

University of Oxford
Department of Experimental Psychology

Amygdala-prefrontal cortex interactions in the primate brain and the neural basis of credit assignment

Davide Folloni



A thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

Christ Church College
Trinity Term, 2019

Amygdala-prefrontal cortex interactions in the primate brain and the neural basis of credit assignment

Davide Folloni

Christ Church College

Trinity Term, 2019

A dissertation submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy at the University of Oxford

Abstract

An important adaption for life in changeable environments, such as those occupied by many primate species, is the ability to learn the contingencies that prevail between events and their consequences. Credit assignment consists in the capacity to correctly associate the value of an outcome to the event or choice that led to it. Interactions between amygdala and prefrontal cortex have been suggested to support credit assignment. **Chapter 1** reviews the anatomy and function of areas supporting value-guided learning and decision making that may contribute to credit assignment mechanisms. This chapter also reviews a group of methodological approaches that have been used or are currently being optimized to investigate these mechanisms. **Chapter 2** focuses on the reconstruction of the anatomical pathways supporting value-guided computations and that may provide the scaffolding for credit assignment to take place. Both macaque monkeys and humans are investigated using the same methodological approach which additionally makes possible direct comparison of anatomical and connectional features in the two species. **Chapter 3** demonstrates the potential of an offline focused transcranial ultrasound stimulation protocol applied at low-intensity to reversibly and minimally-invasively modulate the activity coupling of subcortical and deep cortical areas like the amygdala and anterior cingulate cortex of the primate brain without the necessity for surgical intervention. **Chapter 4** exploits these properties of focused transcranial ultrasound stimulation to investigate the behavioural deficits in credit assignment caused by the disruption of activity in the primate lateral orbitofrontal cortex. The behavioural results reported in this chapter suggest that ultrasound stimulation targeted to this area impairs the capacity to correctly assign credit to a previous choice, specifically following positive feedback, thereby disrupting the impact that the integrated history of choice option-outcome contingencies has on future behaviour. **Chapter 5** highlights the neural basis of credit assignment and the consequence of interfering with activity in lateral orbitofrontal cortex. After stimulation, brain activity encoding the choice-option outcome contingencies through credit assignment was disrupted but the response to reward only was not. **Chapter 6** reviews the results reported throughout the thesis aiming to connect them together in a similar framework and discuss them with reference to the current scientific literature on amygdala-prefrontal cortex connectivity and credit assignment.

Approximate word count: 55,724 words

This work is dedicated to my father who pushed me to pursue my goals until his last moments in life, teaching me the ethics of work and commitment.

Acknowledgments

The majority of the achievements, for a DPhil student, are the results of their capacity to “stand on the shoulders of giants”, adapting a metaphor used by Bernard De Chartres and Isaac Newton. Usually, the term “giants” refers to previous scientists and to their discoveries. However, here, it primarily refers to my DPhil supervisors and mentors Matthew Rushworth, Rogier Mars, Jerome Sallet and Lennart Verhagen. Having four supervisors is no doubt a challenging but very instructive experience. I am extremely thankful for their teaching, constant presence, constructive criticisms, extreme patience, kindness, wisdom, for being such a great inspirational example as scientists as well as for motivating me to do more when it was needed but also for pressuring me to slow down when it was necessary. Beyond these points, I am extremely thankful for their personal presence and humanity when I most needed it because of life events.

Science is the results of academic collaborations. I have been part of three fantastic Labs during my DPhil, full of amazing and hard-working people. I would like to thank all of them, but it would probably take 3 pages. Among all these people, I am especially thankful to Lev Tankelevitch, Nadesha Trudel, Elsa Fouragnan, Marco Wittman, Lea Roumazeilles, Miriam Klein-Flugge and Urs Schuffelgen. Also, thanks to the 4 Ps and to the Wellcome Trust.

I would also like to thank Anne Landau and Michael Winterdahl who accepted to supervise a barely speaking Italian weirdo who was looking for an internship in Denmark back in 2012. Their trust and initial support gave me the opportunity to be here so I will always be immensely thankful to them.

I am extremely thankful to all my family, especially my father Enzo Folloni and mother Teresa Colasanto. The sudden loss of my father at the beginning of this DPhil was a big hit but his motivational words in his last days at the hospital keep guiding me in trying my best at work but most of all in life.

Obviously, friends are an important aspect of life and I would especially like to thank the whole “Squadra Benessere”, especially the invaluable friendship of Andrea Vergnani, Luca Tedeschi and Marco Manfredini. Also, I cannot forget the MSc friends Cristiana Vagnoni, Michael Lohse and Silvia Cerea, as well as Piergiorgio Salvan. Also, thanks to F.R.I.E.N.D.S. and the Avengers.

Most importantly, all I have done would not have been possible without Azzurra Invernizzi, who, since the first days at the Tinbergen Building EP cafeteria, always motivated me to move on when I was feeling blue and who filled my days with happiness, love, enthusiasm, patience and support. I could not have done this without her presence and she definitely is the most important “discovery” of my DPhil.

Contribution to thesis

Advances in scientific research are pushed forward by honest and productive collaborations. The results of this thesis are as well the product of collaborative work with a network of people who partly contributed to the final work.

Chapter 2 reports data and results that have now been published in the following article:

- Folloni, D., Sallet, J., Khrapitchev, A.A., Sibson, N.R., Verhagen, L., Mars, R.B. (2019). *Dichotomous organization of amygdala/temporal-prefrontal bundles in humans and monkeys*. *eLife* 8:e47175. DOI: <https://doi.org/10.7554/eLife.47175>

Chapter 3 reports data and results that have previously been published in the following article:

- Folloni, D., Verhagen, L., Mars, R.B., Fouragnan, E., Constans, C., Aubry, J.-F., Rushworth, M.F.S., Sallet, J. (2019). *Manipulation of Subcortical and Deep Cortical Activity in the Primate Brain Using Transcranial Focused Ultrasound Stimulation*. *Neuron*, 101(6). <https://doi.org/10.1016/j.neuron.2019.01.019>

Chapter 4 and *5* reports unpublished data and results to which the following people contributed during data collection and part of the data analysis:

- Folloni, D., Fouragnan, E., Wittman, M., Roumazeilles, L., Tankelevitch, L., Verhagen, L., Sallet, J., Rushworth M.F.S

Contents

Abstract	1
Acknowledgements	3
Contribution to the thesis	4
Contents	5
Chapter 1: General Introduction	10
1.1 Introduction	10
1.2 Evolution and anatomy	16
1.2.1 Evolution of the primate brain	16
1.2.2 Evolution of the prefrontal cortex	19
1.2.3 Orbitofrontal cortex	25
1.2.4 Amygdala	30
1.2.5 Orbitofrontal cortex–amygdala connectivity	34
1.3 The role of the orbitofrontal cortex and amygdala in associative learning and credit assignment	38
1.3.1 The unsolved orbitofrontal dilemma: what does the orbitofrontal do?	39
1.3.2 A model for value comparison in orbitofrontal cortex	50
1.3.3 Amygdala	55
1.3.4 Amygdala-orbitofrontal cortex interactions	59
1.4 Methodological approaches to investigate anatomy and functions	63

1.4.1 Diffusion Magnetic Resonance Imaging and Tractography	63
1.4.2 Functional Magnetic Resonance Imaging	68
1.4.3 Non-Invasive Brain Stimulation	76

Chapter 2: Dichotomous organization of amygdala/temporal-prefrontal bundles in humans and monkeys

2.1	ABSTRACT	83
2.2	INTRODUCTION	84
2.3	MATERIALS AND METHODS	88
2.3.1	Data acquisition	88
2.3.2	Data preprocessing	91
2.3.3	Probabilistic tractography	92
2.3.4	Connectivity fingerprints	97
2.4	RESULTS	100
2.4.1	Amygdalofugal pathway in the macaque monkey brain	102
2.4.2	Amygdalofugal pathway in the human brain	105
2.4.3	Uncinate fascicle in the macaque monkey brain	108
2.4.4	Uncinate fascicle in the human brain	109
2.4.5	Comparison of the two fiber bundles	110
2.4.6	Prefrontal connectivity and control tracts	112
2.5	DISCUSSION	120

Chapter 3: Manipulation of subcortical and deep cortical activity in the primate brain using transcranial focused ultrasound stimulation

3.1	ABSTRACT	127
-----	----------	-----

3.2	INTRODUCTION	128
3.3	MATERIALS AND METHODS	129
3.3.1	Experimental model and subject details	129
3.3.2	Ultrasound stimulation	130
3.3.3	Macaque MRI acquisition	132
3.3.4	Macaque MRI preprocessing	134
3.3.5	Macaque rs-fMRI connectivity analysis	137
3.3.6	Macaque rs-fMRI statistical inference	139
3.4	RESULTS	143
3.4.1	Stimulation of deep brain structure and resting-state fMRI recording	143
3.4.2	Effects of TUS on subcortical neural activity in the Amygdala	145
3.4.3	Effects of TUS on deep cortical neural activity in anterior cingulate cortex (ACC)	148
3.4.4	Focality of TUS effect	152
3.4.5	Putative auditory effects of offline TUS	155
3.5	DISCUSSION	158

Chapter 4: Causal role of the lateral orbitofrontal cortex in credit assignment

4.1	ABSTRACT	164
4.2	INTRODUCTION	165
4.3	MATERIALS AND METHODS	168
4.3.1	Subjects	168

4.3.2	Behavioural training	169
4.3.3	Experimental task	169
4.3.4	Transcranial focused ultrasound stimulation	176
4.3.5	Reinforcement learning model	179
4.3.6	Statistical analyses	184
4.4	RESULTS	185
4.4.1	Defining choices, rewards, and their conjunction	185
4.4.2	Ultrasound stimulation of lateral OFC disrupts credit assignment	186
4.4.3	Credit assignment deficit impairs subsequent reward-guided adaptive behaviour	190
4.4.4	Ultrasound stimulation of the adjacent anterior PFC does not affect credit assignment or subsequent adaptive behavior	191
4.4.5	Changes in win-stay/lose-switch choice strategy after modulating activity in lateral OFC	194
4.4.6	Lateral OFC/TUS differentially modulates learning from positive and negative outcomes	196
4.4.7	Anterior PFC TUS has no effects on adaptive behavior following either positive nor negative outcomes	198
4.5	DISCUSSION	202

Chapter 5: Neural basis of credit assignment and manipulation of credit assignment-related activity with focused transcranial ultrasound stimulation

5.1	ABSTRACT	209
-----	----------	-----

5.2	INTRODUCTION	210
5.3	MATERIALS AND METHODS	213
5.3.1	Subjects	213
5.3.2	MRI data acquisition	213
5.3.3	MRI data preprocessing	214
5.3.4	Whole-brain fMRI analysis	217
5.3.5	Whole-brain GLMs	218
5.3.6	Region-of-interest (ROI) analyses	222
5.4	RESULTS	223
5.4.1	Whole-brain changes in the representation of the chosen value signal after lateral OFC TUS	223
5.4.2	Lateral OFC TUS disrupts signal encoding the expected value of the chosen option within the amygdala-OFC circuit	225
5.4.3	Lateral OFC TUS does not affect the reward signal	227
5.4.4	Conjunctive differences of expected value of the chosen option and reward between sham and lateral OFC TUS	229
5.4.5	Role of the lateral OFC in adaptive behavior within a volatile environment	231
5.5	DISCUSSION	233
	Chapter 6: General Discussion	240
	Appendix	251
	Bibliography	317

Chapter 1

General Introduction

"It's not hard to make decisions once you know what your values are"

Roy. E. Disney

1.1 Introduction

Where should I go on holiday to finally have some fun and relax? If there is something I have quickly learnt as a DPhil student it is that academic life does not leave much time off, except for a couple of times a year. The consequence is that I need to carefully decide where to book my holidays by taking into account variables like, for instance, during which months of the academic year I will take time off work. This is because I am a big fan of scuba diving destinations but to fully enjoy this sport while on vacation I need to organize trips to different countries depending on when I go. Indeed, I have learnt from experience that going to a novel diving destination during the wrong period of the year may lead to unpleasant surprises such as very windy weather and low underwater visibility.

After a few years of practicing this sport I have learnt that choosing to go to certain holiday locations at the wrong time of the year may not result in

me finding what I had expected to find. If I have time to go only over certain months, then rather than returning to a known destination, exploring alternative vacation options should be considered. On the contrary, sticking always with the same diving spot all the year around, means that you may also be there during the wrong season (i.e. during monsoon season in south-east Asia or west Africa) and will probably also result in unsatisfying experiences primarily characterized by only very few days of scuba diving and many others waiting in the hotel for the rain to stop. After years spent travelling to different destinations and testing them at different times of the year, my current strategy for booking holidays is based on evaluating the past outcome of having chosen to book in one place rather than in another. Based on this past experience, I have now learnt to estimate how likely I am to receive the type of scuba diving vacation I hoped for when choosing to go to one holiday destination rather than another at a certain time of the year. Indeed, such previously acquired knowledge now helps me in flexibly deciding whether to keep exploiting a familiar diving spot based on the known likelihood of a rewarding holiday or making a more risk-taking choice and go exploring a novel destination that may potentially (or not) provide me with an even more enjoyable experience.

Everyday life in the Department involves a constellation of similar, though less exciting, choices. Daily decisions on whether to select a mild or strong coffee at the vending machine (highly influenced by how much caffeine I have already had previously), choosing whether going for lunch at the usual cafeteria or rather trying the new, recently opened, sandwich shop as well as deciding whether to submit the new exciting results to an

“adventurous” high-impact journal or to another one that could ensure a smoother and quicker publication process.

The human brain has evolved as part of primate evolution to be able to deal with complex decisions imposed by the stochastic rules governing our ever-changing world. Obviously, our ancestors did not have to decide where to go in order to maximize their chance of scuba-based enjoyment. However, they did have to make analogous choices when it came to deciding which patch to go hunting in during a specific period of the year or which bay to go diving in when in search of fish to eat. Humans, like macaque monkeys (such as rhesus macaques) and apes (such as chimpanzees and gorillas), are classified by evolutionary taxonomy as being part of the order of primates, more specifically “anthropoid primates” (from Latin *primate*-, from *primus*: “prime, first rank”) or “simians”. Humans, apes, and macaque monkeys, together, are part of a group of mammals that descended from a common ancestor that lived about 30 million years ago (**Figure 1**). They share specific anatomical and functional traits that are directly linked to their behaviours and, consequentially, to their capacity to forage flexibly adapting to the ecological selection pressures they encounter in their environments and, therefore, increase survival rate.

Although humans, apes, and macaque monkeys at some point of evolution split into three separate evolutionary branches, their ancestors still lived in environments governed by similar selection pressures. This ecological similarity influenced and promoted the evolution of brains characterized by many anatomical homologies. Such similarity in anatomy was paralleled by similarities in cognitive processes like learning and

decision-making, which in turn were reflected by similarities in behaviours like foraging. This is, of course, not to say that there is complete similarity between the brains of humans, apes and monkeys and indeed some important differences will be noted at later points in this thesis.




Figure 1. Phylogenetic tree of the primate brain as derived from genotyping examinations. (Reproduced from Heuer et al., 2018).

The aim of this thesis is to explore some of the key circuits for decision-making and learning in the primate brain, from an anatomical,

behavioural, and computational perspective. Initially, I will especially focus on the rhesus macaque monkey (*Macaca mulatta*) and human brains since similarities in major anatomical systems and key learning and decision making behaviours between these two species have been described (Barton and Harvey, 2000; Neubert et al., 2015, 2014; Ongür and Price, 2000; Passingham and Wise, 2012; Price, 2007; Rushworth et al., 2011; Wallis, 2011). The use of the macaque monkey as an animal model is still necessary as its neuroanatomy and physiology are, to date, better known and understood than the human brain. This more thorough description of the macaque monkey brain is the result of the previous careful investigations in this species obtained from many experimental approaches ranging from histology, single-cell electrophysiology, and lesion studies. Although a similar knowledge of the human brain is the current goal of neuroscience, the invasiveness of some of the methods needed precludes many types of investigations in humans. However, one aim of this thesis is to present novel non-invasive or minimally invasive methodological approaches that can be tested in the macaque and then potentially applied to human brain investigation in the future. I will briefly describe some of these approaches in *Chapter 1* as their methodological development and refinement are providing neuroscientists with novel tools to comparatively investigate similarities and differences in anatomy and cognition between macaque monkeys and humans as well as to better elucidate the neural mechanisms underlying learning and decision-making.

Over the following chapters, I will attempt to develop and exploit new neuroimaging approaches to compare the underlying anatomy of specific

brain circuits in the monkey and human brains. I will then report and present the development of a novel technique that can be used to transiently and reversibly manipulate activity in subcortical and deep cortical brain structures transcranially without the necessity to perform any craniotomy. Although this stimulation technique has already been applied to humans to test, primarily, for its capability to induce changes in primary sensory systems (Lee et al., 2016; Legon et al., 2014) or in severe acute medical conditions (Monti et al., 2016), a full mapping and definition of the parameters appropriate for safe application in humans is currently under discussion (Blackmore et al., 2019; Jerusalem et al., 2019). We have therefore developed and exploited a particular powerful form of the approach that as I show later causes neural and behavioural changes over a longer time, of minutes or hours, in contrast to previous techniques. I have applied this approach to the macaque monkey with the goals of (1) successfully modulating activity in deep cortical structures and (2) observing the causal contributions of areas to learning and decision-making. In the experimental chapters presented throughout this dissertation, I will specifically focus on a network of areas interconnecting the amygdala and prefrontal cortex, and that underlie value-guided learning and decision-making.

1.2 Evolution and Anatomy

1.2.1 Evolution of the primate brain

Charles Darwin's central message in his "On the Origin of Species" (Darwin, 1859) constituted a complete revolution in the way we think about the relationship between structure and function in biology and it has had an impact on the way in which we think about the brain and behaviour. His theory on the mechanisms of evolution and the idea humans and monkeys shared an evolutionary inheritance was a shock for the society of the time. According to his hypothesis, over evolution, different ecological pressures in different environments acted in different ways on individuals whose ancestors had been members of a common primate species. This resulted in the selection and promotion of traits and genetic variations that differed from environment to environment; selection favoured traits needed for adaptation and survival in the specific niche that the group of animals occupied. Therefore, different ecological pressures in two different niches will result in the promotion of different traits that will be transmitted and inherited by the descendants. The consequence of this process of "natural selection", is the fine-tuning of both anatomical and cognitive features so that individuals have an increased capacity to face predatory dangers, i.e. foraging during the day in open fields, to ensure a higher likelihood of survival, and reproduction. Dunbar (2009) has argued that in primates such adaptations have included the development of sociality. In summary, the natural selection process has led an initially common primate brain to differentiate into what we now consider as different species such as macaque monkeys, chimpanzees, and humans over millions of years.

As mentioned above, natural selection in the anthropoid primate brain promoted the development of anatomical and behavioural characteristics that could help the individual in the complex daily task of foraging. A distinctive feature of simian anthropoid species, compared to prosimians species (another class of primates mostly nocturnal and, therefore, characterized by different foraging strategies) is their capacity to differentially categorize the objects they encounter in their environment based on physical features such as colours, shape, and size. Moreover, the visual systems of these species have also developed better perception of depth, which further optimized their foraging and defensive abilities in large open fields (Passingham and Wise, 2012). Anatomically, the improved skills in foraging and object recognition had been made possible by the development and specialization of functional modules in the occipital and temporal lobe cortex like area MT/V5 (Kaas, 2006) (**Figure 2**).

Another key “anatomical innovation” strongly linked to the development of specialized visual circuits is to be found in the eye itself. The development of a dense population of cone cells, photoreceptors specialized in transducing light into electrical signals, led to the formation of a portion of the retina of the eye called the central fovea or *fovea centralis*. This small, specialized, portion of retina located on the posterior wall of our eyes is one of the “anatomical advancements” that ecological pressures selected over evolution in order to improve our adaptation to the environment. Indeed, this concentration of cone cells enabled our ancestors to develop extremely acute vision enabling them to distinguish between small features in different objects such as colours or shapes during foraging.

Once again, evolution promoted the specialization of anatomical brain features that, synergistically with other portions of the cerebrum like the prefrontal cortex (PFC), were capable of optimizing daily behaviours necessary for survival and environmental adaptation. Just think about the relevance of this ancient small anatomical modification nowadays when you go grocery shopping or need to decide which Nespresso cup to use for your morning coffee. Alternative ecological pressures in different environments may have instead led to the development of other anatomical features. This is the case of the evolution of the eye anatomy in nocturnal prosimians species like members of the different families of Galagos, also known as bushbabies (Clark, 1985; Diogo et al., 2015; Kaas, 2013).




Figure 2. Representation of the four lobes of the human cerebral cortex.
(Reproduced from Goldberg, 2001).

The development of improved visual abilities in object recognition and an associated refinement of the underlying neural circuit are not the only changes in the anthropoid primate brain induced by ecological pressures. Comparative neuroscience has shed light on several other variations that took place in the brains of humans, apes and macaque monkeys. The neocortex (or isocortex), which constitutes the largest portion of cerebral cortex in primates, has undergone a number of changes over the course of evolution. Alongside changes in gyrification of the cerebral cortex and sulcal patterning (Heuer et al., 2019; Van Essen and Dierker, 2007), and the evolution of specialized sensorimotor regions in parietal and frontal cortex endowed primates with finer skills in reaching, grasping and tool manipulation, which in turn improved their object-based foraging strategies (Donahue et al., 2018; Kaas, 2013, 2008; Passingham, 2008). These patterns of change are apparent even when brain size as well as both absolute and relative scaling are taken into account and controlled for (Donahue et al., 2018; Mars et al., 2017). Although some studies argue against the hypothesis that there has been a general increase in association cortex over the course of primate evolution (Barton and Venditti, 2013; Gabi et al., 2016), several other reports suggest specific parts of the association cortex such as PFC have increased in size (Mars et al., 2017; Passingham and Smaers, 2014; Preuss, 2011).

1.2.2 Evolution of the prefrontal cortex

The evolution of the primate brain has fascinated researchers since Darwin put forwards his theory. A specific portion of the frontal lobe, the PFC, has

drawn particular attention and there is ongoing debate about the nature of anatomical correspondences and changes across species. An extensive description of the full debate is outside the remit of this thesis but for a review on the topic see Sherwood and Smaers (2013).

The first major issue in understanding evolutionary changes in this region arises from the definition of the PFC (a term which was first used by Ferrier and Yeo, 1884). Different definitions of PFC have been used historically and a variety of definitions are still in use today. Rose and Woolsey (1948), for instance, suggested that the PFC can be defined as that portion of frontal cortex that receives projections from the mediodorsal nucleus of the thalamus. Others, described PFC as those frontal areas that receive inputs from dopaminergic neurons (Divac et al., 1978; Porrino and Goldman-Rakic, 1982). These arguably incorrect anatomical descriptions still sometimes persist in current neuroscience (Carlén, 2017) and contribute to inconsistencies in anatomical descriptions across species.

Throughout this dissertation, I will define the PFC as a group of brain areas involved in higher cognitive functions distinguished by their rostral position in the frontal lobe and by an intrinsic anatomical feature: the presence of a layer IV characterized by a dense expression of granule cells. This entails that the PFC is a region only found in primates (**Figure 4**) where it corresponds to a portion of cerebral cortex which lies rostral to the primary motor cortex (Brodmann area 4) and to the agranular premotor cortex (Brodmann area 6). Architectonically, this definition of PFC is consistent with Brodmann's first description of the granular frontal region. These properties made these prefrontal regions distinguishable from other, more caudal, areas

such as the premotor and motor cortex located closer to the central sulcus. By contrast premotor and motor areas were defined as “precentral” (Brodmann, 1909).

Nevertheless, even today, it is still argued that the functions of primate PFC are handled by a region at the front of the brain of other non-primate mammals such as rats even if the region is different in some respects and lacks granule cells (Kolb, 2007; Schoenbaum et al., 2009; Seamans et al., 2008). Although there is considerable uncertainty about the degree to which brain areas in non-primates resemble primate PFC, nowadays, it is, widely, and generally recognized that the PFCs in a variety of anthropoid primates such as macaques, apes, and humans all share many similar cytoarchitectonic regions (Barbas, 2015, 2007; Wise, 2008).




Figure 3. Brodmann's definition of areas part of the prefrontal cortex in the human brain. Lateral, dorsal, ventral (A), medial, including cingulate, and orbital (B) regions of the prefrontal cortex showing a granular cytoarchitecture (Reproduced from Carlén, 2017).

Ecological forces and natural selection have therefore contributed to shape the primate brain over evolution in order to improve successful foraging. Anthropoid primates lived (and still live) in a volatile environment. Events like food shortage within a specific location that in the past might have been associated with a high likelihood of providing fruits to the foraging monkey (or, in the case of human primates, coffee at the usual vending machine just before an important meeting) and seasonal climate changes or competition for resources from a specific patch, all required our ancestors to be able to flexibly adapt their foraging strategies and quickly learn new successful ones to survive. In order to do so, early primates were required to detect specific stimuli or cues in the environment and learn to associate them with their probability of leading to a certain outcome, either aversive or appetitive in nature. However, they also had to be able to learn when changes in such contingencies took place in order to adjust their behaviour accordingly.

Daily life in volatile environments, therefore, drove the development of cognitive skills in order to predict the likely future consequence of the choices that primates had to make during foraging. The fine-grained pattern of connections between granular PFC and areas in the inferior temporal cortex (Saleem et al., 2008) contributed to improving (1) detection of such stimuli but importantly enhanced the primate brain's capacity to (2) associate them with the outcomes that followed their selection or appearance, thereby influencing and guiding future decisions within a similar context, a process that I will refer to as associating or "assigning credit". For instance, consider once again the case of scuba diving holidays. After having gone on a few

trips to various destinations at different times of the year, and after having experienced different levels of enjoyment as a consequence, I have now learned the features of a holiday destination including visual features of the environment and other features such as the timing of its occurrence that are associated with a satisfactory diving experience. Thanks to primate brain evolution, I can now identify other situations likely to yield similar rewards (and save money by avoiding booking scuba diving holidays in less enjoyable destinations).




Figure 4. Comparison of granularity of the frontal cortex in human, macaque monkey and rat brains. Granular (blue and cyan), dysgranular (light blue), agranular (green) cortices and allocortex (yellow) are shown in the brain of humans, monkeys and rats for regions located on the medial (top

row) and orbital (bottom row) portions of the frontal cortex. Human (a) and non-human primate (b) brains show a clear and homologous cortical architecture organized in highly and lightly granular (dense expression of granule cells in layer 4), dysgranular and agranular cortices, the latter typology lying adjacent to the allocortex and subdivided into anterior cingulate, infralimbic and prelimbic cortices. It is possible to observe how the thickness of layer 4 gradually increases from caudal to rostral areas in the frontal cortex on both medial (top row) and orbital (bottom row) surfaces. Cytoarchitectonic differences with the rat brain (c) in these granular properties can be observed (Reproduced from Wallis (2011) but see also Carmichael and Price (1994) and Wise (2008)).

Lesion and neuroimaging studies have repeatedly emphasized the importance of PFC in value-guided decision-making (Jocham et al., 2016; Parker and Gaffan, 1998; Passingham et al., 2000). More specifically, the orbitofrontal (OFC) portion of PFC seems to play a key role in decision-making and in the representation of stimuli and outcomes during associative learning (Noonan et al., 2011; Tremblay and Schultz, 1999; Wilson et al., 2014). However, OFC does not perform these computations in isolation but does so in orchestration with other regions such as the amygdala (Chau et al., 2015; Costa et al., 2016; Izquierdo et al., 2004; Jones and Mishkin, 1972; Lichtenberg et al., 2017; Morrison et al., 2011; Mosher and Rudebeck, 2015; Murray and Wise, 2010; Rudebeck and Murray, 2008; Saez et al., 2017).

In light of this consistent evidence of the role played by OFC-amygdala circuits in value-guided decision-making during associative learning, in this thesis I will focus specifically on these two portions of the brain, their connectivity, and their role in credit assignment.

1.2.3 Orbitofrontal cortex

The term PFC is a broad anatomical definition usually referring to an agglomerate of areas of the primate brain that are involved in a heterogeneous set of higher-order cognitive processes. Due to its size and heterogeneous influences on behaviour, the PFC is often subdivided into several sub-sections on the basis of anatomical location and functional criteria such as *dorsolateral (dlPFC)*, *dorsomedial (dmPFC)*, *rostral (or anterior PFC)*, *ventromedial prefrontal cortex (vmPFC)* in addition to *OFC* (**Figure 5**).

Here, I define the OFC as the group or areas located on the orbital surface of the primate brain that extends from the medial orbital sulcus to the inferior prefrontal convexity, thereby including Brodmann areas 11 (11m and 11), 13, 47/12m and 47/12o (sometimes simply referred to as 12o in macaque brain following the nomenclature employed by Walker (1940)). By contrast I will use the term vmPFC to refer to the subdivision of PFC lying on the medial wall in between the two hemispheres and that spreads along the gyrus rectus including Brodmann areas 14m, 14r and 14c. Posterior to OFC and vmPFC, are three anterior cingulate areas on the cingulate gyrus and cingulate sulcus including areas 24, subgenual area 25, and anterior perigenual cingulate 32. Rostrally, OFC and vmPFC are delimited by the frontopolar cortex (area 10). Because of functional similarities, vmPFC regions along the gyrus rectus and medial OFC structures (medial to the lateral orbital sulcus) are sometimes clustered together and jointly labelled as vmPFC/mOFC (Rushworth et al., 2011).




Figure 5. Anatomy of the prefrontal cortex. The prefrontal cortex is often reported as organized in six subdivisions on the basis of the cytoarchitectonic features of different regions: dorsolateral, dorsomedial, rostral (or anterior PFC), ventromedial and orbitofrontal. However, there is some inconsistency in the way that different authorities use terms. The blue area shown on the parasagittal brain picture on the right is labelled OFC by Szczepanski and Knight (2014) but it would correspond to the vmPFC described in this thesis. By contrast the green region referred to as ventromedial by Szczepanski and Knight (2014) corresponds to part of the anterior cingulate cortex as described in this thesis (Reproduced from Szczepanski and Knight, 2014).

The organization of the ventral PFC can also be described at a finer level on the basis of the cytoarchitectonic features of individual areas. Such approaches are based on comparing the thickness and granularity of the different cortical layers within each area (Barbas, 2007; Mackey and Petrides, 2010; Pandya et al., 1988). Using this approach, the ventral PFC shows a gradual but systemic shift from an agranular to a granular or eulaminate cortical structure as we move from caudal to more rostral areas. This can be appreciated in **Figure 4**. By using this approach, not only the ventral regions but the entire PFC can be primarily subdivided in three major divisions: agranular cortex, dysgranular cortex (characterized by the

formation of an incipient but still poorly developed layer IV) and granular or eulaminate cortex (clearly developed layer IV). Carmichael and Price (1994) described the organization of ventral PFC in even more details using a combination of histochemical and immunohistochemical staining approaches aimed to differentiate portions of ventral PFC on the basis of a range of structural, synaptic and metabolic properties related to neuron, myelinated fibers and neurotransmitters within its regions. By using this approach, the authors were able to distinguish on the basis of their architectonic and morphological properties twenty-two separate cortical region in ventral PFC.

Lastly, the organization of ventral PFC has also been described based on its connectional features (Ongür et al., 2003). Indeed, a unique feature of the ventral PFC (including both OFC and vmPFC) is its dense and complex connections that link these brain regions with limbic, thalamus and brain stem structures. Carmichael and Price (1995) used connectional anatomy data to compare and contrast vmPFC and OFC regions. They argued that medial and orbital frontal cortex could be subdivided into two broad subdivisions that correspond approximately to the vmPFC and OFC subdivisions that I have described. Moreover, OFC has strong connections with several types of sensory association cortex including somatosensory, olfactory, gustatory, auditory and visual (Schmahmann and Pandya, 2006; Carmichael and Price, 1994; Ongur and Price, 2000; Price, 2005). This means that it receives multi-modal information from several areas and can, reciprocally, influence representations in the same areas (Passingham and Wise, 2012). In conjunction vmPFC and OFC are at the centre of a functional

network involved in coordinating and modulating value-guided learning and decision-making.

Independently from the approach taken to describe the organization of OFC/vmPFC, one of its notable features is its strong connectivity with all sensory cortices and limbic structures. The connections include input fibers from several temporal regions such as the amygdala, piriform and inferior temporal cortices (McDonald, 1998; Ongür and Price, 2000; Petrides and Pandya, 1988) and from structures in the brain stem and the mediodorsal nucleus of the thalamus (Ray and Price, 1992). For instance, electrical stimulation of fibers connecting OFC/vmPFC to visceral regions including the brainstem results in several autonomic and mood-related alterations that can be observed in both neural activity and behavior (Lehman et al., 2011; Mayberg et al., 2005; Robinson and Mishkin, 1968). This widespread connectivity profile enables OFC/vmPFC to constantly monitor both internal and external environments independently from the context, that is it enables OFC/vmPFC to continuously receive both visceral and motivational information concerning the current state of the individual as well as information about external events happening in the environment. These functions may depend especially on areas 13 and 11 that have been suggested to be key hubs for information convergence in PFC (Passingham and Wise, 2012). Through their connectivity patterns, areas 13 and 11 are indeed capable of integrating sensory, visceral, and cognitive information into a single representation of the current state of the subject (Rolls and Baylis, 1994; Rolls, 1996; Thorpe et al., 1983; Schoenbaum et al., 2011; Wilson et al., 2014) and use this information to represent alternative choices in a

common, therefore comparable, framework during decision-making (Padoa-Schioppa and Assad, 2008, 2006).

More lateral orbital regions such as monkey area 12 (often referred to as 47 in the human brain), located on the medial wall of the ventrolateral PFC, also share a similar complex connectivity with other sensory and associative areas, especially in the temporal cortex (Amaral and Price, 1984; Carmichael and Price, 1995; Ghashghaei et al., 2007; Ghashghaei and Barbas, 2002; (Schmahmann and Pandya, 2006). The connectional properties of this area support the role of its neurons in associative learning, object discrimination and reward sensitivity (Carmichael and Price, 1995). Indeed, Kobayashi and colleagues (2010) confirm these results by demonstrating that neurons in vIPFC, and particularly some of those within area 12, show outcome-related properties that are sensitive to small differences in reward. These results are in line with the suggestion of Thorpe and colleagues (1983) that neurons in ventrolateral PFC (vIPFC) respond to the presentation of visual stimuli. Moreover, Rich and Wallis (2014) have linked neuronal activity in these regions with the encoding of the value of external stimuli. Lastly, Kennerley and Wallis (2009) suggested that neurons in the more lateral part of vIPFC (although some neurophysiological recordings overlapped with areas 12) encode expected reward information that may eventually influence attentional and working-memory processes.

Recently, comparative approaches that relied either on cytoarchitectonic or connectional classification of OFC regions as well as imaging techniques have suggested an overall similarity in the organization

of ventral PFC architecture in both macaques and humans (Mackey and Petrides, 2010, 2009; Neubert et al., 2015, 2014).

1.2.4 Amygdala

The amygdala is a phylogenetically old brain structure located deep in the medial cortex of the rostral temporal lobe and is often described as one of the main subcortical areas of the limbic system. Anatomically, it is constituted by a heterogenous constellation of several reciprocally connected nuclei and such complexity is conserved across anthropoid primate brains (**Figure 6**). Based on histological criteria, roughly thirteen nuclei have been found to constitute the amygdaloid complex and each of them has its own intra- and extra-amygdalar connections (Freese and Amaral, 2009). In **Figure 6**, it can be observed how basic aspects of morphological organization in the amygdala are preserved across macaque monkeys and humans in a subdivision between a basolateral complex and a central nucleus of the amygdala. The basolateral complex is a group of nuclei comprising primarily lateral, basal, accessory basal, and basomedial nuclei. Specifically, the lateral nucleus is considered to be the main input region of the amygdala, receiving sensory stimuli from sensory cortices and thalamus as well as prefrontal areas (Barbas, 2007; Ghashghaei and Barbas, 2002; LeDoux, 2007). Within the amygdala, nuclei comprising the basolateral complex are strongly interconnected with one another, with the GABAergic intercalated cells, and with the central nucleus (Janak and Tye, 2015; Likhtik et al., 2008).

From a cytoarchitectonic perspective, the amygdala is characterized by a non-laminated architecture suggesting its evolutionary and

developmental origins are distinct to those of other limbic structures such as the hippocampus (Pandya et al., 2015).

To understand the functional role of the amygdala, however, a clear comprehension of its connectivity and anatomy is fundamental. As for the OFC, it is the pattern of connections of this limbic structure with other regions that determines its functions. Indeed, these connections determine the information the amygdala has access to, as well as the control it can have on other regions (Passingham et al., 2002). The connectivity of the amygdala includes projections of its nuclei to striatum, thalamus, bed nucleus of the stria terminalis, hypothalamus and brain stem, and association cortex including parts of PFC (Aggleton, 2000; Aggleton et al., 1980; Catani et al., 2013; Cho et al., 2013; Ghashghaei et al., 2007; McDonald, 1992, 1991).




Figure 6. Evolution of the Amygdala. The amygdala is an almond shaped structure located deep into the temporal lobe and part of the limbic system. The connectivity of this brain region and the internal organisation of its nuclei have been very well conserved over evolution. The figure illustrates the evolutionarily preserved separation into a basolateral complex (BLA) and central nucleus (CeA) in different primate species: the human and macaque (Adapted from Janak and Tye, 2015).

Within the basolateral complex (BLA) it is specifically the lateral nucleus that is considered to be the main input region of the amygdala, receiving sensory stimulus related information from sensory cortex and thalamus. Axonal projections leaving this nucleus primarily innervate more medial nuclei (i.e. basal and accessory basal nuclei) as well as the central nucleus (Aggleton, 2000; Ghashghaei et al., 2007; LeDoux, 2007). On the contrary, the basal nucleus sends primarily feedback projections back to the sensory cortex (Price, 2003). BLA nuclei also shares strong bidirectional excitatory connections with PFC areas, particularly with subgenual and posterior OFC/vmPFC cortices. These fibers have been suggested to regulate affective processes and associative learning (Barbas, 2007; Baxter and Murray, 2002; Ghashghaei et al., 2007; Morrison and Salzman, 2010). The majority of neurons in the central nucleus are GABAergic and are subdivided into lateral (CeL) and medial (CeM) sub-regions. Neurons in the CeM are often considered as “output neurons”, in fact their axons branch widely outside the amygdala towards the paraventricular nucleus of the thalamus, brain stem (including periaqueductal grey) and several other brain regions involved in aversive and autonomic responses. The CeL is itself further divided into two groups of cells projecting to the CeM. These two populations are called CeL-OFF and CeL-ON and are respectively inhibited or excited during aversive stimulation in fear conditioning experiments. Importantly, CeL-OFF and CeL-ON neurons additionally exhibit reciprocal GABAergic inhibition of one another as well as a different latency of response (shorter for CeL-ON), suggesting they may have different weights in regulating aversive responses. A potential mechanism might involve

inhibition of CeL-OFF cells (longer response latency) by CeL-ON neurons. This mechanism may lead to an additional modulation of the expression of avoidance behaviour taking place already within the central nucleus of the amygdala. Furthermore, a portion of CeL neurons receive inputs from the lateral nucleus and send direct projections to the paraventricular nucleus of the thalamus and the periaqueductal grey. This latter network may represent an additional circuit through which CeL neurons can directly modulate aversive responses, thereby constituting a parallel circuit which does not need to synapse first in the CeM (for a full review see Janak and Tye, 2015).

The widespread connectivity of the amygdala with sensory multimodal, limbic, associative and PFC areas, alongside the complex intrinsic microcircuits just described, enables the amygdala to influence several cognitive and affective processes such as reward processing and value-guided decision-making, aversive conditioning, stress and fear responses, emotional processing as well as complex social behaviours (LeDoux, 2007; Whalen et al., 1998). Throughout the first part of this Chapter I have discussed how the anthropoid primate brain evolved to adapt to live in a diurnal environment, where ecological pressures required the early ancestors of macaques, apes and humans to mainly rely on vision as the primary sensory modality. Primate evolution, therefore, involved development and elaboration of occipital-temporal brain areas (Kaas, 2006; Saleem et al., 2008) that could mediate visually guided foraging (i.e. detailed object recognition as well as an improved capability to reach and grasp such objects) and predator detection. Quick and efficient processing of appetitive and aversive cues requires, however, a precise, quick and reciprocal

coordination of activity within visual system and amygdala, as well as between amygdala and sympathetic nuclei in diencephalic, forebrain, prefrontal and brain stem areas. From a comparative perspective, it is therefore fascinating to observe differences between the brain of anthropoid primates and the brain of rodents like rats that are instead nocturnal animals mostly relying on smell to detect dangerous or rewarding stimuli. Rats, macaques and humans, all share similarities in the input connections that their amygdalae receive from sensory cortices such as those in the visual system (Price and Amaral, 1981; Yukie, 2002). Interestingly, however, in macaques and humans (Price, 2003), but not in rats (Krettek and Price, 1977), return projections from the amygdala project back to visual areas beyond those that send direct input projections to the amygdala. Some of these back projections originating in the amygdala even reach the primary visual cortex, area V1 (Amaral and Price, 1984; Iwai and Yukie, 1987).

1.2.5 Orbitofrontal cortex–amygdala connectivity

Anatomical and connectional features of an area are strictly linked to the cognitive processes that are a function of the area. As partly discussed in the previous sub-sections of this chapter, OFC and amygdala are both hubs situated at the center of complex connectional networks. In addition, the two regions OFC and amygdala are also densely and reciprocally interconnected with one another, which again underlie their joint role in coordinating value-guided associative learning and decision-making.

As described above, in the anthropoid primate brain, OFC and amygdala are both connected with all modalities of sensory association

cortex. Additionally, visual and auditory cortices project to the same portion of the amygdala that also projects to the OFC, thereby providing OFC with an additional indirect source of sensory information (Barbas, 1995; Ghashghaei and Barbas, 2002). The densest inputs from and outputs to the amygdala are found in the caudal portions of OFC and extend to vmPFC regions. Specifically, the areas showing the highest levels of interconnections are areas 12o, 13, 14, orbital prefrontal cortex and orbital periallocortex (Ghashghaei et al., 2007; Ghashghaei and Barbas, 2002; Morecraft et al., 1992). A similar caudal-rostral gradient in connectivity strength is also present between cingulate regions on the medial wall of PFC and the amygdala (Ghashghaei et al., 2007). More specifically, the OFC receive dense pattern of inputs from several nuclei of the amygdala which include the lateral, basal, accessory basal, anterior, periamygdaloid, and medial nuclei. Interestingly, some of these inputs to OFC show a topographic organization (Cavada et al., 2000). Among the OFC regions, areas 13 and 12o are among the major recipients of such connections while the vmPFC (Brodmann area 14) and frontal polar area 10 receive a lower number of axons from such areas (Aggleton, 2000; Amaral and Price, 1984; Carmichael and Price, 1995). In the opposite direction, that is OFC projections to the amygdala, it is notable that the OFC innervates the GABAergic neurons in the intercalated masses of the amygdala (LeDoux, 2007). Cells in this part of the amygdala are part of the complex inhibitory micro-circuitries within the amygdala partly reviewed above. For instance, cells in the intercalated masses are known to inhibit the central nucleus of the amygdala (LeDoux, 2007). The functional relevance of these connections can be appreciated when considering the important role

played by the central nucleus in modulating endocrine and autonomic activity through the projection of fibers to hypothalamic nuclei, which in turn project to brain stem regions (Aggleton, 2000; Carmichael and Price, 1995; Janak and Tye, 2015; Price and Amaral, 1981). At the same time, OFC also projects directly onto the central nucleus through another path, creating in this way both direct and indirect circuits regulating neurons in this part of the amygdala (Carmichael and Price, 1995; Ghashghaei et al., 2007; Janak and Tye, 2015). Moreover, such complex connectional networks may additionally empower the OFC with indirect control over autonomic and hormonal systems. Although interesting, the discussion of termination fields in the different cortical layers of OFC and amygdala nuclei is beyond the goal of this dissertation [but see Barbas, (2007) for a detailed review].

Taking a connectional perspective, Price and colleagues (Carmichael and Price, 1995; Ongür et al., 2003; Ongür and Price, 2000) suggested that the OFC is architectonically organized in (1) a medial sensory-visceromotor and a (2) orbital sensory-associative networks and that such organization is preserved across rats, macaque monkeys and humans. Despite the classification of these two networks as separate on the basis of differences in the type of connectional inputs they receive, they are actually strongly interconnected to one another, partly sharing some of the same regions, especially in caudal OFC areas 13 and 12o. One may therefore wonder how amygdala projections contribute to such dichotomous organization and whether they also show a medio-lateral organization. Lateral OFC regions, i.e. area 12o, receives more input connections from the amygdala compared to medial OFC areas. Interestingly, the pattern was the opposite if the

authors looked at the hippocampus which instead targeted more densely regions located on the medial wall of PFC compared to lateral areas (Carmichael and Price, 1995; Ongür et al., 2003).

From a functional perspective, interactions between the OFC and amygdala have a central role in determining the way that rewards influence behavior (Baxter et al., 2000; Murray and Wise, 2010; Rudebeck et al., 2017b, 2013a; Rudebeck and Murray, 2014). The importance of OFC-amygdala interactions in value-guided behaviours was underlined by investigations of the effects of lesions of the uncinate fascicle in monkeys, a tract that connects temporal structures with PFC. Eacott and Gaffan (1992) taught monkeys to select stimuli or actions on the basis of conditional associations with different types of cues. The ability to use a visual stimulus as a conditional cue for selecting other visual stimuli or actions was impaired by a lesion targeted specifically to this tract. The ability to use prior reward information as a cue for selection was, however, unimpaired by the same lesion and may therefore depend on amygdala-frontal interactions originating from fibers not running in the uncinate fascicle. Uncinate fibers may therefore constitute one pathway that primarily conveys visual information between PFC and amygdala, whereas alternative routes may be predominantly responsible for carrying value-related information between prefrontal and temporal structures. The experiment by Eacott and Gaffan (1992) and others (Baxter et al., 2009; Browning and Gaffan, 2008; Bunge and Wallis, 2008; Bussey et al., 2002, 2001; Gutnikov et al., 1997; Passingham et al., 2000; Petrides, 2005; Rushworth et al., 2005; Wise and Murray, 2000) highlight the importance of connections between temporal regions, amygdala, and OFC in

learning arbitrary associations between stimuli, between stimuli and actions, and across temporally extended events; abilities that lie at the heart of much complex primate behaviour.

1.3 The role of the orbitofrontal cortex and amygdala in associative learning and credit assignment

A link between anatomy and function is necessary to understand complex cognitive functions like value-guided decision-making and credit assignment. Credit assignment is the process through which a subject learns to associate an outcome (i.e. food, monetary reward or error) with the choice that produced it. This process is also referred to as contingent learning and supports the capacity to learn the causal relationships between choice options and the following resulting event based on the history of their contingencies. As a result, if credit (or error) is assigned correctly to a choice, this decision will be more likely to be repeated in the future. On the contrary, if an outcome is misassigned to the correct option, it will lead to abnormal decisions in the future.

Because computations like credit assignment are likely to arise from the interactions of multiple areas within a circuit, a switch towards a framework that interpret these computations as arising from the connectivity and functional dialogue between brain structures within this network is necessary. Indeed, although a brain region may have a predominant role in encoding a specific computation, the activity in this area is modulated by a network of brain regions with which the area is interconnected. Functionally,

such connections determine which regions influence an area and which region are influenced by that area, and ultimately enable the latter one to play a key role in specific cognitive functions.

I have earlier described how the pattern of granule cells organization in OFC and its connections to and from the amygdala displays a caudal-rostral gradient. However, we might ask whether such anatomical gradients are mirrored by a similarly clear functional gradient within OFC. Kringelbach and Rolls (2004) suggested that the opposite is the case and that OFC neurons responding to certain categories of stimuli are not clustered within specific regions but are instead intermingled and spread across different areas in OFC. However, the anatomical and functional connections of these neurons with several other areas throughout the brain enable OFC to simultaneously and distinctively represent information about external events (i.e. cues in the environment or physical properties of stimuli during a decision), contingently linking them to outcomes that follows their choices through a credit assignment mechanism. These learnt contingencies are then used to update the value assigned to this events (or stimuli) on the basis of motivation and internal current state during associative learning and value-guided decision-making to guide behaviour (Murray and Izquierdo, 2007; Rudebeck and Murray, 2014; Rushworth et al., 2011).

1.3.1 The unsolved orbitofrontal dilemma: what does the orbitofrontal cortex do?

The anatomical complexity of the OFC has influenced several attempts that have been made to try to understand its core functions. In the last 30-40

years multiple and heterogeneous roles have been assigned to this portion of PFC. All agree that OFC is pivotal for associative learning and value-guided decision-making but disagree on which exact core function it encodes. I will briefly discuss some of them here but this does not aim to be an exhaustive review (see also Hunt and Hayden, 2017; Murray and Izquierdo, 2007; Murray and Rudebeck, 2018; Noonan et al., 2011; Padoa-Schioppa, 2011; Rich and Wallis, 2014; Rudebeck and Rich, 2018; Rushworth et al., 2011; Stalnaker et al., 2015; Wallis, 2007).

Initial case reports of patients with injuries to OFC (Phineas Gage's one is the most famous example) or that had undergone through frontal lobotomies led to the idea that OFC is primarily involved in response inhibition (Ferrier, 1876; Harlow, 1868). This view is still frequently endorsed today. The response inhibition theory was elaborated on the basis of the troublesome symptoms displayed by frontal patients. These symptoms ranged from sudden changes in personality, deficits in behavioural and emotional regulation as observed in social contexts, higher impulsiveness, perseveration and sometimes acquired sociopathic traits (Damasio et al., 1994; Eslinger and Damasio, 1985; Harlow, 1868; Pement, 2013). The belief of OFC involvement in response inhibition was also later supported by initial results obtained with the use of reversal learning tasks. In this type of task, subjects are required to associate a stimulus with a reward over a succession of trials. However, when a criterion level of performance is reached through the task the contingencies between stimuli and outcomes changes: the rewarded stimulus is no longer rewarded and the previously unrewarded stimulus becomes rewarded. The first investigations of the effect

of OFC lesions were consistent in highlighting deficits in the re-learning of such contingencies (Butter et al., 1970; Izquierdo and Jentsch, 2012; Jones and Mishkin, 1972). However, the deficits were shown to be specific to associations that required the linking of a stimulus to an outcome but not to associations requiring creation of an action-outcome association (Rudebeck et al., 2008). Walton and colleagues (2010) used a task in which macaque monkeys had to choose among three stimuli. Congruent with the inhibition response theory they showed that after making a lesion in the part of OFC medial to the lateral orbital sulcus and extending to the ventromedial PFC, the animals were impaired in learning the associations and use their past history to guide behaviour, especially after reversal. However, the fact that task performance was worse even on trials on which animals were not perseverating with previously made responses suggested that a failure to inhibit actions cannot account for the impairment.

In another study, an OFC lesion targeted to a similar orbital region caused animals to behave differently on trials following reward outcomes as opposed to following error outcomes again suggesting that OFC impairments are not simply the result of a difficulty in inhibiting unrewarded actions but suggesting that something about the type of outcome, positive or negative, is important in determining the appearance of these deficits (Rudebeck and Murray, 2008). Moreover, an even more recent study by Rudebeck and colleagues (2013b) suggested that impairment on visual discrimination reversal learning tasks after lesion of the same portion of OFC may only occur when they have been made by aspiration but not when they have been made by excitotoxin injection. If the visual discrimination learning task is

taken as an assay of behavioural inhibition then this suggests that the lesions that may be responsible for causing response inhibition failures are ones which compromise fibers in caudal OFC. Moreover, Wallis (2007) also reviews a number of studies reporting several additional deficits other than response inhibition after OFC lesions.

Results from single-unit recordings of neuronal activity in OFC while animals were performing a reversal learning task led some scientists to suggest that OFC was rather involved primarily in tracking and flexibly representing stimulus-outcome associations (Rolls, 1996; Thorpe et al., 1983). However, like the response inhibition theory, such an account struggles with demonstrations that deficits observed in reversal learning may be due to damage to adjacent white matter connections rather than to the OFC itself. Although affected, monkeys and human patients are indeed still able to perform reversal learning tasks even after specific OFC lesions or following its surgical removal (Baxter et al., 2009; Noonan et al., 2017, 2010).

An alternative popular view of the core function of the OFC from the 1990s is the “somatic marker hypothesis” (Bechara et al., 1996, 1994), which has strong links with the “theory of the emotions” put forward almost 100 years earlier by James (1884) and Lange (Lange and James, 1922). The “somatic marker hypothesis” primarily relies on the observation of the performance of patients with OFC lesions in a neuropsychological test called “Iowa Gambling Task”. This test requires that the subjects choose one card from four decks while trying to maximize the gain obtained from certain types of cards and minimize the loss caused by the selection of other types. Over several consecutive trials, participants should learn which decks are the most

risky ones as they are characterised by some highly rewarding cards but with the majority of them instead leading to frequent losses. Despite the fact that healthy subjects learned to correctly evaluate which decks lead to risky outcomes, OFC patients were unable to learn these associations. Moreover, before choosing from a deck composed by risky cards, OFC patients did not display the higher autonomic response, as indexed by the skin conductance response, usually shown by healthy individuals (Bechara et al., 1999, 1997). Following these results, the “somatic marker hypothesis” assigned to the OFC a role in learning and storing associations between somatic signals and the stimuli that produced them. By using these associations, OFC would then help discriminate between risky (or emotionally arousing) and non-risky choices. One limitation of this theory is that the same patients do not show similar emotional deficits in other paradigms outside the “Iowa Gambling Task”. Moreover, an additional limitation of the “somatic marker hypothesis” is highlighted by the observation that the deficits in this test are dependent on the order of the exploration of the risky contingencies (Fellows and Farah, 2005).

Recently, an alternative account of the OFC has argued that it is the region with the most central role in the encoding and signaling of economic value of potential choices (O'Doherty, 2014). Padoa-Schioppa and Assad (2006) showed that the firing rate of OFC neurons correlated with the subjective value assigned to the potential choices that might be made (“offer values”) and to the subjective value of the option that is ultimately chosen. The relationships between the value assigned to the option and the firing rate of these neurons was linear and independent from variables like the

probability or magnitude of reward in isolation, the identity, quantity or physical features of the chosen option (Padoa-Schioppa and Assad, 2008). The fascinating aspect of this theory is that it supports the hypothesis that stimuli with different properties can be compared by simply converting them into a single common neural currency within the OFC. The OFC would then use this converted abstract information for comparing the different options (Levy and Glimcher, 2012; Montague and Berns, 2002; Plassmann et al., 2007). One limitation in the theorization of this approach is its reliance on correlative measures such as neuronal firing rate or blood oxygen level dependent imaging (BOLD) measurements. Indeed, these types of studies, despite providing important information, do not allow for causal inference of the type that is intrinsic to lesion and neuromodulation experiments. Moreover, this theory perhaps over-simplistically reduces all decision making to the representation of economic value but does not entirely accounts or tackle the role of OFC in emotions or in encoding the saliency of the environment (Esber and Haselgrove, 2011; Ogawa et al., 2013). Lastly, some authors showed that value signals are also found outside OFC (Kennerley et al., 2011; Kennerley and Wallis, 2009).

Schoenbaum and Niv tried to unify the heterogeneous functions assigned to OFC within a unique framework that describe its areas as a network representing an abstract cognitive map inclusive of all the different information currently available in the task space at any given timepoint (Wilson et al., 2014). This abstract model in OFC would additionally describe transitions between consecutive task states and relationships between previously experienced task-relevant variables such as cues, current and

past stimuli, actions, outcomes and future expectations (Takahashi et al., 2013). Notable, such a model would be possible thanks to the variety of multi-modal inputs the OFC receives and would support both model-free and model-based decisions. Rather than considering OFC in isolation this account puts it in the context of a wider network of brain areas including amygdala, hippocampus and striatum (Schoenbaum et al., 2003; Stalnaker et al., 2007; Wilson et al., 2014). A potential limitation to the “cognitive map hypothesis” is that its theorization mostly arises from and is tested in studies in rodents. I have described earlier in this Chapter the lack of granule cells in the OFC of these species, features that anatomically differentiates them from the presence of granular cortex in the primate OFC. This cytoarchitectural dissimilarity may make an application of this theory to monkeys and humans problematic if granule cells in prefrontal areas play a key role in associative learning (Passingham and Wise, 2012). A second limitation of the account is the central importance that it assigns to OFC lesion-induced impairments in reversal learning tasks but, as I have discussed, reversal learning deficits do not really seem linkable to the OFC in primates (Rudebeck et al., 2013b).

An additional perspective has been put forward by Murray and Rudebeck (2018). Value-learning tasks which rely on consecutive decisions across separate stimuli require trial-by-trial updating of the estimated value assigned to each stimulus on the basis of the previously experienced outcomes. This cognitive process is required both in the context of stable and volatile contingencies, although to different degrees. This updating mechanism, however, is not solely based on external and outcome-related information. Indeed, results from reinforcer devaluation tasks in animals, in

which the value assigned to a stimulus is reduced through a satiation process, show that value estimates in OFC are modulated by the internal current, motivational, state of the subject performing the choice. In this framework, OFC is believed to functionally interact with the amygdala (Baxter et al., 2000; Fiuzat et al., 2017; Izquierdo et al., 2005) and the mediodorsal thalamic nucleus (Izquierdo and Murray, 2010) to update current value estimates. This hypothesis has recently been supported by one study in rodents where the terminals of axons projecting from the basolateral nucleus of the amygdala to OFC were silenced resulting in a deficit in using cue-based expectation of reward to drive behaviour (Lichtenberg et al., 2017).

Based on the cytoarchitectonic and immunohistochemical properties of regions in ventral PFC, Murray and Rudebeck (2018; 2014) make one further point and argue for an important functional dissociation within ventral PFC. They assign different roles to (1) the part of OFC that extends from the gyrus rectus to the lateral orbital sulcus (areas 11, 13 and 14) and to (2) the vIPFC that stretches from the lateral orbital sulcus to the portion to the principal sulcus (areas 12, or 47/12 in humans, and 45) (Rudebeck and Murray, 2014). Although both OFC and vIPFC play a central role in value guided decision-making, specifically in stimulus-outcome learning and valuation, their role slightly differs as will be described in the following paragraphs.

Selective satiation and object discrimination studies in macaque monkeys have suggested that the portion of OFC situated between the two orbital sulci, but not more medial regions on the gyrus rectus, is important for updating the value of the a choice option and then for using this updated

estimate during value comparison between choice options (Camille et al., 2011; Izquierdo et al., 2004; Kazama and Bachevalier, 2009; Rudebeck et al., 2017b; Rudebeck and Murray, 2011). Additionally, this strip of OFC could also play an important role in regulating emotional responses (Dias et al., 1996; Pujara et al., 2019; Reekie et al., 2008), particularly to aversive stimuli (Izquierdo et al., 2005). Similar specialization of medial OFC was also replicated in humans using functional magnetic resonance imaging (Howard and Kahnt, 2017). In the same study, activity within more lateral regions of the OFC correlated with estimates of the current value of the preferred stimulus (Howard and Kahnt, 2017).

Rudebeck and Murray suggest that the functions of OFC, however, differ from those of regions in the inferior prefrontal convexity, especially for what concerns the learning of stimulus-outcome contingencies (Iversen and Mishkin, 1970; Rygula et al., 2010). Animal with excitotoxic lesions of vIPFC but not OFC were impaired in tracking the option-specific likelihood of reward, thereby highlighting the role of vIPFC in contingent learning (Rudebeck et al., 2017b).

Summarizing these results, Murray and Rudebeck argue that orbital regions medial to the lateral orbital sulcus represent the current value of a stimulus and therefore its “desirability”, whereas areas lateral to the sulcus track the probability that a specific choice will lead to a positive outcome, and therefore their “availability”. This hypothesis is supported further by the differences in connectivity between these two orbital regions (Schmahmann and Pandya, 2006). Despite the detailed manner in which this account differentiates between OFC sub-regions it is less concerned with functional

specialization with vIPFC (for example between vIPFC sub-regions such as areas 12 and 45 and the parts of area 12 closest to the OFC such as 12o or 12/47o).

