Effort-based Learning and Decision Making

Irma Triasih Kurniawan

Cognitive, Perceptual and Brain Sciences &
Wellcome Trust Centre for Neuroimaging
University College London

Supervisors:
Ray Dolan
Nick Chater

Submitted for the consideration of a PhD in Cognitive Neuroscience

August 2011

"So she was considering in her own mind […], whether the pleasure of making a daisy-chain would be worth the trouble of getting up and picking the daisies..."

Alice in "Alice's Adventures in Wonderland",

Lewis Caroll, 1865, p.11

Declaration

I, Irma T. Kurniawan, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signature: 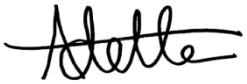

Date: 12 August 2011

Abstract

In the literature on healthy humans, effort is poorly studied and an extension from an animal literature is just emerging. I tested an hypothesis that physical effort is a non-trivial aspect of motivated behaviour; it serves as a cost and interacts with outcomes. To do this I conducted four experimental studies and extended the range of costs to include pain. In my first experiment, I develop a functional Magnetic Resonance Imaging (fMRI) task assessing willingness to expend effort. I show that physical effort discounts value of actions and that activity in dorsal striatum is associated with effort of selected actions. In addition to influencing choice, effort may be influenced by affective outcomes. In my second experiment, I develop a behavioural instrumental learning task examining how reward and punishment influence learning about effortful response. I show that it is easier to expend effort to gain reward and to withdraw effort to avoid punishment, but not the other way around; in other words it is more difficult to expend effort to

avoid punishment and to withdraw effort to gain reward. Results from reinforcement learning modelling account for this tendency in terms of a pavlovian influence on effort. On the one hand, outcome has an influence in effort while, on the other, effort may modulate neural signalling of action anticipation and outcome delivery. In my third experiment, I develop an fMRI cue-predictive instrumental task investigating brain responses for effort anticipation and outcome evaluation. I show that activity in anterior cingulate cortex and dorsal striatum is sensitive to anticipated effort and highlight an effort modulation on activity in ventromedial prefrontal cortex and ventral striatum associated with expected outcomes. Finally, I extend my investigation of costly behaviour in effort to pain by showing an influence of context effects in pain avoidance behaviour.

In summary, within this thesis I demostrate that physical effort as a cost is non-trivial in that it i) discounts value, ii) is sensitive to pavlovian influences, iii) is neurally anticipated and iv) modulates outcome signalling. I show the viability of various experimental paradigms to assess costly behaviours driven by effort and extend this endeavour by studying cost-driven pain avoidance. These experiments forge new research directions for understanding action and decision making as well as show promise for testing aberrant populations that often present with pathology that may reflect under- and over-motivated actions (e.g., apathy and perseveration).

Acknowledgments

My eternal gratitude to my PhD supervisors, Nick Chater and Ray Dolan for the patience, generosity and trusting guidance, and in offering intellectual, psychological and emotional support. I thank Peter Dayan for the kind and welcoming direction for my last two projects. I am grateful for the constructive and friendly collaborations with Marc Guitart-Masip, Debbie Talmi, Ivo Vlaev, Julia Trommershäuser and especially for the constant support from Wako Yoshida and Ben Seymour. I would like to give a special thanks to members of the emotion club and Chater + Lagnado lab who continuously offer excellent intellectual stimulation and patient feedback during my PhD. Thank you to all who gave technical and imaging support at the FIL and CPB and to all my participants.

Thank you to my trusting family the Kurniawans and all NH and Finsy housemates for the loving encouragements and sanity maintenance. A special thanks for 'holding my hands' during these four years to Erica, Matthias, Natassha, David B., and Doran. My time in London would not have been as rich without you.

Finally to JC, I am grateful for the unending love and strength. *Totus tuus.*

Sincerely yours,

Irma Triasih Kurniawan

Collaborations

Chapter 1: Parts of the text are reprinted from a review article as a joint collaboration with Marc Guitart-Masip and Ray Dolan, published in Frontiers in Decision Neuroscience in May 2011.
Chapter 3: Parts of the text are reprinted from a report as a joint collaboration with Ben Seymour, Deborah Talmi, Wako Yoshida, Nick Chater and Ray Dolan, published in Journal of Neurophysiology in 2010.
Chapter 4 and 5: This is a joint collaboration with Marc Guitart-Masip, Peter Dayan, and Ray Dolan, presented as a poster at UCL Neuroscience Symposium 2011.
Chapter 6: Parts of the text are reprinted from a report as a joint collaboration with Ben Seymour, Ivo Vlaev, Julia Trommershäuser, Nick Chater and Ray Dolan, published in Psychological Science in 2010.
Chapter 7: Small parts of the text are reprinted from the review article in Frontiers in Decision Neuroscience and the report in Journal of Neurophysiology.

Publications arising directly from this submission

Kurniawan, I.T., Guitart-Masip, M., & Dolan, R.J. (**2011**). Dopamine and effort-based decision making. *Front. Neurosci.* 5:81. doi: 10.3389/fnins.2011.00081 (Review article in Special issue 'Neurobiology of Choice').

Kurniawan, I.T., Seymour, B., Talmi, D., Yoshida, W., Chater, N., & Dolan, R. J. (**2010**). Choosing to make an effort: the role of striatum in signaling physical effort of a chosen action. *J Neurophysiol 104: 313 – 321.* Originally published online 12 May 2010. DOI:10.1152/jn.00027.2010.

Kurniawan, I.T.*, Seymour, B.*, Vlaev, I., Trommershäuser, J., Dolan, R. J., & Chater, N. (**2010**). Pain relativity in motor control. *Psychological Science 21: 840 – 847.* Originally published online 30 April 2010. DOI: 10.1177/0956797610370160.

*These authors contributed equally to this work

Table of Contents

List of Figures

List of Tables

Chapter 1 Introduction

Understanding healthy cognition, which affords motivated behaviour and the neuropathologies, is a central interest of cognitive neuroscience. A crucial component of motivated behaviour in the contexts of learning and decision making, that is often overlooked, is action cost. Indeed, humans consider the potential costs of an action, as well as their possible rewards, in order to select the best action. Inherent effort costs may also be influenced by outcomes and in turn, may have influence on outcome evaluation. Human basal ganglia and prefrontal cortex are thought to be the crucial substrates in such effort processing.

This thesis aims to dissect the role of physical effort in biasing healthy individuals' choices away from actions which require greater effort, how outcomes influence learning about effortful responses, how basal ganglia-prefrontal circuitry is sensitive to effort anticipation, effort choice and outcomes of effortful responses. The research reported in this thesis provides several new contributions to the field including:1) development of viable experimental paradigms for effort manipulations and measurements, 2) provision of converging evidence for previous non-human animal and clinical work on effort in healthy humans, 3) utilisation of reinforcement learning principles which capture a pavlovian influence on effort, 4) examination of effort processing in affective contexts provided by rewards and punishments and 5) extension of the range of paradigms to another characteristic cost-driven behaviour; that involving pain evasion.

The underlying assumption with which I commenced this work was that effort is costly. With this, one may envision operationalising 'costly' as a subjective hindrance, such that an action is called to be 'effortful' only if it stops being chosen due to its excessive perceived 'cost'. However, this intuition imposes a great limitation in studying effort as it results in a binary output of whether an action is always chosen (not effortful) or never chosen (effortful). An alternative is to operationalise 'costly' as an objective parameter set by the experimenter (such as fixed squeezing force levels). Although the most effortful experimental parameter may never mimic subjective maximal effort, this intuition is experimentally parsimonious as it allows testable hypotheses, for example that different effort or

force levels have differental effects on choice or learning. It is highly likely that the objective manipulation and subjective perception of 'cost' are monotonically related, such that more objective effort, in most cases, would invoke higher subjective effort. I use the latter intuition as the framework in my experiments by using a handgrip to manipulate and measure effort and show that actions with different force requirements have different effects on behavioural and neural measures.

There are two advantageous features in the paradigms developed in this thesis. First, unlike previous work in healthy humans (except Prevost, Pessiglione, Metereau, Clery-Melin, & Dreher, 2010), I employed physical effort, rather than mental effort. Studying how effort is conceptualised in other fields, namely clinical neurology and behaviour ecology, allowed me to determine the likely critical variables in relation to how effort influences behaviour. Auto-activation deficit (AAD), the most severe form of apathy, quantitatively reduces the initiation and execution of actions and this contrasts with a 'cognitive inertia' observed in less severe forms of apathy (Lévy & Dubois, 2006). A foraging literature in animals is concerned with the computation of physical effort costs such as metabolic rates (e.g., Marsh, Schuck-paim, & Kacelnik, 2004) in determining choice of foraging methods (e.g., walking or flying). These observations provided a principled motivation for manipulating physical rather than mental effort.

Secondly, in daily life expending more effort often requires more time. Indeed temporal discounting can be confounded by effort discounting, since effortful actions invariably involve greater time investment (but see Floresco, Tse, & Ghods-Sharifi, 2008 for effort discounting in rodents after controlling for time effects). While it is often experimentally difficult to disentangle the two, I was able to examine effort costs whilst controlling for time effects by equating the grip duration in high, low, and no effort conditions.

1.1 Thesis Summary

In study 1, I seek to extend animal research on effort-based decision making in healthy humans. Here, I validate an effort discounting paradigm where participants choose between effortful gripping which entail varying effort and reward levels and a no effort option with minimum reward. Using fMRI, I show that action choice is influenced by amount of gripping and that activity in dorsal striatum at the time of choice is associated with how much effort the selected action requires. Effort in the context of choice is also associated with a persistence trait which refers to personality tendencies to meet daily challenges.

Having established viability of this effort manipulation in a relatively established choice context, I then explore effort in instrumental learning and manipulate affective context by including appetitive and aversive outcomes. In study 2, I explored a pavlovian impact upon effortful actions in the context of different affective outcomes i.e. rewards and punishments. I extend the orthogonalisation between action and outcome valence (Boureau & Dayan, 2011) by using hand grip actions which reflect either a behavioural activation (squeeze) or withdrawal (release), to either gain reward or avoid punishment. Using this instrumental learning paradigm I demonstrate that it is easier to squeeze to gain reward and to release to avoid punishment than to squeeze to avoid punishment and to release to gain reward. My data is best captured by a reinforcement learning model which characterises this differential action-outcome association as a pavlovian influence of reward and punishments on effort response.

This new paradigm offers a way to assess the role of effort in neural signalling of action anticipation and outcome delivery. In study 3, I investigated modulation of effort during action anticipation and outcome delivery in both reward and punishment contexts. In this fMRI task, participants completed an overlearnt cue-predictive task by squeezing at either low or high effort levels to either win reward or avoid punishment. Supporting an animal literature, I show that activity in the anterior cingulate and dorsal striatum attunes to the level of effort needed for an upcoming action. Moreover, effort that has just been expended modulates activity

during outcome delivery in the ventromedial prefrontal cortex and ventral striatum.

The findings from these studies prove the feasibility of studying physical effort free of temporal contamination. I show support for a natural extension of previous non-human and human findings on neural correlates for effort choice, pavlovian effects on effort learning and neural correlates for effort-outcome interplay. In study 5, I highlight the importance of extending work on effort to pain avoidance. I explore this connection by assessing how magnitude manipulations, such as pain context, influence how our motor system avoids pain.

In summary, my data provide a basic experimental and neuroanatomical framework for human effort-based learning and decision making and an extension to a broader category of costs. I demonstrate the experimental validity of effort-related behaviour, extend previous knowledge about effort and actions in healthy humans, and discover unprecedented potential of effort (instrumental) learning in punishment avoidance. I discuss the implications and contributions of this doctoral work for the field of cognitive and decision neuroscience.

1.2 Literature review

Effort is commonly experienced as a burden, and yet we readily expend effort to reach a desired goal. Many classical and contemporary studies have assessed the effect of effort expenditure on response rates, by varying experimental parameters such as the weight of a lever press, the height of a barrier to scale, or the number of handle turns needed to generate a unit of reward (Collier & Levitsky, 1968; Collier, Hirsch, Levitsky, & Leshner, 1975; Kanarek & Collier, 1973; Kool, Mcguire, Rosen, & Botvinick, 2010; Lewis, 1964; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). There is general agreement that animals, including humans, are disposed to avoid effortful actions. It is paradoxical then that effort is not always treated as a nuisance, and there are instances where its expenditure enhances outcome value as observed in food palatability (Johnson & Gallagher, 2010), likeability (Aronson, 1961) and indeed the propensity to choose a previously effortful option (Friedrich & Zentall, 2004). What is most surprising is the

observation that effort often biases future choice towards effortful actions (Eisenberger, Weier, Masterson, & Theis, 1989).

Laboratory results show that if reward magnitude is held constant then high effort tasks tend to be avoided (Kool et al., 2010). Yet, in daily life most organisms seem superficially indifferent to the varying costs of action and readily choose challenging tasks to achieve a desired goal (Duckworth et al., 2007). Such observations point to the presence of a mechanism that integrates effort costs with benefits in order to implement desired actions (see Floresco, St Onge, Ghods-Sharifi, & Winstanley, 2008 for review on various cost-benefit analyses and Salamone, Correa, Farrar, & Morris, 2007 for an earlier review on dopamine and effort). This perspective has been addressed by optimal foraging theory. It is known that animals will strive to maximise gain whilst minimising energy expenditure (Bautista, Tinbergen, & Kacelnik, 2001; J. R. Stevens, Rosati, Ross, & Hauser, 2005). Thus, ducks choose between walking or flying depending on optimal solution of net gain between energy requirements in walking or flying and the food gained (Bautista et al., 2001).

In what follows I discuss a literature that has endeavored to understand the neural mechanisms of effort and reward integration, including the involvement of dopamine (DA) in effort-based behaviour. This literature points to the basal ganglia (BG), particularly dorsal and ventral striatum, and anterior cingulate cortex (ACC) as the principal substrates in both representing and integrating effort and action implementation.

The regulatory role of dopamine in effort

Over the past three decades, theories concerning the role of midbrain DA on behaviour have changed dramatically. The hedonic hypothesis of DA (Wise, 1980) is now challenged by empirical evidence revealing that global DA depletion (including within the accumbens, a major recipient for DA) does not impair hedonic responses to primary rewards ('liking', TABLE 1-1 for terms) such as orofacial reactions, the preference for sucrose over water, or discrimination among reinforcement (Berridge, Venier, & T. E. Robinson, 1989; Cannon & Palmiter, 2003;

Cousins & Salamone, 1994). On the other hand the same lesions profoundly impair performance of instrumental tasks necessary to obtain rewards that are liked (Berridge & T. E. Robinson, 1998). These observations have led to a formulation that the contribution of DA includes an effect on motivated behaviours towards desired goals, a concept referred to as 'wanting' (Berridge & T. E. Robinson, 1998). 'Wanting' can be expressed in simple instrumental responses, such as button or lever presses or in a more expanded form of behaviours which require an agent to overcome action costs (TABLE 1-1). As demonstrated unequivocally by Salamone and colleagues (Salamone & Correa, 2002; Salamone et al., 2007; Salamone, Correa, Mingote, & Weber, 2003), accumbens DA depletion disrupts instrumental responding if the responses require an energetic cost such as climbing a barrier (Salamone, Cousins, & Bucher, 1994), but leaves reward preference intact when effort is minimal. This has led to an hypothesis that DA plays a role in overcoming "costs" (Phillips, Walton, & Jhou, 2007; Salamone & Correa, 2002).

Alternative views on the role of dopamine in decision-making

There are several alternative views to DA which I summarise in Figure 1-1. Aside from a role in the expression of motivated behaviour, DA is also involved in its acquisition through learning. An influential view on how DA influences behaviour comes from reinforcement learning (Sutton & Barto, 1998). Reinforcement learning offers ways to formalise the process of reward maximisation through learned choices and has a close resonance with the neuroscience of decision-making (Daw & Doya, 2006; Montague & Berns, 2002; Montague, Dayan, & Sejnowski, 1996; Niv, Daw, & Dayan, 2005). In particular, phasic responses of macaque and rodent midbrain dopaminergic neurons to rewards, and reward-associated stimuli, are akin to a reward prediction error signal within reinforcement learning algorithms, responding to unexpected rewards and stimuli that predict rewards but not to fully predicted rewards (Bayer & Glimcher, 2005; Morris, Nevet, Arkadir, Vaadia, & Bergman, 2006; Roesch, Calu, & Schoenbaum, 2007; W. Schultz, Dayan, & Montague, 1997). Moreover, fMRI studies report that the BOLD signal in the striatum, a major target of the dopaminergic system, correlates with the prediction error signals derived from fitting subject's behaviour to a reinforcement learning model (McClure, Berns, & Montague, 2003; O'Doherty, Dayan, Friston, Critchley,

& Dolan, 2003; O'Doherty et al., 2004). In support of such a role for DA in reinforcement learning processes, stimulation of the substantia nigra (using intracranial self-stimulation paradigm) has been shown to induce a potentiation within corticostriatal synapses at the site where nigral output cells terminate, with these effects in turn being blocked by systemic administration of a DA D1/D5 antagonist (J. N. Reynolds, Hyland, & Wickens, 2001). Importantly, the magnitude of potentiation is negatively correlated with the time taken by an animal to learn the self-stimulation paradigm.

DA is also proposed to signal stimulus salience, as opposed to reward prediction error (Redgrave, Prescott, & Gurney, 1999). Redgrave and co-authors have discussed the stereotypical latency and duration of phasic bursts of nigral dopaminergic neurons, as well as the connectivity between nigral dopaminergic neurons and sensory subcortical structures such as the superior colliculus. They argue that activity of DA neurons can be interpreted as reporting biological salient events, either due to novelty or unpredictability. From this perspective, salient events generate short-latency bursts of dopaminergic activity that reinforce actions occurring immediately preceding the unpredictable event. This signal allows an agent to learn that an action caused the salient event (see Redgrave & Gurney, 2006 for an elegant discussion on signal transmission in tecto-nigral and cortico-subcortical pathways for learning of action–outcome associations). According to this view, unpredictable rewarding events are just one among many exemplars of a salient event.

Table 1-1 Useful key terms.

Effort	Strenuous physical or mental exertion typically with the aim of achieving a desired outcome or goal.
Liking	A set of behaviours driven by hedonic or pleasurable properties of a stimulus, such as the smell or taste of a valued food item. Typical liking responses in rodents include orofacial reactions while in humans likeability is operationalised through degrees of attractiveness measured on a Likert scale. A characteristic of likeability is that it needs not be motivational nor sensitive to devaluation procedures.
Wanting	A set of behaviours driven by salient properties of a stimulus often manifests in a disposition to overcome costs in order to obtain an incentive. Wanting often entails actions such as lever pressing in rodents or non-human primates

	to obtain a goal object. One influential hypothesis regarding DA function highlights a role in mediating wanting, but not liking.
Apathy	A mental or behavioural state devoid of motivation with a core feature of lack of self-initiated actions.
Cost-benefit integration	The process of deriving a value of an action based on a combination between potential utility in attaining and disutility incurred in so doing. There is evidence that this type of integration takes place when one is judging whether an action is worth taking, although the mechanisms by which costs and benefits are integrated remain unclear.
Invigoration	To vitalise or increase strength. One hypothesis regarding the role of DA formalises its role as facilitating motivated behaviour by invigorating an organism when faced with increasing demands of effort. This is supported by studies that highlight the effects of a dopaminergic manipulation on effort expenditure.

Finally, Nicola recently suggested that DA is required to flexibly initiate goal-directed instrumental responses (Nicola, 2010). This view is based on observations that the effects of DA depletion in the rat nucleus accumbens (NAc) are dependent on inter-trial interval, such that instrumental responses with short inter-trial intervals are not affected by DA depletion but depletion effects increase as a function of increasing time between responses. Detailed behavioural analysis shows these effects of time are explained by the fact that as the duration between responses increases animals tend to engage with behaviours different from the required instrumental response, with depleted animals unable to flexibly reinitiate execution of the instrumental responses. On the other hand, depleted rats can perform complex sequences of behaviour in situations where these are not interrupted. Such findings suggest that rather than impairing lever presses, DA depletion disrupts an animal's ability to flexibly re-engage with a task after engaging in a task irrelevant behaviour.

Extending reinforcement learning to account for dopamine involvement in effort

The most compelling attempt to link the known role of DA in reward learning to effort is that of Niv and colleagues (Niv, Daw, Joel, & Dayan, 2007) who have developed a model that specifies the vigour (defined as the inverse latency) of action. This model realises a trade-off between two costs: one stemming from the

harder work assumed necessary to emit faster actions and the other from the opportunity cost inherent in acting more slowly. The latter arises out of the ensuing delay to the next, and indeed to all subsequent, rewards. Niv et al. (2007) suggested that agents should choose latencies (and actions) to maximise the rate of accumulated reward per unit time, and showed that the resulting optimal latencies would be inversely proportional to the average reward rate. Based on a review of experimental evidence, Niv et al. (2007) proposed that tonic levels of DA report the average rate of reward, thus tying together prediction error (McClure et al., 2003; Montague et al., 1996; W. Schultz et al., 1997), incentive salience (Berridge & T. E. Robinson, 1998) and invigoration (Salamone & Correa, 2002) theories of DA. As defined by Niv et al. (2007), vigour can be thought of as a specific manifestation of effort expenditure in the time domain. Future work might usefully extend this temporal computational concept of vigour into other aspects of physical effort.

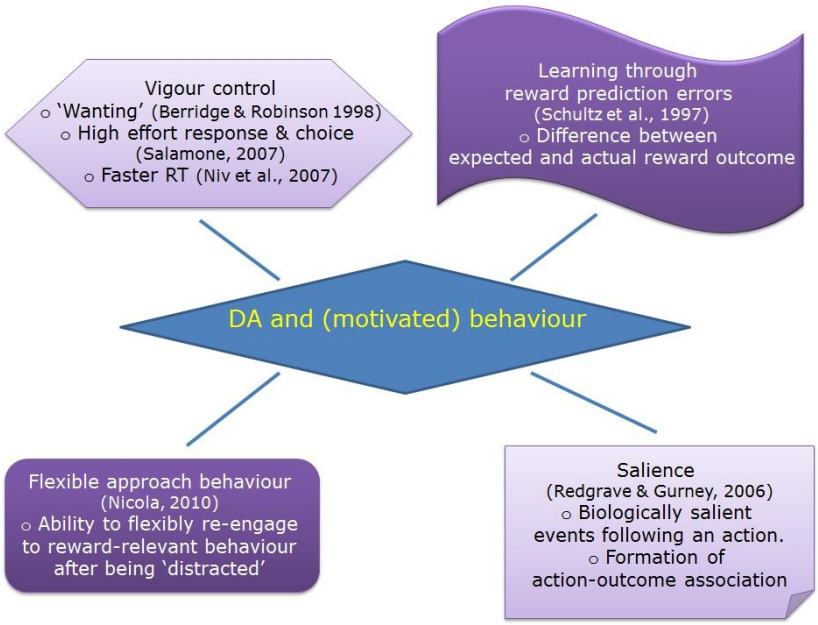

Figure 1-1 The figure illustrates a range of views regarding the role of DA in facilitating motivated behaviour. Clockwise from top left corner: A wealth of evidence shows that DA acts to invigorate an agent's effortful action, integrating ideas about overcoming effort costs, agents' choice for high effort options as well as modelling work on vigour. Another influential view pertains to DA acting as a signal for a prediction in reward as exemplified by its role in reinforcement learning. An alternative view interprets this signal as a saliency signal which allows agents to implement associative learning. Finally, DA may facilitate flexible switching and re-engagement in relation to reward-driven behaviour. The summarised perspectives are not mutually exclusive, nor do they represent the entire literature on DA, but are useful in understanding what support motivated behaviour.

Dopamine and its role in overcoming effort costs

Considerable evidence points to midbrain DA depletion discouraging animals from choosing effortful actions (Aberman & Salamone, 1999; Cousins & Salamone, 1994; Denk et al., 2005; Phillips et al., 2007). A series of experiments in rats has pointed to the crucial role of cortico-subcortical networks for cost-benefit decision making as highlighted in depletion effects (Aberman & Salamone, 1999; Cousins & Salamone, 1994). In these experiments, rats are trained on a T-maze that requires choosing between two actions; one yields high reward (4 pellets of food) but requires higher effort (climb a 30-cm barrier or higher lever press fixed-ratio schedule), the other yields low reward (2 pellets of food) but requires less effort. DA depletion in the NAc changes a rat's preference away from the high effort/high reward option, but does not impact on reward preference when it is readily available, nor does it alter response selection based on reward alone (Cousins & Salamone, 1994; Salamone et al., 1994). This finding has been replicated in other laboratories with a variety of depletion methods (Denk et al., 2005; Floresco, et al., 2008), where some studies point to a stronger effect from depletion in the core as opposed to shell of the NAc (Ghods-Sharifi & Floresco, 2010; Nicola, 2010).

The impact of DA elevation on effort is much less conclusive. Enhancing DA function is commonly realised through injection of amphetamine, an indirect DA agonist that increases synaptic DA levels (but also that of other neuromodulators). Floresco and colleagues (Floresco, Tse, et al., 2008) revealed a dose-dependent effect of amphetamine such that low-doses of amphetamine increased effortful choice, but high dose decreased it. This dose dependent effect is difficult to interpret. First of all, it is unclear what the precise effect of a high dose of amphetamine is on DA concentration level since amphetamine also results in increased extracellular serotonin (5HT) and noradrenaline (Salomon, Lanteri, Glowinski, & Tassin, 2006). Moreover, it is unclear whether a low dose of amphetamine acts by increasing the value of the reward, decreasing the cost of an action, modifying the integration of both, or by affecting other components of behavioural control such as impulsivity (see Pine, Shiner, Seymour, & Dolan, 2010 in relation to the latter). Nevertheless, the data suggest that increasing DA levels

per se does not invariably enhance preference for a high reward/high effort option, ruling out a simple monotonic relationship between DA and effort.

Another study showed an interactive effect of haloperidol, a DA receptor blocker, and amphetamine. While an injection of haloperidol 48 hours before treatment, followed by saline 10 minutes before test, significantly reduced preference for high reward/high effort arm, giving the same haloperidol injection followed by amphetamine 10 minutes before testing blunted the effect of haloperidol, and completely recovered preference for high reward/high effort arm (Bardgett, Depenbrock, Downs, Points, & Green, 2009). Evidence therefore points to amphetamine's ability to overcome the effects of DA blockade induced by haloperidol. However, as indicated amphetamine also increases the levels of 5HT and noradrenaline as well as DA levels making it difficult to completely outrule a possibility that the effect might relate to elevations of other amines aside from DA. We also know that amphetamine increases locomotor activity (Salomon et al., 2006) and it is impossible to dismiss the possibility that a recovered preference for the high-effort arm found might be due to enhanced locomotion.

Recent advances in neurochemical assay techniques, particularly in vivo fast scan cyclic voltammetry, allow detection of DA transients with a temporal resolution of milliseconds in awake behaving animals (D. L. Robinson, Venton, Heien, & Wightman, 2003; Roitman, Stuber, Phillips, Wightman, & Carelli, 2004). Gan and colleagues performed in vivo voltammetry while rats selected between two options in a task where there was an independent manipulation of the amount of reward and effort (Gan, Walton, & Phillips, 2010). These authors found that rats had the expected preference for higher magnitude of reward when costs were held constant and higher preference for options which require less effort when reward magnitude was constant. This study also included a separate set of trials which offered rats either option, while measuring the amount of DA released in the core of the NAc elicited by cues predicting reward and effort. By having this set of non-choice trials the authors ensured that the dopaminergic response was not confounded by the presentation of the second option. Whereas DA release reliably reflected the magnitude of the reward available in these trials, the amount of effort required to obtain the goal was not coded in the amount of DA released in the core of the NAc. This lack of evidence for an effort-dependent dopaminergic signal was surprising

given the extent of prior evidence (discussed above) pointing to a link between DA and the expenditure of effort in overcoming costs.

Overall, there is evidence that DA is required to overcome costs when high levels of effort are necessary to obtain a desired goal. However, the precise mechanism by which DA supports a cost-overcoming function, and how effort is integrated into a dopaminergic modulation of the striatum and prefrontal cortex, is much less clear. In addition, DA depleted animals can engage in high-effort responding given a limited, inflexible set of possible responses but exhibit difficulties and are slower in re-engaging with simple one-lever presses where multiple responses are allowed (Nicola, 2010). Whilst DA may be key to the computation and execution of highly effortful tasks, its role in strategic flexibility (Nicola, 2010) suggests it exerts a more subtle contribution to the complex relationship between task demands and the integration of task-relevant and task-irrelevant behaviour.

I next consider the likely contribution of BG and ACC, and the formation of action-outcome association necessary for motivated behaviour.

Basal ganglia: Anatomy and physiology

The basal ganglia are a set of subcortical nuclei comprising dorsal (putamen and caudate nucleus) and ventral aspects (often synonymous with NAc), the internal (GPi) and external (GPe) segments of globus pallidus, substantia nigra pars compacta (SNc) and reticulata (SNr) as well as the subthalamic nucleus (STN). The BG receives afferents from almost all cortical areas, especially the frontal lobe. Information processed within the BG network is sent via output nuclei (GPi and SNr) to the thalamus, which eventually feeds back to frontal cortex (Alexander & Crutcher, 1990; Bolam, Magill, & Bevan, 2002). This basic circuitry is reproduced in different parallel and integrative corticostriatal loops, with their origin in different frontal domains, and is held to play a critical role in cognitive functions that span motor generation to more cognitive aspects of causal learning, executive function and working memory (Frank, 2005; Frank, Loughry, & R. C. O'Reilly, 2001; Haber & Knutson, 2010; Vitay & Hamker, 2010).

