al. 2003, Delamater 2007, Delamater and Oakeshott 2007). First, the CS may become directly associated with an unconditioned response, such as licking at a reward spout. Second, the CS may become associated with the subjective value of the expected US: *e.g.*, positive in the case of a reward, negative in the case of an aversive event. Third, the CS can become associated with the specific sensory properties of the US, including sound, feel, smell or taste. We examined the available data to determine which of these types of associations was present in the current work.

To see if the neural signals were related to stimulus–response learning (in this case, the CS becoming associated with the unconditioned response of licking), we looked for correlations between neuronal signals and licking behavior using a lick-triggered peristimulus histogram; we did not find any signals that could be attributed to driving the motor behavior. We conclude from this analysis that the motor plans for licking are derived from another "downstream" brain area that presumably could use these signals from OFC and amygdala to correctly guide motor output. This is consistent with the idea that OFC neurons do not simply represent CS-motor associations (Tremblay and Schultz 2000, Wallis and Miller 2003b, Padoa-Schioppa and Assad 2006).

Once CS-motor responses were ruled out, we wanted to investigate whether neural signals were associated with subjective CS value or the sensory properties of the associated US (water). We therefore compared responses to the same CS in block 1 and block 2, in which the subjective value of the CS changed but the sensory properties of the reward remained constant. We found that many amygdala and OFC neurons encoded the relative value of a CS even when the US associated with it had not changed; this implies that these neurons are coding an abstract, subjective value and not the sensory properties of the US or the absolute amount of the specific reward (Figure 4A). This is consistent with previous studies that found that neurons in the OFC and amygdala did not merely represent the sensory properties of a reward associated with a US (Belova et al. 2008, Morrison and Salzman 2009). Therefore, of the three types of associations outlined above, we found that the second – the association between a CS and the subjective value of a US – was most representative for the population of cells we termed "relative value coding."

How is information updated during learning? Does it differ based on whether an integration of information across trial types is required?

Another study performed previously in the Salzman laboratory investigated the dynamics of neural signals in the OFC and amygdala during learning (Morrison et al. 2011). The study found that negative value coding neurons in the amygdala update more quickly than negative value coding neurons in the OFC, while positive value coding neurons in the OFC update more quickly than positive value coding neurons in the

amygdala. After learning, on the other hand, the OFC generally signaled the value of a CS earlier than the amygdala.

In contrast to these earlier findings, in the present study we found that cells in the OFC encoded the relative value of a CS significantly more quickly than cells in the amygdala (Figure 12). There was no significant difference in dynamics between OFC and amygdala when considering reward magnitude coding cells and block coding cells (Figure 13). Relative value requires the integration of information across trial types to adjust value. This might be a more complicated type of value coding than reward magnitude or block coding, which does not need information from other trial types to determine value. This higher order processing might originate in the OFC before being transmitted to the amygdala.

There are some differences in the paradigms used in the current study (discussed in Chapter II) and this other study by our group which might contribute to the different findings. First, the contrast revaluation task had no aversive reinforcement, which may explain why we did not see the amygdala leading: we are not engaging the aversive network, which might update more quickly in the amygdala than in OFC. Second, the contrast revaluation task does not have reversals, so the updating of relative value must integrate information across trial types. This computation might be more complicated than a simple CS-to-value association, which might take place faster in the amygdala.The cells that are only sensitive to changes in US amount – “reward magnitude coding” cells – do not need to integrate information across trial types, as they are simply responding to a change in the US amount associated with a stimulus. Cells sensitive to relative value changes, on the other hand, are integrating information across trial types to track changes

in relative value. If they were only responding to US amounts they would not change their responses to a CS that changes its subjective value, but not the actual magnitude of the US associated with it. By comparing these changes in response across brain areas, we conclude that the relative value cells in the OFC are updating information more quickly than those in the amygdala.

From this result, we can infer that the prefrontal cortex leads the neural computation of relative value compared to the amygdala. This task might engage the OFC's role in executive function and emotional regulation, in which case the OFC might update the relative value of a CS before passing this information to the amygdala. This is consistent with the finding that patients with lesions of the OFC are impaired on tasks that require flexible learning which also require integration of feedback from other trials (Tsuchida, Doll and Fellows 2010).

What other brain areas are involved in providing and using this information?

The OFC and the amygdala do not work in isolation to signal value and guide behavior; rather, they work within an entire network of brain areas that are interconnected. For example, the OFC and amygdala are connected to the anterior cingulate cortex (ACC), which has been shown to encode similar information about value, reward amount, reward proximity, motivation (Shidara and Richmond 2002b, Toda et al. 2012). Meanwhile, internal factors such as satiety could be sent from the ventromedial prefrontal cortex to the OFC and the amygdala (Bouret and Richmond 2010). The OFC is reciprocally connected with the amygdala, hippocampus, striatum,

hypothalamus, and other parts of the prefrontal cortex (Cavada et al. 2000, Haber et al. 1995, Carmichael and Price 1995b) The OFC receives information relevant to motivation and emotion, and has outputs to the hypothalamus and other subcortical structures to participate in regulating autonomic responses related to emotion and motivation (Ongur, An and Price 1998, Rempel-Clower and Barbas 1998). The amygdala has connections to the PFC and higher sensory cortices, the hippocampus, basal ganglia, perirhinal and entorhinal cortices, the basal forebrain, and subcortical structures such as the hypothalamus (Davis 2000). The information about relative value could be passed from the OFC and the amygdala to the ventral striatum to guide actions (Simmons et al. 2007).