A last framework used to understand and describe the role of OFC in associative learning and value-guided decision-making comes from the work of Rushworth and colleagues. Interestingly, this framework may complement some of the limitations of the model suggested by Murray and Rudebeck. Consistently with Murray and Rudebeck (2018), Rushworth uses an anatomy-based approach to describe the OFC. According to Rushworth, OFC can be described as a region composed by a series of parallel areas performing separate and distinct computations with the overall goal of supporting value-guided decision-making and learning (Rushworth et al., 2011). In this account the region on the lateral bank of the lateral orbital sulcus, area 12o or 12/47o is included within the definition of OFC and sometimes referred to as IOFC. However, in the account of Rudebeck and Murray this same region is treated as a part of vIPFC. To avoid confusion I will, therefore, throughout this thesis, refer to this region as IOFC (area 12o/47o). The results from the experiments that I will report in later chapters suggest that it is this region, regardless of whether it is treated as part of OFC or vIPFC, that has a particularly important role in credit assignment.

Using a combination of lesion and neuroimaging studies in macaque monkeys and humans, Rushworth and colleagues demonstrated that the regions of OFC that are located medially to the lateral orbital sulcus encode and represent the value of choices during decision making (Boorman et al., 2009; Fouragnan et al., 2019; Papageorgiou et al., 2017). The region

includes the area adjacent to the lateral wall of the medial orbital sulcus and posterior medial orbital cortex in which Padoa-Schioppa and colleagues reported choice-value related neural activity. This region also overlaps with the portion of vmPFC (area 14) that was suggested by Strait and colleagues to perform value comparison (Strait et al., 2014). Moreover, lesion studies show that decision making is particularly impaired when choice values are close together and therefore difficult to discriminate. Noonan et al. (2010) showed that monkeys with area 14 lesions were impaired in choosing among three options when the value of the intermediate option was similar to the highest one. However, they did not show deficits in learning the value of each single option (Noonan et al., 2010). Related results have been reported in humans' lesion and neuroimaging studies (Chau et al., 2014; Jocham et al., 2016; Noonan et al., 2017).

LOFC (12/47o) regions are instead more concerned with learning the contingencies between the choice option and the outcome that follows its choice. LOFC then uses these contingencies to update the option's value that is represented in medial OFC areas for future value comparison as well as for future outcome predictions (Akaishi et al., 2016). Moreover, IOFC regions also tracks the likelihood of a choice option to lead to a positive or negative outcome in consecutive trials. LOFC does so by integrating the contingencies observed in subsequent trials into an integrated history (Chau et al., 2015; Walton et al., 2010). This is broadly consistent with Murray's and Rudebeck's view of vIPFC (Rudebeck et al., 2017b) but it emphasizes that it is not the entire vIPFC but rather the portion lateral to the orbital sulcus (Chau et al., 2015; Noonan et al., 2012; Rushworth et al., 2011).

Consistently, Chau and colleagues (2014) asked macaque monkeys to perform a value-guided learning task in which the reward rate of the choice options drifted during the experiment and demonstrated that activity in the caudal part of area 12o was associated with a win-stay/lose-switch strategy. However, the hypothesis that area 12o (area 47 in the human brain) is important for learning choice option-outcome associations still remains to be causally tested. In the absence of such a test, it is not yet possible to establish the key role of this portion of OFC in credit assignment.

Several functions in OFC and neural circuits connected with OFC may be represented in OFC at the same time or over slightly different timescales during a choice (Murray et al., 2014) or in parallel within other OFC areas or in nearby PFC networks (Rushworth et al., 2012). Moreover, differences in the volatile versus deterministic relationships of links between certain choice options and the associated outcome in the reward schedules used among the different experiments reviewed above may also determine which mechanisms will primarily drive behaviour in each context.

1.3.2 A model for value comparison in orbitofrontal cortex

A biophysiological plausible model of the neural processes that occur during decision making has been put forward by Wang (2008, 2002). The first attempt was to provide an account of perceptual decision-making mechanisms in posterior parietal cortex (Wang, 2002) and it was later extended to value-guided decision-making (Wang, 2008). This model is described as an attractor-network model in which two attractors (**Figure 7**), each represented by a distinct population of cortical pyramidal neurons,

compete through inhibitory interneurons and lateral inhibitory connections in a winner-take-all competition. Within each pool of neurons, activity is sustained through reverberating excitatory activity which is dependent on recurrent connections mediated by NMDA receptors (Wang, 2008, 2002). Similarly, a model relying on slow reverberating excitatory activity mediated by NMDA receptors was also suggested by Compte and colleagues (2000) to describe the location-tuned increase in neural activity supporting visuo-spatial working-memory in dorsolateral PFC. At the beginning of a decision, each pool of neurons representing a specific stimulus initially reinforces its own activity through reverberating and excitatory connections (Wang, 1999). As soon as the activity in one of the two pool exceeds the activity of the other pool, the network moves towards an attractor state in which the winning pool remains in a high firing rate state, which can be thought of as corresponding to a representation of the chosen option. As a result a choice is made. At the same time, the “losing” pool is inhibited through lateral inhibition mechanisms taking place across the two competing pools as well as through the influence of external GABAergic interneurons.

It should be noted that spontaneous firing of neurons, current excitatory state of the neural network, and other concomitant neural events function as potential noise and may also affect or slightly nudge these neuronal dynamics to an extent that varies as a function of the value difference between the options available. Within the framework of this biophysical model, noisy inputs such as additional un-choosable distractor options can also have an impact on exactly how the network moves to an attractor state and so it may influence exactly which choice is ultimately

taken (Boorman et al., 2011; Chau et al., 2014; Fouragnan et al., 2019; Noonan et al., 2010). The latency in choice evaluation within the attractor model as reflected by the subject's reaction times can be related to the time required for a pool of neurons to reach a particular threshold so that the network enters an attractor state (Webb et al., 2011).

The attractor network model suggested by Wang for decision-making primarily relies on the synaptic release of glutamate and GABA, which in turn control this winner-take-all competition. On the one hand, glutamate drives excitation of the neurons composing one of the pools and leads to the recurrent synaptic reverberation. This molecule influences this process by primarily acting on the kinetically slower NMDA receptors and, to a lesser extent, AMPA receptors (Wang, 1999; Wong and Wang, 2006). Indeed, NMDA receptors are characterized by a longer time-constant of decay for synaptic-current equal to 100 milliseconds (against the faster time-constant for AMPA receptors of 5 milliseconds) which makes this receptor type the most likely candidate to drive such a mechanism. On the other hand, GABA release controls the reciprocal mutual inhibitory activity through its ionotropic GABA_A receptor, which primarily influences the mutual-inhibition mechanisms supporting the selection of the final chosen option (Brunel and Wang, 2001).




Figure 7. Attractor network model in probabilistic decision-making. *A and B represents the pool of pyramidal neurons encoding for two separate stimuli inputs (I_A and I_B , respectively). Neurons in each pool reinforce and maintain their firing rate by recurrent slow synaptic excitation within the population, which is mainly mediated by NMDA receptors. Lateral inhibition between the pools of neurons along with the action of external GABAergic interneurons mediate the mutual inhibition among choice in the winner-take-all competition (Adapted from Hunt and Hayden, 2017).*

The importance of excitatory and inhibitory activity in value-guided decision making was also shown by Jocham and colleagues (2012). The authors combined magnetic resonance spectroscopy and functional magnetic resonance imaging to investigate the relationships between individual differences in baseline levels of glutamate and GABA within OFC and vmPFC regions and performance in value-guided decision-making. The authors found that choice accuracy varied as a function of glutamate and GABA baseline levels, with best performances correlated to low glutamate but high GABA concentrations.

Wang's model has also been used to try to understand the neural mechanism underlying value-guided decision-making processes in neurophysiological investigations. In monkeys, a mechanism for value comparison with features resembling those seen in Wang's artificial networks has been reported in medial OFC and vmPFC by Strait and colleagues (2014). Strait and colleagues showed that neural activity in this region temporally shifted from encoding (1) both reward and the likelihood of reward associated with one option at the beginning of a trial to (2) signaling only the chosen option. Moreover, Hunt and colleagues (2012) demonstrated that mutual-inhibition does occur in human OFC and vmPFC areas and underpins value-comparison during value-guided choices. They tested this hypothesis by analyzing the dynamics of a cortical attractor-network model, similar to Wang's model, in human participants and comparing the estimates of the model with neural data recorded with magnetoencephalography. The data revealed a close similarity between the model and magnetoencephalography activity within OFC and vmPFC as well as parietal cortex. In addition, Hunt and colleagues also found temporal changes in the activity signal within these regions that indicated that they were encoding the value difference between the options being considered during decision making.

A potential hypothesis is that the optimal functioning of the attractor network model in medial OFC and vmPFC may be influenced by contingent learning mechanisms encoded in the more lateral orbital area 12o. For instance, as choice option-outcome contingencies are learned throughout a task, the formation of these association may modulate the excitatory/inhibitory network mechanisms underlying the attractor network

model. According to this hypothesis, a stronger association of a specific choice option with reward may reinforce the excitatory firing of neurons within the pool representing it, thereby resulting in a higher likelihood of this pool to attract the network and win the competition with pools of neurons representing alternative choice options.

1.3.3 Amygdala

The amygdala is a subcortical structure located on the medial bank of the temporal lobe and has been implicated in emotion and social behaviour (Whalen et al., 1998). Moreover, the amygdala has also been suggested to play a role in value-guided learning and decision making alongside OFC regions (Morrison and Salzman, 2010; Saez et al., 2015).

The role of amygdala in behaviour and cognition emerged already in the 19th century when Brown and Sharpey-Schafer (1888) performed a bilateral ablation of the rostral temporal lobe on macaque monkeys and reported remarkable behavioural changes: “ [...] *Prior to the operations he was very wild and even fierce, assaulting any person who teased or tried to handle him. Now he voluntarily approaches all persons indifferently, allows himself to be handled, or even to be teased or slapped, without making any attempt at retaliation or endeavouring to escape. His memory and intelligence seem deficient. [...] Every object with which he comes in contact, even those with which he was previously most familiar, appears strange and is investigated with curiosity. [...] He appears no longer to discriminate between the different kinds of food; e.g., he no longer picks out the currants from a dish of food, but devours everything just as it happens to come [...]*”.

These post-surgical observations highlighted several behavioural deficits in the behavior of the monkey that were later replicated by Klüver and Bucy (1937). A more precise investigation in the role of the amygdala was performed by Weiskrantz (1956) who ablated this structure (along with the temporal pole) in the macaque brain and observed a faster extinction of conditioned responses (acquired before the ablation) but slower learning in other context along with a range of other deficits.

These initial results, along with those from subsequent amygdalectomies performed in macaques, began to direct the attention of researchers towards a role for the amygdala in decision-making and learning (Barrett, 1969; Schwartzbaum and Poulos, 1965). A more specific role for the amygdala in creating stimulus-outcome associations was suggested by Mishkin's work (Jones and Mishkin, 1972; Spiegler and Mishkin, 1981). Interestingly, Aggleton and Passingham (1981) suggested a link between impairment in object discrimination tasks and hypoemotionality after observing the behaviour of monkeys that received partial lesions of the amygdala. The role of the amygdala in modulating the interaction between emotional processes and decision-making was later shown to be present also in the human brain (De Martino et al., 2006). A final refinement in the understanding of the cognitive roles of the amygdala came from Murray's excitotoxic lesion studies in the 1990s (Baxter and Murray, 2002; Murray, 1992). In comparison to previous aspiration lesions, excitotoxic lesions preserved fibers passing close to, or through, the amygdala. The observation that behavioural deficits associated with the amygdala like learning the relationship between choosing a stimulus and its resulting outcome, memory

formation and visual-discrimination learning were no longer present once the lesion was made with an excitotoxic agent (Baxter and Murray, 2002) suggested that some of the previously observed symptoms which followed amygdectomy may have been due to disruption of white matter tracts running near the amygdala, but not even to or from the amygdala, rather than by damage to the amygdala itself.

Later investigations highlighted the importance of the amygdala in updating the value representation of a potential choice in relation to the animal's current motivational state. This hypothesis was supported by the observation that lesion or pharmacological inhibition of this structure resulted in abolition of satiety-induced devaluation of choices (Izquierdo and Murray, 2007; Johnson et al., 2009; Wellman et al., 2005).

Other reward-related choice valuation effects are, however, influenced by the type of outcome that is received after a choice, with positive feedback typically have a bigger impact on adjustment of the behaviour in the following trials of animals with amygdala lesions (Rudebeck and Murray, 2008). Moreover, extinction of previously learned responses were facilitated by lesion of the amygdala (Izquierdo et al., 2005). Deficits in value-guided behaviour following amygdala lesion have also been shown to be present both when stochastic or deterministic reward schedules are used (Costa et al., 2016). The authors demonstrate the contribution of the amygdala to value-guided decision-making by showing that after excitotoxic lesion of this subcortical structure, the animals presented impairments in choice accuracy and in learning from receiving a reward for making a particular choice.

All together these data led Murray and Wise (2010) to hypothesize that the amygdala may influence value-guided learning and decision-making by weighing the current value of an option based on the affective nature of previous outcomes received. Removal of the influence of this region, would therefore inhibit this mechanism and results in a quicker switch strategy in following choices.

An additional elucidation on the complex role of the amygdala comes from a functional magnetic resonance imaging study in macaque monkeys. Animals were required to perform a three-armed reversal learning task in a volatile setting. The best behavioural strategy, in this context, is to stay with an option that led to a reward in the previous trial or that was strongly associated with reward over a recent set of trials and switch to another option if the chosen one did not lead to or was unlikely to give a reward. Activity in the amygdala was predictive only of switching behaviour after negative outcomes (Chau et al., 2015). Such a link between adjustment in behaviour and negative feedback was supported also by a more recent study showing that animals who underwent excitotoxic lesions of the amygdala were affected in their capacity to adjust behaviour after error trials (Rudebeck et al., 2017a).

Above I have extensively described the general agreement that the role of amygdala is in value-guided decision making and learning as well as emotional behaviour. However, the majority of those studies focused on looking at the role of the amygdala in the immediate present, that is its role in the same trial or in very recent ones. This does not elucidate whether the amygdala may have an influence on planning strategies and on behaviours

that may take several trials before they lead to a reward. A relatively recent study (Hernádi et al., 2015) demonstrated that activity in some population of neurons within the amygdala encode goal-oriented behaviour over a longer timeframe. The authors observed that monkeys performing a multi-step choice task were willing to delay the receipt of reward if this strategy led to an higher amount to be received later on. At the neuronal level, this behaviour was associated with changes in prospective activity in two different population of neurons tracking respectively (1) the subjective value of the current plan to be followed to eventually obtain the higher amount of reward and (2) how long (number of trials) this plan would take before the reward could be received. It is worth mentioning that the behavioural planning required in this task was oriented towards obtaining a final reward. These results may complement well the suggestion that one of the amygdala's function is the representation of the current state of the animal or the value of a choice option suggested by Murray and Rudebeck (Murray and Izquierdo, 2007; Murray and Rudebeck, 2018). According to Hernadi and colleagues (2015), the amygdala would also internally generate prospective goal-oriented behaviours and track their value and execution. These results, together, suggest an additional potential role of amygdala in value-guided temporal discounting that may interact with OFC activity to guide decision-making and learning over a long series of trials.

1.3.4 Amygdala-orbitofrontal cortex interactions

Until now I have primarily discussed the separate contribution of OFC and amygdala to value-guided decision-making and learning. However, in the

previous part of the chapter I have discussed the strong and reciprocal interconnections between these complex structures and argued that anatomical connections may be linked to functional properties within a circuit. I will now attempt to briefly describe the functional interactions of OFC and amygdala and their relevance for supporting credit assignment.

OFC and amygdala were initially viewed as separate structures that worked either independently from one another or hierarchically with the latter unidirectionally influencing the former (Aggleton and Passingham, 1981; Butter et al., 1970; Izquierdo et al., 2005). This modular view of brain function partly resembled in the highly modular accounts of visual cortex anatomy and function. However, the application of this view also had several limitations, especially in its capacity to explain why similar behavioural deficits resulted from lesions to either structures. Additional uncertainty on the role of these two structures and their functional interplay was raised by the hypothesis that emotional processes and reward-related mechanisms, both affected by OFC and amygdala lesions, could be interpreted within the same framework and simply as the result of stimulus-reward associations (Rolls, 1999).

An important experiment tried to better understand the role of OFC and amygdala by investigating the deficits following lesions applied to these structures in isolation and comparing them with the effects of their interhemispheric cross-disconnection (Baxter et al., 2000). This surgical procedure aims to investigate the joint functional role of two areas which are part of a circuit in supporting specific cognitive functions and consist in lesioning (or entirely removing) one of the two structure in one hemisphere and the other in the opposite hemisphere. Additional commissurotomy is

performed to prevent interhemispheric exchange of information among the still intact structures. By using a devaluation procedure, the authors showed post-disconnection deficits in updating the value of reward outcomes, similarly to bilateral lesions of either amygdala or OFC (**Figure 8a**) (Baxter et al., 2009, 2000; Izquierdo and Murray, 2007). These impairments were however absent in control animals or when other PFC areas were lesioned (**Figure 8b**) (Baxter et al., 2008).




Figure 8. Effects of bilateral or disconnection lesions of OFC and amygdala. Lesion of amygdala (red), OFC (blue) or disconnection of these two areas affects updating of value representation in macaque monkeys. These effects were specific for these type of lesions as similar deficits were not observed when other control areas (purple and grey) were lesioned (Reproduced from Murray and Wise, 2010).

Interactions between OFC and amygdala in value updating and representation for goal-oriented behaviour were also reported in another

recent study (Fiuzat et al., 2017) although rodent studies have highlighted the primary role of amygdala projections to OFC (Lichtenberg et al., 2017). A similar influence of amygdala on OFC activity was replicated in macaque monkeys where excitotoxic lesions of the first structure caused a reduction in the number of OFC neurons encoding expected and received reward (Rudebeck et al., 2017a).

Although supporting a role of the amygdala in representing acquired choice option-outcome association, the observation that neither value representation in OFC nor outcome-based learning were entirely abolished after amygdala lesion suggests that OFC also plays a key role in these computations (Rudebeck et al., 2017a, 2013a). Little, however, is known of the impact of focal disruption of activity in medial and lateral OFC areas on amygdala activity during credit assignment. Chau and colleagues (2015) showed that IOFC activity is associated with both Win-Stay and Lose-Shift strategies during probabilistic learning. However, amygdala activity only clearly encodes Lose-Switch behaviours. LOFC may therefore play a primary role in learning contingencies between choice options and outcomes and integrate them into a single representation that is then used by the amygdala to represent option-specific value. From an evolutionary perspective, the granular properties of OFC and its widespread connectivity, including the reciprocal connections with the amygdala, may empower the primate brain with more refined credit assignment mechanisms that enables primates to learn the value associated with an option in just a single trial (Noonan et al., 2017, 2010; Passingham and Wise, 2012; Walton et al., 2010).

1.4 Methodological approaches for investigating anatomy and function

Throughout this thesis I will report the results of investigations using several imaging, computational, and neuromodulatory techniques. The choice to switch among them from study to study was driven by the rationale of finding the best tool to investigate each research question that I wanted to tackle. The need to make causal inferences also required the implementation of novel methods. I will here give a brief overview of the approaches that I am going to present over the following chapters.

1.4.1 Diffusion Magnetic Resonance Imaging and Tractography

"[...] Just as Nature deserves praise and admiration for making machines so small, so too the physician who observes them to the best of his ability is worthy of praise, not blame, for he must also correct and repair these machines as well as he can every time they get out of order [...]."

(Marcello Malpighi, 1697)

The central nervous system is composed by two main types of tissue, the grey matter and white matter. Although their separation was initially identified by Andrea Vesalius already in the 16th century, the definition of white matter as a complex tissue formed by fiber bundles with origins and

terminations in different grey matter regions took place only one century later thanks to Marcello Malpighi and Nicolaus Steno (Clarke and O'Malley, 1968; Pandya et al., 2015). The first known nerve system to be investigated in the brain was the complex of fibers running in the occipital lobe adjacent to what nowadays we describe as the primary visual cortex. In the same period, Johann Christian Reil and Friedrich Burdach used blunt dissection approaches in the human brain to further discriminate specific cortico-cortical fiber bundles that were later defined by Gall and Spurzheim as “association” fiber systems (additionally divided into short or long depending on the length of their course). These fiber systems correspond today to a well-known set of tracts: uncinate fascicle, cingulum bundle, arcuate fascicle, superior, and inferior longitudinal fasciculi. In the 17th century, Thomas Willis attempted an interpretation of the role of these fiber bundles and suggested they had a role in generating the spirit of the individual (Pandya et al., 2015; Schmahmann and Pandya, 2006). Some of the later studies that today stand as cornerstones of what we know about the white matter and the role these fibers have on the functions of our brain were carried out by Broca (1861a, 1861b), Wernicke (1874), Onufrowicz (1887), Déjerine and Dejerine-Klumpke (1895), Klingler and Gloor (1960), and Nauta and Gigax (1954).

More recently, anterograde or retrograde “tract-tracing”, allowed to identify in much finer details and high fidelity the fiber bundles connecting several brain areas. Tract-tracing is still today the “gold standard” technique for the investigation of white matter architecture and the reconstruction of axonal projections in the brain as this technique enables a precise tracking and visualization of both origins and terminations of fibers. However,

because it requires the sacrifice of the subject under study, this approach can be performed solely in animals or, with great difficulty in post-mortem human tissue (Petrides and Pandya, 1984; Schmahmann and Pandya, 2006). Moreover, the invasiveness and technical characteristics of this approach make it expensive, time-demanding, and laborious.

A magnetic resonance imaging (MRI) approach has been developed as an alternative method for estimating white matter connections (Basser et al., 1994; Basser et al., 2002). The development of MRI sequences able to non-invasively detect water molecule displacement across the brain opened the road towards the widespread investigation of human connectivity and inter-species comparisons. Indeed, these sequences could be used to investigate human white matter neuroanatomy, as a result of the rate of hindrance they encounter during a set time-interval (Johansen-Berg and Behrens, 2006; Jones, 2009).

Diffusion MRI sequences exploit basic biophysical properties of the body. The diffusion of water in the brain is indeed constrained along the fiber bundles composing the white matter; the diffusion is therefore described as anisotropic because it is larger in one direction than in other. By contrast, in other brain tissue compartments, the grey matter and cerebrospinal fluid, the dispersion of water molecules is isotropic, that is it is equal in any direction and regardless of the diffusion MRI force applied by the sequence pulse. The degree of diffusion can therefore be indexed by a fractional anisotropy estimate at a voxel-level, and recorded simultaneously throughout the entire brain at a relatively low cost and short time. As described above, fractional anisotropy is higher along axonal projections as a consequence of the

preference of water to move along fibers rather than across. By recording these biophysical properties of brain molecules, diffusion MRI enables us to acquire measures on both the organisation and integrity of white matter (Scholz et al., 2009; Takahashi et al., 2002). Additionally, the diffusional properties of water molecules can also be exploited to reconstruct the anatomy and course of fiber bundles. This is possible thanks to the development of mathematical algorithms like tractography, which can be used to reconstruct white matter bundles in the brain without the necessity of any invasive intervention or pharmacological injection (Jbabdi and Johansen-Berg, 2011; Johansen-Berg and Rushworth, 2009; Schmahmann et al., 2007).

Diffusion MRI and tractographic algorithms, therefore, provide us with novel approaches to investigate anatomical features of white matter bundles simply based on the acquisition of imaging data from the subject of interest, thereby providing a common tool that can be used for both within- and between-species investigations (Croxson et al., 2005; Leh et al., 2007; Noonan et al., 2018, 2014). Comparisons of tract-tracing with diffusion MRI tractography have been performed in the macaque monkey brain, repeatedly proving the reliability and reproducibility of the imaging approach (Croxson et al., 2005; Jbabdi et al., 2013; Rushworth et al., 2009). Despite these positive results, when trying to understand similarities and differences among brains of different species, i.e. macaque monkey and human brains, it is still highly advisable to use the same method in both species as some anatomical differences (Thiebaut de Schotten et al., 2012) might be the consequence of using different methods (Mars et al., 2014; Thiebaut de Schotten et al., 2019)

or tractography-related artefacts (Maier-Hein et al., 2017) rather than real connections.

When the primary research question is to compare neuroanatomical structures between-species, an additional important precaution that must be considered is the similarity of the techniques used across species. To achieve this similarity, if macaque data are to be compared with *in-vivo* human data then it may be important to acquire the macaque data at a higher spatial resolution than the human data. This is necessary because of the overall small size of the macaque brain. However, higher spatial resolution can come at the cost of lower signal-to-noise, which can be compensated by longer acquisition times (~24-48 hours) although this is only possible if the *ex-vivo* brain is studied (Miller et al., 2011). Therefore, I have compared data from *in-vivo* human and *ex-vivo* monkey diffusion data sets.

A particular approach, based on tractographic algorithms, tracks fiber bundles by exploiting the fractional anisotropy signal measured by diffusion MRI and then extracting relevant information from the shape of the diffusion probability distribution, namely the principle directions of water displacement, either by eigenvector decomposition or Bayesian probability estimation (Behrens et al., 2007; Johansen-Berg and Rushworth, 2009). This technique is, therefore, useful for performing non-invasive and *in-vivo* tract reconstruction in several primate species (Catani et al., 2002; Mars et al., 2018b; Thiebaut de Schotten et al., 2019). In particular, two different tractographic approaches are currently used: deterministic (Aldroubi and Basser, 1999; Basser et al., 2000) and probabilistic (Behrens et al., 2007) tractography. Both approaches rely on the estimation of molecules' preferred

displacement direction (Behrens and Jbabdi, 2009; Catani et al., 2002). However, in the estimation of the anatomy of the tract, probabilistic tractography takes into consideration some limitations of the deterministic approach, including the level of uncertainty in tracking along consecutive voxels (Behrens et al., 2007). Practically speaking, deterministic tractography provides a thresholded measure in which a voxel is assigned to a tract or not based on a binary principle. On the contrary, probabilistic tractography assigns to each voxel a likelihood estimate of its belonging to a tract. In addition, uncertainty estimations which are part of probabilistic tractography also account for the limitation that the further a tractography streamline is from its original starting point (its seed mask or seed voxel), the less accurate the estimation of its distribution (or of the tract the streamline is part of) becomes. The remarkable importance of considering such uncertainty measures particularly emerges in those regions where the uncertainty of a voxel's principal diffusion direction is particularly high and therefore more likely to produce results characterized by type I or type II error.

One feature of probabilistic tractography is that it provides a grounded basis for meaningful statistical comparisons. It can, therefore, be used to investigate the anatomical pathways in different species.

1.4.2 Functional Magnetic Resonance Imaging

In the previous section I have focused on describing techniques for studying white matter but have put to one side the tissue that is put in communication by the fiber bundles composing the white matter: the grey matter. Almost 30 years ago the journal *Science* published a cover image showing a portion on

the human brain with red-yellow blobs spreading throughout the visual cortex (**Figure 9**). Among a list of articles entitled “*Could the Sun Be Warming the Climate?*” (Kerr, 1991) and “*AIDS Vaccine Meeting: International Trial Soon*” (Cohen, 1991), the article referred to the cover image of the issue titled “*Functional mapping of the human visual cortex by magnetic resonance imaging*”. In this study, Belliveau and colleagues (1991) reported the results of an MRI investigation performed at the Massachusetts General Hospital combining an echo planar imaging (EPI) pulse sequence with simultaneous injection of Gadolinium to visualize pre- and post-stimulation differences in cerebral blood volume in the brain. Although, the EPI sequence had been available since the late ‘70s (Mansfield, 1977) and MRI was already in use for clinical investigations since the ‘80s, this study started the field of fMRI investigations of brain activity *in-vivo*. This novel approach relied on the different magnetic susceptibility properties of oxy- and deoxyhaemoglobin. This property of haemoglobin had been described at the beginning of the century by Pauling and Coryell (1936) and later exploited by Brindle et al. (1979), Thulborn et al. (1982) and other research groups (see Bandettini, 2012) to hypothesise that it could be used in the neuroimaging field to measure changes in an endogenous susceptibility contrast-based signal called BOLD contrast. Using this principle, in-vivo BOLD contrast changes were then shown in animals and humans using MRI (Kwong et al., 1992; Ogawa et al., 1990a, 1990b; Ogawa and Lee, 1990; Turner et al., 1991) (for a full history of the development of MRI see Bandettini, 2012; Ogawa, 2012).



Figure 9. *The beginning of functional brain mapping using MRI. Cover image of the November 1, 1991 issue of Science (Belliveau et al., 1991) obtained by combining gadolinium injection and echo planar imaging (EPI) to map neural activity during visual stimulation.*

The advent of BOLD-dependent functional MRI (fMRI) opened the doors to the advent of the non-invasive investigation of brain activity. The activity-dependent changes in magnetic and paramagnetic properties of haemoglobin in brain blood vessels within a brain area during a task (**Figure 10**) have been suggested to reflect the synchronized extracellular local field potentials (sum of positive and negative post-synaptic potentials) of a

population of neurons within this area, including subthreshold post-synaptic currents, thereby primarily indexing input signal to the area rather than firing related information (Logothetis et al., 2001). The indirect nature of the BOLD response within a brain region, its dependence on local slow fluctuations in blood magnetization as well as the sampling rate delay in MRI cause the temporal resolution of this signal to be relatively slow, often defined as sluggish. In response to the presentation of a stimulus (i.e. a flickering light or an auditory tone), the BOLD signal is usually represented by a Hemodynamic Response Function (HRF) that peaks at about 5-6 seconds after stimulus onset in the human brain (Bijsterbosch et al., 2017; Buckner et al., 1996) and at about 3 seconds in the macaque monkey brain (Chau et al., 2015). Nowadays, BOLD signal is used as an arbitrary index for stimulation-induced activity changes, even though its full physiological mechanisms at the cellular and synaptic levels are still not entirely understood. Despite such limitations, BOLD imaging has enabled non-invasive, *in-vivo* investigation of activity throughout the whole brain.




Figure 10. Physiology of the BOLD response signal in the human brain. Representation of the physiological mechanisms and temporal dynamics of the BOLD signal in the human brain. The haemodynamic response function (HRF) is characterized by an initial dip (inconsistent and usually detected only at high magnetic fields) initially following the stimulus presentation. It is followed by the dominant HRF peak which occurs in humans at about 5-6 seconds, which corresponds to the increased regional blood flow and to the oxyhemoglobin/deoxyhemoglobin ratio detected by the BOLD signal. Finally, the post-stimulus undershoot takes place in the refractory phase at 13-15 seconds. In fMRI, the timecourse recorded from each voxel corresponds to the timing of neural activity after convolution by the HRF (Reproduced from Bijsterbosch et al., 2017).

The possibility of quickly and simultaneously measuring activity from the whole brain using fMRI while subjects were performing a behavioural task enabled investigation of the correlational relationships between neural activity and aspects of behaviour. The time-varying fluctuations in BOLD response induced by the presented stimuli or the task being performed can

be modelled using a General Linear Model (GLM). The application of a GLM to the fMRI data allow to test for the relationships between the imaged brain activity and our statistical hypothesis about the task at the single subject level (first-level GLM) or at the group level (second-level or higher levels GLMs). The fitting of a GLM to the recorded brain activity aims to give a description of the timeseries acquired from each single voxel according to a weighted sum of regressors (EVs). The weight of each EV estimate is directly calculated from the data. By fitting a GLM to the fMRI data it is possible, in this way, to extract effect size estimates associated with each regressor of interest according to our predicted statistical hypotheses as well as the residual noise (estimates of unexplained BOLD response).

Task-based fMRI designs are usually built to maximize their sensitivity to specific statistical hypothesis and, therefore, to the detection of the neural activity underlying them. Several types of designs exist, depending on the nature of the experiment: block, event-related, factorial, parametric designs. By assuming linear scaling between them, parametric designs test for patterns of BOLD response that underlie specific parametric estimates of behaviour, i.e. value comparison processes or credit assignment-related computations as reported in *Chapter 5* (Jenkinson and Chappell, 2017).

Computational models are often used to infer estimates of hidden states that must have been computed by cognitive processes. These estimates of the hidden states, rather than the direct observations of occurrences of behavioural events can then be related to fMRI data to test for their neural activity correlates. Computational models derived, for example, from reinforcement learning theory (Rescorla and Wagner, 1972;

Sutton and Barto, 1998), can be used to describe cognitive states that can then be related to neural activity patterns (or firing patterns) during associative learning and value-guided decision-making (Behrens et al., 2009; Kolling et al., 2012; Schultz et al., 1997). An example of simple Rescorla-Wagner model is the prediction-error learning model (Rescorla and Wagner, 1972) according to which the value of a stimulus is learned and updated in a trialwise manner as a function of the mismatch between the previously predicted value of the stimulus and the current outcome associated with its choice:

$$V_{s,t} = V_{s,t-1} + a(r_{t-1} - V_{s,t-1})$$

In this simple Rescorla-Wagner model, $V_{s,t}$ is the value that the subject assigns to stimulus S at trial t and reflect the expectation or reward linked to that stimulus. R_{t-1} corresponds to the value of the reward received on the previous trial ($t - 1$) as a consequence of choosing the stimulus S . In this model, the prediction error, that is the difference between the received reward for the choice of stimulus S and the expected reward for making that choice ($V_{s,t-1}$) determines the updating of the value of S . The weight of the prediction error is further scaled by the learning rate estimate a .

The complexity of the BOLD signal, its temporal properties and the limitations in the accuracy of its spatial sampling led to the development of several processing approaches aimed to improve its signal-to-noise ratio and therefore the efficiency of the fitted GLMs in extracting reliable parameter estimates for behaviour. From some of these approaches, a new application of fMRI to map brain activity was implemented. This novel approach relied on

the acquisition of BOLD signal from the whole brain while subject laid in the scanner without being required to do any task or action. This method is now better known as “resting-state” fMRI. Although the subject is not performing any task and no external input is applied, during rest the brain is spontaneously active and neural activity in each area is temporally synchronized with activity in other brain regions with a higher rate of activity coupling (temporal correlation among low frequency BOLD fluctuations) with regions which also share stronger anatomical connections. This spontaneous activity at rest is suggested to represent physiological mechanisms like (1) slow cortical potentials, independently of the current subject’s arousal state, and (2) gamma-frequency activity during wakefulness and rapid-eye-movement sleep but not in the slow-wave sleep state. Moreover, the synchronization of these spontaneous BOLD fluctuations is often organized into dissociable large-scale functional networks (He et al., 2008). Measures of activity coupling between different areas are often defined as “functional connectivity” (Biswal et al., 1995; Biswal, 2012). Since their first introduction, the number of approaches taken to analyze resting-state fMRI data has exploded, with applications ranging from cognitive neuroscience to pre-surgery planning. Some of the methodological techniques usually used to examine these data include seed based (ROI-based) approaches, independent component analysis (ICA), methods derived from graph theory, clustering techniques, neural networks and pattern classification approaches. Moreover, analysis can be carried out both on volumetric and surface spaces depending on the primary goal of the study (Bijsterbosch et al., 2017; Lee et al., 2013).

As previously mentioned, activity in different groups of brain structures during rest are primarily organized in specific functional networks, often referred to as resting-state networks and among which probably the most well known and physiologically stable is the default-mode network (Fox et al., 2005). Thus, resting-state fMRI data can be used to investigate the functional architecture of the brain (Glasser et al., 2016), its alterations in psychiatric and neurological conditions (Li et al., 2013; Müller et al., 2011) and its potential correlations with behavioural and environmental variables (Noonan et al., 2014; Sallet et al., 2011).

1.4.3 Non-Invasive Brain Stimulation

The understanding of the functional role of a brain area can be established only through causal inference, that is by manipulating the activity within the area of interest and observing associated behavioural or cognitive changes. Electrophysiological and fMRI methods are correlative measures only. Lesion or neuromodulation approaches, on the other side, can directly disrupt or alter neural activity and provide an optimal tool for causally establishing the role played by regions or networks under investigation. As reviewed above, however, experimental lesion methods are not ethically viable in humans although sometimes it is possible to examine the behavioural impact of natural occurring lesions. This approach is however often limited by the lack of focality of the brain insult.

Minimally-invasive or non-invasive brain stimulation (NIBS) approaches (**Figure 11**) such as transcranial magnetic stimulation (TMS) or transcranial electric stimulation (tES) can be applied in healthy subjects.

They alter neural activity, making it possible to ascertain the causal role of the activity, but their impact is transient and reversible. TMS, modulates neural activity by transmitting a changing magnetic field through the skull which in turn induces electric currents at the target brain area through electromagnetic induction. TES, which comprises a range of different protocols including perhaps the most known transcranial direct current stimulation, relies on the passage of weak electrical currents between two electrodes applied to the skull (Polanía et al., 2018). However, the spatial resolution of TMS and especially tES is limited and none of them can be used to target subcortical areas or cortical regions located deep into the brain. High-intensity focused ultrasound stimulation (HIFU) is another approach well known for causing thermo-induced ablation (Fry et al., 1955; Fry and Fry, 1960) in very narrow strips of cortex when used at high frequencies in patients with certain pathological conditions (for a review see Quadri et al., 2018). Similarly, to deep brain stimulation (DBS), HIFU constitutes permanent invasiveness, making DBS and HIFU not viable approaches for brain mapping in the healthy brain.




Figure 11. Brain stimulation approaches. Representation of invasive and minimally-invasive (green) brain stimulation techniques used in cognitive

neuroscience. The different approaches are reported according to their spatial (horizontal axis) or temporal resolution (vertical axis). Single-pulse TMS = single-pulse transcranial magnetic stimulation; rTMS = repetitive transcranial magnetic stimulation; tDCS = transcranial direct current stimulation; tRNS = transcranial random noise stimulation; tACS = transcranial alternating current stimulation; tFUS (referred to as TUS throughout the main text) = low-intensity transcranial focused ultrasound (Reproduced from Polanía et al., 2018).

In the 1950s, Fry and colleagues (1958) suggested that the transcranial transmission of a certain dosage of acoustic energy to the lateral geniculate nucleus had the potential to reversibly and temporarily suppress activity in the visual system without causing any tissue damage. This intuition opened the way to a potentially novel non-invasive brain stimulation approach which exploits the strong interactions between ultrasound waves and biological tissues (Gavrilov et al., 1996; ter Haar, 2007) to transcranially, reversibly, and minimal-invasively manipulate activity in a focal target independently from its depth in the brain.

Inspired by the hypothesis of using low-intensity, low-frequency focused transcranial ultrasound stimulation (TUS) to modulate neural activity, Tyler et al. (2008) ran a series of experiments that showed the capacity of TUS to modulate electrical activity in neurons through the interactions between the transmitted acoustic energy and voltage-gated ion channels sensitive to sodium and calcium. Soon after, another couple of papers showed behavioural effects of TUS in healthy humans (Legon et al., 2014),

non-human primates (Deffieux et al., 2013) and patients affected by consciousness disorders (Monti et al., 2016).

Intriguingly, some of these articles as well as other conducted in lower species (Dallapiazza et al., 2018; Yoo et al., 2011) suggested that the temporal effects of the TUS extended beyond the time at which the ultrasound pulse was transmitted, thereby highlighting the potential involvement of plastic neural changes which however dissipate after a certain delay the precise length of which length is dependent on the stimulation protocol used. Subsequent clinical application of low-frequency TUS expanded further, for instance, to the use of this approach to transiently open the blood brain barrier for gene-therapy or drug delivery (Burgess et al., 2018; Munoz et al., 2018; Sierra et al., 2017; Wang et al., 2015).

The size of the TUS effect is dependent on the parameters of the TUS protocol being used: pulse repetition frequency, burst duty cycle, inter-stimulus interval, total length of the stimulation (equal to the length of the experiment), waveform of the pulse. The different interactions of the TUS parameters affect the intensity and the spatial specificity of the stimulation as well as the amount of local pressure generated by the delivered acoustic energy (Blackmore et al., 2019). Current research is trying to elucidate the underlying physiological mechanisms that cause this reversible behavioural and temporal effect. As summarized in **Figure 12**, a number of cellular level mechanisms of effect are currently under investigation (Blackmore et al., 2019; Fomenko et al., 2018; Jerusalem et al., 2019; Kubanek, 2018).




Figure 12. Potential mechanisms of TUS. (a) Membrane conformational state; (b) Thermodynamic membrane waves; (c) Direct flexoelectricity; (d) Ion channel opening; (e) Thermal modulation; (f) Sonoporation; (g) Microtubule resonance; (h) Synaptic vesicle modulation (Reproduced from Jerusalem et al., 2019).

One possibility is that ultrasound alters the conformational state of the neuronal membrane and its components (i.e. phospholipid and proteins). By applying pressure on the membrane, the acoustic energy can therefore modify its mechanical properties, thereby affecting its electric charge (**Figure 12a**). A second mechanism through which ultrasound may exert its effect is by interfering with or eliciting electric waves on the neuronal membrane (**Figure 12b**). The pressure applied by TUS may also cause direct flexoelectricity, that is to say that it may deform the phospholipid constituting the cellular membrane thereby modulating the transmembrane potential and thereby generate electric currents (**Figure 12c**) (Chen et al., 2019; Jerusalem et al., 2019; Prieto et al., 2013; Tyler, 2011). A popular hypothesis

is that TUS sound waves interact with mechanosensitive voltage-gated ion channels (i.e. potassium proteins) embedded within the cellular neuronal membrane, even though similar effects have been shown also for non-mechanosensitive channels (i.e. sodium and calcium proteins). Such interaction results in changes in the influx through the ion channels and consequent polarization of the membrane. This hypothesis is consistent with the results of several studies (Kubaneck et al., 2018, 2016a; Moroni et al., 2018; Tyler et al., 2008) (**Figure 12d**). Other potential mechanisms try instead to explain the action of TUS through non-mechanical action. An intuitive one, linked to the use of HIFU to perform ablation in cerebral tissue by heating it, is the thermal hypothesis of TUS, according to which prolonged stimulation or higher intensities may cause a small, not permanent, thermal rise affecting local cellular dynamics (**Figure 12e**). Sonoporation as a consequence of the acoustic force applied has also been suggested. This mechanism hypothesizes that the acoustic force induces the formation of pores in the neuronal membrane which can in turn trigger a cascade effect resulting in changes in membrane permeability and therefore neuron polarization. This mechanism is strictly linked with the better known cavitation effects resulting by application of acoustic radiation force to the membrane (**Figure 12f**). Microtubules have also been involved in the explanation of the mechanisms underlying TUS, as some stimulation protocol use frequencies in the range of microtubule resonance, thereby causing them to fluctuate and affect local electrical activity or synaptic homeostatic dynamics (**Figure 12g**). A final mechanism that has been suggested, and which might be a consequence of the previously revised mechanisms, is the

disruption of neurotransmitter or peptide homeostasis in the synapse (**Figure 12h**). It is possible that several of these mechanisms play a role in mediating TUS-related behavioural and neural effects; they are not mutually exclusive. Different combinations of TUS parameters may result in different set of mechanisms having the primary role in transducing the ultrasound into neural activity and, potentially, behavioural change.

To conclude, the integration of NIBS with some of the anatomical and functional techniques described throughout this chapter constitutes an ideal multi-disciplinary approach necessary to fully understand the relationships between anatomy, brain activity and value-guided processes such as credit assignment. Indeed, throughout this chapter I have extensively discussed how the function of a brain area is influenced by its anatomical properties and connectivity, and how their integration supports its function. Over the next chapters I will use such a multidisciplinary approach to look at the anatomical and functional basis of credit assignment.

First, an understanding of the anatomy of the brain circuits connecting key areas supporting this cognitive process will be addressed in the next chapter. Over the following chapters, a novel method useful to causally establish the functional role that some of these areas may have in credit assignment will be described and combined with computational models to describe how and where in the brain this process arises.

Chapter 2

Dichotomous organization of amygdala/temporal-prefrontal bundles in humans and monkeys

2.1 ABSTRACT

The interactions of anterior temporal structures, and especially the amygdala, with prefrontal cortex are pivotal to learning, decision-making, and socio-emotional regulation. A clear anatomical description of the organization and dissociation of fiber bundles linking amygdala and prefrontal cortex in humans is still poorly understood. Using diffusion imaging techniques, we reconstructed fiber bundles between these anatomical regions in human and macaque brains. First, by studying macaques we assessed which aspects of connectivity known from tracer studies could be identified with diffusion imaging. Second, by comparing diffusion imaging results in humans and macaques we estimated patterns of fibers coursing between human amygdala and prefrontal cortex and compared them with those in the monkey. In posterior prefrontal cortex, we observed a prominent and well-preserved bifurcation of bundles into primarily two fiber systems – an

amygdalofugal path and an uncinate path – in both species. This dissociation fades away in more rostral prefrontal regions.

2.2 INTRODUCTION

The neural circuits centered on the amygdala and temporal lobe on the one hand and the prefrontal cortex on the other are crucial in a variety of complex behaviors, including reward-based learning and decision-making (De Martino et al., 2006; Hunt and Hayden, 2017; Morrison and Salzman, 2010; Murray and Wise, 2010; Rudebeck and Murray, 2014; Rushworth et al., 2011), emotional, and social behavior (Noonan et al., 2014; Volman et al., 2011; Whalen et al., 1998). Recent models increasingly emphasize how the amygdala, a complex subcortical area located in the medial bank of the anterior temporal lobe, and prefrontal cortex (PFC) do not encode isolated computations but have instead complementary roles in these processes (Munuera et al., 2018; Saez et al., 2015; Murray and Wise, 2010). Similar computational divisions of labor have been proposed regarding anterior temporal association cortex and prefrontal cortex (Freedman et al., 2003).

To understand these computational roles requires an understanding of the anatomy underlying the brain systems involved (Marr, 1982). Specifically, we want to understand the principles of connectivity between amygdala/anterior temporal and prefrontal systems, in order to better comprehend how information can flow between the various nodes of these networks. Such principles, for instance, are evident in early models by Carmichael and Price (1995) that argue for a mostly dichotomous

organization of circuits stretches across temporal and frontal lobes. This line of work, however, has mostly been based on invasive tracer data that measure area-to-area connections with high spatial accuracy. These approaches are not available for in-vivo human studies. Recent work on diffusion MRI tractography, the only available method to establish structural connectivity in-vivo, has demonstrated that such an approach to reconstruct area-to-area connections has many pitfalls (Maier-Hein et al., 2017; Reveley et al., 2017). A more feasible strategy is to reconstruct fiber bundles (Mars et al., 2016a; Thiebaut de Schotten et al., 2011). By seeding in the reliably identifiable white matter, it becomes possible to isolate the core of fiber bundles and then follow their trajectory towards the grey matter in both directions. Although this technique has the potential to elucidate principles of connectivity between systems, it relies on a different philosophy than the invasive tracer studies, making it difficult to draw firm comparative conclusion between results obtained using the two techniques. Therefore, we here apply this technique first in the macaque monkey, for which tracer results are available and comparison between the techniques can thus be established, and then to the human.

Based on white matter dissections and tract-tracing studies in non-human primates, anterior temporal regions, including amygdala, and the frontal lobe have been shown to have widespread reciprocal connections (Aggleton et al., 2015; Barbas et al., 2011; Barbas and De Olmos, 1990; Ghashghaei et al., 2007; Ghashghaei and Barbas, 2002; Price, 2003; Timbie and Barbas, 2014). Two pathways are known to primarily connect anterior temporal regions including the amygdala and the frontal lobe: a major

association fiber pathway, the uncinate fascicle (UF) (Amaral and Price, 1984; Déjerine, 1895; Klingler and Gloor, 1960; Kuypers et al., 1965; Nauta, 1961; Thiebaut de Schotten et al., 2012; Yeterian et al., 2012), and a limbic-cortical ventral amygdalofugal pathway from here on also referred simply to as the ‘amygdalofugal’, AmF) (Aggleton et al., 1980; Bachevalier et al., 1985; Lehman et al., 2011; Nauta, 1961; Nieuwenhuys et al., 2008; Russchen et al., 1985). To date the AmF has not been described in great detail in the primate brain in-vivo (Carmichael and Price, 1995; Croxson et al., 2005) and is often omitted in human studies of circuits for social-emotional behavior and decision making (Alm et al., 2015; Ameis and Catani, 2015), even though its role in cognitive and emotional processes in animals is critical (Bachevalier et al., 1985; de Guglielmo et al., 2019). Therefore, we focus our investigation on the translatability of knowledge of macaque amygdala/anterior temporal-prefrontal fiber bundles to the human brain on these two tracts.

We used high-resolution diffusion MRI data in the macaque to ascertain the degree to which the diffusion MRI tractography approach identified aspects of amygdala/anterior temporal-prefrontal connectivity known from previous macaque studies. This demonstrated that the core bundles of the AmF perforate the porous substantia innominata in the basal forebrain. In this inhomogeneous tissue the fractional anisotropy index is not fully informative of tract integrity. While this could lead to false negatives in deterministic tractography, importantly, the probabilistic tractography approach we employ here demonstrates a distribution of streamlines running through this area. Accordingly, here we have adopted a probabilistic

approach to tractography strongly informed by prior anatomical knowledge to reconstruct this tract in a manner similar to previous tract-tracing studies.

Using this approach, we were able to robustly define and dissociate a medial amygdalofugal pathway and an orbital uncinate tract in both the macaque and human brain. Together, these pathways form a larger constellation of amygdala/anterior temporal-prefrontal circuits, but each with a distinct connectional profile interfacing with a unique set of brain regions. While the amygdalofugal pathway predominantly extended in the white matter between the amygdala, nucleus accumbens and subgenual cingulate, ventromedial, and frontopolar regions in ventral PFC, the uncinate pathway primarily coursed in the white matter adjacent to anterior temporal regions, including the amygdala, lateral orbital areas and frontopolar cortex. The relationship and structure of these tracts is preserved across primate species, supporting the translation of insights from non-human primate anatomy to inform our understanding of the human brain.

2.3 MATERIALS AND METHODS

2.3.1 Data acquisition

Human in-vivo diffusion MRI data were provided by the Human Connectome Project (HCP), WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University (Van Essen et al., 2013). The minimally preprocessed datasets of the first twenty-four subjects from the Q2 public data release were used (age range 22-35 years; thirteen females). No data were excluded for analysis. This study did not include an experimental manipulation. Accordingly, there was no pre-selection nor restriction for group allocation. Data acquisition protocols are detailed in Ugurbil et al. (2013) and Sotiropoulos et al. (2013). The diffusion data were collected at a 1.25 mm isotropic resolution across the entire brain on a customized 3T Siemens Skyra scanner using a monopolar Stejskal-Tanner diffusion scheme with a slice-accelerated EPI readout. Sampling in q-space included 3 shells at $b = 1000, 2000, \text{ and } 3000 \text{ s/mm}^2$. For each shell 90 diffusion gradient directions and 6 non-diffusion weighted images ($b = 0 \text{ s/mm}^2$) were acquired with reversed phase-encoding direction for TOPUP distortion correction (Andersson et al., 2003). A subject-specific cortical surface mesh with 10k vertices was created using FreeSurfer on the T1-weighted image (acquired using an MPRAGE sequence at 0.7 mm isotropic resolutions) and aligned to the diffusion space as part of the HCP's minimum preprocessing pipeline (Glasser et al., 2013).

Macaque diffusion MRI data were obtained *ex-vivo* from three rhesus monkeys (*Macaca mulatta*, age range at death 4-14 years, mean age 8 years, standard deviation 6.6; one female) using a 7T magnet with a Varian DirectDrive™ (Agilent Technologies, Santa Clara, CA, USA). All intact post-mortem brain samples available at the time of the investigation were included in this study. No data was excluded for analysis. Immediately after death, the brains were extracted from the cranium, perfusion fixed with formalin and stored. Approximately one week before MRI scanning they were perfused in phosphate buffer solution to enhance their diffusion signal. During scanning, the brains were suspended in fomblin™, which has no magnetic properties itself and therefore allows an unbiased scan. For each brain, nine non-diffusion-weighted ($b = 0 \text{ s/mm}^2$) and a single diffusion-weighted ($b = 4000 \text{ s/mm}^2$) datasets were acquired using a single line readout, 2D Stejskal-Tanner pulse sequence (Stejskal and Tanner, 1965) (TE/TR: 25 ms/10 sec; matrix size: 128 x 128; resolution 0.6 mm x 0.6 mm; 128 slices; slice thickness: 0.6 mm; 131 isotopically distributed diffusion directions; receiver bandwidth: 100 kHz).

In addition, in the same session for each brain a T1-map, resulting in a T1-weighted image, was collected (MPRAGE sequence at TE/TR = 8ms/10s; $T_i = 10.6 \text{ ms}$; Matrix Size = 128 x 128; 128 slices; slice thickness: 0.6 mm). Death and fixation have an impact on the brain, causing a reduction in *ex-vivo* tissue diffusivity. Therefore, to achieve an equivalent diffusion contrast to in-vivo data, the diffusion coefficient was increased from $b = 1000$ to 4000 s/mm^2 (Mars et al., 2016a) (**Figure 1**). Importantly, although diffusion

magnitude is affected and requires compensation, diffusion anisotropy is largely preserved *post-mortem* (D'Arceuil and de Crespigny, 2007).




Figure 1. Methodological pipeline and white matter seed masks used to reconstruct each pathway. (A) We combined high-quality, high-resolution diffusion imaging data from in-vivo human ($n=24$) and post-mortem monkey (*Macaca mulatta*; $n=3$) with probabilistic tractography analyses to reconstruct the course and connectivity pattern of two tracts connecting amygdala and anterior temporal cortex with the frontal lobe: the ventral amygdalofugal pathway and uncinate fasciculus. For comparison, tractographic analyses were also run on three other limbic-PFC tracts: the anterior limb of the internal capsule, the extreme capsule and the cingulum bundle. Anatomical reconstruction of the 5 tracts was made by reference to previously published atlases (Nieuwenhuys et al., 2008; Paxinos et al., 2000), tract-tracing studies (Fudge et al., 2002; Oler et al., 2017; Schmahmann and Pandya, 2006) and diffusion tractography studies (Catani et al., 2013; Croxson et al., 2005). We

aimed to keep all seeding and tracking parameters constant across tracts and species. To constrain the tractography algorithm, a priori anatomical priors were used for the analyses and homologous tract-specific exclusion masks in both species were used in order to obtain high fidelity results. (B-F) White matter seed masks used to reconstruct the ventral amygdalofugal pathway (B), the uncinate fascicle (C), the internal capsule (D), the extreme capsule (E) and the cingulum bundle (F). All seeds are represented for the left hemisphere in both species.

2.3.2 Data preprocessing

Human data were preprocessed according to the HCP minimal preprocessing pipelines (Glasser et al., 2013; Sotiropoulos et al., 2013). Macaque data preprocessing and all subsequent analyses were performed using tools from FSL (Smith et al., 2004) and the in-house MR Comparative Anatomy Toolbox (MrCat). In short, non-diffusion-weighted images were extracted from the full set, averaged, and corrected for spatial RF-field inhomogeneity bias before being linearly registered to the average T1-weighted structural image of the same animal. The T1-weighted images in turn were linearly and non-linearly registered to a dedicated ex-vivo macaque T1w template aligned to the standard macaque F99 space (Van Essen and Dierker, 2007). Combined, these registrations allowed a direct mapping between diffusion-, T1-weighted, and standard space for both human and macaque data.

Diffusion tensors were fitted to each voxel using DTIFIT (Behrens et al., 2003). Voxel-wise crossing-fiber model fitting of diffusion orientations was

performed in both human and macaque datasets using FSL's BedpostX (Behrens et al., 2007). For the human in-vivo data, a multi-shell extension was used to reduce overfitting of crossing fibers due to non-monoexponential diffusion decay (Jbabdi et al., 2012). Up to three fiber orientations per voxel were allowed. This produced voxel-wise posterior distributions of fiber orientations that were subsequently used in probabilistic tractography.

2.3.3 Probabilistic tractography

The principal aim of the tractography analyses was to reconstruct the prefrontal projections of the ventral amygdalofugal pathway (AmF) and uncinate fascicle (UF). Tractography recipes for both were created by reference to previously published atlases (Nieuwenhuys et al., 2008; Paxinos et al., 2000), tract-tracing studies (Fudge et al., 2002; Lehman et al., 2011; Schmahmann and Pandya, 2006), and diffusion tractography studies (Catani et al., 2013; Croxson et al., 2005; Jbabdi et al., 2013). All masks were created in F99 space (Van Essen and Dierker, 2007) in the monkey brain and MNI space in the human brain and subsequently warped to each individual's diffusion space and adjusted by hand to ensure anatomical accuracy.

The seed mask of the AmF was drawn in the sub-commissural white matter perforating the substantia innominata, a region between the dorsal amygdaloid and the bed nuclei of the stria terminalis (BNST). The mask was aimed to carefully reproduced macaque tract-tracing studies (deCampo and Fudge, 2013; Fudge and Haber, 2001; Schmahmann and Pandya, 2006), located medially to the ventral pallidum (VP) and sublenticular extended

amygdala (SLEA) and dorsally to the nucleus basalis of Meynert (NBM) (Mai et al., 2015; Paxinos et al., 2000). In accordance with the local trajectory of the AmF as estimated by tract-tracing the seed mask was constrained to contiguous voxels showing high fractional anisotropy in an anterior-posterior direction (see **Figure 1B**). UF was seeded axially in the anterior temporal lobe, in the white matter rostro-lateral to the amygdala (see **Figure 1C**). After testing various coronal seed masks, an axial seed mask was chosen to account for the strong curve of the fibers from dorsal-ventral to anterior-posterior orientation to orientation as they enter into the frontal lobe (Catani and Thiebaut de Schotten, 2008).

To test the specificity of our tractography results, we reconstructed three alternative tracts projecting to the prefrontal cortex: the anterior limb of the internal capsule (ICa), the extreme capsule (EmC), the cingulum bundle (CB). These tracts were chosen because together they constitute prominent systems connecting thalamic, anterior temporal, and limbic regions with PFC regions and are expected to run or terminate in close proximity to UF and AmF. ICa was seeded in the ventral section of the anterior limb of the internal capsule (see **Figure 1D**) passing between the caudate and pallidum, superior to the anterior commissure, and inferior and lateral to the caudate, in a position similar to that used by Jbabdi et al. (2013). The EmC seed mask was drawn in the sheet of white matter between the putamen and insula (see **Figure 1E**). The seed voxels were placed as laterally as possible within this sheet to favor fibers of the extreme over the external capsule, although due to the resolution of diffusion imaging it is not possible to fully ensure that no fibers of external capsule were included in the reconstructed EmC (Mars et

al., 2016a). EmC seed voxels were placed just superior to the UF, in order to avoid including fibers belonging to this latter tract. The CB seed mask was drawn in a coronal plane capturing the WM dorsal and medial to the corpus callosum, within the cingulate gyrus (see **Figure 1F**). Two additional control tracts were reconstructed in the human brain to address their differentiation from AmF and UF, namely the stria terminalis (ST) and fornix Fx). The ST seed mask was placed in its caudal part, just anterior to the curve this bundle makes in the proximity of the splenium of the corpus callosum (see **Figure 2A**). The Fx seed mask was drawn at the same level as the ST but in a more medial position, in a portion of WM just ventral to the corpus callosum (see **Figure 2B**).



Figure 2. White matter seed masks for two limbic control tracts in the human brain. (A-B) White matter seed masks used to reconstruct the stria terminalis (A) and fornix (B) in the human brain. Similarly to other tracts reconstructed in this article, seeds are represented for the left human hemisphere.

To constrain the tractography algorithm, exclusion masks were drawn in both species to exclude false positive results in areas of high crossing fibers as follows: (1) within the basal ganglia to avoid picking spurious subcortical tracts (except for ICa tracking); (2) posteriorly to the seeds to prevent the projections from running backwards as the prefrontal cortical streamlines were the focus of this study; (3) an axial slice at the level of the superior temporal gyrus to prevent tracts from running in a ventral direction in an unconstrained manner (except for UF, CB, and AmF); (4) on the axial-coronal slices cutting across the thalamus, basal ganglia, and corpus callosum to exclude these subcortical and callosal projections; (5) in the dorsal cingulate cortex to avoid leakage of tracts to nearby bundles as a result of high curvature of the tracts (except for CB tracking); (6) an L-shaped coronal mask from the paracingulate cortex and the inferior frontal sulcus to the vertex to exclude tracking in the superior longitudinal fascicle; (7) another L-shaped mask, only for ST and Fx, drawn at the level of the retrosplenial cortex and dorsal to the corpus callosum to exclude inappropriate leakage into dorsal parietal and dorsal frontal cortices and (8) the opposite hemisphere to only track ipsilateral tracts. An exclusion mask of the CSF was used in each subject to prevent fibers from jumping across sulci during tracking. Streamlines encountering any of the exclusion masks were excluded from the tractography results. A coronal waypoint section was drawn in the frontal lobe at the level of the caudal genu of the corpus callosum to ensure that the fibers emanating from the seeds were projecting to the prefrontal cortex; this prevented leakage of the EmC into the

corticospinal tract. Only for the ST and Fx, two tract-specific coronal masks were placed in the temporal and forebrain white matter.