Neurons in the striatum project either to output nuclei of the BG (GPi and SNr) or to an intermediate relay involving GPe neurons which ultimately project to BG output nuclei. These two populations provide the origin of BG direct and indirect pathways which funnel information, conveyed in parallel to striatum by cortical afferents, to BG output nuclei (Alexander & Crutcher, 1990; Frank, 2005; Frank & Fossella, 2011; Frank, Seeberger, & R. C. O'Reilly, 2004). Under basal conditions, the output nuclei of the BG have a high level of firing and maintain thalamic inhibition that serve to dampen activity in corticostriatal loops (Frank, 2005). The distinct connectivity of direct and indirect pathways (FIGURE 1-2) results in opposite effects: the direct pathway promotes inhibition of BG output nuclei and release of inhibition in thalamic activity whereas the indirect pathway promotes excitation of BG output structure and drives thalamic inhibition.

Anatomical and functional gradients in the striatum

The functional organisation of BG along the direct and indirect pathways, as described above, applies to the full extent of the striatum, forming an integral re-iterated processing matrix which performs common operations across different subdivisions (Wickens, Budd, Hyland, & Arbuthnott, 2007). Although there are suggestions of a dorsal-ventral segregation, the consensus favours a dorsolateral-ventromedial gradient (Voorn, Vanderschuren, Groenewegen, Robbins, & Pennartz, 2004) with no sharp anatomical distinction between dorsal-ventral areas. Indeed, based on the cytology of spiny projection neurons, dopaminergic inputs, and DA-modulated plasticity and inhibition, dorsal and ventral striatum are strikingly similar (Wickens et al., 2007). However, there is evidence for a functional segregation such that dorsolateral striatum, receiving sensorimotor afferents, supports habitual, stimulus-reward associations. This contrasts with ventromedial striatum, receiving afferents from orbito and medial prefrontal cortex, hippocampus and amygdala, which supports formation of stimulus-action-reward associations (Haber & Knutson, 2010; Voorn et al., 2004).

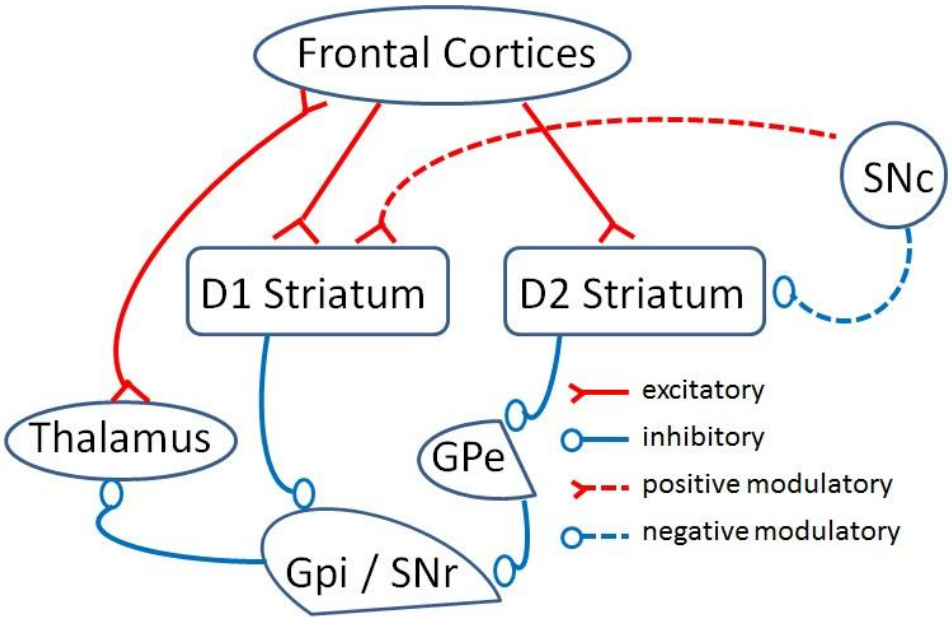

Figure 1-2. A schematic model of direct and indirect pathways of BG (adapted from Frank et al., 2004). The principal input of BG is the striatum, receiving excitatory inputs from most cortical areas. The output nuclei of BG are GPi/SNr, which direct processed information to the thalamus to eventually feed back an excitatory projection to the cortex. Within this circuitry, there are two pathways: a direct pathway expresses D1 receptors and indirect pathway expresses D2 receptors. D1 striatal neurons inhibit GPi/SNr cells forming the direct pathway. D2 striatal cells inhibit an intermediate relay, the GPe which ultimately provides inhibition to GPi/SNr. Under basal conditions, GPi/SNr cells fire at high level and maintain inhibition of the thalamus which in turn dampen corticostriatal loops activity. The different direct/indirect connectivity results in opposite effects: inhibitory effect on GPi/SNr and release of inhibition in thalamic activity by the direct pathway and excitatory effect on GPi/SNr and inhibitory effect on thalamus by the indirect pathway.

A functional gradient in DA signalling is also described in BG (Wickens et al., 2007). DA release is determined by density of DA innervation (densities reduce the distance between release and receptor sites), such that higher innervation densities are necessary for rapid DA signalling. DA clearance is regulated by density of DA transporters (DAT), hence affecting distance and time course of volume transmission. Wickens et al. (2007) have documented greater DA innervation and higher DAT densities in dorsolateral striatum with these densities decreasing along a ventromedial gradient (also Haber & Knutson, 2010). High densities of release sites and DAT result in fast clearance in dorsolateral striatum, which may be related to encoding of discrete events involving reinforced responding, or even automatised and habitualised behaviours. Ventromedially, lower densities of DA innervation and DAT result in slow clearance in NAc core, and even slower

clearance in NAc shell, which may be related to slower time course of action-outcome evaluation (Humphries & Prescott, 2010; Wickens et al., 2007).

Moreover, it is noteworthy that within the ventromedial subdivisions of the striatum, the NAc has interesting particularities. The NAc is subdivided, on the basis of anatomical and histochemical features, into the core and the shell, with the latter more medial and ventral in location than the former (Humphries & Prescott, 2010; Ikemoto, 2007; Voorn et al., 2004). This core/shell distinction is particularly important when considering the role of BG in motivated behaviour.

The NAc core is similar to dorsal striatum (Humphries & Prescott, 2010; but see Nicola, 2007 on role of dorsal-ventral striatum in temporal predictability). Functionally, NAc core seems critical in the translation of raw, unconditioned stimulus value, into a conditioned response. Thus, NAc core plays an important role in conditioned behaviour (Ikemoto, 2007), such as autoshaping in classical conditioning paradigms and conditioned reinforced responses in instrumental learning paradigms. On the other hand, the NAc shell, the most ventromedial aspect of striatum, has unique features compared to the rest of striatum. First, it is involved in unconditioned responding in the appetitive and aversive domains , spanning feeding (Kelley, Baldo, Pratt, & Will, 2005) and maternal behaviour (Li & Fleming, 2003) to defensive treading (S. M. Reynolds & Berridge, 2002). Moreover, the NAc shell is involved in invigorating effects of DA on conditioned behaviours controlled by the NAc core (Parkinson, Olmstead, Burns, Robbins, & Everitt, 1999). Second, the shell is the only striatal subdivision projecting to lateral hypothalamus (Pennartz, Groenewegen, & Lopes da Silva, 1994, and reviewed by Humphries & Prescott, 2010), a key structure in an 'action-arousal' network. Note lateral hypothalamus also exerts an influence over autonomic function and contains orexin-producing cells which influence arousal and energy balance control (see Ikemoto, 2007 for a comprehensive review). Third, whereas amygdala has extensive projections to both the core and shell (Humphries & Prescott, 2010), the NAc shell is the only recipient of hippocampal afferents within the striatal complex (Haber & Knutson, 2010; Wickens et al., 2007). This restricted projection from hippocampus has generated extensive discussion concerning the unique role of ventral BG in spatial navigation, fear-modulated free-feeding, and acquisition of stimulus value through stimulus-outcome pairings (Humphries & Prescott, 2010).

These lines of evidence point to the shell as critical in forming linkages between an object/event in the environment and the agent's natural response towards it.

An alternative interpretation of the anatomical and physiological organisation of the BG is a selection and control model (Gurney, Prescott, & Redgrave, 2001). In this model inputs for selection and control are received separately by striatal D1 receptors and D2-like receptors, respectively. D1 transmission is then projected as inhibition to GPi/SNr which acts as an action selection output, whereas D2 transmission inhibits GPe which acts as an output layer for a control mechanism. The control output layer, in turn modulates action selection: GPe inhibits activity in GPi/SNr output nuclei. Akin to inhibitory mechanisms described in the direct/indirect BG model, this selection/control BG model also describes inhibitory relationships between nuclei in BG. It is not clear what the thalamic inhibitory/excitatory impacts are on movement. Nevertheless this model highlights an important role for BG in action selection and control. More recently, Nicola (2007) has discussed the potential role of NAc in such a model, particularly in disinhibiting motor efferents for one action and inhibiting motor efferents for another, thereby allowing action selection.

Basal ganglia and effort-related processes

To facilitate execution of motivated behaviour, one needs to internally represent action costs and benefits. While the animal literature significantly informs our knowledge about brain structures subserving motivated behaviour, it is unknown how effort processing is supported in the human brain. Using fMRI, Croxson and colleagues investigated where in the human brain effort and reward are represented (Croxson, Walton, J. X. O'Reilly, Behrens, & Rushworth, 2009). Participants saw a discriminative stimulus signaling an action with a particular cost and benefit and then completed a series of finger movements using a computer mouse, to gain secondary reinforcers. The cost, in terms of effort and time, increased as more finger movements were completed, whilst the benefit increased as the secondary reinforcer was larger. When anticipating these actions, striatum activity correlated with both anticipated costs and anticipated reward of effortful actions.

More recent fMRI studies have replicated an involvement of striatum in effort-related processes, reporting higher dorsolateral striatal activity for choosing low compared to high effort options in a physical effort task (Kurniawan et al., 2010; chapter 3) and higher ventral striatal activity in a low cognitive demand block compared to a high cognitive demand block in a mental effort task (Botvinick, Huffstetler, & McGuire, 2009). Whilst, it is still unclear whether physical and cognitive mechanisms of effortful actions reflect similar psychological and neural processes, together these studies provide support for the importance of striatum in effort-related processes. In the following section, I assess the type of association formed when an organism performs a motivated, goal-directed, behaviour.

Encoding action and its outcomes

Linking a chosen action to its outcome is central for optimal goal-directed behaviour. When a monkey travels a distance to forage for food, not only does it need to link contextual cues to food consumption, for example associating a tree full of ripe fruits with eating fruits, it also needs to associate the action (climbing a tree) with the consequences of the action, namely the energetic cost of climbing. Neurons in primate dorsal striatum, can be categorised into those that encode the action made by the monkey (direction of saccade made) and neurons sensitive to the outcome of the monkey's choice (reward/ unrewarded) (B. Lau & Glimcher, 2007). However, these neurons do not appear to support the kind of action-outcome association required for goal-directed behaviour.

Using reinforcement learning models, similar to those used to characterise activity in DA neurons, Samejima and colleagues reported neurons in the striatum whose activity correlated with the value of an action (Samejima, Ueda, Doya, & Kimura, 2005). These action value neurons are important because they track the value of say, a left handle turn in a probabilistic two-choice task, independent of whether the monkey ultimately selects the action, and thus provide input information for action selection. Furthermore, in a subsequent study, Lau and Glimcher found action value neurons, including neurons which traced the value of the chosen action, in the striatum (B. Lau & Glimcher, 2008). These chosen value neurons

show enhanced activity when the tracked action has a higher value and, on this basis, was subsequently chosen. Using similar reinforcement learning models, human fMRI studies also report that BOLD signal in the dorsal striatum correlate with the relative advantage of taking one action over an alternative (O'Doherty et al., 2004).

Effortful action and its outcomes: implication of the ACC

These action and chosen value representations in the striatum are precisely the kind of association between action and outcome required for goal directed behaviour. However, the unanswered question is where does the information needed for this computation come from? One possibility is ACC, a region suggested to represent this action-outcome association (Rushworth, Behrens, Rudebeck, & Walton, 2007). For example, Hadland and co-workers trained macaque monkeys to pull a joystick upward after receiving a type of food, say a peanut, in order to obtain a second peanut and to turn a joystick to the side after obtaining a different food type, say a raisin, to receive a second raisin (Hadland, Rushworth, Gaffan, & Passingham, 2003). They found that while control monkeys could select an action based on this reward-response association, monkeys with a lesion to ACC were impaired in selecting the correct response. Interestingly, the impairment was not due to an inability to make an association between visual cues and reward as tested in a second visual discrimination task, but instead was specific to an inability to utilise reward-action association to make the correct response. In a different experiment, monkeys with ACC lesion were impaired in selecting a set of response when the correct responses were determined by an integration across past contingencies between action and reward (Kennerley, Walton, Behrens, Buckley, & Rushworth, 2006). In addition, using fMRI, human ACC was found to be most active when participants had to simultaneously internally generate a sequence of actions whilst monitoring the outcome of their actions (Walton, Devlin, & Rushworth, 2004).

Lesions studies with rodents using the T-maze consistently show impairments in effort-based decision making following removal of ACC. As with DA depletion experiments, these lesions result in a shift of preference away from an option with

a larger food reward that requires scaling a high barrier, thus requiring more effort. This reduced preference for larger/effortful arm was not due to lethargy or immobility as it is immediately restored when both arms have equal effort costs (Denk et al., 2005; Floresco & Ghods-sharifi, 2007; Walton et al., 2009 but see Floresco et al., 2008 for a discussion the extent to which ACC plays a role in effort-based tasks).

Human ACC lesions provide a more subtle interpretation for the role of ACC in effort processing. Naccache and colleagues tested a patient with a large lesion to left mesial frontal region including the left ACC using a, cognitively demanding, Stroop task (Stroop, 1935). This patient could not verbally recognise nor express discriminatory skin conductance responses in difficult trials where greater mental effort was required, but could perform as well as healthy controls. This case study suggests dissociability of objective cognitive performance from a physiological response and from the subjective appraisal of mental effortfulness (Naccache et al., 2005, but see McGuire & Botvinick, 2010 for the involvement of lateral prefrontal cortex, instead of ACC, in a closer inspection of subjective experience of mental effort through intentional and behavioural avoidance from mentally challenging tasks).

The ACC is implicated in a host of cognitive processes, ranging from cognitive control to suppression of prepotent responses such as in Stroop or go-nogo tasks, tasks that induce negative emotions, and tasks that predict delivery of painful stimuli. In a recent review (Shackman et al., 2011), the authors discussed a challenge in advancing knowledge of its functional organisation being the complexity of its anatomical organisation and variability across individuals. For example, a tertiary sulcus in dorsal ACC, the paracingulate sulcus, is present in one-third of the population, and its presence causes location change of architectonic Brodmann area 32', and a volumetric reduction of Brodmann areas 24a' and 24b'. Consequently, spatially normalised cingulate premotor regions differ across subjects, and an unmodeled cingulate sulcal variability may inflate the spread of activation clusters found across studies, rendering complex a clear functional dissociation within ACC.

Bush and co-authors proposed the rostro-ventral cingulate could be functionally segregated into cognitive and affective components located to dorsal and ventral ACC, respectively (Bush, Luu, & Posner, 2000). This segregation seems too broad. Shackman and co-authors (2011) using a sample of almost 200 neuroimaging experiments that included negative affect, pain and cognitive control reported strongly overlapping activation clusters in dorsal ACC, or what they termed as middle cingulate cortex (MCC), challenging a strict segregationist view of ACC (see FIGURE 1-3). These authors also pointed to evidence that the dorsal ACC might be involved in affective control, including autonomic regulation (Critchley et al., 2003) and pain processing, suggesting these findings may reflect an agent's need for behavioural control when habitual responses are not sufficient under uncertain action-outcome contingencies.

Anatomically, the ACC projects to striatum, particularly the caudate nucleus and portions of ventral striatum (Haber & Knutson, 2010). Moreover, ACC has bilateral connections to motor and prefrontal cortex fulfilling a role as a hub where action and outcome associations might be represented. In human and non-human primates, the ventral cingulate has strong interconnections with ventral striatum including the NAc, whilst the dorsal cingulate connects more strongly to dorsal striatum including putamen and caudate (Beckman, Johansen-Berg, & Rushworth, 2009; Kunishio & Haber, 1994), potentially facilitating transmission of reward-related information. Furthermore, dorsal ACC is interconnected with premotor cortex and a more posterior part constitutes the cingulate motor area (Beckman et al., 2009) implicated in action selection (Picard & Strick, 2001). Shima & Tanji (1998) reported that cingulate motor areas in monkeys respond to selection of voluntary movement based on reward, supporting a role in linking internally-generated action to reward. Indeed, a working hypothesis is that ACC could support adaptive control, integrating aversive, biologically relevant information in order to bias motor regions towards a contextually appropriate action (Shackman et al., 2011).

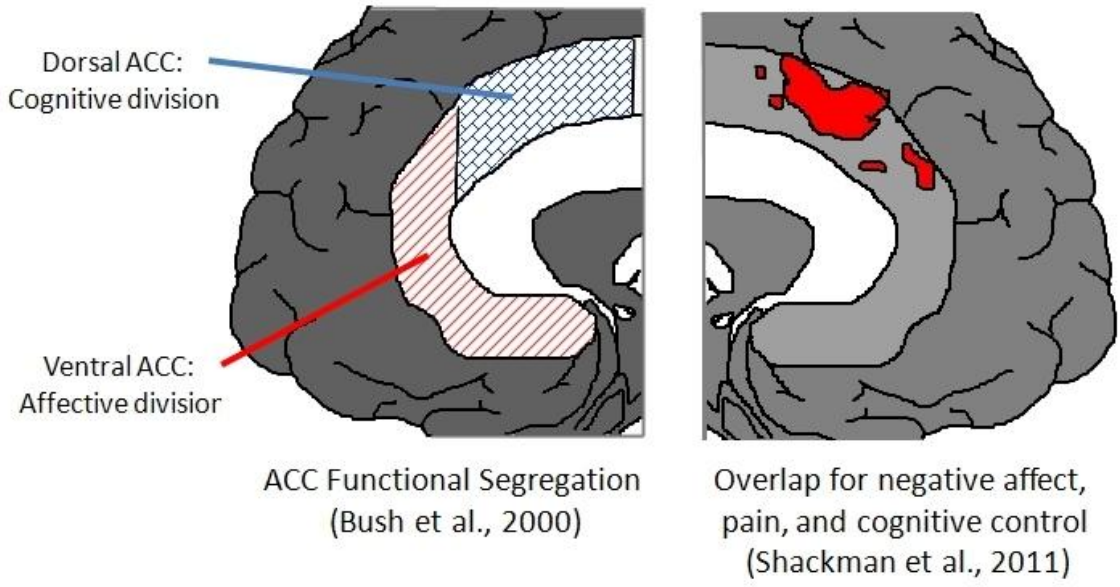

Figure 1-3. Views on the psychological function of ACC. Left: ACC function has been suggested as anatomically segregated into a dorsal cognitive division and a ventral affective division (Bush et al., 2000). Right: More than a decade later, a meta-analysis on almost 200 fMRI experiments suggested a strong overlap in clusters of activation in studies of cognitive control, negative affect, and pain (Shackman et al., 2011). Figures adapted from Shackman, et al. (2011).

This wide-ranging anatomical connectivity between BG, ACC and other cortical regions provide a neuroanatomical foundation for establishing action and outcome representations, of a type needed for motivated behaviour. Normal function of this circuitry can be inferred to facilitate willingness to execute effortful actions. On the other hand, disruption of this circuitry, as in people with apathy (see TABLE 1-1), would discourage execution of such actions. This account has a resemblance to phenomena in a case study of a patient with a lesion to mesial prefrontal cortex (which included ACC) that led to profound apathy (Eslinger & Damasio, 1985). This patient was severely impaired in execution of real-life events such as holding a job, although various measures of logical reasoning, general knowledge, planning, and social and moral judgments proved intact. The authors discussed how the lesion did not impact on pure action execution, but on the analysis and integration of the costs and benefits pertaining to real-life situations.

Conclusion

I provide evidence for an intimate interplay between ACC, BG, and dopaminergic pathways in enabling animals, including humans, to choose and execute effortful action. I suggest that effort may act as a discounting factor for action value, and that integrative mechanisms between cost and benefit facilitate a willingness to incur costs. Our review of reinforcement learning, empirical findings on the relationship between dopaminergic coding and cost-benefit parameters of an action, and the organisation of BG and ACC point to these latter structures as critical in linking a stimulus to an action and the consequences of that action. Notably, patients with apathy often manifest a pathology that disrupts this ACC-BG network. This fractures a link between action and outcomes resulting in lack of drive to execute potentially valuable actions.

My review highlights the psychological and neural mechanisms through which an organism is willing and capable of executing an effortful act to attain a goal. The core process appears to involve coding of specific action requirements, an analysis and integration of costs and benefits, and a decision to expend effort and to implement an action. I do not dissect a potentially important distinction between cognitive and physical types of effort (Kool et al., 2010; Kurniawan et al., 2010; Prevost et al., 2010). Future research might usefully endeavor to examine how one makes a trade-off between both effort types and examine how we determine when investing in one type of effort (mental) is more appropriate than investing in the other (physical).

Chapter 2 General methods

Two studies in this thesis used functional imaging and thus I briefly describe and discuss the basic principles of fMRI and specific behavioural and fMRI methodology used in studies 1-4 on effort. To maintain completeness of study 5 on pain, methods for it are separately described in CHAPTER 6.

2.1 Basic principles of fMRI

Blood-oxygenation level dependent (BOLD) contrast

BOLD contrast is ubiquitously used in cognitive neuroscience as a proxy for neural activity associated with cognitive functions. However, BOLD signal does not measure brain activity directly and measure neuronal processes by assuming a tight relationship between brain regional perfusion and neuronal changes. How is this so? Neuronal activity due to information processing such as sensory, motor, or cognitive processes causes changes in metabolic demand, in this case oxygen consumption. A change in oxygen consumption is physiologically indicated by a change in de-oxygenated haemoglobin concentration which translates into a change in magnetisation properties. Deoxyhaemoglobin is paramagnetic and so changing the concentration of it will change the magnetic resonance (MR) signal picked up by a magnetic resonance (MR) scanner. The change in MR signal caused by altered deoxyhaemoglobin concentration is what is refered to as BOLD contrast.

Metabolic consumption and neuronal activity

We know that energy source in the brain mostly comes from adenosine triphospate (ATP). The ATP budget in the rodent brain is allocated more for restoring unequal distributions of ions caused by action potentials, and depolarisation and hyperpolarisation of cell membranes postsynaptically (respectively termed *excitatory* and *inhibitory* postsynaptic potentials; EPSP and IPSP) than for things like protein synthesis. In other words, events related to information processing tax brain metabolic energy much more than cellular housekeeping functions (Huettel, Song, & McCarthy, 2009). This supports the assumption that an MRI brain metabolic index (i.e. BOLD) tells us something about cognition; it also gives

support for the use of BOLD signal as a proxy for neural mechanisms underlying cognitive functions.

Strengths

Today fMRI is by far the most widely used cognitive neuroscientific method for whole-brain investigation in humans. With its progressive development, it can provide images with a sub-millimetre voxel resolution, giving it a much better spatial resolution compared to other methods e.g. electroencephalography, magnetoencephalography, and positron emission tomography (PET) (Logothetis, 2008). Its advantage over PET include its non-invasive nature and the absence of a need for radioactive tracers, a feature that endows it with an ability to test a much wider population than equivalent methods.

Weaknesses

As described, fMRI provides an indirect measure of neuronal activity which constrains interpretation of the relationship between BOLD increase and stimulus-related events.

Commonly, increase in BOLD signal in a region such as the striatum invokes the interpretation that this region is 'active' for that cognitive event. One caveat to this 'language' is that although an increase in BOLD signal could be driven by an overall increased spiking rate of cells in the relevant microcircuit, it could also occur as a result of a balanced, proportional increases in excitatory and inhibitory conductance, a net excitation, or even an increased inhibition (Logothetis, 2008).

A finding that BOLD signal increases as a function of a stimulus parameter (e.g. reward magnitude) may not lend as strong a basis as findings that a cellular measure (e.g. single-cell spike rate), in which case the interpretation would be that these cells positively track the size of potential reward. What we know from single-cell recordings is that there may be roughly a 50:50 ratio between neurons in the same region which have positive and those with negative correlation between their firing rate and a stimulus property, say, effort size (e.g., Kennerley, Dahmubed, Lara, & Wallis, 2008). Therefore, a macro method like fMRI which reflects metabolic demand over averaged neuronal activity in a region may not be able to selectively pick up a pure population of neurons which have positive correlation with effort size. What could potentially be interpreted from an fMRI finding is that

the relevant region is sensitive to a stimulus parameter, and that the direction of this 'tracking' is to be adjudicated based upon convergent methods and findings.

An important limitation to fMRI is its susceptibility to motion artefacts. This becomes critical in my attempt at imaging the brain while subjects perform a motor vigour task in the scanner. I addressed this issue by i) minimising motion during scanning through (almost excessive) padding around subjects' head and the arm used for squeezing and (almost excessive) explicit instructions and constant reminders, ii) by correcting any deformed images through unwarping (see imaging analysis section below), and iii) by taking into account motion-related BOLD signal through entering motion parameters as parameters in all brain analyses.

2.2 Specific methodology

2.2.1 Image acquisition

In studies 2 and 4, I used a 3T Siemens TRIO system (Siemens, Erlangen, Germany) with 12-channel head coil to acquire both T1-weighted anatomical images and T2*-weighted MRI transverse echoplanar images (EPIs) (64x64mm, TE = 30 ms, TR study 2/ TR study 4 = 2.72 s/ 3.36 s) with BOLD contrast. The EPI sequence was optimised for maximising signal in inferior brain regions (Weiskopf, Hutton, Oliver Josephs, & Deichmann, 2006). Each EPI comprised forty (study 2) or forty-eight (study 4) 3-mm-thick contiguous axial slices taken every 3 mm, positioned to cover the whole orbitofrontal cortex, striatum, up to the anterior cingulate and motor cortices. In total, 180 - 212 (study 2) or 212-220 (study 4) volumes were acquired for each participant in one session. The first five (study 2) or four (study 4) volumes were discarded to allow for T1 equilibration effects. The field maps were acquired between the second and third scanning sessions. For the structural images I acquired a standard high-resolution T1-weighted anatomical image with acquisition matrix 256x240, TR/TE/Flip Angle = 7.92ms/ 2.48ms/ 16°, voxel size 1 x 1 x 1 mm, 176 axial slices (Deichmann, Schwarzbauer, & Turner, 2004).

2.2.2 Imaging analysis

Data were analysed using Statistical Parametric Mapping (SPM8b; Wellcome Trust Centre for Neuroimaging, London, UK, http://www.fil.ion.ucl.ac.uk/spm). Five preprocessing steps involved intra-modal realignment and unwarping, inter-modal co-registration, segmentation, normalisation, and smoothing.

Realignment and unwarping

All EPI volumes were re-aligned to the first volume to correct for inter-scan movement. Images were unwarped using fieldmaps to remove unwanted gripping-related variance without removing variance attributable to the motor task (Andersson, Hutton, John Ashburner, Turner, & Friston, 2001).

Co-registration

The mean motion-corrected image was co-registered to individual's T1 images using a 12-parameter affine transformation. To correct for different acquisition times, the signal measured in each slice was shifted relative to the acquisition of the lower slice using sinc interpolation in time.

Segmentation

Individual T1 images were segmented based on grey and white matter, a method fairly robust and accurate in creating spatial normalisation parameters for the EPI and anatomical images (John Ashburner & Friston, 2004).

Spatial normalisation

To allow across-subject comparison, the co-registered EPI and T1 volumes were normalised using segmentation parameters, based on the Montreal Neurological Institute (MNI) reference brain in Talairach space (Talairach & Tournoux, 1988) and re-sampled to 3 x 3 x 3 mm^3 and 1x1x1 mm^3 voxels, respectively.

Spatial smoothing and filtering

All normalised images are smoothed with an isotropic 8 mm full-width half-maximum Gaussian kernel to account for inter-subject differences and allow valid statistical inference according to Gaussian random field theory (K. J. Friston, J Ashburner, et al., 1995; K. J. Friston, Holmes, et al., 1995). The time series in each

voxel were high-pass filtered at 1/128 Hz to remove low-frequency confounds and scaled to a grand mean of 100 over voxels and scans within each session.

Statistical modelling

I performed random-effect, event-related, statistical analyses. In every General Linear Model (GLM), I convolved each regressor (described in the relevant chapters) with a canonical hemodynamic response function and its temporal derivatives. Motion parameters from pre-processing were entered into the design matrix to further account for BOLD noise related to gripping.

In studies 2 and 4, a separate GLM for each participant was specified by creating separate regressors representing different events for each of the scanning sessions. In study 2, I ran simple t-tests between regressors testing the contrasts of interest (main effects and interaction) at first level for each individual. Consistency across the resulting maps of sensitivity for each participant was tested in a series of one-sample t-test as group analyses. In study 4, each regressor-of-interest was contrasted against baseline activity at first level for each individual, and these t-contrasts were brought over into the second level and entered into a series of *F* tests as group analyses to assess for main effects and interactions.

2.2.3 Effort manipulations and measurements

Grip device

In studies 1-4, I utilised a pneumatic handgrip device as effort manipulation. Participants either used their dominant hand when completing a behavioural task, or their right hand when being scanned (all fMRI participants were right-handed). The handgrip device was molded from two plastic cylinders that compressed an air tube that was connected to a transducer (Honeywell, Morristown, NJ) to convert air pressure into a voltage output. Thus, variation in air compression within the cylinders due to the force applied resulted in different voltage signals, and these are linearly proportional to exerted grip force. The signal was recorded (Spike2, Cambridge Electronic Design) and transmitted to MATLAB 6.5 (www.mathworks.com).