Could relative value be explained by range adaptation?

One theory that has been proposed to explain relative value is "range adaption", where the neurons use the current range of rewards to determine a specific range of firing rate. Two studies have demonstrated the existence of such mechanisms (Padoa-Schioppa 2009, Kobayashi et al. 2010). In the study by Schultz and colleagues, animals performed a task in which a fixation cue predicted the standard deviation of the probability distribution of juice volumes, while the expected mean volume was kept constant. A subsequent cue specified the exact juice volume obtained for a correct saccade response. Population responses of orbitofrontal neurons that reflected the predicted juice volume showed adaptation of their firing rate to the reward distribution. In the study by Padoa-Schioppa, monkeys chose between different juices, and their choice patterns provided a measure of subjective value. Value ranges were varied from session to session and, in

each session, OFC neurons encoded values in a linear way. The neuronal activity range did not depend on the value range. Thus, Padoa-Schioppa argued, the activity of each neuron adapts to the range values it encodes, but does not depend on other, previously available goods (Padoa-Schioppa and Assad 2008b). In both studies, the authors found that the encoding of value undergoes range adaptation such that a given range of firing rates represents different ranges of values in different behavioral contexts.

Neurons can be selective to multiple aspects of this task: value, CS identity, block, etc. Testing the range adaptation hypothesis requires being able to manipulate value independently from other parameters. The task used in this work was not designed for testing this hypothesis and is not well suited for it. Indeed, the range of firing rates in each of our blocks is determined by only two conditions, or trial types. Each of these conditions corresponds to one level of value (given by the associated reward amount) but also to one CS, and cells in the amygdala and in OFC can be selective to one or the other as well as both. In addition, cells can be selective to the block itself and the contribution of the block factor to the firing rate might not add linearly with the other factors. This makes the investigation of how the range of firing rates varies exclusively with respect to the range of values challenging.

What are possible uses of the neural signals in OFC and amygdala?

During the contrast revaluation task, the monkeys could have used a neural representation of relative value coding to guide their behavior (licking patterns). The firing rate of relative value coding cells in the OFC and amygdala cells was significantly

correlated with the licking behavior ($p_{Fisher} < 0.05$). Consistent with this finding, we noted that the monkeys' licking corresponded to the relative value and not the absolute reward magnitude of a CS. The fact that behavior tracks the relative value of the stimulus implies that there must be signals in the brain also tracking the relative value of that stimulus to be part of the circuit for driving behavior. As mentioned above, there might be multiple brain areas – in addition to amygdala and OFC – involved in this process. For example, one brain area could store the current value of a stimulus in a specific environmental condition and another could be involved in the motor output pathways that are driving the licking behavior.

In addition to cells that encode the relative value of a CS, we found cells that only encode the absolute magnitude of the reward associated with the CS (Figure 4B). These cells might be important for multiple functions: 1) keeping track of recent rewards on an absolute scale, 2) computing the total value of reward available, and 3) setting the range of rewards available on an absolute scale. These signals are sometimes found in the same cells that code relative value, or can sometimes be found in separate cells.

It is important to note that relative value cannot be calculated without a reference. This means that reward history, an integration of rewards in all trial types, must be computed and compared to the current reward. We found neurons in the amygdala and OFC that encode the total, or overall, available reward in a given block of trials (Figure 4C). This neural representation, which we referred to as "block coding", can be used to compute relative value by comparing the overall reward to the current reward – the output of that computation being relative value coding.

The central finding of the current study – the existence of cells that encode relative value in amygdala and OFC – could help provide a neurobiological basis for the results from many previous experiments in rodents, monkeys, and humans. For example, in economics, they have shown that humans' decisions are altered by contextual factors such as framing effects (*e.g.*, the presentation of a problem in terms of losses or gains) and reference points – an internal zero point where everything above is considered a gain and everything below a loss (Tversky and Kahneman 1981). This decision process could be guided by cells that are sensitive to relative value and that integrate information about the context to determine the value of a stimulus.

Turning to rodents, relative value coding cells could be useful during a contrast task in which a rewarding stimulus decreases in subjective value when a better reward enters the environment (Flaherty et al. 1994). In rats, intake of a sweet solution is suppressed when the rat knows that an even sweeter solution will soon become available (Flaherty et al. 1994). This task requires that cells process information about other rewards in the environment and adjust the valuation of the current reward in comparison. The relative value cells identified in the current study could guide this process.