Probabilistic tractography was performed in each individual's diffusion space. All seeding and tracking parameters were kept constant across tracts and only adjusted between species to deal with differences in brain size and data resolution. Human tractography parameters were: maximum of 3,200 steps per sample; 10,000 samples; step size of 0.25 mm; curvature threshold of 0.2. Macaque tractography parameters were: maximum 3,200 steps per sample; 10,000 samples; step size of 0.1 mm; curvature threshold of 0.2.

A visitation map or 'tractogram' was constructed for each individual in order to allow comparison of these maps between tracts, subjects, and species. Each tractogram was log-transformed to account for the exponential decrease of visitation probability with distance and normalised by dividing each voxel's value by the 75th percentile value across the tractogram, thereby removing potential bias of differences in numbers of streamlines between tracts and across species. In both human and macaques, the focus of the investigation was on the organization of the tracts in prefrontal cortex. Prefrontal streamlines of each tract were averaged in a species-specific group-template tractogram. As such, the center of the group is defined as the mean of the individual subject tractograms. For the purposes of visualization only, the normalized tractograms were subsequently thresholded with minimum and maximum values equal to 0.5 and 2, respectively.

We further asked how streamlines origination in either the amygdala or temporal pole were organized with respect to each other within the AmF and UF bundles. To this end, we performed a within-bundle connectivity-

gradient analysis. We placed seeds in the amygdala (Amyg) and the temporal pole (TP) and used the core of the AmF and UF tracts as waypoints (i.e. these waypoint masks were identical to the seed masks of the previous analyses). This resulted in four normalized tractograms, masked by the bundles as described above: AmF_{Amyg} , AmF_{TP} , UF_{Amyg} , and UF_{TP} . A gradient ratio describing the relative occurrence of Amyg and TP streamlines was defined for the AmF bundle as: $ratio_{AmF} = \frac{AmF_{Amyg} - AmF_{TP}}{AmF_{Amyg} + AmF_{TP}}$, and for the UF bundle as: $ratio_{UF} = \frac{UF_{Amyg} - UF_{TP}}{UF_{Amyg} + UF_{TP}}$. In other words, these indices quantify the relative probability that streamlines coursing through the AmF and UF bundles have originated from either the amygdala (positive values) or the temporal pole (negative values).

2.3.4 Connectivity fingerprints

The differential distribution of streamlines of the reconstructed tracts with areas in prefrontal cortex were quantified by creating so-called connectivity fingerprints (Mars et al., 2018a; Passingham et al., 2002) of each tract with regions of interest (ROIs) in prefrontal cortex. Such fingerprints can be used to compare tracts' prefrontal projections with one another, but also to compare the tracts across species. Human and macaque prefrontal ROIs were created based on previously published coordinates (Neubert et al., 2015, 2014; Sallet et al., 2013) or drawn by hand on the basis of anatomical atlases (Mai et al., 2015; Paxinos et al., 2000; Petrides et al., 2012; Yeterian et al., 2012) (see **Figure 3**). Each coordinate was projected onto a subject-specific cortical surface mesh with 10,000 vertices (10k) in the human and to a template 10k surface in the macaque. Here, on the surface representation,

these coordinates were expanded to circular ROIs following the cortical contours, with a 10 mm radius in humans and a 3 mm radius in macaques. The resulting ROIs were then projected back to each subject's native brain image in volumetric space. The creation of ROIs on the surface rather than in volumetric space allows us to account for individual idiosyncrasies in sulci and gyri anatomy, ensuring that we have truly cortical rather than spherical ROIs. The white matter/gray matter border of each ROI was then extracted and used as target region in the connectivity fingerprint. This approach allows to maximize the estimation of streamlines projecting to each region while minimizing the reduction on signal consequent to tissue-related poor anisotropic signal characteristic of ROIs created in gray matter only. The fingerprints are then constructed by counting for each tract the average number of streamlines hitting the voxels of each ROI. To allow comparison across tracts and species, these values were normalized by dividing each tract-fingerprint by its maximum values across the selected ROIs. The center of the group is defined as the mean of the individual subject connectivity fingerprints. Connectivity fingerprints can then be compared by calculating the Manhattan distance between them (Mars et al., 2016b).




Figure 3. Target regions of interest used to estimate prefrontal connectivity fingerprint for each tract. (A-B) Regions of interest (ROI) were created on the cortical surface of the macaque (A; F99 space) and human (B; MNI space) brain. Each ROI was then projected to the volumetric gray matter / white matter border and used for the connectivity fingerprint analyses. All ROI labels as in Figure 7.

2.4 RESULTS

Our goal was to reconstruct the anatomical organization of the amgydalofugal and uncinate bundles in white matter (WM) within amygdala/anterior temporal regions (ATR) and prefrontal territory. In particular, we were interested in assessing the course of these pathways towards the prefrontal cortex and their interaction with a wide range of prefrontal regions of interest. Given the difficulty of reconstructing these pathways, we employed the following strategy. First, we reconstructed the tracts in the macaque monkey brain, so that they can be compared to known tracer results (**Figure 4A-C**, left and middle panels). This demonstrated whether our tractography techniques are valid for these pathways. Then, we applied the same tracking strategy to the human, using the seeding and masking procedures applied successfully in the macaque (**Figure 4A-C**, right panels). Finally, we compared the tracts' projections directly and across species and with respect to a number of control tracts (**Figure 7**).




Figure 4. Histology-informed approach for the reconstruction of the frontal limb of the ventral amygdalofugal pathway. (A-C, left and middle panels) Coronal sections depicting side-by-side the reconstruction of the frontal limb of the amygdalofugal pathway (AmF) in the macaque brain using probabilistic tractography on diffusion MRI data (middle panels) and published reconstructions of the tract using radio-labeled fibers identified using tract-tracing injections in macaque (left panels; adapted from (Oler et al., 2017)). The AmF tractography reconstruction showed marked correspondence in anatomy and course of the fibers with the histological tracing data. AmF fibers leave the amygdala caudal to the anterior commissure (ac) and turn medially towards the medial wall of the hemisphere (A) and rostrally towards the frontal lobe by coursing ventrally to

the ventromedial striatum (B-C). **(A-C, right panels)** The histology-informed reconstruction of the AmF was then used to guide the tractographic algorithm in the human brain, including the seed, exclusion and waypoint masks were used in both species (see Methods for details on the masks). The human AmF follows a course matching the results observed in the macaque brain (A-C, middle panels) with fibers projecting both medially (A) and rostrally (B-C). T1-weighted image are shown in radiological convention. ac, anterior commissure; C, caudate nucleus; GP, globus pallidus; ic, internal capsule; P, putamen; L, left hemisphere; R, right hemisphere.

2.4.1 Amygdalofugal pathway in the macaque monkey brain

The frontal body of the AmF was represented by a thin bundle of fibers running in the WM adjacent to medial subcortical and cortical brain structures (**Figure 5A-E** and see **Figure 6A-C**). In a posterior section of the macaque brain, fibers running in the medially projecting limb of the AmF pathway coursed in the WM dorsal to the amygdala lying between this structure and the medial longitudinal fissure at the level of the posterior anterior commissure (AC) and sublenticular extended amygdala area (SLEA) (**Figure 4A-B, Figure 5A** and see **Figure 6A-B**). Here, a first bundle of AmF fibers ran in the WM territory in between the nucleus basalis of Meynert, lentiform nucleus, ventral pallidum, and the piriform cortex (see **Figure 6A-B**). The medial AmF fibers are known to separate into different sets of projections (Lehman et al., 2011; Nauta, 1961; Novotny, 1977; Oler et al., 2017) that extended ventrally to innervate the anterior hypothalamic areas at the level of the paraventricular-supraoptic nuclei and dorsally to connect the amygdala with the thalamus, the BNST, and septum. All of these fibers were visible in

our data leaving the seed location, but the frontal waypoint mask means that the pathways to the prefrontal cortex, which was the focus of the current investigation, were the most prominent (**Figure 5A** and see **Figure 6A-B**). This second bundle of the AmF ran in the WM surrounding the substantia innominata, ventral pallidum, ventral bank of the AC and more rostrally in the region ventral to the ventral striatum/nucleus accumbens (VS/NAc; **Figure 4C** and **Figure 5A-E**). From tract-tracing, we know that some fibers connecting ventral PFC and amygdala are intermingled with the much more prominent anterior commissure (Lehman, et al. 2011). The AC is very prominent in the principal fiber direction of diffusion MRI data, but nevertheless some weak AmF fibers overlapping with AC were evident in our tractography data (**Figure 4B**, **Figure 5** and see **Figure 6B**).

It must be noted that over its course in the basal forebrain and ventral to the striatum, the AmF is likely to interdigitate with other connectional systems including the rostral fibers of the thalamic projections, the medial forebrain bundle, the diagonal band of Broca and the dopaminergic fibers connecting the ventral tegmental area with prefrontal cortex. Eventually, the anterior fibers of the AmF extended in the orbital periallocortex on the medial wall of the hemisphere and in the PFC by coursing in WM adjacent to the subgenual anterior cingulate cortex (sACC or Brodmann area 25; **Figure 5A**, **B**, **D** and see **Figure 6C**).




Figure 5. Anatomy and organization of the Amygdalofugal Pathway and Uncinate Fascicle in the basal ganglia and frontal lobe of the macaque monkey and human brain. (A) amygdalofugal pathway (AmF, red) and uncinate fascicle (UF, blue) are displayed on eight coronal brain slices. AmF and UF show the same medial-lateral organization in both macaque (left) and human (right) brains. (B-C) Sagittal view of the macaque (B) and human (C) brains displaying the preservation of the dorsal-ventral and medial-lateral organization of the two tracts across both species in caudal subcortical regions of the brain but their merge in caudal orbitofrontal cortex WM. (D-E) The medial-lateral organization of the two tracts is shown on axial sections of macaque (D) and human (E) brains. The AmF follows a course medial to the UF with its fibers wrapping the full extent of the ventral striatum and nucleus accumbens in both macaques (D) and humans (E). Here, the main body of the UF runs instead more laterally, primarily in the WM adjacent to the

ventral-lateral striatum and the insular cortex. The tracts merge into an intermingled connectional system in the posterior OFC (area 13) WM and for the full extent of the ventromedial prefrontal cortex. Macaque coordinates are in F99 space and human coordinates are in MNI space. (F) Schematic template representing the organization of AmF (red) and UF (blue) in the ventral prefrontal cortex in the human brain. Dashed line (red) shows a bundle subcortical AmF fibers running underneath and through the ventral striatum.

Near the sACC, the AmF split again into two sub-sections. The first sub-section ran dorsally alongside the cingulum bundle (CB) in anterior cingulate sulcus and perigenual WM. The second branch instead stretched between the sACC and the frontopolar cortex passing subjacent to medial orbitofrontal cortex (OFC) between the gyrus rectus and the medial orbital gyrus (**Figure 5B** and **Figure 5D**). Because diffusion MRI cannot discern between monosynaptic and polysynaptic connections, the identified fibers could represent either type of projections and could run in either direction.

2.4.2 Amygdalofugal pathway in the human brain

As in the macaque, we observed that in the human brain the AmF pathway ran adjacent to the lateral and basal nuclei of the amygdala and curves in the more dorsal WM within the substantia innominate and SLEA bordering the central and medial nuclei where additional fibers entered the main body of the tract (**Figure 4A-C** right panels, **Figure 5A** and see **Figure 6A-B**). Here, the AmF extended subcortically along the medial wall and gives rise to two

limbic bundles. Consistent with the macaque, the first set of AmF fibers constituted a medially projecting path that followed a trajectory resembling a reversed S-shaped coursing in the WM between the caudal amygdala, periventricular regions both superior and inferior to the AC and bed nuclei of stria terminalis, hypothalamus and thalamus. A second portion of the ventral AmF pathway coursed in the WM surrounding the substantia innominata and SLEA by interdigitating with AC fibers or running through the sub-striatal WM surrounding the VS/NAc (**Figure 5A** and see **Figure 6A-B**).

The human AmF in the PFC predominantly coursed alongside the sACC although a smaller contingent of fibers bordered the caudal OFC WM. In ventral PFC, the main body of the AmF kept running along the medial bank of the cortex in a position adjacent to ventromedial PFC (vmPFC) in the gyrus rectus and medial OFC (**Figure 5A**, **Figure 5C**, **Figure 5E** and see **Figure 6C**). Recently, Neubert et al. (2014) suggested that the frontopolar cortex (FPC) in the human brain is composed of functionally specialized medial and lateral sub-regions to a greater extent than in the macaque. The AmF projections kept running close to the medial wall in the FPC but fewer projections also extended laterally to reach the lateral FPC (**Figure 5C** and **Figure 5E**).

The architecture of the human AmF closely resembled that of the macaque brain, suggesting that this tract is evolutionarily preserved; both its medial and anterior branches were very similar in the human and macaque monkey brains (**Figure 5A-E**) despite the limited WM in the latter species (Schoenemann et al., 2005). Minor differences arose when we look at FPC;

in the human few AmF fibers headed towards the lateral FPC and the majority of the tract innervated medial FPC (**Figure 5C** and **E**).




Figure 6. Comparison of the anatomical architecture of the ventral amygdalofugal pathway and uncinate fascicle in macaque monkey and human brains. (A-C) Anatomy of macaque and human amygdalofugal pathway on three coronal slices showed from caudal (A) to rostral (C). (C-F) Anatomy of macaque a human uncinate fascicle in a sagittal slice and in two coronal slices showed from caudal (E) to rostral (F). Macaque coordinates are in F99 space and human coordinates are in MNI space. Note the macaque T1-weighted image (A-F, left panels) contrast-to-noise ratio is inverted because the brain sample was scanned ex-vivo, conversely to the human in-vivo samples (A-F, right panels). Ex-vivo acquisition of T1-weighted images are characterized by longer relaxation times of brain tissue which causes the white/grey matter contrast to be inverted compared to standard in-vivo acquisitions. A, amygdala; BNST, bed nucleus of stria terminalis; H, hypothalamus; Ins, insular cortex; NBM, nucleus basalis of Meynert; sACC, subgenual anterior cingulate cortex (Brodmann area 25); S, septum; SLEA, sublenticular extended amygdala; VP, ventral pallidum; all other labels as in Figure 4 and Figure 7. Tractograms are thresholded with minimum and maximum values equal to 0.5 and 2, respectively

2.4.3 Uncinate fascicle in the macaque monkey brain

The UF is a hook-shaped tract containing fibers connecting the temporopolar cortex and amygdala with orbital gyri and ventrolateral PFC in a bidirectional manner. On a ventral section, the macaque UF extended along the lateral surface of the temporal lobe where it made a dorsal C-shaped curve through the superior temporal gyrus of the temporal lobe, coursing medially to the Sylvian fissure (**Figure 5D** and see **Figure 6D**). Although in this region the principal fiber direction runs in a superior-inferior plane, some fibers branched along a medial-lateral axis towards the amygdala, thereby interconnecting fibers adjacent to ATR, and lateral prefrontal cortex (**Figure 5A, D** and see **Figure 6E**). Other fibers also ran in a more caudal position along the parahippocampal gyrus and rhinal cortex (**Figure 5D** and see **Figure 6D-E**). At the level of the limen insulae, UF fibers ran in a WM region ventrally to the claustrum, ventrolaterally to the putamen and globus pallidus, and medially to the insular cortex. Here, the most dorsal fibers of the UF intermingled with the ventral axons of the EmC as well as with ventrolateral PFC WM (**Figure 5A** and see **Figure 6D-F**).

In monkey PFC, the UF strongly innervated the WM bordering the caudolateral OFC, the opercular and inferior frontal areas although a subsection of fibers extended next to the subgenual cingulate area and the medial bank of posterior OFC (**Figure 5A-B** and see **Figure 6F**). Here, the more medial UF fibers joined with the AmF and CB fibers to innervate the ACC WM (**Figure 5B**). Overall, the main body of the UF in ventral PFC stretched throughout lateral, central, and medial orbital gyri giving rise to a

complex network in OFC which extended into the frontopolar cortex (**Figure 5D**).

2.4.4 Uncinate fascicle in the human brain

The human UF had a preserved C-shaped structure lying between rostromedial temporal regions and ventrolateral PFC (**Figure 5C** and see **Figure 6D**). Similarly to macaque temporal UF, it was anatomically organized predominantly along a dorsal-ventral gradient but with lateral fibers extending towards amygdala nuclei and parahippocampal and rhinal areas (**Figure 5E** and see **Figure 6D-E**).

In the intersection between prefrontal and temporal cortices, it coursed in close proximity to the insular cortex and regions in the caudal territory of the inferior frontal gyrus (**Figure 5A, E** and see **Figure 6D-F**). More medially, human UF encompassed the lateral ventral striatum and kept a predominant lateral course in the ventral PFC (**Figure 5A** and see **Figure 6D-F**), whereas a bundle of UF fibers stretched alongside the cingulate AmF and CB (**Figure 5C**). Here, the main body of the UF kept running in a lateral position sending widespread innervation to all the ventrolateral PFC WM. Furthermore, a subsection of fibers coursed also more medially, adjacent to sACC and more prominently to posterior OFC (**Figure 5A, E** and see **Figure 6F**).

In the anterior territory of ventral PFC the human UF lay along both medial and lateral orbital areas in a widespread fashion. A major bundle of fibers occupied a substantial part of the lateral OFC WM and adjacent cortical areas located on the ventral surface of the brain. Another lateral branch of the UF extended more rostrally in the WM flanking the lateral FPC,

whereas medial FPC WM was primarily reached by the more medially coursing UF fibers (**Figure 5C, E** and see **Figure 6F**).

Again, as for the AmF pathway, the anatomy of the UF as well as its interdigitation with AmF in the human brain was highly conserved in monkey and human brains (**Figure 5A-E** and see **Figure 6D-F**). There may be differences in FPC organization between humans and macaques (Neubert et al., 2015) and it was notable that the UF innervated medial and lateral FPC WM in the human brain in equal proportion (**Figure 5E**).

2.4.5 Comparison of the two fiber bundles

The main bodies of the human AmF and UF were anatomically segregated and constituted two anatomically distinct connectional systems in ATR, and striatal WM. Alongside the dorsolateral amygdala, AmF interdigitated with more lateral temporo-frontal UF axons. At this level, it was possible to identify in both species a medial-lateral organization of the two bundles with the former leaving the amygdala WM more caudally and the latter's fibers projecting to and from this subcortical structure in a more anterior position (**Figure 5**).

More anteriorly, the two tracts maintained a clear and evolutionary preserved organization along a medial-lateral gradient, but also started to develop a dorsal-ventral relationship (**Figure 5A-E**). As previously described, in this region the AmF coursed medial to the ventral striatum, just underneath the anterior limb of the ICA and medial to the ventromedial striatum, specifically inferiorly to the internal and external globus pallidus, the ventral head of the caudate and the medial bank of the putamen (**Figure 5A-E** and

see **Figure 6A-B**). Conversely, the UF approached the basal ganglia more laterally. A medial portion of UF fibers interdigitated with lateral AmF axons and courses ventral to the AmF in the WM surrounding the putamen. In contrast to the more ventromedially positioned course, UF fibers ran alongside the full ventrolateral extent of the basal ganglia. In the WM lateral to the putamen and ventral to the claustrum, its fibers interlaced with the ventral axons of the EmC (**Figure 5A-E** and see **Figure 6E-F**).

The relative mediolateral positioning of the AmF and UF extended more rostrally within the WM adjacent to the sACC (**Figure 5A** and **Figure 5D-E**); a major component of the AmF lay adjacent to sACC (see **Figure 6C**) while only a medial bundle from the UF lay adjacent to sACC (**Figure 5A** and see **Figure 6F**). The relative positioning of the two pathways then began to partially rotate with the UF moving to a location ventral to the AmF as the anatomical distance between the two fiber bundles reduced in the subgenual WM. In the WM adjacent to infralimbic cortex (area 25) the segregation of the two tracts began to fade as the dorsomedial UF fibers approached the ventrolateral AmF axons (**Figure 5A, F** and **G**).

Although the AmF and UF were clearly distinct entities in more posterior parts of the brain their clear separation disappeared in the posterior OFC. Here, the UF and AmF merged into a single connectional system. This anatomical organization was present and conserved across both human and macaques in relation to several homologous landmarks in both species (**Figure 5B-E**; see **Figure 6C** and **F**).

These results describe a preserved dichotomous organization of the AmF and UF along a predominantly medial-lateral axis across both species.

As described in the introduction, the choice of these two fiber bundles was guided by tract-tracing results suggesting their connectivity of prefrontal regions with predominantly ATR fibers, respectively. Although, as argued, tractography is not suitable for an investigation of full area-to-area connectivity, so instead we sought to investigate the contributions of connectivity with these grey matter areas by means of a connectivity-gradient analysis. Specifically, we were interested to see if we could observe a similar medial-lateral gradient of connectivity within the fiber bundles, driven by amygdala and anterior temporal polar connectivity, respectively. In different analyses, we quantified the relative probability of streamlines coursing through the core of the AmF and UF to have originated in the amygdala or in the temporal pole (see Methods). The resultant connectivity-gradient ratio is represented in **Figure 8**, separately for AmF (left panel) and for UF (right panel). These analyses showed that a strong medial-lateral organization in the human brain is present not only across but also within tracts. This is particularly evident for the UF that contained streamlines originating from both the anterior temporal lobe and the amygdala, but with clear prevalence in the lateral and medial direction, respectively. A similar pattern is visible in the AmF, although it contained only very few streamlines originating in the temporal pole.

2.4.6 Prefrontal connectivity and control tracts

We have described the distinct courses of the AmF and UF along their course between the vicinity of the ATR and the prefrontal cortex. To quantify the streamlines of the tracts to prefrontal areas, we can describe their

connectivity fingerprints with prefrontal regions of interest (Mars et al., 2018a). This allows a quantitative comparison of the areal connectivity of AmF and UF. To examine the specificity of the results, we also compared their connectivity fingerprints to those of three control tracts connecting nearby systems: the anterior limb of the internal capsule (ICa), the extreme capsule (EmC), and the cingulum bundle (CB) (**Figure 7**).




Figure 7. Comparison of connectivity fingerprints of the amygdalofugal pathway, uncinate fascicle and adjacent white matter tracts in the human (left) and macaque (right) frontal lobe. (A) The pattern of AmF fibers is preserved across the two species and predominantly targets medial areas in subgenual, orbital, ventral, and polar PFC. (B) Overall UF axons show high similarity between the two species but stronger innervation to the

opercular and insular cortex is observable in the human UF. (C-E) Macaque-human comparison of the prefrontal fibers organization of three connectional systems located in the white matter dorsal to the AmF (internal capsule; C), dorsal to the UF (extreme capsule; D) or interacting with both tracts in the infralimbic and prelimbic cortex (cingulum bundle; E). 9m, dorsomedial PFC; 8m, caudal bank of the dorsomedial PFC; 32, anterior cingulate cortex; 25, subgenual cingulate cortex; 24, dorsal anterior cingulate cortex; preSMA, pre-supplementary motor area; IFGtr, inferior frontal gyrus, pars triangularis; FOp, frontal operculum; Ins, insular cortex; 14m, ventromedial PFC; 13, caudal orbitofrontal cortex; 11, rostral orbitofrontal cortex; 47 (human) or 47/12 (macaque), lateral orbitofrontal cortex (also named lateral convexity of PFC, Murray and Rudebeck, 2018; Rudebeck and Murray, 2014); 10l, lateral frontopolar cortex; 10m, medial frontopolar cortex; 44, caudal bank of the ventrolateral PFC; 45, rostral bank of the ventrolateral PFC. Shades represent standard error of the mean (SEM).

Comparing the connectivity fingerprints of AmF and UF show that in both species, a medial-lateral organization in streamlines density between the two tracts was evident (**Figure 7A-B**). While both bundles were primarily reaching ventral prefrontal cortex in both species, AmF was more likely to reach medial areas such as 25, 14m, 11 and 10m (**Figure 7A**), whereas UF primarily reached more lateral areas such as the inferior frontal gyrus, insular cortex, frontal operculum, area 47 (or the macaque homolog 47/12), and area 45 (**Figure 7B**). However, UF fibers in the rostral ventral PFC were also in contact with more medial OFC (area 11) and FPC (10m).

Some species-specific differences were also evident. Human AmF fibers were denser in areas 11 and 10m, whereas macaque AmF showed

this pattern in area 32 and caudal vmPFC area 14m (**Figure 7A**). Inter-species differences in cingulate fibers may partly be the results of species-specific idiosyncrasies in neuroanatomical geometries, especially when diffusion MRI is used to reconstruct tracts running in the WM adjacent to brain structure presenting different degrees of curvatures in different species. UF also exhibited some contrasting inter-species differences. In the human brain, the numbers of streamlines reaching lateral orbital, ventrolateral, insular, and opercular areas were higher compared to the macaque brain, suggesting a more widespread pattern of connectivity in human prefrontal cortex (**Figure 7B**). This may reflect prefrontal cortical expansion in humans compared with other old-world primates (Smaers et al., 2017). As discussed above, in the most anterior part of prefrontal cortex, AmF and UF tend to intermingle and it is therefore not surprising to see that both reached the frontal pole (area 10m), including the human lateral frontal pole (area 10l) for which the homology in the macaque monkey brain is debated (Neubert et al., 2015).

The connectivity fingerprints of the control tracts were markedly distinct from those of both AmF and UF (**Figure 7C-E**). The anterior limb of the internal capsule carries fibers originating and terminating in the thalamus and brainstem (Lehman et al., 2011). Frontally, its streamlines predominantly reached more dorsomedial regions than AmF and UF, including areas 32, 8m, and 9m (**Figure 7C**). The connectivity pattern of EmC was even more widespread in relation to areas of lateral and dorsal prefrontal cortex, including areas 45, 8m, 9m, and 32 (**Figure 7D**). The cingulum bundle showed high streamline count medial areas (**Figure 7E**). Interestingly,

similarly to the UF, some branches of these control tracts exhibited a more widespread pattern of connectivity in the human compared to macaque brain. Indeed, human EmC streamline counts were higher in more caudal areas of dorsal PFC such as the preSMA (**Figure 7D**). CB fibers, instead, showed a similar trend in frontopolar, orbital, subgenual and midcingulate cortex compared to the macaque (**Figure 7E**).




Figure 8. Gradient organization of ATR-prefrontal fibers. White matter gradients showing the organization of fibers coursing between ATR and prefrontal areas. In the left panel we represent the connectivity gradient ratio of the AmF bundle, indexing the relative probability of AmF streamlines originating in either the amygdala (red-yellow) or temporal pole (blue). The color scale represents the full range of the gradient ratio, from solely originating in the amygdala (yellow), to solely originating in the temporal pole (light blue). In the right panel we represent the same gradient ratio for the UF bundle.

The connectivity fingerprints can be used to quantitatively compare the connectivity profiles across species. We can create a dissimilarity matrix between species by calculating the Manhattan distance for each possible

pair of tracts. This measure indicates how different two connectivity fingerprints are, whereby the best matching pairs are defined as the set of fingerprints that have the smallest distance values (Mars et al., 2016b). When comparing the fingerprints (averaged across the two hemispheres) between species, we saw that the smallest differences for each tract were on the diagonal of the distance matrix (**Figure 9B**), indicating that each tract had a connectivity fingerprint most similar to its counterpart in the other species. AmF and UF were each other's second best choice, but matched best with their counterpart. This similarity is not unexpected given close proximity in ventral PFC and considering the strong interconnection and reciprocal horizontal projections of OFC regions, yet despite these correspondences the dissimilarity matrix confirms the two pathways are distinguishable.

These results suggest that the AmF and UF streamlines reach different parts of prefrontal cortex and, moreover, that their connectivity profiles are distinct from those of nearby fiber bundles. This dissimilarity with respect to the control tracts was also reflected in their course through the white matter (**Figure 9A**). To provide a more complete picture of the organization of AmF and UF in the human brain, a clear topographic organization of these two fibers bundle relative to the three tracts described above in the human brain was observed.

Indeed, the course of both tracts could be dissociated from ICa, EmC and CB but also with respect to two other tracts, namely the stria terminalis (ST) and fornix (Fx), that lie and course in temporal and prefrontal portions of WM closely adjacent to AmF and UF (**Figure 5A**). Because prefrontally they primarily innervate basal forebrain and limbic regions and have poor or

absent projections to more rostral PFC areas (Aggleton, 2008; Catani et al., 2002; Kitt et al., 1987; Oler et al., 2017), with only some ST fibers reaching the posterior vmPFC (Lehman et al., 2011), ST and Fx were not included in the dissimilarity matrix analysis (**Figure 5B**).




Figure 9. Organization of ATR-prefrontal tracts. (A) Topographic organization of the AmF and UF with respect to the other three surrounding WM tracts depicted in panel B. Moreover, it is possible to see the organization of AmF and UF with respect to three bundles that occupy bordering white matter regions in the human brain, namely the Stria Terminalis (ST), Fornix (Fx) and the anterior limb of the Internal Capsule. Despite the closeness of the close proximity of these 5 bundles, they occupy different portion of white matter and can be dissociated to one another. (B) Dissimilarity matrix of the prefrontal course of the tracts of interest across species. The matrix shows a quantitative comparison of the connectivity fingerprints of the five frontal tracts that were reconstructed both in the human and monkey brains. Human tracts are listed on the horizontal axis; macaque tracts are listed on the vertical axis. Pair-wise dissimilarity is

defined as the “Manhattan distance” between each pair of pathways across the two species. The greatest similarity was found when comparing each tract with its homologue in the other species (bright diagonal in the dissimilarity matrix). AmF, ventral amygdalofugal pathway; CB, cingulum bundle; ICa, anterior limb of the internal capsule; EmC, extreme capsule; UF, uncinate fascicle; ST, Stria Terminalis; Fx, Fornix.

2.5 DISCUSSION

In this study we identify and define two closely related but distinct connectional systems linking amygdala and PFC using diffusion MRI tractography in primates. The structure, course, and projections of these tracts in the macaque monkey closely match previous tracer injection studies (Fudge et al., 2002; Fudge and Haber, 2001; Kunishio and Haber, 1994; Lehman et al., 2011). Moreover, we showed that both anatomy and connectional profiles of these two tracts are preserved across humans and monkeys. We demonstrated that both the AmF and UF run separately and in parallel when exiting the medial temporal lobe but then merge to form a single fiber complex in the ventral PFC.

Tract-tracing studies provide important information about where specific axons in fiber bundles terminate or originate in the brain (Petrides and Pandya, 2007). However, the technique is invasive and requires the sacrifice of the subject, limiting its use in human anatomy. Diffusion MRI tractography excels when tracing in homogenous white matter and can reconstruct many of the major white matter tracts. However, it can fail in areas in which fiber density is low or different fibers cross. These properties are exhibited by the porous heterogenous tissue in the basal forebrain that the amygdalofugal pathway traces through, as evidenced by invasive tracer injection and histology studies. It would be especially difficult to track in these circumstances using algorithms that do not utilize the full probability distribution of the fiber directions. Perhaps because of these difficulties, previous studies using tractography to investigate the limbic system might have overlooked or undervalued the prefrontal projections of the AmF

(Catani et al., 2013; Croxson et al., 2005; Kamali et al., 2016). In contrast, to overcome these challenges, we employed a strategy of first performing probabilistic tractography based on crossing-fiber resolved diffusion models in the macaque monkey to develop a protocol that was able to reconstruct the AmF fibers as known from previous invasive tracer studies in that species (Oler et al., 2017). This strategy was successful. We were able to reconstruct the course of the tract and confirmed the AmF as primarily a subcortical limbic tract, but with prefrontal fibers mostly extending along medial areas such as 25 and 14m. This pattern was similar in the human.

The division between AmF and UF bundles maps onto the tracer-based division of the ventral PFC into a medial region and a more lateral/orbital region as originally suggested by Price and colleagues (Carmichael and Price, 1995; Ongür et al., 2003). Following the logic that the connections of an area are a crucial determinant of its function, this suggested a different functional role of these subdivisions with the medial network strongly connected with visceromotor areas and the lateral orbital network interconnected with multi-modal sensory inputs (Ongür and Price, 2000). This was confirmed by dissociations apparent from lesion studies with medial frontal lesions showing an impairment in the capacity to maintain sustained arousal when anticipating rewards (Rudebeck et al., 2014), but temporal and ventral lateral frontal lesions leading to an impairment in learning arbitrary associations between stimuli and outcomes or stimuli and actions (Browning and Gaffan, 2008; Eacott and Gaffan, 1992; Gutnikov et al., 1997). Our results show divergences across AmF and UF in ventral PFC that supports this dissociation and demonstrate that it generalizes to the

human brain. Fibers in AmF and UF may constitute the primary white matter scaffolding from which these two networks may arise.

The importance of understanding a dissociation between AmF and UF may be further supported by considering the directionality of influence of these two bundles as described by more invasive investigations: while the UF is a bidirectional tract with many projections from prefrontal cortex terminating in medial and lateral basal nuclei of the amygdala, the AmF predominantly carries efferent projections from the amygdala to subcortical and prefrontal targets (Bickart et al., 2012). Potentially, this dissociation can help to inform whether abnormal neural processes central to certain psychiatric disorders, are more related to efferent or afferent projections of the amygdala (Bramson et al., 2019); Murray and Rudebeck, 2018; Volman et al., 2013). Unfortunately, one of the limitations of probabilistic tractography is that it cannot infer directionality of the reconstructed bundles and is unable to differentiate monosynaptic from polysynaptic connections. For this reason, the combination of tracer and imaging studies to translate results from macaque to human brains is an essential approach better understand human neuroanatomy in a non-invasive fashion.

It is important to note, however, that the dissociation between medial and lateral circuits is not complete, with the circuits joining up in anterior OFC and rostral prefrontal cortex. An integration between the two processing streams has been suggested both on anatomical and functional grounds. Anatomically, more rostral parts of prefrontal cortex are suggested to be more interconnected (Carmichael and Price, 1995). Functionally, such an integration would subserve learning and decision-making computations by

allowing the integration of learned stimulus-outcome contingencies with the internal state of the organism. This is evident in the proposal of a “common currency” representation comprising the value of different objects irrespective of the physical properties (Montague and Berns, 2002; Padoa-Schioppa and Assad, 2008). The rostral interactions of AmF and UF in areas 11, 13 and 14m may therefore represent the anatomical support for such computations.

It is important to point out that our use of diffusion-weighted MRI tractography necessitated some different protocol choices than a conventional tract-tracing study would do. A number of recent studies have criticized tractography for its perceived lack of reliability in areal to areal connections (Reveley et al., 2015; Donahue et al., 2016). Therefore, rather than replicating the tracer approach of seeding in circumscribed areas and assessing which grey matter regions are reached, we choose to seed in the core of the white matter bundles. Here, the tractography algorithm is less impaired and not affected by gyral bias or superficial white matter. Only in follow-up analyses do we investigate whether tractography seeded in these bundles reaches grey matter targets or whether grey matter seeded tractography reaches certain bundle. By placing the bundle at the core of the analyses, our results become more reliable. It also allows us to more easily create tractography recipes for the fiber bundles of interest that can then be shared with other researchers (e.g., Bramson et al., 2019).

The importance of considering neural circuitry, extending beyond individual brain regions, is increasingly evident when studying the contributions of ventral PFC in emotion, learning, and decision-making. Combining controlled focal lesion and transient perturbation studies with

whole-brain neuroimaging highlights the critical role of remote connections to the impact of local perturbations (Folloni et al., 2019; Froudust-Walsh et al., 2018; Verhagen et al., 2019). Moreover, it has recently been demonstrated that some lesion effects might be better explained by damage to connections as opposed to damage to single regions (Rudebeck et al., 2013b; Thiebaut de Schotten and Foulon, 2018) and differences in connectivity have already been argued to play a crucial role in regional specialization in other brain areas (Krubitzer, 2007). Given the inverse problem in which different changes in brain activity can have multiple causes (O'Reilly et al., 2013), an understanding of connectional anatomy is therefore essential to understand these network effects. As such, this work can help to move from a view of ventral PFC as a constellation of isolated, “functionally specialized” regions to a more network-based approach, thereby helping to more appropriately describe the underlying neurobiology of psychiatric disorders and promote effective treatment options.

Conventionally, surgical interventions, including deep brain stimulation for depression, addiction, and obsessive-compulsive disorders were primarily targeted at grey matter tissue. For instance, aspects of area 25 anatomy are associated with the presence of psychiatric disorders (Drevets et al., 1997; Murray et al., 2011). Indeed, chronic deep brain stimulation in this area improves chronic dysphoric conditions (Mayberg et al., 2005). However, with evidence accumulating that the most effective electrodes in deep brain stimulation are situated in or nearby white matter fiber tracts, rather than in grey matter proper, recent deep brain stimulation approaches are explicitly targeting white matter fiber structures (Baldermann et al., 2019; Lozano et

al., 2019). Considering the effective electrical stimulation of area 25 (Johansen-Berg et al., 2008), this region is innervated and visited by the nearby AmF, suggesting the hypothesis that the physiological effects may be primarily ascribed to the AmF pathway and the regions it interconnects. Riva-Posse and colleagues (2014) put forward the hypothesis that UF, forceps minor and cingulum bundle primarily mediate the therapeutic effects of DBS in depression. However, the WM region that they defined as UF and forceps minor is better represented by the course of the AmF, as shown in this Chapter. Moreover, the authors did not reconstruct the AmF in their study, which may be due to the challenges in tracking the complex course of this tract. This latter limitation, therefore, makes it difficult to completely exclude the AmF from their conclusions.

For this study we used high quality diffusion-weighted MRI data, obtained by scanning ex-vivo macaque samples for multiple hours, and by using in-vivo human diffusion MRI data with an acquisition duration of one hour. Both these datasets relied on unconventionally strong diffusion gradients, high number of diffusion directions, and high spatial resolution. However, the tractography recipes developed as part of this study have already been shown to be applicable and successful in human diffusion MRI data of more conventional quality. In fact, the approach and results achieved here, including the division between these two amygdala-prefrontal tracts, are behaviorally relevant in the context of social-emotional actions, advancing our understanding of the neurobiology of emotion regulation (Bramson et al., 2019).

To conclude, two principles underlie the organization of amygdala-PFC connectivity in the human brain. They are both evolutionarily conserved and shared with the macaque monkey. First a “connectional segregation principle” means that AmF and UF occupy a relative mediolateral position with respect to one another. Second, a “connectional integration principle” means that the tracts merge together as they enter and progress through the WM in the PFC to create an intricate network of projections from both cortical and subcortical regions. The two connectional systems described here may support distinct but synergic brain functions in PFC and this hypothesis should be investigated in future to understand how regions in the temporal lobe and PFC interact during affective, social, and cognitive tasks, how they could be disturbed in psychiatric disorders, and how this can be alleviated using targeted brain stimulation approaches.

Chapter 3

Manipulation of subcortical and deep cortical activity in the primate brain using transcranial focused ultrasound stimulation

3.1 ABSTRACT

The causal role of an area within a neural network can be determined by interfering with its activity and measuring the impact. Many current reversible manipulation techniques have limitations preventing their application particularly in deep areas of the primate brain. Here we demonstrate focused transcranial ultrasound stimulation (TUS) protocol impacts activity even in deep brain areas: a subcortical brain structure, the amygdala (experiment 1), and a deep cortical region, anterior cingulate cortex (ACC, experiment 2), in macaques. TUS neuromodulatory effects were measured by examining relationships between activity in each area and the rest of the brain using functional magnetic resonance imaging (fMRI). In control conditions without sonication, activity in a given area is related to activity in interconnected

regions but such relationships are reduced after sonication, specifically for the targeted areas. Dissociable and focal effects on neural activity could not be explained by auditory confounds.

3.2 INTRODUCTION

To establish the functional role of a brain area it is necessary to examine the impact of disrupting or altering its activity. It has recently been proposed that this might be accomplished with low-intensity focused transcranial ultrasound stimulation (TUS) (Tufail et al., 2011; Yoo et al., 2011). While a TUS impact on behavior has been described (Deffieux et al., 2013), little however, is known about its impact on neural activity and if its effects persist after stimulation is terminated. Moreover, current reversible and minimally-invasive neuromodulatory techniques are limited in their capacity to directly reach subcortical structures or areas located deep below the cortical surface of the brain. We show here that in the macaque, TUS modulates neural activity and does so even in subcortical nuclei such as the amygdala and deep cortical regions such as anterior cingulate cortex (ACC). Moreover, we demonstrate a protocol that exerts an “offline” effect that lasts for an extended period of tens of minutes after an initial stimulation period of 40s. This extended period of action is important because it means that its neural effect substantially outlasts any potential direct acoustic or somatosensory effects that might occur during the stimulation period itself (Guo et al., 2018; Sato et al., 2018). We also confirm this by showing that the stimulation protocol was not

associated with any similarly sustained impact on the activity of the auditory system.

Finally we demonstrate that a considerable degree of focality is possible with TUS. The peak and extent of the TUS neuromodulatory effect closely matched those of the ultrasonic intensity as estimated by simulations of the acoustic wave propagation. When TUS is applied to amygdala its impact is most apparent in amygdala rather than in more distal regions or those between the stimulation cone and the target area. The same is true of ACC TUS; its impact is most apparent in ACC where the acoustic intensity is highest. The focal impact of offline TUS in deep brain structures may underlie the specific patterns of behavioral impairment recently reported when the same protocol was used in awake behaving animals (Fouragnan et al., 2019).

3.3 MATERIALS AND METHODS

3.3.1 Experimental model and subject details

For this study, resting-state fMRI (rs-fMRI) and anatomical MRI scans were collected for 11 healthy macaques (*Macaca mulatta*, NCBITaxon:9544, nine males, two females, mean age: 7.3 years, mean weight: 10.3 kg). Resting state fMRI from four animals were acquired post amygdala TUS (experiment 1; n=4); rs-fMRI from three animals were acquired post ACC TUS (experiment 2; n=3); and rs-fMRI from nine animals were acquired without stimulation (control; n=9). All animals were purchased from a UK breeding center (Centre for Macaques, Porton Down, UK). All animals were socially

housed and kept under a 12:12 light/dark cycle. No water or food regulation was needed for the conduction of this project. All procedures were conducted under licenses from the United Kingdom (UK) Home Office in accordance with The Animals (Scientific Procedures) Act 1986. In all cases they complied with the European Union guidelines (EU Directive 2010/63/EU).

All suitable animals available at the time of experimentation took part in this study. Accordingly, there was no pre-selection nor restriction for group allocation. No suitable and available datasets were excluded: data from all TUS sessions where the sonication could be focused on the target coordinates were included. Sample sizes could not be predetermined statistically in the absence of a prior literature reporting relevant expected effect sizes; instead we adopted sample sizes similar to those reported in previous publications detailing interventional macaque functional magnetic resonance imaging (fMRI) studies (O'Reilly et al., 2013). Data collection and analysis were not performed blind to the conditions of the experiments.

3.3.2 Ultrasound stimulation

A single-element ultrasound transducer (H115-MR, diameter 64 mm, Sonic Concept, Bothell, WA, USA) with a 51.74 mm focal depth was used with region-specific coupling cones filled with degassed water and sealed with a latex membrane (Durex) to assess TUS of amygdala (experiment 1) and ACC (experiment 2) (**Figure 1**). The ultrasound wave frequency was set to the 250 kHz resonance frequency and 30 ms bursts of ultrasound were generated every 100 ms (duty cycle 30%) with a digital function generator (Handyscope HS5, TiePie engineering, Sneek, the Netherlands). Overall, the

stimulation lasted for 40s. These TUS parameters were chosen to ensure an efficient targeting of the regions of interest making and to prevent potential distortions of the ultrasound waves that may otherwise been caused by their interaction with the bone tissue as they cross through the skull. Moreover, the offline application of this protocol allowed us to rule out the potential influence of acoustic and somatosensory stimulation from the resulting neural effects of TUS. A 75-Watt amplifier (75A250A, Amplifier Research, Souderton, PA) was used to deliver the required power to the transducer. A TiePie probe (Handyscope HS5, TiePie engineering, Sneek, The Netherlands) connected to an oscilloscope was used to monitor the voltage delivered. The recorded peak-to-peak voltage was constantly maintained throughout the stimulation. Voltage values per session ranged from 128 to 134 V. It corresponded to a peak negative pressure ranging from 1.15 to 1.27 MPa respectively as measured in water with an in house heterodyne interferometer (Constans et al., 2017). The acoustic wave propagation of our focused ultrasound protocol was simulated at 130 V peak-to-peak voltage using finite element models of an entire monkey head to obtain estimates for the pressure amplitude, peak intensity, and spatial distribution (Constans et al., 2017). 3D maps of the skull were extracted from a monkey CT scan (0.36 mm isotropic resolution). Based on these numerical simulations, the maximum spatial peak pulse average intensity (I_{sppa}) in a focal region was estimated to be 64.9 W/cm² (spatial peak temporal average intensity (I_{spta}) = 19.5 W/cm²) in the amygdala and 18.8 W/cm² (I_{spta} = 5.63 W/cm²) in ACC with a maximum pressure of 1.44 MPa in amygdala and 0.78 MPa in ACC. One train was applied to each of the more laterally situated amygdalae but a

single train was applied to the midline structure (ACC) in experiments 1 and 2 respectively.

Each individual animal's structural MR image was registered to its head with a frameless stereotaxic neuronavigation system (Rogue Research, Montreal, CA). By recording the positions of both the ultrasound transducer and the head with an infrared tracker it was then possible to co-register the ultrasound transducer with respect to the MRI scan of the brain to position the transducer over the targeted brain region, either ACC (MNI coordinates $x = 0$, $y = 15$, $z = 6$) or amygdala (MNI coordinates $x = -10$, $y = 1$, $z = -11$; $x = 9$, $y = 1$, $z = -11$). The ultrasound transducer / coupling cone montage was placed directly onto previously shaved skin on which conductive gel (SignaGel Electrode; Parker Laboratories Inc.) had been applied to ensure ultrasonic coupling between the transducer and the animal's head. In the non-stimulation condition (control), all procedures (anesthesia, pre-scan preparation, fMRI scan acquisition and timing), with the exception of actual TUS, mirrored the TUS sessions.

3.3.3 Macaque MRI acquisition

Resting state fMRI and anatomical MRI scans were collected under inhalational isoflurane anesthesia using a protocol which has previously proven successful (Neubert et al., 2015; Sallet et al., 2013) in preserving whole-brain functional connectivity as measured with BOLD signal. In the case of the TUS conditions, fMRI data collection began only after completion of the TUS train (delay between ultrasound stimulation offset and scanning onset: 37.5 minutes; SEM: 2.21 minutes). Anesthesia (Table 1) was induced

using intramuscular injection of ketamine (10 mg/kg), xylazine (0.125-0.25 mg/kg), and midazolam (0.1 mg/kg). Macaques also received injections of atropine (0.05 mg/kg, intramuscularly), meloxicam (0.2 mg/kg, intravenously). The anesthetized animals were placed in a sphinx position and placed in a horizontal 3T MRI scanner with a full-size bore. Scanning commenced 1.53 hours (SEM: 4 minutes) and 2.38 hours (SEM: 4 minutes) after anesthesia induction in TUS and control sessions, respectively. In both cases data collection commenced when the clinical peak of ketamine had passed. Anesthesia was maintained, in accordance with veterinary recommendation, using the lowest possible concentration of isoflurane to ensure that macaques were anesthetized. The depth of anesthesia was assessed and monitored using physiological parameters (heart rate and blood pressure, as well as clinical checks before the scan for muscle relaxation). During the acquisition of the functional data, the inspired isoflurane concentration was in the range 0.8–1.1%, and the expired isoflurane concentration was in the range 0.7-1%. Isoflurane was selected for the scans as it has been demonstrated to preserve rs-fMRI networks (Neubert et al., 2015; Jérôme Sallet et al., 2013; Vincent et al., 2007). Macaques were maintained with intermittent positive pressure ventilation to ensure a constant respiration rate during the functional scan, and respiration rate, inspired and expired CO₂, and inspired and expired isoflurane concentration were monitored and recorded using VitalMonitor software (Vetronic Services Ltd.). Core temperature and peripheral capillary oxygen saturation (SpO₂) were also constantly monitored throughout the scan.

A four-channel phased-array coil was used for data acquisition (Dr. H. Kolster, Windmiller Kolster Scientific, Fresno, CA, USA). Whole-brain BOLD fMRI data were collected from each animal for up to 78 minutes. All fMRI data were collected using the following parameters: 36 axial slices; in-plane resolution, 2 x 2 mm; slice thickness, 2 mm; no slice gap; TR, 2000 ms; TE, 19 ms; 800 volumes per run. A minimum period of 10 days elapsed between sessions.

A structural scan (average over up to three T1-weighted structural MRI images) was acquired for each macaque in the same session as the functional scans, using a T1-weighted magnetization-prepared rapid-acquisition gradient echo sequence (0.5 x 0.5 x 0.5 mm voxel resolution).

3.3.4 Macaque MRI preprocessing

The preprocessing and analysis of the MRI data (Verhagen et al., 2019) was designed to follow the HCP Minimal Processing Pipeline (Glasser et al., 2013), using tools of FSL (Jenkinson et al., 2012; <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki>), HCP Workbench (<https://www.humanconnectome.org/software/connectome-workbench>), and the Magnetic Resonance Comparative Anatomy Toolbox (MrCat; www.neuroecologylab.org).

The T1w images were processed using tools of FSL in an iterative fashion cycling through brain-extraction (BET), RF bias-field correction, and linear and non-linear registration (FLIRT and FNIRT) to the *Macaca mulatta* F99 atlas (Van Essen and Dierker, 2007). The application of robust and macaque-optimized versions of BET and FAST also provided segmentation

into grey matter, white matter, and cerebral spinal fluid compartments. Segmentation of subcortical structures was obtained by registration to the D99 atlas (Reveley et al., 2017).

The first 5 volumes of the functional EPI datasets were discarded to ensure a steady RF excitation state. EPI timeseries were motion corrected using MCFLIRT. Given that the animals were anesthetized and their heads were held in a steady position, any apparent image motion, if present at all, is caused by changes to the B0 field, rather than by head motion. Accordingly, the parameter estimates from MCFLIRT can be considered to be 'B0-confound parameters' instead. Each timeseries was checked rigorously for spikes and other artefacts, both visually and using automated algorithms; where applicable slices with spikes were linearly interpolated based on temporally neighboring slices. This procedure identified an epoch with strong noise contributions at the end of the last run of a single rs-fMRI session following amygdala TUS. Accordingly, for one animal the last 254 volumes (out of a total of 2,400) were removed from further analysis. Brain extraction, bias-correction, and registration was achieved for the functional EPI datasets in an iterative manner, similar to the preprocessing of the structural images with the only difference that the mean of each functional dataset was registered to its corresponding T1w image using rigid-body boundary-based registration (FLIRT). EPI signal noise was reduced both in the frequency and temporal domain. First, the functional time series were high-pass filtered at 2000s. Temporally cyclical noise, for example originating from the respiratory apparatus, was removed using band-stop filters set dynamically to noise peaks in the frequency domain. Remaining temporal noise was described by

the mean time course and the first two subsequent principal components of the white matter (WM) and cerebral spinal fluid (CSF) compartment (considering only voxels with a high posterior probability of belonging to the WM or CSF, obtained in the T1w image using FAST). The B0-confound parameter estimates were expanded as a second degree Volterra series to capture both linear and non-linear B0 effects. Together the WM and CSF expanded B0 confound parameters were regressed out of the BOLD signal for each voxel.

The cleaned time course was then low-pass filtered with a cut-off at 10 seconds. The cleaned and filtered signal was projected from the conventional volumetric representation (2mm voxels) to the F99 cortical surface (~1.4mm spaced vertices) using Workbench command “myelin-style” mapping, while maintaining the subcortical volumetric structures. The data was spatially smoothed using a 3mm FWHM gaussian kernel, while taking into account the folding of the cortex and the anatomical boundaries of the subcortical structures. Lastly, the data were demeaned to prepare for functional connectivity analyses.

To represent subject effects, the timeseries from the three runs were concatenated to create a single timeseries per animal per intervention (control, ACC TUS, amygdala TUS). To represent group effects the run-concatenated timeseries of all animals were combined using a group-PCA approach (Smith et al., 2014) that was set to reduce the dimensionality of the data.

3.3.5 Macaque rs-fMRI connectivity analysis

Although the blood oxygen level dependent (BOLD) signal recorded with fMRI does not provide an absolute measure of activity it does provide a relative measure of activity change in relation to external events or activity recorded from other brain areas. This means that we cannot easily use BOLD to capture a measure such as activity in a brain area averaged over time. However, what we can do is to examine how BOLD responses in one area, such as the one that we are sonicating, relate to BOLD in another area using approaches similar to those employed previously (Margulies et al., 2009; Neubert et al., 2015; Sallet et al., 2013; Vincent et al., 2007).

To construct a region-of-interest (ROI) for ACC, a circle of 4mm radius was drawn on the cortical surface around the point closest to the average stimulation coordinate (**Figure 1**), in both the left and the right hemisphere. The same procedure was used to define other bilateral cortical regions of interest, based on literature coordinates (Neubert et al., 2015; Sallet et al., 2013; Neubert et al., 2014), to serve as targets for the fingerprint and spatial extent analyses (**Figure 2**, sub-panel i). The amygdala ROI was constructed for each animal individually through non-linear registration of their T1w image to the D99 template and by subsequently resampling the (subcortical) D99 macaque atlas in native space (Reveley et al., 2017).

In order to make a statistical comparison of the functional coupling of the amygdala or ACC in the control and TUS conditions it is problematic to compare coupling at each and every other point in the brain because there is a risk of false positive effects if multiple comparisons are made. Given the limited sample sizes possible with non-human primate experiments,

however, there is a risk of false negative results if stringent correction for multiple comparisons is undertaken at the whole-brain level. Indeed, here we avoid these pitfalls and reproduce whole-brain functional connectivity maps unthresholded to report on the full extent of the effects. Importantly, statistical inference was drawn on a limited set of regions beyond the amygdala known to be interconnected with macaque amygdala or ACC from anatomical tracing studies (Amaral and Price, 1984) and to exhibit, again in macaques, activity coupling with the amygdala under anesthesia (Neubert et al., 2015). An additional consideration was that some of the areas were connected to (Van Hoesen et al., 1993) and exhibited activity coupling with ACC (Neubert et al., 2015). Finally some areas, such as primary motor cortex (M1) and posterior intraparietal sulcus (pIPS) were chosen because, by contrast, they have limited connections and coupling with amygdala or ACC.

Coupling between the activity of each region of interest and the rest of the brain was estimated by calculating the Fisher's z-transformed correlation coefficient between each point in the ROI and all other datapoints. The resulting 'connectivity-maps' were averaged across all points in the ROI, across both hemispheres. Accordingly, the final maps represent the average coupling of a bilateral ROI with the rest of the brain. The fingerprints are obtained by extracting the average coupling with each target ROI and averaging across the two hemispheres.

To assess the impact of TUS not only for the target areas but across the whole brain, we indexed for every point in the brain its activity coupling with the same *a priori* defined constellation of regions used throughout the analyses (**Figure 2G** and **2H**), but now excluding the sonicated areas. We

quantify the impact of TUS by comparing for each point the average coupling with this set of regions in the control condition with the coupling observed following amygdala TUS and following ACC TUS. This approach resulted in two ‘heat-maps’ that show the peak location and extent of the brain activity impacted by TUS over amygdala and ACC.

Independently of the nature of the mechanisms underlying TUS (see Dallapiazza et al., 2018; Fomenko et al., 2018; Tyler et al., 2018), if TUS affects brain activity in a specific manner then it would be expected that the normal relationship seen at rest between the activity in the stimulated region and activity elsewhere will change. This does not mean that activity induced by TUS is diffused across the brain or that it is induced in one area and then “spreads” to others. It means simply that the relationship between activity in one area and another is changing. Measurements of activity throughout any area of tissue that is similarly affected by the TUS may become more highly correlated with one another. However, if the stimulated tissue becomes less responsive to other inputs from elsewhere in the brain then the relationship between activity in the stimulated region and elsewhere will decrease.

3.3.6 Macaque rs-fMRI statistical inference

Statistical inference on the fingerprints was performed using non-parametric permutation tests on cosine similarity metrics describing how similar or dissimilar pairs of fingerprints are (Mars et al., 2016b; Verhagen et al., 2019). The cosine similarity metric considers the shape of the fingerprint as a whole (but not its mean amplitude) and performs one test per pair of fingerprints, negating the necessity for correcting for multiple comparisons across

fingerprint targets. In contrast to conventional parametric tests, this approach does not rely on assumptions about the shape of the distribution but will acknowledge dependencies between target ROIs in the fingerprint; as such this approach will avoid inflation of type I error. For each test we ran 10,000 permutations across individual fMRI runs to accurately approximate with high accuracy the true probability of rejecting the null-hypothesis of permutable conditions in this sample.

To examine the spatial extent of the neuromodulatory impact of TUS on activity coupling we extracted for every point in the brain, both subcortically and on the cortical surface, its average coupling strength with the fingerprint targets (**Figure 2**, sub-panel i), excluding the amygdala and ACC. This approach allowed the creation of a quantified spatial map of the difference in average coupling between the control state and amygdala TUS, and between the control state and ACC TUS. For regions affected by TUS this difference will be large, while for all other regions this difference is close to zero.

Statistical inferences on the anesthesia levels and associated physiological parameters were drawn in the context of generalized linear mixed-effects (GLME) models. These models considered the intercept, the TUS condition (control, amygdala, ACC), and the resting-state fMRI run index (1, 2, or 3) as fixed effects and the intercept and slope grouped per animal as random effects with possible correlation between them (as implemented in MATLAB, Mathworks, Natick, USA). The models were assumed to adhere to a normal distribution of the data and were fitted using Maximum-Pseudo-Likelihood estimation methods where the covariance of

the random effects was approximated using Cholesky parameterization. Statistical significance was set at $\alpha = 0.05$, two-tailed, and estimated using conventional analyses of variance (ANOVA).

TABLE 1		
REAGENT or RESOURCE	SOURCE	IDENTIFIER
Chemicals, Peptides, and Recombinant Proteins		
Isoflurane – ISOFLU 250ml	Centaur	30135687
Ketamine – Narketan 10% 10ml INJ CD(SCH4)1 1-MCD	Centaur	03120257
Midazolam – Hypnoval amps 10mg/2ml	Centaur	23191407
Atropine – Atrocare INJ 25ml	Centaur	01300236
Meloxicam – Metacam INJ 10ml 5mg/ml DOGS/CATS	Centaur	02500456
Ranitidine 50mg/2ml x5 INJ	Centaur	30294115
Saline	DPAG, University of Oxford	N/A
Formalin	DPAG, University of Oxford	N/A
SignaGel Electrode Gel	Parker Laboratories	#15-25
Experimental Models: Organisms/Strains		
Macaca mulatta, 9 males, 2 females, between 4- 11 years old, between 7-15 kg, socially housed	MRC, Centre for Macaques	NCBITaxon:9544
Software and Algorithms		
MATLAB 2017a	Mathworks	RRID: SCR_001622
FMRIB Software Library v5.0	FMRIB, WIN, Oxford, UK	RRID:SCR_0028 23
Connectome Workbench	The Human Connectome Project and Connectome Coordination Facility	RRID:SCR_0087 50
Magnetic Resonance Comparative Anatomy Toolbox	Neuroecology Lab	https://github.com /neuroecology/MrCat
Other		
Transducer H-115MR 250kHz SN:018	Sonic Concepts	www.sonicconcepts.com
Transducer H-115MR 250kHz SN:017	Sonic Concepts	www.sonicconcepts.com
Amplifier Model 75A250A – 75Watts – 10khz 250MHz	Amplifier Research	www.arworld.us
Tie Pie Handyscope HS5 SN: 32239	Tie Pie	https://www.tiepie.com
Brainsight frameless stereotaxic neuronavigation system	Rogue Research	RRID:SCR_0095 39
MRI compatible frame	Crist Instruments	http://www.cristinstrument.com/products/stereotax/stereotax-primate
four-channel phased-array coil	Windmiller Kolster Scientific	https://www.wkscientific.com/#mri-coils

3.4 RESULTS

3.4.1 Stimulation of deep brain structure and resting-state fMRI recording

On each session of TUS application, a 40s train of pulsed ultrasound (250 kHz) comprising 30ms bursts of ultrasound every 100ms was directed to the target brain region using a single-element transducer in conjunction with a region-specific coupling cone filled with degassed water. To control for any confounds resulting from concomitant ultrasound stimulation and neural signal recording (Guo et al., 2018; Sato et al., 2018), recordings of neural activity only begun approximately 30 minutes after the end of TUS application when any potential auditory or somatosensory effects of stimulation were dissipated. We therefore refer to this stimulation protocol as an “offline” protocol.