Visual stimuli were presented using Cogent 2000 (http://www.fil.ion.ucl.ac.uk/ and http://www.icn.ucl.ac.uk/) and Cogent Graphics

(John Romaya at the LON at the Wellcome Trust Centre for Neuroimaging, at UCL). I constructed a squeeze stimulus in a red vertical bar which was a direct translation of the recorded grip force signal as veridical, real-time visual feedback for squeezing.

To estimate the reliability of the two grip devices used in studies 1-4, I asked eight local staff in the department to simply grip as hard as they could using their dominant hand on two days, separated by one week. Participants gave verbal consent to participate, no financial reimbursement was provided. Participants were asked three times to produce maximum force, each time for a period of 3-5 secs during which the highest value was taken as one data point. The highest of the three measurement times was treated as the maximum force for the day. The difference between their maximum force between days 1 and 2 for both devices were non-significant. This suggests reliability of the grip devices across days, and gives validation for my instruction to 'squeeze as hard as you can'.

At the start and end of each study, I asked participants to produce maximum forces, and calibrated grip levels based on the pre-task maximum value. The difference values between maximum force at beginning and end were not significant, suggesting that behaviour in the tasks was not influenced by fatigue. Before completing the experimental blocks, participants were shown the calibrated squeeze stimulus and had the opportunity to try squeezing guided by the vertical bar. Participants were explicitly instructed to only use one hand, to never switch hands or use both hands when squeezing. All associated *t* values are in TABLE 2-1.

Table 2-1 *T* values for calibration of grip device show no difference between maximum force on day 1 and 2, or before and after experimental tasks.

No.	Device/ Study	T values (max start – max end)
1.	Grip device 1 (study 1-2)	$t(7) = 1.99$, $p = 0.08$
2.	Grip device 2 (study 3-4)	$t(7) = 0.46$, $p = 0.65$
3.	Study 1	$t(13) = 1.84$, $p = 0.08$ (eleven missing data)
4.	Study 2 behavioural task	$t(16) = 0.04$, $p = 0.96$
5.	Study 2 scanning task	$t(17) = 0.68$, $p = 0.50$
6.	Study 3	$t(18) = .76$, $p = .45$
7.	Study 4 training day	$t(20) = 1.74$, $p = .09$
8.	Study 4 scanning day	$t(19) = 1.08$, $p = .29$ (one missing data)

2.2.4 Statistics

Throughout studies 1-5, I conducted two- or three-way within-subjects Analysis of Variance (ANOVA) using SPSS 11.5, 13.0 or 19.0. *F* values were calculated under the assumption of sphericity, and I report Greenhouse-Geisser *F* values when sphericity was violated. To allow for comparison with other studies, I report effect size of all *F* tests in partial eta squared (η_p^2), a way to gauge the strength of association between the independent and dependent variables by partialling out other factors from the total non-error variance (Pierce, Block, & Aguinis, 2004).

Chapter 3 Effort and Choice (study 1 & 2)

Abstract

The possibility that we will need to invest effort influences our future choice behaviour. Indeed deciding whether an action is actually worth taking finds its pathological expression in human apathy or inertia. There is a well- developed literature on brain activity related to anticipation of effort, but how effort impacts on actual choice is less well understood. Here, I investigated choice behaviour and brain activity, using fMRI, in two studies where healthy participants are required to make decisions between effortful gripping, where the factors of force and reward were varied, and an option of merely holding a grip device for a minimal monetary reward. Behaviourally, I show that force level influences the likelihood of choosing an effortful grip. I observed greater activity in the putamen when participants opt for a low effort option compared with when they opt for high effort option. The results suggest that effort discounts the value of an action, and second, over and above a nonspecific role in movement anticipation and salience, the putamen plays a crucial role in choice computations that entail effort costs.

3.1 Introduction

The cost involved in an action is an important determinant of choice behaviour (Kennerley, Dahmubed, Lara, & Wallis, 2009). A number of animal and human experiments have examined how effort determines choice, and crucially, how the brain integrates effort into an action value (Croxson et al., 2009; Floresco & Ghods-sharifi, 2007; Floresco, Tse, et al., 2008; Kennerley et al., 2009; Rudebeck et al., 2008; Salamone et al., 1994; Walton, Croxson, Rushworth, & Bannerman, 2005; Walton et al., 2009). Other costs are better understood as, for example, discounting prospects whose outcomes entail possible pain or loss (Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006; Seymour, Daw, Dayan, Singer, & Dolan, 2007; Talmi, Dayan, Kiebel, Frith, & Dolan, 2009; Talmi, Seymour, Dayan, & Dolan, 2008), and temporal delay (Kable & Glimcher, 2007; McClure, Ericson, Laibson, Loewenstein, & J. D. Cohen, 2007; Pine et al., 2010; Rudebeck, Walton, Smyth, Bannerman, & Rushworth, 2006). The neurobiology of effort choice remains relatively underexplored.

In CHAPTER 1**,** I have described an extensive literature based upon animal studies which implicates regions such as NAc and ACC in effortful choice, but very limited human studies have examined the neural representation of physical effort *to choose* an action. In non-choice contexts, the striatum and the ACC are activated when participants anticipate an upcoming action that entails effort (Croxson et al., 2009). More specifically, activity in the striatum is correlated with the anticipated effort for an action.

To the best of my knowledge, the only other empirical work investigating human choice to invest in physical effort is that of Prevost et al. (2010). In their study, Prevost and colleagues conducted a careful delineation of neural regions subserving effort and delay costs. In line with previous animal work and that of Croxson et al. (2009), they report that BOLD signal in the ACC negatively tracks the subjective value of actions that require the investment of physical effort. In their study, effort hyperbolically discounts the subjective value such that larger effort gives lower subjective value. Thus, in light of growing evidence for involvement of the striatum and ACC in value based decision-making and effort anticipation, here I hypothesised ACC and striatal involvement in action choice, where a neural computation entails an integration of effort as a cost.

In this chapter I report behavioural and neuroimaging data on effort-based decision making. I employed a simple effort-based choice task where participants decide between holding a grip device and effortful gripping. The holding option entailed no effort and a minimal reward. The gripping option varied across two factors, namely monetary reward and force levels (percent of individual maximum force) indicated by a visual stimulus. In study 1, I report behavioural evidence for effort-discounting using parametric levels of effort and reward. In study 2, I reduced the effort and reward levels and adapted the paradigm to look for striatal and ACC involvement when humans make choices which entail physical effort. In the imaging analysis, brain activity was time-locked to events at the time of choice, in order to index activity associated with, and effort modulation on, the decision to grip. I hypothesised activity in striatum and ACC would be associated with biasing choice away from actions that entail greater physical effort.

3.2 Methods

Participants

All participants were recruited through the psychology participant database at Univesity College London (UCL) and the study was approved by the UCL ethics committee.

Sixty-six healthy individuals participated in the behavioural experiments. I excluded data from twenty-four participants as these were tested during pilot phase under various experimental designs. I report data from 25 and 17 participants for behavioural data in study 1 and 2, respectively (*M* age = 26 (SD = 6 years) for 17 participants). All sixty-six participants were paid £10 - £20 depending on duration of experiment.

Eighteen right-handed healthy individuals (five females, *M* age = 27 (*SD* = 3) years) participated in the fMRI experiment. One participant was excluded from the analysis of brain activity due to excess motion artifact, but was included in the behavioural analysis. These participants were paid £25 - £30 depending on duration of experiment.

Stimuli

The stimuli potentially requiring effortful gripping is referred as 'grip' stimuli; and the stimulus requiring non-effortful holding of the hand-grip device as the 'hold' stimulus. As with the visual stimuli used in Croxson et al. (2009), 'grip' stimuli comprised of red circles with two black lines (see FIGURE 3-3). Where the vertical lines are located in the circle indicated effort with two levels (leftmost is lowest effort, rightmost is highest effort), while the location of the horizontal lines indicated reward levels (bottom is lowest reward, top is highest reward). The 'hold' (no grip) stimulus is a red circle with a horizontal line at the bottom representing a fixed low reward and no vertical line.

I used a similar squeeze stimulus described in CHAPTER 2 in the shape of a thermometer with a yellow horizontal line to indicate the squeezing target, set at a thermometer height corresponding to the chosen effort level (e.g., 80% of thermometer height for 80% effort level; FIGURE 3-3). This moving thermometer was presented after a 'grip' choice whilst a 'frozen' thermometer was presented after a 'hold' choice.

3.3 Study 1: 5 x 5 Effort by Reward design

As I did not know the optimal experimental parameters for this task, I trialed various reward and effort parameters ranging between 1 to 20 pence and 30% to 90% of maximum force, respectively, with a fixed 1 or 2 pence reward level for the 'hold' option. I first report choice data from 25 subjects with changing parameters across subjects (FIGURE 3-2A). Within these 25, I report data from 10 subjects who underwent the same experimental parameters of 2, 3, 6, 9 and 12 pence reward (and 1 pence for 'hold' option), and 50, 60, 70, 80 and 90% effort levels (FIGURE 3-2B). See APPENDICES for precise effort and reward values for the first 15 subjects.

Procedure

Participants undertook, successively, force calibration, training and experimental blocks, and completed post-scan questionnaires before being debriefed, and reimbursed for their participation.

The experimental task comprised of 150 trials, which were split into 3 blocks of 50 trials and this gave 6 repetitions of each of the 25 unique choices

between 'grip' vs. 'hold' action, the order of presentation for these 25 stimuli pairs is pseudo-random.

There were two training blocks. In the first, participants learned how to use the gripper, by pressing a spacebar to commence the trial, squeezing and reaching the yellow target line during presentation of the thermometer cue. Subjects trained to reach the line within 2 secs and stay at that effort level for another 4 secs (FIGURE 3-1A). The trial is aborted every time participants gripped too slow, too fast, or released the gripper before 6 secs have elapsed. This approach was motivated by a need to have tight control of the amount of effort exerted at each trial, and to control for time differences between low and high effort levels. Participants completed 15 practice trials (3 grips in each effort level) to ensure that that they were able to complete the gripping successfully.

In the second training block, participants learned the values of each 'grip' and 'hold' action by completing single-stimulus training trials which comprised of a cue presentation and a button press before squeezing, and a reward outcome presentation after squeezing. There were 52 trials (2 repetitions for 1 'hold' action and 25 'grip' actions from 5x5 effort and reward levels).

At the beginning of each training trial, either a 'hold' or 'grip' cue was randomly presented on either side of the screen until participants made a button press with their non-dominant hand. Following a 'hold' cue, they saw a 'frozen' thermometer for 6 secs (participants typically just held the hand-grip), whereas following a 'grip' cue, they saw the thermometer for 6 secs while squeezing to reach the target. FIGURE 3-1A depicts grip trajectory for one subject in trials with varying levels of effort. As seen here, subject 1 is able to commit to the squeezing criteria for each effort level. Once 6 secs elapsed, participants saw the corresponding reward outcome as indicated by the cue. For the 'grip' trials, if they fail to squeeze as trained, the outcome is zero pence and the trial is aborted. Participants were not informed about the precise effort and reward amounts, but learned the stimulus-effort-reward contingencies from experience (Hertwig et al. 2004).

Behavioural Choice Task. At the start of each experimental trial, a fixation cross appears for 200 ms, followed by 'grip' and 'hold' cues, randomly presented on left and right of the fixation cross. Participants choose one of these cues with a button press, using their non-dominant hand. They have up to 4 secs to respond, otherwise they miss that trial and the next trial then commences immediately. The two cues remain on the screen for 1 sec before a 'GET READY' message appears for

500 ms. The thermometer appears and subjects have to squeeze for 6 secs before the next trial begins. Participants do not see the monetary outcome in every trial, but only at the end of a block. FIGURE 3-1B depicts behaviour of two participants in the first 50 trials; subject 1 is able to commit to his/her choices, while subject 2 received several aborted trials due to failure to squeeze according to criteria.

Recalibration. Around halfway through the second block, the grip force was recalibrated unobtrusively. This is done by measuring the baseline force on a trial when subjects chose the 'hold' option. For all subjects (except two subjects) this took place between the 80th-85th trial. Then immediately after this trial, there was a surprise trial wherein subjects are told to get ready to grip for £2 reward. They then saw the thermometer with the yellow line at the top of the meter, telling them to grip as maximum. The experiment continued as usual with the new baseline and maximum force values. We did not implement recalibration in study 2 to avoid mistaken recalibrated values unnecessarily.

Results

Throughout this chapter, the main behavioural measure is how often participants accepted the 'grip' action in percentages (% choice to 'grip'). I averaged across 25 participants with slightly varying effort and reward parameters and entered 'grip' choice into a 5 (effort) x 5 (reward) ANOVA. As seen in FIGURE 3-2A, I found significant effects of effort, $F(1.92,46.24) = 34.01$, $p < .00001$, $\eta_p^2 = .58$, reward, $F(1.37,32.97) = 15.12$, $p = .0001$, $\eta_p^2 = .38$, and effort by reward interaction, $F(7.30,175.32) = 2.67$, $p = .01$, $\eta_p^2 = .10$.

Because effects might be driven by non-systematic effort and reward manipulations, I separated the last 10 participants who experienced the same experimental manipulations and ran two-way ANOVA on their % 'grip' choice (FIGURE 3-2B), revealing significant effects of effort and reward, $F(1.82,16.39) = 13.83$, $p = .0003$, $\eta_p^2 = .60$; $F(1.25,11.31) = 7.73$, $p = .013$, $\eta_p^2 = .60$, respectively, but a non-significant interaction, $p = .07$. This parametric design provides evidence for effort discounting. Put simply, a willingness to exert effort is not only governed by a reward manipulation, but also by an effort manipulation.

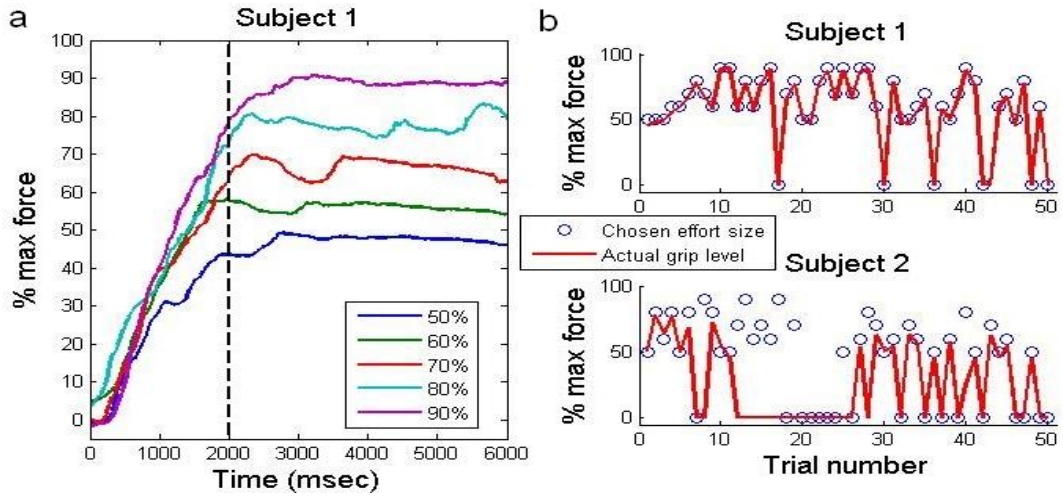

Figure 3-1 a) Example of squeezing time course for one subject for trials with different effort levels. To successfully execute the squeeze, subjects have to reach the target line within 2 secs and then maintain that squeeze level within a margin for another 4 secs. Black dotted line indicates the time the target level is reached. b) Example of choice-execute match/mismatch for the first 50 trials from two subjects. Blue circle indicates participants' chosen effort level; red line indicates actual squeeze level for that trial, averaged over the last 4 secs of thermometer presentation. Top: Subject 1 shows a good match between choice and squeeze execution. Bottom: Subject 2 failed to execute their chosen effort levels in the first twenty trials, decided to go for the no-grip option for subsequent trials, but managed to match their effortful choices in the later trials.

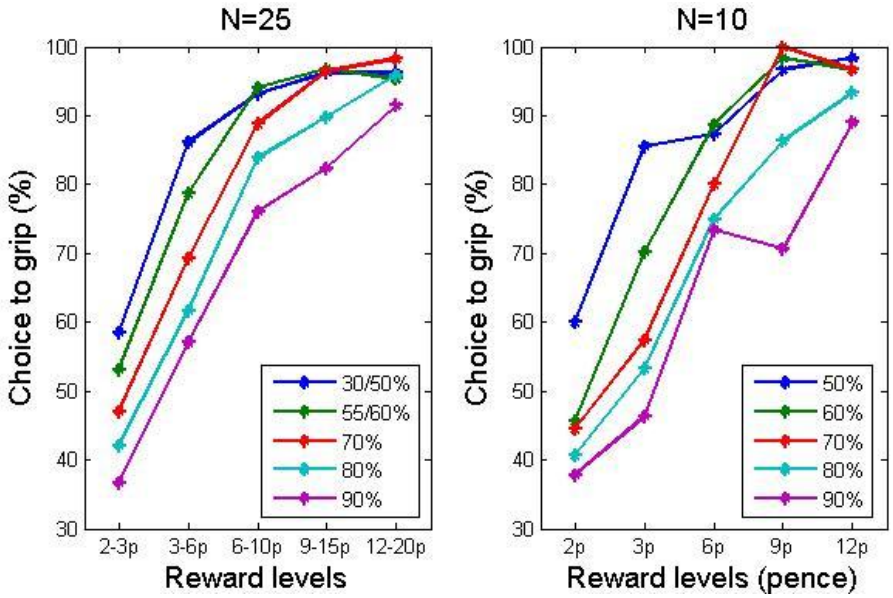

Figure 3-2 Proportion of 'grip' choices across reward and effort levels. a) I tested 25 subjects with parametric levels, but actual values were changing across subjects. Nevertheless I show significant effects of effort and reward and effort by reward interaction. b) Within these 25, 10 participants received the same effort and reward parameters, and effort and reward effects were significant, but interaction was n.s.

Discussion

In study 1, I demonstrated the viability of the effort-based choice paradigm which tightly manipulates hand force and controls for temporal confounds. In this task, participants were able to distinguish different effort and reward parameters, make and execute choices based on these parameters. It was important that the choices participants made were genuine and that they reflected an approximation of their ability to execute their selected actions. Only two out of 25 participants in study 1 had more than 10% aborted trials due to failure to squeeze according to the criteria, and this assured choices based on effort integration. Having established feasibility I then adapted this task to an fMRI environment to allow recording of BOLD responses that reflect sensitivity to effortful choices.

3.4 Study 2: 2 x 2 Effort by Reward design

In study 2, I used two levels of effort and two levels of reward. To avoid monotony I varied the effort and reward levels, trial-by-trial, by adding a pseudo-random value to base values of effort (40% and 85% maximum force) and reward (3 and 11 pence) of each stimulus. These values were drawn from a normal distribution with mean zero and one unit of standard deviation, ranged from -5.2 to 5.4% for effort and -2.6 to 2.7 pence for reward. I used a fixed, 2 pence, reward for the 'hold' option.

Procedure

Behavioural subjects entered the testing room, and completed the tasks while sitting upright, whereas fMRI subjects lay on the scanner bed to undergo, successively, force calibration, training, four experimental blocks and a final structural scan. Calibration and training blocks were completed as participants lay on the scanner bed outside the magnet, while experimental blocks and the structural scan were completed as participants lay inside the magnet. While BOLD data is recorded, participants completed the experimental task with a rest period (up to 3 mins) between the blocks. Participants completed post-scan questionnaires outside the scanner at end of experiment.

Here, I adapted the choice task to the fMRI set up by separating the choice and squeeze events into CHOICE and EXECUTE mini blocks (clearly prompted at the start each period) to remove motor preparatory brain activity in anticipation of

gripping. In CHOICE periods, participants made a series of twelve consecutive choices between a 'grip' and 'a hold' cue (FIGURE 3-3). In EXECUTE periods, participants executed their preceding selected actions by gripping (or simply holding) a hand device at the corresponding effort level to receive the corresponding reward amount. To further de-correlate brain signal for choice from that for execution, the fMRI participants only executed 75% of the choices, randomly selected from the preceding twelve CHOICE trials. A pair of the CHOICE and the EXECUTE periods was repeated five times in each block.

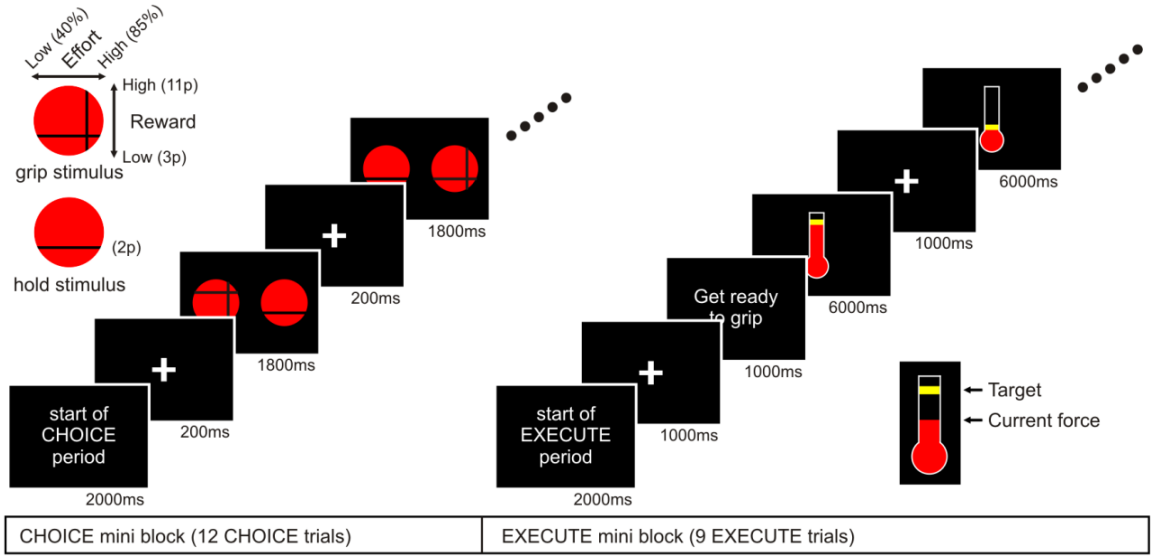

Figure 3-3 Left-top: Grip and hold stimuli. Grip stimulus: a horizontal line indicates reward levels (in pence), a vertical line indicates effort levels (in % maximum grip). In study 2, I added a random value to effort and reward levels of each grip stimulus; values in brackets show the averages. Hold stimulus: a horizontal line indicates a fixed reward value in pence. Middle: A schematic of the task. CHOICE period: in each CHOICE trial, a fixation cross appears, followed by a 'grip' and a 'hold' stimulus. Participants had to make a decision to grip or to hold. There were twelve CHOICE trials; each grip stimulus was presented pseudo-randomly. At the end of each CHOICE period, the computer randomly selects nine out of twelve participants' choices from the preceding CHOICE period to be executed. EXECUTE period: immediately following the 12th CHOICE trial, the EXECUTE period comprising of nine trials; either a grip or a hold trial, commences. In the grip trials, a thermometer with a target level was displayed to guide squeezing. In the hold trials, a 'frozen' thermometer was presented. Each participant carried out five sets of CHOICE and EXECUTE period in total. Bottom-right: A thermometer stimulus is used to guide squeezing during EXECUTE period. The red 'mercury' indicates current force level; yellow horizontal line indicates target level. Figure taken from Kurniawan et al., JNeurophysiol, 2010, Am Physiol Soc, used with permission.

Overall, the behavioural participants completed 180 'choice' and 'execute' trials, split in three blocks of 60 'choice' and 60 'execute' trials. The fMRI

participants completed 240 'choice' trials and 180 'execute' trials, split in four scan sessions of 60 'choice' trials and 45 'execute' trials.

To improve fMRI efficiency, I shortened the choice event, such that participants only had 1800 ms to select either 'grip' or 'hold' stimulus. At the beginning of each EXECUTE trial, a fixation cross appeared for 200 ms, followed by a message to get ready to grip in the case of a 'grip' execution trial, and the thermometer for 6 secs. As before, the trial aborts if participants do not reach the target within 2 secs after thermometer onset or if they release the hand-grip before 6 secs expire with a reward outcome of zero pence. On average in the fMRI experiment, 2% (SD = 0.7%) of all trials were aborted; these trials were included in the fMRI analysis. To reduce noise caused by no-go signal during a 'hold' execution trial, participants did not see any prompt message, and instead were immediately presented with a 'frozen' thermometer for six seconds.

Questionnaires. Immediately after the experimental task, participants completed a 20-item persistence scale that measures individual propensity to work harder when facing daily challenges (e.g., 'I usually push myself harder than most people do'; APPENDICES) (Cloninger et al. 1993) and made ratings of how much they like the effort-reward combinations of the 'hold' and 'grip' cues. Participants also responded to two manipulation check questions for reward and effort on paper.

Imaging analysis

To highlight activity correlating with anticipated effort, and with the choice to grip or hold, I defined four regressors-of-interest representing four event types that varied in effort level and participants' choice (low effort vs. high effort and grip vs. hold) at choice onset: grip-low effort (gripLE), grip-high effort (gripHE), hold-low effort (holdLE), and hold-high effort (holdHE). Furthermore, to assess activity correlating with reward, I entered a trial-by-trial reward value (3 or 11 pence $\pm$ a random value) as a parametric modulator for each of the four regressors. I entered two regressors-of-no interest from the 'grip' and 'hold' trials in the EXECUTE periods at thermometer onset with 6 secs duration; suprathreshold activity for grip > hold contrast in execute periods is found in left primary motor cortex (see APPENDICES).

I computed a set of contrasts for each participant, testing the main effects of choice, effort, and an interaction. As I found persistence correlated with

behavioural choice in the scanner (reported below), I entered persistence score as a covariate at the second level and ran a whole-brain analysis, thresholded at p = .001 uncorrected, >5 voxels, to search for areas active in response to choice (grip vs. hold), effort (low vs. high), choice-effort interaction, and simple effects of effort at both choices (gripHE vs. gripLE and holdHE vs. holdLE).

3.4.1 Behavioural Results

I first report behavioural data from behavioural (N = 17) and fMRI participants (N = 18).

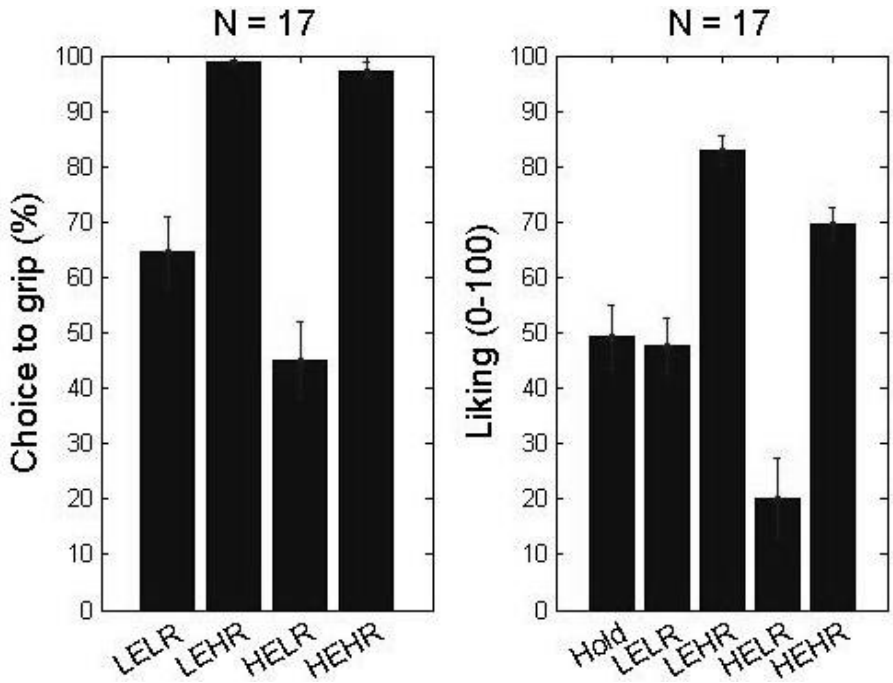

Figure 3-4 Choice and subjective ratings from behavioural participants (N =17) (mean ± SEM). LELR = Low effort-low reward, LEHR = Low effort-high reward, HELR = High effort-low reward, HEHR = High effort-high reward.

Choice. I replicate the behavioural choice effects seen in study 1. In both sets of subjects (FIGURE 3-4 (left) & FIGURE 3-5A (dark shade)), acceptance rate for gripping is significantly higher for actions with low than high effort and for actions with high than low reward. An effort by reward interaction for behavioural subjects was significant, but not for fMRI subjects. For the same low reward, the behavioural subjects chose low effort significantly more than high effort actions,

whereas choice difference between low and high effort actions for high reward is n.s. Statistical values are shown in TABLE 3-1.