These cells could also be involved in devaluation assessments. In the Pavlovian version of the classic reinforcer devaluation task, via conditioning, a cue comes to predict food. Then, the normally rewarding food is devalued (in the absence of the cue) by pairing the food with illness or satiation. After devaluation, normal animals spontaneously reduce responding in the presence of the cue that predicts availability of the "devalued" food. The amygdala and OFC are both critical for reinforcer devaluation (Hatfield et al. 1996, Pickens et al. 2003, Johnson et al. 2009, Malkova et al. 1997,

Izquierdo and Murray 2004) (Murray and Izquierdo 2007, Izquierdo and Murray 2007) (Izquierdo et al. 2004), (Pickens et al. 2003, Pickens et al. 2005). These studies strongly suggest that when a CS-US association has a specific value and then the value of the US is changed, the amygdala and OFC are involved in updating the association so that the CS, as well as the US, is devalued. The same mechanisms could be involved as in the contrast revaluation task, in which a CS-US association is formed for a specific reward magnitude, and then a bigger reward enters the environment, requiring an integration of information about the new reward and the current one; in each case, the requisite updating of CS value may take place in the OFC and the amygdala using the relative value cells.

The amygdala and OFC encoding relative value during this paradigm agrees with theories of efficient statistics: it is biologically efficient to re-scale because only a limited capacity of information can be inferred from changes in firing rate. In contrast, the number of possible reward values has no absolute limits. If a neuron's limited output were allocated evenly to represent the large – perhaps infinite – number of possible reward values, then that neuron's activity would allow for little, if any, discrimination between rewards. Moreover, these results suggest that, in different contexts, the system rescales its calculation of reinforcement value based on its relationship to the environment. This neural rescaling could underlie the psychological phenomenon of "normalized happiness," which can be described as the following: we predict that our lives would improve if given X (a new job, more money, power, etc.); when we actually obtain X, our overall happiness remains unchanged (after a possible brief burst of increased happiness), because now our whole value spectrum has rescaled to include X.

We are constantly shifting our value ranges depending on life experiences. For example, if someone takes a trip to an island and their worst recent experience is sitting in traffic, they would assign sitting in traffic to the bottom of the value scale and a trip to an island as the top of value scale. If they are then diagnosed with cancer, they would assign being diagnosed with cancer to the bottom of the value scale, with a trip to an island still at the top of the value scale. Now sitting in traffic would be viewed as much less negative than in the previous situation. Thus, we stretch out along a value scale all events in recent history and we are constantly coming to a new baseline. This explains why we are able to rise to the occasion during tough situations while at other times losing our car keys can make us extremely upset. This could also underlie why lottery winners are not happier than controls (Brickman, Coates and Janoff-Bulman 1978); or why the blind, the cognitively disabled, and the physically disabled are not less happy than other people (Cameron et al. 1973). There are also findings that sex, race, age, income, education, family life-cycle stage, and other demographic classification variables accounted for relatively little variance in general happiness in two independent national surveys (Andrews and Withey 1976).

A framework for understanding these phenomena can be provided by “adaptation level theory” (Helson and Bevan 1964). This theory proposes that the effect of the current level of stimulation depends on whether it is greater or less than the level of stimulation to which a subject’s previous history has accustomed them. If something very good happens it will lessen the pleasure of mundane events and then habituation will decrease the pleasure of the very good thing that happened. The reverse is also true for extremely

bad events: mundane events now increase the pleasure associated with them and habituation erodes the negative value associated with the bad event.

The current work helps expand our knowledge of how the brain represents relative changes and absolute changes of rewards available in the environment: either between US amounts or between blocks. It also elucidates the dynamics of neural signals in the OFC and amygdala, providing evidence that cells in the OFC encode the relative value of a CS significantly more quickly than cells in the amygdala (Figure 12). This task might utilize the OFC's role in executive function and emotional regulation because of the need to integrate information across trial types to determine relative value. There was no significant difference in dynamics between OFC and amygdala when considering reward magnitude coding cell and block coding cells Figure 13). These non-relative value coding cells are only sensitive to changes in US amount and do not need to integrate information across trial types; thus, this type of neural signal does not engage the OFC earlier than the amygdala. Overall, these neural signals could underlie decision processes in adapting contexts where value rescaling is necessary.

Chapter III. Future Directions

The work described in Chapter II leaves open many related questions that could provide avenues for future research. It also opens up a number of new questions that could serve as a basis for follow-up experiments.

The two main areas I will address below are 1) the limitations of single unit electrophysiology in awake behaving monkeys and 2) our lack of understanding of the underlying neural mechanism of relative value encoding. The first set of questions can start to be addressed with some technological advances and the second through experimental rather than technological means.

Limitations of single unit electrophysiology in awake behaving monkeys

There are many powerful elements to the experimental set-up of awake-behaving electrophysiology. It can be very informative to have an animal performing a task, and to be able to record moment-to-moment changes in firing rate from single neuron, but there are many limitations with this technique as well that must be addressed especially in areas like the OFC and amygdala.

In electrophysiology, we can only record the firing rate of a neuron that we have an electrode next to. This comes with a number of limitations addressed below. First, we do not know what type of cell we are recording from, since no columnar or laminar organization according to cell type has been found in either area, and all classic types of

analysis to identify cell type, such as inter-spike interval and waveform analysis, have not brought about fruitful answers to these questions. Second, we only record a few neurons on a given day and verify the responses to a specific task. We try to adjust for day-to-day noise and fluctuations by recording over many sessions and averaging them together. However, a more advantageous solution would be to record from a very large number of neurons at a time to infer what a given brain area is coding overall. Could there be a read-out system where the number of cells coding positive value or negative value determines the overall value? Or could it be about which population produces a stronger signal? Third, we don't know how downstream areas use the signals that we recorded. Are they able to separate multiple signals from one firing rate much like we do with a regression analysis? Fourth, we don't know the circuitry of how cells are wired together in the OFC and the amygdala.