Frameless stereotaxic neuronavigation was used to position the transducer over the target brain area taking into consideration the focal depth of the sonication (**Figure 1**; experiment 1: amygdala n=4; experiment 2: ACC n=3; relatively deep brain regions known to be interconnected and co-active during similar cognitive processes such as social cognition (Munuera et al., 2018; Noonan et al., 2014). A single train was applied sequentially to each amygdala, in experiment 1 and to the midline structure, ACC, in experiment 2.

The impact of TUS was determined by examining brain activity over an 80-minute period starting approximately 30 minutes after the 40s stimulation train began (see Materials and Methods). Activity was recorded not just from the stimulated site but from across the entire brain using

functional MRI (fMRI). FMRI data from the stimulated animals were compared with data from an additional group of control individuals (n=9) that had received no TUS. Note that depth of anesthesia and the delay between sedation induction and data acquisition were similar between the TUS and the control groups (0.7-0.8% and 0.7-1% range of expired isoflurane concentration, 1.53 and 2.38 hours, respectively; Materials and Methods). FMRI data were acquired at 3T under isoflurane anesthesia and processed using established tools and protocols (Verhagen et al., 2019; see Methods). The anesthesia protocol has previously been shown to preserve regional functional connectivity measurable with fMRI (Sallet et al., 2013; Neubert et al., 2015).




Figure 1. Stimulation targets. Stimulation target position is shown for each individual animal (colored dots) on sagittal and coronal views for TUS targeted at amygdala (A, B) and ACC (D, E). Acoustic intensity field (W/cm^2) generated by the ultrasound beam in the brain is shown for one example animal per TUS target, amygdala (C) and ACC (F). The target position can be delineated with accuracy in all animals in panels A, B, D, and E by using each individual's own MRI scan. As a result, the activity and functional connectivity of the target areas can be examined accurately in each animal (see subsequent figures). However, some slight imprecision in the estimation in the acoustic intensity maps in panels c and f may occur; this is because group average targets are used in conjunction with the computed tomography X-ray scan of a single individual during the modelling.

3.4.2 Effects of TUS on subcortical neural activity in the amygdala

To examine the spatial specificity of TUS effects and to investigate the capacity of TUS to stimulate subcortical structures we investigated its effects on the coupling of amygdala activity with activity in other brain areas. Even at rest in the control state, BOLD activity in one area is correlated with BOLD activity in other areas and such relationships are most prominent when the areas are monosynaptically connected, although some residual connectivity is mediated by indirect connections (O'Reilly et al., 2013). The pattern of activity coupling for any given area reflects its unique constellation of projections and interactions, sometimes called its “connectivity fingerprint” (Passingham et al., 2002).

In the control state, amygdala activity was coupled with activity in cingulate, ventral prefrontal, and orbitofrontal cortex, striatum, and the

anterior temporal lobe (**Figure 2A** and **2G**). To limit the risk of false positive/negative results, we focused our analysis on a limited set of brain regions known to be interconnected with the amygdala based on previous studies (Neubert et al., 2015; Neubert et al., 2014; Sallet et al., 2013) and compared the overall patterns or “fingerprint” of coupling (amygdala TUS vs. control) using cosine similarity metrics in a non-parametric statistical framework (see Methods for details).

The amygdala’s activity coupling was significantly changed after amygdala TUS (non-parametric permutation test, $p = 0.0020$; **Figure 2B** and **2G**). A whole-brain quantitative analysis revealed that this effect of amygdala TUS was most apparent in the amygdala, and not anywhere else in the brain (see ‘Focality of TUS’, **Figure 3A**).

A second way to establish the specificity of TUS effects within the network is to examine whether the amygdala connectivity effects seen after amygdala TUS are found after ACC TUS. This was not the case; ACC TUS left most of amygdala’s coupling pattern unaffected (non-parametric permutation test, $p = 0.1346$; **Figure 2C** and **2G**) although not surprisingly ACC TUS led to alteration in amygdala’s coupling with ACC.

Finally, to further establish the nature of amygdala TUS effects within the network we investigated the activity coupling patterns of five further control areas. We investigated three regions adjacent to the amygdala and found their functional connectivity was unaltered (**Figure 3B**). We also examined an area with a very distinct constellation of projections – ventral premotor area F5c – and again found no change (see **Figure 4**). Similarly, analyses of the temporal variability of the BOLD signal (see **Figure 5M**) did

Figure 2. Amygdala and ACC functional coupling changes after stimulation. (A-C) Activity coupling between amygdala (in yellow on the coronal view) and the rest of the brain in no stimulation/control condition (A), after amygdala TUS (B), and after ACC TUS (C). (D-F) Activity coupling between ACC (outlined in black) and the rest of the brain in no stimulation/control condition (D), after amygdala TUS (E), and after ACC TUS (F). Hot colors indicate positive coupling (Fisher's z). Functional connectivity from TUS-targeted regions are highlighted by black boxes. Each type of TUS had a selective effect on the stimulated area: amygdala coupling was strongly changed by amygdala TUS only (B) and ACC coupling was strongly changed by ACC TUS only (F). (G) Connectivity fingerprint representation of the strength of activity coupling between amygdala and other brain areas in control animals (blue), after amygdala TUS (yellow), and after ACC TUS (red). (H) Activity coupling between ACC and the rest of the brain in control animals (blue), after ACC TUS (red), and after amygdala TUS (yellow). Each type of TUS had a selective effect on the stimulated area: amygdala coupling was strongly affected by amygdala TUS (the yellow line is closer to the center of the panel than the blue line) and ACC coupling was strongly disrupted by ACC TUS (the red line is closer to the center of the panel than the blue line). Standard error of the mean is indicated by shading around each line. (I) The regions-of-interest constituting the fingerprints depicted on lateral, medial, orbital, and dorsal views.

3.4.3 Effects of TUS on deep cortical neural activity in anterior cingulate cortex (ACC)

To examine the specificity of TUS effects further and to investigate the capacity of TUS to stimulate deep cortical structures we investigated the effects of ACC TUS on ACC activity. In control animals, ACC activity at rest was coupled with activity in strongly connected areas: dorsal, lateral, and orbital prefrontal cortex (PFC), frontal pole, mid and posterior cingulate

(**Figure 2D** and **2H**). After ACC TUS the ACC coupling pattern was altered (non-parametric permutation test, $p = 0.0210$; **Figure 2F** and **2H**). A parsimonious interpretation is that normally the activity that arises in ACC is a function of the activity in the areas that project to it, but this is no longer the case when ACC's activity is artificially driven or diminished by TUS. Because these interactions with other areas determine the information ACC receives from elsewhere in the brain and the influence it exerts over other areas, ACC TUS should alter ACC's computation and induce specific changes in behavior (Fouragnan et al., 2019).

Similar to the analyses of spatial extent of amygdala TUS effects, we quantified the change in coupling induced by ACC TUS not only in ACC itself, but for every point in the brain. This analysis revealed that ACC TUS affected primarily the ACC (See 'Focality of TUS', **Figure 3A**).




Figure 3. Spatial extent of the TUS neuromodulatory effect and its impact on areas neighboring the stimulated region. (A) Amplitude and spatial extent of the impact of amygdala TUS (top row) and ACC TUS (bottom row) on the coupling of each point in the brain with the same set of a priori defined areas used in Figure 2G and 2H. Hot colors indicate a strong decrement in coupling after TUS compared to the control state (delta Fisher's z). The effect of TUS on activity coupling was restricted to the amygdala after amygdala TUS (top row) and to the ACC and regions immediately ventral along the ultrasound trajectory following ACC TUS (bottom row). (B) The whole brain coupling of the amygdala target region (i, also shown in Figure 2) and regions along (ii,iii) or immediately surrounding the trajectory of the ultrasound stimulation beam (iv), in the control condition and after amygdala TUS. There were no major changes in the coupling of these off-target regions and the rest of the brain. (C) The whole brain coupling of the ACC target region (i) and surrounding regions near the ACC target (ii-v) in both the ACC TUS and control conditions. Some changes in coupling can be seen along the stimulation trajectory in the area just ventral to the target (v) and also in an area which is unlikely to have been hit directly by the ultrasound beam (iv). These areas are strongly anatomically connected with the targeted area. Seed regions in panels B and C are indicated with black outlines.

The specificity and selectivity of the effects are further underscored by the results observed when mapping the coupling pattern of areas interconnected with the stimulated ACC region. First, we examined the activity coupling pattern of the amygdala – an area with which ACC is monosynaptically interconnected (Amaral and Price, 1984; Van Hoesen et al., 1993) and functionally coupled (Neubert et al., 2015). Not surprisingly, there was some evidence that amygdala-ACC coupling had changed as a function of ACC TUS, as had coupling with a third area – caudal orbitofrontal

cortex – with which both ACC and amygdala are strongly interconnected. However, other aspects of the amygdala’s coupling pattern were relatively unaltered by ACC TUS; although there was a trend in the non-parametric permutation test for amygdala connectivity to change after ACC TUS ($p = 0.0744$; **Figure 2C** and **2H**) it was clear that there was a significant difference between ACC and amygdala TUS effects (non-parametric permutation test, $p = 0.0428$).

Just as amygdala TUS did not affect the connectional profile of F5c, a region outside the interconnected network of the stimulated areas, ACC TUS also did not affect F5c’s coupling (see **Figure 4**). Again, below, we explain additional control analyses that confirmed that the TUS effect could not have been mediated via the auditory cortex.

The spatial and connectional specificity of the observed effects make it unlikely that the TUS induced modulations were mediated by general physiological effects, such as those related to anesthesia level and duration (see Methods for details).




Figure 4. Whole-brain functional connectivity between stimulated and not stimulated areas with the rest of the brain. Panels a, b, and c show activity coupling between a control area, the caudal ventral premotor area F5c, and the rest of the brain in no stimulation/control condition (**A**), after amygdala TUS (**B**), and after ACC TUS (**C**). Hot colors indicate positive coupling (Fisher's z). Compared to a no stimulation condition (A), neither amygdala TUS nor ACC TUS (B, C) affected the whole-brain coupling activity of F5c which has weak anatomical connections with ACC and amygdala.

3.4.4 Focality of TUS effect

To examine the focality of the TUS's effect on activity coupling two additional sets of analyses were conducted. The first set of analyses assessed the impact of TUS on the connectivity fingerprint not only for the target areas but for every point in the brain resulting in 'heat-maps' of TUS's impact, while the second set focused in detail on the areas surrounding the target areas or located between the stimulation cone and the target area (**Figure 3**; see **Figure 6**).

In the first set of analyses, following amygdala TUS the strongest neuromodulatory effects were observed in the amygdala itself, and only in the amygdala (**Figure 3A**, top row). Following ACC TUS the extent of the neuromodulation was limited to the ACC and regions immediately ventral to it along the ultrasound beam (**Figure 3A**, bottom row). In fact, the spatial maps of TUS impact on activity coupling are strikingly in correspondence with the spatial maps of estimated sonication intensity (**Figure 1C** and **1F**). This correspondence is specific and sensitive: it includes particulars of the wave propagations, such as how the ultrasound wave targeted at amygdala

reflects on the basal bone, while in the ACC TUS condition considerable acoustic energy is also deposited immediately ventral to the target along the trajectory, partly due to sound waves reflecting on the orbital bone.




Figure 5. Temporal changes of TUS effects on amygdala and ACC functional coupling. Amygdala functional coupling (A,B,C) and ACC (D,E,F) functional coupling across 3 consecutive runs is displayed after no stimulation (amygdala: A,B,C; ACC: G,H,I) and TUS (amygdala: D,E,F; ACC: J,K,L). TUS effects on the whole-brain coupling of each stimulated region persisted throughout the full length of scanning. Interestingly, TUS effects seems to show slightly decrease over the three runs with amygdala and ACC functional coupling resembling more their correspondent coupling in the no stimulation condition. (M) Analysis of the dynamic range of the BOLD signal across all voxels in the control state (top row), after amygdala TUS (middle row), and after ACC TUS (bottom row) revealed similar levels of variance in activity in all three cases.

To further qualify and examine the extent of the ultrasonic intervention, in the second set of analyses we assessed the impact of TUS on whole-brain activity coupling of control areas surrounding the sonication target or along the trajectory of the ultrasound beam (**Figure 3B** and **3C**). Confirming the spatial maps of TUS impact (**Figure 3A**), and matching the estimated contours of the acoustic intensity (**Figure 1F**), there were no major changes in the activity coupling of areas situated between the target and the transducer (**Figure 3B**, sub-panels ii-iii and **Figure 3C**, sub-panel iii), while as estimated following ACC TUS there were some changes in the connectional profile of a region along the stimulation trajectory just ventral to the ACC target (**Figure 3C**, sub-panel v).

Finally, we examined whether areas outside the directly sonicated region, but strongly connected to it, might exhibit a network-derived effect of TUS. Changes to the connectional profile of an area just posterior (**Figure 3C**, sub-panel iv) to the target region (**Figure 3C**, sub-panel ii), could be suggestive of such a network-effect. Indeed, additional analyses of areas sharing the same spatial proximity but lacking the same anatomical connectedness confirmed these changes were likely due to the anatomical connections this region shared with the target area (see **Figure 6**).




Figure 6. Effects of TUS on regions outside the target area are mediated by the strength of anatomical connectivity rather than a result of spatial proximity. Panels A, B, and C of the figure show whole-brain activity coupling in control and ACC TUS conditions for ACC (**A**) and two areas at an equal Euclidian distance from the ACC target region: SMA (**B**) and area 24ab (**C**). Hot colors indicate positive coupling (Fisher's z). The activity coupling of SMA, an area weakly connected with the target area ACC, and the rest of the brain is predominantly preserved after TUS. However, the whole-brain coupling of an area more strongly connected with the target ACC region (Van Hoesen et al., 1993), area 24ab, is more influenced by TUS to the ACC target region.

3.4.5 Putative auditory effects of offline TUS

It has recently been suggested that TUS's impact on neural activity is mediated by its auditory impact (Guo et al., 2018; Sato et al., 2018). Several

considerations suggest that it might not be possible to explain away the current findings as the result of an auditory artefact. First, the auditory impact of TUS is likely a function of specific features of its frequency and pulse type, and especially of the frequency used to modulate the ultrasonic carrier wave. Second, the auditory stimulation associated with the TUS application ceased after the sonication but the neural activity measurements were initiated tens of minutes later. Third, TUS of each area, ACC and amygdala, had specific effects that were distinct to one another.




Figure 7. Effect of amygdala and ACC TUS on the functional coupling of primary auditory cortex. (A) Mediation via the auditory cortex cannot explain the effects seen after either amygdala (yellow) or ACC (red) TUS. (B) ACC TUS (red line) had no effects on the functional coupling of A1. Amygdala TUS (yellow line) affected the relationship between A1's activity and activity in several areas that are linked to the A1 via the amygdala including the amygdala itself, lateral orbitofrontal cortex area 47/12o and ACC.

Nevertheless, we also carried out a fourth line of inquiry and examined whether it is plausible that an auditory effect could have mediated the TUS effects on amygdala and ACC. To quantify this probability, we correlated any TUS effects on primary auditory cortex (A1) connectivity with TUS effects on the targeted regions (**Figure 7A**). TUS effects on the auditory cortex after both amygdala ($r=0.1084$, $p=0.7007$) and ACC ($r=0.1474$, $p=0.6000$) sonication are unrelated to the TUS effects at each target site and are therefore unlikely to have mediated effects seen at the stimulation sites. However, it is possible that TUS over amygdala or ACC had an impact on A1 connectivity separately from its impact on the stimulated sites themselves (**Figure 7B**). While A1 connectivity is not impacted by ACC TUS (**Figure 7B**, non-parametric permutation test, $p=0.6871$), amygdala TUS did have a significant impact on A1 connectivity (**Figure 7B**, non-parametric permutation test, $p = 0.0002$). Closer inspection revealed that this was due to a diminution solely in A1's interactions with the amygdala itself and two areas with which the amygdala is itself strongly connected with: ACC and orbitofrontal cortex. Differential effects of ACC and amygdala TUS on A1 connectivity might be driven by some direct, albeit weak, connections of amygdala with A1 (Yukie, 2002). Similarly, given amygdala's strong connections to ACC and orbitofrontal cortex, it is perhaps not surprising that amygdala sonication might affect A1's interactions with them. Importantly, these circumscribed effects on A1 connectivity are not predictive of the effects elsewhere. In summary, the alteration seen in the A1 fingerprint is a poor match to the alteration seen in the amygdala fingerprint after amygdala TUS or in the ACC fingerprint after ACC TUS.

3.5 DISCUSSION

In these investigations we combined TUS with resting-state fMRI to examine the impact of modulating activity in subcortical and deep cortical areas of the primate brain. Experiments 1 and 2 revealed dissociable effects of amygdala and ACC TUS. The dissociable nature of the effects and the fact that they were observed more than an hour after the 40s stimulation period suggests they are not mediated by the stimulation's auditory impact (Guo et al., 2018; Sato et al., 2018). In each case effects were apparent as reductions in activity coupling between the stimulated area and other regions with which it is normally interconnected; after TUS, a brain area's activity appears to be driven less by activity in the areas with which it is connected and more so by the artificial modulation induced by TUS.

Any impact that TUS exerts on the auditory system is likely to depend on the precise details of the sonication frequency, pulse/modulation frequency, and pulse shape and might be specific to other features of the preparation such as anesthesia level (Airan and Pauly, 2018). Here we employed an ultrasound frequency of 250kHz that we pulse modulated at 10Hz: as such we ensured that both the ultrasound carrier wave and the wave envelope frequency are well outside of the macaque hearing range. This can be contrasted against more conventional protocols where the ultrasound is pulse modulated at ~1kHz, within the audible range of both rodents and primates. Moreover, the “offline” stimulation protocol we employed also made it less likely that the auditory system was stimulated at the time that neural activity was recorded; neural activity was only measured many minutes after the cessation of a 40s period of TUS.

The administration of certain doses of anaesthetic agents influence the neural signal. Although the anaesthesia protocol used in this experiment was effective and previously showed to preserve resting-state fMRI networks (Neubert et al., 2015; Neubert et al., 2014; O'Reilly et al., 2013; Sallet et al., 2013; Vincent et al., 2007), in a separate set of scans acquired in combination with TUS (*unpublished data*), we observed that the modulatory effects of TUS were harder to observe when the doses of anaesthesia were higher than those reported in Chapter 3. This could have been due to a (1) weaker BOLD signal, resulting from the higher doses of anaesthesia, or (2) by its direct interaction with the sound waves at the level of the cell membrane, thereby resulting into a reduction in the modulatory effects of TUS on the polarisation of the neurons. The former hypothesis is the most likely in these *unpublished datasets* since the higher doses of anaesthesia were administered during the acquisition phase and, therefore, only after TUS was applied. Future investigations on the mechanisms underlying TUS should try to elucidate how sound waves interact with different anaesthetic agents and what is the resulting impact on neural plasticity.

Our aim in the current study was to examine whether TUS can modulate neural activity. The results obtained demonstrate that TUS can exert a relatively focal and circumscribed impact on neural activity. However, as a consequence of using a recording technique that is sensitive to a number of neurophysiological processes, it was not possible to establish the precise nature of the neurophysiological process that mediated the fMRI signal effects that we observed. It is possible that TUS may act not simply by immediately inducing or reducing activity in neurons but by modulating their

responsiveness to other neural inputs; thus its effect may have been more easily detected by an analysis strategy such as the current one that focused on measuring the relationship between activity in the stimulated area and elsewhere. As with other repetitive neurostimulation protocols it is also possible that TUS's offline effects are partly driven by the induction of plastic changes, with long-term-potential/depression-like characteristics, and again this might have implications for how its effects are best detected.

Several molecular mechanisms describing how low-intensity ultrasound stimulation modulates neuronal activity have been suggested. However, recent investigations on the interactions between sound pressure waves and brain tissue suggest that ultrasound primarily exerts its modulatory effects through a mechanical action on cell membranes, notably affecting ion channel gating (Kubanek et al., 2016; Tyler et al., 2008). While the precise mechanisms are being determined (Fomenko et al., 2018; Tyler et al., 2018) the current results suggest TUS may be suitable as a tool for focal manipulation of activity in many brain areas in primates. Specifically, they show that TUS may even be used to manipulate activity in subcortical structures in monkeys.

TUS's capacity to stimulate subcortical and deep cortical areas in primates, therefore, opens the prospect of advanced non-invasive causal brain mapping. To date, non-invasive manipulation of brain activity in humans can be done reversibly only using neuromodulation methods such as transcranial magnetic stimulation and transcranial current stimulation. However, the spatial resolution of some of these techniques is limited (Polanía et al., 2018). Even more critically, application of these techniques is

constrained to the surface of the brain as their efficacy falls off rapidly with depth.

Before it becomes possible to use repetitive TUS to study the human brain in a routine manner, a number of considerations must be borne in mind. It will be important to establish the safety of the technique. In another recent study we have shown that TUS of the type used here causes no permanent damage to tissue on histological analysis (Verhagen et al., 2019). Structural MRI scans collected shortly after TUS in the present study showed no evidence of transient edema (see **Figure 8**). While such results are encouraging, further studies may be needed to establish if this remains true even after a greater number of TUS sessions, after TUS sessions of longer duration, or after TUS at a greater intensity. Care may need to be taken with the assessment of each new protocol that is devised. It should also be noted that its neural effects may be sustained over a period of time that is substantially longer than in many laboratory experiments (Verhagen et al., 2019); care will therefore need to be taken in deciding when a human participant might leave the laboratory and travel home. In addition, under some circumstances sonication appears to impact on the meninges (Verhagen et al., 2019) and the full nature of this impact may need to be established. Not only does this have safety implications but it also suggests that the impact of TUS to a brain area is best assessed by comparison to the impact of TUS to an appropriate control site.




Figure 8. Lack of effects of TUS on brain tissue. Coronal T1-weighted images collected before and after TUS (in this example targeted bilaterally to the amygdala) representing the stimulated brain area. Neither structural changes nor evidence of transient edema were found following TUS targeted to the amygdala bilaterally in two exemplar animals.

In summary, based on the results reported here, TUS can be used to transiently and reversibly alter neural activity in subcortical and deep cortical areas with high spatial specificity. To date, it is the most promising neuromodulatory technique to reach areas deep below the dorsolateral surface of the brain in a minimally invasive and focal manner, thereby providing it with the potential for causally mapping brain functions within and across species. While it may currently lack the capacity to target specific neurons, as do some optogenetic and chemogenetic techniques (Yizhar et

al., 2011), it may provide a method for investigating brain areas that may make it suitable for use with primate species, which are rarely investigated with such techniques even though many brain areas are particularly well developed or only present in primates (Passingham and Wise, 2012). With care it may even be possible to employ offline TUS protocols in investigations of human brain function.

Chapter 4

Causal role of the lateral orbitofrontal cortex in credit assignment

4.1 ABSTRACT

Lateral orbitofrontal cortex has been suggested to be a key area for “credit assignment” in the primate brain, that is activity in this region enables the formation of associations between choice options and outcomes during contingent learning. These contingencies are then used to guide future behaviour in value-guided learning tasks. Here, we used focused transcranial ultrasound stimulation to focally and reversibly disrupt activity in the lateral orbitofrontal cortex of macaque monkeys and causally test for the role played by this region in credit assignment. After ultrasound stimulation, animals consistently showed deficits in performing this computation as well as in their ability to update their options value estimates in order to adapt future decisions on the basis of the integrated past history of choice option-outcome contingencies. Conversely, credit assignment deficits were absent after disrupting activity within a neighboring area. Overall, these results suggest

that disruption of lateral orbitofrontal cortex specifically affects credit assignment in primates.

4.2 INTRODUCTION

Adaptation to a volatile environment requires being able to flexibly learn contingent associations between visual features of choice options and outcomes, also described as reward value, and using these learned associations to guide subsequent behavior. This type of reward-guided learning is often referred to as “credit assignment” and it is a process that allows animals to learn choice option-outcome contingencies and, thereby, to update the reward value of a choice option once its outcome has been witnessed. For instance, if choosing an option resulted in a reward, then the value assigned to the chosen option should be positively updated; the choice value should increase. As a result the likelihood of the animal choosing the same option in following trials should increase. By contrast if the choice is not rewarded, then the choice value should decrease and the animal should be less likely to take the choice in the future. This is sometimes referred to as the law of effect (Thorndike, 1911). In a stochastic environment, animals should, therefore, integrate the outcomes of a series of choices into a stimulus-specific value and use this value to guide future choices. Moreover, healthy animals should also be capable of crediting the correct choice for the outcome; it is important that the choice that actually led to the reward, rather than an earlier or later choice that happened to be made at approximately the same time, is credited with the outcome. Sometimes, however, we know that

the outcome is incorrectly linked not just to the contingent choice but to other choices made just before and after; such failures of credit assignment are sometimes referred to as “spread of effect” (Thorndike, 1933).

Orbitofrontal cortex (OFC) and adjacent ventral prefrontal cortex have been shown to play a key role in encoding stimulus-outcome associations as a function of the history of outcomes in past choices (Murray and Rudebeck, 2018; Noonan et al., 2010; Walton et al., 2010). Moreover, Walton and colleagues (2010) observed that animals with OFC lesions showed a deficit in correctly encoding value to a specific stimulus. After lesions to these areas, animals wrongly assigned the outcome not only to the option chosen in the current decision but also to other choice options chosen in recent and subsequent trials.

Walton and Noonan and colleagues made aspiration lesions in the OFC between the medial orbital and lateral orbital sulci (Noonan et al., 2010; Walton et al., 2010). Because the lesions were made by aspiration it is possible that they damaged not just the OFC itself but also adjacent white matter. It is known that white matter sparing lesions, such as those made by excitotoxin injection have more selective effects than aspiration lesions (Rudebeck et al., 2013b). Rudebeck and colleagues have reported that excitotoxic lesions to the same OFC region do not impair credit assignment but that excitotoxic lesions of ventrolateral prefrontal cortex do (Rudebeck et al., 2017b).

Functional Magnetic Resonance Imaging (fMRI) experiments suggest that the critical region may be in the border region between what is classical thought of as OFC and ventrolateral cortex in a region sometimes referred to

as lateral OFC (lOFC; Neubert et al., 2015; Noonan et al., 2011) or area 12o (Carmichael and Price, 1994; Chau et al., 2015). In conjunction fMRI and computational modelling provide powerful methods for showing correlations between behaviour and brain activity. For example parameters, such as value, associated with a choice, can be linked to activity in specific brain areas or networks. Chau and colleagues (Chau et al., 2015) suggested the key region might lie in the lOFC/area 12o because activity here changes whenever outcomes lead to appropriate value adjustments and changes in behavior. For example lOFC/12o is active when a choice leads to reward and the choice will be repeated again in the future and when a choice leads to no-reward and the choice will not be repeated again in the future. There is significantly less activity when the choice leads to reward but the choice will not be repeated again in the future and when a choice leads to no-reward and the choice will be repeated again in the future. However, to determine whether such activation is truly necessary for credit assignment, the impact of disrupting it must be assessed.

To address this issue, here we used offline focused Transcranial Ultrasound Stimulation (TUS) in 4 rhesus macaque monkeys while they performed a probabilistic three-arm reversal learning bandit task in which the values of the three options were uncorrelated with one another. TUS allows us to transiently and reversibly disrupt activity in a specific brain region to investigate its specific impact on behaviour. As this neuromodulatory approach does not require making a permanent lesion of a specific brain area, TUS allows us to test the same animal across the different conditions of interest (control, stimulation targeted to the region of interest, stimulation

targeted to a control area) and to perfectly alternate and counterbalance sessions from each of them. We have recently showed the spatial focality of the effects of TUS on neural activity (Folloni et al., 2019; Verhagen et al., 2019) and behaviour (Fouragnan et al., 2019). Control (sham) and TUS sessions were interspersed within each animal and their order was counterbalanced across animals. The ultrasound wave frequency was set to the 250 kHz resonance frequency and 30 milliseconds bursts of ultrasound were generated every 100 milliseconds. Overall, the stimulation lasted for 40 seconds and was applied to the lateral OFC in both hemispheres (experiment 1) and to another control region, the anterior frontal polar cortex (anterior PFC; experiment 2) immediately before the animals performed the task.

4.3 MATERIALS AND METHODS

4.3.1 Subjects

Four male rhesus monkeys (*Macaca mulatta*), here referred to as D, E, F and G, were tested in the experiment (mean weight 12.65 kg; 9.9 years of age). They were socially housed (group of 4 animals) and kept on a 12 hours light dark cycle, with access to water for 12-16 hours on testing days and with *ad libitum* water otherwise. Prior to task training, the animals underwent aseptic surgery to implant an MRI-compatible head post (Rogue Research). The head post was necessary to prevent head movement during the TUS and imaging data acquisition. All procedures were conducted under licenses from the United Kingdom (UK) Home Office in accordance with the UK Animals (Scientific Procedures) Act 1986 and with the European Union guidelines

(EU Directive 2010/63/EU).

4.3.2 Behavioural training

All four animals were initially trained to sit in a Magnetic Resonance Imaging (MRI)-compatible chair in a sphinx position and perform a computer-based three-armed reversal learning bandit task while inside a custom-made mock scanner while mock MRI noise was played in the background for the full length of the session. Such mock conditions were implemented to simulate the MRI scanning environment and train the animals to tolerate the testing conditions. Similarly, before each training sessions the animals were tested in a mock TUS session in which no ultrasonic waves were transmitted through the skull. All animals were trained to reach two custom-made infrared touch sensors (one located in front of the left hand and one located in front of the right hand) to respond to abstract stimuli presented on the left and right of a black screen located in front of the animal. Touching the left and right sensors effected choices of the left and right hand options respectively. Once the animals reached a learning criterion, the animals were moved to the actual MRI scanner where, they performed the task with head fixation at over 70% accuracy. MRI data acquisition started once animals performed above 70% accuracy for more than 3 consecutive times in the MRI scanner environment.

4.3.3 Experimental task

Testing was conducted in a horizontal bore 3 Tesla MRI scanner with the animals' head fixed to the MRI-compatible chair. Animals were required to

learn the continually changing reward probabilities associated with three choice options and to adapt their behavior to maximize the amount of reward received. On each trial two out of the three possible stimuli were presented. Animals had to learn the reward probability associated with each choice option by choosing between the two options presented on the screen at each trial and monitoring whether or not a reward was received for making the choice. The three stimuli used to represent the three options were novel in each testing session (**Figure 1**) and each testing session took place on a different day. The task consisted in a probabilistic reward-guided reversal learning task inspired by tasks originally designed to study reinforcement learning. Choice options were allocated pseudo-randomly to the right- and left-hand sides of the screen and the animals responded by reaching a right or left infrared sensor placed in front of each of their hands. The rewards were delivered probabilistically, and the probabilities associated with the three options fluctuated during the entire session, with the probability of two of the options reversing toward the middle of a session (**Figure 2**).




Figure 1. Experimental task. In each session, three novel abstract stimuli, each associated to a specific reward rate, were used as options the animals could choose from in order to obtain a reward. By making consecutive

choices throughout the task, the animals were required to learn the reward likelihood associated with each choice option. To test for the causal role of area IOFC/12o in credit assignment and the spatial specificity of the behavioural effects after TUS, three conditions were implemented: (1) no stimulation (sham, control condition) or (2) a 40s train of ultrasound stimulation targeted to either IOFC/12o (experimental condition, experiment 1) or (3) to the control stimulation region anterior PFC (additional control condition, experiment 2). Either sham or ultrasound stimulation were applied to the animal before the experimental task in a separate room next to the MRI scanner where the ultrasound and neuronavigational apparatus were set up. Each animal performed up to 200 trials per session in each experiment and session always took place on separate days. Half-way through the experiment, the reward rate of two of the stimuli reversed. **DECISION PHASE:** On each trial two out of the three stimuli were presented on a black screen in a dark room after an inter trial interval (ITI) characterized by a random delay the length of which was jittered. The animals were required to choose one of these choice options by reaching and touching the infrared sensors positioned on the same side (left or right) as the option on the screen they desired to choose. The positions of each choice option on either the left or the right part of the screen were fully pseudorandomized. **OUTCOME PHASE:** Choices could lead to either a reward or no reward. The outcome phase lasted 1.5 seconds and the chosen stimulus that led to a reward remained visible on the screen during this phase. On the contrary, if the selected stimulus did not lead to a reward then it disappeared from the screen immediately after the choice and none of the choice options remained on the screen. Each reward was composed of two 0.6 ml drops of fruit juice mixed with water (25% / 75% ratio) delivered by a spout placed near the animal's mouth during the experiment. The session was finished once the animal performed all the 200 trials.

To control for reward schedule-induced biases in the animals' choice preferences, two macaques were presented with one reward schedule and the other two with a second reward schedule with matching reward rates for the three stimuli. Thus, in both reward schedules the probability range for option A was 90 to 10%, the probability range for option B was 70 to 30% and the probability range for option C was 10 to 90%. As mentioned above, although three choice options were presented in each session, only two of them could be chosen on each trial (**Figure 1**). This manipulation requires the animals to learn the contingency between the selected choice option and the subsequent outcome. Thus, to maximize the amount of reward received, the animals are required to use these choice option-outcome associations to guide behavior. Moreover, the probabilistic nature of the task required animals to continually monitor the outcomes that followed choices of each option chosen. For instance, in a given trial the macaque may have chosen the stimulus associated with the highest likelihood of reward but may not have been rewarded for making that choice. The choice's valuation should reflect not just what happened on that trial but also whether outcomes were delivered recently when the same choice was taken. The animals had also to keep track of outcomes received in previous trials because only two of the three options were presented on any given trial. This task is, therefore, ideal for investigating the capacity of animals to integrate the past history of choice option-outcome conjunctions in order to guide future behavior. Furthermore, the probabilistic nature of the stimulus presentation on each trial prevents the animal from predicting which option will be presented in the next trial.




Figure 2. Reward Schedules. The two plots show the two probabilistic reward schedules used in experiments 1 and 2 in all conditions. Their presentation was counterbalanced across the 4 animals. The two schedules represent the reward rates (likelihood of delivering reward if chosen) associated with option 1 (blue), 2 (yellow), or 3 (red) on each trial throughout each of the 200-trial sessions. Half-way through the task (around trial 100) the reward likelihoods associated with option 1 and option 2 were reversed to test for the animals' capacity to rapidly adapt their behaviour to changes in the contingencies between choice options and rewards and therefore their ability to correctly assign credit to a visual stimulus and flexibly update previous representation of choice option-outcome associations.

After making their decision, if an option selected led to a reward (as per the programmed reward contingencies associated with each option (**Figure 2**), the unselected option disappeared and the chosen option remained on the screen and a juice reward was delivered. If an option selected led to no reward, no juice was delivered. The outcome phase lasted 1.5 s. Each reward was composed of two 0.6 ml drops of blackcurrant Ribena juice (25% blackcurrant Ribena, 75% water) delivered by a spout placed near the animal's mouth during scanning. Each animal performed up

to 200 trials per session.

Each animal underwent two experiments for a total of 20 sessions in each one (4 animals × 5 sessions). Experiment 1 consisted in two separate conditions: one condition employed TUS targeted to IOFC/12o. The other condition in which no stimulation was applied (sham). To ensure that the behavioural effects of TUS were not induced by non-focal modulation of activity in prefrontal cortex, we ran an additional control experiment (experiment 2) but this time targeting a separate prefrontal area before the animals underwent the same task. Experiment 2 also consisted of two separate conditions: one control condition in which no stimulation was applied (sham) and an active control condition in which TUS was applied to a control target area in the anterior prefrontal cortex (anterior PFC). In each experiment, each animal performed 5 sessions of the sham condition that were paired with 5 session of the TUS condition (5 IOFC/12o TUS sessions in experiment 1 and 5 anterior PFC TUS sessions in experiment 2) in the MRI scanner. No statistical methods were used to pre-determine sample sizes but our sample sizes are similar to those reported in previous publications (Chau et al., 2015; Fouragnan et al., 2019; Kaskan et al., 2017; Papageorgiou et al., 2017). The experiment was controlled by Presentation software (Neurobehavioral Systems Inc., Albany, CA).




Figure 3. Transcranial Ultrasound Stimulation sites in IOFC/12o and anterior PFC. The same 4 macaque monkeys (D, E, F, G) participated in all the conditions split between the two experiments. Sham or TUS targeted to either IOFC/12o (top row: coronal experiment 1) or anterior PFC (experiment 2) were performed immediately before the task in a separate room using a single-element transducer with a diameter of 64 millimeters and a focal depth of 51.74 millimeters in combination with a coupling cone filled with degassed water and sealed with a latex membrane. Each area, both during sham and TUS conditions, was targeted using a Brainsight Neuronavigation system after previous registration of each animal's T1-weighted image to the system. Stimulation was delivered at a wave frequency of 250 kHz. The TUS parameters used in the two experiments generated a beam of ultrasonic

waves with a 10 mm-wide diameter in the rhesus macaque brain (Folloni et al., 2019; Verhagen et al., 2019; Fouragnan et al., 2019). Given the limited size of the beam of ultrasound, stimulation of the IOFC/12o was applied bilaterally in experiment one, whereas a single train was applied on the midline through the inter-hemispheric fissure to target the anterior PFC (experiment 2). Sham sessions were interleaved with TUS sonication days and completely mirrored a typical stimulation session (the experimental laboratory setting was the same as were the procedures involved in preparing for stimulation such as neuro-navigation and cone positioning and application of the cone to the shaved skin on the head of the animal) except that sonication was not triggered. LOFC/12o: lateral orbitofrontal cortex; anterior PFC: anterior prefrontal cortex; sham: no stimulation.

4.3.4 Transcranial focused ultrasound stimulation

A single-element ultrasound transducer (H115-MR, diameter 64 mm, Sonic Concept) with a 51.74 mm focal depth was used with a coupling cone filled with degassed water and sealed with a latex membrane (Durex). The ultrasound wave frequency was set to the 250 kHz resonance frequency and 30 ms bursts of ultrasound were generated every 100 ms with a digital function generator (Handyscope HS5, TiePie Engineering). Overall, the stimulation lasted for 40 s. A 75-W amplifier (75A250A, Amplifier Research) was used to deliver the required power to the transducer. A TiePie probe connected to an oscilloscope was used to monitor the voltage delivered. The recorded peak- to-peak voltage was constant throughout the stimulation session. The recorded peak- to-peak voltage was constant throughout the stimulation session. Voltage values per session ranged from 130 to 136 V and corresponded to a peak negative pressure ranging from 1.17 to 1.29

MPa, respectively, as measured in water with an in house heterodyne interferometer (Constans et al., 2017).

TUS was applied immediately before each testing session using what has been referred to as an “offline” procedure. This protocol has previously been shown to modulate neural activity for up to 2 hours (see chapter 3; Folloni et al., 2019; Verhagen et al., 2019; Fouragnan, et al., 2019). Monkeys sat in an MRI-compatible chair in a sphinx position with their heads fixed to the chair through the implanted head post. Each animal’s head was registered to a frameless stereotaxicBrainsight Neuronavigation system (Rogue Research, Montreal, Canada) in advance of the session using a previously acquired T1-weighted image. The anatomical localization of fiducial markers used for each animal was also acquired at the same time as the collection of the T1-weighted image for that monkey. The registration of the fiducial markers to the T1-weighted image of each animal made anatomically precise spatial navigation within the brain of each subject possible and, therefore, the same exact area could be identified and stimulated in each session. The transducer was positioned using the Brainsight Neuronavigation system so that the focal spot was centered on the targeted brain region. In experiment 1, the targeted brain region was the caudal portion IOFC/12o (MNI coordinates $x = -16$, $y = 8$, $z = -3$; $x = 16$, $y = 8$, $z = -3$), in a region overlapping with the IOFC/12o region-of-interest in Chau et al., (2015). In experiment 1, to ensure complete modulation of activity in IOFC/12o, bilateral stimulation was applied to the same region in each hemisphere. In the control experiment 2, TUS was instead applied to a portion of the anterior PFC (MNI coordinates $x = 0$, $y = 23$, $z = 11$)

corresponding to the target previously stimulated in a separate resting-state fMRI experiment (Verhagen et al., 2019). As in the protocol reported by Verhagen et al., TUS was applied once for 40 s on the midline through the interhemispheric fissure. As we previous showed (Verhagen et al., 2019), the TUS parameters used in the current study generated a beam of ultrasonic waves with a 10 mm-wide diameter that targeted anterior PFC in both hemispheres with a single sonication.

The ultrasound transducer/coupling cone montage was directly positioned on previously shaved skin on which conductive gel (SignaGel Electrode; Parker Laboratories Inc.) had been applied. The coupling cone filled with degassed water and gel applied to the skin was used to ensure ultrasonic coupling between the transducer and the animal's head. Following stimulation, the monkeys were moved to the MRI scanner for testing. The TUS procedure lasted 20 minutes on average.

In each experiment, sham sessions were interleaved with TUS sonication days and completely mirrored a typical stimulation session (the experimental laboratory setting was the same as were the procedures involved in preparing for stimulation such as neuro-navigation and cone positioning and application of the cone to the shaved skin on the head of the animal) except that sonication was not triggered.

To test for the specificity of IOFC/12o TUS, in experiment 2, 20 sham and 20 anterior PFC TUS (4 animals $\times$ 5 sessions) sessions were collected using the same experimental design as in experiment 1. TUS and control days were interleaved in one of two pseudo-random orders that were counterbalanced across animals in each experiment (order 1: S, S, R, T, R,

T, R, T, R, T, R, T, R, S, R, S, R, S; order 2: T, T, R, S, R, S, R, S, R, S, R, S, R, T, R, T, R, T) where T, S and R stand for TUS, sham, and rest days respectively. No statistical methods were used to pre-determine sample sizes but our sample sizes are similar to those reported in previous publications. Data collection and analysis were not performed blind to the conditions of the experiments.

4.3.5 Reinforcement learning model

For every session of experiment 1 and experiment 2, we fitted a reinforcement learning (RL) model so that we could estimate the trial-by-trial reward estimates held by the animals. The reinforcement learning (RL) model applied to the behavioural data is explained here (for an extended explanation on model fitting and comparison with other reinforcement learning models, i.e. with a standard simple RL model (*RLsimple*) see the manuscript Wittman et al., in appendix). In the current RL model, each choice option was described by a different value estimate ($Q(A)$, $Q(B)$, $Q(C)$). Each of these value estimates reflected whether a reward was received or not when a given choice option was selected in a trial. As described in the “Experimental task” section of this chapter, novel abstract stimuli were used in each session to prevent any influence of previous learning, from an earlier session, being exerted on current behaviour. All Q-values were initialized at 0.5 at the start of each session and the chosen stimulus was updated based on the discrepancy between outcome (“0” or “1” for reward and no reward, respectively) and Q-value, scaled by a learning rate α . By contrast, at each trial, the Q-value estimate of the unchosen choice option remained

unchanged. For example, if option A was chosen on trial t, then its prediction error (PE) and value update would be calculated based on the reward r received as follows:

$$\textbf{Equation 1: } PE_t(A) = r_t - Q_t(A)$$

$$\textbf{Equation 2: } Q_{t+1}(A) = Q_t(A) + \alpha PE_t(A)$$

In addition to choice values, the decision variable (DV) also reflected which choices had been made recently regardless of whether or not they had been rewarded. Choices taken recently were more likely to be repeated on the current trial regardless of their reward value. Such terms are sometimes referred to as “choice stickiness” or “choice trace” terms. The current model incorporated two choice trace terms. One reflected which location (left or right) had been chosen recently regardless of the identity of the option chosen. Choices were more likely to be made again to the same location. The other choice trace reflected the identities of the choice options that had been chosen recently regardless of their location on the left or the right. Note that the choice option presented on the left or right side of the screen was independent from the option identity (i.e. option A, B or C) and was assigned pseudorandomly. In a simpler model, *RLsimple* (Wittman et al., in appendix), which did not include any choice or reward memory traces, the DV for each trial was simply calculated as the value of the stimulus presented on the right side minus the value of the stimulus presented on the left side plus an additional additive constant reflecting a side bias (SB).

$$\textbf{Equation 3: } DV_{RLsimple,t} = Q_t(right) - Q_t(left) - SB$$

Then, for any model m (here $m=RLsimple$), the decision variable was filtered through a standard softmax function to calculate the probability of a rightward choice.

$$\textbf{Equation 4: } p_t(right) = \frac{1}{1 + e^{-\beta \times DV_{m,t}}}$$

The probability of the observed choice on each trial was calculated as:

$$\textbf{Equation 5: } p_t(choice) = \begin{cases} p_t(right), & \text{if right option is chosen} \\ 1 - p_t(right), & \text{otherwise} \end{cases}$$

Finally, three free parameters were included: the learning rate α (bound between 0 and 1), the inverse temperature parameter β (bound between 0 and positive Infinity) and the side bias SB (bound between -1 and 1). The current model included additional memory trace parameters associated with choice (history of past choice stimuli identities and choice locations) and reward. The two choice traces, choice location trace (*CL-trace*) and choice stimulus trace (*CS-trace*), modified the basic DV (Wittman et al., in appendix) by adding terms related to the history of past choice stimuli and choice locations. The reward trace parameter modified the calculation of prediction errors instead.

Choice and reward memory traces were initialized at zero at the start

of each session. The choice location trace (*CL-trace*) was updated on every trial based on the discrepancy between the actual choice location (L), coded as “-1” or “1” (for right and left side, respectively), and the *CL-trace*, scaled by a learning rate α_{CL} .

$$\textbf{Equation 6: } CL\text{-}trace_{t+1} = CL\text{-}trace_t + \alpha_{CL}(L_t - CL\text{-}trace_t)$$

The choice stimulus trace (*CS-trace*) decayed exponentially from one trial to the next one with a rate determined by a free parameter λ_{CS} (Bogacz et al., 2007; Sutton and Barto, 1998). Every time an option was chosen in one trial, the *CS-trace* associated with that option was set to 1 at the end of the trial. For example, the decay of the *CS-trace* for a stimulus A on trial t was calculated as:

$$\textbf{Equation 7: } CS\text{-}trace_t(A) = \lambda_{CS}CS\text{-}trace_{t-1}(A)$$

CL-trace and *CS-trace* required only knowledge of the choice location and choice stimulus, respectively, and ignored the sequence of outcomes experienced over the course of a session.

The reward trace (*R-trace*) was updated on every trial based on the discrepancy between *R-trace* and the observed outcome, scaled by a learning rate α_R , and independently from the choice taken.

$$\textbf{Equation 8: } R\text{-}trace_{t+1} = R\text{-}trace_t + \alpha_R(r_t - R\text{-}trace_t)$$

The free parameters λ_{CS} , α_{CL} and α_R were all bound between 0 and 1. The influence of *CL-trace*, *CS-trace* and *R-trace* on learning and choice was scaled by weight parameters, which could assume values that could be positive, zero, or negative. The magnitude and direction of influence of the memory traces were therefore determined empirically during model fitting (Wittman et al., in appendix). *CL-trace* and *CS-trace* were added to the *DV*. The *CL-trace* could be added directly, since it was already coded in terms of spatial location similarly to the *DV* itself (although inverted). For the *CS-trace*, the influence on choice was determined by the difference in *CS-trace* associated with the stimuli on the right and left (referred to as $CS\text{-}bonus_t$ in **Equation 9**). Note that the *CS-trace* difference was added to the *DV* after the trial-wise decay but before the update of the chosen stimulus. In the current model, the *DVs* contained both *CL-trace* and *CS-trace* in addition to the components from *RLsimple* previously described (see also Wittman et al., in appendix).

Equation 9: $CS\text{-}bonus_t = CS\text{-}trace_t(right) - CS\text{-}trace_t(left)$

Equation 10: $DV_{CL+CS,t} = DV_{RLsimple,t} + w_{CL}CL\text{-}bonus_t + w_{CS}CS\text{-}bonus_t$

According to **Equation 10**, if w_{CL} assumes positive values then the *DV* in the current trial will be biased towards the same location as the direction of previous choices. If w_{CS} assumes positive values, the *DV* will oscillate towards the stimulus with the highest *CS-trace*. Positive values of w_{CL} and

w_{CS} therefore describe a tendency of the animals to repeat predominant previous choice locations and predominant previous choice stimuli, respectively.

Finally, the *R-trace* was scaled by a weight parameter w_R and their conjoint estimate influenced calculation of the prediction error. For example, if option *A* was chosen on trial *t*, the corresponding *PE* would be calculated as:

Equation 11: $PE_{m,t}(A) = r_t + w_R R-trace_t - Q_t(A)$

4.3.6 Statistical analyses

To analyze the behavior of the animals in the three experiments and estimate the effects of no ultrasound stimulation (sham), IOFC/12o TUS, and anterior PFC TUS on reward-guided learning and decision-making we used multiple linear or logistic regressions as implemented in MATLAB. For logistic functions, we used a logit link with categorical predictors. All regressors were normalized to ensure between-session and between-condition commensurability of the regression coefficients. For each session, one β regression weight was extracted for each regressor. These were then tested for statistical significance across all animals/sessions. All data were shown as mean with standard error of the mean (mean $\pm$ s.e.m.). Probabilities of $P < 0.05$ were considered as significant.

4.4 RESULTS

4.4.1 Defining choices, rewards, and their conjunction

We first wanted to investigate whether TUS targeted to IOFC/12o, but not to anterior PFC, induced a disruption in credit assignment defined as the learning of association between a visual stimulus that constituted, a choice option and the outcome it led to. In addition, this analysis allowed us to estimate whether modulation of activity in IOFC/12o but not in anterior PFC affected contingent learning. As described in the introduction of this chapter, we hypothesized that IOFC TUS would affect the capacity of the animals to use the history of past choice option-outcome conjunctions to flexibly adapt future choices to the volatile environment of the task.

We used linear mixed-effects models to examine animals' choices as a function of recent choice-outcome conjunctions and compared the results between sham and IOFC/12o TUS (experiment 1) or between IOFC/12o and the control area anterior PFC TUS (experiment 2). In each case, at every trial, the animal's history of choice option-outcome associations, choices, and rewards were modelled separately as a function of a linear combination of two fixed effects variables, condition [*sham*, *TUS*] and choice history [*1 trial in the past (Trial-1)*, *2 trials in the past (Trial-2)*, *3 trials in the past (Trial-3)*], and one random-effects (*subject*, 4 levels) variable:

Equation 12: $choice = 1 + condition + choice\ history + (1 | subject)$

Equation 13: $reward = 1 + condition + choice\ history + (1 | subject)$

Equation 14: $choice - outcome = 1 + condition + choice\ history +$
 $(1 | subject)$

Equation 15: $\frac{STAY}{SWITCH} index = 1 + condition + options + (1 | subject)$

The separate modelling of both within-subject and between-subject variance by treating the monkeys as random-effects allows extrapolation of the results beyond the 4 monkeys specifically sampled in the current study to the wider population.

4.4.2 Ultrasound stimulation of lateral OFC disrupts credit assignment

To test for changes in the impact that the history of choices, rewards, and choice-reward conjunctions exerted on subsequent decisions following disruption of IOFC/12o activity, we used a series of linear mixed-effects models (**Equations 12-14**) using condition and choice history as fixed effects and subject as random-effects. We first observed that after IOFC/12o TUS, the animals were impaired in using contingent associations between choice options and outcomes to guide future behavior (**Figure 4,5**). Following TUS, animals were still influenced by the history of choice options in isolation (**Figure 4,5**; $t_{(177)} = 0.56$, $p = 0.58$) and by the history of outcomes (rewards) in isolation (**Figure 4,5**; $t_{(177)} = -0.76$, $p = 0.45$) but they were impaired in learning the contingent relationships between choices and outcomes; there was a significant effect of condition (sham, IOFC/12o: **Figure 4,5**; $t_{(177)} = -3.06$, $p < 0.01$).




Figure 4. LOFC/12o TUS modulates credit assignment. Influence of previous choices (vertical axis) and outcomes (horizontal axis) and choice-outcome conjunctions (diagonal from upper left to lower right) from previous trials ($t - 1$, $t - 2$, $t - 3$) on subsequent choices. Bright orange colours indicate

a strong positive influence while blue colours indicate an absence of influence. In the control sham case, the bright orange diagonal indicates that subsequent choices were influenced by the past history of choice-outcome conjunctions as opposed to the history of choices or history of outcomes in isolation. The same was true in the control anterior PFC TUS case. The absence of the orange diagonal in the IOFC/12o TUS case, however, indicates an absence of the same influence. Instead, after IOFC/12o TUS animals' choices were simply influenced by the most recent choice and outcome (top left corner is bright orange). IOFC/12o: lateral orbitofrontal cortex; anterior PFC: anterior prefrontal cortex; sham: no stimulation.

Together, these results suggest a role for IOFC/12o in integrating identity of a choice option with information about whether choosing it had led to the receipt of a reward. In addition IOFC/12o TUS decreases this influence. We, therefore, next examined how the influence of past choice-outcome conjunctions changed after IOFC/12o TUS as a function of time. The impact of disruption increased as trials further in the past were considered (**Figure 4,5**; $t_{(177)} = -4.48$, $p < 0.0001$).




Figure 5. LOFC/12o TUS only disrupts choice-reward association but not the separate representation of choice and reward. On the left are represented, on a T1-weighted template image, the stimulation targets in IOFC/12o (experiment 1, top row) and anterior PFC (experiment 2, bottom row). Influence of previous choices (top and bottom rows, right column “choice”) and outcomes (top and bottom rows, left column “rewards”) and choice-reward associations (top and bottom rows, middle column “choice-reward associations”) from previous trials ($t-1$, $t-2$, $t-3$) on subsequent choices. The column “rewards”, “choices” and “choice-reward associations” correspond to the estimates represented on the horizontal, vertical axes and diagonal (from upper left to lower right), respectively, in Figure 4. Comparison of “rewards”, “choices” and “choice-reward associations” estimates between the control sham (blue) and either IOFC/12o TUS (red, top row) or the control anterior PFC TUS (green, bottom row) are shown. In experiment 1, a significant difference in the capacity of the animals to encode the history of “choice-reward associations” was found. This contrasted to a preserved ability of the animals to encode either “choices” or “rewards” in isolation after IOFC/12o TUS. No differences between control sham and the control anterior PFC TUS was found for any of the credit assignment estimates. LOFC/12o: lateral orbitofrontal cortex; anterior PFC: anterior

prefrontal cortex; sham: no stimulation. For plotting formatting purposes, IOFC/12o is indicated as IOFC and anterior PFC is reported as aPFC in the legend of this figure.

4.4.3 Credit assignment deficit impairs subsequent reward-guided adaptive behaviour

In the behavioural task that the animals learned (**Figure 2**), an adaptive strategy to maximize reward consists of staying with choices that have been rewarded in the past and switching away from choices that were not rewarded in the past. We, therefore, analyzed the strategy of the animals to STAY with or SWITCH from choices across consecutive trials. In the model used here, if an animal chose to STAY with the same option in the following trials, and the chosen option was rewarded, then a value of 1 was assigned to the decision. Conversely, if the animal decided to switch to another option in the subsequent trials after the previous chosen option was rewarded, then a value of -1 was assigned. However, because the task relies on stochastic reward schedules where only two out of three stimuli were presented and available for choice on any given trial, it is possible that the option chosen in one trial will not be presented in the following one. In these trials we therefore included a contextual modulation of the STAY/SWITCH value assigned to the behavior of the animal. Based on this contextual modulation, if on any given trial the animal chose an option that was not presented in the previous one but was chosen two trials in the past, then a value of +0.5 was assigned to the decision. Conversely, if the animal did not choose the same option, then

a value of -0.5 was allocated. We then used a linear mixed-effects model (**Equation 15**) to test whether any difference in sticking with or switching from one choice preference was induced by IOFC/12o TUS. Overall, after IOFC/12o stimulation the animals were impaired in flexibly adapting their subsequent behavior (**Figure 6**; $t_{(177)} = -3.48$, $p < 0.001$) and such deficit was present irrespective of the specific choice identity (**Figure 6**; $t_{(177)} = 0.85$, $p = 0.4$).

4.4.4 Ultrasound stimulation of the adjacent anterior PFC does not affect credit assignment or subsequent adaptive behavior

In experiment 2 we targeted the active control area anterior PFC to examine the anatomical and functional specificity of the TUS protocol and the role assigned to IOFC/12o on the basis of the results from experiment 1. Because of the anatomical proximity of anterior PFC to IOFC/12o, this region was a good site to investigate the focality of the stimulation. For example, it is possible that the behavioural deficits observed in experiment 1 might not be specific to IOFC/12o TUS but instead they might reflect widespread disruption of prefrontal lobe activity or by unfocused noise spreading throughout the frontal and motor system that would be caused by the application of TUS to any adjacent area. If this were the case, then impairments in certain features of credit assignment and reward-guided adaptive behavior should be observed also after stimulating the neighboring anterior PFC using the same protocol. Moreover, anterior PFC is also involved in decision-making processes such as counterfactual choice and learning (Boorman et al., 2011) that are performed alongside credit

assignment. For these functional and anatomical reasons, anterior PFC represented an ideal active control site in our experiments.

The same analysis previously performed to investigate the effects of TUS targeted to IOFC/12o on credit assignment was used here but this time using the anterior PFC TUS sessions as experimental condition instead of the IOFC/12o TUS ones. Similarly to the analysis for experiment 1 (**Equations 12-14**), we used a series of linear mixed-effects models using condition and choice history as fixed effects and subject as random-effects. In contrast with the behavior observed after TUS was applied to IOFC, TUS targeted to the active control region anterior PFC did not cause any impairment in the animals' capacity to encode and use conjoint choice-outcome information to guide subsequent stimulus choices (**Figure 4,5**; $t_{(177)} = -0.4$, $p = 0.69$). Additionally, and matching the results from experiment 1, neither the influence of choice history in isolation (**Figure 4,5**; $t_{(177)} = -0.18$, $p = 0.86$) nor reward history in isolation (**Figure 4,5**; $t_{(177)} = -0.28$, $p = 0.78$) were affected by anterior PFC TUS.




Figure 6. Adaptive flexible behaviour following a reward. Left and right panels represent a comparison of the animals' frequencies of staying with the same option or switching to another one in the trial following receipt of a reward for each of the three options. Left panel compares sham control (blue) with IOFC/12o TUS (red) sessions. Right panel compares sham control (blue) with the control anterior PFC TUS (green) sessions. The stay/switch index was estimated as follows: staying with the same option in the trial following reward receipt = + 1; switching to another option in the subsequent trial after the previous choice was rewarded = - 1; staying with an option not available (not presented on the screen) in the previous trial but chosen the last trial it appeared on the screen and that led to a reward = + 0.5; switching to another option not available (not presented on the screen) in the previous trial but chosen the last trial it appeared on the screen and that led to a reward = - 0.5. Significant differences between sham control and IOFC/12o TUS conditions, but not between sham control and the control anterior PFC TUS condition, were observed. LOFC/12o: lateral orbitofrontal cortex; anterior PFC: anterior prefrontal cortex; sham: no stimulation. For plotting formatting purposes, IOFC/12o is indicated as IOFC and anterior PFC is reported as aPFC in the legend of this figure.

4.4.5 Changes in win-stay/lose-switch choice strategy after modulating activity in lateral OFC

In the next analysis we took a complementary approach to examine credit assignment as a function of whether outcomes were positive or negative. A simple heuristic is for the animals to repeat a choice (STAY) if it was rewarded in the previous trial (WIN) but they should make a different choice (SWITCH) if no reward was received (LOSE). The learning rule underlying this strategy (Kelley et al., 1962) relies on the observation of the outcome of the previous decision to adapt the following one.

To examine how the animals adapted their choice preferences on the basis of the outcome received in the previous trials, we classified the trials into two categories: (1) whether the animals stayed with an option, after receiving a reward in the previous trial (WIN-STAY) and (2) whether the animals changed their choice, switch, after not receiving a reward in the previous trial (LOSE-SWITCH). According to this classification, if a “WIN-STAY” choice was made in a trial, then a value of 1 was assigned. Conversely, if a “LOSE-SWITCH” choice was performed, a value of -1 was assigned instead. Finally, values assigned to “WIN-STAY” and “LOSE-SWITCH” were combined to create an integrated “WIN-STAY/LOSE-SWITCH” index.