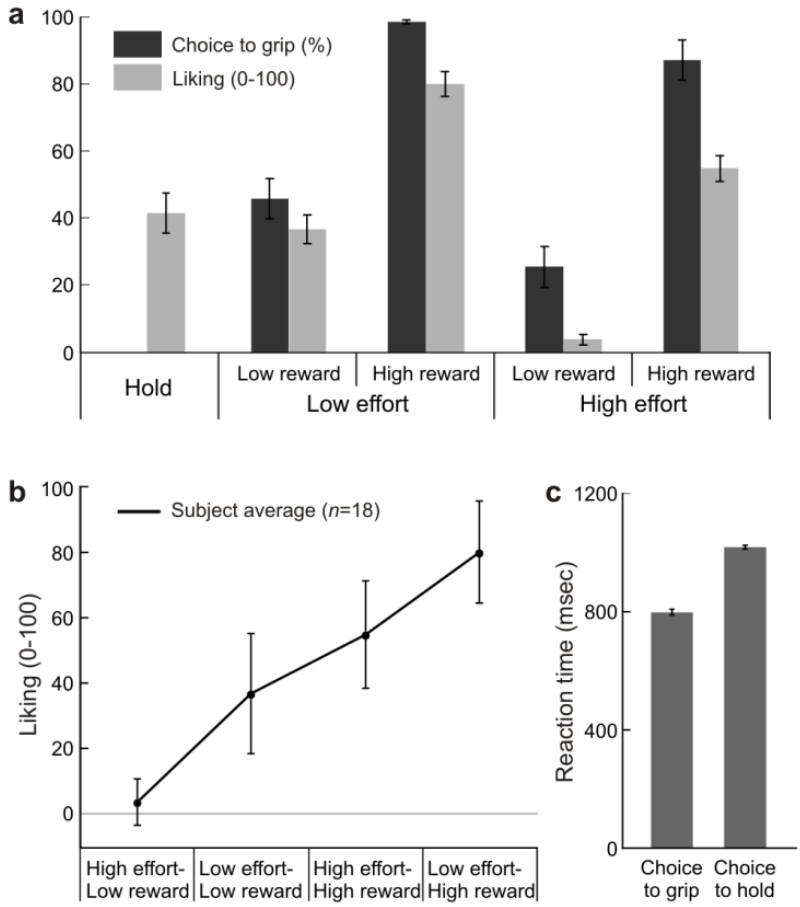

Figure 3-5 Behavioural choice, subjective rating, and RT's. a) Proportion of trials where participants chose to grip (dark shade) and their subjective rating (light shade) for each option. Participants chose to grip more often when the reward offered was high than when it was low, and when the effort anticipated was low than when it was high. The interaction was non-significant. Liking (light shade) was higher for options with high reward than for options with low reward, higher for options with low effort than for options with high effort, and comparable between 'hold' and low effort-low reward. The interaction was non-significant. b) The same liking data to (a), showing that on average, the order of rating from lowest to highest is: high effort-low reward, low effort-low reward, high effort-high reward, and low effort-high reward. c) RT's were slower for choice to hold than for choice to grip. (Mean + SD). Figure taken from Kurniawan et al., JNeurophysiol, 2010, Am Physiol Soc, used with permission.

Subjective rating. As seen in FIGURE 3-4 (right) and FIGURE 3-5A (light shade), participants rated the 'hold' and low effort-low reward option comparably. I computed a difference score between liking for 'hold' (as baseline) and each of the 'grip' options for each participant, and found effort and reward main effects;

subjective liking was significantly higher for high than low reward and for low than high effort.

Table 3-1 Two-way effort by reward ANOVA results for behavioural (N = 17) and fMRI (N = 18) participants in study 2. LELR = Low effort-low reward, LEHR = Low effort-high reward, HELR = High effort-low reward, HEHR = High effort-high reward.

No.	Effect	F / t values	N
1.	Choice: Low > High Effort	$F(1,16) = 87.41$, $p < .00001$, $\eta_p^2 = .84$	17
2.	Choice: High > Low Reward	$F(1,16) = 5.50$, $p = .03$, $\eta_p^2 = .25$.	17
3.	Choice: Effort x Reward	$F(1,16) = 4.50$, $p = .04$, $\eta_p^2 = .21$.	17
4.	Choice: LELR > HELR	$t(16) = 2.26$, $p = .03$.	17
5.	Choice: LEHR = HEHR	$t(16) = 1.22$, $p = .23$.	17
6.	Choice: Low > High Effort	$F(1,17) = 13.07$, $p = .002$, $\eta_p^2 = .43$.	18
7.	Choice: High > Low Reward	$F(1,17) = 105.08$, $p < .0001$, $\eta_p^2 = .86$.	18
8.	Liking: Low > High Effort	$F(1,16) = 14.82$, $p = .001$, $\eta_p^2 = .48$.	17
9.	Liking: High> Low Reward	$F(1,16) = 52.74$, $p < .00001$, $\eta_p^2 = .76$.	17
10.	Liking: Effort x Reward	$F(1,16) = 5.84$, $p = .02$, $\eta_p^2 = .26$.	17
11.	Liking: LELR > HELR	$t(16) = 3.64$, $p = .002$.	17
12.	Liking: LEHR > HEHR	$t(16) = 3.26$, $p = .004$.	17
13.	Liking: High > Low Reward	$F(1,17) = 173.41$, $p < .0001$, $\eta_p^2 = .91$	18
14.	Liking: Low > High Effort	($F(1,17) = 86.61$, $p < .00001$, $\eta_p^2 = .83$	18

Again, effort by reward interaction for liking in behavioural subjects was significant, but not in fMRI subjects. In behavioural subjects, a reduced likeability for high compared to low effort action is greater when the reward is low than when it is high. Based on the fMRI group-averaged liking scores, I could describe the order of subjective liking for these actions, from lowest to highest: high effort-low reward, low effort-low reward, high effort-high reward, and low effort-high reward (see FIGURE 3-5B). These findings suggest a fair generalisability to common views on effortful and rewarding actions whereby actions with more effort and less reward are less liked.

Response times. Overall, fMRI participants took significantly longer in choosing to grip than to hold ($t(17) = 28.95$, $p < .0001$; FIGURE 3-5C). I ran a separate ANOVA to formally test the effects of effort and reward on response times (RTs). A 2 x 2 (effort x reward) ANOVA revealed that, regardless of choice (grip/ hold), RTs were slower for low ($M = 994$ ($SD = 25$ ms)) than high reward ($M = 764$ ($SD = 60$ ms)), $F(1,17) = 566.59$, $p < .0001$, $\eta_p^2 = .97$; and for high ($M = 882$ ($SD =$ 126 ms)) than low effort ($M = 876$ ($SD = 125$ ms)), $F(1,17) = 4.61$, $p < .046$, $\eta_p^2 = .21$. There was no significant interaction. Participants represented each option by taking account of both its effort and reward. I ran a separate imaging analysis with RTs as a covariate of-no-interest at the first level analysis, and this analysis did not change the main findings reported below. I found no difference in RTs from behavioural participants, Fs < .7, ps > .30.

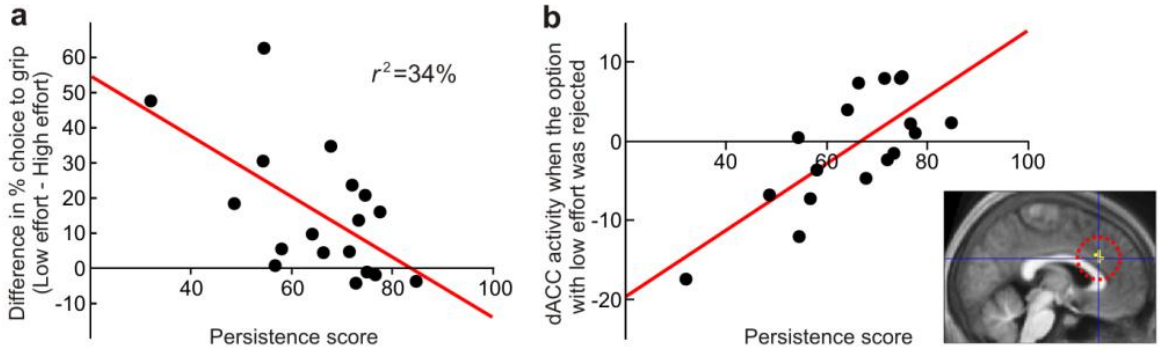

Figure 3-6 Persistence, behavioural choice, and dACC signal. a) Persistence is negatively correlated with the effect of effort on choice (N = 18). Regardless of reward, low persistence is associated with a higher preference for options with low effort, whereas high persistence is associated with indifference between options with low effort and options with high effort. b) Activity in the dACC when the rejected option entailed low effort is positively correlated with persistence (p <.001 unc., 11 voxels; N = 17). Figure taken from Kurniawan et al., JNeurophysiol, 2010, Am Physiol Soc, used with permission.

Persistence. A persistence trait is linked to self-directedness (Cloninger et al. 1993), a characteristic especially lacking when an individual suffers from apathy. I calculated a correlation between persistence scores and the effects of effort, reward, and interaction on choice. In the fMRI participants, I found that the main effect of effort on choice, regardless of reward level, was negatively correlated with persistence, $r = .59$, $r^2 = 34\%$, $p = .01$. As persistence score decreases, there was a greater difference between choice to grip an option with low effort compared to choice to grip an option with high effort: i.e., less persistent participants much preferred low compared to high effort, while those with high persistence (or less apathy) chose to grip options with low and high effort equally often (FIGURE 3-6A).

Correlations with reward and interaction effects on choice and those in behavioural subjects were non-significant.

Manipulation checks. In the fMRI experiment, I checked whether participants understood the reward and effort amounts indicated by the cues, by presenting isolated visual markings for reward and effort. For reward, participants were shown two red circles each with a horizontal line, one at the top and one the bottom, and responded to the question 'how much money does the horizontal line on the circle mean?' For effort, participants were presented with two thermometer cues each with a yellow line, one at the top and one at the bottom, and responded to the question 'how much money do you think is considered a fair pay for gripping at the yellow line 10 times in a row?'

Responses to the reward item show the desired effect: participants estimate the amount of reward for high and low reward cues reasonably accurately (FIGURE 3-7A) and the difference is significant, $t(17) = 18.93$, $p = .00001$). Likewise, the effort manipulation check also show that high and low effort levels are perceived differently: the estimate for an expected fair pay to squeeze ten times in a row at high effort was significantly greater than that for low effort, $t(17) = 2.80$, $p = .012$ (FIGURE 3-7A).

Additionally I checked if the decision to accept an action with low reward was largely driven by a simple reward comparison between reward in the 'grip' cue and the fixed reward of the 'hold' cue (2 pence), regardless of effort levels. If this is so, then by looking at the trial-by-trial random values added to the 'grip' cue, it should be possible to detect that the greater the random value that is subtracted from the 'grip' reward value, the more likely the action is to be rejected. To do this I looked specifically at both low reward conditions (LELR and HELR) and calculated the average random values associated with accepted and rejected trials. FIGURE 3-7B shows that the random values associated with accepted and rejected low effort-low reward and high-effort-low reward trials do not differ significantly, ps > .30 (FIGURE 3-7B).

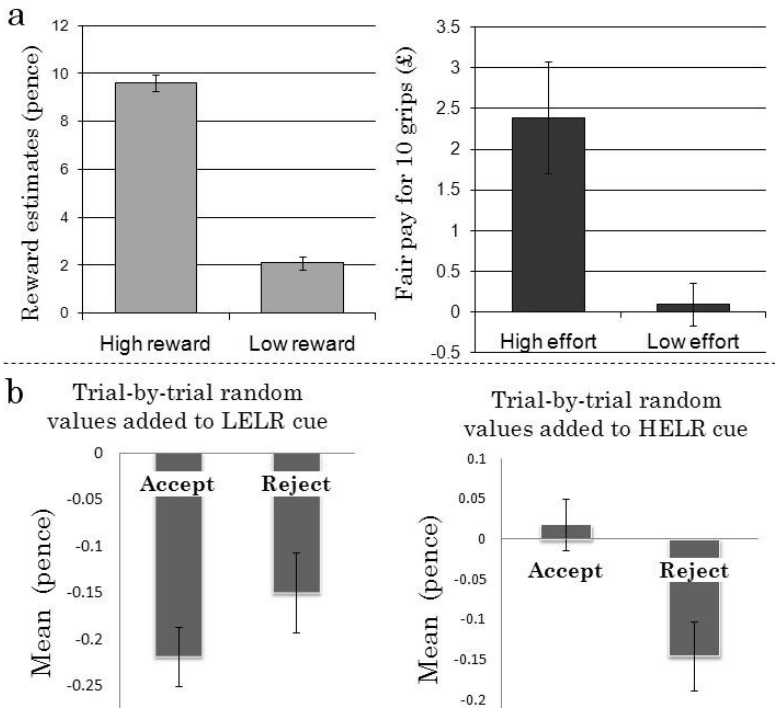

Figure 3-7 Additional behavioural measures for fMRI subjects (N = 18). a) Successful manipulation for effort and reward. b) Proportion of acceptance and rejection rate of low effort-low reward (LELR; left) and high effort-low reward (HELR; right) cues as a function of trial-by-trial random values added to reward. Random values do not affect acceptance rate for both LELR and HELR.

3.4.2 fMRI results

I sought to extend my findings regarding the influence of effort on behavioural measures and look for brain regions where activity reflects a bias in choice away from effortful actions. To do this, I examined BOLD response when participants chose to grip or to hold, and when the required effort was high or low. I added trial-by-trial reward level as a parametric modulator for each regressor and persistence score as a subject-by-subject parametric regressor at second level.

Choice-related activity

The main effect of choice (choice to grip > choice to hold) was associated with activity in the anterior part of right superior frontal gyrus (Z = 3.49, x = 18, y = 53, z = -2, 7 voxels; TABLE 3-2). I did not find any suprathreshold activity for choice to hold > choice to grip. No supra-threshold clusters were evident for the main effect of effort or interaction between choice-effort.

Effort-related activity

I next explored activity modulated by effort level for trials where participants chose to grip, *chose grip* trials, and for trials where participants chose to hold, *chose hold* trials, separately. Particularly, using a whole-brain analysis, I looked for striatal and ACC activity associated with effort information of the option. I also explored contrasts that were modulated by persistence trait.

Table 3-2 MNI coordinates of regions the activity of which is correlated with choice (thresholded at $p = 0.001$, unc., > 5 voxels).

Region	Nearest Brodmann Areas	Coordinates (mm) x	y	z	Z value	No. of voxels	P
Contrast: Choice to Grip > Choice to Hold							
Superior Frontal Gyrus	10	+18	+53	-2	3.49	7	.0001 (unc.)
Middle Parietal Lobe	7, 19	+21	-52	+25	3.32	5	.0001 (unc.)

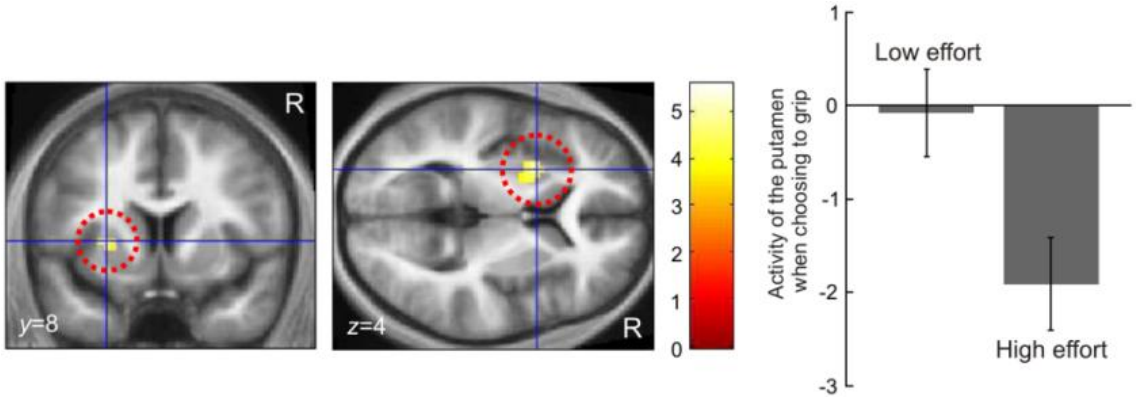

Figure 3-8 Activity in left putamen is higher when participants chose to grip an option which involved low effort than when they chose to grip an option which involved high effort (cluster-corrected FWE p <.01, 51 voxels). Bar graph depicts the parameter estimates for this contrast for visual illustration. Figure taken from Kurniawan et al., JNeurophysiol, 2010, Am Physiol Soc, used with permission.

Table 3-3 MNI coordinates of regions the activity of which is correlated with effort (thresholded at $p = 0.001$, unc., > 5 voxels).

Region	Nearest Brodmann Areas	Coordinates (mm) x	y	z	Z value	No. of Voxels	P
Contrast: GripLE > GripHE							
Putamen	N/A	-27	+8	+4	4.04	51	.01 (corr.)
Putamen	N/A	-21	+20	-2	3.81		
Primary Somatosensory Cortex	1	-57	-19	+43	3.64	13	.0001 (unc.)
Primary Motor Cortex	4p	-33	-19	+49	3.60	32	.0001 (unc.)
Primary Somatosensory Cortex	3b	-42	-25	+49	3.32		
Cingulate Motor Area	23, 24	+12	-28	+46	3.51	6	.0001 (unc.)
Supplementary Motor Area	6, 4a	+3	-16	+55	3.31	8	.0001 (unc.)
Supramarginal Gyrus	7, 40	-51	-40	+34	3.61	12	.0001 (unc.)
Supramarginal Gyrus	7, 40	-45	-43	+28	3.34		.
Middle Temporal Gyrus	39	-57	-52	+19	3.36	8	.0001 (unc.)
Contrast: HoldHE > HoldLE							
Mid-brain	N/A	+9	-25	-8	3.57	7	.0001 (unc.)
Putamen	N/A	-33	-13	-5	3.39	5	
Mid. Temporal Gyrus	37	+60	-34	-8	3.32	6	
Contrast: Persistence x HoldLE							
Anterior Cingulate Cortex	24	+3	+26	+25	3.70	11	.0001 (unc.)

Posterior part of Middle Temporal Gyrus	19	+51	-76	+13	3.40	8	.0001 (unc.)

For the *chose grip* trials I observed significant striatal activity when participants chose to grip a low compared to when they chose to grip a high effort option (gripLE > gripHE; FIGURE 3-8). This activity extended dorsally towards the caudate with a peak in the left putamen (Z = 4.04, x = -27, y = 8, z = 4, 51 voxels), and survived a more stringent threshold (cluster corrected FWE p .01). Regardless of reward level, the dorsal aspect of the putamen signaled effort information of the chosen action, with lower effort invoking greater signal. In the same contrast, I also found activity in the left motor cortex (Z = 3.6, x = -33, y = -19, z = 49, 32 voxels), right cingulate motor area (Vogt, 2005) (Z = 3.51, x = 12, y = -28, z = 46, 6 voxels), and right SMA (Z = 3.31, x = 3, y = -16, z = 55, 8 voxels). The reverse contrast (gripHE > gripLE) did not show any supra-threshold activity. The statistics of the activations are summarised in TABLE 3-3 .

For the *chose hold* trials, on the other hand, I did not find any suprathreshold activity with a contrast of trials where the rejected option involved low or high effort (holdLE > holdHE). The reverse contrast (holdHE > holdLE) yielded an enhanced activity in midbrain, in the vicinity of ventral thalamus (Z = 3.57, x = 9, y = -25, z = -8, 7 voxels; TABLE 3-3) for rejecting options with high effort compared to rejecting options with low effort.

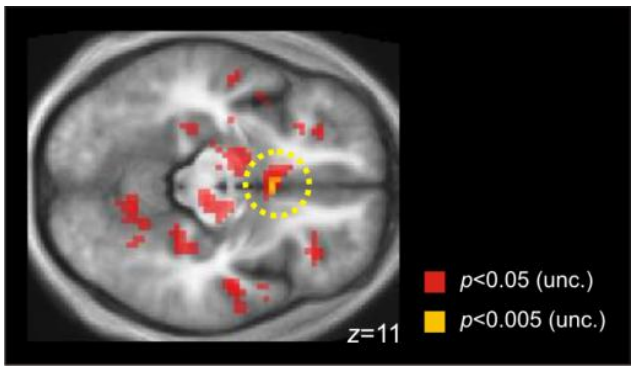

Figure 3-9 Reward level is positively correlated with activity in bilateral nucleus accumbens when participants chose to grip an option which involved high effort. Activation displayed in pink is thresholded at p < .005 (unc., 5 voxels), activation displayed in yellow is thresholded at p < .05 (unc., 582 voxels). Figure taken from Kurniawan et al., JNeurophysiol, 2010, Am Physiol Soc, used with permission.

Finally, I tested if persistence modulates effort-related activity, using the behavioural persistence scale as a covariate. I found no effect on activity associated with effort-related choices to grip. However I did find an effect on choices to hold, such that persistence significantly modulated activity in right dorsal ACC when participants rejected an option with low effort (Z = 3.7, x = 3, y = 26, z = 25, 11 voxels; TABLE 3-3). Thus, the more persistent a subject is, the greater the activation in dorsal ACC when rejecting an option that entailed low effort FIGURE 3-6B. This was the only significant correlation between persistence and the BOLD response to each condition.

Table 3-4 MNI coordinates of regions the activity of which is correlated with reward (thresholded at p = 0.001, unc., > 5 voxels, except for the last contrast; thresholded at p = 0.005, unc., > 5 voxels).

Region	Nearest Brodmann Areas	Coordinates (mm)			Z value	No. of voxels	*P*
		x	y	z			
Contrast: Reward x Choice to Grip > Choice to Hold							
Inferior Temporal Gyrus	37	-51	-58	-5	4.07	10	.0001 (unc.)
Supplementary Motor Area	6	-3	-19	+55	3.61	8	.0001 (unc.)
Contrast: Reward x GripHE							
Nucleus Accumbens	N/A	0	+11	-11	2.84	5	.002 (unc.)

Activity reflecting reward modulation

With reward level as a parametric modulator, I found a significant correlation with activity in the supplementary motor area (SMA) (Z = 3.61, x = -3, y = -19, z = 55, 8 voxels; TABLE 3-4) for the contrast chose grip > chose hold trials. No suprathreshold activity was found for reward modulation in other contrasts. However, driven by a strong prediction that NAc may be involved in reward processing (Knutson, Taylor, Kaufman, R. Peterson, & Glover, 2005), I lowered the threshold to p < .005 (unc., >5 voxels) and found a small, but significant cluster at the vicinity of NAc (Z = 2.84, x = 0, y = -11, z = 11, 5 voxels) that positively

correlated with reward only in trials where they opted an option with high effort (FIGURE 3-9).

3.5 General Discussion

The present chapter report behavioural data and brain activations involved in choosing an action based on physical effort.

I show that effort acts on behaviour in a manner that reflects discounting the value of an action, an effect reflected in lower ratings and lower preference for options with high effort. I also found that effort interacts with reward in influencing willingness to choose effortful gripping and participants' likeability, but this finding did not hold in the fMRI participants. Whether effort has a simple subtractive effect or a more complex interactive effect with reward is an empirical question. Most studies that attempt to investigate cost-benefit analysis assumes a simple or hyperbolic, subtractive effect (Bautista et al., 2001; Croxson et al., 2009; Prevost et al., 2010), although recently Talmi and co-workers have demonstrated how pain-reward integration, as another form of cost-benefit thinking, could be approximated by an interactive model (Talmi et al., 2009).

My findings support previous laboratory and field experiments with animals including humans, highlighting a sensitivity to action costs namely higher fixed reinforcement schedule in lever presses, weight of levers, higher metabolic requirements, longer travelling distance in foraging, a higher physical response requirement of climbing when compared to walking and hand force investment (Bautista et al., 2001; Eisenberger et al., 1989; Prevost et al., 2010; Salamone et al., 1994; J. R. Stevens et al., 2005; Walton et al., 2009, 2006). My behavioural findings also accord with a human observation study of pedestrian walking efficiency (Bitgood, 2006).

I expected involvement of ACC in the choice for effortful behaviour. For example, rodent experiments report that rats who expend effort for a larger gain preoperatively, choose an effortless, small reward, after a lesion in the ACC (Floresco & Ghods-sharifi, 2007; Walton, Bannerman, & Rushworth, 2002; Walton et al., 2009). Monkey single-cell recordings (Kennerley et al., 2009) and human imaging experiments with passive action valuation (Croxson et al., 2009) or mental load (Botvinick et al., 2009) also report enhanced ACC activity with increasing

effort. I contrast my study with that by Prevost et al. (2010) who found ACC is associated with subjective value (roughly the inverse of effort). Both studies utilised the investment of physical effort and examined the process of deciding whether the offered costly action is worth more than a default option. Nevertheless an important distinction is in the possibility that BOLD signal found in ACC during 'choice' (i.e. cue presentation) in their study reflects a mixture of a more abstract effort computation and a motor anticipatory process. It is true that motor anticipation could importantly contribute to the more abstract value comparison or effort integration putatively involved when making a choice. Nevertheless through the separation of choice and execution events, my study explicitly intended to decontaminate motoric processes from a purely abstract decision process. Perhaps ACC is important in a situation where the decision to act is accompanied by an imminent execution of the selected action, more ubiquitously found in previous work.

I found an association between persistence and the effect of effort on choice, which suggests the task captures a tendency to persist in everyday tasks thus strengthening interpretability and generalisability. Although this correlation could be driven by other traits, such as obedience to experimenter or social desirability, there are good reasons to think otherwise. In the task, participants knew that the experimenters could not see their actual choices during the experiment, and this is likely to eliminate desirability biases. Moreover, the correlation with persistence was selective to the effect of effort, not to reward effect, nor did it correlate with effort-reward interaction. Nevertheless, the generalisability of the task is subject to further testing.

Overall, choice and reward recruited frontal circuitries. I observed a modulation of activity for reward in the supplementary motor area (SMA) when participants opted to grip regardless of actual effort levels. SMA region has been previously implicated in movement planning (Shima & Tanji, 1998), which suggests that the choice to grip may evoke a representation of the outcome of the chosen action, which in these instances is correlated with reward expectation.

Croxson et al (2009) identified activity in the striatum, including the putamen, corresponding with net value (cost in terms of time and effort divided by reward) of an upcoming action. This led me to hypothesise involvement of striatum in effort-based choices in humans. I designed the experiment such that motor preparatory activity did not contaminate BOLD response during choice events (see

FIGURE 3-3). Notably, I found that the putamen was more active during anticipation of low relative to high effort, a finding that argues against traditional views of the putamen as being solely involved in pure motoric aspects of movement execution (e.g., Marchand et al., 2008; Prodoehl, Corcos, & Vaillancourt, 2009), and instead points to a role in higher order aspects of action valuation (Tobler, O'Doherty, Dolan, & W. Schultz, 2007) that in my study pertains to a consideration of effort cost.

Previous rodent studies provide evidence for involvement of NAc (ventral striatum) in effort-related responses (Salamone, Correa, Farrar, & Mingote, 2007). A direct comparison of the regional anatomy of the striatum is difficult between humans and rodents. In humans there is good evidence of anatomical and functional dissociation between dorsal (dorsal caudate-putamen) and ventral (NAc, ventral putamen/caudate and olfactory tubercle) (e.g., O'Doherty et al., 2004), but the connectivity of dorsal and ventral striatum share a similar parallel organisation (Haber, Fudge, & McFarland, 2000). The dorsal striatum has a stronger role in action learning and choice (as compared to passive prediction), which is a central way that effort impacts upon behaviour in my task. Croxson and co-workers (2009) found a large cluster of activation spanning across the dorsolateral and ventromedial aspects of the striatum that correlated with the net value of an upcoming action, consistent with the notion that broad regions of the striatum may be sensitive to the cost of an action. My finding of involvement of putamen along with previous work, provide converging evidence that the striatum is implicated in effort-related choices in human and across species.

An important caveat to the interpretation of putamen activity as related to economic cost is that I do not see positive activity related to financial reward per se in this region. First, my imaging analysis was not designed to assess a simple difference in activity for high versus low reward. Second, I failed to identify a significant modulation of reward in the effort contrasts despite many previous demonstrations elsewhere for reward-related activity in this region (Croxson et al., 2009; Knutson et al., 2005; Pessiglione et al., 2007; Schmidt et al., 2008). One possible, and intriguing, explanation for this failure is that it may relate to a relative lack of salience of reward, as compared to effort, in the task. Even so, high reward still had a strong effect on behaviour in the task, and modulated brain activity for other contrasts in the SMA and in the NAc.

Choosing to make a physical effort in my study reflects a critical evaluation of whether an action is worth taking, a pertinent cognitive process that may be lacking in DA-depleted conditions such as Parkinson's disease and apathy. Indeed evidence in rodents suggests that DA antagonism biases preference away from expending effort for a larger gain after controlling for time effects (Floresco et al., 2008). This evaluation also captures an individual propensity to persist through daily challenges.

Chapter 4 Pavlovian effects on learning (study 3)

Abstract

People and non-human animals tend to be active when attaining rewards and when withdrawing from punishments. How difficult is it for people to learn to overcome this tendency through learning, and to learn to be active in order to avoid punishment, or to learn to withdraw to obtain a reward? This experiment explores this question in the context of vigour of actions, i.e., where activation corresponds to a strong response (i.e., squeezing a gripper) and withdrawal corresponds to a weak response (i.e., releasing a gripper). A factorial design, in which reward-punishment was crossed with squeeze-release, showed that learning is poorer when people are attempting to overcome their 'natural' tendency to invigorate towards rewards and withdraw from punishments. Specifically, we observed i) worse performance to squeeze to avoid punishments than to obtain rewards and ii) worse learning to release to obtain rewards than to avoid punishments. These results can be modelled using a reinforcement learning mechanism, combined with prior biases arising from Pavlovian mechanisms. We discuss the nature of actions and how they relate to the nature of affective outcomes. The data speak to a wider conceptualisation of influences on motivated action beyond that of minimal behavioural activation.

4.1 Introduction

Understanding motivated, effortful, behaviour calls for separating actions from their reinforcers. In principle, an orthogonality exists between the valence of a potential outcome and the nature of action required to realise that outcome (Boureau & Dayan, 2011). This orthogonality (FIGURE 4-1) implies two separate continua; stretching from top to bottom on the ordinate is behavioural invigoration / inhibition axis and from right to left on the abscissa is axis for appetitive (reward) / aversive (punishment) outcomes.