One way to address the first and fourth of these limitations is through utilizing optogenetics to activate neuronal firing using light-activated ion channels (Boyden et al. 2005, Arenkiel et al. 2007, Miesenbock 2009). In the future we will have the ability to genetically target through cell type specific promoters the expression of light-activated channels to different classes of cells. We can record from neurons during a task and then at the end light activate and be able to classify the cell types depending on whether it is activated. The manipulation will be done at the end of the experiment, so that firing rate is not affected during the experiment. This will help address which cell we are recording from, and can help us understand the underlying circuits.

The second limitation can be addressed by to simultaneously recording from many more neurons, across multiple brain areas, in order to gain an understanding of the

neural circuitry underlying a behavior, and not just the activity of single neurons in isolation.

Underlying neural mechanism of relative value encoding

Although we found that neurons in OFC and amygdala that track the relative value of a stimulus, we do not know how the brain computes this value. There are many theories for how this relative valuation occurs, which I will discuss below, but we have no conclusive evidence for confirming any of the theories.

A theory proposed to explain this phenomenon is range adaption, where neurons use the current range of rewards to determine a specific range of firing rate. In order to use this model to explain these results, the range would have to be computed on a block-by-block basis. In the contrast revaluation task, the range of rewards in block 1 goes from small to medium. In this case, for example, the medium reward could be coded with a 30 Hz firing rate and the small reward with a 10 Hz firing rate. In block 2, the range of rewards goes from medium to large; in this example, the large reward could be coded with a 30 Hz firing rate and the medium reward with a 10 Hz firing rate. This would appear as though the neuron was devaluing the medium reward even though, instead, it is simply adapting its firing range to the current reward range.

Testing the range adaptation hypothesis requires being able to manipulate value independently from other parameters. The task used in this work was not designed for testing this hypothesis and is not well suited for it. Indeed, the range of firing rates in each of our blocks is determined by only two conditions, or trial types. Each of these

conditions corresponds to one level of value (given by the associated reward amount) but also to one CS, and cells in the amygdala and in OFC can be selective to one or the other as well as both. In addition, cells can be selective to the block itself and the contribution of the block factor to the firing rate might not add linearly with the other factors. This makes the investigation of how the range of firing rates varies exclusively with respect to the range of values challenging.

Another theory to explain relative value is that neurons are taking either a ratio or a difference between the two rewards and computes a relative value from that measure. A fractional rule would predict that neural responses to reward will remain unaffected as long as the *ratio* between the reward values remains unchanged; in contrast, a difference rule would predict that responses to the rewards will change if the *difference* between the reward values changes. The ratio of the rewards in all three of our blocks is equal but the difference is not. If neurons were coding the ratio between the rewards the range of firing rates should remain constant throughout the blocks as was proposed by the range adaptation theory but this was not found to be the case. One way the firing rate range expansion of block 2 could be explained is that neurons could be tracking the difference between rewards which increases in block 2.

It may be that the full range of values can be coded on an absolute scale despite the limited number of neurons and firing rate constraints. This idea could be tested by giving extremely large and small rewards, or by having many differently sized rewards to test the limits of the system to see if the absolute value scale can be maintained. On the other hand, it remains true that the number of possible reward values has no absolute limits. If a neuron's limited output were allocated evenly to represent the large – perhaps

infinite – number of possible reward values, then that neuron's activity would allow for little if any discrimination between rewards.

References

Andrews, F. M. & S. B. Withey. 1976. Social indicators of well-being. New York: Plenum Press.

Arenkiel, B. R., J. Peca, I. G. Davison, C. Feliciano, K. Deisseroth, G. J. Augustine, M. D. Ehlers & G. Feng (2007) In vivo light-induced activation of neural circuitry in transgenic mice expressing channelrhodopsin-2. *Neuron,* 54**,** 205-18.

Barbas, H. (2007) Specialized elements of orbitofrontal cortex in primates. *Ann N Y Acad Sci,* 1121**,** 10-32.

Barbas, H. & G. J. Blatt (1995) Topographically specific hippocampal projections target functionally distinct prefrontal areas in the rhesus monkey. *Hippocampus,* 5**,** 511-33.

Baxter, M. G., A. Parker, C. C. Lindner, A. D. Izquierdo & E. A. Murray (2000) Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. *J Neurosci,* 20**,** 4311-9.

Belova, M. A., J. J. Paton, S. E. Morrison & C. D. Salzman (2007) Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron,* 55**,** 970-84.

Belova, M. A., J. J. Paton & C. D. Salzman (2008) Moment-to-moment tracking of state value in the amygdala. *J Neurosci,* 28**,** 10023-30.

Bouret, S. & B. J. Richmond (2010) Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J Neurosci,* 30**,** 8591-601.