We investigated alteration in adaptive behavior by using a linear mixed-effects model with two fixed effects variable, condition [*sham*, *lOFC/12o TUS*] and choice history [*1 trial in the past (Trial-1)*, *2 trials in the past (Trial-2)*, *3 trials in the past (Trial-3)*], and one random-effects (*monkeys*, *4 levels*) variable:

Equation 16: $\text{choice strategy} = 1 + \text{condition} + \text{choice history} + (1 \mid \text{subject})$

After TUS targeted to the IOFC/12o, we observed changes in the choices made by the animals as a function of the feedback received for the previous decision. In the sham condition, animals were able to adjust their future choices based on the nature of the feedback received. However, IOFC/12o TUS caused a significant difference in the implementation of the WIN-STAY/LOSE-SWITCH strategy in comparison with sham (**Figure 7**; $t_{(177)} = -2.98$, $p < 0.01$).




Figure 7. Win-Stay/Lose-Switch behaviour in control sham and IOFC/12o TUS sessions. Comparison of the animals' frequency of staying with the same option after receiving a reward (Win-Stay) or switching to another one after not receiving a reward (Lose-Switch). The y axis represents an integrated Win-Stay/Lose-Switch estimate for the control sham

(blue) and the IOFC/12o TUS (red) over three consecutive trials (x axis). Significant differences in Win-Stay/Lose-Switch behaviour were found between sham and IOFC/12o TUS. For plotting formatting purposes, IOFC/12o is indicated as IOFC in the legend of this figure.

4.4.6 Lateral OFC/12o TUS differentially modulates learning from positive and negative outcomes

After observing a disruption in adapting future choices according to the feedback received, we wanted to determine whether these effects were present independently from the nature of the feedback received or whether the impairment in behavioural adaptation was specific to behaviour following one type of outcome, reward or non-reward. We, therefore, further classified the trials into four categories according to the nature of the feedback received and the consequent choice preference: (1) whether the animals stayed with one option, thereby repeat the same choice preference, after receiving a reward in the previous trial (WIN-STAY), (2) whether the animals switched to another choice option, after receiving a reward in the previous trial (WIN-SWITCH), (3) whether the animals stayed with an option even though no reward was received in the previous trial (LOSE-STAY), (4) whether the animals changed their choice after not receiving a reward in the previous trial (LOSE-SWITCH). Modulation of IOFC/12o might disrupt just one type of behaviour following just either positive outcomes (reward), or negative outcomes (i.e. lack of reward). We then used a series of linear mixed-effects models as previously done for the integrated WIN-

STAY/LOSE-SWITCH (WSLS) strategy (**Figure 7**) using the same fixed and random effects variables:

$$\textbf{Equation 17: } WIN - STAY = 1 + \textit{condition} + \textit{choice history} + (1 | \textit{subject})$$

$$\textbf{Equation 18: } WIN - SWITCH = 1 + \textit{condition} + \textit{choice history} + (1 | \textit{subject})$$

$$\textbf{Equation 19: } LOSE - STAY = 1 + \textit{condition} + \textit{choice history} + (1 | \textit{subject})$$

$$\textbf{Equation 20: } LOSE - SWITCH = 1 + \textit{condition} + \textit{choice history} + (1 | \textit{subject})$$

When TUS was targeted to IOFC/12o, we observed a reduction in the animals' ability to update their estimate of a choice option's value after trials in which a reward was received for choosing it ("WIN TRIALS") and the same choice was made ("STAY") in the following trial (**Figure 8A**; $t_{(177)} = -3.92$, $p < 0.001$). However, no significant effect of IOFC/12o was present in trials in which a reward was not obtained ("LOSE TRIALS") but the same choice was re-performed ("STAY") in the subsequent trial (**Figure 8B**; $t_{(177)} = 0.5$, $p = 0.62$). Next, we looked at the capacity of the animals to update their value estimate of a choice option as indexed by switching choice ("SWITCH") after trials in which reward was or was not received. After IOFC/12o TUS, animals switched more often to select another choice option in the following trial despite having just been rewarded ("WIN TRIALS") for making that choice (**Figure 8C**; $t_{(177)} = 3.56$, $p < 0.001$). On the contrary, no differences between IOFC TUS and sham were present when we looked at the strategy of the

animals to switch to another choice option a negative feedback (“LOSE TRIALS”) was received in the previous decision (**Figure 8D**; $t_{(177)} = 0.69$, $p = 0.49$).

4.4.7 Anterior PFC TUS has no effects on adaptive behavior following either positive nor negative outcomes

Once again, we examined the specificity of the OFC TUS effects by running another set of analyses on the anterior PFC TUS data collected in the control experiment 2. We classified the trials and ran a series of linear mixed-effects models mirroring the ones used to examine the win-stay/lose switch strategy in experiment 1 (**Equations 15-20**) but this time comparing sham with sessions in which TUS was applied to the anterior PFC.

We first looked at changes in the integrated WSLs estimates as described above. In contrast to IOFC/12o TUS, after anterior PFC TUS the animals did not exhibit any deficit in adaptive behaviour (**Figure 9**; $t_{(177)} = -1.13$, $p = 0.26$). Consistently, when we carried a more detailed analysis of behaviour as a function of the outcome of the previous trial (reward/non-reward) and the behaviour on the next trial (stay/switch), we observed a lack of any effect of anterior PFC : “WIN-STAY” strategy (**Figure 10A**; $t_{(177)} = -1.54$, $p = 0.13$), “LOSE-STAY” strategy (**Figure 10B**; $t_{(177)} = 0.15$, $p = 0.88$), “WIN-SWITCH” strategy (**Figure 10C**; $t_{(177)} = 1.3$, $p = 0.2$), “LOSE-SWITCH” strategy (**Figure 10D**; $t_{(177)} = 0.07$, $p = 0.94$).

In summary, IOFC/12o affected adaptive switch/stay behaviour but anterior PFC TUS exerted no similar effect.




Figure 8. Comparison of “stay” and “switch” behaviour after positive and negative feedback in control sham and IOFC/12o TUS sessions. **Left column:** Animals’ frequency of (top row, **A**) staying in the current trial with the same option as that chosen in the previous trial after receiving positive feedback (receipt of reward, (Win-Stay) or (bottom row, **B**) staying with the same option in the current trial after receiving negative feedback, (no reward), in the previous trial (Lose-Stay). **Right column:** Animals’ frequency of (top row, **C**) switching in the current trial to another option compared to the one chosen in the previous trial even though that choice had led to a positive feedback on a previous trial (Win-Switch) or (bottom row, **D**) switching to another option in the current trial after receiving negative feedback in the previous one (Lose-Switch). Stay and switch behaviours after either positive or negative feedback are displayed with reference to choices made 1, 2 or 3 trials in the past (Trial -1, Trial -2, Trial -3) and compared between control sham (blue) and IOFC/12o TUS (red). A significant difference in Win-Stay and Win-Switch but not Lose-Stay and Lose-Switch behaviours were found. For plotting formatting purposes, IOFC/12o is indicated as IOFC in the legend of this figure.




Figure 9. Win-Stay/Lose-Switch behaviour in control sham and control anterior PFC TUS sessions. Comparison of the animals' behaviour of staying with the same option after receiving a reward (Win-Stay) or switching to another one after not receiving a reward (Lose-Switch). The y axis represents an integrated Win-Stay/Lose-Switch estimate for the control sham (blue) and the control anterior PFC TUS (green) over three consecutive trials (x axis). No differences in Win-Stay/Lose-Switch behaviour were found between sham and anterior PFC TUS. For plotting formatting purposes, anterior PFC is indicated as aPFC in the legend of this figure.




Figure 10. Comparison of “stay” and “switch” behaviour after positive and negative feedback in control sham and control anterior PFC TUS sessions. Left column: Animals’ frequency of (top row, **A**) staying in the current trial with the same option chosen as in the previous trial after receiving positive feedback (receipt of reward), (Win-Stay) or (bottom row, **B**) staying with the same option in the current trial after receiving negative feedback (no reward), in the previous one when choosing the same option (Lose-Stay). **Right column:** Animals’ frequency of (top row, **C**) switching in the current trial to another option compared to the one chosen in the previous trial even though that choice had led to positive feedback (Win-Switch) or (bottom row, **D**) switching to another option in the current trial after receiving negative feedback in the previous one (Lose-Switch). Stay and switch behaviours after either positive or negative feedback are displayed with reference to choices made 1, 2 or 3 trials in the past (Trial -1, Trial -2, Trial -3) and compared between control sham (blue) and the control anterior PFC (green). No differences for any of the Win-Stay and Win-Switch but not Lose-Stay and Lose-Switch behaviours were found across the control conditions. For plotting formatting purposes, IOFC/12o is indicated as IOFC in the legend of this figure. For plotting formatting purposes, anterior PFC is indicated as aPFC in the legend of this figure.

4.5 DISCUSSION

Making a choice in a volatile environment implies not only evaluating choice options and comparing them to one another but in addition it requires updating the estimates of those values when the contingencies between the choices and outcomes change. To successfully adapt to environmental changes and maximize the likelihood of receiving a positive outcome from their choices, primates such as humans and macaques learn the relationship between the sensory identity of a stimulus that is chosen and the subsequent outcome. To do this it is critical that the outcomes that are assigned or credit to the stimulus really are ones that were caused by that choice. Moreover, they must be capable of tracking these stimulus-outcome associations over an extended period of time and therefore taking into account the history of past contingencies when making a choice again in the future. However, because primates' environments are often stochastic and not deterministic, they must be able to flexibly re-learn such associations in the face of changes in reward contingencies.

Here, we show that macaque monkeys are able to learn the discrete sensory identities of the choice options available and the outcomes resulting from choosing them (**Figure 4,5**). Moreover, they were also able to learn the contingencies between the identity of the chosen stimuli and their associated outcomes and used the history of these association to guide future behaviour in subsequent trials (**Figure 4,5**). Previous literature has consistently suggested that one network or regions centered in IOFC/12o is involved in representing and learning choice option-outcome associations (Chau et al., 2015; Murray and Rudebeck, 2018; Rudebeck et al., 2017b; Rushworth et

al., 2011).

Rather than using aspiration lesions (Noonan et al., 2010; Walton et al., 2010) to investigate the causal role of frontal lobe areas in credit assignment we used a minimally invasive technique to temporally disrupt brain activity. Using TUS we were instead able to investigate the role of IOFC/12o in credit assignment by disrupting activity in this region temporarily so that it was possible to alternate sham and stimulation sessions within the same animal, thereby reducing the likelihood of any rapid post-lesion anatomical and functional reorganization that may follow small cortical lesions (Rorden and Karnath, 2004). Moreover, a standard lesion approach requires acquiring pre- and post-surgical experimental data in a “block design-like” fashion and it is possible that if an animal struggles to perform a task over a sustained period of time then it may find the task aversive and only engage with it to only a low degree. Such problems, which may make data interpretation more difficult, are reduced with a temporary interference technique that is applied in alternation with control SHAM stimulation days. By using an offline approach to ultrasonic modulation of neural activity in awake behaving animals we were able to disrupt behaviour even after the stimulation period ceased. However, by alternating no stimulation and stimulation sessions on different days, we showed that the TUS-induced impairment on credit assignment is only temporary and performance returns to baseline by the time of the next control session. We have therefore been able to show in awake behaving animals that TUS exerts a behavioural effect that is consistent with the neural changes that were reported in *Chapter 3* in anaesthetized macaque monkeys. These results are consistent with previous

TUS literature investigating its temporal dynamics (Folloni et al., 2019; Verhagen et al., 2019; Fouragnan et al., 2019; Dellapiazza et al., 2018). The sustained effects and non-surgical feasibility of this offline ultrasound protocol make it a potentially novel approach for neuromodulation and neurorehabilitation in neurological and psychiatric patients.

Interestingly, TUS directed to IOFC/12o not only perturbed credit assignment but it left intact other cognitive processes such as the encoding of the sensory identity of the options or receipt of reward (**Figure 5**). Such specificity in the effects of the disruption suggest a key role of IOFC/12o in associative learning but suggest that the other two computations may be encoded in other portions of cortex. On the one hand, the visual association cortex of the temporal lobe (area TE), hippocampus and rhinal cortex play a primary role in visual discrimination, conditional learning, and memory retrieval and may interact with IOFC/12o through the uncinate fascicle and cortico-striatal circuit (Eacott and Gaffan, 1992; Gaffan, 1994). On the other hand, subcortical loops centered on amygdala, striatum and mesolimbic regions may provide IOFC/12o with information about reward and saliency of the stimulus at a particular moment (Cox and Witten, 2019; Haber, 2011; Lichtenberg et al., 2017; Schultz, 1998). Credit assignment in IOFC/12o may therefore depend on the information's encoded in these circuits but their actual integration into a representation of choice option-outcome association may be especially a function of this area. These results are consistent with previous investigations on this area (Chau et al., 2015; Rudebeck et al., 2017b; Walton et al., 2010) and additionally show that we could achieve similar behavioural perturbation but with only temporal disruption by using

TUS, thereby refining current invasive methodological approaches used for causally investigations on the roles of specific brain areas in cognition.

We show that TUS targeted to IOFC/12o disrupts the capacity of the animal to adapt their choices through learning of the choice option-outcome contingencies. Notably, by using a conservative statistical approach we demonstrate that these effects were specific for behavioural adaptation following the receipt of positive but not negative feedback (**Figure 6, 7, 8**). This result is consistent with the report of outcome related IOFC/12o activity described by Chau et al., (2015). However, it partially contrasts with the work of Chau et al. as, after stimulation, we did not find any impairment in learning from negative feedbacks. Such discrepancy may arise from the different statistical approach used. We indeed implemented more stringent statistical analysis accounting for both between and within-subject variability compared to Chau et al. (2015). An alternative explanation could reside in a role for IOFC/12o in learning from negative outcome during credit assignment which becomes evident only when considered within a network. Indeed, this region may contribute to such learning through its connectivity with other areas but may not constitute the key region supporting the core of this computation. A dissociation in the timing of encoding within orbital regions (although in this experiment recording was performed in area 13) and amygdala during appetitive versus aversive reversal learning was also suggested by Morrison and colleagues (2011). At the time of reversal, the authors showed that neurons in the amygdala updated quicker than those in area 13, starting to respond earlier to stimuli predicting an aversive feedback. On the contrary, neurons in area 13 started to respond earlier than those in the amygdala for

reward-predicting stimuli.

Thus while activity related to negative outcome-guided learning may be observed in IOFC/12o, it is also observed elsewhere and so its disruption may not be sufficient to disrupt negative outcome-guided learning. In line with this hypothesis are the results from Chau et al. suggesting orchestration of activity in IOFC/12o and amygdala in guiding future behavioural adaptation following negative feedbacks as well as an increased activity coupling between these two areas on lose-switch trials. This hypothesis would be consistent with another set of results showing a primary role for the amygdala in encoding stimulus-reward association during learning from negative outcomes (Rudebeck et al., 2017a).

In summary, synthesizing the results from this Chapter with Chau's and Rudebeck's results, a potential model aiming to describe the core of the circuit underlying credit assignment might be that the IOFC/12o is primarily involved in driving reward-guided learning whereas the amygdala mainly influences learning following negative feedback such as lack of reward. On the one hand, the IOFC/12o may primarily rely on positive outcomes to update option-outcome contingencies and therefore the expected value of the available options. On the other hand, the amygdala may be important in driving the behaviour towards a switch in the currently adopted strategy used to achieve a reward rather than merely encoding a negative outcome. A similar role for the amygdala in behavioural adaptation may potentially arise as a result of its capacity to represent the temporal saliency of an option by integrating information about external changes in option-outcome contingencies as well changes in its current motivational state. Access to

both types of information may result from its widespread direct connectivity with prefrontal, limbic and brainstem structures (Aggleton, 1993; Aggleton et al., 1980; Barbas and De Olmos, 1990; Carmichael and Price, 1995; Ongür and Price, 2000; Schmahmann and Pandya, 2006). Future research should aim to investigate whether this hypothesis is true by targeting the amygdala with TUS and investigate related changes in adaptive behaviour following lack of receipt of a reward for a choice previously made.

Neuromodulatory approaches like TUS provide great opportunities for gaining insight into how the brain works and what is the impact of certain computations on behaviour. However, experiments that rely on these techniques to test the causal importance of different brain areas for a given behaviour require strict methodical considerations to be taken into account at the time of designing an experiment and interpreting its results (for a detailed overview see Polania et al., 2018). Behavioural effects following stimulation targeted to a certain region could, for instance, be not functionally specific to the stimulated area itself or may not arise from modulation of the neural computation under examination. In addition to collecting data from a matched no stimulation/sham condition while the animal performs the same task in the same experimental setting, the adoption of control conditions like stimulation of a control brain area is recommended. On the one hand, it is true this additional condition makes the experimental more laborious, costly and time demanding. On the other hand, the addition of this control condition to the design may help to demonstrate that behavioural changes are specific to manipulation of activity in a specific area. We therefore ran a second control experiment (experiment 2) in which we compared credit assignment between

a no stimulation condition and a control stimulation condition in which the TUS was targeted to a control prefrontal area, anterior PFC, a region close to IOFC/12o using the same stimulation protocol. We consistently found a lack of any effect on credit assignment (**Figure 4,5**) and WSLS strategy after stimulation of anterior PFC when compared with a no stimulation condition (**Figure 6,9,10**).

In summary, the two experiments described in this chapter suggest that only IOFC/12o TUS affects credit assignment and that such effects are not present after modulating activity in the adjacent control area anterior PFC, therefore supporting evidence of a causal role of primate IOFC/12o in forming stimulus-outcome associations based on the history of past outcomes.

Chapter 5

Neural basis of credit assignment and manipulation of credit assignment-related activity with focused transcranial ultrasound stimulation

5.1 ABSTRACT

Neural activity somewhere in the ventral prefrontal cortex, is believed to form and update contingent relationships between a choice option and the value associated with its outcome, thereby supporting credit assignment. Previous lesion studies in macaque monkeys put forward this hypothesis. However, the lesions used in these experiments were targeted to a larger portion of cortex. A subsequent fMRI study in macaque monkeys narrowed the potential portion of cortex responsible for credit assignment to the caudal part of the lateral orbitofrontal cortex. Here we disrupted activity in this region using focused transcranial ultrasound stimulation to causally test this hypothesis and to observe the effects of this manipulation on the activity

underlying subsequent choices. After stimulation, activity associated with the value of the chosen option, but not activity encoding reward only, was decreased within a circuit including lateral orbitofrontal cortex, ventromedial prefrontal cortex and amygdala. These neural effects were in line with the behavioural disruption presented in Chapter 4. These neural results suggest that activity in lateral orbitofrontal cortex play an important role in credit assignment.

5.2 INTRODUCTION

The current debate on the neurobiological basis of credit assignment agrees on the key functional role of the inferior prefrontal convexity (Rushworth et al., 2011; Passingham and Wise, 2012; Murray and Rudebeck, 2018). A former hypothesis, assigned to this portion of cortex a predominant role in representing information related to negative feedback such as punishment or errors and contrasted with a role in the representation of positive outcomes in medial and ventromedial prefrontal cortex (PFC) (Kringelbach and Rolls, 2004). However, subsequent investigations have shown that such a strict distinction as a function of valence (positive/negative specialization) does not adequately describe the role played by medial and lateral regions in ventral PFC (Rich and Wallis, 2014). Indeed, neurons in medial orbitofrontal cortex (OFC) and gyrus rectus/14 are responsive to both rewarding/appetitive and punishment/aversive outcomes in both macaques (Morrison and Salzman, 2009) and humans (Plassmann et al., 2010; Tom et al., 2007). In addition, Noonan et al. (2011), showed that lesions targeting either (1) medial OFC

and gyrus rectus/14 or (2) lateral OFC did not have a differential impact on adaptive behavior, as indexed by switching choice preference in subsequent trials, following receipt of either rewards or errors. This lack of segregation between encoding of appetitive and aversive outcomes in medial and lateral portion of OFC is also consistent with the results shown in *Chapter 4*; after ultrasonic disruption of lateral orbitofrontal cortex (lOFC/12o) activity, macaque monkeys show an impairment in learning and adapting future behavior after receiving positive outcomes but not negative outcomes.

We are now starting to better understand the role of different areas of OFC in value-guided learning and decision-making (Wallis, 2011) and an alternative functional dissociation between regions located either medially or laterally to the lateral orbital sulcus seems to be arising. Experiments using tasks which modulate the value of an option through satiety now emphasize the importance of areas 11 and 13 in the central region of the OFC between the medial and lateral sulci (Murray et al., 2015; Murray and Rudebeck, 2018) in the updating of goal values. Electrophysiological studies in the gyrus rectus may be especially important for evaluating and comparing the expected value of different choice options through competition via mutual inhibition in both macaques (Rudebeck and Murray, 2011; Strait et al., 2014) and humans (Hunt et al., 2012). Lateral regions such as lateral OFC and more anatomically lateral structures of the inferior prefrontal convexity (vIPFC), located ventrally to the sulcus principalis are primarily engaged and recruited during associative learning when the of value of separate choice options must be learned through credit assignment (as defined in *Chapter 4*) as a result of the conjoint history of recent choices and their associated

outcomes (Walton et al., 2010; Passingham and Wise, 2012; Rudebeck et al., 2017b).

Currently, however, there is disagreement on which portion of vIPFC/OFC is driving credit assignment. Lesions to this portion of cortex leads to credit assignment deficits (Walton et al., 2010; Noonan et al., 2010; Rudebeck et al., 2017b). However, Chau and colleagues (2015) used functional Magnetic Resonance Imaging (fMRI) in behaving macaque monkeys while they performed a probabilistic reversal learning task and showed an increase in activity within a more narrow region on the lateral bank of the caudal orbital sulcus here referred to as IOFC/12o, when macaques were required to adapt their future decisions based on the outcome of the choices they had previously made.

In Chapter 4 we pushed Chau's results a step further by disrupting activity in this region using focused transcranial ultrasound stimulation (TUS) and demonstrating a causal role of IOFC/12o in credit assignment and adaptive behaviour. Here we combined TUS with functional MRI in 4 macaque monkeys and recorded blood oxygen level dependent (BOLD) activity while they performed the probabilistic three-arm reversal learning bandit task described in *Chapter 4* and compared such activity with that recorded in a no stimulation (sham) condition. We modelled the behavioural and neural data using a computational model (see Chapter 4 and appendix for a detailed description) that enabled us to test whether TUS directed to IOFC/12o modulated the neural activity underpinning credit assignment and contingent learning, here indexed by the neural activity associated with the expected value of the chosen option. Moreover, if IOFC/12o TUS had an

impact on credit assignment activity, we also tested whether the direction of such neural effects was in line with the behavioural deficits. Using whole-brain functional MRI we were further able to look whether TUS effects extended beyond IOFC/12o and modulated neural activity in a larger network of areas involved in value-guided learning and decision-making such as vmPFC/mOFC and amygdala.

5.3 MATERIALS AND METHODS

5.3.1 Subjects

Four male rhesus monkeys (*Macaca mulatta*), here referred to as D, E, F and G, were involved in the experiment (mean weight 12.65 kg; 9.9 years of age). They were socially housed (group of 4 animals) and kept on a 12 hours light dark cycle, with access to water for 12-16 hours on testing days and with *ad libitum* water otherwise. Prior to task training, the animals underwent aseptic surgery to implant an MRI-compatible head post (Rogue Research). The head post was necessary to prevent head movement during the focused transcranial ultrasound stimulation (TUS) and imaging data acquisition. All procedures were conducted under licenses from the United Kingdom (UK) Home Office in accordance with the UK Animals (Scientific Procedures) Act 1986 and with the European Union guidelines (EU Directive 2010/63/EU).

5.3.2 MRI data acquisition

Awake animals were head-fixed in a sphinx position in an MRI-compatible

chair. Functional MRI (fMRI) data were collected using a whole-body 3T MRI scanner with a custom-made four-channel phased-array local receive coil in conjunction with a radial transmit coil (Dr. H. Kolster, Windmiller Kolster Scientific, Fresno, CA, USA). Whole-brain BOLD fMRI data were acquired using a gradient-echo T2* echo planar imaging (EPI) sequence with 1.5 mm³ isotropic voxel size resolution, repetition time (TR) = 2.28 s, echo time (TE) = 30 ms and flip angle = 90°, and reference images for artifact corrections were also collected. Proton-density-weighted images using a gradient-refocused echo (GRE) sequence (TR = 10 ms, TE = 2.52 ms, flip angle = 25°) were acquired as references for body motion artifact correction. T1-weighted magnetization-prepared rapid gradient echo (MP-RAGE) images (0.5 mm³ isotropic voxel size resolution, TR = 2.5 ms, TE = 4.01 ms, inversion pulse time (TI) = 1.1 s, and flip angle of 8°) were acquired in separate anesthetized scanning sessions (anesthesia induction via intramuscular injection of 10 mg/kg ketamine, 0.125 - 0.25 mg/kg xylazine, and 0.1 mg/kg midazolam) prior to the beginning of the testing sessions.

5.3.3 MRI data preprocessing

Magnetic Resonance images were preprocessed and analyzed using tools from the FMRIB Software Library (FSL) (1) and the Magnetic Resonance Comparative Anatomy Toolbox (MrCat; www.neuroecologylab.org). The T1-weighted images were processed in an iterative fashion cycling through a macaque-optimised implementation of FSL's brain-extraction tool (BET), RF

bias-field correction, and linear and non-linear registration (FLIRT and FNIRT) to the *Macaca mulatta* McLaren template in F99 space as implemented in MrCat. The GRE image was used to carry out the T2* EPI image reconstruction by an offline-SENSE reconstruction method (Dr. H. Kolster, Windmiller Kolster Scientific, Fresno, CA, USA).

A T1-weighted group template of the 4 animals included in the study was constructed by cycling through a two-iteration approach of: (1) registration to the initial McLaren template in F99 space; (2) group averaging; (3) registration to the new group template. Specifically, tools from Advanced Normalisation Tools (ANTs), as implemented in MrCat, were used for the creation of the group template. At each step the resulting group template was registered to the source template, thus avoiding drift and retaining registration to F99 space. All coordinates reported (in millimeters) refer to F99 space and results are visualized on the group template.

Despite head-fixation during the EPI acquisition, occasional movements of the animals' limb and body may cause a distortion in the main (B0) magnetic field in a time-varying manner, thereby producing non-linear motion-related artefacts in the phase-encoding direction (anterior-posterior) on a slice by slice basis, like ghosting artefacts, misaligned slices and signal loss. Additional causes of magnetic field distortion and signal loss in non-human primate imaging may be induced by magnetic susceptibility differences in tissue, spatially inhomogeneous magnetic field and spatial differences in the proximity of the head from the RF coils. To correct for these artefacts, using a processing pipeline implemented in MrCat, each slice was registered, first linearly, then non-linearly, to a robust reference

based on EPI volumes from the same timeseries with least distortion. This registration and slice alignment processes were performed volume-by-volume. To avoid overfitting, the degrees of freedom were constrained in several ways: because distortions are the strongest along the phase-encoding direction, here anterior-posterior, only distortions along this direction were considered; registration was initialized using priors from temporally neighboring slices; low-order solutions were preferred over high-order registration (rigid > affine > non-linear); non-linear degrees of freedom were regularized using b-splines. Additional volumes concomitant with any occasional event in the scanning room that may have caused a potential disruption in the main magnetic field were also removed.

Subsequently, the slice-registered average functional image (EPI) was iteratively linearly and non-linearly registered to the high-resolution structural reference (T1-weighted) of each subject, which was then registered to the group-specific template using FMRIB Software Library (FSL) and ANTs tools (Avants et al., 2011; Jenkinson et al., 2012; Smith et al., 2004). Automated brain-extraction of the EPI timeseries was based on brain masks obtained in the high-resolution structural space. Finally, to increase signal-to-noise ratio EPI images were spatially smoothed using a 3 mm kernel (full width at half maximum) and temporally high-pass filtered (cut-off 100 s) to remove slow-moving trends. First-level whole-brain analyses (see below) were performed on the motion-corrected low-resolution images in the original acquisition/native space for each session; the statistical maps for each contrast were then transformed from acquisition/native to the group-template standard space with a single resampling by ANTs for group-level analysis.

5.3.4 Whole-brain fMRI analysis

After transforming each statistical map for each of the four animals from native space to the group-template standard space, group-level statistical analyses of the functional data were performed adopting General Linear Models (GLM) and tools implemented in FSL (Jenkinson et al., 2012; Smith et al., 2004). Analysis at the group-level were performed on the whole-brain using mixed-effects analysis based on FSL's FMRIB's Local Analysis of Mixed Effects (FLAME) 1 + 2. To increase the number of fMRI sessions available and improve the quality of the analysis, sham sessions from both experiment 1 (n = 20; five sessions per animal) and experiment 2 (n = 20; five sessions per animal), thereby creating a pool of 40 no stimulation sessions that were contrasted with 20 IOFC TUS sessions (five sessions per animal). As a consequence of post-preprocessing abnormalities in the EPI signal, one IOFC TUS session was excluded and a final sample of 19 stimulation sessions were used at the group-level. For each GLM, the weight of sham and IOFC TUS sessions was adjusted accordingly (sham = 1; IOFC TUS = 2.1) in order to match the influence of the two groups on the final contrasts despite the different number of sessions in the two conditions. The large number of sessions made it possible to estimate robust mixed-effects statistics and model within-subject and between-subjects variances separately, thereby treating each monkey as a random effect. This modelling approach allows consideration of each monkey as if it was randomly sampled from the general population, therefore making it possible to make inferences not just about the sample of individuals tested but the population from which the samples were drawn. Nuisance regressors were used to capture

additional noise-related variance in the EPI data. These nuisance regressors were intended to: (1) identify noisy volumes after the slice alignment preprocessing step and potential signal loss at the volume level, (2) capture non-linear effects of the motion-related magnetic field distortions, and (3) detect potential remaining variance in motion-related magnetic field distortions.

5.3.5 Whole-brain GLMs

All analyses at the group-level were performed on the whole-brain using mixed-effects analysis based on FSL's FLAME 1 + 2 treating subject as random effects as described above. All parametric regressors were normalized to ensure between-session, between-subject, and between-condition commensurability of the regression coefficients. For each session, one β regression weight was extracted for each regressor. These were then tested for statistical significance across all animals/sessions. Cluster-based inference was performed in all analyses based on a cluster-defining threshold of $Z > 2.3$ with a Gaussian Random Field (GRF)-corrected threshold of $p = 0.05$. According to this model, the following GLMs were implemented.

GLM 1 – Representation of the expected value of the chosen option

The first fMRI analysis was set up on the basis of the behavioural results described in *Chapter 4* that reflected an impairment in credit assignment after IOFC/12o TUS. This result suggest a deficit in the animal's capacity to credit the value of the outcome received after making a decision to the correct

choice option, thereby leading to an impairment in the animal's capacity to update the value estimate of the chosen option and to use such a representation to guide behavior in the following trials. We, therefore, use blood-oxygen-level-dependent imaging, or BOLD-contrast imaging, as measured by fMRI, to investigate whether changes in the neural signal correlating with the "CHOSEN VALUE" were present after IOFC TUS. A Haemodynamic Response Function (HRF) of 3 seconds with a gamma convolution function was used in the following GLMs to capture the peak of the BOLD signal in the rhesus macaque monkey brain that resembled that used in previous studies (Chau et al., 2015; Papageorgiou et al., 2017; Fouragnan et al., 2019). Regressors were built using the reinforcement learning model described in *chapter 4* and in *appendix* (Wittman et al.). *GLM1* included an unmodulated decision constant regressor (*DEC*) time-locked to the onset of the decision (when stimuli appeared on screen) with a duration set to the reaction times of each trial. Two additional binary constant regressors were time-locked at the time of feedback onset to capture variance in the BOLD signal associated with the type of outcome, namely receipt (*REW*) or non-receipt (*NOREW*) of reward. Both constant regressors had a duration equal to a Dirac Delta Function of 100 milliseconds. To capture variance explained by distortions in the magnetic field caused by the response action, two unconvolved categorical regressors for leftwards (*UNC_left*) and rightwards (*UNC_right*) responses were time-locked at the time of feedback with a duration of TR (2.28 seconds). All the following regressors were convolved with the HRF described above. A parametric regressor aimed to capture choice location (*c_Clo*) was normalized and time-

locked at the time of response. Three fully parametric regressors were time-locked at the time of decision and were modulated by the expected value of the chosen (*ChoV*), unchosen (*UncV*), and unpresented (*UnpV*) options. Two parametric regressors were intended to capture choice location trace (*locT*) and the choice trace associated with the unpresented option (*unpCT*) were normalized and time-locked at the time of decision. An additional parametric regressor modelled the comparison between the choice traces associated with chosen and unchosen options (*choT-uncT*) was also time-locked at the time of decision and normalized. Two parametric regressors, time-locked at the time of feedback onset and normalized, aimed to capture the variance in the BOLD signal explained by the reward trace in trials in which reward was received (*rewT_reward*) or not (*rewT_noreward*) for making a choice. Lastly, 4 contrasts were additionally set up and looked respectively at: 1) comparison between the expected value of the chosen option (*choV*) and the expected value of the unchosen option (*uncV*); 2) comparison between BOLD signal associated with reception of a reward (*REW*) and lack of reception of a reward (*NOREW*); 3) reward prediction error calculated using a standard simple reinforcement learning approach (*RLsimple* summarized in chapter 4 and described in detail in the appendix in Wittman et al.); 4) reward prediction error modulated by a reward trace signal *R-trace model* summarized in chapter 4 and described in detail in the appendix in Wittman et al.).

$$\begin{aligned}
 \mathbf{fMRI\ GLM1} = & \beta_1 DEC + \beta_2 choV + \beta_3 uncV + \beta_4 unpV + \beta_5 choT - uncT + \\
 & \beta_6 unpCT + \beta_7 locT + \beta_8 REW + \beta_9 NOREW + \beta_{10} c_{Clo} + \beta_{11} rewTreward + \\
 & \beta_{12} rewTnoreward + \beta_{13} UNC_{left} + \beta_{14} UNC_{right} + \varepsilon
 \end{aligned}$$

GLM 2 – Representation of Win-Stay / Lose-Switch signal

In the second fMRI analysis we aimed to identify eventual neural changes associated with the use of WIN-STAY / LOSE-SWITCH strategy after stimulation of the IOFC/12o. The HRF convolution function used was the same as for *GLM1*. A similar unmodulated decision constant regressor (*DEC*) with a duration set to the reaction times of each trial was used but time-locked to the feedback onset in the current GLM. Three parametric regressors, capturing variance in the BOLD signal underlying the encoding of a WIN-STAY / LOSE-SWITCH strategy at the previous trial (*WLST1*), two trials in the past (*WLST2*), and three trials in the past (*WLST3*) were normalized and time-locked at the time of feedback. Two additional constant regressors were time-locked at the time of feedback onset with a duration of 100 milliseconds to capture variance in the BOLD signal associated with the receipt (*REW*) or non-receipt (*NOREW*) of reward. Two unconvolved categorical regressors for leftwards (*UNC_left*) and rightwards (*UNC_right*) responses were time-locked at the time of feedback with a duration of TR (2.28 seconds) to capture motion-induced distortion in the magnetic field during the response.

$$\begin{aligned} \textbf{fMRI GLM2} = & \beta_1 DEC + \beta_2 WSLST1 + \beta_3 WSLST2 + \beta_4 WSLST3 + \beta_5 REW \\ & + \beta_6 NOREW + \beta_7 UNC_{left} + \beta_8 UNC_{right} + \varepsilon \end{aligned}$$

5.3.6 Region-of-interest (ROI) analyses

Additional analyses were performed on ROIs centered on regions showing the highest significant differences in ventral prefrontal cortex activity between sham and IOFC/12o TUS sessions. As these analyses rely on coordinates estimated from contrast activation maps resulting from previously mixed-effects whole-brain GLMs, they are here reported primarily for visualization purposes.

ROIs were manually defined as a 3-voxel mask on the group-template image in F99 space (Van Essen et al., 2012), warped back to each individual session's EPI image. Activity within the mask was averaged across all voxels and the mean time-course was extracted. ROIs were placed over the IOFC/12o (left hemisphere: $x = -15.59$, $y = 10.06$, $z = -1.51$); gyrus rectus / area 14 (right hemisphere: $x = 2.52$, $y = 20.62$, $z = 3.02$) and amygdala (left hemisphere: $x = -10.56$, $y = 4.52$, $z = -9.05$). Average values of activity within each ROI were extracted from each session. A total number of 40 sham and 19 IOFC/12o sessions were used. Effect sizes for single sessions and average within each condition [sham, IOFC/12o TUS] were extracted for visualization as mean and standard error of the mean (*s.e.m.*) estimates.

5.4 RESULTS

5.4.1 Whole-brain changes in the representation of the chosen value signal after lateral OFC TUS

We used fMRI data to examine the potential existence of changes in the neural signal that may underlie the observed behavioural differences in credit assignment and win-stay/lose-switch strategy described in *Chapter 4*. At the first-level, we ran *GLM1* (*fMRI GLM1*) for each session within sham and stimulation conditions separately. Then, at the second-level, we ran a two-sample whole-brain fMRI GLM, time-locked at the decision phase, to investigate whether differences in neural activity underlying the expected value estimate signal of the chosen option were present in the brain after TUS targeted to the IOFC/12o. To control for both between-subjects and within-subject variance, in *GLM1* at the second-level we set *CONDITION* [sham, IOFC/12o TUS] as fixed-effect and *SUBJECT* [4 monkeys] as random-effect. To estimate differences in BOLD signal, 40 sham sessions were contrasted with 19 IOFC/12o TUS sessions in this GLM. To prevent any artefactual result arising from the different number of sessions between the two conditions, the weight of sessions within each condition was adjusted according to their respective number in sham and TUS (see Materials and Methods for full details).




Figure 1. TUS disruption of IOFC/12o is associated with a decrease in activity encoding the expected value of the chosen option. A whole-brain analysis testing for voxels associated with a decrease in neural activity correlated with the trial-by-trial estimates of the expected value of the chosen option after IOFC/12o TUS. The fMRI analysis was time-locked to the decision phase of each trial (in red-yellow: cluster-corrected, $Z > 2.3$, $P < 0.05$; control sham sessions, $n = 40$; IOFC/12o TUS sessions, $n = 19$; Weight (w) of each session in both conditions was adjusted to $w = 1$ for sham sessions and $w = 2.1$ for IOFC/12o TUS sessions). Differential values in whole-brain BOLD signal among the two conditions were estimated as SHAM > IOFC/12o TUS. Reduced activity encoding the expected value of the chosen option was observed after IOFC/12o TUS in orbitofrontal, insular and amygdala structures.

Consistent with the behavioural changes described in Chapter 4, we observed differences in neural activity associated with perturbation of IOFC/12o. Significant clusters of activity correlating with trial-by-trial estimates of the expected value of the chosen value were present across a circuit extending between ventral prefrontal cortex OFC, and amygdala. Specifically, stronger activity in the sham condition relative to the IOFC/12o

TUS was identified in the gyrus rectus / area 14 (**Figure 1A**; panels *i* and *iii*), medial OFC (**Figure 1A**; panels *i*, *ii*, *iii*), central OFC (**Figure 1A**; panels *iii* and *v*), lateral orbital sulcus (**Figure 1A**; panels *iv* and *v*), IOFC/12o (**Figure 1A**; panels *iv*, *v* and *vi*) and amygdala (**Figure 1B**) ($Z > 2.3$, whole-brain cluster-wise correction, $P < 0.05$).

Table 1 Regions-of-interest analysis for neural correlates of Expected Value Estimates of the Chosen Option				
ROI	Hemisphere	X	Y	Z
Gyrus rectus / area 14	Left	2.52	20.62	3.02
Lateral OFC / area 12	Right	-15.59	10.06	-1.51
Amygdala	Left	10.56	4.52	-9.05

5.4.2 Lateral OFC TUS disrupts signal encoding the expected value of the chosen option within the amygdala-OFC circuit

To illustrate the significant changes in activity before and after IOFC/12o TUS, in a second set of analyses we extracted the average time-courses in three ROIs within the amygdala-OFC network. The ROIs were centered on coordinates showing the peak differences in activity correlating with the expected chosen value in the gyrus rectus, IOFC/12o, and amygdala (**Table 1**; see Materials and Methods for voxels coordinates and full details of the ROI analysis). Note that the aim of these figures is visualization of the effects

of stimulation within the amygdala-OFC network rather than further statistical inference (further statistical testing of effects extracted from an ROI obtained from the same contrast would be circular). The visualization confirmed that there were decrements in chosen value-related activity across all three regions investigated (**Table 1; Figure 2**).




Figure 2. Regions-of-interest analysis within the amygdala-OFC network for the expected value of the chosen option. ROI analyses (right panels) were performed in the gyrus rectus/14 (top row), left IOFC/12o (middle row) and left amygdala (bottom row) based on the whole-brain fMRI analysis (Figure 1) that showed a decrease, after bilateral IOFC/12o TUS (red), in activity correlated with the expected value of the chosen option (left panels) compared to control sham (blue). Because these ROI analyses are

directly driven by the whole-brain fMRI analysis in GLM, they consistently show larger expected value signals in sham control sessions compared to IOFC/12o TUS. Left panels: The fMRI analysis was time-locked to the decision phase of each trial (in red-yellow: cluster-corrected, $Z > 2.3$, $P < 0.05$; control sham sessions, $n = 40$; IOFC/12o TUS sessions, $n = 19$; Weight (w) of each session in both conditions was adjusted to $w = 1$ for sham sessions and $w = 2.1$ for IOFC/12o TUS sessions). Differential values in whole-brain BOLD signal among the two conditions were estimated as SHAM > IOFC/12o TUS.

5.4.3 Lateral OFC TUS does not affect the reward signal

After observing post-IOFC/12o TUS changes in the neural activity underlying the expected value of the chosen option, in another analysis we looked at whether TUS targeted to this region also affected a different signal – the activity pattern that differentiated receipt or non-receipt of reward. Thus the reward signal was estimated as the difference between the reward and no reward regressors (reward - no reward). Time-courses for this contrast were extracted from ROIs within the same amygdala-OFC network reported in *Figure 2*. Therefore, the ROIs were extracted from the coordinates showing the peak differences in the activity underlying the expected chosen value: gyrus rectus, IOFC/12o, and amygdala (**Table 1**; see Materials and Methods for voxels coordinates and full details of the ROI analysis). At the behavioural level, as showed in *Chapter 4*, we had found that the encoding of reward was not affected by IOFC/12o TUS, whereas the ability of the animal to optimally perform credit assignment was. Therefore, we should not expect changes in the reward-related activity in this analysis but disruption of the encoding of

the expected value of the chosen option. This is indeed what we observed. IOFC/12o TUS did not affect the neural encoding of reward, showing results consistent with what had already been observed behaviourally and reported in *Chapter 4* (Table 1; **Figure 3**). The brain did not show activity changes even in anatomically highly connected areas such as the amygdala.




Figure 3. Regions-of-interest analysis within the amygdala-OFC network for the reward signal. ROI analyses (right panels) looking at neural activity encoding reward (reward minus no reward contrast) were performed for sham (blue) and IOFC/12o (red) TUS sessions in the same areas that showed a decrease, after bilateral IOFC/12o TUS, in activity correlated with the expected value of the chosen option. The ROI analyses were, therefore, performed again in the gyrus rectus/14 (top row), left IOFC/12o (middle row) and left amygdala (bottom row). Sham and IOFC/12o TUS session did not show significant differences in the encoding of reward

as can be seen in the whole-brain fMRI analysis (left panels). Left panels: The fMRI analysis was time-locked to the feedback phase of each trial (in red-yellow: cluster-corrected, $Z > 2.3$, $P < 0.05$; control sham sessions, $n = 40$; IOFC/12o TUS sessions, $n = 19$; Weight (w) of each session in both conditions was adjusted to $w = 1$ for sham sessions and $w = 2.1$ for IOFC/12o TUS sessions). Differential values in whole-brain BOLD signal between the two conditions were estimated as SHAM > IOFC/12o TUS.

5.4.4 Conjunctive differences of expected value of the chosen option and reward between sham and lateral OFC TUS

During learning the representation of each choice-specific expected value estimate must be updated whenever an outcome is received for selecting that option. After the animal makes a choice, the value of the chosen option is updated by the prediction error between its expected value in the previous trial and the actual outcome received in the current trial (reward).

In the next analysis we identified brain regions where differences in activity between the no stimulation condition and IOFC/12o TUS condition were both related to the value of the choice taken during decision making (chosen value) and to the outcome that followed (reward/non-reward). A conjunction test was used to locate brain areas in which differences in activity between the two conditions were related both to the representation of expected value (time-locked at the decision phase) and outcome (time-locked at the time of feedback: receipt/non-receipt of reward). Specifically, for the reward variable we looked at the difference between the BOLD signal correlating with reward and no reward (*GLM1*). We also carried out a second-level GLM comparing the conjunction maps obtained from the first

separate first level analyses of the sham and IOFC/12o TUS conditions. Threshold was set at 2.3 for the z-statistics and $P = 0.05$.




Figure 4. Conjunction analysis. Conjunction analysis of BOLD signal correlations with the value estimate of chosen option (difference between sham and IOFC/12o TUS) and BOLD signal correlations with “reward minus no reward”. Threshold of the conjunction analysis was set at 2.3 for the z-statistics and $P = 0.05$.

Differences in activations in the medial OFC, lateral OFC, lateral orbital sulcus, orbital periallocortex and amygdala were positively correlated with both the “expected value signal” and “reward minus no reward” contrast. The conjunction effect was significantly stronger in the sham than the IOFC/12o TUS condition in all these areas (**Figure 4**).

5.4.5 Role of the lateral OFC in adaptive behavior within a volatile environment

In Chapter 4 we showed that animals exhibited an impairment in adapting their behavior contingently to the receipt of positive or negative outcomes. This research question was addressed by looking at patterns of staying or switching with a given choice as a function of whether a reward had or had not been received for choosing it. We therefore ran a second whole-brain fMRI analysis (*GLM2*; see Materials and Methods for full details) aiming to look at differences between sham and IOFC/12o TUS sessions in the encoding of a WIN-STAY/LOSE-SWITCH (WSLS) strategy.

As was the case for *GLM1*, *GLM2* (*fMRI GLM2*) was first run at the first-level for each session within the two conditions separately. At the second-level we ran a whole-brain fMRI analysis time-locked at the feedback onset phase, to search changes in the neural signal encoding WSLS behavior. To control for both between-subjects and within-subject variance, also in *GLM2* at the second-level we set CONDITION [sham, IOFC/12o TUS] as a fixed-effect and SUBJECT [4 monkeys] as a random-effect. 40 sham sessions were contrasted with 19 lateral OFC TUS sessions in this GLM and the weight of sessions within each condition was adjusted according to their respective number in sham and TUS (see Materials and Methods for full details). Differences in neural activity between sham and lateral OFC TUS in the past trial were examined in this second-level GLM.

Activation differences between sham and IOFC/12o TUS at the time of outcome were observed in several regions. In the sham condition, compared to the IOFC/12o TUS condition, the WSLS signal correlated with stronger

activity in medial and lateral OFC. Intriguingly, the stronger activation in IOFC/12o maps well onto the sham-IOFC/12o TUS in chosen value encoding already reported (GLM1). The directionality of the strength of such activation is also consistent across the two analysis, with BOLD signal within this region being lower in the IOFC/12o TUS condition (**Figure 5**). Although activation differences could still be observed for decisions consecutive to the next one, they were inconsistent and spatially scattered already on unthresholded maps.




Figure 5. TUS disruption of IOFC/12o is associated with a decrease in activity encoding the Win-Stay/Lose-Switch signal. A whole-brain analysis testing for voxels associated with a decrease in neural activity correlated with the trial-by-trial estimates of the Win-Stay/Lose-Switch signal (WSLS) after IOFC/12o TUS. The fMRI analysis was time-locked to the outcome phase of each trial (in red-yellow: cluster-corrected, $Z > 2.3$, $P < 0.05$; control sham sessions, $n = 40$; IOFC/12o TUS sessions, $n = 19$; Weight (w) of each session in both conditions was adjusted to $w = 1$ for sham sessions and $w = 2.1$ for IOFC/12o TUS sessions). Differential values in whole-brain BOLD signal among the two conditions were estimated as SHAM > IOFC/12o TUS and looked at the impact of the WSLS signal over choices made in the consecutive trial.

5.5 DISCUSSION

Everyday survival requires the constant making of decisions between different potential choices associated with different outcomes given a constantly changing internal motivational state. Such ecological pressure necessitates the primate brain to represent the current value of each option at the time of decision and to constantly compare the values of possible options in the environment before making a decision. In addition, it is necessary to update value estimates whenever the choice-outcome contingencies change. To do this, the ability to associate a given option with the precise consequences that follow its choice is essential. Overall, in this *Chapter* we show that disruption of activity in IOFC/12o with TUS was associated with changes in the underlying neural activity encoding cognitive computations supporting credit assignment. Such neural results are consistent with the behavioural deficits in value-guided learning and decision-making induced by IOFC/12o TUS.

The minimally-invasive TUS enabled us to interleave stimulation and control sessions in a pseudo-random order and to counterbalance such order across animals. We showed that the offline protocol used in this study can produce clear disrupting effects both on behaviour (*Chapter 4*) and neural activity (current *Chapter*). The temporal and reversible properties of TUS enabled us to control for any potential “order effect” or “post-surgical re-organization” that might in theory occur when aspiration, excitotoxic or pharmacological lesions are made within the same animals and pre-operative post-operative comparisons are made (Machado and Bachevalier, 2006; Meunier et al., 1999; Wilson et al., 2014).

One alternative approach is to instead compare the behaviour of separate groups of animals, one with lesions to a certain brain area and one without, while they perform the same task. Although controlling for the “order effect”, this methodological strategy necessitates careful matching and controlling for individual differences across sham and lesion groups.

On the other hand, pharmacological approaches that aim to temporally inactivate a certain area are limited by other variables such as the size of the target region. The amount of agent to be injected in a region in order to have an effect is dependent on the size of the region itself and so care must be taken to avoid systemic toxic effects. To summarize, TUS does not constitute only a remarkable refinement in the welfare of the animals but its temporal and reliable properties, in addition to with its potential to use more conservative statistics, makes it an innovative tool able to improve both the internal (reliable cause-and-effect relationship between manipulation and resulting behaviour) and external (how generalizable to the wider population the results of the study are) validities of a study. Such technical improvement is notable at a time in which neuroscience is facing criticisms arising from difficulties in replicating previous neuroimaging and computational experiments (Fanelli, 2018; Gilmore et al., 2017; Masouleh et al., 2019).

Despite the initial promising results, it must be considered that TUS, because of its relative novelty, remains relatively poorly understood in many ways. There is evidence that it may affect ion channels in the neuron cell membrane (Blackmore et al., 2019; Jerusalem et al., 2019; Kubanek et al., 2018) but a full picture of precisely how this occurs is yet to emerge. In summary, while studies such as the current one demonstrate that it can

affect neural activity and behaviour in relatively precise ways, the factors that might limit its applicability have yet to be determined. For instance, the application of TUS may show a predominant disruption of activity in the targeted area but additional cascade effects within the close anatomical circuit may be difficult to control for, as shown in *Chapter 3*. Moreover, the physiological status of the animal or neural noise induced by factors external (i.e. skull tissue heterogeneities) to the TUS may also affect its action. Although intriguing and consistent, temporal and spatial dynamics of TUS-induced effects still require a better understanding, especially in relation to the specific stimulation protocol and parameters used in each experiment (and species) (see Blackmore et al., 2019; Jerusalem et al., 2019; Fomenko et al., 2018 for a longer discussion). Therefore, invasive lesion methods should still be considered the “gold standard” approach to causally define the functional role of a brain area or to test hypothesis previously made using minimally-invasive approaches like TUS.

Comparison of neural activity encoding the expected value of the chosen option between sham and IOFC/12o TUS identified changes in activation after TUS within an amygdala-OFC circuit. Areas IOFC/12o and 13, which cover lateral and medial banks of the orbital sulcus respectively, showed a significant decrease in activity after TUS. Changes in the same directions extended also more rostrally in the gyrus rectus/14 and rostral OFC area 11. Caudally, a significant reduction of activity was identified in the rostral amygdala. These results are consistent with previous research identifying these structures as core hubs supporting value-guided decision-making (Morrison and Salzman, 2009; Padoa-Schioppa and Assad, 2006;

Rudebeck and Murray, 2011; Saez et al., 2017; Tremblay and Schultz, 1999; Tsujimoto et al., 2009). Notably, TUS effects on choice value extended beyond the stimulated area IOFC/12o. This is expected as the TUS-induced deficits in the capacity of the animal to learn associations between choice options and outcomes have repercussion on the ability of the animal to represent and update correctly the expected value of choices. This manifests as a deficit in choice-reward contingent learning that was reported in *Chapter 4*. Decreases in activity within the amygdala-OFC circuit at the neural level was associated with a deficit in integration of the outcomes received during recent trials (**Figure 1,2**). To further confirm the functional focality of the neuromodulation effects when TUS was targeted to IOFC/12o, neither behavioural nor neural changes within the amygdala-OFC circuit were observed in the animals' ability to encode represent the receipt of reward (**Figure 3**).

We next investigated the whether the disruption of neural activity was specific to the expected value of the chosen option or was affecting the encoding of reward as well. We hypothesised that the reward signal should not be affected after IOFC/12o TUS as no post-stimulation changes were observed at the behavioural level (*Chapter 4*). This set of analyses showed that the lack of significant differences in reward was present also at the neural level (**Figure 3**). This result is consistent with the behavioural effects observed in *Chapter 4* and highlights a specificity in the cognitive and neural effect of TUS. Indeed, we replicated at the neural level that credit assignment but not reward only was affected by modulation of activity in IOFC/12o.

In reinforcement learning computational models like the one used in this study, the value estimate for a choice (the expected probability of receiving a reward when choosing that option) is updated through a prediction error estimate modulated by learning rate. The prediction error estimate is the difference between the prior probability estimate associated with the chosen option and the actual value associated with the outcome received at that trial (Rescorla and Wagner, 1972; Sutton and Barto, 1992; Wittman et al., Appendix). This model therefore assumes a strict relationship between the expected value of a choice option and the recent receipt of reward/no reward convergence of activity representing these two quantities with a given brain structure. We therefore tested whether there was any brain region showing differences, with and without TUS, in the conjunctive encoding of the value of the choice taken at the decision phase and the encoding of the subsequent outcome. We found activation difference in an amygdala-OFC circuit (**Figure 4**) similar to the one encoding the expected value of the chosen option (**Figure 1,2**).

In a final analysis we looked for neural changes underlying the deficits in the adaptive use of a WSLS strategy in a stochastic environment following IOFC/12o TUS observed in *Chapter 4*. Indeed, we previously showed that if the animals are unable to correctly learn contingencies between a certain option and the outcome following its choice, they are also unable to adopt an optimal WSLS strategy. In an additional whole-brain fMRI analysis we compared the neural activation associated with WSLS across the two conditions searching for differences on its impact in next decisions following TUS. A striking result was a strong reduction in activity encoded exactly in

area IOFC/12o when we looked at the influence on the following trial (**Figure 5**). The role of this area in WSLs found here replicates what Chau et al. (2015) found using a similar reversal learning task. In addition, our results suggest that this signal may causally drive credit assignment and WSLs as it shows a reduction after IOFC/12o TUS that is also associated with concomitant behavioural deficits. Whether the activity change can be linked more specifically to a win-stay or a lose-switch strategy will be addressed in future investigations.

WSLS and credit assignment are cognitive processes related to one another and could be described as model-free and model-based strategies, respectively. WSLs is a strategy used to guide behaviour on the basis of outcome feedback during the initial exploration phase of a task and enables the animals to start linking each choice option with its relative probability to lead or not to a reward in the future. This model free strategy, therefore, enables an initial exploration of the contingencies between choice options and reward in the task. Through a WSLs strategy, the animals start assigning credit to each choice option in a flexible way as they experience several trials. Once enough trials have been experienced and the animals have created a stable model of the world representing the contingencies between each option and the associated outcomes, then a model-based strategy is used to guide behaviour and integrate feedbacks over consecutive past choices. This relationship between WSLs and credit assignment may explain the consistency disruption observed for both strategies in our experiment after IOFC/12o TUS.

The study presented in the current chapter focused on the neuromodulatory effects of TUS targeted to IOFC/12o on contingent learning and showed that when activity in this region is disrupted transcranially with ultrasound, the underlying neural activity driving credit assignment is affected. However, as we show in Wittman et al. (Appendix), non-contingent learning computations also plays a role in determining value-guided learning and modulates contingent learning effects. More specifically, in Wittman et al. (appendix) we defined non-contingent learning primarily as the influence of a memory trace state representing the average reward rate of the environment the animal has been in over the recent trials. Future investigations should also look at the modulatory effects of IOFC/12o TUS on these non-contingent learning mechanisms. Changes in the rate of reward availability in the environment over a short number of consecutive choices may mean that this memory trace has an impact on associative learning within the amygdala-OFC circuit may alter its impact.

Chapter 6

General Discussion

The investigation of orbitofrontal cortex (OFC) and amygdala as independent brain structures has provided us with an important understanding of the role these regions have in credit assignment and, more generally, in cognitive computations underlying value-guided associative learning (Baxter et al., 2000; Costa et al., 2016; Hernádi et al., 2015; Hunt et al., 2012; Izquierdo et al., 2005; Jocham et al., 2012; Luk and Wallis, 2013; Murray and Rudebeck, 2018; Namburi et al., 2015; Noonan et al., 2011; Papageorgiou et al., 2017; Rudebeck et al., 2013a; Rushworth et al., 2011; Wilson et al., 2014; Zangemeister et al., 2016). However, it is becoming increasingly clear that only their consideration as two structures that are part of the same circuit can lead to a better understanding of how their interactive dynamics enables the primate brain to correctly assign value to a choice option during learning and to successfully use this information to guide future behaviour (Murray and Wise, 2010; Chau et al., 2015; Rudebeck et al., 2017b; Lichtenberg et al., 2017). In this thesis, I aimed to investigate the synergic communication of OFC and amygdala in the primate brain by investigating their interactions at the anatomical (*Chapter 2*), functional (*Chapter 3*), behavioural (*Chapter 4*) and neural (*Chapter 5*) level.

Most of our current knowledge of human brain anatomical connections relies on invasive and laborious investigations in the macaque monkey model. In *Chapter 2*, I started off by reconstructing the pathway interconnecting the amygdala and surrounding rostral temporal regions with areas in OFC and in the entire prefrontal cortex (PFC) within macaque and human brains. Using non-invasive diffusion magnetic resonance imaging (MRI) in both species, I showed that two white matter pathways project between amygdala/temporal and PFC areas in both macaque monkeys and humans: the ventral amygdalofugal pathway and the uncinate fascicle. These results are consistent with what has already been reported by previous tract-tracing investigations in the post-mortem monkey brain (Lehman et al., 2011; Oler et al., 2017). Here, however, I reconstructed and compared for the first time the anatomical architecture and connectivity in PFC of these two tracts in both monkeys and humans. In line with previous tracer density estimations (Ghashghaei and Barbas, 2002; Ghashghaei et al., 2007; Barbas et al., 2007), I demonstrated that amygdala-OFC connectivity is stronger in more caudal OFC and subgenual regions. Furthermore, I extensively described the medial-lateral organization of the two tracts. The ventral amygdalofugal pathway coursed along the medial bank of the ventral PFC. Conversely, the uncinate fascicle primarily ran in a more lateral position at the level of subcortical and posterior orbital areas but then branches through both medial and lateral regions in the rostral part of ventral PFC. Notably, this organization is preserved in both monkeys and humans and may have important functional implications for value-guided associative learning and credit assignment. Indeed, the connectional features of the two tracts

suggest that the ventral amygdalofugal pathway primarily carries information from areas involved in the processing of arousal and outcomes (Morrison et al., 2011; Namburi et al., 2015; Rudebeck and Murray, 2014; Saez et al., 2017), thereby potentially providing an anatomical path bridging amygdala and OFC and supporting an exchange of information on the current internal state of the subject. Conversely, the uncinate fascicle primarily interconnects temporal regions involved in processing visual information (Browning and Gaffan, 2008; Clark et al., 2012; Eacott and Gaffan, 1992; Eldridge et al., 2016; Sugase-Miyamoto et al., 2008; Sugase-Miyamoto and Richmond, 2005) with lateral and rostral PFC areas, including lateral OFC (IOFC/area12), where activity is associated with value representation and credit assignment (Rudebeck et al., 2017b; Chau et al., 2015; Walton et al., 2010; Noonan et al., 2017). The medial-lateral organization of the ventral amygdalofugal pathway and uncinate fascicle maps onto the separation of ventral PFC into a medial visceromotor and an orbital sensory network (Carmichael et al., 1995; Ongur and Price, 2000). I showed that the two tracts, similarly to the two networks described by Price and colleagues, are clearly distinct but yet show several interconnections in medial OFC and ventromedial PFC regions. Moreover, the convergence of the two tracts in OFC may play a key role in the representation of choice options in a common value currency (Padoa-Schioppa and Assad, 2006; Padoa-Schioppa and Assad, 2008) or in creating a representation of task variables as an abstract cognitive map (Schoenbaum et al., 2011; Wilson et al., 2014). Lastly, the different types of information carried by the two tracts, as a result of differences in their connectional features and interaction in OFC may support

credit assignment. In the future studies examining disconnection of either tracts and observation of the subsequent behavioural deficits following each lesion might help elucidate the functional significance of the connectional results reported in *Chapter 2*.