Previous work has suggested asymmetric couplings between invigoration and outcome, such that, actions that fall anywhere in the top-right and bottom-left quadrants (invigoration/appetitive and inhibition/aversive couplings) seem to be more strongly and readily established than do actions that fall under inhibition/appetitive and invigoration/aversive pairs. Such an action architecture facilitates efficiency in learning the correct action so as to produce the appropriately valenced outcome. For example, an appetitive stimulus requires some behavioural activation in order to acquire it, whereas an aversive stimulus if distal, usually requires withholding an action in order evade it. My aim was to carve this affect-effect architecture to include actions that entail an expenditure of vigour.

Traditionally, classical conditioning paradigms (i.e., stimulus-outcome pairing) involve pavlovian indices which typically imply engagement of appetitive and aversive systems as responses to rewarding and punishing stimuli, respectively. Operant, instrumental conditioning paradigms (i.e., response-outcome pairing) typically involve behavioural engagement such as an approach response, and behavioural withdrawal such as avoid or withhold response. This predicts an overlap and potential interference between pavlovian and instrumental indices reflecting behavioural biases for certain actions and outcomes (Boureau & Dayan, 2011).

In a classical conditioning paradigm where an illuminated key predicts reward delivery, pigeons spontaneously peck on the illuminated key and move towards this conditioned stimulus (CS). In contrast, pigeons peck less on CS-, which predicts reward omission, and move away from the CS- (Wasserman, Franklin, & Hearst, 1974). Note this paradigm does not require pigeons to make

any response to receive rewards, and the only relevant action is to approach the food outlet to consume the grains of food. Nevertheless the appetitive system engaged by reward conditioning seems to spill over and engages the instrumental system, even inappropriately. This is more clearly demonstrated in a pavlovian-instrumental transfer (PIT) paradigm where learned stimulus-outcome associations (pavlovian) interfere with previously learned stimulus-response associations (instrumental). In these paradigms, after separate pavlovian and instrumental conditioning sessions, subjects are given an extinction instrumental block, along with presentations of the pavlovian CSs previously associated with outcomes. What is generally found in extinction block is that subjects tend to emit an action that was paired with an outcome more often when the pavlovian CS+ for that outcome was presented, even though the CS was never paired with the action.

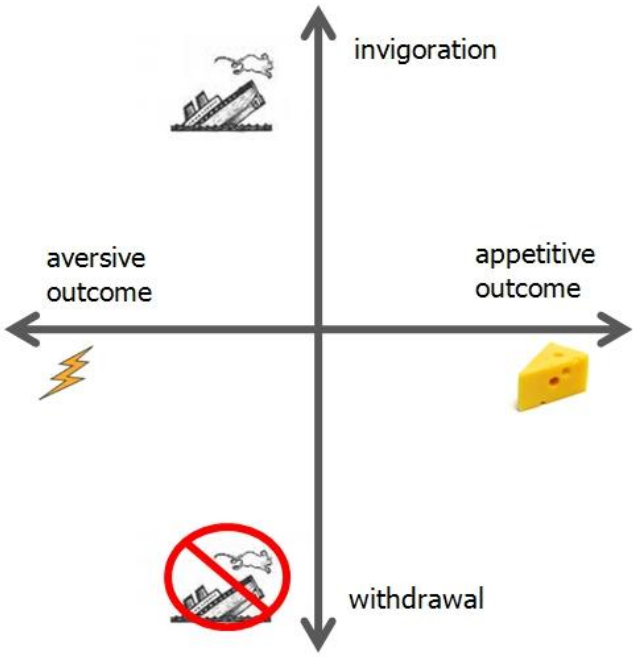

Figure 4-1 The affect-effect plot (adapted from Boureau & Dayan, 2011) appropriated for vigorous actions. Vertical axis contains an activation spectrum from behavioural withdrawal (bottom) to behavioural invigoration exemplified in squeezing. Horizontal axis contains affective results from rewarding (right) to punishing outcomes (left).

This spillover between pavlovian and instrumental indices seems to be specific to the valence of both outcomes (e.g. appetitive) and actions (e.g. approach). In a conditioning paradigm where food is only obtainable when one locomotes away from it, chicks fail to acquire this conditioned avoid response and instead keep approaching the food (Hershberger, 1986). In addition, in an atypical instrumental

contingency where a cue indicates that pecking would omit reward delivery, pigeons are unable to withhold such an approach response. Instead, they show persistent, ineffective pecking even for as long as 15 days (D. R. Williams & H. Williams, 1969). This demonstrates that an existing stimulus-reward association between a CS (illuminated key) and unconditioned stimulus (US) such as food is important in generating this maladaptive persistent pecking, even though the response-reward association counters the stimulus-reward association. This very 'pavlovian' stimulus-outcome association may facilitate/interfere instrumental learning conditioning where an association between response-outcome is formed in a valence-compatible or -incompatible manner.

More recently, Guitart-Masip and colleagues (Guitart-Masip et al., 2011) demonstrated in humans that when anticipating appetitive outcomes such as winning money, participants are more prepared to emit ('go') than withhold ('nogo') a response, whereas in anticipation of aversive outcomes such as losing money participants are better at withholding than emitting a response.

Moreover, actions that are instrumental in attaining reward or evading punishment often require substantial effort. I extended the behavioural activation axis to actions with either effort expenditure or withdrawal. In this task, participants underwent an instrumental conditioning paradigm where they learn trial-by-trial to discriminate four cues predicting effort expenditure (*squeeze*) or withdrawal (*withdraw*) to either obtain money (*win*) or avoid losing money (*avoid loss*). I use a similar action by outcome valence design to that of Guitart-Masip et al., (2011). In a *squeeze* condition, participants do not simply emit a minimal effort response but a vigorous one where 80% maximum force is expended. Instead of simply withholding an action, in *withdraw* condition participants intentionally reduce the amount of effort currently expended to minimal level.

If vigorous action learning is governed by pavlovian influence of outcomes, I predict an interaction between action and outcome valence. On the other hand, if there are equal pavlovian influences on instrumental effort, I predict a simple effect of costly effort, that is learning would be more efficient when participants *withdrew* than expend effort.

4.2 Method

Participants, stimuli and procedure

Nineteen participants (12 females, mean age = 25 (4.70) years) were recruited through the psychology participant database at UCL. Six participants used reward and punishment values of £1 and -£1 hypothetical money, whereas the rest used values of £0.20 and -£0.20 real money. No difference was found between these groups, thus I collapsed across these groups in all analyses. All participants were paid £5 -10 (mostly based on performance). The study was approved by the UCL ethics committee.

Cues were four fractal stimuli, sub-imposed behind the squeeze stimulus at the centre of screen. Each fractal stimulus was randomly assigned to one of four contingencies, crossing between action (squeeze vs. withdraw) and valence (win vs. avoid losing).

Before the learning task, participants completed two blocks of twelve training trials using a practice image, which is a different image from the experimental stimuli. During training, subjects were instructed to perform each of *squeeze* and *withdraw* responses twelve times. No outcome feedback was presented. The aim was to train participants to make both responses that satisfy the criteria within 1500 msec of cue onset. The criterion for a *squeeze* response was to reach 80% maximum force at least once, within 1500 msec, the criterion for a *withdraw* response was to release grip force to minimum level.

In each experimental trial I present one of four stimuli which participants learned trial-by-trial to either squeeze or withdraw, where each correct action yielded either a positive outcome 80% of the time or none 20% of the time (win condition), or none 80% of the time or a negative outcome 20% of the time (avoid loss condition). Incorrect action always yielded nothing in win (and a negative outcome in avoid loss) trials.

Overall, participants underwent 60 continuous, fully-randomised repetitions of four conditions, presented with a 5-seconds rest every twelve trials and a 15-seconds rest after the first and second 100 trials. Overall, there were 240 trials.

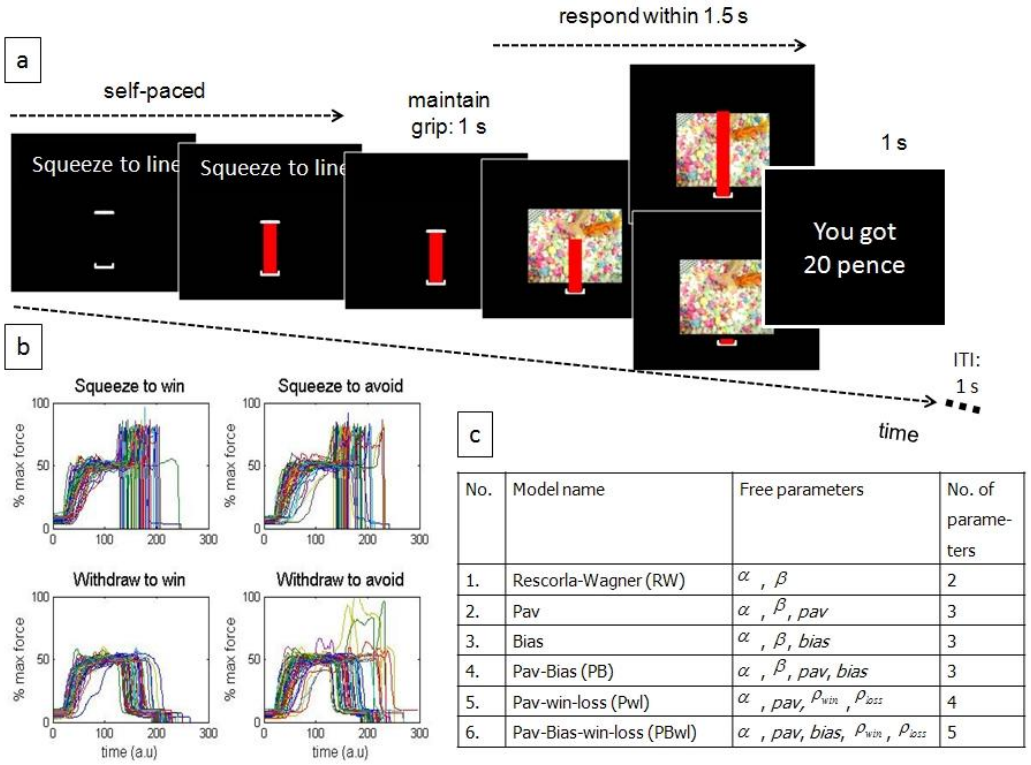

No.	Model name	Free parameters	No. of parame-ters
1.	Rescorla-Wagner (RW)	α , β	2
2.	Pav	α , β, *pav*	3
3.	Bias	α , β, *bias*	3
4.	Pav-Bias (PB)	α , β, *pav, bias*	3
5.	Pav-win-loss (Pwl)	α , *pav*, ρ_{win} , ρ_{loss}	4
6.	Pav-Bias-win-loss (PBwl)	α , *pav, bias*, ρ_{win} , ρ_{loss}	5

Figure 4-2 a) A schematic of a single trial. Participants self-initiated a trial by squeezing a hand grip to move a visual red bar stimulus to reach a 50% force line on screen and stay on that line for 1 sec. Immediately following this, one of four cues is presented which gave participants 1500 msec to respond either by squeezing to reach the top level (80%), or by completely withdrawing force. The monetary outcome is presented at the end of 1500 msec. Correct responses are probabilistically rewarded (gain in win condition and zero in avoid loss condition), whilst incorrect responses are never rewarded (zero in win condition and loss in avoid loss condition). b) Grip data for one participant in all conditions. Each line represents one trial with time on x-axis from onset of 'Squeeze to line' screen until end of grip period. Participants start around zero force with natural noise in the grip device, they then voluntarily reaches the 50% force and maintains for 1 sec to initiate cue, followed by either a squeeze to 80% force (top) or force withdrawal (bottom). As seen here, this participant made more incorrect squeeze responses when they were supposed to withdraw to avoid loss (bottom right); a.u. = arbitrary units. c) A list of models used to describe the hypothesized underlying decision processes and their free parameters.

As seen in FIGURE 4-2A, participants self-initiated a trial by reaching a horizontal line on screen (individually-calibrated as 50% of max force) using a hand grip. Participants had to maintain force at this line for 1 sec. If participants under- or over-squeezed before 1 sec had elapsed, the trial was aborted and re-started, which happened < 5% of all trials in all participants, except in one subject who had aborted trials 26% of the time. After this level had been maintained for 1 sec, one of the four fractal cues was randomly presented which prompts participants to choose either to make a *squeeze* or *withdraw* response. There was a blank screen for 1 sec, followed by a message indicating the monetary outcome of that trial for 1 sec. ITI

was 1 sec. Participants saw the total outcome at the end of experiment and were paid a maximum of £5 based on their total earning.

As an example, FIGURE 4-2B shows grip data for one participant in all conditions. Each line represent gripping trajectory across time starting from onset of instruction screen to self-initiate trial until offset of squeeze stimulus. Grip value starts around zero, increases and stays at 50% level for 1 sec followed by either a *squeeze* (80%) or *withdraw* (around 0%) response to the cue. This participant was able to discriminate the cues and only made few incorrect responses when required to withdraw to avoid loss; overall accuracy for this participant was above 93%.

Modelling

I modified a standard reinforcement learning model to capture my behavioural data as previously used to model pavlovian to instrumental interactions (Guitart-Masip, Talmi, & Dolan, 2010; Huys et al., 2011). I first describe the model with the lowest Bayesian Information Criterion (BIC) value (Schwarz, 1978), model PBwl (pav-bias-win-loss) and then the alternative models (FIGURE 4-2C).

Let s_t be stimulus presented at trial *t*, and a_t be the action (choice) on that trial. An action can be one of two types: squeeze or withdraw. Let also r_t be the reinforcement obtained, either positive or zero in win, or zero or negative in avoid loss. I define the action weight, $W_t(a_t)$ for each action. The weight of a *withdraw* action, $W_t(withdraw_t)$ is equal to the action value $Q_t(s_t, withdraw_t)$ associated with the action *withdraw* in the presence of stimulus s_t. The weight of a *squeeze* action $W_t(squeeze_t)$, is an update of i) the action value Q_t associated with the *squeeze* action $Q_t(s_t, squeeze_t)$ in the presence of stimulus s_t, and ii) a fixed, time-invariant Pavlovian term, *pav*, multiplied by stimulus value V_t associated with stimulus s_t, allowing positive outcomes to boost action value for *squeeze* and negative outcomes to damp down action value for *squeeze*, and iii) a fixed, time-invariant bias term, *bias,* boosting a *squeeze* action, constrained in positive numbers.

$$W_t(withdraw_t) = Q_t(s_t, withdraw_t) \quad (1)$$

$$W_t(squeeze_t) = Q_t(s_t, withdraw_t) + pav.V_t(s_t) + bias \quad (2)$$

I assume participants' choice are based on some comparison between the action weights with some stochasticity, based on a softmax distribution, such that probability of choosing say a *squeeze* action given stimulus s_t, $p(squeeze_t \mid s_t)$ for the winning model is:

$$p(squeeze_t \mid s_t) = \frac{1}{1+e^{-(W_t(withdraw)-W_t(squeeze))}} \quad (2)$$

Outcome was always immediately following the action, thus a Rescorla-Wagner rule was applied to compute the expectations with a fixed learning rate constrained between 0 and 1, α. I only update the values associated with the chosen action and use the same learning rate for each individual to update action and stimulus values Q_t and V_t. As implemented in Huys et al. (2011), the immediate, intrinsic value of rewards and punishments may be different, so I added two outcome sensitivity parameters, ρ_{win} and ρ_{loss}. I used these terms to update the action and stimulus values.

$$Q_{t+1}(s_t, a_t) = Q_t(s_t, a_t) + \alpha.(\rho_t.r_t - Q_t(s_t, a_t)) \quad (3)$$

$$V_{t+1}(s_t) = V_t(s_t) + \alpha.(\rho_t.r_t - V_t(s_t)) \quad (4)$$

Where ρ_t is ρ_{win} if r_t >0 and ρ_t is ρ_{loss} if r_t <0

To find the optimal solution for these free parameters, I conducted nonlinear optimisation which calculates the smallest negative log likelihood function of choice (akin to maximum likelihood estimate (MLE) (Daw, 2009) using iterations with 30 different starting points.

Alternative models are as follows: Model Pwl (pav-win-loss; model 5) assumes no *bias* for squeezing. Model Pav (model 2) assumes no *bias* and no different outcome sensitivity for win and loss, ρ_{win}, ρ_{loss}. Instead, model Pav included β as the slope of softmax function to govern choice probability (constrained as positive numbers). Note that replacing ρ with β is mathematically equivalent; modifying the sensitivity to outcome is simply changing the scale of the function to be more stretched/ dispersed (as rho gets higher, there is less stochasticity) and this gives the same effect as changing the slope of the function. Thus $p(squeeze_t \mid s_t)$ for model Pav is:

$$p(squeeze_t \mid s_t) = \frac{1}{1+e^{-\beta(W_t(withdraw)-W_t(squeeze))}} \quad (5)$$

Model PB (pav-bias; model 4) is an extension of model Pav (model 2) with a *bias* for squeezing. Model Bias (model 3) assumes no *pav* term, whereas model RW (Rescorla-Wagner; model 1) does not assume *pav* nor *bias* terms.

I report negative log likelihoods (-LL; lower values indicate better fit of the model), both pure and penalised for number of free parameters (BIC). I also report a pseudo-r^2 statistic (Camerer & Ho, 1999), defined as (r - l)/r where r and l are, the log values of data likelihood under the model and under purely random choices (0.5 for each trial) (TABLE 4-1).

4.3 Results

Raw data: p(correct)

Learning was evident as revealed in a 2 x 2 x 6 (Action x Valence x Block) repeated measures ANOVA on proportion of correct responses. There was a significant main effect of block, $F(3.05,54.96) = 20.22$, $p < .0001$, $\eta_p^2 = .52$ and a significant action by valence interaction, $F(1,18) = 7.70$, $p = .012$, $\eta_p^2 = .30$. There was a consistent block-by-block performance improvement, ps < .005 (FIGURE 4-3A). I also found a marginally significant action by valence by block interaction at the transition into the 4th block, $p = .058$.

To find out what drives this (marginally) significant three-way interaction, I averaged over each subject's proportion of correct trials across the first three blocks and compared these against trials in block 4 to test for learning in two-way Valence x Block (block 1-3 vs. 4) ANOVAs at *squeeze* and *withdraw* actions separately (FIGURE 4-3B).

In *squeeze* trials, main effects of valence and block were significant. Participants were significantly more accurate in win than avoid loss trials, regardless of learning, $F(1,18) = 9.21$, $p = .007$, $\eta_p^2 = .34$ and they did better in block 4 than in previous blocks, regardless of valence, $F(1,18) = 18.60$, $p < .0001$, $\eta_p^2 = .50$. Valence by block interaction was not significant, $F < .8$, $p > .40$.

In *withdraw* conditions, there was a significant main effect of block, there were more correct trials in block 4 than there were in previous blocks, $F(1,18) = 13.61$, $p = .002$, $\eta_p^2 = .43$, and a significant Valence x Block interaction, $F(1,18) =$

6.16, $p = .023$, $\eta_p^2 = .25$, Valence effect was non-significant, $F < 2$, $p > .18$. Following up the interaction, I found a significant simple effect of block in avoid loss, but not in win trials: subjects significantly improved their performance to make a withdraw response to avoid loss, $t(18) = 4.57$, $p < .0001$, but not to win, $t < 1$, $p = .37$.

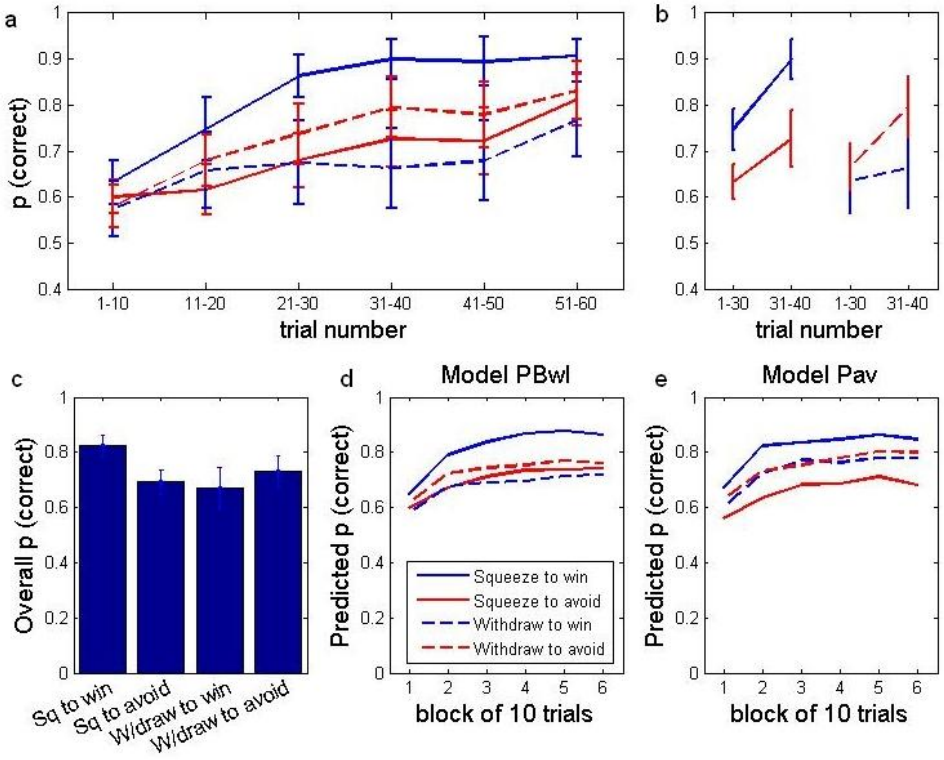

Figure 4-3. Group-averaged proportion of successful trials, observed across six blocks (a), with follow-up tests between block 1-3 vs. 4 (b), overall (c) and predicted by model PBwl (d) and Pav (e). a) There was significant performance improvement at each block, regardless of action and valence. I found a significant action by valence interaction and a marginally significant action by valence by block interaction. b) Separate follow-up Valence x Block ANOVAs for squeeze and withdraw trials show that squeeze performance improved in both win and avoid loss trials (valence by block interaction was n.s.). Withdraw performance only improved in avoid loss, but not in win trials, valence by block interaction was significant, $p = 0.023$. Solid lines denote squeeze actions, dotted lines denote withdraw actions; blue lines denote win, red lines denote avoid loss trials. c) Overall, there was a significant action by valence interaction, and this was driven by higher accuracy when squeezing to win than to avoid loss. d-e) Prediction of proportion of succesful trials by model PBwl and Pav. Visual inspection suggests that model PBwl gives a better match between observation and prediction than model Pav does.

The action by valence interaction I found in the learning data was also evident in overall performance. I entered the overall proportion of correct responses, into a 2 x 2 (Action x Valence) repeated-measures ANOVA and found a significant action by valence interaction, $F(1,18) = 7.76$, $p = .012$, $\eta_p^2 = .30$ (FIGURE 4-3C). This interaction was driven by a significantly better performance for squeezing to win than to avoid loss, $t(18) = 3.05$, $p = .007$. Correct *withdraw*

responses to win and to avoid loss were not significantly different. No main effect of action or valence was found, *F*s < 1.2, *p*s > .25. The learning data suggest a potential pavlovian role of outcomes in the *withdraw* domain, whilst overall performance data suggest a pavlovian role in the *squeeze* domain. Combined, these learning and overall performance results suggest a potential interdependence between action and valence where rewards preferentially facilitate squeezing, and punishments preferentially support learning to withdraw effort.

Raw data: p(stay)

I can test whether the tendency to make the same response (p(stay) is influenced by what happened in the previous trial; whether participants just made a correct response that was rewarded, a correct response that was unrewarded (due to probabilistic outcomes), or an incorrect response. To do this, I calculated the proportion of making the same response (p(stay)) at trials *t+1* and *t+2* after rewarded correct, unrewarded correct, and (unrewarded) incorrect trial *t*. I entered p(stay) into two separate 2 x 2 x 3 Action x Valence x Trial-t (rewarded/ unrewarded correct and incorrect trials) repeated measures ANOVAs.

For p(stay) at trial *t+1,* there was a main effect of trial-t, p(stay) following incorrect trials were significantly lower than following correct trials (FIGURE 4-4A, *F* values in TABLE 4-1) and a main effect of valence, participants had a stronger tendency to make the same response when trial *t* was a win trial than when it was an avoid loss. Action effect was n.s., *p* >.3. I also found an action by trial-t interaction, followed up by separate one-way ANOVAs and bonferroni t-tests for *squeeze* and *withdraw* trials (collapsed over valence trials) to see if p(stay) was influenced by the fact that trial *t* was rewarded/unrewarded/incorrect. Both one-way ANOVAs at *squeeze* and *withdraw* conditions were significant: what happened at trial *t* influences probability of making the same response at trial *t+1* . Follow-up Bonferroni t-tests show that p(stay) for *squeeze* and *withdraw* responses after correct trials were significantly higher than after an incorrect trial, and p(stay) *squeeze* following a rewarded correct trial was higher than an unrewarded correct trial, but p(stay) *withdraw* was the same following rewarded and unrewarded correct trials.

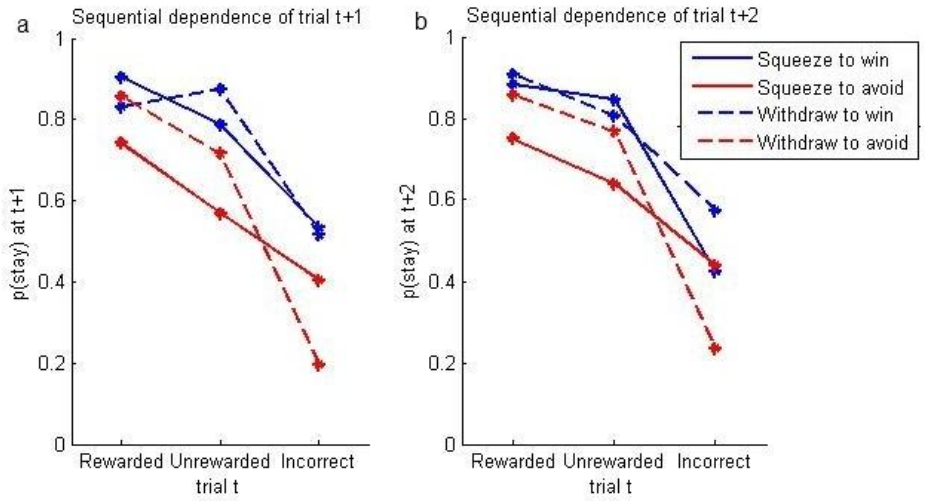

Figure 4-4. Group-averaged probability of staying at trial t+1 (a) and trial t+2 (b) for rewarded and unrewarded correct and (unrewarded) incorrect trials t. Solid lines denote squeeze actions, dotted lines denote withdraw actions; blue lines denote win trials, red lines denote avoid loss trials.

For p(stay) at trial *t+2* (FIGURE 4-4B, *F* values in

Table 4-1), there was a main effect of trial-t, p(stay) following incorrect trials were significantly lower than following correct trials, and a main effect of valence, participants had a stronger tendency to make the same response at trial *t+2* when trial *t* was a win than avoid loss trial. Action effect and interactions were n.s., $p > .3$. Note I ran this analysis on 14 participants as there were missing data for the other 5 participants. All *F* and *t* values are in TABLE 4-1.

Raw data: predicting current action

Furthermore I tested if previous actions or outcomes can be used to predict action at trial *t*. Although this analysis does not address whether there is an interaction between action and valence in a way which instantiates a pavlovian influence of outcome on actions, it may suggest the importance of participants' previous actions and outcomes. To do this, I ran three logistic regressions to predict whether action at trial t a_t, was a *squeeze* or *withdraw* response. The three models included these predictors of the same cue: a_{t-1} and a_{t-2} (action model), $outcome_{t-1}$ and $outcome_{t-2}$ (outcome model) and a_{t-1}, a_{t-2}, $outcome_{t-1}$, and $outcome_{t-2}$ (action + outcome model). I calculated regression weights for each model in each individual (group-averaged values in FIGURE 4-5), and found that in less than 50% of participants these regression weights were significant. When entered into a series of paired-samples

t-test, I could not find any significant difference between predictors in trial *t-1* and *t-2* in all three models, action and outcome predictors in action+outcome model were also not significantly different from each other, *p*s > .15.