Boyden, E. S., F. Zhang, E. Bamberg, G. Nagel & K. Deisseroth (2005) Millisecond-timescale, genetically targeted optical control of neural activity. *Nat Neurosci,* 8**,** 1263-8.

Brickman, P., D. Coates & R. Janoff-Bulman (1978) Lottery winners and accident victims: is happiness relative? *J Pers Soc Psychol,* 36**,** 917-27.

Bromberg-Martin, E. S., O. Hikosaka & K. Nakamura (2010) Coding of task reward value in the dorsal raphe nucleus. *J Neurosci,* 30**,** 6262-72.

Cameron, P., D. G. Titus, J. Kostin & M. Kostin (1973) The life satisfaction of nonnormal persons. *J Consult Clin Psychol,* 41**,** 207-14.

Camille, N., A. Tsuchida & L. K. Fellows (2011) Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *J Neurosci,* 31**,** 15048-52.

Cardinal, R. N., J. A. Parkinson, J. Hall & B. J. Everitt (2002) Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci Biobehav Rev,* 26**,** 321-52.

Carmichael, S. T. & J. L. Price (1994) Architectonic subdivision of the orbital and medial prefrontal cortex in the macaque monkey. *J Comp Neurol,* 346**,** 366-402.

--- (1995a) Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J Comp Neurol,* 363**,** 615-641.

--- (1995b) Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *Journal of Comparative Neurology,* 363**,** 615-641.

--- (1995c) Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol,* 363**,** 642-664.

Cavada, C., T. Company, J. Tejedor, R. J. Cruz-Rizzolo & F. Reinoso-Suarez (2000) The anatomical connections of the macaque monkey orbitofrontal cortex. A review. *Cerebral Cortex,* 10**,** 220-42.

Corrado, G. S., L. P. Sugrue, H. S. Seung & W. T. Newsome (2005) Linear-Nonlinear-Poisson models of primate choice dynamics. *J Exp Anal Behav,* 84**,** 581-617.

Cromwell, H. C. & W. Schultz (2003) Effects of expectations for different reward magnitudes on neuronal activity in primate striatum. *J Neurophysiol,* 89**,** 2823-38.

Davidson, R. J. (2002) Anxiety and affective style: role of prefrontal cortex and amygdala. *Biol Psychiatry,* 51**,** 68-80.

Davis, M. 2000. The role of the amygdala in conditioned and unconditioned fear and anxiety. In *The Amygdala: A Functional Analysis,* ed. J. Aggleton, 213-287. New York: Oxford University Press.

de Araujo, I. E., E. T. Rolls, M. I. Velazco, C. Margot & I. Cayeux (2005) Cognitive modulation of olfactory processing. *Neuron,* 46**,** 671-9.

Deaner, R. O., A. V. Khera & M. L. Platt (2005) Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Curr Biol,* 15**,** 543-8.

Delamater, A. R. (2007) The role of the orbitofrontal cortex in sensory-specific encoding of associations in pavlovian and instrumental conditioning. *Ann N Y Acad Sci,* 1121**,** 152-73.

Delamater, A. R. & S. Oakeshott (2007) Learning about multiple attributes of reward in Pavlovian conditioning. *Ann N Y Acad Sci,* 1104**,** 1-20.

Dorris, M. C. & P. W. Glimcher (2004) Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron,* 44**,** 365-78.

Elliott, R., Z. Agnew & J. F. Deakin (2008) Medial orbitofrontal cortex codes relative rather than absolute value of financial rewards in humans. *Eur J Neurosci,* 27**,** 2213-8.

Everitt, B. J., R. N. Cardinal, J. A. Parkinson & T. W. Robbins (2003) Appetitive behavior: impact of amygdala-dependent mechanisms of emotional learning. *Ann N Y Acad Sci,* 985**,** 233-50.

Flaherty, C. F., J. Turovsky & K. L. Krauss (1994) Relative hedonic value modulates anticipatory contrast. *Physiol Behav,* 55**,** 1047-54.

Gallagher, M., R. W. McMahan & G. Schoenbaum (1999) Orbitofrontal cortex and representation of incentive value in associative learning. *J Neurosci,* 19**,** 6610-4.

Ghashghaei, H., Barbas, H (2002) Pathways for emotion: interactions of prefrontal and anterior temporal pathways in the amygdala of the rhesus monkey. *Neuroscience,* 115**,** 1261-79.

Ghashghaei, H. T. & H. Barbas (2002) Pathways for emotion: interactions of prefrontal and anterior temporal pathways in the amygdala of the rhesus monkey. *Neuroscience,* 115**,** 1261-79.

Ghashghaei, H. T., C. C. Hilgetag & H. Barbas (2007) Sequence of information processing for emotions based on the anatomic dialogue between prefrontal cortex and amygdala. *Neuroimage,* 34**,** 905-23.

Gilbert, P. E. & R. P. Kesner (2002) The amygdala but not the hippocampus is involved in pattern separation based on reward value. *Neurobiol Learn Mem,* 77**,** 338-53.

Gottfried, J. A., J. O'Doherty & R. J. Dolan (2003) Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science,* 301**,** 1104-7.

Haber, S. N., K. Kunishio, M. Mizobuchi & E. Lynd-Balta (1995) The orbital and medial prefrontal circuit through the primate basal ganglia. *Journal of Neuroscience,* 15**,** 4851-67.