Although providing important anatomical information and a common tool to study multiple species using the same approach, diffusion MRI and tractographic algorithms do not allow investigation of the functional properties of a network. In *Chapter 3*, I therefore decided to reconstruct the activity coupling (also known as functional connectivity) of the amygdala and anterior cingulate cortex with regions mainly located in PFC. As extensively discussed throughout *Chapters 1* and *3*, current neuromodulatory approaches do not allow efficient targeting of deep cortical or subcortical areas without the use of surgical intervention. A common practice is, therefore, to using lesions and pharmacological inactivations to investigate the causal contribution of specific brain regions to behaviour. This invasive approach, however, requires a trade-off between using (1) a larger number of animals (especially if lesions of a control region needs to be performed in order to determine the focality of the effects and (2) using a lower number of animals by renouncing such a control condition. Focused transcranial ultrasound stimulation (TUS) represents a good alternative for the causal investigation of the functional contribution of brain areas to behaviour. Indeed, its minimally-invasive and reversible properties allow for more conservative and powerful within-subject statistics by requiring a lower number of animals and permitting the use of the same animal across multiple experimental conditions. The potential of sound waves to produce acoustic

energy has been known since the ancient historical studies of Pythagoras, Aristotle and Galileo Galilei (Cartwright et al., 2001; Jagannathan et al., 2009) but it was only in the 1880 that Pierre and Jacques Curie demonstrated the generation of piezoelectricity, that is the creation of electric polarization, as a result of application of pressure to crystals (Curie and Curie, 1880), a principle that was later used by Fry and colleagues for ultrasound stimulation of the central nervous system (Fry et al., 1958, 1955; Fry and Fry, 1960). In *Chapter 3*, I exploited the mechano-electric potential of acoustic energy to interact with neural tissue at low frequency (250kHz) to investigate, in two separate experiments, whether TUS was capable of modulating the activity coupling of amygdala and anterior cingulate cortex in the primate brain. By implementing a 40 seconds stimulation protocol applied offline (before the data collection began) in anaesthetized monkeys, I showed that the neuromodulatory effects were primarily focal to the stimulated region in both experiments. Moreover, I showed and discussed why the observed stimulation effects were unlikely to be the consequence of auditory or somatosensory stimulation, as suggested by Guo (2018) and Sato (2018). Notably, the same TUS parameters but applied to different structures in the brain resulted in a different proportion of ultrasound-induced acoustic force applied to brain tissue. These effects are likely due to a heterogeneous interaction of the sound waves with bone tissues at the stimulation targets, which results in a different diffraction and reverberation of standing waves along the trajectory of the ultrasound beam. This observation requires careful consideration when designing a TUS experiment as the interaction of sound waves with other areas along or beyond the trajectory

between the ultrasound transducer and the target may act as confound and result in unforeseen behavioural deficits. In all the TUS experiments reported in this thesis, I used a single-element piezoelectric transducer for the ultrasound stimulation as its efficacy in modulating neural activity and behaviour had previously been shown (Constans et al., 2017; Deffieux et al., 2013; Wattiez et al., 2017). However, in future studies some of the limitations highlighted above may be dealt with by using piezoelectric transducers composed by an array of transducers or adopting target-specific custom lenses in order to improve focality of the stimulation and reduction in the diffraction and perturbation of the transmitted sound waves caused by skull and bone tissue heterogeneities. Moreover, in both regions, the effects lasted for up to two hours after the stimulation suggesting a neuroplastic mechanism as has sometimes previously been argued (Yoo et al., 2011; Dellapiazza et al., 2018; Legon et al., 2014; Verhagen et al., 2019). Although some recent studies indicate that the TUS effect may be mediated by the mechanical action of acoustic energy on the cell membranes (Kubanek et al., 2018, 2016a; Tyler et al., 2008), the precise nature of the primary mechanism driving these plastic effects is still under discussion and should be elucidated in future investigations (Blackmore et al., 2019; Fomenko et al., 2018; Jerusalem et al., 2019; Kubanek, 2018; Tyler et al., 2018).

The understanding of the relationships among anatomy, neural activity and behaviour is the primary goal of neuroscience. The use of correlational measurements like neuronal recordings and neuroimaging approaches can lead to the formulation of hypotheses about brain function. However, they do not allow for causal inference: correlations between brain activity and

behaviour do not demonstrate that the neural activity is causally responsible for the behaviour. Neural activity changes in one area observed using such methods, for instance, may not be sufficient or necessary to induce changes in behaviour. The results showed in *Chapter 3* demonstrate a remarkable potential of TUS in successfully modulating neural activity in a transcranial and temporary fashion even in regions deep into the brain in subcortical areas and in areas deep within sulci, in animals with complex cortical gyrification such as primates. In *Chapters 4* and *5* I exploited these properties and used TUS to investigate whether a specific portion of the caudal ventral PFC, the lateral OFC (IOFC/12o), plays a causal role in credit assignment. The potential importance of this portion of OFC was suggested by previous researchers using either functional MRI (fMRI) measurements (Chau et al., 2015) or lesions of the OFC (Walton et al., 2010; Noonan et al., 2010) and ventrolateral PFC in macaques (Rudebeck et al., 2017b) and human patients (Noonan et al., 2017). By combining TUS with whole brain fMRI while macaque monkeys performed a reversal learning object discrimination task, I was able to observe, within the same monkeys, the impact of TUS applied bilaterally to IOFC/12o on credit assignment and compare it with both a no stimulation and an active control stimulation condition. Disruption of IOFC/12o activity only resulted in credit assignment deficits, leaving the encoding of rewards and choice options alone unaffected. This result is consistent with what has already been reported in the work of Rudebeck (2017b), and Noonan (2017), however in those studies the lesions occupied a larger portion of ventral PFC. In line with this observation, I also showed that after the IOFC/12o stimulation, the animals

were unable to flexibly adapt their behaviour after receiving a reward but did not show any impairment in using negative feedback to guide future behaviour. This result partly disagrees with Chau's claim that activity in IOFC/12o guides a change in the choice selection strategy even when a reward is not received in the previous trial. Rudebeck and collaborators (2017), however, showed that macaque monkeys with an excitotoxic lesion of the entire amygdala presented deficits in learning from negative feedback especially during the initial phase of the task. Interestingly, this post-lesion behavioural impairment was also linked to a decrease in the number of OFC neurons encoding the association between one choice option and its associated reward magnitude. Unfortunately, Rudebeck et al. (2017) recorded from neurons in more medial OFC areas but not in IOFC/12o, therefore we do not know whether IOFC/12o also showed a similar response. A hypothetical model summarizing the behavioural results reported both by these authors and here in this thesis in *Chapter 4* is that activity in IOFC/12o may play a primary role in credit assignment by updating the likelihood of an option to lead to a reward after positive feedback is received, whereas amygdala activity may be primarily important in driving a choice strategy switch towards an alternative choice option following contextual changes in value-based contingencies. Throughout this thesis I have shown that this hypothesis may be supported by the strong anatomical (*Chapter 2*) and functional (*Chapter 3*) interactions between these two regions. Moreover, any such potential function for the amygdala in driving value-guided aversive responses could be made possible by the inhibitory microcircuits known to exist between amygdala nuclei and through connections with other brain

stem, striatal, and cingulate areas (Ghashghaei et al., 2007; Janak and Tye, 2015; Namburi et al., 2016, 2015; *Chapter 2*).

The use of fMRI made it possible to look for neural changes not only at the stimulation site but also throughout the whole brain. I was therefore able to look at how the brain responded before and after TUS within the entire amygdala-OFC circuit. In *Chapter 5*, I showed that TUS-induced disruption of activity in IOFC/12o specifically led to a decrease in the activity encoding the expected value of a choice option within IOFC/12o itself, in the anterior amygdala and in more medial OFC and ventromedial PFC areas. However, the way that the brain responded to rewards and non-rewards did not change after IOFC/12o TUS even in interconnected areas such as the amygdala and ventromedial PFC. Overall, the results from *Chapter 5* support the hypothesis that disrupting IOFC/12o activity does affect the way the brain uses positive feedbacks (i.e. reward) to construct a value estimate that is linked to particular choice options, that is impairs credit assignment. Future experiments will need to investigate whether areas within the amygdala-OFC differentially or synergistically respond after either positive or negative feedback. Remarkably, the focality of TUS effects in targeting only specific cognitive computations (credit assignment) but not others (reward or choice identity) show the functional specificity of this approach and its capacity to focally disrupt primarily cognitive mechanisms that are encoded by the targeted structures only. Moreover, the reversible, minimally-invasive and temporary properties of this technique have great potential for bridging the gap between non-human and human primate research. Indeed, TUS may represent a tool that could be used in both species to test for the existence of

homologous cognitive computations within the same brain circuits, regardless of their anatomical depth and distance from the scalp.

Future research should also either aim to integrate within the amygdala-OFC circuit, or contrast, the contribution that other areas, such as cingulate regions on the medial wall of the hemisphere or regions located on the inferior prefrontal convexity, may have in value-guided associative learning mechanisms. Anterior cingulate and anterior PFC regions contribute to counterfactual choice and choice strategy switching. Their role in exploration (Kolling et al., 2016, 2012; Wittmann et al., 2016) cost-benefit trade-off (Klein-Flügge et al., 2016) and counterfactual decision making (Boorman et al., 2011) has already been shown and could support computations linked to or interacting with credit assignment mechanisms. For instance, in a recent paper (Fouragnan et al., 2019) I co-authored, we showed that anterior cingulate cortex ranked the value of the unchosen and unrepresented options during associative learning. Moreover, activity in the anterior cingulate cortex was predictive of a switch in the choice preference in future trials, specifically when a switch towards a better alternative option was considered adaptive, that is the option with the higher value among the unchosen and unrepresented ones. We show that TUS targeted to anterior cingulate cortex particularly disrupted such behavioural switches (Fouragnan et al., 2019). Similar computations are likely to have an impact on credit assignment when contextual changes in contingencies take place.

Lastly, a description of OFC from a more connectivity-based perspective that takes into consideration its synergic interaction with other areas rather than the area in isolation is also likely to lead to a clearer

understanding of its role in orchestrating value-guided behaviour and credit assignment.

Appendix

Global reward state affects learning, the raphe nucleus, and anterior insula in monkeys

Marco K Wittmann, Elsa Fouragnan, Davide Folloni, Miriam C Klein-Flügge, Bolton Chau, Mehdi Khamassi, Matthew F S Rushworth

Abstract

People and other animals learn the values of choices by observing the contingencies between them and their outcomes. However, decisions are not guided by choice-linked reward associations alone; macaques also maintain a memory of the general, average reward rate – the global reward state – in an environment. Remarkably global reward state affects the way that each choice outcome is valued and influences future decisions so that the impact of both choice success and failure is different in rich and poor environments. Successful choices are more likely to be repeated but this is especially the case in rich environments. Unsuccessful choices are more likely to be abandoned but this is especially likely in poor environments. Functional magnetic resonance imaging (fMRI) revealed two distinct patterns of activity, one in anterior insula and one in the dorsal raphe nucleus, that track global reward state as well as specific outcome events.

Introduction

Humans and animals make a multitude of choices every day. Ideally the choice taken should be the one with the highest value – the one most likely to yield positive results. It is widely held that choice value estimates reflect past experience of the consequences of making the same choices; if a choice has led to a positive outcome in the past then it has a higher value than one lacking such a consequence^{1–3}. A choice outcome that is better than previously expected generates a positive prediction error (PE) and, as a consequence, increases the choice's value. Conversely, negative PEs decrease a choice's value estimate. Therefore, implicit in such models is the notion that the experience of *choice-reward conjunctions* determines choice valuations.

While substantial evidence suggests that this is indeed the case and that such a process can be captured in reinforcement learning (RL) models^{2–5} there is evidence that other aspects of choice experience are also important determinants of whether they will be taken again. For example, there is a tendency for both animals and humans simply to repeat previous choices^{6,7}; taking the choice on one occasion makes it more likely that it will be taken again on another. In addition, there is evidence for the existence of non-conjunctive effects – the forging of inappropriate links between the “wrong” choice and reward^{8–14}.

Here we first report analyses of decision-making behavior in four macaques. We show that in addition to conjunctive choice-reward associations, unlinked memories of both choice and reward persist over several trials and influence current and future choices the monkeys take.

Rewards related to the latter effect – an influence of reward experience that is unlinked to any particular choice – we term *global reward state* (GRS). It has a striking impact on behavior:

Animals stayed increasingly with rewarded choices if those were encountered in high GRS, while they abandoned poor choices particularly when encountered in low GRS. Drawing on previous models^{15–18}, we constructed a new RL model that captures both conjunctive choice-reward associations and GRS effects to explain the monkeys' behavior.

Finally, we applied the RL model to neural activity recorded with fMRI. Activity in two areas, the anterior agranular insular cortex (Ia) adjacent to posterior orbitofrontal cortex and the dorsal raphe nucleus (DRN) reliably reflected specific instances of rewards and non-rewards but also the GRS. The temporal pattern of activity in Ia suggested Ia gradually integrates new rewards into a longer-term global estimate of value. By contrast, the DRN pattern suggested a role in regulating the impact that any negative event will have as a function of GRS.

These results shed light on the interactions and influences of the DRN, and therefore possibly serotonin – and frontal cortex on behavioral flexibility^{19,20}. The idea that orbitofrontal cortex is critical for behavioral flexibility has been an influential one but it has recently been challenged by a demonstration that orbitofrontal cortex lesions do not impair classic tests of behavioral inhibition and flexibility^{21,22}. The Ia results suggest a way to reconcile these apparently contradictory findings. In addition, they suggest new ways of thinking about the impact of serotonergic treatments in the clinic. Rather than simply reflecting the likelihood of behavioral inhibition *per*

se, activity in DRN and Ia reflect how an adverse event will be interpreted in the light of the current reward context and it is this combination of factors – the recent outcome and the GRS – that determines whether change in behavior will occur.

Results

We analyzed behavioral data from four macaque monkeys in three experiments using a binary-choice probabilistic bandit task (overall 65 sessions of 200 trials each). In addition, we also analyzed neural data from one of the three experiments (25 sessions)^{23,24}. Choice options were allocated pseudorandomly to the right and left side of the screen and monkeys responded towards congruently located right or left sensors. Choice outcomes were either a drop of juice or nothing (Fig.1a). New stimuli were used for every session. The rewards were delivered probabilistically and the probabilities of two of the options reversed towards the middle of a session and monkeys' choice frequencies followed this reversal (Fig.1b).

Importantly, each day the task contained three new choice stimuli, but only two of them were choosable on each trial (Fig.1c). This was crucial in several respects. First, it made it possible to dissociate choice repetition effects linked to either a choice's location or target stimulus. Second, selective choice presentation and the probabilistic nature of the task ensured that effects of the GRS were dissociable from conjunctive choice-reward effects and that the GRS would fluctuate continuously over the course of a session (Fig.1d). Finally, it ensured that the GRS would be dissociable from

reversals in the relative rates of reward associated with different options (compare Fig.1b,d).

Behavioral effects of linked and unlinked choice and reward memories

To discover potential effects of unlinked memory traces of choice and reward on decision making, we conceptualized monkeys' binary choices as stay/switch decisions to either continue a current course of action or switch to an alternative²⁵. We constructed a detailed logistic general linear model (GLM) for which we concatenated data per monkey per experiment, pooling data over all three experiments (n = 12 monkey data sets; 4 macaques; see Methods). For every trial t, we identified the chosen stimulus C and examined whether the animals would choose C again the next time C was offered (irrespective of whether C was offered on the immediate next or on a subsequent trial; Fig.1e). We tested whether the stay/switch decision was predicted, first, by the conjunctive choice-reward history of C (CxR-history), second, by the (reward-unlinked) choice history of C (C-history) and, third, by the (choice-unlinked) reward history (R-history; Supplementary Fig.1a). With the R-history regressor, we were able to test whether the GRS, irrespective of the pure choice history and the conjunctive choice-reward history, had an influence on stay/switch decisions. Note that this GLM also controls for the history of choice and reward of the upcoming alternative to C (Supplementary Fig.1b).

< insert figure 1 about here >

We found that all three sets of regressors, CxR-history, C-history and R-history, significantly promoted stay behavior (Fig.1f). Firstly, macaques stayed with options if choosing them had been rewarded recently; reward on trial t increased the likelihood of repeating the choice made on that trial (one-sample t -test; CxR-history(t): $t_{11} = 8.883$; $p < 0.001$). Secondly, macaques repeated choices they had made most recently irrespective of reward (C-history($t-1$): $t_{11} = 6.496$; $p < 0.001$). The strength of these effects decreased with time (interaction effect in 5 [recency] x 2 [history type] ANOVA: $F(4,44) = 5.61$; $p < 0.001$). Finally, and most intriguingly, R-history also had a significantly positive effect on choice (R-history: $t_{11} = 4.711$; $p < 0.001$). That means that irrespective of directly reinforced choices and choice repetitions, animals repeated choices that they had previously made at times of a high GRS. In turn, they switch away from a choice - even if that specific choice has been rewarded lately – more often if it is encountered when in a low GRS.

These results imply the possibility that the GRS alters the animals' response to rewards received for a current choice. To investigate this more directly, we regressed all effects of the previous GLM (Fig.1f), except those of CxR-history(t) and R-history, out of the choice data. We then examined the residual choice probabilities, binned, first, by whether animals received reward on trial t and, second, by the binarized (low or high) GRS (median split of R-history; Fig.1g). Animals were more likely to stay after a win and more likely to switch after a loss. However, in addition, staying after a win was more likely when R-history was high compared to when it was low and *vice versa* switching away from a loss was more likely in low reward

environments (2 x 2 ANOVA; interaction: $F_{1,11} = 31.68$; $p < 0.001$). This indicates, remarkably, that the monkeys not only evaluate an option as a function of the outcomes contingent on its choice, but that they do so more strongly when those outcomes are consistent with the overall reward context (i.e. unlinked influence): In a high GRS, negative feedback has a relatively weaker impact on an options' value than it does in a low GRS; switching is less likely to occur after non-reward in a high as opposed to low GRS.

RL model incorporating unlinked choice memory and global reward state

After providing evidence for the unlinked influence of choice memories and GRS on decision making, we went on to formalize a possible computational mechanism for the observed effects in an RL framework. We used an expectation-maximization algorithm for hierarchical fitting over sessions^{26–29} and conducted a nested model comparison over all three data sets collapsed.

In addition to a standard RL architecture, we added unlinked choice and reward memory traces in the model. A choice-location trace (CL-trace) and choice-stimulus traces (CS-traces) captured recency-weighted averages of past choice locations and past choice stimuli, respectively^{4,30} (Fig.2a,b). Thus CL-trace and CS-trace capture two aspects of choices – their locations and their identities. A reward trace (R-trace) captured the recency-weighted GRS. R-trace corresponded to the “disembodied” average of past rewards, irrespective of the choice stimuli or choice locations linked to it (Fig.2c).

< insert figure 2 about here >

In the simplest RL implementation (RLsimple), the decision variable (DV) was calculated as the difference in Q-value between the left and right option, plus an added side bias, filtered through a standard softmax function (Fig.2d). We extended RLsimple by adding the choice traces scaled by separate weights. In the final model, consistent with the behavioral analyses (Fig.1f), the weights associated with both CL and CS traces were positive, indicating that animals directed choices towards previously chosen locations and stimuli irrespective of the reward linked to those choices (Supplementary Fig.3). The choice traces directly changed the model's policy. By contrast, as R-trace is unlinked to specific choices, it must interact with the learning process if it is to affect decision making. For this reason, we added R-trace to the PE calculation scaled by a weight parameter w_R . This meant that the PE was no longer conceived as simply the difference of expectation and outcome. Instead the PEs could be enhanced or diminished by low or high levels of global rewards. The directionality of the R-trace effect depended on w_R which was empirically fitted and allowed to range between -1 and 1. Such use of R-trace is inspired by models of average reward rate learning^{15,31}. When w_R is positive (as was overwhelmingly the case in the current experiment) then the effect of the GRS is facilitative: outcomes are experienced as even more positive in a high GRS compared to low GRS. Negative w_R on the other hand induce contrast effects during learning.

In our model comparison, we considered models with every possible permutation of the components as well as the full model (Fig.2e,f). We

calculated the integrated Bayesian information criterion (BICint)²⁷ pooled over all experiments (Fig.2e). We observed a steady decrease of BICint (indicating better model fit) when including memory traces. The full model (RL+cl+cs+rt) won the model comparison. As an additional validation, we calculated the exceedance probability, i.e., the posterior probability that a model is the most likely model used by the population among a given set of models, and again found that the full model was the best one (exceedance probability = 1.0; Fig.2f)³². Note that both results hold also when testing each of the three experiments separately (Supplementary Fig.2). Moreover, we used such experiment-wise fitted models to specifically compare the full model with the second best one (RL+cl+cs). We computed the session-wise difference in log model evidence (derived via Laplace approximation) (Fig.2g,h,i). This approximates the log Bayes factors, i.e. the ratio of log posterior probability of the two models given the data. Over all data sets, the majority of sessions were in favour of the best over the second-best model (47 out of 65; binomial test against 0.5; $p < 0.001$).

RL learning dynamics induced by a representation of global reward state

Animals learn the value of choices based on their reward consequences. Our behavioral (Fig.1f,g) and modelling results (Fig.2e,f) suggest that in our data, in addition to these choice-reward conjunctions, the GRS affects animals' learning of choice values. In the model, memory traces of unrelated past rewards persist and are incorporated into PEs experienced for current choices. This implies that the value of choices is updated partly as a

consequence of previous rewards that were not caused by these choices. However, the direction of the influence of GRS critically depends on w_R . In principle, the model allows GRS to have a negative, contrasting effect ($w_R < 0$), no effect at all ($w_R = 0$) or a positive effect ($w_R > 0$). We therefore examined the fitted parameter values of w_R . We found that the posterior density of the fitted w_R was overwhelmingly positive for all three experiments (Fig.3a). The integral over the posterior distribution that was above zero was 83% for the fMRI experiment, and 87% and 95% for behavioral experiments 1 and 2 (Fig.3b). This finding is well in line with the behavioral analyses (Fig.1f). It indicates that rather than being contrasted against the GRS, value updates are facilitated by the GRS. The implication is that even when reward is absent on a particular trial, a high GRS means there will be no negative PE. Equally in a very low GRS, reward receipt may not result in a positive PE. An intuitive interpretation of a positive w_R is that the GRS serves as a proxy for how good future states will be. In this view GRS is reminiscent of future reward expectations used in multi-state applications of temporal difference algorithms which are also weighted by a positive discount factor^{4,33}. While such learning mechanisms may be ineffective in laboratory bandit-tasks, they may well be adaptive in the natural environments that animals and humans have evolved to live in, where choices often have serial dependencies^{34–36}.

To further discern the effects of GRS on value updating in our data, we examined the PEs of our full model directly. As R-trace is added to the PE (via the positive w_R), both positive and negative PEs are shifted in a positive direction when R-trace is high (Fig.3c; sessions are fitted to the

mean of the posterior). The implications of this become more transparent when reframing the PEs as absolute PEs – reflecting the degree to which outcomes suggest a need for re-evaluation regardless of direction (sometimes referred to as “surprise”²⁶). Then it becomes clear that a positive w_R leads to positive value updates becoming stronger in a high GRS and negative value updates becoming stronger in a low GRS (median split of R-trace; 2 x 2 ANOVA interaction; $F_{1,64} = 849$; $p < 0.001$). This means that the model effectively implements asymmetric learning rates for positive and negative outcomes that change dynamically as a function of GRS; the impact of newly encountered outcomes is scaled up or down depending on the GRS. This maps directly onto the pattern of stay-switch choices observed in the initial behavioral analysis, showing that a high GRS promotes win-stay and a low GRS promotes lose-switch behavior (Fig.1g).

< insert figure 3 about here >

Our RL model allows rewards not only to influence concurrent choices, but also choices that occur later in time and are logically unrelated (via R-trace). For this reason, we examined whether our model also predicts some previously reported cases of credit misassignment in which rewards are not correctly assigned to the choices that cause them^{8,11,13}. We focused on forward “spread of reward”, where past rewards are mistakenly credited to subsequent choices^{11,13}. We used our model to simulate choices on our task. We varied w_R (for illustration even beyond the bounds used above) and then analyzed the simulated choices with a GLM akin to that used in previous

work to investigate spread of reward^{11,13}. These simulations show a strong effect of choice-reward conjunctions: Reward at time t increases the likelihood of repeating the choice made at time t . However, we find that also rewards received before, at time point $t-1$ systematically impact how likely the agent is to repeat the choice made at time t . Positive w_R leads to a repetition of *subsequent* choices if they occur after reward, because that reward contributes to a more positive GRS increasing subsequent PEs (Fig.3e,f,g,h,i). The inverse was the case for negative w_R . This suggests that spread of reward can be caused by a representation of the GRS.

Rather than biasing decision making exclusively, the GRS may influence reaction times (RTs) as well^{31,38}. We predicted the trialwise negative logarithm of RTs using linear regression (Fig.3j). Q-values of the chosen and unchosen options had no significant influence (both $t_{64} < 1.27$; both $p > 0.2$), and neither did their difference ($t_{64} = 1.57$; $p > 0.13$). Instead, animals responded quicker when stimuli were offered that had been picked more recently: when the chosen CS-trace ($t_{64} = 6.63$; $p < 0.001$) and or the unchosen CS-trace was high ($t_{64} = 3.08$; $p = 0.003$). Finally, the single most significant predictor of RT in the GLM was R-trace ($t_{64} = 8.07$; $p < 0.001$). Therefore, quicker responding was not a function of specific choice-reward associations presented in a trial but of the GRS.

In sum, R-trace in our model encapsulates the persistence of previous unrelated rewards in memory and implements asymmetric and dynamic learning from current outcomes. The dynamics of PEs emerging from this are consistent with the behaviorally observed pattern of stay/switch choices (Fig.1) and create patterns of learning akin to forward “spread of reward”^{11–13}.

As it is unlinked to particular choices, *R-trace* might act as a global signal indicating the current or expected state of the animal. Consistent with this interpretation, the reward trace not only influenced learning, but was also tightly linked to the overall readiness of the animal to engage in the task.

Global reward state representation in Ia and DRN

We started our MRI analysis by asking whether there is a brain region that codes the GRS in a manner consistent with our RL model (RL+cl+cs+rt). In a first GLM (*GLM1*), we found the strongest evidence of encoding of the GRS – *R-trace* encoding – in bilateral anterior insular cortex at the time of choice feedback (cluster-corrected at $Z > |2.6|$, $p = 0.05$ as all whole-brain results reported in this manuscript), just posterior to medial orbitofrontal cortex, which we refer to as agranular insula (Ia; Fig.4a; Supplementary Table 1)³⁹. It is notable that precisely this region of the macaque brain has recently been implicated in discrimination reversal learning (Sallet et al., submitted) and may be important for balancing the influence of reward context versus specific reward outcomes. Consistent with this idea, we not only observed an effect of *R-trace* on bilateral Ia activity but also large signals reflecting specific outcome events (reward versus non-reward) (Fig.4b). The reward signals were located slightly posterior to the *R-trace* signals although they overlapped, particularly in left Ia.

A key feature of our RL model is that learning is driven both by current outcomes and by the GRS. Those two sources are integrated with the same sign in the value update: an option's value increases if it is rewarded but also if GRS is high. Therefore, it is notable that current outcome and *R-trace* both

affect activity in Ia and that they do so with the same sign. Another key feature of the model is that *R-trace* has a positive effect on choice value updates during both rewarded and unrewarded trials. This main effect makes positive PEs more positive and negative PEs less negative as a function of the GRS (Fig.3c,d), which in turn leads animals to stay more with their current choice if the GRS is high (Fig.1f,g). For this reason, we examined, whether, similarly, Ia carries a positive representation of *R-trace* both during rewarded and unrewarded trials. We used a new GLM (*GLM2*) and applied a leave-one-out procedure to the previously identified peak coordinates. We extracted the effect sizes (COPE images) and performed a square-root transformation to de-weight outlying data points (see Methods). We found that, strikingly, Ia represented *R-trace* during both rewarded and non-rewarded trials in the left (rewarded: $t_{24} = 2.07$; $p = 0.049$; unrewarded: $t_{24} = 3.68$; $p = 0.01$) as well as in the right hemisphere (rewarded: $t_{24} = 3.44$; $p = 0.002$; unrewarded: $t_{24} = 3.64$; $p = 0.001$; Fig. 4c,d). There was no significant difference between conditions in left ($t_{24} = 1.23$; $p = 0.23$) or right Ia ($t_{24} = 0.73$; $p = 0.47$).

< insert figure 4 about here >

Ia, particularly left Ia (because of the extensive region in which strong *R-trace* and outcome effects overlapped), seemed to be the most likely brain region to integrate past and current rewards in a manner predicted by the computational model. We therefore examined the BOLD time course in Ia in more detail as a function of both *R-trace* and outcome time-locked to the

feedback onset (see Methods; Fig.5a,b,c,d,e). In order to best understand neural activity in relation to behavioral and RL analyses, we binned activity by *R-trace* (median split; low/high) and outcome (rewarded and unrewarded; *R* and *no R*) (compare with Fig.3c). We found a striking temporal pattern of activity in left Ia. Initially, BOLD activity clustered positively and negatively as a function of *R-trace*. However, 4-7 seconds after feedback onset (and thus at quite a late timepoint given that the monkey hemodynamic response function typically peaks after 3-4 seconds) Ia activity reflect the current outcome (reward or no reward; Fig.5d). This became very clear by entering residual BOLD signals at every time point in a 2 x 2 ANOVA with the factors *R-trace* and outcome and examining the time course as a function of main effects and interaction. Highly consistent with our model predictions, we observed no interaction of *R-trace* and outcome at any point in the time course (Fig.5e). We repeated the analysis for the right Ia and observed the same temporal pattern (Fig.5f). Even though, on the whole brain level, the outcome activity did not overlap directly with right Ia, we found a significant effect of outcome when examining the right Ia time course at the time of the contralateral peak of this effect (time = 6.15 seconds; $F(1,24) = 7.2$; $p = 0.013$).

< insert figure 5 about here >

Complementary to our initial analysis, we repeated our GLM on the whole-brain level so that it was possible to identify areas that only represented *R-trace* when a trial was rewarded or when it was unrewarded.

In the latter case we found a single cluster of activity outside Ia. The cluster was located in the brainstem with a clear peak in a location consistent with the DRN (Fig. 6a). Just as in Ia, overlapping within the same DRN region we also found strong positive coding of current outcomes (ROI analysis: $t_{24} = 3.49$; $p = 0.002$). However, unlike Ia, DRN activity only reflected *R-trace* on non-reward outcome trials, but not on rewarded trials ($t_{24} = 0.51$; $p = 0.62$). There was a significant difference between conditions ($t_{24} = 2.96$; $p = 0.007$; Fig.6b; using a leave-one-out DRN ROI to avoid bias). We found some evidence for the opposite pattern as well as coding of current outcomes in dorsal anterior cingulate cortex (dACC; Supplementary Fig.4). We note that neurophysiological recordings provide evidence that not only the rewards currently being received, but also the rewards previously received are encoded in anterior cingulate cortex^{40–43}.

< insert figure 6 about here >

As expected from the whole brain analyses, the time courses of effects in DRN and dACC looked strikingly different to Ia and there were outcome $\times$ *R-trace* interaction effects at some points (Fig.6c,d; Supplementary Fig.4a-e). Both DRN and dACC showed a significantly earlier outcome encoding compared to Ia or dACC (Supplementary Figure 4f). Therefore, while DRN and dACC represent outcomes at the time of feedback, Ia's neural activity is better described as reflecting ongoing integration of new rewards into a longer-term representation of the choice-unlinked value of each animal's global state.

Coding of value and choice history in prefrontal cortex

In the first part of the fMRI analysis, we reported activity reflecting GRS and specific reward outcomes. In this final section, we again consider neural signals that parallel computations in our RL model. Now, however, we examine activity reflecting the decision variable (DV) that resulted from this learning process and that combines choice values (informed by both the GRS and choice contingent reward learning) and reward-unlinked history of past choices. In contrast to analyses of the GRS, which is – by definition – unrelated to specific choices, we now examined variables in the reference frame of choice. We tested whether any brain region integrated choice evidence in the same comprehensive manner as was apparent in behavior. We regressed the DV (coded as “chosen – unchosen”; DV_{total} of the full model $RL+cl+cs+rt$) at the time of decision against the BOLD signal ($GLM3$). We found a large cluster of activity focused in vmPFC/mOFC but spreading into adjacent parts of prefrontal cortex (Fig.7a). As in previous reports, these effects were negatively signed (i.e. negative relationship between higher relative choice evidence and BOLD signal)^{24,44}.

< insert figure 7 about here >

One possibility is that this activity is solely driven by the relative reward expectation associated with the chosen option because this is part of DV_{total} . Alternatively, prefrontal activity might reflect this but in addition integrate the weight that the previous history of choices has on the subsequent decision. To address this question, we set up a second GLM

(GLM4) that broke DV_{total} apart into its two component parts, DV_{value} and DV_{choice} . DV_{value} is the difference in Q-value between chosen and unchosen options and DV_{choice} is the remaining part of DV_{total} and comprises the weighted choice location and choice stimulus traces (CL-trace and CS-trace, Fig.2d). For DV_{value} , we again found a cluster of activity in vmPFC/mOFC that overlapped with DV_{total} (Fig.7b). However, the effects were considerably smaller than for the full decision variable. To test this difference formally, we performed a neural model comparison in vmPFC ²⁴ using a leave-one-out procedure to avoid bias (see Methods). We regressed two identical GLMs against vmPFC BOLD with the only difference that one of them included DV_{total} and the other one DV_{value} instead. We found that the exceedance probability favored the DV_{total} model by an extensive margin (exceedance probability = 0.999; Fig.7c). This suggests that vmPFC/mOFC integrates different types of evidence such as reward expectation and choice history into a compound decision variable rather than comparing reward associations alone. Note that DV_{choice} is a measure of choice persistence. Higher values indicate an alignment of current choice with recent ones. The sign of these signals in vmPFC and adjacent OFC was similar to that of the value difference signals in vmPFC; BOLD signals were stronger when animals opted for high value choices (which typically mean staying with their previous choices) and when they opted for previously chosen choices regardless of reward (Supplementary Fig.5).

Lastly, we compared the signatures of reward processing across our three ROIs: Ia (bilaterally), DRN and vmPFC. While non-contingent GRS

signals were confined to Ia and DRN, value comparison signals were vice versa specific to vmPFC (Fig.8, Supplementary Fig.6).

< insert figure 8 about here >

Discussion

Animals' choices are driven by past choice-reward conjunctions. Here, we find that they are also heavily influenced by the global reward state (GRS) as well as by the history of past choices *per se*. Our analysis of the macaque BOLD signal suggests critical roles for Ia in integrating new outcomes with the GRS, for the DRN in the relative weighting of outcomes (especially negative ones) in light of the GRS, and for vmPFC in the integration of different sources of information into a multifaceted decision variable.

Our behavioral analyses conceptualize the macaque's choices as a stay/leave-type of decision^{45–47} and show that animals have a strong tendency to repeat rewarded choices, but also, more generally, simply to repeat previous choices regardless of whether they were rewarded. Crucially, in addition, the GRS had a very specific effect on behavior: regardless of a choice's specific history of reward association, animals tended to stay with choices made when GRS was high but switch away from choices when GRS was low (Fig.1).

Our winning RL model conceptualized these effects and thereby supported and extended our behavioral findings (Fig.2). Critically, the RL model offered insight into a possible learning mechanism (Fig.3). The model's estimate of the GRS (i.e. R-trace) summarized the current average

reward rate and biased PEs in a way that facilitated positive value updates in high GRS but relatively depressed them in low GRS. Our model accords well with previous models that have included estimates of the average reward rate^{15,17,18,31}. A feature of the model is that it allows the direction of influence of the GRS on learning to be determined empirically by fitting w_R . That influence can be positive or negative or absent. The positive effects of GRS observed in our study are in line with previous reports of spread of effect^{8,11–13}. They suggest an RL mechanism mediating the “misassignment” of previous rewards to subsequent choices. Negative w_R , on the other hand, implements a temporal contrast effect whereby the value of new outcomes is de-weighted in a high GRS, because they are referenced to the already high value of the environment. Such contrast effects have been reported in other contexts^{17,18,31,48}. Therefore, our model suggests a common algorithmic implementation of a variety of GRS related effects on choice. The model, moreover, captures how GRS effectively modulates the speed of learning. While the GRS influences learning rates in other models³⁷, it affects value updates directly in our model. This means that it can reproduce dynamic changes in value update rates and asymmetries in value updates without explicitly adjusting learning rates. The model might therefore also be related to observations that learning of a specific outcome value can be privileged if its valence accords with prior beliefs about the valence of the general environment in which the option is encountered; for example human subjects who are optimistic about their prospects – people who estimate that the GRS is high – are more influenced by positive outcomes than negative ones^{49–52}.

Our neural results suggest a critical role for Ia³⁹ in the processing of the GRS (Fig.4,5). Reports of neurophysiological recordings in Ia are rare and to our knowledge how single neuron activity is modulated by choice and reward history is unknown. This might be investigated by performing experiments analogous to those performed in areas such as dACC^{41–43,53–55}. However, human fMRI studies provide evidence that rewards are represented with multiple time constants in Ia^{18,56}. Reward related responses have been observed in macaque Ia using fMRI⁵⁷. The neural effect of GRS in Ia was analogous to the constant positive effect of GRS found in the behavioral analysis and present in our RL model (in the model it is also present during both rewarded and unrewarded outcomes). Therefore, activity in Ia closely mimicked the way by which our model combined the influence of reward context and specific reward outcomes to guide value updates. Notably the outcome effect in Ia was conspicuously late in comparison to several other signals in Ia and elsewhere suggesting that Ia might integrate new rewards into the GRS at a slow rate, rather than guiding moment-to-moment adjustments in choice behavior.

Our results suggest a way to begin reconciling claims that orbitofrontal cortex either is or is not critical for flexible behavior in reversal learning tasks. On the one hand excitotoxic lesions of orbitofrontal cortex in rodents and the marmoset – a new world monkey – impair choice-reward discrimination reversal learning^{21,58–60}. By contrast excitotoxic lesions of orbitofrontal cortex in an old world monkey – the macaque – do not²². One possibility is that a critical region for mediating behavioral flexibility is, or is homologous to, macaque Ia. Ia is on the border between anterior insula and posterior

orbitofrontal cortex and is sometimes treated as a very posterior part of OFC. The links between macaque Ia and rodent orbitofrontal cortex have previously been emphasized⁶¹ and there may be links between macaque Ia and parts of marmoset OFC⁶². Two recent observations are consistent with this suggestion (Sallet, et al., submitted). First, macaque Ia, as opposed to more anterior parts of macaque orbitofrontal cortex, is the area exhibiting most extensive grey matter change during choice-reward discrimination reversal learning. Second, aspiration lesions of macaque orbitofrontal cortex, which it is agreed do disrupt choice-reward discrimination reversal learning²², have a profound impact on Ia. A critical requirement of such a task is to learn a particular pattern of counterweighting of the impact that GRS and specific outcome information has on learning; this may explain why Ia is taxed when reversal tasks are learned.

The fact that serotonin depletion in marmoset orbitofrontal cortex^{19,20} also disrupts reversal learning is also consistent with the present finding that both Ia and DRN – an origin of serotonergic innervation of cortex – possess complementary GRS signals. At a behavioral level we found that an adverse outcome – an absence of reward – had different effects depending on the GRS, which is encoded in DRN and Ia. DRN neurons in rodents encode positive and negative PEs⁶³ and manipulation of DRN activity can affect the learning rate by which observed outcomes impact value updates⁶⁴. Such up- and down-regulation of the impact of newly observed outcomes is apparent in the BOLD signal we observe during negative outcomes. The negative response to reward omission appears to be even more negative in states of high GRS. In the light of this pattern of activity, it is possible to

consider that serotonin-based treatments for depression reflect a change in GRS encoding that in turn alters the impact that any specific adverse event will have on a person.

Finally, vmPFC/mOFC integrated stimulus-reward associations as well as choice history into a compound DV. Behaviorally, the animals have a strong tendency to repeat choices regardless of reward delivery. One way to interpret the representation of choice history in vmPFC/mOFC is to relate it to reports of the encoding of attentional variables in prefrontal cortex^{65–67}. Attention is often directed towards the course of action that is currently pursued. It is possible that the animals focus their attention towards a specific choice over several trials. This leads them to preferentially pick options that they have selected before and thus mediates choice repetition. In other words, animals might not only direct their attention towards their current course of action, but also sometimes pursue the course of action that they are currently attending.

Online Methods

Subjects. Four male rhesus monkeys (*Macaca mulatta*) were involved in the experiment. They weighed 10.4–11.9 kg and were 7 years of age. They were group housed and kept on a 12 hr light dark cycle, with access to water 12–16 hr on testing days and with free water access on non-testing days. All procedures were conducted under licenses from the United Kingdom (UK) Home Office in accordance with the UK The Animals (Scientific Procedures) Act 1986 and with the European Union guidelines (EU Directive 2010/63/EU).

Behavioral Training. Prior to the data acquisition, all animals were trained to work in an MRI compatible chair in a sphinx position that was placed inside a custom mock scanner simulating the MRI scanning environment. They were trained to use custom-made infra-red touch sensors to respond to abstract symbols presented on a screen and learned the probabilistic nature of the task until reaching a learning criterion. The animals underwent aseptic surgery to implant an MRI compatible head post (Rogue Research, Mtl, CA). After a recovery period of at least 4 weeks, the animals were trained to perform the task inside the actual MRI scanner under head fixation. The imaging data acquisition started once they performed at more than 70% accuracy (choosing the option with the highest expected value) for at least another three consecutive sessions in the scanner.

Experimental task. Animals had to choose repeatedly between different stimuli that were novel in each testing session (Figure 1). Each session comprised 200 trials. We used a probabilistic reward-based learning task.

The task consisted of a series of choices, on each trial, between two stimuli drawn out of a larger pool of three. The position of the two available options on the left and right side of the screen were pseudorandomized. Animals had to choose any symbols by touching one of two infra-red sensors placed in front of their two hands corresponding to the stimuli on the screen. After making their decision, if the correct option was selected, the unselected option disappeared, and the chosen option remained on the screen and a juice reward was delivered. If an incorrect choice was made, no juice was delivered. The outcome phase lasted 1.5 seconds. Each reward was composed of two 0.6 ml drops of blackcurrant juice delivered by a spout placed near the animal's mouth during scanning. The experiment was controlled by Presentation software (Neurobehavioral Systems Inc., Albany, CA). We used an intertrial-interval of 5-7 seconds for the MRI experiment, but there was no temporal jitter between decision and outcome phase (except reaction time) to keep the animals engaged in the task.

The experimental task was the same in all three experiments as were the four monkeys, but data collection for the experiments was separated by several months and acquired over the course of a two-year period. For the MRI experiment, each animal performed five to seven sessions (25 sessions overall). For the two behavioral experiments, each animal performed 5 sessions (20 sessions per behavioural experiment).

Behavioral data analysis. We analysed behavioral data using Matlab 2018a version 9.4.0 and Jasp version 0.9.0.1.

Basic RL architecture. All reinforcement learning (RL) models shared the same basic architecture. Three value estimates ($Q(A)$, $Q(B)$, $Q(C)$) tracked the rewards received for choosing each of the three stimuli that were presented in each session. Note that each session used new stimuli to avoid carry over effects. All Q-values were initialized at 0.5 at the start of each session and the chosen stimulus was updated based on the discrepancy between outcome (“0” or “1” for reward and no reward, respectively) and Q-value, scaled by a learning rate α . Unchosen Q-value estimates remained unchanged. For example, if option A was chosen on trial t , then its prediction error (PE) and value update would be calculated based on the reward r as follows:

$$1) PE_t(A) = r_t - Q_t(A)$$

$$2) Q_{t+1}(A) = Q_t(A) + \alpha PE_t(A)$$

For all models, the decision variable (DV) reflected the evidence for making a rightward choice. Note that the identity of the left and right choice (whether they were option A , B or C) was pseudorandom. For the simplest model, *RLsimple*, which did not include any choice or reward memory traces, the DV for each trial was simply calculated as the value of the stimulus presented on the right side minus the value of the stimulus presented on the left side plus an additional additive constant reflecting a side bias (SB).

$$3) DV_{RLsimple,t} = Q_t(right) - Q_t(left) - SB$$

Then, for any model m (here $m=RLsimple$), the decision variable was filtered through a standard softmax function to calculate the probability of a rightward choice.

$$4) p_t(right) = \frac{1}{1 + e^{-\beta \times DV_{m,t}}}$$

The probability of the observed choice on each trial was calculated as:

$$5) p_t(choice) = \begin{cases} p_t(right), & \text{if right option is chosen} \\ 1 - p_t(right), & \text{otherwise} \end{cases}$$

The simplest model, *RLsimple*, comprised three free parameters: the learning rate α (bound between 0 and 1), the inverse temperature parameter β (bound between 0 and positive Infinity) and the side bias SB (bound between -1 and 1). Subsequent models shared the basic architecture of model *RLsimple* and comprised the same free parameters with the same bounds plus additional free parameters. Models including choice traces modified the basic DV by adding terms related to the history of past choice stimuli and choice locations. By contrast, models including a reward trace modified the calculation of prediction errors.

Modelling of memory traces. All memory traces were initialized at zero at the start of a session and updated in the following ways. The choice location trace (*CL-trace*) was updated on every trial based on the discrepancy between the actual choice location (L), coded as “-1” or “1” (for right and left side, respectively), and the *CL-trace*, scaled by a learning rate α_{CL} .

$$6) \text{ } CL\text{-}trace_{t+1} = CL\text{-}trace_t + \alpha_{CL}(L_t - CL\text{-}trace_t)$$

The choice stimulus trace (*CS-trace*) decayed exponentially from one trial to the next one with a given rate determined by a free parameter λ_{CS} ^{4,30}. However, the *CS-trace* for the chosen option was set to 1 at the end of the trial. For example, the decay of the *CS-trace* for a stimulus A on trial t was calculated as:

$$7) \text{ } CS\text{-}trace_t(A) = \lambda_{CS}CS\text{-}trace_{t-1}(A)$$

The reward trace (*R-trace*) was updated on every trial based on the discrepancy between *R-trace* and the observed outcome, scaled by a learning rate α_R .

$$8) \text{ } R\text{-}trace_{t+1} = R\text{-}trace_t + \alpha_R(r_t - R\text{-}trace_t)$$

Note that the *R-trace* calculation was independent of the specific choices taken and hence only knowledge about the factual sequence of outcomes was required to calculate *R-trace*. Similarly, *CL-trace* and *CS-trace* required only knowledge of the choice location and choice stimulus, respectively, ignoring the sequence of outcomes experienced over the course of a session. λ_{CS} , α_{CL} and α_R were all bound between 0 and 1.

All three memory traces exerted their influence on learning and choice scaled by weight parameters. In all cases, the weight parameters could be positive, zero, or negative, meaning that the magnitude and direction of influence of the memory traces were determined empirically during model fitting. *CL-trace* and *CS-trace* were added to the *DV*. The *CL-trace* could be added directly, since it was already coded in terms of spatial location similarly to the *DV* itself (although inverted). For the *CS-trace*, the influence on choice was determined by the difference in *CS-trace* between right and left stimulus. Note that the *CS-trace* difference was added to the *DV* after the trial-wise decay but before the update of the chosen stimulus. Below the *DVs* for models that contained only one of the two or both choice traces in the *DV* in addition to the components from *RLsimple*.

$$9) \text{ } CL\text{-}bonus_t = -CL\text{-}trace_t$$

$$10) \text{ } CS\text{-}bonus_t = CS\text{-}trace_t(right) - CS\text{-}trace_t(left)$$

Then for Models *CL*, *CS*, and *CL+CS*, we have respectively:

$$11) DV_{CL,t} = DV_{RLsimple,t} + w_{CL}CL\text{-}bonus_t$$

$$12) DV_{CS,t} = DV_{RLsimple,t} + w_{CS}CS\text{-}bonus_t$$

$$13) DV_{CL+CS,t} = DV_{RLsimple,t} + w_{CL}CL\text{-}bonus_t + w_{CS}CS\text{-}bonus_t$$

Note that positive values of w_{CL} indicate that the DV of the current trial would be biased towards the same location as the direction of previous choices, and positive values of w_{CS} indicate that the DV would be swayed towards the stimulus with the highest *CS-trace*. In other words, positive values of w_{CL} and w_{CS} reflect a tendency to repeat predominant previous choice locations and predominant previous choice stimuli, respectively.

In any model m that included *R-trace*, *R-trace* was inserted directly into the calculation of the prediction error scaled by a weight parameter w_R , which was allowed to range between -1 and 1. For example, if option *A* was chosen on trial t , the corresponding *PE* would be calculated as:

$$14) PE_{m,t}(A) = r_t + w_R R\text{-}trace_t - Q_t(A)$$

Model fitting. All RL modelling was conducted in Matlab (version 2018a). We used an iterative expectation-maximization (EM) algorithm to fit the models^{27,28}. We estimated models both over all experiments and for each experiment separately. This ensured that our results were valid over all data sets but also within each experiment. It also enabled us to use the experiment-appropriate parameter estimates for the analysis of the MRI data.

During the expectation step, we calculated the maximum posterior likelihood (NPL_i) estimate obtained with the parameter vector $\mathbf{h}_i$ of each session i ($i \in \{1..N\}$), given the observed choices and given group-level Gaussian distributions over the parameters with a mean vector μ and standard deviations σ .

$$15) NPL_i = \max_h [\sum_t (\log(p_t(\text{choice}|\mathbf{h}))) + \sum_{\mu, \sigma} \log(\text{normpdf}(\mathbf{h}|\mu, \sigma))]$$

$$16) \mathbf{h}_i = \text{argmax}_h [\sum_t (\log(p_t(\text{choice}|\mathbf{h}))) + \sum_{\mu, \sigma} \log(\text{normpdf}(\mathbf{h}|\mu, \sigma))]$$

The former part expresses the likelihood of the observed choices given a vector of free parameters and the second part captures the likelihood of these parameters given a normal group level distribution. We initialized the group-level Gaussians as uninformative priors with means of 0.1 (plus some added noise) and variance of 100. During the maximization step, we

recomputed ***mu*** and ***sigma*** based on the estimated set of ***h_i*** and their Hessian matrix ***H_i*** (as calculated with Matlab's *fminunc*) over all *N* sessions.

$$17) \mathbf{mu} = \frac{1}{N} \sum_i \mathbf{h}_i$$

$$18) \mathbf{sigma}^2 = \frac{1}{N} \sum_i [\mathbf{h}_i^2 + \text{diag}(\text{pinv}(\mathbf{H}_i))] - \mathbf{mu}^2$$

where the diagonal terms of the inverted Hessian matrix (computed in Matlab with *diag(pinv(H_i))*) give the second moment around ***h_i***, approximating the variance, and thus the inverse of the uncertainty with which the parameter can be estimated ²⁸.

We repeated expectation and maximization steps iteratively until convergence of the posterior likelihood *NPL_i* summed over the group or a maximum of 800 steps. Convergence was defined as a change in *NPL_i* < 0.001 from one iteration to the next. Note that bounded free parameters (for example the learning rates) were transformed from the Gaussian space into the native model space via appropriate link functions (e.g. a sigmoid function in the case of the learning rates) to ensure accurate parameter estimation near the bounds.

Model comparison. We compared fitted models by calculating their integrated BIC (*BIC_{int}*) (Huys et al., 2012). For this, we drew *k=2000* samples

of parameter vector $\mathbf{h}_i$ per session i from the Gaussian population distributions using the final estimates of μ and σ , and computed the negative log likelihood ($NLL_{i,k}$) of each sample and session using the equation (corresponding to the first part in equation 15)

$$19) NLL_{i,k} = -\sum_t \log(p_t(choice|\mathbf{h}_{i,k}))$$

Next, we integrated the $NLL_{i,k}$ over samples k and sessions i and calculated BIC_{int} based on the integrated log-likelihood ($iLog$) in the following way:

$$20) iLog = \sum_i \log(\sum_{k=1}^{2000} e^{-NLL_{i,k}} / 2000)$$

$$21) BIC_{int} = -2 \times iLog + Np \times \log(\sum_i Nt_i)$$

Np refers to the number of free parameters per model and Nt_i refers to the number of trials per session i .

As a second index of model fit, we used the Laplace approximation to calculate the log model evidence (LME) per session i based on NPL_i (see equation 15):

$$22) LME_i = -NPL_i - \frac{1}{2} \log(\det(\mathbf{H}_i)) + \frac{Np}{2} \log(2\pi)$$

We submitted the *LME* scores to `spm_BMS`³² to compute the ‘exceedance probability’, the posterior probability that one model is the most likely model used by the population among a given set of models. In addition, we computed the session-wise difference in *LME* between two candidate models to approximate log Bayes factors, i.e. the ratio of posterior probability of the models given the data^{68,69}.

Model simulations. We ran model simulations to illustrate how an RL model including reward traces can effectively lead to the assignment of rewards to choices that occur after (rather than before) that reward was obtained. We simulated sequences of choices for every session of our three experiments using the true underlying reward schedule. Then, we concatenated sessions from the same monkey within each experiment (just as had been done in the behavioral analysis in Figure 1) and applied a behavioral GLM to the simulated choices. The resulting 12 (3 experiments with 4 monkeys) beta weights per regressor were averaged and are shown in Fig.3e,f,g,h,i.

To isolate the effects of R-trace on decision making in our simulations, we removed effects related to choice repetition and side bias by setting λ_{CS} , w_{CS} , α_{CL} , w_{CL} and SB to zero. Moreover, we set the remaining model parameters to standard values ($\alpha = 0.35$, $\beta = 1$, $\alpha_R = 0.5$) and systematically

varied the effect of R-trace in the prediction error calculation (see equation 14). For this, in separate iterations of the simulation, we set w_R to -1, 0, 1 and 2. This determines whether R-trace has a negative effect, no effect, or a positive effect in the prediction error calculation. We derived choice probabilities from the simulated Q-values by using the softmax equation above (equation 4) and used these choice probabilities to generate simulated choices.

The logistic GLM we applied to the simulated choices was similar to previous studies ^{11,13}. The GLM considered the last 5 choices and the last 5 outcomes. For each of the three choice stimuli, we applied a separate GLM and combined their beta weights to a covariance-weighted mean. We used 25 regressors in the GLM; one for each choice-outcome conjunction. The conjunctions were coded as 1, -1 or 0 depending on whether the choice of interest was made and rewarded, another choice was made and rewarded or no reward occurred, respectively. As in previous reports, regressors relating to the last time points of choices or rewards (e.g. Choice(t-5) and reward(t-1)) were confound regressors and are therefore not shown in figure 3.

Imaging Data Acquisition. Awake-animals were head-fixed in a sphinx position in an MRI-compatible chair. We collected fMRI using a 3T MRI scanner and a four-channel phased array receive coil in conjunction with a radial transmission coil (Windmiller Kolster Scientific Fresno, CA). FMRI data were acquired using a gradient-echo T2* echo planar imaging (EPI) sequence with 1.5 x 1.5 x 1.5 mm³ resolution, repetition time (TR) = 2.28 s,

Echo Time (TE) = 30 ms, flip angle = 90, and reference images for artifact corrections were also collected. Proton-density-weighted images using a gradient-refocused echo (GRE) sequence (TR = 10 ms, TE = 2.52 ms, flip angle = 25) were acquired as reference for body motion artifact correction. T1-weighted MP-RAGE images (0.5 x 0.5 x 0.5 mm³ resolution, TR = 2,5 ms, TE = 4.01 ms) were acquired in separate anesthetized scanning sessions.

fMRI data preprocessing. FMRI data were corrected for body motion artefacts by an offline-SENSE reconstruction method ⁷⁰ (Offline_SENSE GUI, Windmiller Kolster Scientific, Fresno, CA). The images were aligned to an EPI reference image slice-by-slice to account for body motion and then aligned to each animal's structural volume to account for static field distortion ⁷¹ (Align_EPI GUI and Align_Anatomy GUI, Windmiller Kolster Scientific, Fresno, CA). The aligned data were processed with high-pass temporal filtering (3-dB cutoff of 100s) and Gaussian spatial smoothing (full-width half maximum of 3mm). The data that were already registered to each subject's structural space were then registered to the CARET macaque F99 template ⁷² using affine transformation.

fMRI whole brain analysis. We employed a univariate approach within the general linear model (GLM) framework to perform whole-brain statistical analyses of functional data as implemented in the FMRIB Software Library version 5.0.11 ⁷³ where each of the psychological regressors was convolved with a hemodynamic response function (HRF) specific for monkey brains

^{74,75}. Using this framework we initially performed a first-level fixed effects analysis to process each individual experimental run which were then combined in a second-level mixed-effects analysis (FLAME 1) treating session as a random effect (we also had a similar number of sessions across subjects) and using family-wise error cluster correction ($z > 2.6$ and $p = 0.05$).

Our first level GLMs had several features in common. When using regressors derived from the modelling, they were always taken from the full model $RL+cl+cs+rt$ and were estimated session-wise based on the mean of the Gaussian parameter distributions. All GLMs employed the same basic set of regressors to capture variance in the data. A decision constant (*DEC*) was time-locked to the onset of the decision screen and the duration was set to the reaction time for that trial. We used two constant regressors time-locked to the time of feedback, one for rewarded outcomes (*FBrew*) and one for unrewarded outcomes (*FBnrew*). Both had a default duration of 0.1 seconds. All subsequent regressors in all GLMs were set to this default duration unless otherwise noted. In addition, to account for movement related scanning artefacts, we used unconvolved regressors for leftwards and rightwards responses (*Cleft_unc* and *Cright_unc*), time-locked to the respective responses, with the duration set to the TR (2.28 seconds). Except for *Cleft_unc* and *Cright_unc*, all regressors in all GLMs were convolved with the monkey-specific HRF (mean lag of 3s, standard deviation of 1.5s). Finally, we used a binary regressor to indicate choice location (1 for right, 2 for left), time-locked to choice (*Cloc*). As all other subsequent binary or parametric regressors, *Cloc* was z-scored before entering it in the GLM.

Our first GLM focused on the memory traces specific to our RL model (*GLM1*). Time-locked to *DEC*, it contained the Q-value difference between the chosen and the unchosen option. This quantity reflects the value driven part of the decision variable and we refer to it as $DV_{value}(chosen-unchosen)$.

$$23) DV_{value}(chosen - unchosen) = Q(chosen) - Q(unchosen)$$

In addition, the GLM contained the Q-value of the unpresented option (*Q-unp*), the CL-trace (*CL-trace*), the CS-trace of the unpresented option (*CS-unp*), as well as the CS-trace comparison between the chosen and the unchosen option (*CS(chosen-unchosen)*). Moreover, we included the R-trace in the GLM (*R-trace*), which was time-locked to the onset of the feedback screen, as in the model, the reward trace affects the prediction error at the time of learning. These additional regressors were z-scored and had a default duration of 0.1 seconds. Moreover, we calculated the contrast between rewarded and unrewarded outcomes (*FBrew-FBnrew*) and refer to this contrast as the *outcome* effect. For the second level analysis of the *outcome* contrast specifically, we employed a restrictive whole-brain pre-threshold mask to exclude movement artefacts related to the consumption of the juice rewards. However, the reported effects remain significant even without this mask.