Table 4-1 F and t values for sequential dependence analysis between trial *t* and trial *t+1 and t+2.*

No.	Name of effect	*F* or *t* values	*p* values	Partial eta sq
Action x Valence x Trial-t (rewarded correct/unrewarded correct /incorrect) on p(stay) at trial t+1				
1.	Correct > incorrect (main effect of trial-t)	*F* (2,7) = 52.32	< .00001	.93
2.	Win > avoid loss (main effect of valence)	*F* (1,8) = 47.23	.0001	.85
3.	Action by trial-t interaction	*F* (2,7) = 6.59	.02	.65
4.	One-way ANOVA of trial-t in *squeeze* condition	*F* (2,12) = 14.91	.001	.71
5.	One-way ANOVA of trial-t in *withdraw* condition	*F* (2,12) = 16.83	<.0001	.73
6.	Correct>incorrect in *squeeze* condition	*t* (13)=5.06	<.0001	
7.	Rewarded correct >unrewarded correct in *squeeze* condition	*t* (13)=2.261	.042	
8.	Correct>incorrect in *withdraw* condition	*t* (13)=4.65	<.0001	
9.	Rewarded correct >unrewarded correct in *withdraw* condition	*t* (13)=1.22	.24	
Action x Valence x Trial-t (rewarded correct/unrewarded correct /incorrect) on p(stay) at trial t+2				
10.	Correct > incorrect (main effect of trial-t)	*F* (2,7) = 110.66	< .00001	.97
11.	Win > avoid loss (main effect of valence)	*F* (1,8) = 11.26	.009	.58

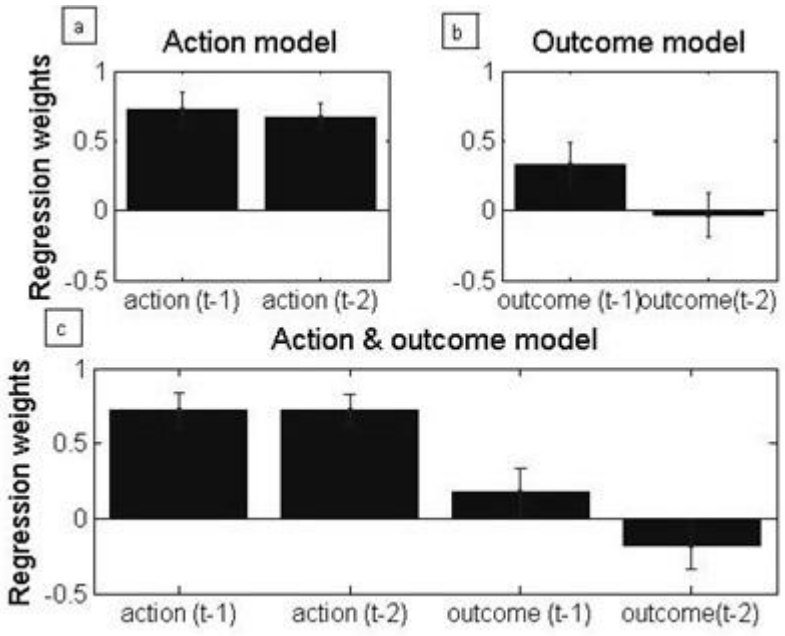

Figure 4-5. Group-averaged regression weights for predicting action at trial t, based on actions and outcomes at trial t-1 and t-2, for (a) action, (b) outcome, and (c) action+outcome models. Paired-samples t-tests revealed no significance between t-1 and t-2.

Modelling results

Behaviourally I found that state values have pavlovian influence on i) block-by-block learning, such that learning is more evident when participants withdrew effort to avoid loss than to win, and ii) overall performance, such that performance is better when squeezing to win than to avoid loss (FIGURE 4-3). To begin describing the underlying decision processes for my choice data, and to specifically test for a pavlovian influence in effort-related actions, I built and ran 6 models, each either assumes or does not assume a pavlovian influence of outcome on action (*pav* term). In addition I also vary whether a model has a bias for squeezing, and whether it has different outcome sensitivity for wins and losses.

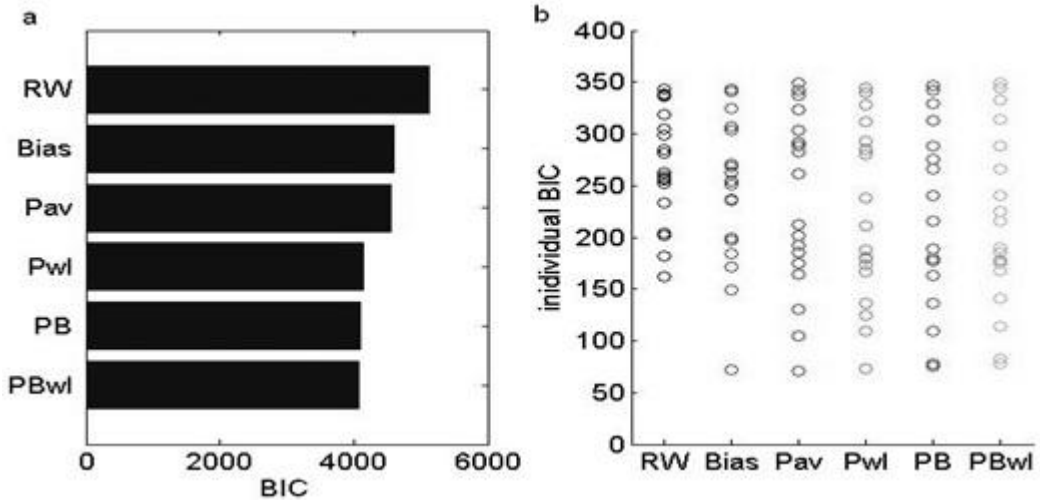

Figure 4-6. Bayesian Information Criterion (BIC) scores for 6 models summed across participants (a) and individually plotted (b).

Model PBwl (pav-bias-win-loss) is the winning model, showing the lowest BIC score than other models (FIGURE 4-6A). This model qualitatively supports the action by valence interaction I found behaviourally, but also includes a *bias* term and different reward and loss sensitivity parameters. While model Pav did not show a low BIC score, a random effects, group level Bayesian model selection (BMS) procedure revealed that model Pav and model PBwl show equally larger exceedance probabilities, such that these models are more likely than others (exceedance probability of .40 and .49, respectively). This BMS procedure takes into account individual BIC values and distribution of BIC values (Stephan, Penny, Daunizeau, Moran, & Friston, 2009).

Inspection of individual BIC scores for both models (FIGURE 4-6B) helps explain why model Pav seem to approach PBwl despite its high BIC score. Model Pav seems to divide participants into two groups, those with low and high BIC scores, whereas model PBwl seem to be more normally distributed with high frequency around the lower values. However, the split point based on BIC scores for model Pav does not correspond to my learner/nonlearner split using behavioural data (non-learner: performance under 75% correct in the last 20 trials of at least 1 condition) described below. TABLE 4-2 shows how model PBwl has better quality of behavioural fits than model Pav does.

Table 4-2 Poorer quality of behavioural fits to 4,560 choices from 19 participants for model Pav than those for model PBwl, shown with negative log likelihood (-LL), pseudo-r2, and Bayesian Information Criterion (BIC).

	Model Pav	Model PBwl
-LL	2122.93	1772.61
Pseudo-r^2	0.328345	0.43918
Number of parameters	3	5
BIC	4558.26	4065.88

A likelihood ratio to compare model PB (pav-bias) with its nested model PBwl (pav-bias-win-loss) showed a mean log likelihood ratio of 8.4712 which significantly favours model PBwl, $p < .05$ with a chi-square cumulative distribution test. This model is also significantly better than model Pwl (pav-win-loss), mean log likelihood ratio = 5.5280, $p < .05$. Model PB (pav-bias) and Pwl (pav-win-loss) did not differ significantly with each other.

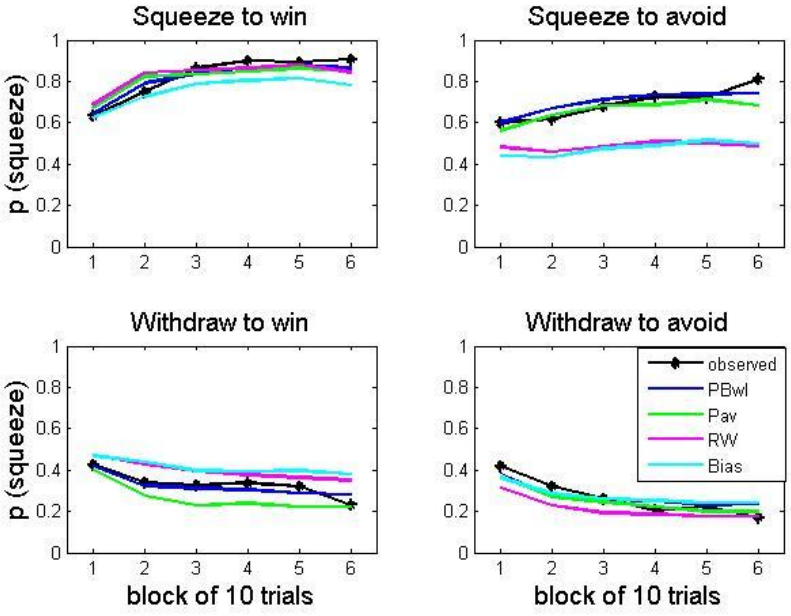

Figure 4-7 Model predictions for probability of effort deployment (p(squeeze)) in 6 blocks of 10 trials by 4 different models: model RW, Pav, Bias, and PBwl in each condition. Black line denotes observed data. Particularly, the winning model PBwl seems to match data better than other models (Pav, RW & Bias) in squeeze to avoid and withdraw to win conditions.

Surrogate data based on these known decision processes provide some qualitative support for model PBwl (FIGURE 4-3D-E, FIGURE 4-7). Visual inspection suggests that all models seem be able to predict *squeeze to win* and *withdraw to*

avoid conditions as well as the winning model PBwl does, but the three models seem to do poorly in predicting *squeeze to avoid loss* and *withdraw to win* conditions. These two conditions contain the conflict between pavlovian tendencies of positive, approach actions and negative, inhibitory actions.

Relation between parameter estimates and individual differences in learning

I next sought if estimates from model PBwl are different in participants who learned and did not learn the task. I categorised a non-learner if s/he did not attain 75% accuracy in the last 20 trials of 1 of 4 conditions. I realise this is an arbitrary cut-off, nonetheless it is one way to categorise if a subject managed to make correct responses at the end of experiment. This cutoff split the group into 10 learners and 9 non-learners. Independent-samples t-tests revealed no significant difference in all model parameters (α , *pav*, *bias* , ρ_{win} and ρ_{loss}), even after excluding several participants whose estimates did not reach reasonable values (2 non-learners, value > 100), *p*s > .1.

4.4 Discussion

I show a pavlovian influence of affective outcomes on instrumental, vigorous actions. Here, I orthogonalised affect and effect to reveal asymmetric associations between invigoration and appetitive/aversive outcomes. Specifically, this study is the first to extend the activation/inhibition axis into an effort expenditure/ withdrawal spectrum and to use reinforcement learning concepts to describe decision processes that may underlie such asymmetries.

Behaviourally, I found evidence for differential pavlovian effects of outcomes on actions. That is, expending effort seems to have an appetitive advantage, whereas effort withdrawal seems to benefit from avoidance from punishment. At the core of this asymmetry is a pavlovian notion that, on the one hand, the appetitive system which more readily associates rewards to a neutral stimulus would also more readily associate an invigorated action to the same neutral stimulus. On the other hand, the aversive system which more readily associates

punishment to a neutral stimulus would invoke a stronger association between a neutral stimulus with a withdrawal of effort.

I instantiated this pavlovian idea into the models by adding the stimulus value associated with outcomes to the action value of expending effort. In the avoid loss conditions, even though participants made all correct choices, the running average of stimulus values would still be negative as they would receive money loss in a fraction of the trials. In other words, when a cue is mostly associated with rewards (win conditions), its stimulus value would be added to the action value for expending effort, whereas even when a cue is only occasionally associated with punishments and with a neutral outcome otherwise (avoid loss conditions), its stimulus value would be subtracted from the action value of expending effort. Model PBwl (pav-bias-win-loss) specifies this pavlovian term and it shows the best evidence for the observed data. In addition, the winning model also specifies a squeezing bias and allows for differential outcome sensitivity which will be discussed below.

Pavlovian term

I discussed some kind of spillover between pavlovian and instrumental indices at the start of this chapter. In situations where expending effort is required to earn reward (*squeeze to win*), the stimulus-outcome/action-outcome spillover facilitates learning and correct choices. Likewise, in situations where punishment cessation is achieved through effort withdrawal (*withdraw to avoid loss*), this spillover faciliates correct responding. In contrast, when expending effort is required to avoid punishment (*squeeze to avoid loss*) and withdrawal gains reward (*withdraw to win*), this spillover may impede learning the correct response. The *pav* term addresses this spillover and is evident in how well the models' surrogate data match the observed data. Compared to other models (RW, Bias, Pav), surrogate data of model PBwl show better prediction for block-by-block learning data in the latter two conditions.

The model assumes an incremental effect of pavlovian influence and bias on action value for squeezing (eq. 2). I found no evidence for a squeezing bias (except for 1 subject). In principle, though, the pavlovian term could interact with action value and have a multiplicative effect, such that in cases where pavlovian effect is weak, it causes diminution of action value, whereas in cases where pavlovian effect

is large, it exagerrates the value of squeezing. Nevertheless, assuming an additive effect is more parsimonious; thus I used this assumption in the models. Further work may address whether a multiplicative effect could explain data better.

Bias for squeezing

As mentioned above, I found no evidence for a squeezing bias. Nonetheless two subjects' bias data are noteworthy. First, I estimated an extremely high bias parameter for one subject who never learnt to withdraw. This became a sanity check that the bias term was instantiated correctly, as it strongly reflects this subject's raw data. Second, I found another subject who shows a negative bias value, but has raw data which clearly suggests a stronger sensitivity to reward than punishment, and this is confirmed by the rho values. This suggests that bias alone does not explain learning in the paradigm. Indeed, in go/nogo paradigms where an action has minimal effort cost, a bias for 'go' is perhaps more detectable, but this putative bias may be obscured by the fact that squeezing in this task is much more costly than a simple button press. The models which only included bias did not fit the data very well, although once *pav* and rho are included in the model, having bias seems to fit the data better than not adding it (model PBwl vs. Pwl). Future work should explore different amounts of effort costs and the extent to which effort costs do minimise an existing bias to act.

Sensitivity for reward and punishment

The effect of outcome valence can take two forms, first a main effect of valence which could manifest in outcome sensitivity parameters such as rhos for rewards and punishments, second an interaction with action which could manifest in a pavlovian term. In addition to estimating *pav*, the winning model PBwl allows for different sensitivy to reward and punishment. This is similar to participants in Huys et al. (2011) who showed greater sensitivity to reward than punishment in a task with deterministic outcomes. Here, there is better model evidence when I allow for separate updating of action and stimulus value for positive and negative outcomes (model PBwl), than when I simply specify one softmax temperature for each individual (model PB). Different outcome sensitivity may appear like a simple opponency between positive and negative outcomes, regardless of actions. Nevetherless, midbrain DA neurons have been reported to be sensitive to both

positive and negative outcomes (Matsumoto & Hikosaka, 2009), and in addition, this opponency has been elegantly expanded to conflate with action, which again points to the interaction between affect and effect and has been discussed previously to implicate neutotransmitters 5HT and DA (Daw, 2002; Guitart-Masip et al., 2011; Huys et al., 2011). Unfortunately, the current analyses are not able to pick apart the relative importance of each parameter in the model to determine which form of valence effect dominates the data.

Nature of squeeze and withdraw actions

I realise that *squeeze* and *withdraw* responses in this task are not simple analogues to approach and avoid responses, and they may not be direct extensions of response emission and inhibition either (Guitart-Masip et al., 2011). However, Huys and colleagues have conceptualised (one) go approach and (two) go withdraw actions in their analysis of pavlovian and instrumental interaction (Huys et al., 2011). I likened the binary actions to behavioural activation and inhibition. In this task, it is reasonable to assume a general preparedness to squeeze every time participants self-commenced a trial. A *squeeze* choice simply allows a release of that prepared invigoration hence a form of behavioural activation, whilst *withdrawal* may serve like an inhibition or withdrawal of that preparedness to expend effort.

In addition, the original orthogonalisation between affect and effect was driven by the notion of 'appetitive actions' and that experimental conditions involving the possibility to earn reward are gratuitously termed 'approach' actions (Niznikiewicz & Delgado, 2011). Thus, it is unclear whether an exemplar of activation such as squeezing counts as an 'appetitive action', or how such an action that is not neccesarily 'appetitive', could have a bias for reward, as was found in this study.

Moreover, one cannot fully rule out that the squeeze/appetitive and withdraw/aversive associations could be due to the visuo-motor dynamics the participants experienced in the task. Indeed, visually participants saw a red bar going up as they squeezed which might invoke a sense of gain, whereas they saw the same bar coming down as they withdrew which could feel like losing.

Influence of past trials on current choices

I made some attempt to run frequency analyses and regression models to dissect the data without any assumptions of underlying decision processes. After making a correct *squeeze* response, being rewarded in the last trial seems to facilitate sticking to the same *squeeze* choice on the current trial compared to not being rewarded. However, after making a correct *withdraw* response, subjects are as likely to stick to the same *withdraw* choice on the current trial regardless of reward delivery. This may suggest that probability of sticking to the same choice is more dependent upon reward in a *squeeze* condition, but not in a *withdraw* condition. Unfortunately, the regression models could not reach reliable results to clarify the roles of past actions and outcomes on influencing current choice. It would be useful to address this issue with a kernel analysis which still assumes some influence of previous distant trials although with more forgetting.

Several caveats and future work

I have observed pavlovian effects separately in block-by-block (in *withdrawal*) and overall performance (in *squeeze*). My discussion thus far has not necessarily treated them as separate, but it is plausible that this separate effect may reflect and be driven by different decision mechanisms. My model is not able to distinguish this. A recent attempt (not reported here) using a model comparison method (Huys et al., 2011) shows the best model evidence for a model which specifies separate pavlovian terms for squeeze and withdraw actions. I aim to carefully characterise this new modelling result in future work.

Further modelling work can use Bayesian posterior maximization (Daw, 2009) using regularisation to underweight participants who show irregular learning data. This would address the issue that the current Bayesian model selection results gave high exceedance probability to model Pav which did not have a good BIC value. I could also refine the analysis, such as running ANOVA on surrogate data, to test if predicted data by model PBwl is significantly better than that by models RW, bias, and pav (FIGURE 4-7), and assessing reaction time data and grip force data at the time of cue presentation.

To sum, I expanded the behavioural invigoration of affect-effect architecture to vigorous actions. By orthogonalising effort deployment/withdrawal and affective valence, I observed a pavlovian effect on both kinds of action and have

approximated the observed data with decision processes which involve a pavlovian effect, bias for squeezing, and separate affective sensitivity. These terms together gave the best model evidence for the data. My data gave a fresh insight into characterising the nature of actions and vigour, and how action and outcomes are associated under pavlovian influences.

Chapter 5 Modulation on outcome delivery (study 4)

Abstract

Prior to taking an action, we anticipate how much effort the action will require. After taking that action, we evaluate the affective outcomes delivered. Effort magnitude, the vigour of the action, and whether the action yields a reward or avoids a punishment, are all likely to influence action anticipation and outcome evaluation. ACC and dorsal striatum (dSTR) are known to play a role in anticipation of effort and stimulus-response association while ventromedial prefrontal cortex (vmPFC) and ventral striatum (vSTR) are known to signal outcome evaluative processes. It remains unclear if effort and outcome valence influence neural signals for outcome evaluation. Using fMRI, I conducted a cue predictive task wherein participants anticipate and execute vigorous actions and were also presented with outcomes for their actions. I manipulated actions, to entail low or high effort,and the valence of the outcome so that an action yielded a reward or avoided a punishment. When an action is anticipated, activity in ACC and dSTR is sensitive to the action's effort level but not to outcome valence. When an action has been completed, activity in vmPFC and insula is sensitive to the action's outcome valence but not to effort size. Importantly, I manipulated expectation such that participants occasionally did not receive the expected outcomes, and thus experience a negative prediction error. Here, I found dissociation in effort and valence modulation of expectation, such that activity in vSTR and vmPFC for expected outcomes is modulated by effort, while insular activity for negative prediction error (undelivered outcomes) is modulated by outcome valence. These findings confirm involvement of ACC in anticipating effort, and provide new insight into a neural modulation of effort on outcome evaluation.

5.1 Introduction

A consequence of instrumental learning is the continuation of cue responding and the maintenance of reinforcement delivery. During this response-outcome cycle, one anticipates the upcoming action, executes the action and then monitors whether the expected outcome is delivered. Indeed, once stable performance is achieved, it is assumed that there is coding of effort and outcome magnitudes, or perhaps coding of an integrated value that combines effort and outcome. Indeed, previous work has shown that presenting a cue that reliably predicts an upcoming action, which embodies effort costs and rewards, implicates ACC and striatum (Croxson et al., 2009; Gan et al., 2010; Kennerley et al., 2009). These brain regions seem to represent action parameters which could be useful in making adaptive action selection given a larger range of effortful actions available (Bautista et al., 2001).

Recent evidence suggests that appetitive and aversive outcomes may have a different relationship to action (Boureau & Dayan, 2011; Daw, 2002; Guitart-Masip et al., 2011; Huys et al., 2011). Indeed, as discussed in CHAPTER 4, I observed differential effect of rewards and punishments on effort expenditure. Here, I examine if neural anticipatory responses to an incoming action, and the neural evaluation of outcome delivery and omission, are sensitive to effort and outcome valence.

In this task, participants performed an overlearnt instrumental task where one of four cues reliably predicts a low or high effort action, and leads either to a probabilistic reward or a probabilistic avoidance of loss under a correct response (FIGURE 5-1). Using fMRI, I recorded BOLD responses to cue and outcome presentation enabling the examination of effort and outcome valence influences in action anticipation and outcome evaluation, respectively.

While most previous studies examining effort have used appetitive stimuli (Cousins & Salamone, 1994; Croxson et al., 2009; Ghods-Sharifi & Floresco, 2010; Prevost et al., 2010; Rudebeck et al., 2006), none have examined exertion of effort for active avoidance from punishment, or instantiated such active avoidance in actions with different effort sizes. I predicted involvement of ACC and striatum in effort and reward anticipation as found previously, but now tested if similar activity could be observed when avoiding punishment. Moreover, I expected

activity in regions typically responsive to either appetitive or aversive values such as orbitofrontal cortex, ventral striatum, and insula (Seymour, Singer, & Dolan, 2007), and explore any modulation by expended effort during outcome evaluation.

5.2 Method

General Task Description

The task required participants to respond to one of four cues which predicted either a low of high effort level, and either a probabilistic win or a probabilistic avoidance of loss. After exhaustive training on day 1, participants then completed this task in the scanner on day 2. At the start of each trial, participants saw a cue, then after a jittered period (see FIGURE 5-1A), based on the cue presented, they had to squeeze a hand bar to reach one of two effort targets (25% or 65% max effort; FIGURE 5-1B) within 1.5 sec. Performance was >90% correct. After receiving a visual feedback indicating that they satisfied the squeezing target and time criteria, they saw an outcome which was 20, 0, or -20 pence. Outcome probabilities for correct responses were 80/20, such that outcome was 20 pence 80% of the time, and 0 pence otherwise in win condition, while in the avoid loss condition, outcome was 0 pence 80% of the time and -20 pence otherwise.

Participants, stimuli and procedure

Twenty one participants (10 females, mean age = 22 (3) years) were recruited through a participant database at UCL. I excluded two participants in the imaging analysis as they were tested on different scanning parameters. All participants came to the lab on two consecutive days (day 1: training, day 2: scanning, roughly 24h apart) and were told they would receive payment at the end of second day based on their performance on both days. The reward scheme was adjusted such that all subjects received £30 for the time spent in the lab. The study was approved by the UCL ethics committee.

Cues were four fractal stimuli. Each cue is presented at the centre of screen with the squeeze stimulus superimposed on it. Each fractal stimulus (FIGURE 5-1D) was randomly assigned to one of four contingencies, crossing between effort (25% vs. 65% max force) and outcome valence (win vs. avoid loss).

General procedure. On day 1, participants underwent a learning and a testing block. On day 2, participants lay on the scanner bed to undergo, successively, force calibration, a practice block, four experimental blocks, a structural scan and force re-measurement. Calibration was completed as participants lay on the scanner bed outside the magnet, while the rest of the conditions were completed inside the magnet.

Training day. Participants completed a full learning block (as in Chapter 4) and a short testing block to test that they have learnt the contingencies. If they did not perform well when tested, they were explicitly told the contingencies, completed a short learning block and another testing block. By the end of day 1, all participants knew the cue-condition contingencies and that these stayed the same on the second day.

*Scanning day***.** On day 2, I made sure they remembered the contingencies to ensure stable, non-learning performance in the scanner. In the scanner, participants underwent 4 blocks of scanning sessions with a rest period between sessions. Each scanning block has 20 continuous, fully-randomised repetitions of four conditions, presented with a 5-secs rest every twelve trials. Overall, there were 320 trials lasting for 45 minutes.

As seen in FIGURE 5-1A, at the start of each session, a fixation cross was presented for 1 sec, followed by one of four fractal cues for 1 sec. Another fixation cross then appeared for a jittered anticipation period between 0.5-3.5 sec. Following this, a 'squeeze' instruction appeared which gave participants 1.5 sec to respond by squeezing either to the low or high effort level (indicated on screen by two tick marks). Within this 1.5 sec, participants had to reach either effort level within 1 sec and maintain grip force at that level for another 0.5 sec. Following this, they saw another fixation cross for 1 sec and the monetary outcome (20/0 pence for reward condition, 0/-20 pence for avoid lose condition) for 1 sec. There was a jittered ITI between 750-1500 ms before the next trial commenced.

I intend to disambiguate signal associated with anticipation of effort and valence from the actual movement execution. To do this, I introduced probabilistic execution such that although seeing a similar display of a bar reaching the squeeze target line, in half the trials participants did not have to execute the trials and were told that the computer will 'squeeze' for them. This essentially de-correlated the anticipation signal from squeezing signal. I was aware that during these

computer-executed trials, there would be a 'nogo' signal for inhibiting a squeeze response.

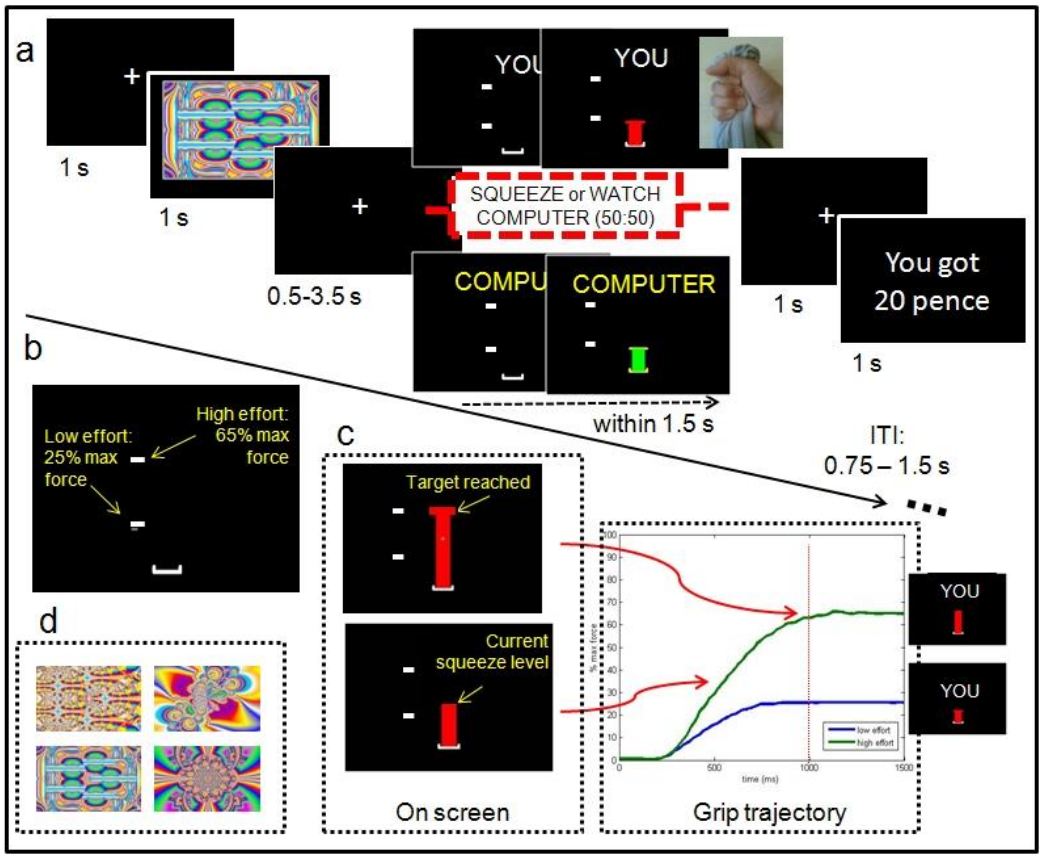

Figure 5-1 a) A schematic of one trial. Left-top to right-bottom: the first three screens showed a fixation cross for 1 sec, one of four fractal cues for 1 sec, and a jittered anticipation period between 0.5-3.5 sec (fixation cross). Following this, participants saw a 'squeeze' instruction which gave them 1.5 sec to respond by squeezing either to the low or high effort level (indicated on screen by two tick marks. Here, only low effort is shown). Then, they saw a fixation cross (1 sec) followed by a monetary outcome (20/0 pence for win condition, 0/-20 pence for avoid loss condition) for 1 sec. Inter-trial interval was jittered between 750-1500 ms before the next trial commenced. In 50% of the trials, participants did not execute squeezing but instead saw a green bar moving upwards indicating computer executed trials. b) The two tick marks were shown in both high and low effort trials during the squeeze period, the lower for low effort and higher for high effort targets. c) Here depicted an illustration of what participants see on screen (left) and the grip trajectories during 1.5 sec (right). The green line shows a trial of high effort, blue line for low effort squeeze, vertical red line indicates the cut-off time to reach the squeeze target line. Roughly, participants start at zero force level and slowly increase the force level, reaching the target at 1 sec and maintaining grip force for another 0.5 sec. d) For each participant, the four fractal stimuli were randomly assigned to the experimental conditions.

Nevertheless, I made it explicit to participants that 50% of the time they would not have to squeeze and that they only found out whether it was a computer-executed trial at the end of the jittered anticipation period, ~1.5-4.5 s after fractal stimulus was presented. This creates a situation where the fractal cue simply codes

the level of effort (and its associated outcome), with equal uncertainty for 'go' and 'nogo' responses, and that as the grip cue appears this uncertainty becomes a certain 'go' in half of the trials, and a certain 'nogo' in the other half. I accept that any activity which arose at the fractal cue could be a mixture between a certain 'go' and a certain 'nogo' signal, but I am confident that this signal is decorrelated from movement anticipation.