Hassani, O. K., H. C. Cromwell & W. Schultz (2001) Influence of expectation of different rewards on behavior-related neuronal activity in the striatum. *J Neurophysiol,* 85**,** 2477-89.

Hatfield, T., J. S. Han, M. Conley, M. Gallagher & P. Holland (1996) Neurotoxic lesions of basolateral, but not central, amygdala interfere with Pavlovian second-order conditioning and reinforcer devaluation effects. *J Neurosci,* 16**,** 5256-65.

Heeger, D. J. (1992) Normalization of cell responses in cat striate cortex. *Vis Neurosci,* 9**,** 181-97.

Helson, H. & W. Bevan. 1964. An investigation of variables in judgement of relative area. *J Exp Psychol,*.

Hikosaka, K. & M. Watanabe (2000) Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. *Cereb Cortex,* 10**,** 263-71.

Hosokawa, T., K. Kato, M. Inoue & A. Mikami (2007) Neurons in the macaque orbitofrontal cortex code relative preference of both rewarding and aversive outcomes. *Neurosci Res,* 57**,** 434-45.

Ikeda, T. & O. Hikosaka (2003) Reward-dependent gain and bias of visual responses in primate superior colliculus. *Neuron,* 39**,** 693-700.

Itoh, H., H. Nakahara, O. Hikosaka, R. Kawagoe, Y. Takikawa & K. Aihara (2003) Correlation of primate caudate neural activity and saccade parameters in reward-oriented behavior. *J Neurophysiol,* 89**,** 1774-83.

Izquierdo, A. & E. A. Murray (2004) Combined unilateral lesions of the amygdala and orbital prefrontal cortex impair affective processing in rhesus monkeys. *J Neurophysiol,* 91**,** 2023-39.

--- (2007) Selective bilateral amygdala lesions in rhesus monkeys fail to disrupt object reversal learning. *J Neurosci,* 27**,** 1054-62.

Izquierdo, A., R. K. Suda & E. A. Murray (2004) Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *J Neurosci,* 24**,** 7540-8.

Janssen, P. & M. N. Shadlen (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci,* 8**,** 234-41.

Johnson, A. W., M. Gallagher & P. C. Holland (2009) The basolateral amygdala is critical to the expression of pavlovian and instrumental outcome-specific reinforcer devaluation effects. *J Neurosci,* 29**,** 696-704.

Kennerley, S. W., T. E. Behrens & J. D. Wallis (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci,* 14**,** 1581-9.

Kennerley, S. W. & J. D. Wallis (2009a) Encoding of reward and space during a working memory task in the orbitofrontal cortex and anterior cingulate sulcus. *J Neurophysiol,* 102**,** 3352-64.

--- (2009b) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci,* 29**,** 2061-73.

Kobayashi, S., O. Pinto de Carvalho & W. Schultz (2010) Adaptation of reward sensitivity in orbitofrontal neurons. *J Neurosci,* 30**,** 534-44.

Kondo, H., K. S. Saleem & J. L. Price (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.

Kringelbach, M. L., J. O'Doherty, E. T. Rolls & C. Andrews (2003) Activation of the human orbitofrontal cortex to a liquid food stimulus is correlated with its subjective pleasantness. *Cereb Cortex,* 13**,** 1064-71.

La Camera, G. & B. J. Richmond (2008) Modeling the violation of reward maximization and invariance in reinforcement schedules. *PLoS Comput Biol,* 4**,** e1000131.

Lau, B. & P. W. Glimcher (2005) Dynamic response-by-response models of matching behavior in rhesus monkeys. *J Exp Anal Behav,* 84**,** 555-79.

--- (2008) Value representations in the primate striatum during matching behavior. *Neuron,* 58**,** 451-63.

Lauwereyns, J., K. Watanabe, B. Coe & O. Hikosaka (2002) A neural correlate of response bias in monkey caudate nucleus. *Nature,* 418**,** 413-7.

Leon, M. I. & M. N. Shadlen (1999) Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron,* 24**,** 415-25.

Louie, K., L. E. Grattan & P. W. Glimcher (2011) Reward value-based gain control: divisive normalization in parietal cortex. *J Neurosci,* 31**,** 10627-39.

Malkova, L., D. Gaffan & E. A. Murray (1997) Excitotoxic lesions of the amygdala fail to produce impairment in visual learning for auditory secondary reinforcement but interfere with reinforcer devaluation effects in rhesus monkeys. *J Neurosci,* 17**,** 6011-20.

McCoy, A. N., J. C. Crowley, G. Haghighian, H. L. Dean & M. L. Platt (2003) Saccade reward signals in posterior cingulate cortex. *Neuron,* 40**,** 1031-40.

McDonald, A. J. (1998) Cortical pathways to the mammalian amygdala. *Progress in Neurobiology,* 55**,** 257-332.

Miesenbock, G. (2009) The optogenetic catechism. *Science,* 326**,** 395-9.

Miller, E. K. & J. D. Cohen (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci,* 24**,** 167-202.