In our second GLM (*GLM2*), we used the same set of regressors as in *GLM1* with one difference. We split up the *R-trace* regressor in rewarded and unrewarded trials. This means we time-locked this effect to only rewarded and only unrewarded trials (*R-trace (R)* and *R-trace (no R)*), respectively,

and the parametric variation of *R-trace* was z-scored within the respective trial set.

Our third GLM (*GLM3*) investigated decision related activity by regressing the decision variable of the RL model against the BOLD signal in addition to the set of basic regressors described above. No further regressors were included in the GLM. We coded it in terms of chosen minus unchosen, i.e. as relative evidence for the chosen option ($DV(chosen-unchosen)$). It was time-locked to *DEC* onset. For brevity, $DV(chosen-unchosen)$ is just referred to as DV_{total} in the main text.

The final whole-brain GLM (*GLM4*) broke apart the decision variable in one element related to value and one related to the previous choice history, $DV_{value}(chosen-unchosen)$ (same as in *GLM1* and *GLM2*) and $DV_{choice}(chosen-unchosen)$ (see Fig. 2d). These two regressors were the only ones added to the basic set of regressors and they were coded in terms of chosen minus unchosen, similar to the decision variable in *GLM 1*. Specifically, the two quantities were defined as:

$$24) DV_{choice}(chosen - unchosen) = DV(chosen - unchosen) - DV_{value}(chosen - unchosen)$$

$DV_{value}(chosen-unchosen)$ and $DV_{choice}(chosen-unchosen)$ were both time-locked to *DEC*. For brevity, $DV_{value}(chosen-unchosen)$ and $DV_{choice}(chosen-unchosen)$ are just referred to as DV_{value} and DV_{choice} in the main text.

Region of interest (ROI) analyses. We analyzed ROI data using Matlab 2018a version 9.4.0, Jasp version 0.9.0.1, and SPM version 12. For all ROI analyses, we used spherical ROIs with 2mm radius centered on peak coordinates reported in Supplementary Table 1. First, to examine *R-trace* related effects during trials in which either reward was or was not delivered, we read out FSL's *contrast of parameter estimate* (COPE) maps in independently selected ROIs by warping our spherical mask in session specific space and averaging the parameter estimates over the mask. To account for outlying data points, we square-root transformed the data. In case of negative COPE values, we used the absolute value and added a minus after the transformation. For ROI analyses of COPE effects where the ROI focused on a similar effect to the defining whole-brain contrast, ROIs were selected via a leave-one-out method to avoid selection bias. For this, we calculated the whole-brain contrast of interest over all sessions but kept one session out at a time. ROI peak coordinates for the contrast were then identified based on the (incomplete) group average (that did not include the left-out session) and applied to the left-out session. We repeated this procedure for all sessions and applied statistical tests to the resulting independent ROIs.

We used the ROI timecourse analyses to examine the temporal evolution of neural signals in our ROIs. To this end, we extracted the preprocessed BOLD time series and up-sampled it by a factor of 10 (using spline intrapolation) and aligned the time series trial-wise. On every time point, we applied a set of z-scored regressors to the data. The time course GLM (*tGLM1*) performed the following calculations to each time point of the

upsampled data in a session by session manner. First, we used a regressor set comprising $DV_{value}(chosen-unchosen)$, $DV_{choice}(chosen-unchosen)$, $Q-unp$, $CS-unp$, $CL-trace$ and $Cloc$. After applying this regressor set to the data time-locked to the feedback phase onset with a linear GLM, we calculated the residuals of the BOLD timecourse. Then, we binned the BOLD residuals by *outcome* (rewarded or unrewarded) and by *R-trace* (high or low; median split). On the group level we performed two statistical procedures with the binned BOLD residuals. First, we applied a 2x2 analysis of variance (ANOVA) to it on every time point and calculated main effects of *outcome* and *R-trace* as well as the interaction. Second, for each session, we calculated the peak times of relevant ANOVA effects to be able to compare when effects were strongest both within and across brain regions. As dACC represented *R-trace* preferentially during rewarded trials and DRN preferentially represented *R-trace* during no reward trials, we calculated separate measures of the *R-trace* effect for rewarded and unrewarded trials for dACC, DRN, and Ia. For every time step and session, we subtracted the residual BOLD time course of *R-trace* low from *R-trace* high separately for rewarded and unrewarded trials, resulting in the effects *R-trace (R)* and *R-trace (no R)*, respectively. Moreover, we calculated the overall difference of rewarded minus unrewarded bins to obtain the effect of *outcome*. Within a plausible time window of two to six seconds (given the relatively fast macaque hemodynamic response function ²³) from feedback onset, we determined the session-specific time of the peak of the relevant effect. As we are not testing the size of the peak effect but only the time for each session at which it occurs, restricting the effects to a positive or negative direction

does not bias the results. For each effect, we constrained the analysis to identify the timing of effect peaks in the direction concordant with the group effect. This means for dACC, we looked for the time of the most negative *R-trace* effect during rewarded trials as well as the most negative *outcome* effect. For Ia bilaterally, we looked for the time of the most positive peak for each effect as both effects were positive. For the DRN, we took the time of the most positive *outcome* effect and the time of the most negative *R-trace* effect during unrewarded trials.

Neural model comparison. To assess whether vmPFC signalled a combined decision variable rather than value difference alone, we performed a neural model comparison^{17,24}. Based on the DV_{total} signal from GLM3, we used a leave-one-out procedure to select session-specific vmPFC ROIs and avoid biased selection. Then, we first regressed GLM3 against the BOLD signal in vmPFC. Second we did the same for GLM4, but crucially without the inclusion of DV_{choice} . This means that both GLMs contained identical regressors, with the exception that the first one contained DV_{total} and the second one DV_{value} instead. The loglikelihood estimates from the two GLMs were then fed into a Bayesian model selection random-effects analysis (using the `spm_BMS` routine from SPM12³²), which computed the exceedance probability of each GLM. This analysis indicated which GLM best explained the neural data.

Figures




Figure 1. Task design and behavior. (a) Monkeys performed a probabilistic binary choice task for juice reward. (b) Smoothed choice frequency (moving window of 20 trials) followed the reversal in reward probabilities of two of the

three options (blue and yellow). (c) One session comprised three options indicated by different stimuli, only two of which were choosable each trial (experimenter determined, pseudo-random). (d) The smoothed overall average reward rate (moving window of 20 trials) fluctuated continuously over the course of the session. (e) Illustration of GLM analysis. For every trial we predicted whether the chosen option (C in this example) would be repeated on the next occasion it was presented using three sets of regressors. The first set reflected the reward probability of C if it was chosen (CxR-history), the second set reflected the tendency to choose C in the past if it was one of the offered options (C-history), and the third reflected the global reward state unrelated to a specific choice (R-history). (f) Significantly positive effects of CxR-history, C-history, and R-history on the decision to stay with an option. (g) Residual probability of making a stay or switch choice. Stay/switch behavior was consistent with the outcome (win/lose) on trial t , but the effect was facilitated if the outcome was in accordance with global reward state: Switching away from an option was more likely if R-history was low, while staying with a rewarded option was more likely in high reward contexts. (in panels b,d, data related to the MRI experiment is shown and error bars are SEM across sessions; $n = 25$; panels f,g pool over all experiments and error bars are SEM across monkey data sets; $n = 12$; *, $p < 0.05$).




Figure 2. RL modelling of choice and reward traces. Illustration of choice location traces (CL-trace, **a**), choice stimulus traces (CS-trace, different colors indicate different choice stimuli, **b**) and reward traces (R-trace, **c**). (**d**) Model architecture. Black colors indicate components used by all models. The other colors refer to choice and reward traces as illustrated in panels **a**, **b**, **c**. Note that CL-trace and CS-trace effects are added directly to the DV of the model, whereas R-trace affects the prediction error calculation. (**e**) Nested model comparison using BICint (lower values indicate better fit). The full model including R-trace explains the data best. x-labels are the same for panels **e** and **f**. (**f**) Exceedance probability also favors the full model among our set of candidate models. (**g**, **h**, **i**) Experiment specific comparisons of Bayes factors between full model and second-best model. Bars are individual sessions and direction of the bars (left vs right) indicate better model fit in favor of the full (going to the right) compared to the second-best model (going to the left). ($n = 65$ sessions).




Figure 3. R-trace related learning dynamics. (a) Posterior density over w_R for the FMRI experiment as well as the two behavioral experiments. (b) Percentage of posterior distribution shown in panel a that is bigger than zero. (c) Average PEs binned by current outcome (Reward/No reward) and R-trace (low/high; median split) show the offset of PE coding generated by positive w_R . (d) Absolute PEs plotted as in panel c illustrate stronger positive value updates when R-trace is high and stronger negative value updates when R-trace is low. (e,f,g,h) Choice simulations using different levels of w_R (-1,0,1,2, respectively) are analyzed with a GLM that estimates the effect of each possible choice-reward-pair going 4 trials back on subsequent choices. Lighter colors indicate more positive beta weights, i.e. subsequent choice probability is increased by reward delivered for the respective conjunction. For example, when w_R is set to zero in panel f, the bright white diagonal indicates that the next choice the agent takes is influenced by the conjunction of choice and reward on the previous trial (t-1), the trial before that (t-2) and so on. However, as w_R increases (g, h), the precise history of choice-reward conjunction has less effect and “illusory” conjunctions between choices made on one trial and reward occurrences on other trials become more likely to influence the next choice that is taken (lighter colors at lower left in g and h). (i) Effects of rewards delivered at various time points relative to time point 0 on the probability of repeating the choice made at time point zero. Effects are derived from averaging beta weights shown in panels e-h along the diagonal with different offsets. The area shaded in red corresponds to the mean effects along the red squares in panels e-h and indicates that rewards have contrasting ($w_R < 0$) or facilitating effects on repeating subsequent choices ($w_R > 0$) depending on w_R . (j) Beta weights of linear regression on the negative log of reaction times suggest that animals respond quicker when the choices that are offered on the current trial were recently pursued and also when R-trace is high. (in all panels except i, error bars are SEM across sessions; $n=65$; for panel i, error bars are SEM across monkey data sets, $n=12$; panels a-d and j show data from the three experiments, whereas panels e-i show simulated data; green star symbols in panel j indicate outlier beyond the y-limits of the plot (3 outlying data points); *, $p < 0.05$).




Figure 4. Bilateral Ia represents global reward state and current reward. (a) Positive effects of R-trace bilaterally in Ia. (b) Both regions also showed concurrent overlapping or closely adjacent positive activation related to the current outcome (reward – no reward). (c,d) In accordance with our RL model, in both regions, R-trace exerted a positive effect during rewarded as well as unrewarded trials. ROIs were selected using a leave-one-out procedure. (Crosshairs highlight peak coordinates for subsequent ROI analysis; Errorbars are SEM across sessions; $n=25$; dot color in c,d indicate monkey identity; all MRI results cluster-corrected at $Z>2.6$, $p=0.05$; *, $p < 0.05$).




Figure 5. Timecourse of bilateral la activity (a,b,c) Residual BOLD effects in left la binned by R-trace (low/high; median split) and outcome (reward/no reward: R and no R) at three time points after feedback delivery (indicated as grey bars in panel d; feedback occurs at time 0sec and the approximate peak of the macaque hemodynamic response function is 3-4 seconds later). (d) Time course of the four effects shown in panels a,b,c. Note that the BOLD signal clusters initially as a function of R-trace (low/high) and then regroups as a function of outcome (no reward/reward). Zero is time of feedback onset. (e) F-statistic for 2x2 (R-trace x outcome) ANOVA applied at each time step to the data shown in panel d. Rows indicate main effects and interaction effect. Circles are scaled F-values. Larger circles indicate higher F-values. Yellow indicate p-values of effects are bigger than 0.05. Green circles indicate effects for a given row and time point where $p < 0.05$. (f) Same analysis as in panel e applied to the right la replicate pattern of effects found in left la. (Errorbars are SEM across sessions; $n=25$; *, $p < 0.05$).




Figure 6. R-trace and outcome signals. (a) During unrewarded trials, DRN encoded R-trace negatively. This effect is shown at $z > 3.3$ for spatial specificity. Reward feedback occurs at time 0sec and the approximate peak of the macaque hemodynamic response function is 3-4 seconds later (b) This was not the case during rewarded trials (reward/no reward: R and no R). ROIs were selected using a leave-one-out procedure. (c,d) Time course analyses of DRN activity related to outcome and R-trace coding. Same conventions as in Fig.5d,e. (Crosshairs highlight peak coordinates for subsequent ROI analysis; Errorbars are SEM across sessions; $n=25$; dot color in b indicates monkey identity; all MRI results cluster-corrected at $Z > 2.6$, $p=0.05$; *, $p < 0.05$).




Figure 7. Decomposition of vmPFC and DRN activity in value and choice. (a) Effects of DV_{total} in vmPFC/mOFC. (b) DV_{value} , the part of DV_{total} reflecting the value difference between chosen and unchosen options, correlated with a similar vmPFC area, but signal sizes were weaker. (c) Bayesian model comparison suggests vmPFC activity is better described by DV_{total} than by DV_{value} . ROIs were determined using a leave-one-out procedure. The dashed red line indicates an exceedance probability of 0.95 ($n=25$ sessions; all MRI results cluster-corrected at $Z>2.6$, $p=0.05$; *, $p < 0.05$).




Figure 8. Computations of learning and choice in Ia, DRN and vmPFC.

Check marks indicate observed effects with references to the respective figure panel. “x” indicates absent effects as detailed in Supplementary Figure 6. **(a)** Bilaterally in Ia as well as in DRN, we found a representation of the GRS, i.e. a choice-unlinked general representation of recent reward. We did not observe such a non-contingent general reward signal in vmPFC. **(b)** In vmPFC, however, we found a value difference signal capturing the difference in contingent reward expectation of the chosen compared to the unchosen option (DV_{value}). These signals were clearly absent in Ia and DRN. The activity patterns in Ia and DRN suggest they are more related to learning than to decision making.

References Wittman et al. Appendix


1. Rescorla, R. A. & Wagner, A. R. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Class. Cond. II Curr. Res. Theory* 2, 64–99 (1972).
2. Schultz, W. Behavioral theories and the neurophysiology of reward. *Annu. Rev. Psychol.* 57, 87–115 (2006).
3. Thorndike, E. L. *Animal Intelligence; Experimental Studies*. (New York: Macmillan, 1911).
4. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction*. 28, (MIT Press, 1998).
5. Averbeck, B. B. & Costa, V. D. Motivational neural circuits underlying reinforcement learning. *Nat. Neurosci.* 20, 505–512 (2017).
6. Akaishi, R., Umeda, K., Nagase, A. & Sakai, K. Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron* 81, 195–206 (2014).
7. Lau, B. & Glimcher, P. W. Dynamic response-by-response models of matching behavior in rhesus monkeys. *J Exp Anal Behav* 84, 555–79 (2005).
8. Jocham, G. et al. Reward-Guided Learning with and without Causal Attribution. *Neuron* 90, 177–90 (2016).
9. Noonan, M. P. et al. Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc Natl Acad Sci U A* 107, 20547–52 (2010).
10. Noonan, M. P., Chau, B., Rushworth, M. F. & Fellows, L. K. Contrasting effects of medial and lateral orbitofrontal cortex lesions on credit assignment and decision making in humans. *J Neurosci* (2017). doi:10.1523/JNEUROSCI.0692-17.2017
11. Rudebeck, P. H., Saunders, R. C., Lundgren, D. A. & Murray, E. A. Specialized Representations of Value in the Orbital and Ventrolateral Prefrontal Cortex: Desirability versus Availability of Outcomes. *Neuron* 95, 1208–1220.e5 (2017).
12. Thorndike, E. L. A Proof of the Law of Effect. *Science* 77, 173–5 (1933).
13. Walton, M. E. M. et al. Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* 65, 927–39 (2010).
14. Lesaint, F., Sigaud, O. & Khamassi, M. Accounting for Negative Automaintenance in Pigeons: A Dual Learning Systems Approach and Factored Representations. *PLoS ONE* 9, e111050 (2014).
15. Daw, N. D. & Touretzky, D. S. Long-term reward prediction in TD models of the dopamine system. *Neural Comput* 14, 2567–83 (2002).
16. Dayan, P. Prospective and retrospective temporal difference learning. *Netw. Bristol Engl.* 20, 32–46 (2009).
17. Palminteri, S., Khamassi, M., Joffily, M. & Coricelli, G. Contextual modulation of value signals in reward and punishment learning. *Nat Commun* 6, 8096 (2015).
18. Wittmann, M. K. et al. Predictive decision making driven by multiple time-linked reward representations in the anterior cingulate cortex. *Nat Commun* 7, 12327 (2016).
19. Clarke, H. F., Dalley, J. W., Crofts, H. S., Robbins, T. W. & Roberts, A. C. Cognitive inflexibility after prefrontal serotonin depletion. *Science* 304, 878–880 (2004).
20. Clarke, H. F., Walker, S. C., Dalley, J. W., Robbins, T. W. & Roberts, A. C. Cognitive inflexibility after prefrontal serotonin depletion is behaviorally and neurochemically specific. *Cereb. Cortex N. Y. N* 1991 17, 18–27 (2007).

21. Roberts, A. C. Primate orbitofrontal cortex and adaptive behaviour. *Trends Cogn. Sci.* 10, 83–90 (2006).
22. Rudebeck, P. H., Saunders, R. C., Prescott, A. T., Chau, L. S. & Murray, E. A. Prefrontal mechanisms of behavioral flexibility, emotion regulation and value updating. *Nat Neurosci* 16, 1140–5 (2013).
23. Chau, B. K. et al. Contrasting Roles for Orbitofrontal Cortex and Amygdala in Credit Assignment and Learning in Macaques. *Neuron* 87, 1106–18 (2015).
24. Fouragnan, E. F. et al. The macaque anterior cingulate cortex translates counterfactual choice value into actual behavioral change. *Nat. Neurosci.* 22, 797–808 (2019).
25. Rushworth, M. F., Kolling, N., Sallet, J. & Mars, R. B. Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr Opin Neurobiol* 22, 946–55 (2012).
26. Fouragnan, E., Queirazza, F., Retzler, C., Mullinger, K. J. & Philiastides, M. G. Spatiotemporal neural characterization of prediction error valence and surprise during reward learning in humans. *Sci. Rep.* 7, 4762 (2017).
27. Huys, Q. J. et al. Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding. *PLoS Comput Biol* 7, e1002028 (2011).
28. Huys, Q. J. et al. Bonsai trees in your head: how the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Comput Biol* 8, e1002410 (2012).
29. Queirazza, F., Fouragnan, E., Steele, J. D., Cavanagh, J. & Philiastides, M. G. Neural correlates of weighted reward prediction error during reinforcement learning classify response to cognitive behavioral therapy in depression. *Sci. Adv.* 5, eaav4962 (2019).
30. Bogacz, R., McClure, S. M., Li, J., Cohen, J. D. & Montague, P. R. Short-term memory traces for action bias in human reinforcement learning. *Brain Res.* 1153, 111–21 (2007).
31. Niv, Y., Daw, N. D., Joel, D. & Dayan, P. Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacol. Berl* 191, 507–20 (2007).
32. Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J. & Friston, K. J. Bayesian model selection for group studies. *Neuroimage* 46, 1004–17 (2009).
33. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* 275, 1593–9 (1997).
34. Enomoto, K. et al. Dopamine neurons learn to encode the long-term value of multiple future rewards. *Proc Natl Acad Sci U A* 108, 15462–7 (2011).
35. Kolling, N., Wittmann, M. & Rushworth, M. F. S. Multiple neural mechanisms of decision making and their competition under changing risk pressure. *Neuron* 81, 1190–1202 (2014).
36. Shidara, M. & Richmond, B. J. Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* 296, 1709–11 (2002).
37. Farashahi, S. et al. Metaplasticity as a Neural Substrate for Adaptive Learning and Choice under Uncertainty. *Neuron* 94, 401–414.e6 (2017).
38. Wang, A. Y., Miura, K. & Uchida, N. The dorsomedial striatum encodes net expected return, critical for energizing performance vigor. *Nat. Neurosci.* 1–12 (2013). doi:10.1038/nn.3377
39. Carmichael, S. T. & Price, J. L. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J Comp Neurol* 363, 615–641 (1995).

40. Bernacchia, A., Seo, H., Lee, D. & Wang, X. J. A reservoir of time constants for memory traces in cortical neurons. *Nat Neurosci* 14, 366–72 (2011).
41. Donahue, C. H., Seo, H. & Lee, D. Cortical signals for rewarded actions and strategic exploration. *Neuron* 80, 223–34 (2013).
42. Kennerley, S. W., Behrens, T. E. & Wallis, J. D. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci* 14, 1581–9 (2011).
43. Seo, H., Barraclough, D. J. & Lee, D. Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cereb Cortex* 17 Suppl 1, i110–7 (2007).
44. Papageorgiou, G. K. et al. Inverted activity patterns in ventromedial prefrontal cortex during value-guided decision-making in a less-is-more task. *Nat Commun* 8, 1886 (2017).
45. Hayden, B. Y., Pearson, J. M. & Platt, M. L. Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14, 933–9 (2011).
46. Kolling, N. et al. Value, search, persistence and model updating in anterior cingulate cortex. *Nat Neurosci* 19, 1280–5 (2016).
47. Stephens, D. W. & Krebs, J. R. *Foraging theory*. (Princeton University Press, 1986).
48. Crespi, L. P. Quantitative Variation of Incentive and Performance in the White Rat. *Am. J. Psychol.* 55, 467–517 (1942).
49. Garrett, N. et al. Losing the rose tinted glasses: neural substrates of unbiased belief updating in depression. *Front Hum Neurosci* 8, 639 (2014).
50. Korn, C. W., Prehn, K., Park, S. Q., Walter, H. & Heekeren, H. R. Positively biased processing of self-relevant social feedback. *J. Neurosci. Off. J. Soc. Neurosci.* 32, 16832–16844 (2012).
51. Nicolle, A., Symmonds, M. & Dolan, R. J. Optimistic biases in observational learning of value. *Cognition* 119, 394–402 (2011).
52. Sharot, T., Korn, C. W. & Dolan, R. J. How unrealistic optimism is maintained in the face of reality. *Nat Neurosci* 14, 1475–9 (2011).
53. Donahue, C. H. & Lee, D. Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. *Nat Neurosci* 18, 295–301 (2015).
54. Seo, H. & Lee, D. Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci* 27, 8366–77 (2007).
55. Seo, H., Cai, X., Donahue, C. H. & Lee, D. Neural correlates of strategic reasoning during competitive games. *Science* 346, 340–3 (2014).
56. Meder, D. et al. Simultaneous representation of a spectrum of dynamically changing value estimates during decision making. *Nat Commun* 8, 1942 (2017).
57. Kaskan, P. M., Dean, A. M., Nicholas, M. A., Mitz, A. R. & Murray, E. A. Gustatory responses in macaque monkeys revealed with fMRI: Comments on taste, taste preference, and internal state. *NeuroImage* 184, 932–942 (2019).
58. Schoenbaum, G., Setlow, B., Nugent, S. L., Saddoris, M. P. & Gallagher, M. Lesions of orbitofrontal cortex and basolateral amygdala complex disrupt acquisition of odor-guided discriminations and reversals. *Learn. Mem. Cold Spring Harb. N* 10, 129–140 (2003).
59. Schoenbaum, G., Saddoris, M. P. & Stalnaker, T. A. Reconciling the roles of orbitofrontal cortex in reversal learning and the encoding of outcome expectancies. *Ann. N. Y. Acad. Sci.* 1121, 320–335 (2007).

60. Stalnaker, T. A., Cooch, N. K. & Schoenbaum, G. What the orbitofrontal cortex does not do. *Nat. Neurosci.* 18, 620–627 (2015).
61. Wise, S. P. Forward frontal fields: phylogeny and fundamental function. *Trends Neurosci* 31, 599–608 (2008).
62. Passingham, R. E. & Wise, S. P. *The Neurobiology of the Prefrontal Cortex: Anatomy, Evolution, and the Origin of Insight.* (OUP Oxford, 2012).
63. Matias, S., Lottem, E., Dugué, G. P. & Mainen, Z. F. Activity patterns of serotonin neurons underlying cognitive flexibility. *eLife* 6, e20552 (2017).
64. Iigaya, K., Fonseca, M. S., Murakami, M., Mainen, Z. F. & Dayan, P. An effect of serotonergic stimulation on learning rates for rewards apparent after long intertrial intervals. *Nat. Commun.* 9, 2477 (2018).
65. Hunt, L. T. et al. Triple dissociation of attention and decision computations across prefrontal cortex. *Nat. Neurosci.* 21, 1471–1481 (2018).
66. Lim, S. L., O'Doherty, J. P. & Rangel, A. The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. *J Neurosci* 31, 13214–23 (2011).
67. Lopez-Persem, A., Domenech, P. & Pessiglione, M. How prior preferences determine decision-making frames and biases in the human brain. *Elife* 5, (2016).
68. Bornstein, A. M., Khaw, M. W., Shohamy, D. & Daw, N. D. Reminders of past choices bias decisions for reward in humans. *Nat Commun* 8, 15958 (2017).
69. Daw, N. Trial-by-trial data analysis using computational models. *Decis. Mak. Affect Learn. Atten.* ... 1–26 (2009).
70. Kolster, H. et al. Visual field map clusters in macaque extrastriate visual cortex. *J Neurosci* 29, 7031–9 (2009).
71. Kolster, H., Janssens, T., Orban, G. A. & Vanduffel, W. The retinotopic organization of macaque occipitotemporal cortex anterior to V4 and caudoventral to the middle temporal (MT) cluster. *J Neurosci* 34, 10168–91 (2014).
72. Van Essen, D. C. et al. Mapping visual cortex in monkeys and humans using surface-based atlases. *Vis. Res* 41, (2001).
73. Smith, S. M. et al. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23 Suppl 1, S208-19 (2004).
74. Kagan, I., Iyer, A., Lindner, A. & Andersen, R. A. Space representation for eye movements is more contralateral in monkeys than in humans. *Proc Natl Acad Sci U S A* 107, 7933–8 (2010).
75. Nakahara, K., Hayashi, T., Konishi, S. & Miyashita, Y. Functional MRI of macaque monkeys performing a cognitive set-shifting task. *Science* 295, 1532–6 (2002).


Supplementary Figures Wittman et al., Appendix




Supplementary Figure 1. Behavioral choice GLM. (a) Details of regressor construction relating to Figure 1e. Left: for each trial, two out of three options were offered (A, B, and C). Black circles indicate choices and (crossed through) drops on the left indicate (no) rewards. From this information, three sets of regressors were constructed. We used them to analyze for each trial t whether the animal would subsequently stay with the same choice made in this trial or switch. The three sets of regressors are illustrated on the right side: Firstly, to capture conjunctive reward effects (green, CxR-history), we entered regressors that indicated whether choices of C on trial t and also on the previous trials were rewarded or not (CxR(t), CxR($t-1$), ...). Note that “ t ” here refers to actual occurrences of choices of C and not necessarily to consecutive trials. We only selected trials on which C was actually chosen for

this regressor as only those are informative about the conjunctive choice-reward history of C. In other words, CxR-history regressors were coded for past trials in which C was chosen and set to 1/0 for rewarded/unrewarded outcomes. Reward obtained in conjunction with C should increase the probability of a stay choice. However, the stay/switch choice should also depend, inversely, on the conjunctive reward history of the alternative option that was offered as an alternative to C. To account for these effects, we created analogous regressors for the alternative option (see green weights on the right of panel b). Then, for simplicity, we aggregated the beta weights from both sets of regressors – those relating to C and those relating to the alternative option with which it was subsequently paired (see panel b for details on the aggregation). Secondly, we built regressors to account for the recent choice history of C ($C(t-1)$, etc.). For this, we considered whether in past trials in which C was among the offered options, C was chosen or not, irrespective of any reward receipt. Again, “t” refers only to a subset of trials (the ones in which C was among the offered options), but a different subset compared to the CxR regressors. As the reward history for C was captured by the CxR regressors, C-history reflects the strength of the choice trace for C irrespective of whether C was recently rewarded or not. C-history regressors were coded for past trials in which C was offered and set to 1/0 for choosing/not choosing C on that trial. Again, we constructed analogous regressors for the alternative option that would be presented with C on the next trial and aggregated the results. Finally, we took the simple average reward on the six trials before t as an index of the overall current levels of reward (R-history). This time, “t” refers to the actual six most recent trials as these are most relevant for estimating reward history regardless of whether C is chosen or rewarded. However, we note that the results do not change with changes in history length. **(b)** Complete set of beta weights for the GLM analysis in Fig.1f. As highlighted in panel a and in the main text, we also accounted for reward related and choice repetition related effects of the upcoming alternative option (effects on the right-hand side). For Fig.1f we sign-reversed the beta weights relating to the upcoming alternative option and averaged them with CxR-history and C-history of the current choice. We averaged the beta weights of CxR-history(t-1) of the current choice (second


bar in the panel) with the beta weights from CxR-history($t-1$) of the upcoming alternative (13th bar in the panel), the same for $t-2$ and so on. Following the same approach, we averaged the beta weights of C-history($t-1$) of the current choice and the (again sign reversed) C-history($t-1$) of the upcoming alternative choice, and so on for $t-2$, $t-3$, $t-4$ and $t-5$.



Supplementary Figure 2. Supplementary model comparison for individually fitted data sets. BICint and exceedance probability for the fMRI data set (**a,d**) and the two behavioral data sets (**b,e** and **c,f**, respectively).




Supplementary Figure 3. Parameter distributions for all free parameters of the full model $RL+cl+cs+rt$. The model was fitted individually for each experimental data set. See main text for details on panel h showing the weight of R-trace in the prediction error calculation.




Supplementary Figure 4. (a) R-trace, calculated over all trials (GLM1), identified a significant negative cluster of activation in dorsal anterior cingulate cortex (dACC). (b) Such processing of recent rewards was complemented with (negative) signals relating to current outcome that were correlated with dACC activity in a more anterior portion of the cingulate sulcus. (c) ROI analyses in dACC (see crosshair in panel a) revealed that R-trace signals in dACC were significant only during rewarded ($t_{24} = 3.32$; $p = 0.003$), but not during unrewarded trials ($t_{24} = 1.00$; $p = 0.33$), with a significant difference between them ($t_{24} = 2.96$; $p = 0.007$). Note that we used a leave-one-out procedure to identify session-specific ROIs. This pattern of effects is opposite to the pattern observed in DRN. (d,e) Time course analyses of dACC activity. Same conventions as in figures 5 and 6. (f) Time to peak for R-trace and outcome ANOVA effects for bilateral Ia, DRN and dACC. One intriguing difference in coding schemes between, on the one hand, Ia and, on the other hand, DRN and dACC was the late timing of outcome encoding in Ia. Rather than exerting an effect after 3-4 seconds, i.e. the approximate peak of the macaque hemodynamic response function (HRF), outcome signals emerged after 5 seconds in both left and right Ia. On average, Ia outcome signals emerged significantly later than the average R-

trace peak in Ia ($t_{24} = 4.26$; $p < 0.001$), or the R-trace signals in DRN (in unrewarded trials: $t_{24} = 4.82$; $p < 0.001$) and dACC (in reward trials $t_{24} = 3.07$; $p = 0.005$), all of which emerged close in time to the expected HRF peak. Strikingly, such late processing of reward outcomes in Ia also differed from much earlier outcome signals in DRN ($t_{24} = 7.25$; $p < 0.001$) and dACC ($t_{24} = 5.89$; $p < 0.001$). This might suggest that outcomes in Ia are represented at a time at which they no longer define the current state of the animal but have themselves become part of the GRS. Note that dACC and DRN R-trace effects are only shown for rewarded and unrewarded trials, respectively, because R-trace effects were only observed under these conditions in these ROIs. (Crosshairs highlight peak coordinates for subsequent ROI analysis; Errorbars are SEM across sessions; $n=25$; dot color in c indicate monkey identity; all MRI results cluster-corrected at $Z > 2.6$, $p=0.05$; *, $p < 0.05$).



Supplementary Figure 5. We attempted to specify the type of choice history representation present in prefrontal cortex more precisely. As shown in the behavioral analyses and modelling, the monkeys in our experiments had a strong tendency, irrespective of reward, to stay with previously picked stimuli. The experimental design allows dissociation of this tendency, indexed by the CS-trace variable from the tendency to just repeat choice location. We regressed the relative strength of the chosen choice stimulus trace (CS-trace(chosen-unchosen)) against BOLD activity (GLM1). CS-trace(chosen-unchosen) captures the degree to which the animals stay with their previous choice stimulus or switch away to a stimulus they have not picked for a while. We found a cluster of activity overlapping but slightly adjacent to previous vmPFC/mOFC activity in a posterior part of medial orbitofrontal cortex (mOFC). Again, chosen and unchosen option effects manifested with the same negative and positive signs as observed for chosen and unchosen value effects in vmPFC/mOFC. (Errorbars are SEM across sessions; $n=25$; MRI results cluster-corrected at $Z>2.6$, $p=0.05$).



Supplementary Figure 6. Summary analysis of sqrt-COPE images in previously established ROIs. (a) Examining R-trace, we have found evidence for a global reward state representation in la bilaterally and in DRN (see references to main figures below the check marks). In vmPFC, however, we did not find a R-trace signal, neither during rewarded ($t_{24} = -1.56$; $p = 0.132$) nor unrewarded trials ($t_{24} = -1.56$; $p = 0.132$). We used the vmPFC ROI identified by DV_{total} . We did not use a leave-one-out-procedure to identify the ROI since the tested contrast was independent of the ROI defining contrast. (b) In contrast, representations of the value comparison, the chosen minus unchosen value or DV_{value} , were very clear in vmPFC (see reference to main text figure). However, we did not find such signals in the left la ($t_{24} = -1.34$; $p = 0.193$), right la ($t_{24} = 1.50$; $p = 0.148$), or DRN ($t_{24} = 0.99$; $p = 0.331$). la and DRN ROIs were identified by the main effect of R-trace for la and by the effect of R-trace during unrewarded trials in DRN; again, not using a leave-one-out procedure since we were examining effects that were independent of ROI selection. (Errorbars are SEM across sessions; $n=25$; dot color in c indicate monkey identity).

Contrast	Region	Peak Coordinates x/y/z (in mm)	Z-Value
R-trace (positive contrast; GLM1)	Right agranular insula	19.5/4.5/-5	4.58
	Left agranular insula	-18/3.5/-1	4.43
	Cerebellum	14/-35/-20.5	4.14
R-trace (negative contrast; GLM1)	Dorsal anterior cingulate cortex (dACC)	-1/14.5/13.5	3.94
	Right dorsolateral prefrontal cortex	14/6.5/11.5	3.85
	Right posterior superior temporal sulcus / posterior parietal cortex	18.5/-33.5/13.5	3.85
	Left posterior superior temporal sulcus/ posterior parietal cortex	-24/-25/2	4.04
Outcome (positive contrast; Reward < no Reward; GLM1)	Dorsal anterior cingulate cortex (dACC)	0/19/11	4.31
Outcome (positive contrast; Reward > no Reward; GLM1)	Left ventral premotor cortex	-22/-6/8	6.48
	Right ventral premotor cortex	23/-5.5/8.5	5.65
	Right Amygdala	11.5/-5/-11	6.1
	Left Amygdala	-11.5/-5.5/-6.5	5.37
	Left Putamen	-13.5/0/-4	5.3
	Right Putamen	14/-2/-2.5	4.71
	Left insula	-22/0.5/2	4.94
	Right insula	18.5/1.5/-1	3.95
	Medial dorsal Thalamus	2/-9/2	4.61
	Ventral tegmental area	1/-8.5/-8.5	4.57
	Pons	1.5/-14/-14.5	4.65

Supplementary Table 1. Peak coordinates of activation clusters. Coordinates are given according to the CARET macaque F99. Regions of interest analyses were performed in regions highlighted in bold.

Bibliography

- Aggleton, J. (Ed.). (2000). *The Amygdala: A Functional Analysis*. Oxford, New York: Oxford University Press.
- Aggleton, J. P. (1993). The contribution of the amygdala to normal and abnormal emotional states. *Trends in Neurosciences*, 16(8), 328–333.
- Aggleton, J. P., Burton, M. J., & Passingham, R. E. (1980). Cortical and subcortical afferents to the amygdala of the rhesus monkey (*Macaca mulatta*). *Brain Research*, 190(2), 347–368.
- Aggleton, John P. (2008). EPS Mid-Career Award 2006. Understanding anterograde amnesia: Disconnections and hidden lesions. *Quarterly Journal of Experimental Psychology* (2006), 61(10), 1441–1471.
- Aggleton, John P., & Passingham, R. E. (1981). Syndrome produced by lesions of the amygdala in monkeys (*Macaca mulatta*). *Journal of Comparative and Physiological Psychology*, 95(6), 961–977.
- Aggleton, J. P., Wright, N. F., Rosene, D. L., & Saunders, R. C. (2015). Complementary Patterns of Direct Amygdala and Hippocampal Projections to the Macaque Prefrontal Cortex. *Cerebral Cortex* (New York, N.Y.: 1991), 25(11), 4351–4373.
- Airan, R. D., & Pauly, K. B. (2018). Hearing out Ultrasound Neuromodulation. *Neuron*, 98(5), 875–877.
- Akaishi, R., Kolling, N., Brown, J. W., & Rushworth, M. (2016). Neural Mechanisms of Credit Assignment in a Multicue Environment. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 36(4), 1096–1112.
- Aldroubi, A., & Basser, P. (1999). Reconstruction of vector and tensor fields from sampled discrete data. *Contemporary Mathematics*.
- Alm, K. H., Rolheiser, T., Mohamed, F. B., & Olson, I. R. (2015). Fronto-temporal white matter connectivity predicts reversal learning errors. *Frontiers in Human Neuroscience*, 9, 343.
- Amaral, D. G., & Price, J. L. (1984). Amygdalo-cortical projections in the monkey (*Macaca fascicularis*). *The Journal of Comparative Neurology*, 230(4), 465–496.

- Ameis, S. H., & Catani, M. (2015). Altered white matter connectivity as a neural substrate for social impairment in Autism Spectrum Disorder. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 62, 158–181.
- Andersson, J. L. R., Skare, S., & Ashburner, J. (2003). How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. *NeuroImage*, 20(2), 870–888.
- Avants, B. B., Tustison, N. J., Song, G., Cook, P. A., Klein, A., & Gee, J. C. (2011). A Reproducible Evaluation of ANTs Similarity Metric Performance in Brain Image Registration. *NeuroImage*, 54(3), 2033–2044.
- Bachevalier, J., Parkinson, J. K., & Mishkin, M. (1985). Visual recognition in monkeys: Effects of separate vs. combined transection of fornix and amygdalofugal pathways. *Experimental Brain Research*, 57(3), 554–561.
- Baldermann, J. C., Melzer, C., Zapf, A., Kohl, S., Timmermann, L., Tittgemeyer, M., Huys, D., Visser-Vandewalle, V., Kühn, A.A., Horn, A., Kuhn, J. (2019). Connectivity Profile Predictive of Effective Deep Brain Stimulation in Obsessive-Compulsive Disorder. *Biological Psychiatry*.
- Bandettini, P. A. (2012). Twenty years of functional MRI: The science and the stories. *NeuroImage*, 62(2), 575–588.
- Barbas, H. (1995). Anatomic basis of cognitive-emotional interactions in the primate prefrontal cortex. *Neuroscience and Biobehavioral Reviews*, 19(3), 499–510.
- Barbas, H., & De Olmos, J. (1990). Projections from the amygdala to basoventral and mediodorsal prefrontal regions in the rhesus monkey. *The Journal of Comparative Neurology*, 300(4), 549–571.
- Barbas, H. (2007). Flow of information for emotions through temporal and orbitofrontal pathways. *Journal of Anatomy*, 211(2), 237–249.
- Barbas, H. (2015). General cortical and special prefrontal connections: Principles from structure to function. *Annual Review of Neuroscience*, 38, 269–289.
- Barbas, H., Zikopoulos, B., & Timbie, C. (2011). Sensory pathways and emotional context for action in primate prefrontal cortex. *Biological Psychiatry*, 69(12), 1133–1139.

- Barrett, T. W. (1969). Studies of the function of the amygdaloid complex in *Macaca mulatta*. *Neuropsychologia*, 7(1), 1–12.
- Barton, R. A., & Harvey, P. H. (2000). Mosaic evolution of brain structure in mammals. *Nature*, 405(6790), 1055–1058.
- Barton, Robert A., & Venditti, C. (2013). Human frontal lobes are not relatively large. *Proceedings of the National Academy of Sciences of the United States of America*, 110(22), 9001–9006.
- Basser, P. J., Mattiello, J., & LeBihan, D. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1), 259-267.
- Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., & Aldroubi, A. (2000). In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4), 625–632.
- Basser, P. J., Jones, & D. K. (2002). Diffusion-tensor MRI: theory, experimental design and data analysis - a technical review. *NMR in Biomedicine*, 15(7-8), 456-67.
- Baxter, M. G., Parker, A., Lindner, C. C., Izquierdo, A. D., & Murray, E. A. (2000). Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 20(11), 4311–4319.
- Baxter, Mark G., Gaffan, D., Kyriazis, D. A., & Mitchell, A. S. (2008). Dorsolateral prefrontal lesions do not impair tests of scene learning and decision-making that require frontal-temporal interaction. *The European Journal of Neuroscience*, 28(3), 491–499.
- Baxter, Mark G., Gaffan, D., Kyriazis, D. A., & Mitchell, A. S. (2009). Ventrolateral prefrontal cortex is required for performance of a strategy implementation task but not reinforcer devaluation effects in rhesus monkeys. *The European Journal of Neuroscience*, 29(10), 2049–2059.
- Baxter, Mark G., & Murray, E. A. (2002). The amygdala and reward. *Nature Reviews. Neuroscience*, 3(7), 563–573.
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1–3), 7–15.

- Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 19(13), 5473–5481.
- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science (New York, N.Y.)*, 275(5304), 1293–1295.
- Bechara, A., Tranel, D., Damasio, H., & Damasio, A. R. (1996). Failure to respond autonomically to anticipated future outcomes following damage to prefrontal cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, 6(2), 215–225.
- Behrens, T. E. J., Berg, H. J., Jbabdi, S., Rushworth, M. F. S., & Woolrich, M. W. (2007). Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34(1), 144–155.
- Behrens, T. E. J., Woolrich, M. W., Jenkinson, M., Johansen-Berg, H., Nunes, R. G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S. M. (2003). Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50(5), 1077–1088.
- Behrens, T. E. J., Hunt, L. T., & Rushworth, M. F. S. (2009). The computation of social behavior. *Science (New York, N.Y.)*, 324(5931), 1160–1164.
- Behrens, T. E. J., & Jbabdi, S. (2009). Chapter 15—MR Diffusion Tractography. In Heidi Johansen-Berg & T. E. J. Behrens (Eds.), *Diffusion MRI* (pp. 333–351).
- Belliveau, J. W., Kennedy, D. N., McKinstry, R. C., Buchbinder, B. R., Weisskoff, R. M., Cohen, M. S., Vevea, J.M., Brady, T.J., Rosen, B. R. (1991). Functional mapping of the human visual cortex by magnetic resonance imaging. *Science (New York, N.Y.)*, 254(5032), 716–719.
- Bickart, K. C., Hollenbeck, M. C., Barrett, L. F., & Dickerson, B. C. (2012). Intrinsic amygdala-cortical functional connectivity predicts social network size in humans. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 32(42), 14729–14741.
- Bijsterbosch, J., Smith, S. M., & Beckmann, C. F. (2017). *Introduction to Resting State fMRI Functional Connectivity*. Oxford, New York: Oxford University Press.

- Biswal, B. B. (2012). Resting state fMRI: A personal history. *NeuroImage*, 62(2), 938–944.
- Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34(4), 537–541.
- Blackmore, J., Shrivastava, S., Sallet, J., Butler, C. R., & Cleveland, R. O. (2019). Ultrasound Neuromodulation: A Review of Results, Mechanisms and Safety. *Ultrasound in Medicine & Biology*, 45(7), 1509–1536.
- Bogacz, R., McClure, S. M., Li, J., Cohen, J. D., & Montague, P. R. (2007). Short-term memory traces for action bias in human reinforcement learning. *Brain Research*, 1153, 111–121.
- Boorman, E. D., Behrens, T. E. J., Woolrich, M. W., & Rushworth, M. F. S. (2009). How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron*, 62(5), 733–743.
- Boorman, E. D., Behrens, T. E., & Rushworth, M. F. (2011). Counterfactual choice and learning in a neural network centered on human lateral frontopolar cortex. *PLoS Biology*, 9(6), e1001093.
- Bramson, B., Folloni, D., Verhagen, L., Hartogsveld, B., Mars, R. B., Toni, I., & Roelofs, K. (2019). Human lateral Frontal Pole contributes to control over social-emotional behaviour. *BioRxiv*, 584896.
- Brindle, K. M., Brown, F. F., Campbell, I. D., Grathwohl, C., & Kuchel, P. W. (1979). Application of spin-echo nuclear magnetic resonance to whole-cell systems. Membrane transport. *Biochemical Journal*, 180(1), 37–44.
- Broca, P. (1861a). Perte de la parole: Ramollissement chronique et destruction partielle du lobe antérieur gauche du cerveau. *Bulletins de La Société d'anthropologie*, 2, 235–238.
- Broca, P. (1861b). Remarques sur le siège de la faculté du langage articulé, suivies d'une observation d'aphémie (perte de la parole). *Bulletin de La Société Anatomique*, 6, 330–357.
- Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*.

- Brown, S., & Sharpey-Schafer, E. A. (1888). XI. An investigation into the functions of the occipital and temporal lobes of the monkey's brain. *Philosophical Transactions of the Royal Society of London. (B.)*, 179, 303–327.
- Browning, P. G. F., & Gaffan, D. (2008). Impairment in object-in-place scene learning after uncinate fascicle section in macaque monkeys. *Behavioral Neuroscience*, 122(2), 477–482.
- Brunel, N., & Wang, X. J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, 11(1), 63–85.
- Buckner, R. L., Bandettini, P. A., O'Craven, K. M., Savoy, R. L., Petersen, S. E., Raichle, M. E., & Rosen, B. R. (1996). Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 93(25), 14878–14883.
- Bunge, S. A., & Wallis, J. D. (2008). *Neuroscience of Rule-Guided Behavior*. Oxford University Press.
- Burgess, M. T., Apostolakis, I., & Konofagou, E. E. (2018). Power cavitation-guided blood-brain barrier opening with focused ultrasound and microbubbles. *Physics in Medicine and Biology*, 63(6), 065009.
- Bussey, T. J., Wise, S. P., & Murray, E. A. (2001). The role of ventral and orbital prefrontal cortex in conditional visuomotor learning and strategy use in rhesus monkeys (*Macaca mulatta*). *Behavioral Neuroscience*, 115(5), 971–982.
- Bussey, T. J., Wise, S. P., & Murray, E. A. (2002). Interaction of ventral and orbital prefrontal cortex with inferotemporal cortex in conditional visuomotor learning. *Behavioral Neuroscience*, 116(4), 703–715.
- Butter, C. M., Snyder, D. R., & McDonald, J. A. (1970). Effects of orbital frontal lesions on aversive and aggressive behaviors in rhesus monkeys. *Journal of Comparative and Physiological Psychology*, 72(1), 132–144.
- Camille, N., Tsuchida, A., & Fellows, L. K. (2011). Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 31(42), 15048–15052.

- Carlén, M. (2017). What constitutes the prefrontal cortex? *Science (New York, N.Y.)*, 358(6362), 478–482.
- Carmichael, S. T., & Price, J. L. (1994). Architectonic subdivision of the orbital and medial prefrontal cortex in the macaque monkey. *The Journal of Comparative Neurology*, 346(3), 366–402.
- Carmichael, S. T., & Price, J. L. (1995). Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *The Journal of Comparative Neurology*, 363(4), 615–641.
- Cartwright, J. H. E., González, D. L., & Piro, O. (2001). Pitch perception: A dynamical-systems perspective. *Proceedings of the National Academy of Sciences of the United States of America*, 98(9), 4855–4859.
- Catani, M., Dell'acqua, F., & Thiebaut de Schotten, M. (2013). A revised limbic system model for memory, emotion and behaviour. *Neuroscience and Biobehavioral Reviews*, 37(8), 1724–1737.
- Catani, M., Howard, R. J., Pajevic, S., & Jones, D. K. (2002). Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *NeuroImage*, 17(1), 77–94.
- Catani, M., & Thiebaut de Schotten, M. (2008). A diffusion tensor imaging tractography atlas for virtual in vivo dissections. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 44(8), 1105–1132.
- Cavada, C., Compañy, T., Tejedor, J., Cruz-Rizzolo, R. J., & Reinoso-Suárez, F. (2000). The anatomical connections of the macaque monkey orbitofrontal cortex. A review. *Cerebral Cortex (New York, N.Y.: 1991)*, 10(3), 220–242.
- Chau, B. K. H., Kolling, N., Hunt, L. T., Walton, M. E., & Rushworth, M. F. S. (2014). A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nature Neuroscience*, 17(3), 463–470.
- Chau, B. K. H., Sallet, J., Papageorgiou, G. K., Noonan, M. P., Bell, A. H., Walton, M. E., & Rushworth, M. F. S. (2015). Contrasting Roles for Orbitofrontal Cortex and Amygdala in Credit Assignment and Learning in Macaques. *Neuron*, 87(5), 1106–1118.
- Chen, H., Garcia-Gonzalez, D., & Jérusalem, A. (2019). Computational model of the mechano-electrophysiological coupling in axons with application to neuromodulation. *Physical Review E*, 99(3), 032406.

- Cho, Y. T., Ernst, M., & Fudge, J. L. (2013). Cortico-amygdala-striatal circuits are organized as hierarchical subsystems through the primate amygdala. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(35), 14017–14030.
- Clark, A. B. (1985). Sociality in a nocturnal “solitary” prosimian: Galago crassicaudatus. *International Journal of Primatology*, 6(6), 581–600.
- Clark, A. M., Bouret, S., Young, A. M., & Richmond, B. J. (2012). Intersection of Reward and Memory in Monkey Rhinal Cortex. *Journal of Neuroscience*, 32(20), 6869–6877.
- Clarke, E., & O'Malley, C. D. (1968). *The Human Brain and Spinal Cord: A Historical Study Illustrated by Writings from Antiquity to the Twentieth Century*. Los Angeles: Berkeley.
- Cohen, J. (1991). AIDS vaccine meeting: International trials soon. *Science (New York, N.Y.)*, 254(5032), 647.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X. J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex (New York, N.Y.: 1991)*, 10(9), 910–923.
- Constans, C., Deffieux, T., Pouget, P., Tanter, M., & Aubry, J.-F. (2017). A 200–1380-kHz Quadrifrequency Focused Ultrasound Transducer for Neurostimulation in Rodents and Primates: Transcranial In Vitro Calibration and Numerical Study of the Influence of Skull Cavity. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 64(4), 717–724.
- Costa, V. D., Dal Monte, O., Lucas, D. R., Murray, E. A., & Averbeck, B. B. (2016). Amygdala and Ventral Striatum Make Distinct Contributions to Reinforcement Learning. *Neuron*, 92(2), 505–517.
- Cox, J., & Witten, I. B. (2019). Striatal circuits for reward learning and decision-making. *Nature Reviews. Neuroscience*, 20(8), 482–494.
- Croxson, P. L., Johansen-Berg, H., Behrens, T. E. J., Robson, M. D., Pinski, M. A., Gross, C. G., Richter, W., Richter, M. C., Kastner, S., Rushworth, M. F. S. (2005). Quantitative investigation of connections of the prefrontal cortex in the human and macaque using probabilistic diffusion tractography. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 25(39), 8854–8866.

- Curie, J., & Curie, P. (1880). Développement par compression de l'électricité polaire dans les cristaux hémiedres à faces inclinées. *Bulletin de Minéralogie*, 3(4), 90–93.
- Dallapiazza, R. F., Timbie, K. F., Holmberg, S., Gatesman, J., Lopes, M. B., Price, R. J., Miller, G.W., Elias, W. J. (2018). Noninvasive neuromodulation and thalamic mapping with low-intensity focused ultrasound. *Journal of Neurosurgery*, 128(3), 875–884.
- Damasio, H., Grabowski, T., Frank, R., Galaburda, A. M., & Damasio, A. R. (1994). The return of Phineas Gage: Clues about the brain from the skull of a famous patient. *Science (New York, N.Y.)*, 264(5162), 1102–1105.
- D'Arceuil, H., & de Crespigny, A. (2007). The effects of brain tissue decomposition on diffusion tensor imaging and tractography. *NeuroImage*, 36(1), 64–68.
- Darwin, C. (1859). *On the Origin of Species by Means of Natural Selection, Or, The Preservation of Favoured Races in the Struggle for Life*.
- de Guglielmo, G., Kallupi, M., Pomrenze, M. B., Crawford, E., Simpson, S., Schweitzer, P., Koob, G.F., Messing, R.O., George, O. (2019). Inactivation of a CRF-dependent amygdalofugal pathway reverses addiction-like behaviors in alcohol-dependent rats. *Nature Communications*, 10(1), 1238.
- De Martino, B., Kumaran, D., Seymour, B., & Dolan, R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science (New York, N.Y.)*, 313(5787), 684–687.
- deCampo, D. M., & Fudge, J. L. (2013). Amygdala projections to the lateral bed nucleus of the stria terminalis in the macaque: Comparison with ventral striatal afferents. *The Journal of Comparative Neurology*, 521(14), 3191–3216.
- Deffieux, T., Younan, Y., Wattiez, N., Tanter, M., Pouget, P., & Aubry, J.-F. (2013). Low-Intensity Focused Ultrasound Modulates Monkey Visuomotor Behavior. *Current Biology*, 23(23), 2430–2433.
- Déjerine, J. J., & Dejerine-Klumpke, A. (1895). *Anatomie des centres nerveux*.
- Dias, R., Robbins, T.W., & Roberts, A.C. (1996). Dissociation in prefrontal cortex of affective and attentional shifts. *Nature*, 380(6569), 69–72.

- Diogo, R., Muchlinski, M. N., & Hartstone-Rose, A. (2015). Chapter 4—Comparative Anatomy of Primates. In M. P. Muehlenbein (Ed.), *Basics in Human Evolution* (pp. 43–55).
- Divac, I., Björklund, A., Lindvall, O., & Passingham, R. E. (1978). Converging projections from the mediodorsal thalamic nucleus and mesencephalic dopaminergic neurons to the neocortex in three species. *The Journal of Comparative Neurology*, 180(1), 59–71.
- Donahue, C. J., Glasser, M. F., Preuss, T. M., Rilling, J. K., & Van Essen, D. C. (2018). Quantitative assessment of prefrontal cortex in humans relative to nonhuman primates. *Proceedings of the National Academy of Sciences of the United States of America*, 115(22), E5183–E5192.
- Drevets, W. C., Price, J. L., Simpson, J. R., Todd, R. D., Reich, T., Vannier, M., & Raichle, M. E. (1997). Subgenual prefrontal cortex abnormalities in mood disorders. *Nature*, 386(6627), 824–827.
- Dunbar, R. I. M. (2009). The social brain hypothesis and its implications for social evolution. *Annals of Human Biology*, 36(5), 562–572.
- Eacott, M. J., & Gaffan, D. (1992). Inferotemporal-frontal Disconnection: The Uncinate Fascicle and Visual Associative Learning in Monkeys. *The European Journal of Neuroscience*, 4(12), 1320–1332.
- Eldridge, M. A. G., Lerchner, W., Saunders, R. C., Kaneko, H., Krausz, K. W., Gonzalez, F. J., Ji, B., Higuchi, M., Minamimoto, T., Richmond, B. J. (2016). Chemogenetic disconnection of monkey orbitofrontal and rhinal cortex reversibly disrupts reward value. *Nature Neuroscience*, 19(1), 37–39.
- Esber, G. R., & Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: A model of attention in associative learning. *Proceedings of the Royal Society B: Biological Sciences*, 278(1718), 2553–2561.
- Eslinger, P. J., & Damasio, A. R. (1985). Severe disturbance of higher cognition after bilateral frontal lobe ablation: Patient EVR. *Neurology*, 35(12), 1731–1741.
- Fanelli, D. (2018). Opinion: Is science really facing a reproducibility crisis, and do we need it to? *Proceedings of the National Academy of Sciences*, 115(11), 2628–2631.

- Fellows, L. K., Farah, M. J. (2005). Different underlying impairments in decision-making following ventromedial and dorsolateral frontal lobe damage in humans. *Cerebral Cortex*, 15(1), 58-63.
- Ferrier, D. (1876). *The functions of the brain*.
- Ferrier, D., & Yeo, G. F. (1884). XIX. A record of experiments on the effects of lesion of different regions of the cerebral hemispheres. *Philosophical Transactions of the Royal Society of London*, 175, 479–564.
- Fiuzat, E. C., Rhodes, S. E. V., & Murray, E. A. (2017). The Role of Orbitofrontal-Amygdala Interactions in Updating Action-Outcome Valuations in Macaques. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 37(9), 2463–2470.
- Folloni, D., Verhagen, L., Mars, R. B., Fouragnan, E., Constans, C., Aubry, J.-F., Rushworth, M.F.S., Sallet, J. (2019). Manipulation of Subcortical and Deep Cortical Activity in the Primate Brain Using Transcranial Focused Ultrasound Stimulation. *Neuron*, 101(6):1109-1116.e5..
- Fomenko, A., Neudorfer, C., Dallapiazza, R. F., Kalia, S. K., & Lozano, A. M. (2018). Low-intensity ultrasound neuromodulation: An overview of mechanisms and emerging human applications. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, 11(6), 1209–1217.
- Fouragnan, E. F., Chau, B. K. H., Folloni, D., Kolling, N., Verhagen, L., Klein-Flügge, M., Tankelevitch, L., Papageorgiou, G.K., Aubry, J.-F., Sallet, J., Rushworth, M. F. S. (2019). The macaque anterior cingulate cortex translates counterfactual choice value into actual behavioral change. *Nature Neuroscience*, 22(5), 797–808.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27), 9673–9678.
- Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2003). A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 23(12), 5235–5246.

- Freese, J. L., & Amaral, D. G. (2009). Neuroanatomy of the primate amygdala. In *The human amygdala* (pp. 3–42). New York, NY, US: Guilford Press.
- Froudish-Walsh, S., Browning, P. G., Young, J. J., Murphy, K. L., Mars, R. B., Fleysher, L., & Croxson, P. L. (2018). Macro-connectomics and microstructure predict dynamic plasticity patterns in the non-human primate brain. *ELife*, 7.
- Fry, F. J., Ades, H. W., & Fry, W. J. (1958). Production of reversible changes in the central nervous system by ultrasound. *Science (New York, N.Y.)*, 127(3289), 83–84.
- Fry, W. J., & Fry, F. J. (1960). Fundamental Neurological Research and Human Neurosurgery Using Intense Ultrasound. *IRE Transactions on Medical Electronics*, ME-7(3), 166–181.
- Fry, W. J., Fry, F. J., Barnard, J. W., Krumins, R. F., Brennan, J. F., & Fry, W. J. (1955). Ultrasonic Lesions in Mammalian Central Nervous System. *Science*, 122(3179), 1091–1091.
- Fudge, J. L., & Haber, S. N. (2001). Bed nucleus of the stria terminalis and extended amygdala inputs to dopamine subpopulations in primates. *Neuroscience*, 104(3), 807–827.
- Fudge, J. L., Kunishio, K., Walsh, P., Richard, C., & Haber, S. N. (2002). Amygdaloid projections to ventromedial striatal subterritories in the primate. *Neuroscience*, 110(2), 257–275.
- Gabi, M., Neves, K., Masseron, C., Ribeiro, P. F. M., Ventura-Antunes, L., Torres, L., Mota, B., Kaas, J.H., Herculano-Houzel, S. (2016). No relative expansion of the number of prefrontal neurons in primate and human evolution. *Proceedings of the National Academy of Sciences of the United States of America*, 113(34), 9617–9622.
- Gaffan, D. (1994). Interaction of Temporal Lobe and Frontal Lobe in Memory. In A. M. Thierry, J. Glowinski, P. S. Goldman-Rakic, & Y. Christen (Eds.), *Motor and Cognitive Functions of the Prefrontal Cortex* (pp. 129–138). Springer Berlin Heidelberg.
- Gavrilov, L. R., Tsurulnikov, E. M., & Davies, I. A. (1996). Application of focused ultrasound for the stimulation of neural structures. *Ultrasound in Medicine & Biology*, 22(2), 179–192.