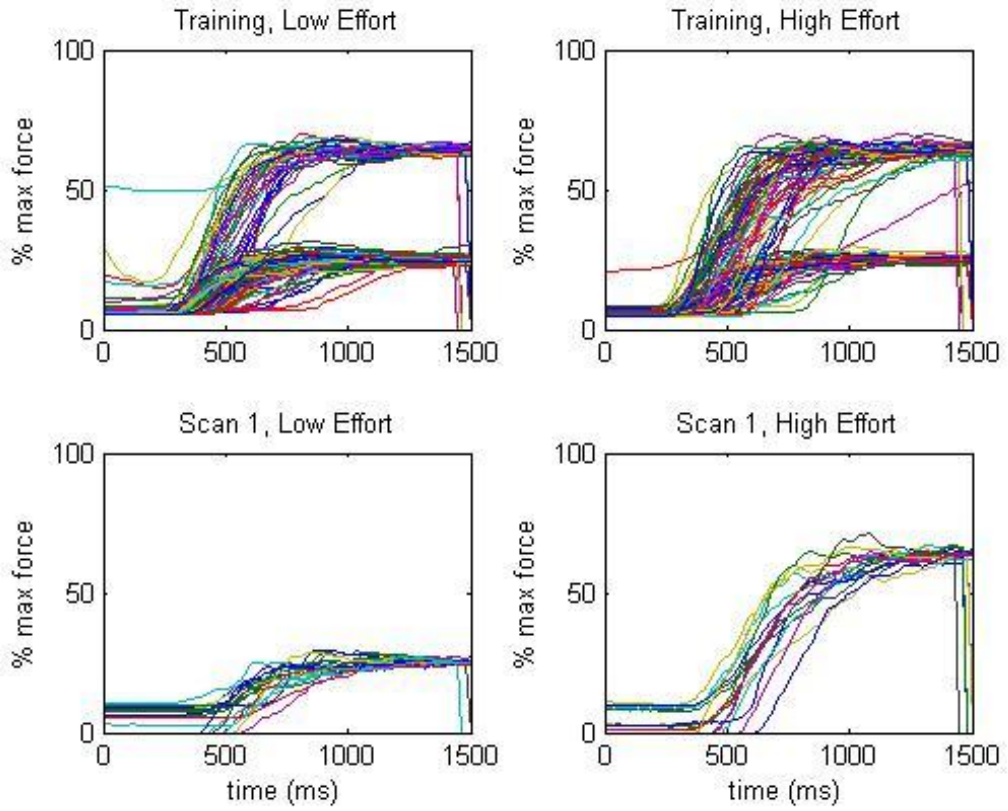

Figure 5-2 Example of one subject squeeze levels over training and scanning sessions. Each coloured line represents 1 trial. Squeeze level starts around zero with minimal natural noise from the squeeze device. Squeeze level then increases and reaches the target (25% or 65% max force) before 1 sec, and stays at that level until 1.5 sec lapsed. As seen on top row, this participant made incorrect responses during training: he squeezed 65% for the low effort cue, and squeezed for 25% for the high effort cue, but made no incorrect responses in the scanning session as shown on bottom row. The bottom row shows trials from only scan session 1, so there are fewer lines.

During the squeeze period I presented both effort target levels (tick marks). Thus, participants had to use their memory to decide which effort level to reach. Once the red bar reached the line and stayed there for 0.5 sec, the tick mark reached would turn red, indicating that they had successfully squeezed according to this 1.5 sec time rule. However, this did not indicate that they had chosen the correct effort level, given the cue. I specifically designed the grip task this way,

such that during training, participants did not experience motoric and visual uncertainties about achieving/not achieving the squeeze target. This is to ensure that any lack of learning could only be due to unformed/ impaired associations between cue, discrete effort levels (low vs. high), and outcome valence. In FIGURE 5-2, I show an example of one participant's squeeze behaviour. Here, squeezing diverged into either the low (25%) or the high (65%) effort levels, and at no other force levels. This demonstrates that this participant's response was motorically apt, and that there was always a discrete choice between squeezing at low or high effort. FIGURE 5-2 also illustrates that participant was making incorrect choices in training block, but was performing accurately in the scanning session.

Imaging analysis

I specified separate first level general linear models (GLM) for each participant by creating sets of regressors time-locked to i) fractal cue (action anticipation) and ii) outcome cue (outcome evaluation), with four scanning sessions concatenated into one.

To highlight activity correlating with *anticipation* of effort and valence, I defined four regressors-of-interest representing four event types at cue onset that varied in effort level and outcome valence: low effort-win (LowWin), low effort-avoid loss (LowAvoid), high effort-win (HighWin), and high effort-avoid loss (HighAvoid). To highlight activity reflecting *outcome evaluation* which correlated with effort, valence, and expected outcome, I defined eight regressors-of-interest representing eight event types at outcome onset that varied in effort level, outcome valence, and whether they received the expected (80% or the time) or the unexpected outcome (20% of the time). As participants were >95% correct in squeeze responses, their expected outcome was the better outcome (20 or zero pence in win or avoid loss conditions), and their unexpected outcome was the worse outcome (zero or -20 pence in win and avoid loss conditions). These regressors are called LowWinExpect, LowWinUnexpect, LowAvoidExpect, LowAvoidUnexpect, HighWinExpect, HighWinUnexpect, HighAvoidExpect and HighAvoidUnexpect. I entered three regressors-of-no interest for own squeeze periods (low and high effort separately), and computer-executed squeeze periods (collapsing low and high effort trials).

Two separate second level *F*-tests were specified. To do this, I computed a set of contrasts at first level for each participant for each of the relevant regressors-of-interest against baseline and fed the t-contrasts into second level F-tests. The first is a two-way Effort x Valence *F*-test with regressors-of-interest at cue onset, and the second is a three-way Effort x Valence x Expected Outcome *F*-test with regressors-of-interest at outcome onset. I ran a priori region-of-interest (ROI) analyses using anatomically defined masks for bilateral ACC and bilateral striatum (4 masks created using the software Marsbar, Brett et al., 2002; http://marsbar.sourceforge.net/; FIGURE 5-5A) and voxel-based, whole-brain analyses to look for main effects and interaction effects. To examine further involvement of suprathreshold regions from the whole-brain analysis, I created 4mm spherical ROI masks at the peaks of each suprathreshold cluster and extracted the signal in these ROIs. I then ran repeated measures ANOVA on the extracted signal. Any interaction effects found in these tests are orthogonal to the original clusters from which I derived the ROIs. I thresholded results at *p* .001 uncorrected, but only report *p* values with family-wise error (FWE) correction.

5.3 Results

Behaviourally, participants reached a good accuracy level in the training session, >65% correct, and a stable performance, >95% correct, in the scanner (FIGURE 5-3A). No effects of effort, valence, and effort by valence interaction on overall performance were significant.

In block-by-block training data (FIGURE 5-3B), there is evidence of learning on the training day, $F(5,16) = 7.12$, $p = .001$, $\eta_p^2 = .69$, a block by valence interaction between block 3 and previous blocks, $F(1,20) = 5.67$, $p = .027$, $\eta_p^2 = .22$ and a block by effort interaction between block 4 and previous blocks, $F(1,20) = 5.23$, $p = .03$, $\eta_p^2 = .20$. None of the 2x2 effort by valence follow-up ANOVAs showed significance in blocks 3 and 4.

Time to reach squeeze target

Participants learnt the squeeze timing criterion well. They reached the target before 1 sec, but the pattern of squeezing on day 1 is different to that on day 2 (FIGURE 5-3C). Subjects took their time in reaching the target before 1 sec on

scanning day, showing evidence of learning. I entered RT*reach* in a 2 x 2 x 2 Day x Effort x Valence repeated measures ANOVA and found that RT*reach* was different across the two days (significant main effect of day, day by effort, day by valence interaction effects, *p*s < .02, and marginally significant day x effort x valence, p = .059).

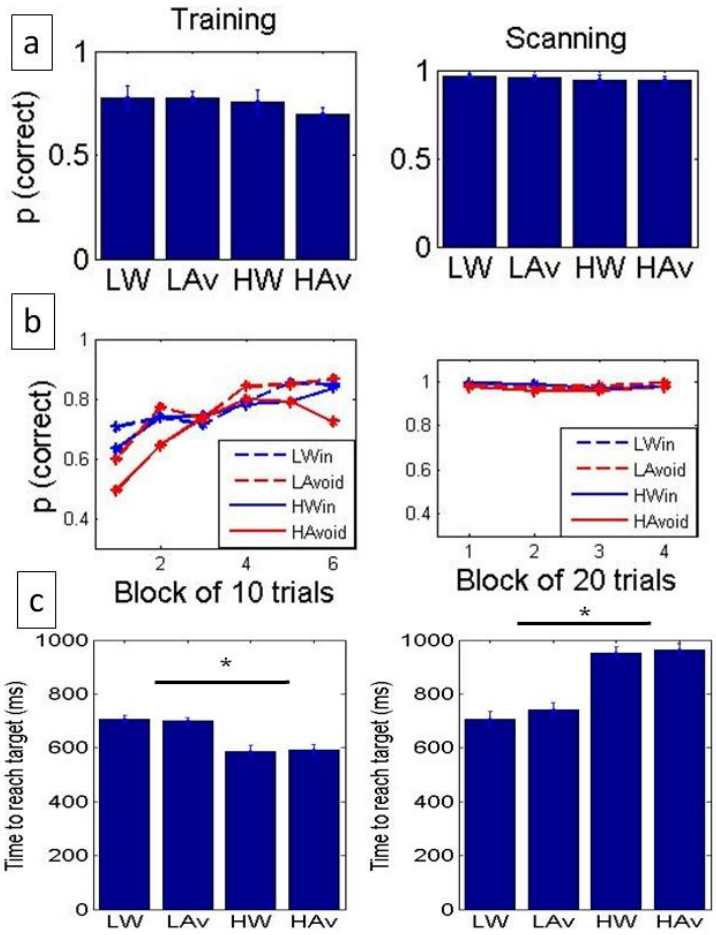

Figure 5-3 Behaviour on training (left) and scanning (right) days. a) Overall proportion of successful trials was >65% on training and >95% on scanning day. b) Participants show increasing block-by-block performance on training and stable performance on scanning day. c) The time it takes for participants to reach each target effort level (25% and 65% max force). On training day, participants took longer to reach low than high targets. On scanning day, they took longer to reach high than low targets, and they were also faster to reach targets to win than to avoid loss. LW/Lwin = Low effort to win, LAv/LAvoid = Low effort to avoid loss, HW/HWin = High effort to win, HAv/HAvoid = High effort to avoid loss.

To follow up, I ran separate 2x2 Effort x Valence ANOVAs on both days. On day 1, there was a main effect of effort, participants took longer to reach low than high effort, $F(1,20) = 21.58$, $p < .0001$, $\eta_p^2 = .51$, valence and interaction effects were n.s. On day 2, there were main effects of effort and valence, RT was longer to reach

high than low effort, $F(1,20) = 82.29$, $p < .0001$, $\eta_p^2 = .80$, and RT to win was faster than to avoid loss, $F(1,20) = 9.12$, $p = .006$, $\eta_p^2 = .31$, interaction was n.s. Note there was an opposite effect of effort in both conditions, while participants reached the high effort target faster in training, they reached the same target slower during scanning, just under 1 sec (FIGURE 5-3C) which indicates motoric mastery.

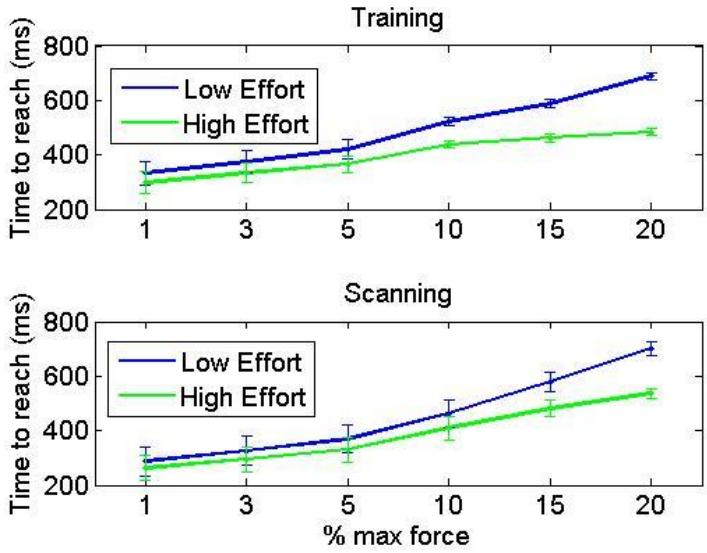

Figure 5-4 Squeeze acceleration in low (blue) and high (green) effort condition on training (top) and scanning (bottom) days. The shallower the slope is, the faster the acceleration from reaching 1% to reaching 20% max force. Here we see squeeze accelerate at a higher rate in high effort condition.

Squeeze acceleration

I sought difference in squeeze acceleration before reaching targets. To do this, I calculated how long participants needed to reach 1, 3, 5, 10, 15, and 20% max force levels. I observed no difference between valence conditions. Thus, I collapsed these conditions and entered RT into a 2 (day) x 6 (force levels reached), x 2 (effort conditions) ANOVA, yielding a significant day by effort interaction, $p = .017$. I than ran separate 2 x 6 Effort x Force level ANOVAs on training and scanning days yielding significant main effects of effort, force levels, and interaction on both days, ps < .0001. FIGURE 5-4 shows how long it took participants to reach 1, 3, 5, 10, 15 and 20% max force in low (blue) and high (green) effort conditions. What is most informative is whether squeeze acceleration is slower in low than in high effort condition.

A series of paired-samples *t*-tests between low and high effort conditions showed significant effects across all force bins on both days, such that participants took longer squeezing in low effort condition than they did for high effort condition, *t*s > 4.00, *p*s < .0001, except for reaching 1% on scanning day, $p = .06$. Visual inspection suggested that the differences in latency to reach levels that are <10% maximum force between the two effort conditions are small (below 100 msec), but at 15% and 20% maximum force, the data reflected a much longer latency when approaching the low effort target than when squeezing at the same levels for a high effort target. Participants seem to squeeze quickly to reach the high effort target and accelerate much faster than they do in low effort trials. This may suggest more control exerted in low effort conditions (FIGURE 5-4).

Anticipatory brain responses for effort and valence

I focused on BOLD response at onset of the fractal stimulus, indicating anticipation of action. Given robust past evidence for involvement of rodent ACC and striatum in effort processing, I ran a priori ROI analyses to examine whether the ACC and striatum responded to anticipation of effort, valence, or both (FIGURE 5-5A). These ROIs show a main effect of effort in bilateral ACC and left dorsolateral striatum. Responses to the cue in bilateral dorsal ACC and left putamen (shown in FIGURE 5-5B) reveal higher activity when participants anticipated performing an action at high effort than when a low effort action was anticipated (TABLE 5-1). I only found suprathreshold activity in this high > low effort contrast, but no difference in these ROIs for low > high effort contrasts, or for a valence or an interaction effect.

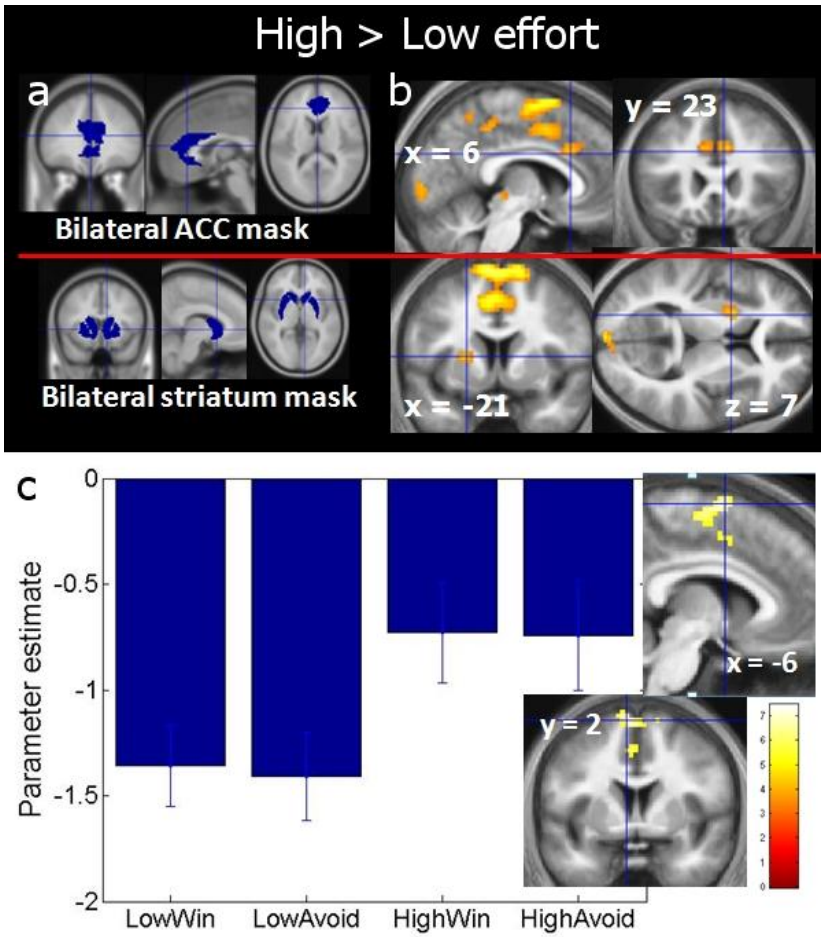

Figure 5-5 Main effect of effort. a) To test an apriori hypothesis on striatal and ACC involvement, I used anatomically-defined ROI masks on these regions bilaterally. b) ROI analyses showed that activity in both ACC and dorsolateral striatum was higher when cue indicated an incoming high effort action than when it indicated an action with low effort, small-volume and FWE corrected *p*s < .01. c) Voxel-based, whole brain analysis show higher activity in a large cluster around supplementary motor area for high > low effort contrast. Bars show averaged parameter estimates for the cluster (mean ± SEM, 147 voxels, *p* FWE-corrected <.05).

I next conducted a voxel-based, whole brain analysis which revealed a robust main effect of effort in a cluster involving bilateral supplementary motor area, premotor and primary motor, and somatosensory areas (147 voxels, *p* FWE-corrected <.05; TABLE 5-1). Seen in FIGURE 5-5C, averaged parameter estimates in this cluster indicate higher activity when anticipating high than low effort actions. No suprathreshold activity is found for low > high effort contrasts, valence main effects, or interaction effects.

Table 5-1 MNI coordinates of regions the activity of which is higher for anticipated high effort than anticipated low effort based on ROI and whole-brain analyses (*p* reported for ROI analysis is after small-volume correction, all *p* reported is FWE corrected at peak level).

Region	Nearest Brodmann Areas	Coordinates (mm)			Z value	No. of voxels	P
		x	y	z			
ROI analysis							
Left ACC	24/32	-6	+17	+28	3.70	13	.018
Right ACC	24/32	+6	+23	+28	3.82	17	.01
Left Putamen	NA	-21	+5	+7	4.24	30	.007
Left SMA	6	-6	+2	67	6.15	147	<.0001
Whole-brain analysis							
Left SMA	6	-9	-13	58	5.32		0.002
Left SMA	24	-3	+2	+46	5.17		0.004
Left premotor area	6/4	-24	-13	+70	5.41	46	0.001
Left primary motor	6/4	-21	-19	+58	5.27		0.002
Left primary visual	17	-6	-85	-8	5.31	4	0.002
Right primary motor	4	+21	-25	+64	4.84	6	0.017
Left somatosensory	1/2	-18	-40	+58	4.72	6	0.029
Left somatosensory	1/2	-21	-40	+57	4.70		0.031
Left primary visual	17	-3	-97	+10	4.65	1	0.038

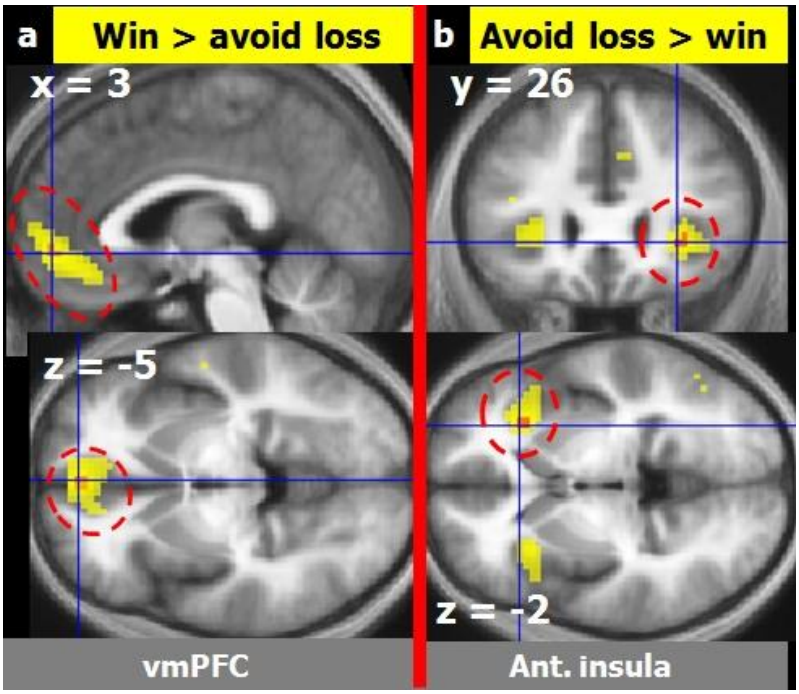

Figure 5-6 Brain response to outcome valence during outcome phase. a) Activity in vmPFC is higher when participants just completed trials that allowed them to win, than trials that only allowed them to avoid loss. b) The reverse contrast yielded activation in anterior insula. Yellow clusters are thresholded at p .001 uncorrected, red clusters survived FWE-correction.

Outcome evaluation for effort, valence, and expected outcome

I found main effects of valence and expected outcome during outcome phase, but no main effects of effort. For valence effects, FIGURE 5-6A shows that, at the time of outcome, vmPFC responds more strongly in trials where the fractal stimulus just indicated an action to win compared to an action to avoid loss. The reverse contrast shows that at the time of outcome, activity in the anterior insula is higher in avoid loss trials than in win trials (FIGURE 5-6B; both clusters p FWE-corrected <.05). Note that even though I am looking at brain responses during outcome phase, this main effect of valence simply reflects the valence context of the trials; that is a context where the response allowed participants to win or to avoid loss. As the outcome delivery was probabilistic, this effect of valence does not take into account the actual outcome presented, whether it was expected or unexpected.

Thus, I looked at a main effect of expected outcome by contrasting response to outcomes that were expected and those unexpectedly omitted (see methods).

Table 5-2 MNI coordinates of regions the activity at outcome phase of which is higher for win conditions than avoid loss conditions (all *p* reported is FWE corrected at peak level).

Region	Nearest Brodmann Areas	Coordinates (mm)			*Z* value	No. of voxels	*P*
		x	y	z			
Contrast: Win > avoid loss							
Right vmPFC	10	+3	+53	-5	4.73	5	.02
Right vmPFC	10	+12	+53	+1	4.52	1	.047
Contrast: Avoid loss > win							
Right anterior insula	NA	+30	+26	-2	4.72	5	.021

Brain response correlated with expected outcome

Voxels in bilateral ventral striatum (ventral putamen) and left vmPFC were more active when seeing an expected than unexpected outcome (FIGURE 5-7A, TABLE 5-3). I then examined involvement of these regions in outcome expectation by creating three 4mm spherical ROI masks at the peaks of these cluster and extracted the signal in these ROIs. To test for modulatory effects of effort or valence on outcome evaluation, I ran a 2x2x2 Effort x Valence x Expected outcome repeated measures ANOVA on the extracted signal (see methods).

In left vSTR ROI (peak at [-18 8 -8]), I found no main effect of effort, but a significant valence effect such that there was a stronger response in win than avoid loss trials, $F(1,18) = 7.23$, $p = .015$, $\eta_p^2 = .28$, and a two-way interaction between effort and expected outcome, $F(1,18) = 6.00$, $p=.025$, $\eta_p^2 = .25$.

To follow up, I averaged the extracted signal values across valence conditions, and ran t-tests to look for simple effects of expected outcome in low and high effort separately. This effort by expected outcome interaction is driven by a diminished effect of expected outcome in the high effort trials. T-tests show that following a low effort action, vSTR show a stronger response to expected than unexpected outcomes, $t(18) = 7.53$, $p < .0001$, whereas following a high effort action, this effect was only marginally significant, $t(18) = 2.06$, $p = .054$ (FIGURE 5-7B top).

Table 5-3 MNI coordinates of regions the activity of which reflects main effect of expected outcome based on whole-brain analyses (all *p* reported is FWE corrected at peak level). R= right, L= left, IFG= inferior frontal gyrus, dACC= dorsal anterior cingulate cortex.

Region	Nearest Brodmann Areas	Coordinates (mm) x	y	z	Z value	No. of voxels	*P*
Expected > unexpected							
Left (ventral) putamen	NA	-18	+8	-8	5.73	19	<.001
Left mid orbital gyrus	10/12	-9	+47	-5	5.00	14	.006
Right ventral putamen	NA	+21	+8	-8	4.99	10	.006
Right posterior insula	NA	+54	-22	+16	4.98	19	.006
L.parahippocampal gyr.	NA	-30	-37	-11	4.92	3	.009
Left insula	NA	-54	-4	10	4.88	6	.010
Right mid orbital gyrus	12	+6	+32	-11	4.88	13	.011
Left cerebellum	NA	-24	-49	-50	4.70	2	.022
L. sup. frontal gyrus	8	-21	+38	+43	4.55	1	.041
Right (dl) putamen	NA	+27	-13	+13	4.55	1	.042
Unexpected > expected							
Right insula	NA	+33	+23	-2	7.32	530	<.0001
R. IFG (p. opercularis)	44	+54	+17	+37	6.11		<.0001
R. IFG (p. triangularis)	44	+51	+23	+25	6.09		<.0001
Left insula	NA	-30	+23	-5	6.35	76	<.0001
R. inferior parietal cortex	40	+42	-52	+46	6.09	123	<.0001
R. middle temporal gyrus	21	+57	-28	-5	5.25	9	.002
L. inferior parietal cortex	39/ 40	-33	-55	+40	4.83	5	.013
Left cerebellum	NA	-12	-79	-29	4.68	4	.025
R. dACC	6	+6	+32	+46	4.53	2	.045

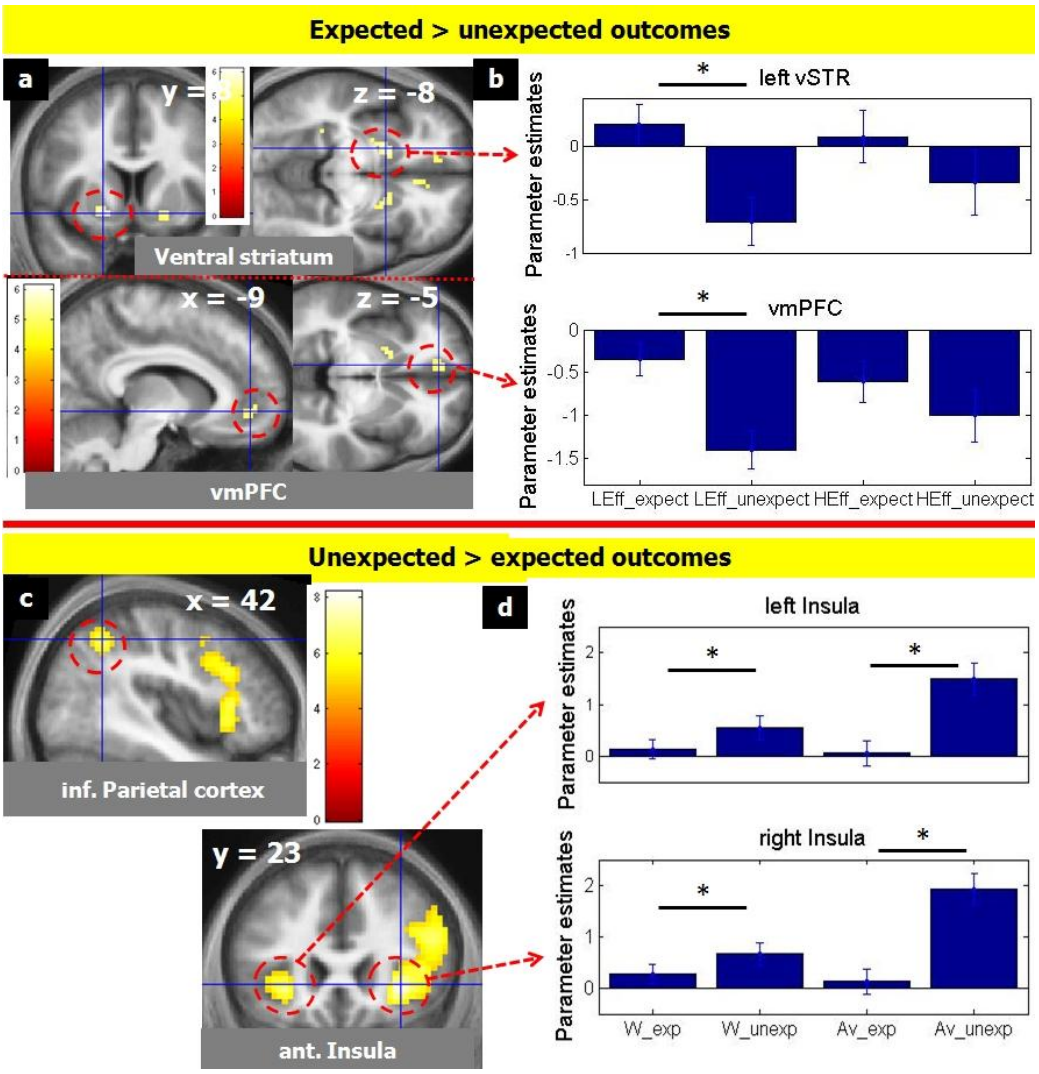

Figure 5-7 Brain response to expected and unexpected outcomes during outcome phase. a) Regions in ventral striatum (putamen) and vmPFC were more active when participants saw an expected, better outcome than an unexpected, worse outcome (p FWE-corrected <.05). b) Parameter estimates of extracted signal in left vSTR (top) and left vmPFC (bottom) show effort by expected outcome interaction such that the effect of expected outcome is more pronounced if it is an outcome of low effort action than if it is an outcome of high effort action. The effects of expected outcome in both ROIs are significant only in low effort action. c) Brain response to unexpected outcome shows a large cluster in bilateral insula expanding to right inferior frontal gyrus (pars opercularis and triangularis; left image) and in the right parietal cortex (right image), p FWE-corrected <.05. d) Extracted signal in bilateral insula ROIs shows valence by unexpected outcome interaction. In these regions, brain response is significantly stronger to unexpected than expected outcomes in both valence conditions, although is more pronounced in avoid loss than in win trials. Bars show mean ± SEM. LEff_expect = Low effort, expected outcome, HEff_expect = High effort, expected outcome, LEff_unexpect = Low effort, unexpected outcome, HEff_unexpect = High effort, unexpected outcome, W_exp = Win trial, expected outcome, W_unexp = Win trial, unexpected outcome, Av_exp = Avoid loss trial, expected outcome, Av_unexp = Avoid loss trial, unexpected outcome.