Minamimoto, T., G. La Camera & B. J. Richmond (2009) Measuring and modeling the interaction among reward size, delay to reward, and satiation level on motivation in monkeys. *J Neurophysiol,* 101**,** 437-47.

Morrison, S. E., A. Saez, B. Lau & C. D. Salzman (2011) Different time courses for learning-related changes in amygdala and orbitofrontal cortex. *Neuron,* 71**,** 1127-40.

Morrison, S. E. & C. D. Salzman (2009) The convergence of information about rewarding and aversive stimuli in single neurons. *J Neurosci,* 29**,** 11471-83.

--- (2011) Representations of appetitive and aversive information in the primate orbitofrontal cortex. *Ann N Y Acad Sci,* 1239**,** 59-70.

Murray, E. A. & A. Izquierdo (2007) Orbitofrontal cortex and amygdala contributions to affect and action in primates. *Ann N Y Acad Sci,* 1121**,** 273-96.

Musallam, S., B. D. Corneil, B. Greger, H. Scherberger & R. A. Andersen (2004) Cognitive control signals for neural prosthetics. *Science,* 305**,** 258-62.

Nakahara, H., H. Itoh, R. Kawagoe, Y. Takikawa & O. Hikosaka (2004) Dopamine neurons can represent context-dependent prediction error. *Neuron,* 41**,** 269-80.

Nakamura, K., A. Mikami & K. Kubota (1992) Activity of single neurons in the monkey amygdala during performance of a visual discrimination task. *J Neurophysiol,* 67**,** 1447-63.

Nishijo, H., T. Ono & H. Nishino (1988a) Single neuron responses in amygdala of alert monkey during complex sensory stimulation with affective significance. *J Neurosci,* 8**,** 3570-83.

--- (1988b) Topographic distribution of modality-specific amygdalar neurons in alert monkey. *J Neurosci,* 8**,** 3556-69.

O'Brien, T. 1990. *The things they carried : a work of fiction*. Boston: Houghton Mifflin/Seymour Lawrence.

Ongur, D., X. An & J. L. Price (1998) Prefrontal cortical projections to the hypothalamus in macaque monkeys. *J Comp Neurol,* 401**,** 480-505.

Ongur, D., A. T. Ferry & J. L. Price (2003) Architectonic subdivision of the human orbital and medial prefrontal cortex. *J Comp Neurol,* 460**,** 425-49.

Ongur, D. & J. L. Price (2000) The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cereb Cortex,* 10**,** 206-19.

Padoa-Schioppa, C. (2009) Range-adapting representation of economic value in the orbitofrontal cortex. *J Neurosci,* 29**,** 14004-14.

--- (2011) Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci,* 34**,** 333-59.

Padoa-Schioppa, C. & J. A. Assad (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature,* 441**,** 223-6.

--- (2008a) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat Neurosci,* 11**,** 95-102.

--- (2008b) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat Neurosci,* 11**,** 95-102.

Paton, J., M. Belova, S. Morrison & C. Salzman (2006a) The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature,* 439**,** 865-870.

Paton, J. J., M. A. Belova, S. E. Morrison & C. D. Salzman (2006b) The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature,* 439**,** 865-70.

Pavlov, I. P. 1927. *Conditioned Reflexes*. London: Oxford Univ. Press.

Penick, S. & P. R. Solomon (1991) Hippocampus, context, and conditioning. *Behav Neurosci,* 105**,** 611-7.

Pickens, C. L., M. P. Saddoris, M. Gallagher & P. C. Holland (2005) Orbitofrontal lesions impair use of cue-outcome associations in a devaluation task. *Behav Neurosci,* 119**,** 317-22.

Pickens, C. L., M. P. Saddoris, B. Setlow, M. Gallagher, P. C. Holland & G. Schoenbaum (2003) Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. *J Neurosci,* 23**,** 11078-84.

Plassmann, H., J. O'Doherty, B. Shiv & A. Rangel (2008) Marketing actions can modulate neural representations of experienced pleasantness. *Proc Natl Acad Sci U S A,* 105**,** 1050-4.

Platt, M. L. & P. W. Glimcher (1999) Neural correlates of decision variables in parietal cortex. *Nature,* 400**,** 233-8.

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.

Rempel-Clower, N. L. & H. Barbas (1998) Topographic organization of connections between the hypothalamus and prefrontal cortex in the rhesus monkey. *J Comp Neurol,* 398**,** 393-419.

Rescorla, R. A. & A. R. Wagner. 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In *Classical Conditioning II: Current Research and Theory,* eds. A. H. Black & W. F. Prokasy, 64-99. New York: Appleton Century Crofts.

Roesch, M. R. & C. R. Olson (2003) Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. *J Neurophysiol,* 90**,** 1766-89.

--- (2004) Neuronal activity related to reward value and motivation in primate frontal cortex. *Science,* 304**,** 307-10.

--- (2005) Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J Neurophysiol,* 94**,** 2457-71.

Rolls, E. T. & L. L. Baylis (1994) Gustatory, olfactory, and visual convergence within the primate orbitofrontal cortex. *J Neurosci,* 14**,** 5437-52.