- Ghashghaei, H. T., & Barbas, H. (2002). Pathways for emotion: Interactions of prefrontal and anterior temporal pathways in the amygdala of the rhesus monkey. *Neuroscience*, 115(4), 1261–1279.
- Ghashghaei, H. T., Hilgetag, C. C., & Barbas, H. (2007). Sequence of information processing for emotions based on the anatomic dialogue between prefrontal cortex and amygdala. *NeuroImage*, 34(3), 905–923.
- Gilmore, R. O., Diaz, M. T., Wyble, B. A., & Yarkoni, T. (2017). Progress toward openness, transparency, and reproducibility in cognitive neuroscience. *Annals of the New York Academy of Sciences*, 1396(1), 5–18.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615), 171–178.
- Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80, 105–124.
- Goldberg, M. E. (2001). Parietal Lobe. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social & Behavioral Sciences* (pp. 11051–11054).
- Guo, H., Hamilton, M., Offutt, S.J., Gloeckner, C.D., Li, T., Kim, Y., Legon, W., Alford, J.K., Lim, H.H., (2018). Ultrasound Produces Extensive Brain Activation via a Cochlear Pathway. *Neuron* 98, 1020-1030.e4.
- Gutnikov, S. A., Ma, Y. Y., & Gaffan, D. (1997). Temporo-frontal disconnection impairs visual-visual paired association learning but not configural learning in Macaca monkeys. *The European Journal of Neuroscience*, 9(7), 1524–1529.
- Haber, S. N. (2011). Neuroanatomy of Reward: A View from the Ventral Striatum. In J. A. Gottfried (Ed.), *Neurobiology of Sensation and Reward*. Retrieved from
- Harlow, J. M. (1868). *Recovery from the passage of an iron bar through the head*. Boston: Clapp.

- He, B. J., Snyder, A. Z., Zempel, J. M., Smyth, M. D., Raichle, M. E. (2008). Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. *Proceedings of the National Academy of Sciences of the United States of America*, 105(41), 16039-16044.
- Hernádi, I., Grabenhorst, F., & Schultz, W. (2015). Planning activity for internally generated reward goals in monkey amygdala neurons. *Nature Neuroscience*, 18(3), 461–469.
- Heuer, K., Gulban, O. F., Bazin, P.-L., Osoianu, A., Valabregue, R., Santin, M., Herbin, M., Toro, R. (2019). Evolution of neocortical folding: A phylogenetic comparative analysis of MRI from 34 primate species. *Cortex*, 118, 275–291.
- Howard, J. D., & Kahnt, T. (2017). Identity-Specific Reward Representations in Orbitofrontal Cortex Are Modulated by Selective Devaluation. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 37(10), 2627–2638.
- Hunt, L. T., & Hayden, B. Y. (2017). A distributed, hierarchical and recurrent framework for reward-based choice. *Nature Reviews. Neuroscience*, 18(3), 172–182.
- Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F. S., & Behrens, T. E. J. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature Neuroscience*, 15(3), 470–476, S1-3.
- Iversen, S. D., & Mishkin, M. (1970). Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity. *Experimental Brain Research*, 11(4), 376–386.
- Iwai, E., & Yukie, M. (1987). Amygdalofugal and amygdalopetal connections with modality-specific visual cortical areas in macaques (*Macaca fuscata*, *M. mulatta*, and *M. fascicularis*). *The Journal of Comparative Neurology*, 261(3), 362–387.
- Izquierdo, A., & Jentsch, J. D. (2012). Reversal learning as a measure of impulsive and compulsive behavior in addictions. *Psychopharmacology*, 219(2), 607–620.
- Izquierdo, A., & Murray, E. A. (2007). Selective bilateral amygdala lesions in rhesus monkeys fail to disrupt object reversal learning. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(5), 1054–1062.

- Izquierdo, A., & Murray, E. A. (2010). Functional interaction of medial mediodorsal thalamic nucleus but not nucleus accumbens with amygdala and orbital prefrontal cortex is essential for adaptive response selection after reinforcer devaluation. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 30(2), 661–669.
- Izquierdo, A., Suda, R. K., & Murray, E. A. (2004). Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 24(34), 7540–7548.
- Izquierdo, A., Suda, R. K., & Murray, E. A. (2005). Comparison of the effects of bilateral orbital prefrontal cortex lesions and amygdala lesions on emotional responses in rhesus monkeys. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 25(37), 8534–8542.
- Jagannathan, J., Sanghvi, N. T., Crum, L. A., Yen, C.-P., Medel, R., Dumont, A. S., Sheehan, J.P., Steiner, L., Jolesz, F., Kassell, N. F. (2009). High-intensity focused ultrasound surgery of the brain: Part 1--A historical perspective with modern applications. *Neurosurgery*, 64(2), 201–210; discussion 210-211.
- James, W. (1884). What is an emotion ? *Mind*, 1(34), 188–205.
- Janak, P. H., & Tye, K. M. (2015). From circuits to behaviour in the amygdala. *Nature*, 517(7534), 284–292.
- Jbabdi, S., & Johansen-Berg, H. (2011). Tractography: Where do we go from here? *Brain Connectivity*, 1(3), 169–183.
- Jbabdi, S., Lehman, J. F., Haber, S. N., & Behrens, T. E. (2013). Human and monkey ventral prefrontal fibers use the same organizational principles to reach their targets: Tracing versus tractography. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(7), 3190–3201.
- Jbabdi, S., Sotiropoulos, S. N., Savio, A. M., Graña, M., & Behrens, T. E. (2012). Model-based analysis of multi-shell diffusion MR data for tractography: How to get over fitting problems. *Magnetic Resonance in Medicine: Official Journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 68(6), 1846–1855.

- Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, 62(2), 782–790.
- Jenkinson, M., & Chappell, M. (2017). *Introduction to Neuroimaging Analysis*. Oxford, New York: Oxford University Press.
- Jerusalem, A., Al-Rekabi, Z., Chen, H., Ercole, A., Malboubi, M., Tamayo-Elizalde, M., Verhagen, L., Contera, S. (2019). Electrophysiological-mechanical coupling in the neuronal membrane and its role in ultrasound neuromodulation and general anaesthesia. *Acta Biomaterialia*.
- Jocham, G., Brodersen, K. H., Constantinescu, A. O., Kahn, M. C., Ianni, A. M., Walton, M. E., Rushworth, M.F.S., Behrens, T. E. J. (2016). Reward-Guided Learning with and without Causal Attribution. *Neuron*, 90(1), 177–190.
- Jocham, G., Hunt, L. T., Near, J., & Behrens, T. E. (2012). A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nature Neuroscience*, 15(7), 960–961.
- Johansen-Berg, H., Gutman, D. A., Behrens, T. E. J., Matthews, P. M., Rushworth, M. F. S., Katz, E., Lozano, A.M., Mayberg, H. S. (2008). Anatomical connectivity of the subgenual cingulate region targeted with deep brain stimulation for treatment-resistant depression. *Cerebral Cortex (New York, N.Y.: 1991)*, 18(6), 1374–1383.
- Johansen-Berg, Heidi, & Behrens, T. E. J. (2006). Just pretty pictures? What diffusion tractography can add in clinical neuroscience. *Current Opinion in Neurology*, 19(4), 379–385.
- Johansen-Berg, Heidi, & Rushworth, M. F. S. (2009). Using diffusion imaging to study human connectional anatomy. *Annual Review of Neuroscience*, 32, 75–94.
- Johnson, A. W., Gallagher, M., & Holland, P. C. (2009). The basolateral amygdala is critical to the expression of pavlovian and instrumental outcome-specific reinforcer devaluation effects. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 29(3), 696–704.
- Jones, B., & Mishkin, M. (1972). Limbic lesions and the problem of stimulus—Reinforcement associations. *Experimental Neurology*, 36(2), 362–377.

- Jones, D. (2009). Gaussian Modeling of the Diffusion Signal. In *Diffusion MRI* (pp. 37–54).
- Kaas, J. H. (2006). Evolution of the neocortex. *Current Biology: CB*, 16(21), R910-914.
- Kaas, J. H. (2008). The evolution of the complex sensory and motor systems of the human brain. *Brain Research Bulletin*, 75(2–4), 384–390.
- Kaas, J. H. (2013). The evolution of brains from early mammals to humans. *Wiley Interdisciplinary Reviews. Cognitive Science*, 4(1), 33–45.
- Kamali, A., Sair, H. I., Blitz, A. M., Riascos, R. F., Mirbagheri, S., Keser, Z., & Hasan, K. M. (2016). Revealing the ventral amygdalofugal pathway of the human limbic system using high spatial resolution diffusion tensor tractography. *Brain Structure & Function*, 221(7), 3561–3569.
- Kaskan, P. M., Costa, V. D., Eaton, H. P., Zemskova, J. A., Mitz, A. R., Leopold, D. A., Ungerleider, L.G., Murray, E. A. (2017). Learned Value Shapes Responses to Objects in Frontal and Ventral Stream Networks in Macaque Monkeys. *Cerebral Cortex (New York, NY)*, 27(5), 2739–2757.
- Kazama, A., & Bachevalier, J. (2009). Selective aspiration or neurotoxic lesions of orbital frontal areas 11 and 13 spared monkeys' performance on the object discrimination reversal task. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 29(9), 2794–2804.
- Kelley, H. H., Thibaut, J. W., Radloff, R., & Mundy, D. (1962). The development of cooperation in the 'minimal social situation'. *Psychological Monographs: General and Applied*, 76(19), 1–19.
- Kennerley, S. W., Behrens, T. E. J., & Wallis, J. D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature Neuroscience*, 14(12), 1581–1589.
- Kennerley, S. W., & Wallis, J. D. (2009). Encoding of Reward and Space During a Working Memory Task in the Orbitofrontal Cortex and Anterior Cingulate Sulcus. *Journal of Neurophysiology*, 102(6), 3352–3364.
- Kerr, R. A. (1991, November 1). Could the sun be warming the climate?

- Kitt, C. A., Mitchell, S. J., DeLong, M. R., Wainer, B. H., & Price, D. L. (1987). Fiber pathways of basal forebrain cholinergic neurons in monkeys. *Brain Research*, 406(1–2), 192–206.
- Klein-Flügge, M. C., Kennerley, S. W., Friston, K., & Bestmann, S. (2016). Neural Signatures of Value Comparison in Human Cingulate Cortex during Decisions Requiring an Effort-Reward Trade-off. *Journal of Neuroscience*, 36(39), 10002–10015.
- Klingler, J., & Gloor, P. (1960). The connections of the amygdala and of the anterior temporal cortex in the human brain. *The Journal of Comparative Neurology*, 115, 333–369.
- Klüver, H., & Bucy, P. C. (1937). 'Psychic blindness' and other symptoms following bilateral temporal lobectomy in Rhesus monkeys. *American Journal of Physiology*, 119, 352–353.
- Kolb, B. (2007). *Do All Mammals Have a Prefrontal Cortex?* 443–450.
- Kolling, N., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2012). Neural mechanisms of foraging. *Science (New York, N.Y.)*, 336(6077), 95–98.
- Kolling, N., Wittmann, M. K., Behrens, T. E. J., Boorman, E. D., Mars, R. B., & Rushworth, M. F. S. (2016). Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience*, 19(10), 1280–1285.
- Kobayashi, S., Pinto de Carvalho, O., & Schultz, W. (2010). Adaptation of reward sensitivity in orbitofrontal neurons. *The Official Journal of the Society for Neuroscience*, 30(2), 534–544.
- Krettek, J. E., & Price, J. L. (1977). Projections from the amygdaloid complex to the cerebral cortex and thalamus in the rat and cat. *Journal of Comparative Neurology*, 172(4), 687–722.
- Kringelbach, M. L., & Rolls, E. T. (2004). The functional neuroanatomy of the human orbitofrontal cortex: Evidence from neuroimaging and neuropsychology. *Progress in Neurobiology*, 72(5), 341–372.
- Krubitzer, L. (2007). The magnificent compromise: Cortical field evolution in mammals. *Neuron*, 56(2), 201–208.
- Kubaneck, J., (2018). Neuromodulation with transcranial focused ultrasound. *Neurosurgical Focus* 44, E14.

- Kubaneck, J., Shi, J., Marsh, J., Chen, D., Deng, C., & Cui, J. (2016). Ultrasound modulates ion channel currents. *Scientific Reports*, 6(1).
- Kubaneck, J., Shukla, P., Das, A., Baccus, S. A., & Goodman, M. B. (2018). Ultrasound Elicits Behavioral Responses through Mechanical Effects on Neurons and Ion Channels in a Simple Nervous System. *The Journal of Neuroscience*, 38(12), 3081–3091.
- Kunishio, K., & Haber, S. N. (1994). Primate cingulostriatal projection: Limbic striatal versus sensorimotor striatal input. *The Journal of Comparative Neurology*, 350(3), 337–356.
- Kuypers, H. G., Szwarcbart, M. K., Mishkin, M., & Rosvold, H. E. (1965). OCCIPITOTEMPORAL CORTICOCORTICAL CONNECTIONS IN THE RHESUS MONKEY. *Experimental Neurology*, 11, 245–262.
- Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences of the United States of America*, 89(12), 5675–5679.
- Lange, C. G., & James, W. (1922). *A series of reprints and translations. The emotions, Vol. 1. The emotions, Vol. 1.*
- LeDoux, J. (2007). The amygdala. *Current Biology: CB*, 17(20), R868-874.
- Lee, M. H., Smyser, C. D., & Shimony, J. S. (2013). Resting-state fMRI: A review of methods and clinical applications. *AJNR. American Journal of Neuroradiology*, 34(10), 1866–1872.
- Lee, W., Kim, H.-C., Jung, Y., Chung, Y. A., Song, I.-U., Lee, J.-H., & Yoo, S.-S. (2016). Transcranial focused ultrasound stimulation of human primary visual cortex. *Scientific Reports*, 6, 34026.
- Legon, W., Sato, T. F., Opitz, A., Mueller, J., Barbour, A., Williams, A., & Tyler, W. J. (2014). Transcranial focused ultrasound modulates the activity of primary somatosensory cortex in humans. *Nature Neuroscience*, 17(2), 322–329.
- Leh, S. E., Ptito, A., Chakravarty, M. M., & Strafella, A. P. (2007). Fronto-striatal connections in the human brain: A probabilistic diffusion tractography study. *Neuroscience Letters*, 419(2), 113–118.

- Lehman, J. F., Greenberg, B. D., McIntyre, C. C., Rasmussen, S. A., & Haber, S. N. (2011). Rules ventral prefrontal cortical axons use to reach their targets: Implications for diffusion tensor imaging tractography and deep brain stimulation for psychiatric illness. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 31(28), 10392–10402.
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22(6), 1027–1038.
- Li, R., Wu, X., Chen, K., Fleisher, A. S., Reiman, E. M., & Yao, L. (2013). Alterations of directional connectivity among resting-state networks in Alzheimer disease. *AJNR. American Journal of Neuroradiology*, 34(2), 340–345.
- Lichtenberg, N. T., Pennington, Z. T., Holley, S. M., Greenfield, V. Y., Cepeda, C., Levine, M. S., & Wassum, K. M. (2017). Basolateral Amygdala to Orbitofrontal Cortex Projections Enable Cue-Triggered Reward Expectations. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 37(35), 8374–8384.
- Likhtik, E., Popa, D., Apergis-Schoute, J., Fidacaro, G. A., & Paré, D. (2008). Amygdala intercalated neurons are required for expression of fear extinction. *Nature*, 454(7204), 642–645.
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843), 150–157.
- Lozano, A. M., Lipsman, N., Bergman, H., Brown, P., Chabardes, S., Chang, J. W., Matthews, K., McIntyre, C.C., Schlaepfer, T.E., Schulder, M., Temel, Y., Volkmann, J., Krauss, J. K. (2019). Deep brain stimulation: Current challenges and future directions. *Nature Reviews. Neurology*.
- Luk, C.-H., & Wallis, J. D. (2013). Choice Coding in Frontal Cortex during Stimulus-Guided or Action-Guided Decision-Making. *Journal of Neuroscience*, 33(5), 1864–1871.
- Machado, C. J., & Bachevalier, J. (2006). The impact of selective amygdala, orbital frontal cortex, or hippocampal formation lesions on established social relationships in rhesus monkeys (*Macaca mulatta*). *Behavioral Neuroscience*, 120(4), 761–786.

- Mackey, S., & Petrides, M. (2009). Architectonic mapping of the medial region of the human orbitofrontal cortex by density profiles. *Neuroscience*, 159(3), 1089–1107.
- Mackey, Scott, & Petrides, M. (2010). Quantitative demonstration of comparable architectonic areas within the ventromedial and lateral orbital frontal cortex in the human and the macaque monkey brains. *The European Journal of Neuroscience*, 32(11), 1940–1950.
- Mai, J. K., Majtanik, M., & NHMRC, G. P. A. (2015). *Atlas of the Human Brain* (4 edition). Amsterdam: Academic Press.
- Maier-Hein, K. H., Neher, P. F., Houde, J.-C., Côté, M.-A., Garyfallidis, E., Zhong, J., Chamberland, M., Yeh, F.-C., Lin, Y.-C., Ji, Q., Reddick, W.E., Glass, J.O., Chen, D.Q., Feng, Y., Gao, C., Wu, Y., Ma, J., Renjie, H., Li, Q., Westin, C.-F., Deslauriers-Gauthier, S., González, J.O.O., Paquette, M., St-Jean, S., Girard, G., Rheault, F., Sidhu, J., Tax, C.M.W., Guo, F., Mesri, H.Y., Dávid, S., Froeling, M., Heemskerk, A.M., Leemans, A., Boré, A., Pinsard, B., Bedetti, C., Desrosiers, M., Brambati, S., Doyon, J., Sarica, A., Vasta, R., Cerasa, A., Quattrone, A., Yeatman, J., Khan, A.R., Hodges, W., Alexander, S., Romascano, D., Barakovic, M., Auría, A., Esteban, O., Lemkaddem, A., Thiran, J.-P., Cetingul, H.E., Odry, B.L., Mailhe, B., Nadar, M.S., Pizzagalli, F., Prasad, G., Villalon-Reina, J.E., Galvis, J., Thompson, P.M., Requejo, F.D.S., Laguna, P.L., Lacerda, L.M., Barrett, R., Dell'Acqua, F., Catani, M., Petit, L., Caruyer, E., Daducci, A., Dyrby, T.B., Holland-Letz, T., Hilgetag, C.C., Stieltjes, B., Descoteaux, M. (2017). The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications*, 8(1), 1349.
- Malpighi, M. (1697). *Opera posthuma*. Churchill.
- Mansfield, P. (1977). Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 10(3), L55–L58.
- Margulies, D. S., Vincent, J. L., Kelly, C., Lohmann, G., Uddin, L. Q., Biswal, B. B., Villringer, A., Castellanos, F.X., Milham, M.P., Petrides, M. (2009). Precuneus shares intrinsic functional architecture in humans and monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, 106(47), 20069–20074.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman.

- Mars, R. B., Foxley, S., Verhagen, L., Jbabdi, S., Sallet, J., Noonan, M. P., Neubert, F.-X., Andersson, J.L., Croxson, P.L., Dunbar, R.I.M., Khrapitchev, A.A., Sibson, N.R., Miller, K.L., Rushworth, M. F. S. (2016a). The extreme capsule fiber complex in humans and macaque monkeys: A comparative diffusion MRI tractography study. *Brain Structure & Function*, 221(8), 4059–4071.
- Mars, R. B., Neubert, F.-X., Verhagen, L., Sallet, J., Miller, K. L., Dunbar, R. I. M., & Barton, R. A. (2014). Primate comparative neuroscience using magnetic resonance imaging: Promises and challenges. *Frontiers in Neuroscience*, 8, 298.
- Mars, R. B., Passingham, R. E., & Jbabdi, S. (2018a). Connectivity Fingerprints: From Areal Descriptions to Abstract Spaces. *Trends in Cognitive Sciences*, 22(11), 1026–1037.
- Mars, R. B., Sotiropoulos, S. N., Passingham, R. E., Sallet, J., Verhagen, L., Khrapitchev, A. A., Sibson, N., Jbabdi, S. (2018b). Whole brain comparative anatomy using connectivity blueprints. *ELife*, 7.
- Mars, R. B., Verhagen, L., Gladwin, T. E., Neubert, F.-X., Sallet, J., & Rushworth, M. F. S. (2016b). Comparing brains by matching connectivity profiles. *Neuroscience and Biobehavioral Reviews*, 60, 90–97.
- Masouleh, S., Eickhoff, S. B., Hoffstaedter, F., & Genon, S. (2019). Empirical examination of the replicability of associations between brain structure and psychological variables. *ELife*, 8.
- Mayberg, H. S., Lozano, A. M., Voon, V., McNeely, H. E., Seminowicz, D., Hamani, C., Schwalb, J.M., Kennedy, S. H. (2005). Deep brain stimulation for treatment-resistant depression. *Neuron*, 45(5), 651–660.
- McDonald, A. J. (1991). Organization of amygdaloid projections to the prefrontal cortex and associated striatum in the rat. *Neuroscience*, 44(1), 1–14.
- McDonald, A. J. (1992). Projection neurons of the basolateral amygdala: A correlative Golgi and retrograde tract tracing study. *Brain Research Bulletin*, 28(2), 179–185.
- McDonald, A. J. (1998). Cortical pathways to the mammalian amygdala. *Progress in Neurobiology*, 55(3), 257–332.

- Meunier, M., Bachevalier, J., Murray, E. A., Málková, L., & Mishkin, M. (1999). Effects of aspiration versus neurotoxic lesions of the amygdala on emotional responses in monkeys. *The European Journal of Neuroscience*, 11(12), 4403–4418.
- Miller, K. L., Stagg, C. J., Douaud, G., Jbabdi, S., Smith, S. M., Behrens, T. E. J., Jenkinson, M., Chance, S.A., Esiri, M.M., Voets, N.L., Jenkinson, N., Aziz, T.Z., Turner, M.R., Johansen-Berg, H., McNab, J. A. (2011). Diffusion imaging of whole, post-mortem human brains on a clinical MRI scanner. *NeuroImage*, 57(1), 167–181.
- Montague, P. R., & Berns, G. S. (2002). Neural Economics and the Biological Substrates of Valuation. *Neuron*, 36(2), 265–284.
- Monti, M. M., Schnakers, C., Korb, A. S., Bystritsky, A., & Vespa, P. M. (2016). Non-Invasive Ultrasonic Thalamic Stimulation in Disorders of Consciousness after Severe Brain Injury: A First-in-Man Report. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, 9(6), 940–941.
- Morecraft, R. J., Geula, C., & Mesulam, M. M. (1992). Cytoarchitecture and neural afferents of orbitofrontal cortex in the brain of the monkey. *The Journal of Comparative Neurology*, 323(3), 341–358.
- Moroni, M., Servin-Vences, M. R., Fleischer, R., Sánchez-Carranza, O., & Lewin, G. R. (2018). Voltage gating of mechanosensitive PIEZO channels. *Nature Communications*, 9(1), 1–15.
- Morrison, S. E., Saez, A., Lau, B., & Salzman, C. D. (2011). Different time courses for learning-related changes in amygdala and orbitofrontal cortex. *Neuron*, 71(6), 1127–1140.
- Morrison, S. E., & Salzman, C. D. (2009). The convergence of information about rewarding and aversive stimuli in single neurons. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 29(37), 11471–11483.
- Morrison, S. E., & Salzman, C. D. (2010). Re-valuing the amygdala. *Current Opinion in Neurobiology*, 20(2), 221–230.
- Mosher, C. P., & Rudebeck, P. H. (2015). The amygdala accountant: New tricks for an old structure. *Nature Neuroscience*, 18(3), 324–325.
- Müller, R.-A., Shih, P., Keehn, B., Deyoe, J. R., Leyden, K. M., & Shukla, D. K. (2011). Underconnected, but How? A Survey of Functional

- Connectivity MRI Studies in Autism Spectrum Disorders. *Cerebral Cortex*, 21(10), 2233–2243.
- Munoz, F., Aurup, C., Konofagou, E. E., & Ferrera, V. P. (2018). Modulation of Brain Function and Behavior by Focused Ultrasound. *Current Behavioral Neuroscience Reports*, 5(2), 153–164.
- Munuera, J., Rigotti, M., & Salzman, C. D. (2018). Shared neural coding for social hierarchy and reward value in primate amygdala. *Nature Neuroscience*, 21(3), 415–423.
- Murray, E. A. (1992). Medial temporal lobe structures contributing to recognition memory: The amygdaloid complex versus the rhinal cortex. In *The amygdala: Neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 453–470). New York, NY, US: Wiley-Liss.
- Murray, E. A., & Izquierdo, A. (2007). Orbitofrontal cortex and amygdala contributions to affect and action in primates. *Annals of the New York Academy of Sciences*, 1121, 273–296.
- Murray, E. A., Moylan, E. J., Saleem, K. S., Basile, B. M., & Turchi, J. (2015). Specialized areas for value updating and goal selection in the primate orbitofrontal cortex. *ELife*, 4.
- Murray, E. A., & Rudebeck, P. H. (2018). Specializations for reward-guided decision-making in the primate ventral prefrontal cortex. *Nature Reviews. Neuroscience*, 19(7), 404–417.
- Murray, E. A., & Wise, S. P. (2010). Interactions between orbital prefrontal cortex and amygdala: Advanced cognition, learned responses and instinctive behaviors. *Current Opinion in Neurobiology*, 20(2), 212–220.
- Murray, E. A., Wise, S. P., & Drevets, W. C. (2011). Localization of dysfunction in major depressive disorder: Prefrontal cortex and amygdala. *Biological Psychiatry*, 69(12), e43-54.
- Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, 17(12), 1661–1663.
- Namburi, P., Al-Hasani, R., Calhoon, G. G., Bruchas, M. R., & Tye, K. M. (2016). Architectural Representation of Valence in the Limbic System.

- Namburi, P., Beyeler, A., Yorozu, S., Calhoon, G. G., Halbert, S. A., Wichmann, R., Holden, S.S., Mertens, K.L., Anahtar, M., Felix-Ortiz, A.C., Wickersham, I.R., Gray, J.M., Tye, K. M. (2015). A circuit mechanism for differentiating positive and negative associations. *Nature*, 520(7549), 675–678.
- Nauta, W. J. (1961). Fibre degeneration following lesions of the amygdaloid complex in the monkey. *Journal of Anatomy*, 95, 515–531.
- Nauta, W. J., & Gyax, P. A. (1954). Silver impregnation of degenerating axons in the central nervous system: A modified technic. *Stain Technology*, 29(2), 91–93.
- Neubert, F.-X., Mars, R. B., Sallet, J., & Rushworth, M. F. S. (2015). Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 112(20), E2695-2704.
- Neubert, F.-X., Mars, R. B., Thomas, A. G., Sallet, J., & Rushworth, M. F. S. (2014). Comparison of Human Ventral Frontal Cortex Areas for Cognitive Control and Language with Areas in Monkey Frontal Cortex. *Neuron*, 81(3), 700–713.
- Nieuwenhuys, R., Voogd, J., & Huijzen, C. van. (2008). *The Human Central Nervous System: A Synopsis and Atlas* (4th ed.).
- Noonan, M. P., Kolling, N., Walton, M. E., & Rushworth, M. F. S. (2012). Re-evaluating the role of the orbitofrontal cortex in reward and reinforcement. *The European Journal of Neuroscience*, 35(7), 997–1010.
- Noonan, M. P., Mars, R. B., & Rushworth, M. F. S. (2011). Distinct roles of three frontal cortical areas in reward-guided behavior. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 31(40), 14399–14412.
- Noonan, M. P., Mars, R. B., Sallet, J., Dunbar, R. I. M., & Fellows, L. K. (2018). The structural and functional brain networks that support human social networks. *Behavioural Brain Research*, 355, 12–23.
- Noonan, M. P., Walton, M. E., Behrens, T. E. J., Sallet, J., Buckley, M. J., & Rushworth, M. F. S. (2010). Separate value comparison and learning

mechanisms in macaque medial and lateral orbitofrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 107(47), 20547–20552.

- Noonan, M. P., Chau, B. K. H., Rushworth, M. F. S., & Fellows, L. K. (2017). Contrasting Effects of Medial and Lateral Orbitofrontal Cortex Lesions on Credit Assignment and Decision-Making in Humans. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 37(29), 7023–7035.
- Noonan, M. P., Sallet, J., Mars, R. B., Neubert, F. X., O'Reilly, J. X., Andersson, J. L., Mitchell, A.S., Bell, A.H., Miller, K.L., Rushworth, M. F. S. (2014). A neural circuit covarying with social hierarchy in macaques. *PLoS Biology*, 12(9), e1001940.
- Novotny, G. E. (1977). A ventral tract from the bed nucleus of the stria terminalis to the amygdaloid complex in the monkey (*Macaca fascicularis*): Possible pathway in the regulation of ovulation [proceedings]. *The Journal of Physiology*, 266(1), 13P-14P.
- O'Doherty, J. P. (2014). The problem with value. *Neuroscience and Biobehavioral Reviews*, 43, 259–268.
- Ogawa, M., van der Meer, M. A. A., Esber, G. R., Cerri, D. H., Stalnaker, T. A., & Schoenbaum, G. (2013). Risk-responsive orbitofrontal neurons track acquired salience. *Neuron*, 77(2), 251–258.
- Ogawa, S., & Lee, T. M. (1990). Magnetic resonance imaging of blood vessels at high fields: In vivo and in vitro measurements and image simulation. *Magnetic Resonance in Medicine*, 16(1), 9–18.
- Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990a). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, 87(24), 9868–9872.
- Ogawa, S., Lee, T. M., Nayak, A. S., & Glynn, P. (1990b). Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magnetic Resonance in Medicine*, 14(1), 68–78.
- Ogawa, Seiji. (2012). Finding the BOLD effect in brain images. *NeuroImage*, 62(2), 608–609.
- Oler, J. A., Tromp, D. P. M., Fox, A. S., Kovner, R., Davidson, R. J., Alexander, A. L., McFarlin, D.R., Birn, R.M., Berg, B., deCampo, D.M., Kalin, N.H., Fudge, J. L. (2017). Connectivity between the central

- nucleus of the amygdala and the bed nucleus of the stria terminalis in the non-human primate: Neuronal tract tracing and developmental neuroimaging studies. *Brain Structure & Function*, 222(1), 21–39.
- Ongür, D., & Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral Cortex (New York, N.Y.: 1991)*, 10(3), 206–219.
- Ongür, D., Ferry, A. T., & Price, J. L. (2003). Architectonic subdivision of the human orbital and medial prefrontal cortex. *The Journal of Comparative Neurology*, 460(3), 425–449.
- Onufrowicz, W. (1887). Das balkenlose Mikrocephalengehirn Hofmann. *Archiv für Psychiatrie und Nervenkrankheiten*, 18(2), 305–328.
- O'Reilly, J. X., Croxson, P. L., Jbabdi, S., Sallet, J., Noonan, M. P., Mars, R. B., Browning, P.G.F., Wilson, C.R.E., Mitchell, A.S., Miller, K.L., Rushworth, M.F.S., Baxter, M. G. (2013). Causal effect of disconnection lesions on interhemispheric functional connectivity in rhesus monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, 110(34), 13982–13987.
- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: A good-based model. *Annual Review of Neuroscience*, 34, 333–359.
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090), 223–226.
- Padoa-Schioppa, C., & Assad, J. A. (2008). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nature Neuroscience*, 11(1), 95–102.
- Pandya, D. N., Seltzer, B., & Barbas, H. (1988). Input-output organization of the primate cerebral cortex. In *Comparative Primate Biology* (2nd ed., Vol. 4, pp. 39–80). New York.
- Pandya, D., Seltzer, B., Petrides, M., & Cipolloni, P. B. (2015). *Cerebral Cortex: Architecture, Connections, and the Dual Origin Concept*. Oxford, New York: Oxford University Press.
- Papageorgiou, G. K., Sallet, J., Wittmann, M. K., Chau, B. K. H., Schüffelgen, U., Buckley, M. J., & Rushworth, M. F. S. (2017). Inverted activity patterns in ventromedial prefrontal cortex during value-guided decision-making in a less-is-more task. *Nature Communications*, 8.

- Parker, A., & Gaffan, D. (1998). Memory after frontal/temporal disconnection in monkeys: Conditional and non-conditional tasks, unilateral and bilateral frontal lesions. *Neuropsychologia*, 36(3), 259–271.
- Passingham, R. (2008). *What is special about the human brain?* Oxford, New York: Oxford University Press.
- Passingham, R. E., Toni, I., & Rushworth, M. F. (2000). Specialisation within the prefrontal cortex: The ventral prefrontal cortex and associative learning. *Experimental Brain Research*, 133(1), 103–113.
- Passingham, R. E., & Wise, S. P. (2012). *The neurobiology of the prefrontal cortex: Anatomy, evolution, and the origin of insight* (1st ed). Oxford, United Kingdom: Oxford University Press.
- Passingham, R. E., & Smaers, J. B. (2014). Is the prefrontal cortex especially enlarged in the human brain allometric relations and remapping factors. *Brain, Behavior and Evolution*, 84(2), 156–166.
- Passingham, R. E., Stephan, K.E., & Kötter, R. (2002). The anatomical basis of functional localization in the cortex. *Nature Reviews Neuroscience*, 3(8), 606–616.
- Pauling, L., & Coryell, C. D. (1936). The Magnetic Properties and Structure of Hemoglobin, Oxyhemoglobin and Carbonmonoxyhemoglobin. *Proceedings of the National Academy of Sciences of the United States of America*, 22(4), 210–216.
- Paxinos, G., Huang, X. F., & Toga, A. W. (2000). *The Rhesus Monkey Brain in Stereotaxic Coordinates*. Academic Press.
- Pemment, J. (2013). Psychopathy versus sociopathy: Why the distinction has become crucial. *Aggression and Violent Behavior*, 18(5), 458–461.
- Petrides, M., & Pandya, D. N. (1984). Projections to the frontal cortex from the posterior parietal region in the rhesus monkey. *The Journal of Comparative Neurology*, 228(1), 105–116.
- Petrides, M., & Pandya, D. N. (1988). Association fiber pathways to the frontal cortex from the superior temporal region in the rhesus monkey. *The Journal of Comparative Neurology*, 273(1), 52–66.
- Petrides, M. (2005). Lateral prefrontal cortex: Architectonic and functional organization. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 360(1456), 781–795.

- Petrides, Michael, & Pandya, D. N. (2007). Efferent association pathways from the rostral prefrontal cortex in the macaque monkey. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(43), 11573–11586.
- Petrides, Michael, Tomaiuolo, F., Yeterian, E. H., & Pandya, D. N. (2012). The prefrontal cortex: Comparative architectonic organization in the human and the macaque monkey brains. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 48(1), 46–57.
- Plassmann, H., O'Doherty, J. P., & Rangel, A. (2010). Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 30(32), 10799–10808.
- Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(37), 9984–9988.
- Polanía, R., Nitsche, M. A., & Ruff, C. C. (2018). Studying and modifying brain function with non-invasive brain stimulation. *Nature Neuroscience*, 21(2), 174–187.
- Porrino, L. J., & Goldman-Rakic, P. S. (1982). Brainstem innervation of prefrontal and anterior cingulate cortex in the rhesus monkey revealed by retrograde transport of HRP. *The Journal of Comparative Neurology*, 205(1), 63–76.
- Preuss, T. M. (2011). The human brain: Rewired and running hot. *Annals of the New York Academy of Sciences*, 1225 Suppl 1, E182-191.
- Price, J. L., & Amaral, D. G. (1981). An autoradiographic study of the projections of the central nucleus of the monkey amygdala. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 1(11), 1242–1259.
- Price, J. L. (2003). Comparative aspects of amygdala connectivity. *Annals of the New York Academy of Sciences*, 985, 50–58.
- Price, J. L. (2005). Free will versus survival: brain systems that underlie intrinsic constraints on behavior. *Journal of Comparative Neurology* 493, 132–139.

- Price, J.L. (2007). Definition of the orbital cortex in relation to specific connections with limbic and visceral structures and other cortical regions. *Annals of the New York Academy of Sciences*, 1121, 54–71.
- Prieto, M. L., Ömer, O., Khuri-Yakub, B. T., & Maduke, M. C. (2013). Dynamic response of model lipid membranes to ultrasonic radiation force. *PloS One*, 8(10), e77115.
- Pujara, M. S., Rudebeck, P. H., Ciesinski, N. K., Murray, E. A. (2019). Heightened Defensive Responses Following Subtotal Lesions of Macaque Orbitofrontal Cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 39(21), 4133-4141.
- Quadri, S. A., Waqas, M., Khan, I., Khan, M. A., Suriya, S. S., Farooqui, M., & Fiani, B. (2018). High-intensity focused ultrasound: Past, present, and future in neurosurgery. *Neurosurgical Focus*, 44(2), E16.
- Ray, J. P., & Price, J. L. (1992). The organization of the thalamocortical connections of the mediodorsal thalamic nucleus in the rat, related to the ventral forebrain-prefrontal cortex topography. *The Journal of Comparative Neurology*, 323(2), 167–197.
- Reekie, Y. L., Braesicke, K., Man, M. S., & Roberts, A. C. (2008). Uncoupling of behavioral and autonomic responses after lesions of the primate orbitofrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 105(28), 9787-2792.
- Rescorla, R. A., & Wagner, A. R. (1972). 3 A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement. 18.
- Reveley, C., Gruslys, A., Ye, F. Q., Glen, D., Samaha, J., E Russ, B., Saad, Z., K Seth, A., Leopold, D.A., Saleem, K. S. (2017). Three-Dimensional Digital Template Atlas of the Macaque Brain. *Cerebral Cortex (New York, N.Y.: 1991)*, 27(9), 4463–4477.
- Rich, E. L., & Wallis, J. D. (2014). Medial-lateral organization of the orbitofrontal cortex. *Journal of Cognitive Neuroscience*, 26(7), 1347–1362.
- Riva-Posse, P., Choim K. S., Holtzheimer, P. E., McIntyre, C. C., Gross, R. E., Chaturvedi, A., Crowell, A. L., Garlow, S. J., Rajendra, J. K., Mayberg, H. S. (2014). Defining critical white matter pathways mediating successful subcallosal cingulate deep brain stimulation for treatment-resistant depression. *Biological Psychiatry*, 76(12), 963-769.

- Robinson, B. W., & Mishkin, M. (1968). Penile erection evoked from forebrain structures in *Macaca mulatta*. *Archives of Neurology*, 19(2), 184–198.
- Rolls, E. T., & Baylis, L. L. (1994). Gustatory, olfactory, and visual convergence within the primate orbitofrontal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 14(9), 5437–5452.
- Rolls, E. T. (1996). The orbitofrontal cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 351(1346), 1433–1443; discussion 1443–1444.
- Rolls, E. T. (1999). *The Brain and Emotion*. Oxford University Press.
- Rorden, C., & Karnath, H.-O. (2004). Using human brain lesions to infer function: A relic from a past era in the fMRI age? *Nature Reviews. Neuroscience*, 5(10), 813–819.
- Rose, J. E., & Woolsey, C. N. (1948). The orbitofrontal cortex and its connections with the mediodorsal nucleus in rabbit, sheep and cat. *Research Publications - Association for Research in Nervous and Mental Disease*, 27 (1 vol.), 210–232.
- Rudebeck, P. H., Mitz, A. R., Chacko, R. V., & Murray, E. A. (2013a). Effects of amygdala lesions on reward-value coding in orbital and medial prefrontal cortex. *Neuron*, 80(6), 1519–1531.
- Rudebeck, P. H., & Murray, E. A. (2008). Amygdala and orbitofrontal cortex lesions differentially influence choices during object reversal learning. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 28(33), 8338–8343.
- Rudebeck, P. H., & Murray, E. A. (2011). Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 31(29), 10569–10578.
- Rudebeck, P. H., & Murray, E. A. (2014). The orbitofrontal oracle: Cortical mechanisms for the prediction and evaluation of specific behavioral outcomes. *Neuron*, 84(6), 1143–1156.
- Rudebeck, P. H., Putnam, P. T., Daniels, T. E., Yang, T., Mitz, A. R., Rhodes, S. E. V., & Murray, E. A. (2014). A role for primate subgenual cingulate cortex in sustaining autonomic arousal. *Proceedings of the National Academy of Sciences of the United States of America*, 111(14), 5391–5396.

- Rudebeck, P. H., & Rich, E. L. (2018). Orbitofrontal cortex. *Current Biology: CB*, 28(18), R1083–R1088.
- Rudebeck, P. H., Ripple, J. A., Mitz, A. R., Averbeck, B. B., & Murray, E. A. (2017a). Amygdala Contributions to Stimulus-Reward Encoding in the Macaque Medial and Orbital Frontal Cortex during Learning. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 37(8), 2186–2202.
- Rudebeck, P. H., Saunders, R. C., Lundgren, D. A., & Murray, E. A. (2017b). Specialized Representations of Value in the Orbital and Ventrolateral Prefrontal Cortex: Desirability versus Availability of Outcomes. *Neuron*, 95(5), 1208–1220.e5.
- Rudebeck, P. H., Saunders, R. C., Prescott, A. T., Chau, L. S., & Murray, E. A. (2013b). Prefrontal mechanisms of behavioral flexibility, emotion regulation and value updating. *Nature Neuroscience*, 16(8), 1140–1145.
- Rushworth, M. F. S., Boorman, E., & Mars, R. B. (2009). Chapter 20—Comparing Brain Connections in Different Species using Diffusion Weighted Imaging. In Heidi Johansen-Berg & T. E. J. Behrens (Eds.), *Diffusion MRI* (pp. 445–460).
- Rushworth, M. F. S., Buckley, M. J., Gough, P. M., Alexander, I. H., Kyriazis, D., McDonald, K. R., & Passingham, R. E. (2005). Attentional selection and action selection in the ventral and orbital prefrontal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 25(50), 11628–11636.
- Rushworth, M. F. S., Kolling, N., Sallet, J., & Mars, R. B. (2012). Valuation and decision-making in frontal cortex: One or many serial or parallel systems? *Current Opinion in Neurobiology*, 22(6), 946–955.
- Rushworth, M. F. S., Noonan, M. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal Cortex and Reward-Guided Learning and Decision-Making. *Neuron*, 70(6), 1054–1069.
- Russchen, F. T., Amaral, D. G., & Price, J. L. (1985). The afferent connections of the substantia innominata in the monkey, *Macaca fascicularis*. *The Journal of Comparative Neurology*, 242(1), 1–27.
- Rygula, R., Walker, S. C., Clarke, H. F., Robbins, T. W., & Roberts, A. C. (2010). Differential contributions of the primate ventrolateral prefrontal and orbitofrontal cortex to serial reversal learning. *The Journal of*

Neuroscience: The Official Journal of the Society for Neuroscience, 30(43), 14552–14559.

- Saez, A., Rigotti, M., Ostojic, S., Fusi, S., & Salzman, C. D. (2015). Abstract Context Representations in Primate Amygdala and Prefrontal Cortex. *Neuron*, 87(4), 869–881.
- Saez, R. A., Saez, A., Paton, J. J., Lau, B., & Salzman, C. D. (2017). Distinct Roles for the Amygdala and Orbitofrontal Cortex in Representing the Relative Amount of Expected Reward. *Neuron*, 95(1), 70-77.e3.
- Saleem, K. S., Kondo, H., & Price, J. L. (2008). Complementary circuits connecting the orbital and medial prefrontal networks with the temporal, insular, and opercular cortex in the macaque monkey. *The Journal of Comparative Neurology*, 506(4), 659–693.
- Saleem, K. S., & Logothetis, N. K. (2012). *A Combined MRI and Histology Atlas of the Rhesus Monkey Brain in Stereotaxic Coordinates*. Academic Press.
- Sallet, J., Mars, R. B., Noonan, M. P., Andersson, J. L., O'Reilly, J. X., Jbabdi, S., Jbabdi, S., Croxson, P.L., Jenkinson, M., Miller, K.L., Rushworth, M. F. S. (2011). Social network size affects neural circuits in macaques. *Science (New York, N.Y.)*, 334(6056), 697–700.
- Sallet, J., Mars, R. B., Noonan, M. P., Neubert, F.-X., Jbabdi, S., O'Reilly, J. X., J.X., Filippini, N., Thomas, A.G., Rushworth, M. F. (2013). The Organization of Dorsal Frontal Cortex in Humans and Macaques. *Journal of Neuroscience*, 33(30), 12255–12274.
- Sato, T., Shapiro, M. G., & Tsao, D. Y. (2018). Ultrasonic Neuromodulation Causes Widespread Cortical Activation via an Indirect Auditory Mechanism. *Neuron*, 98(5), 1031-1041.e5.
- Schmahmann, J. D., & Pandya, D. N. (2006). *Fiber Pathways of the Brain*. Oxford University Press, Oxford, United Kingdom.
- Schmahmann, J. D., Pandya, D. N., Wang, R., Dai, G., D'Arceuil, H. E., de Crespigny, A. J., & Wedeen, V. J. (2007). Association fibre pathways of the brain: Parallel observations from diffusion spectrum imaging and autoradiography. *Brain: A Journal of Neurology*, 130(Pt 3), 630–653.
- Schoenbaum, G., Roesch, M. R., Stalnaker, T. A., & Takahashi, Y. K. (2009). A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nature Reviews. Neuroscience*, 10(12), 885–892.

- Schoenbaum, G., Setlow, B., Saddoris, M. P., & Gallagher, M. (2003). Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. *Neuron*, 39(5), 855–867.
- Schoenbaum, G., Takahashi, Y., Liu, T.-L., & McDannald, M. A. (2011). Does the orbitofrontal cortex signal value? *Annals of the New York Academy of Sciences*, 1239, 87–99.
- Schoenemann, P. T., Sheehan, M. J., & Glotzer, L. D. (2005). Prefrontal white matter volume is disproportionately larger in humans than in other primates. *Nature Neuroscience*, 8(2), 242–252.
- Scholz, J., Klein, M. C., Behrens, T. E. J., & Johansen-Berg, H. (2009). Training induces changes in white matter architecture. *Nature Neuroscience*, 12(11), 1370–1371.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80(1), 1–27.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, 275(5306), 1593–1599.
- Schwartzbaum, J. S., & Poulos, D. A. (1965). Discrimination behavior after amygdectomy in monkeys: Learning set and discrimination reversals. *Journal of Comparative and Physiological Psychology*, 60(3), 320–328.
- Seamans, J. K., Lapish, C. C., & Durstewitz, D. (2008). Comparing the prefrontal cortex of rats and primates: Insights from electrophysiology. *Neurotoxicity Research*, 14(2–3), 249–262.
- Sherwood, C. C., & Smaers, J. B. (2013). What's the fuss over human frontal lobe evolution? *Trends in Cognitive Sciences*, 17(9), 432–433.
- Sierra, C., Acosta, C., Chen, C., Wu, S.-Y., Karakatsani, M. E., Bernal, M., & Konofagou, E. E. (2017). Lipid microbubbles as a vehicle for targeted drug delivery using focused ultrasound-induced blood-brain barrier opening. *Journal of Cerebral Blood Flow and Metabolism: Official Journal of the International Society of Cerebral Blood Flow and Metabolism*, 37(4), 1236–1250.
- Smaers, J. B., Gómez-Robles, A., Parks, A. N., & Sherwood, C. C. (2017). Exceptional Evolutionary Expansion of Prefrontal Cortex in Great Apes and Humans. *Current Biology: CB*, 27(5), 714–720.

- Smith, S. M., Hyvärinen, A., Varoquaux, G., Miller, K. L., & Beckmann, C. F. (2014). Group-PCA for very large fMRI datasets. *NeuroImage*, 101, 738–749.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage, Mathematics in Brain Imaging* 23, S208–S219.
- Sotiropoulos, S. N., Jbabdi, S., Xu, J., Andersson, J. L., Moeller, S., Auerbach, E. J., Glasser, M.F., Hernandez, M., Sapiro, G., Jenkinson, M., Feinberg, D.A., Yacoub, E., Lenglet, C., Van Essen, D.C., Ugurbil, K., Behrens, T.E.J., WU-Minn HCP Consortium. (2013). Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *NeuroImage*, 80, 125–143.
- Spiegler, B. J., & Mishkin, M. (1981). Evidence for the sequential participation of inferior temporal cortex and amygdala in the acquisition of stimulus-reward associations. *Behavioural Brain Research*, 3(3), 303–317.
- Stalnaker, T. A., Cooch, N. K., & Schoenbaum, G. (2015). What the orbitofrontal cortex does not do. *Nature Neuroscience*, 18(5), 620–627.
- Stalnaker, T. A., Roesch, M. R., Calu, D. J., Burke, K. A., Singh, T., & Schoenbaum, G. (2007). Neural correlates of inflexible behavior in the orbitofrontal-amygdalar circuit after cocaine exposure. *Annals of the New York Academy of Sciences*, 1121, 598–609.
- Stejskal, E. O., & Tanner, J. E. (1965). Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient. *The Journal of Chemical Physics*, 42(1), 288–292.
- Strait, C. E., Blanchard, T. C., & Hayden, B. Y. (2014). Reward Value Comparison via Mutual Inhibition in Ventromedial Prefrontal Cortex. *Neuron*, 82(6), 1357–1366.
- Sugase-Miyamoto, Y., Liu, Z., Wiener, M. C., Optican, L. M., & Richmond, B. J. (2008). Short-term memory trace in rapidly adapting synapses of inferior temporal cortex. *PLoS Computational Biology*, 4(5), e1000073.
- Sugase-Miyamoto, Y., & Richmond, B. J. (2005). Neuronal signals in the monkey basolateral amygdala during reward schedules. *The Journal*

of Neuroscience: The Official Journal of the Society for Neuroscience, 25(48), 11071–11083.

Sutton, R. S., Barto, A. G., & Williams, R. J. (1992). Reinforcement learning is direct adaptive optimal control. *IEEE Control Systems Magazine*, 12(2), 19–22.

Sutton, Richard S., & Barto, A. G. (1998). *Introduction to Reinforcement Learning* (1st ed.). Cambridge, MA, USA: MIT Press.

Szczepanski, S. M., & Knight, R. T. (2014). Insights into human behavior from lesions to the prefrontal cortex. *Neuron*, 83(5), 1002–1018.

Takahashi, M., Hackney, D. B., Zhang, G., Wehrli, S. L., Wright, A. C., O'Brien, W. T., Uematsu, H., Wehrli, F.W., Selzer, M. E. (2002). Magnetic resonance microimaging of intraaxonal water diffusion in live excised lamprey spinal cord. *Proceedings of the National Academy of Sciences*, 99(25), 16192–16196.

Takahashi, Y. K., Chang, C. Y., Lucantonio, F., Haney, R. Z., Berg, B. A., Yau, H.-J., Bonci, A., Schoenbaum, G. (2013). Neural estimates of imagined outcomes in the orbitofrontal cortex drive behavior and learning. *Neuron*, 80(2), 507–518.

ter Haar, G. (2007). Therapeutic applications of ultrasound. *Progress in Biophysics and Molecular Biology*, 93(1), 111–129.

Thiebaut de Schotten, M., Croxson, P. L., & Mars, R. B. (2019). Large-scale comparative neuroimaging: Where are we and what do we need? *Cortex*, 118, 188–202.

Thiebaut de Schotten, M., Dell'Acqua, F., Forkel, S. J., Simmons, A., Vergani, F., Murphy, D. G. M., & Catani, M. (2011). A lateralized brain network for visuospatial attention. *Nature Neuroscience*, 14(10), 1245–1246.

Thiebaut de Schotten, M., Dell'Acqua, F., Valabregue, R., & Catani, M. (2012). Monkey to human comparative anatomy of the frontal lobe association tracts. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 48(1), 82–96.

Thiebaut de Schotten, M., & Foulon, C. (2018). The rise of a new associationist school for lesion-symptom mapping. *Brain: A Journal of Neurology*, 141(1), 2–4.

Thorndike, E. L. (1933). A proof of the law of effect. *Science*, 77, 173–175.

- Thorndike, E. L. (1911). *Animal intelligence; experimental studies*.
- Thorpe, S. J., Rolls, E. T., & Maddison, S. (1983). The orbitofrontal cortex: Neuronal activity in the behaving monkey. *Experimental Brain Research*, 49(1), 93–115.
- Thulborn, K. R., Waterton, J. C., Matthews, P. M., & Radda, G. K. (1982). Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochimica Et Biophysica Acta*, 714(2), 265–270.
- Timbie, C., & Barbas, H. (2014). Specialized pathways from the primate amygdala to posterior orbitofrontal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 34(24), 8106–8118.
- Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science (New York, N.Y.)*, 315(5811), 515–518.
- Tremblay, L., & Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature*, 398(6729), 704–708.
- Tsujimoto, S., Genovesio, A., & Wise, S. P. (2009). Monkey Orbitofrontal Cortex Encodes Response Choices Near Feedback Time. *Journal of Neuroscience*, 29(8), 2569–2574.
- Tufail, Y., Yoshihiro, A., Pati, S., Li, M. M., & Tyler, W. J. (2011). Ultrasonic neuromodulation by brain stimulation with transcranial ultrasound. *Nature Protocols*, 6(9), 1453–1470.
- Turner, R., Le Bihan, D., Moonen, C. T., Despres, D., & Frank, J. (1991). Echo-planar time course MRI of cat brain oxygenation changes. *Magnetic Resonance in Medicine*, 22(1), 159–166.
- Tyler, W. J. (2011). Noninvasive neuromodulation with ultrasound? A continuum mechanics hypothesis. *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, 17(1), 25–36.
- Tyler, W. J., Lani, S. W., & Hwang, G. M. (2018). Ultrasonic modulation of neural circuit activity. *Current Opinion in Neurobiology*, 50, 222–231.
- Tyler, W. J., Tufail, Y., Finsterwald, M., Tauchmann, M. L., Olson, E. J., & Majestic, C. (2008). Remote Excitation of Neuronal Circuits Using Low-Intensity, Low-Frequency Ultrasound. *PLoS ONE*, 3(10), e3511.

- Uğurbil, K., Xu, J., Auerbach, E. J., Moeller, S., Vu, A. T., Duarte-Carvajalino, J. M., Lenglet, C., Wu, X., Schmitter, S., Van de Moortele, P.F., Strupp, J., Sapiro, G., De Martino, F., Wang, D., Harel, N., Garwood, M., Chen, L., Feinberg, D.A., Smith, S.M., Miller, K.L., Sotiropoulos, S.N., Jbabdi, S., Andersson, J.L.R., Behrens, T.E.J., Glasser, M.F., Van Essen, D.C., Yacoub, E., WU-Minn HCP Consortium. (2013). Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *NeuroImage*, 80, 80–104.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium (2013). The WU-Minn Human Connectome Project: an overview. *Neuroimage* 80, 62–79.
- Van Essen, David C. (2002). Surface-based atlases of cerebellar cortex in the human, macaque, and mouse. *Annals of the New York Academy of Sciences*, 978, 468–479.
- Van Essen, David C., & Dierker, D. L. (2007). Surface-Based and Probabilistic Atlases of Primate Cerebral Cortex. *Neuron*, 56(2), 209–225.
- Van Essen, David C., Glasser, M. F., Dierker, D. L., & Harwell, J. (2012). Cortical parcellations of the macaque monkey analyzed on surface-based atlases. *Cerebral Cortex (New York, N.Y.: 1991)*, 22(10), 2227–2240.
- Van Hoesen, G. W., Morecraft, R. J., & Vogt, B. A. (1993). Connections of the Monkey Cingulate Cortex. In *Neurobiology of Cingulate Cortex and Limbic Thalamus* (pp. 249–284).
- Verhagen, L., Gallea, C., Folloni, D., Constans, C., Jensen, D. E., Ahnine, H., Roumazeilles, L., Santin, M., Ahmed, B., Lehericy, S., Klein-Flugge, M., Krug, K., Mars, R.B., Rushworth, M.F., Pouget, P., Aubry, J.-F., Sallet, J. (2019). Offline impact of transcranial focused ultrasound on cortical activation in primates. *ELife*, 8.
- Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, A. Z., Baker, J. T., Van Essen, D. C., Zempel, J.M., Snyder, L.H., Corbetta, M., Raichle, M. E. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, 447(7140), 83–86.
- Volman, I., Roelofs, K., Koch, S., Verhagen, L., & Toni, I. (2011). Anterior Prefrontal Cortex Inhibition Impairs Control over Social Emotional Actions. *Current Biology*, 21(20), 1766–1770.

- Volman, I., Verhagen, L., den Ouden, H. E. M., Fernández, G., Rijpkema, M., Franke, B., Toni, I., Roelofs, K. (2013). Reduced serotonin transporter availability decreases prefrontal control of the amygdala. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(21), 8974–8979.
- Walker, A. E. (1940). A cytoarchitectural study of the prefrontal area of the macaque monkey. *Journal of Comparative Neurology*, 73(1), 59–86.
- Wallis, J. D. (2007). Orbitofrontal cortex and its contribution to decision-making. *Annual Review of Neuroscience*, 30, 31–56.
- Wallis, J. D. (2011). Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nature Neuroscience*, 15(1), 13–19.
- Walton, M. E., Behrens, T. E. J., Buckley, M. J., Rudebeck, P. H., & Rushworth, M. F. S. (2010). Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron*, 65(6), 927–939.
- Wang, S., Olumolade, O. O., Sun, T., Samiotaki, G., & Konofagou, E. E. (2015). Noninvasive, neuron-specific gene therapy can be facilitated by focused ultrasound and recombinant adeno-associated virus. *Gene Therapy*, 22(1), 104–110.
- Wang, X. J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 19(21), 9587–9603.
- Wang, X.-J. (2002). Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. *Neuron*, 36(5), 955–968. Wang, X.-J. (2008). Decision Making in Recurrent Neuronal Circuits. *Neuron*, 60(2), 215–234.
- Wattiez, N., Constans, C., Deffieux, T., Daye, P.M., Tanter, M., Aubry, J.-F., Pouget, P., (2017). Transcranial ultrasonic stimulation modulates single-neuron discharge in macaques performing an antisaccade task. *Brain Stimulation* 10, 1024–1031.
- Webb, T. J., Rolls, E. T., Deco, G., & Feng, J. (2011). Noise in Attractor Networks in the Brain Produced by Graded Firing Rate Representations. *PLoS ONE*, 6(9).

- Weiskrantz, L. (1956). Behavioral changes associated with ablation of the amygdaloid complex in monkeys. *Journal of Comparative and Physiological Psychology*, 49(4), 381–391.
- Wellman, L. L., Gale, K., & Malkova, L. (2005). GABAA-mediated inhibition of basolateral amygdala blocks reward devaluation in macaques. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 25(18), 4577–4586.
- Wernicke, C., Emminghaus, H., & King's College London. (1874). *Der aphasische Symptomencomplex: Eine psychologische Studie auf anatomischer Basis*.
- Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked Presentations of Emotional Facial Expressions Modulate Amygdala Activity without Explicit Knowledge. *Journal of Neuroscience*, 18(1), 411–418.
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal Cortex as a Cognitive Map of Task Space. *Neuron*, 81(2), 267–279.
- Wise, S. P., & Murray, E. A. (2000). Arbitrary associations between antecedents and actions. *Trends in Neurosciences*, 23(6), 271–276.
- Wise, Steven P. (2008). Forward frontal fields: Phylogeny and fundamental function. *Trends in Neurosciences*, 31(12), 599–608.
- Wittmann, M. K., Kolling, N., Akaishi, R., Chau, B. K. H., Brown, J. W., Nelissen, N., & Rushworth, M. F. S. (2016). Predictive decision making driven by multiple time-linked reward representations in the anterior cingulate cortex. *Nature Communications*, 7, 12327.
- Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 26(4), 1314–1328.
- Yeterian, E. H., Pandya, D. N., Tomaiuolo, F., & Petrides, M. (2012). The cortical connectivity of the prefrontal cortex in the monkey brain. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 48(1), 58–81.
- Yizhar, O., Fenno, L.E., Davidson, T.J., Mogri, M., Deisseroth, K., (2011). Optogenetics in Neural Systems. *Neuron* 71, 9–34.

- Yoo, S.-S., Bystriksy, A., Lee, J.-H., Zhang, Y., Fischer, K., Min, B.-K., McDannold, N.J., Pascual-Leone, A., Jolesz, F. A. (2011). Focused ultrasound modulates region-specific brain activity. *NeuroImage*, 56(3), 1267–1275.
- Yukie, M. (2002). Connections between the amygdala and auditory cortical areas in the macaque monkey. *Neuroscience Research*, 42(3), 219–229.
- Zangemeister, L., Grabenhorst, F., & Schultz, W. (2016). Neural Basis for Economic Saving Strategies in Human Amygdala-Prefrontal Reward Circuits. *Current Biology: CB*, 26(22), 3004–3013.