Rolls, E. T., H. D. Critchley, R. Mason & E. A. Wakeman (1996) Orbitofrontal cortex neurons: role in olfactory and visual association learning. *J Neurophysiol,* 75**,** 1970-81.

Rolls, E. T. & F. Grabenhorst (2008) The orbitofrontal cortex and beyond: from affect to decision-making. *Prog Neurobiol,* 86**,** 216-44.

Romanski, L. M., J. F. Bates & P. S. Goldman-Rakic (1999) Auditory belt and parabelt projections to the prefrontal cortex in the rhesus monkey. *J Comp Neurol,* 403**,** 141-57.

Rudebeck, P. H., T. E. Behrens, S. W. Kennerley, M. G. Baxter, M. J. Buckley, M. E. Walton & M. F. Rushworth (2008) Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J Neurosci,* 28**,** 13775-85.

Salzman, C. D., J. J. Paton, M. A. Belova & S. E. Morrison (2007) Flexible neural representations of value in the primate brain. *Ann N Y Acad Sci*.

Satoh, T., S. Nakai, T. Sato & M. Kimura (2003) Correlated coding of motivation and outcome of decision by dopamine neurons. *J Neurosci,* 23**,** 9913-23.

Schultz, W. (2002) Getting formal with dopamine and reward. *Neuron,* 36**,** 241-63.

Schultz, W., P. Dayan & P. R. Montague (1997) A neural substrate of prediction and reward. *Science,* 275**,** 1593-9.

Shidara, M. & B. J. Richmond (2002a) Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science,* 296**,** 1709-11.

--- (2002b) Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science,* 296**,** 1709-11.

Simmons, J. M., S. Ravel, M. Shidara & B. J. Richmond (2007) A comparison of reward-contingent neuronal activity in monkey orbitofrontal cortex and ventral striatum: guiding actions toward rewards. *Ann N Y Acad Sci,* 1121**,** 376-94.

Simmons, J. M. & B. J. Richmond (2008) Dynamic changes in representations of preceding and upcoming reward in monkey orbitofrontal cortex. *Cereb Cortex,* 18**,** 93-103.

Skinner, B. F. 1938. *The behavior of organisms: an experimental analysis.* Englewood Cliffs, NJ: Prentice Hall.

Stuphorn, V., T. L. Taylor & J. D. Schall (2000) Performance monitoring by the supplementary eye field. *Nature,* 408**,** 857-60.

Sugase-Miyamoto, Y. & B. J. Richmond (2005) Neuronal signals in the monkey basolateral amygdala during reward schedules. *J Neurosci,* 25**,** 11071-83.

Sugrue, L. P., G. S. Corrado & W. T. Newsome (2005) Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat Rev Neurosci,* 6**,** 363-75.

Sutton, R. & A. Barto. 1998. *Reinforcement Learning*. Cambridge, Massachusetts: MIT Press.

Takikawa, Y., R. Kawagoe, H. Itoh, H. Nakahara & O. Hikosaka (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res,* 142**,** 284-91.

Thorndike, E. L. 1911. *Animal intelligence: experimental studies*. New York: MacMillan.

Thorpe, S. J., E. T. Rolls & S. Maddison (1983) The orbitofrontal cortex: neuronal activity in the behaving monkey. *Exp Brain Res,* 49**,** 93-115.

Tobler, P. N., C. D. Fiorillo & W. Schultz (2005) Adaptive coding of reward value by dopamine neurons. *Science,* 307**,** 1642-5.

Toda, K., Y. Sugase-Miyamoto, T. Mizuhiki, K. Inaba, B. J. Richmond & M. Shidara (2012) Differential encoding of factors influencing predicted reward value in monkey rostral anterior cingulate cortex. *PLoS One,* 7**,** e30190.

Tremblay, L. & W. Schultz (1999) Relative reward preference in primate orbitofrontal cortex. *Nature,* 398**,** 704-8.

--- (2000) Reward-related neuronal activity during go-nogo task performance in primate orbitofrontal cortex. *J Neurophysiol,* 83**,** 1864-76.

Tsuchida, A., B. B. Doll & L. K. Fellows (2010) Beyond reversal: a critical role for human orbitofrontal cortex in flexible learning from probabilistic feedback. *J Neurosci,* 30**,** 16868-75.

Tversky, A. & D. Kahneman (1981) The framing of decisions and the psychology of choice. *Science,* 211**,** 453-8.

Waber, R. L., B. Shiv, Z. Carmon & D. Ariely (2008) Commercial features of placebo and therapeutic efficacy. *JAMA,* 299**,** 1016-7.

Walker, A. E. (1940) A cytoarchitectural study of the prefrontal area of the macaque monkey. *J. Comp. Neurol.,* 73**,** 59-86.

Wallis, J. D. & E. K. Miller (2003a) Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci,* 18**,** 2069-81.

--- (2003b) Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci,* 18**,** 2069-81.

Wellman, L. L., K. Gale & L. Malkova (2005) GABAA-mediated inhibition of basolateral amygdala blocks reward devaluation in macaques. *J Neurosci,* 25**,** 4577-86.

Worthy, D. A., W. T. Maddox & A. B. Markman (2008) Ratio and difference comparisons of expected reward in decision-making tasks. *Mem Cognit,* 36**,** 1460-9.