In right vSTR ROI (peak at [21 8 -8]), I only found a significant valence effect such that there was a stronger response in win than avoid loss trials, $F(1,18) = 4.7$, $p < .04$, $\eta_p^2 = .20$, no main effect of effort or any two-way interaction was significant.

In left vmPFC ROI (peak at [-9 47 -5]), there was no main effect of effort, but response in win trials was significantly stronger than that in avoid loss trials, $F(1,18) = 9.02$, $p = .008$, $\eta_p^2 = .33$. There was also a marginally significant two-way interaction between effort and expected outcome, $F(1,18) = 4.30$, $p = .053$, $\eta_p^2 = .19$. This marginally significant effort by expected outcome interaction is driven by a non-significant effect of expected outcome in the high effort trials. T-tests show that following a low effort action, vmPFC was sensitive to expected outcomes, $t(18) = 5.04$, $p < .0001$, whereas high effort does not modulate expectation, n.s. (FIGURE 5-7B bottom).

Brain response correlated with negative prediction error

In the reverse contrast, where brain response to an unexpected omission of outcome (negative prediction error) was stronger than to an expected outcome delivery, I found a large cluster involving bilateral insula and extending into right inferior frontal gyrus (both pars opercularis and triangularis). This contrast also yielded higher activity in right parietal cortex. I then also ran an ROI analysis using 4mm masks peaking at bilateral insula to test for modulatory effects of effort or valence on brain response to unexpected outcome.

Within the left insula ROI (peak at [-30 23 -5], FIGURE 5-7C), as found in the whole-brain analysis, response in avoid loss trials were significantly stronger than that in win trials, $F(1,18) = 12.43$, $p = .002$, $\eta_p^2 = .40$. There was also a significant interaction between valence and unexpected outcome, $F(1,18) = 20.89$, $p < .0001$, $\eta_p^2 = .53$. There was no main effect of effort. To follow up this interaction, I averaged the extracted signal values across effort conditions, and ran t-tests to look for simple effects of unexpected outcome in win and avoid loss trials separately. This significant valence by unexpected outcome interaction is driven by a stronger effect of unexpected outcome in the avoid loss trials, $t(18) = 5.13$, $p < .0001$ than in the win trials, $t(18) = 3.12$, $p = .006$ (FIGURE 5-7D top).

Within the right insula ROI (peak at [30 23 -5], FIGURE 5-7C), I found a significant avoid loss > win effect, $F(1,18) = 18.15$, $p < .0001$, $\eta_p^2 = .50$, and a

significant valence by unexpected outcome interaction, $F(1,18) = 29.25$, $p < .0001$, $\eta_p^2 = .61$, but a non-significant effect of effort. Similar to left insula, this valence by unexpected outcome interaction in right insula is driven by a stronger effect of unexpected outcome in avoid loss trials, $t(18) = 7.99$, $p < .0001$, than that in win trials, $t(18) = 2.31$, $p = .032$ (FIGURE 5-7D bottom).

5.4 Discussion

The results suggest that brain activity is only sensitive to size of *anticipated* effort, but not to reward or punishment. In contrast, once the action is completed, brain activity does not respond to effort just exerted, but is instead sensitive to outcome valence (i.e. the possibilities to win or to lose), and to modulation of effort and valence on outcome monitoring. That is, response to expected outcomes is weaker after actions with high than low effort, whereas response to unexpected outcomes is stronger in trials where they could only avoid losing compared to trials where they have the possibility to win.

ACC, dorsal striatum and SMA for anticipated high versus low effort

I found higher activity in the ACC and dorsal striatum when an action with large effort is anticipated. First, this provides converging evidence to previous involvement of ACC in rodents and humans (Croxson et al., 2009; Floresco & Ghods-sharifi, 2007; Prevost et al., 2010; Walton et al., 2002). Indeed, lesions to rodent ACC impair willingness to scale an effortful barrier in order to gain large reward (Rudebeck et al., 2006). However, it is worth noting that, unlike most of previous effort studies, my current task does not involve explicit cost-benefit analysis. What this suggests is that value comparison during cost-benefit analysis might require a representation of the kind I observed in this task, one about effort size.

Second, dorsal striatum in rodents supports stimulus-response (S-R) associations and the results demonstrate sensitivity to the specific cost parameter of the action when such a cue-action representation is useful to indicate an upcoming action. However, in CHAPTER 3, I report that the same voxels in left dorsal putamen show a stronger response when an action with low effort, compared

to high effort, was chosen. This is puzzling given BOLD signal in the same voxels in these two studies is coding effort size in opposite directions.

The investment of effort in this thesis is operationalised and inspected with a tight control using extensive training blocks. This was done with the intention to ensure that every choice (study 2) or anticipation (study 4) phase is a genuine process which closely matches the actual effort investment. Such control invokes habit–based actions. My results demonstrate that such cognitive process of choice or anticipation of effort recruits dorsal striatum, a substructure of the striatum which primarily receives sensorimotor information (Voorn et al., 2004), and facilitates habitualised, automatised behaviour (Wickens et al., 2007).

In the choice task in CHAPTER 3, the BOLD response at the time of choice may reflect both the cognitive representation of effort and a result of value comparison. It is possible that this putamen activity is a correlate for how valuable the chosen action is, which would reflect a stronger signal in actions with low effort (high value) than high effort (low value). We know that striatum has been implicated extensively in valuation and value-based decision making (Rangel, Camerer, & Montague, 2008), and this result provides support for this literature. The whole-brain contrast between anticipating high and low effort actions showed a large activation in the supplementary motor area, an area that supports behaviour such as movement generation (Picard & Strick, 2001). I have made extra steps to exclude as much motor execution and preparatory signal as possible by de-correlating activity during cue presentation and action execution, and including self and computer executed squeeze periods as regressors-of-no interest (see methods). These steps are taken to ensure that BOLD response to cue presentation reflect pure action anticipation. However, I cannot fully exclude the possibility that the contrast for high versus low effort still contain a general motoric preparatory signalling.

vmPFC and ventral striatum reflects sensitivity to goodness of event

There is consensus that vmPFC and ventral striatum are part of circuitry implicated in appetitive value (Seymour, Singer, et al., 2007), with some proposing the role for vmPFC in subjective value and biasing choice (Kable & Glimcher, 2007). In my task, vmPFC was more responsive in reward than in punishment context. When examined further, vmPFC, now together with ventral striatum seem to selectively respond to outcomes that are expected. This selective signal to

expected outcomes persists only following an action with low effort. In fact, following an action with high effort, these regions fail to distinguish whether an action resulted in expected or unexpected outcomes.

Both vmPFC and ventral striatum seem to indicate the goodness of the trial, with strongest activity for a desired, expected outcome (a reward in win condition, and punishment avoidance in avoid loss condition) which resulted from a less demanding action (low effort). Indeed after having executed a more strenuous, high effort action, sensitivity to the goodness of the outcome diminished. This effort modulation of signalling of desired, expected events is novel and brings a new focus on neural scaling of outcome signal by action costs. The finding is in line with the dominant views concerning the functional roles played by vmPFC and ventral striatum in signalling desired, positive events.

Insula reflects sensitivity to punishments

I found that insula is more responsive in punishment than in a reward context. The same insula cluster also responds stronger to unexpected omissions of outcomes than to expected outcomes. This selective signal to unexpected outcomes is modulated by valence such that, it is stronger in a punishment than a reward context. Both the avoid loss > win and valence-modulated unexpected > expected outcomes contrasts point to the possibility that insula, here, is indicating the badness of the event. In my task, the unexpected omission of reward in a *win* condition can be strictly thought as negative prediction error, but the unexpected omission of safety signal in *avoid loss* condition (when they expected to receive zero pence for a correct response) is simply an unexpected punishment. I found the largest activity over these trials in cases where subjects received an unexpected punishment. I also observed a noteable effect in unexpected omission of reward in the insula. The insula is extensively implicated in representation of aversive value (Sarinopoulos, Dixon, Short, Davidson, & Nitschke, 2006; Seymour, Daw, et al., 2007; Seymour, Singer, et al., 2007; Talmi et al., 2009). While it is useful and appropriate to divide the trials and run the factorial three-way ANOVA as described above, this interaction found in insula may ultimately express sensitivity to aversive events.

Furthermore, from the unexpected > expected outcome contrast, I also found activity in the inferior parietal cortex. In light of previous role of this region in switching strategies, activity in the parietal cortex may serve as an evaluative

signal which informs the subject whether s/he is wrong and that s/he needs to change strategy. Activity in this region has been suggested previously to contribute a switching strategy (Rushworth, Paus, & Sipila, 2001).

Several caveats and further analyses

It is surprising that I could not detect any suprathreshold BOLD response to anticipated reward or punishment, considering ubiquitous findings of reward- or value-correlated BOLD activity in regions such as striatum and OFC at the time of cue (e.g., Croxson et al., 2009; Knutson et al., 2005).

I specifically contrast my paradigm with that used by Croxson and colleagues (2009). Reward size in their paradigm was visually and explicitly marked by the location of a horizontal line on a circular cue. This is different to the process through which my participants may have formed a cue-outcome association. My participants went through a trial-by-trial learning in which making the correct action was what was relevant to them. A stronger association between a cue and the correct action is formed when a cue-action link is reinforced, i.e. results in better outcomes (win or loss avoidance), but this reinforcing effect may not necessarily result in a cue-to-action+outcome association, but a sequential cue-action-outcome one. Indeed, as reviewed by Schoenbaum and colleagues, dorsal striatum in rodents and putamen in primates have been broadly implicated in cue-action associations (Stalnaker, Calhoon, Ogawa, Roesch, & Schoenbaum, 2010).

Another possible explanation for the absence of outcome-related activity is that the monetary outcome here may not be as salient as the vigour the participants had to produce shortly following cue presentation. Using the same effort device in CHAPTER 3, I also failed to find a main effect of reward at cue despite seeing robust reward effects on behaviour. This could be directly tested by increasing the outcome sizes in further experiments, and by devising ways to estimate relative salience between effort and outcome.

Second, my interpretation for the effort by expected outcome interaction is limited with the fact that I did not exclude outcome onsets that happened after computer-executed squeezing periods. I cannot rule out the possibility that this effort modulation on vmPFC and ventral striatal signal for expected outcomes simply reflects the modulation of having seen a tall or a short 'squeeze bar' on the screen, rather than having squeezed with a large or small hand force. I aim to

refine my analyses by creating a design matrix which excludes all outcome onsets in computer-executed trials. This way I can be sure that effort modulation was due to participants' own experience of exerting large or small force.

Third, I am aware that I have not excluded trials in which participants made an incorrect response, thus activity to outcomes that are expected and unexpected may be contaminated by trials in which an expectation was to receive an outcome for an incorrect response. Nonetheless these trials are so few that any effect due to incorrect responses would not have been strong enough to change the current interpretation.

Fourth, the follow-up tests for the significant two-way Effort x Expected outcomes and Valence x Unexpected outcomes interactions could be significantly improved by creating new second-level F-tests which would contain the averaged regressors across low and high effort in ventral striatum and vmPFC and across *win* and *avoid loss* in insula. Instead, I simply manually calculated an average of the extracted signal outside SPM8. Although the current ROI follow-up tests are orthogonal from the whole-brain results, the proposed analysis would be more conventional.

Fifth, I have qualitatively categorised expected/ unexpected outcomes, resulting from less/ more demanding actions to be in goodness-badness spectrum. I am aware this is rather an informal, yet psychologically valid description of events in the task. I would like to adopt a more formal approach to describing the value of each trial/event, for example with a reinforcement learning approach which mathematically characterise trial-by-trial action values and the state values associated with them.

Finally, I computed the time it took to reach different grip level criteria to assess grip acceleration between high and low effort conditions (FIGURE 5-4). It would be useful to also do this across valence conditions. Additionally, a potentially more refined analysis to test if people exert different vigour as a function of effort and valence is to model the slope of the increase in grip trajectory.

Summary and conclusions

In sum, I found no behavioural evidence for valence modulation on effort deployment, but I have found neural findings relevant to action anticipation and outcome evaluation. Activity in the ACC and dorsal striatum is higher for anticipating high compared to low effort, but reward and punishment contexts do

not seem to be relevant at the time of action anticipation. When action has been completed, effort levels do not seem to be relevant, but the goodness and badness of events seems to evoke differential activity in the ventral striatum/ vmPFC and insula, respectively.

This study provides support for the role of ACC in signalling effort. It also provides data that inform a general understanding of the neural underpinnings for processing affective events. In this case a monetary outcome (deemed 'desirable' or 'undesirable' in the task) which results from an action that is either less or more demanding, evokes activity in regions for appetitive and aversive values. Note that the interpretation of the findings relies heavily on the assumption that i) losing money is aversive, that ii) winning money is appetitive and iii) that exerting effort is costly where a larger effort bears more cost. I have reason to think that this is the case based on previous work (Kurniawan et al., 2010; CHAPTER 3), where I showed that behavioural choice and psychological liking are associated with squeezing and monetary earnings, and that choice is associated with a broader personality trait to persist with daily challenges.

This work forges new avenues for exploring brain responses for committing an action and its association with rewards and punishments. A more refined design that allows examination of different actions and different ways in which effort can be exerted is likely to shed light into how affective events are associated with actions and the boundaries in which forming of an action-outcome association becomes impaired.

But pleasure (as opposed to pain) cannot be the only factor affecting my decision to act ..."

Karol Wojtyla, Love and Responsibility, p.36, 1960.

Chapter 6 Context and Pain (study 5)

6.1 Pain and effort

In 1960 Karol Wojtyla, an actor and a philosopher, who later became better known as Pope John Paul II, identified the importance of pain in influencing our decisions about which is the best course of action. He wrote this just a few years after Stevens' (S. S. Stevens, 1957) attempt to provide psychophysical measures for various sensations such as loudness and heaviness, although at that time failing to include pain sensation. A decade previously Mosteller & Nogee (Mosteller & Nogee, 1951) attempted to create a laboratory measures for 'utility', at that time ignoring notions of cost-benefit tradeoffs.

What became clear in his later philosophical and theological work is that what he meant with 'pain' was not only the primary visceral cost that we experience as we receive a sudden electrical jolt by touching a power source, but also emotional suffering and the enduring of physically challenging demands that we encounter in life. The latter resembles the construct I explore in this thesis: effort. Indeed in everyday language, (physical and emotional) pain and effort seem interchangeable as both contain aversive value and are more or less traded-off against benefits such as a job salary, a top-of-mountain ecstasy, or eternal life. In what follows I briefly discuss ways in which effort is distinct from physical pain, but also ways in which effort can relate to physical pain. This subsection provides a background rationale for my final study which investigates the influence of context on pain avoidance.

Physical pain plays a major role in shaping behaviours related to health and disease. As the body's primary aversive stimulus, pain signals imminent or actual physical harm, evokes a feeling of unpleasantness, and constitutes a potent signal that helps to shape future behaviour toward minimising injury (Craig, 2003; Fields, 2004). Physical pain can be defined in laboratory settings as any primary, visceral sensation caused by aversive events such as electrical currents, focal heat stimulation, or sharp pricks or pinches on the skin. The aversive and robust bottom-up quality of pain alongside with abundant evidence for its malleability to top-down control have triggered and maintained long-standing bodies of knowledge

on issues such as placebo effects, pain rehabilitation, analgesia or chronic pain syndromes.

Empirically, effort is distinct from pain. Despite the recurrent theme in previous chapters that effort is costly we do not avoid effort to the extent that we avoid pain. Indeed, effort is an abstract concept which has not yet implicated such robust neural and physiological signatures as pain does (Tracey & Mantyh, 2007). Neverthless, due to the scarce empirical work on either, work on understanding pain-reward integration has used previous discussions on effort-reward integration as a framework for cost-benefit tradeoffs (Talmi et al., 2009) and work on 'suffering' for the sake of charity-giving has conflated both pain and effort manipulations (Olivola, 2010).

Effort is intimately linked to pain. In physical rehabilitation settings, perceived effort and pain ratings are simultaneously used as metrics for rehabilitation training efficacy. For example, both ratings of effort and pain were acquired to assess peripheral control on movement (Hollander et al., 2010). In this study, the authors compared venous occlusion on an arm during light-weight biceps exercise with a non-occluded medium-weight biceps exercise, to test if perception of effort and pain can be influenced by peripheral sensation that is caused by venous occlusion. Participants had to perform a number of arm flexion exercises and make verbal reports on effort and pain. They found that both ratings rose to a medium level ('6' on a BORG scale) at a similar rate as a result of arm exercises, and demonstrated peripheral control of effort and pain ratings.

In occupational health, reports of pain are also associated with perceived effort. As reviewed previously (Tam & Yeung, 2006), cases of body pain are robustly associated with perception of physical exertion during work, such that workers who required treatments for their lower back pain due to work demands (e.g. lifting) also perceived higher exertion rate when tested on various lifting measures. It could be that perceived effort becomes a cognitive signal for behavioural modification to avoid pain occurrence (Tam & Yeung, 2006).

Colloquially, the exertion of effort is often implicated as a source of pain in various body parts. Indeed loss of grip strength, which is also an effort measure used in previous chapters, seems to be ubiquitous in individuals reporting chronic pain (Lohman, Thorpe, Prior, George, & J. P. Kim, 2008). Effortful breathing causes significantly greater pain than pain at rest in post-upper abdominal operative patients (Kimball et al., 2008). While minimal activities such as walking

may not cause pain in healthy individuals, it is certainly true for certain individuals, such as those with obesity. The more challenged our body parts feel in completing physical movements or tasks, the more likely we are to report exertion or effort, and this is linked to subsequent reports for pain (e.g. Karason et al., 2005). On the flip side, faked effort as a pain index is a controversial issue in medico-legal settings, and tests have been developed to distinguish submaximal effort exertion in malingerers who try to claim legal benefits for chronic pain (Lohman et al., 2008). The biological mechanism for how effort causes pain is not straightforward. Indeed, simple analgesic manipulations to modulate the endogenous opioid system (using codeine) does not attenuate muscle pain ratings after strenuous grip exercises (Cook et al., 2000). Nevertheless, in most settings, effort exertion does eventually lead to subjective pain.

The above highlights the intimacy between the experience of effort and pain and the importance of extending research on action costs to the pain domain. Below I report my first attempt to study pain, instead of effort, in assessing the influence of a context manipulation on pain avoidance as an important exemplar for cost-driven actions.

6.2 Relative magnitude influences on pain avoidance

Abstract

Motivational theories of pain highlight its role in people's choices of actions that avoid bodily damage. By contrast, little is known regarding how pain influences action implementation. To explore this poorly understood area, I conducted a study wherein participants had to rapidly point to a target area to win money while avoiding an overlapping penalty area that would cause pain in their contralateral hand. I found that pain intensity, and target-penalty proximity, repelled participants' movement away from pain and that motor execution was influenced not by absolute pain magnitudes but by relative pain differences. My results indicate that the magnitude and probability of pain have a precise role in guiding motor control and that representations of pain that guide action are, at least in part, relative rather than absolute. Additionally, my study shows that the implicit monetary valuation of pain, like many explicit valuations (e.g., patients' use of rating scales in medical contexts), is unstable, a finding that has implications for pain treatment in clinical contexts.

6.2.1 Introduction

Pain dominates the shaping of health and illness-related behaviours, providing an imminent aversive signal for harm. Traditional studies of motivational aspects of pain have concentrated on either subjective rating of unpleasantness (in humans) or aversive classical and instrumental conditioning (primarily in other species; Dayan & Seymour, 2008; Price, 2000). Although these approaches have yielded considerable insight into how pain influences action *choice*, few studies have investigated how pain influences action *implementation*. Both action choice and action implementation are central themes in theories of optimal control: action choice is formalised, for example, by reinforcement learning theory (Seymour et al., 2004), while action implementation is formalised by theories of motor control. To see how both factors operate, imagine that you burn your arm while removing bread from an oven. The ensuing pain might influence both your decision to use the oven in the future and the movements you will make when reaching into the oven again. Pain's influence on action implementation, although ubiquitous in ecological contexts, remains poorly understood.

From a functional point of view, pain is often viewed as helping to guide behaviour in an effort to balance an agent's long-term interests and immediate goals. Conventional ideas about the motivational role of pain are based on the assumption that pain provides a signal of an approximate but absolute quantity of ascending nociceptive input (leaving aside descending modulatory influences that arise in specific circumstances; Fields, 2004). Optimality requires that pain signals provide an absolute measure of potential bodily damage. For example, from an evolutionary or economic and nutritional standpoint, people should stop gathering or eating a food at exactly the point when the risk of bodily damage outweighs that food's caloric value. Successfully making this type of trade-off via the proxy of experienced pain requires that instances or predictions of bodily damage map consistently onto subjective pain—that is, such ideas assume that pain is absolute rather than relative.

However, recent studies on explicit decision making when pain is a factor have produced striking results that call into question this assumption about the absolute nature of pain. For example, when people bid money to avoid painful electrical stimuli in an auction paradigm, the financial value they were willing to pay for pain relief was influenced by the amount of a different pain they had

recently experienced (Vlaev, Seymour, Dolan, & Chater, 2009). This finding supports theories about relative judgment in explicit affective valuation, as well as theories in perceptual domains such as vision and audition (Garner, 1954; Laming, 1984, 1997). However, it remains possible that these results reflect a relativistic process related to the construction of explicit valuations rather than a more fundamental property of pain perception itself. This possibility motivated my experimental approach in the present study, in which I exploited a motor task that obviates the need for explicit judgments (Maloney, Trommershäuser, & Landy, 2007) but nevertheless provides a metric of sensitivity to pain intensity.

In recent motor-control experiments, participants making rapid pointing movements in situations involving risk chose visuomotor strategies that maximised gain (Trommershäuser, Landy, & Maloney, 2006; Trommershäuser, Maloney, & Landy, 2003a, 2003b, 2008). In these studies, participants pointed at configurations similar to the ones shown in FIGURE 6-1B. If participants hit the target area, they won a small monetary reward, but if they hit an overlapping or abutting penalty circle, they incurred a small monetary loss. Results showed that participants optimised their mean pointing response according to changes in penalty value. The distance by which participants avoided the penalty region was indicative of how "bad" they rated the monetary loss. Participants chose pointing strategies that maximised expected gain.

Extending this approach, one can estimate how aversive a shock would be to participants in terms of monetary units by presenting two overlapping regions, one carrying monetary gain and one carrying immediate shock, and measuring how far participants' finger points are repelled from the shock region. A region that carries a higher shock level should repel finger pointing farther than a region that carries a milder shock level. This approach provides an ideal system in which to study the role of pain as a disincentive in motor planning and to test the hypothesis that relative coding of pain intensity is a core property of pain representation.

6.2.2 Method

Participants, apparatus and materials

Seventeen volunteers (9 males and 8 females; mean age = 24 years, *SD* = .74) were recruited through the psychology participant database at UCL. All participants

were right-handed or ambidextrous. They gave written consent to participate in the study, were paid between £20 and £32 (depending on performance), and were debriefed after the experiment. The study was approved by the UCL ethics committee.

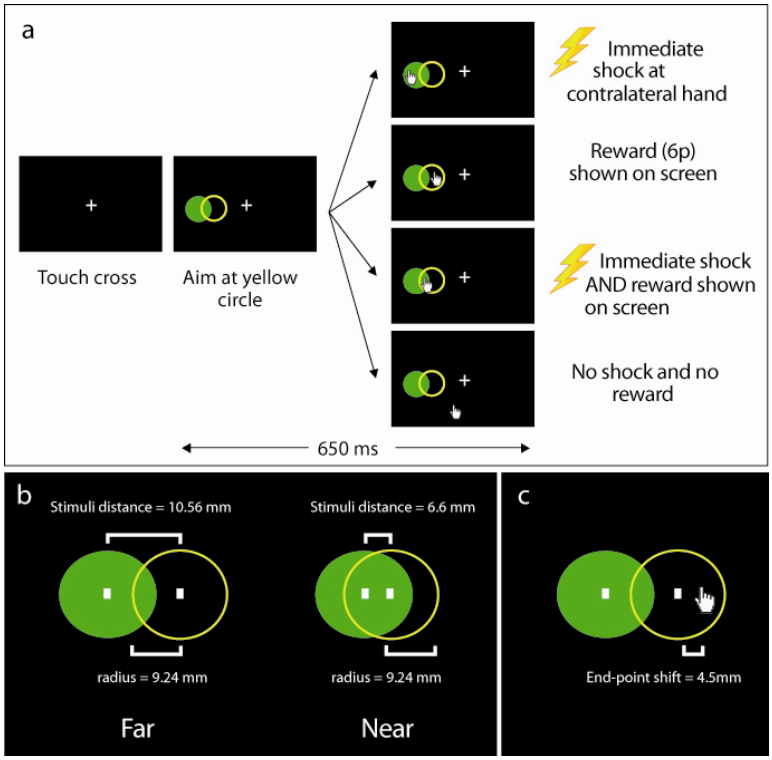

Figure 6-1 Illustration of the experimental stimuli, sequence of events in a trial, and main dependent variable. The stimulus compound consisted of an open yellow circle (the target circle) and a filled coloured circle (the penalty circle, shown here in green). The hand images indicate the end points of participants' pointing movements. a) Participants had to touch a central cross to make the stimulus compound appear, after which they had 650 ms to respond to the stimulus. If participants touched the penalty region, they received an electric shock. If they touched the target region, a monetary reward was shown on the screen. Participants received both pain and reward if they touched the overlapping region, and they received neither pain nor reward if they touched the screen outside the target and penalty regions. b) Two stimuli configurations in the far and near conditions. White squares show the centres of the circles. c) The measured end-point shift for a given trial was the horizontal distance between the end point of the participant's response and the centre of the target region. The illustrations are not to scale.

The MATLAB toolbox used was Psychophysics Toolbox Version 2.54 (Brainard, 1997; Pelli, 1997). Participants sat 70 cm from a 25-in. touch screen (Keytec, Inc., Garland, TX). Electrical pain stimuli were delivered and controlled by three DS7 Stimulators (Digitimer, Hertfordshire, United Kingdom), which have been fully approved for clinical use. These apparati have been used for various pain

experiments (Mobbs et al., 2007; Vlaev et al., 2009). Electrical pain stimulates a broader range of nociceptive and nonnociceptive afferents than, for example, laser or thermal noxious stimulation. Electrical pain offers researchers an advantage over other forms of stimuli because it is largely free of the confounding effects of stimulus habituation or sensitisation (McMahon & Koltzenburg, 2005).

General task description

I trained participants to rapidly touch (within 650 ms) a small target area on a computer screen (Gepshtein, Seydell, & Trommershäuser, 2007; Trommershäuser, Gepshtein, Maloney, Landy, & Banks, 2005). Participants earned money by hitting the target area, which carried a fixed known reward of approximately 6 pence per hit (paid at the end of the experiment). Hitting the penalty area resulted in immediate administration of a shock (low, medium, or high level).

Participants received both money and a shock if they hit the overlapping region of the target and penalty areas (FIGURE 6-1A). The magnitude of pain varied between trial blocks, and participants learned the magnitude in each block only when they hit the penalty region. Participants received no money or shock if they did not respond within 650 ms, in which case they see a message "too late" on the screen.

I manipulated the target-penalty distance (near: 6.6 mm; far: 10.56 mm) and the shock level associated with each penalty (low, medium, and high pain). End-point shift—the distance between the center of the target circle and the end point of a pointing movement (FIGURE 6-1C)—was the critical dependent variable. The idea behind the experiment was that penalties should have the effect of repelling a participant's end points away from the penalty region to a degree dependent on the movement inaccuracy for that individual participant. Specifically, a higher pain level and a near penalty region would be more aversive than a lower pain level and a far penalty region (Trommershäuser et al., 2006) and would therefore result in larger end-point shifts.

To test for absolute versus relative pain encoding, I presented two shock strengths during each trial block (low-medium, medium-high, and low-high). On each trial, the relative intensity of the shock was indicated by the colour of the penalty area. That is, participants were told that the colour of the penalty area

indicated whether the higher or lower shock intensity in that block was in effect, but experience alone informed them of the actual intensity.

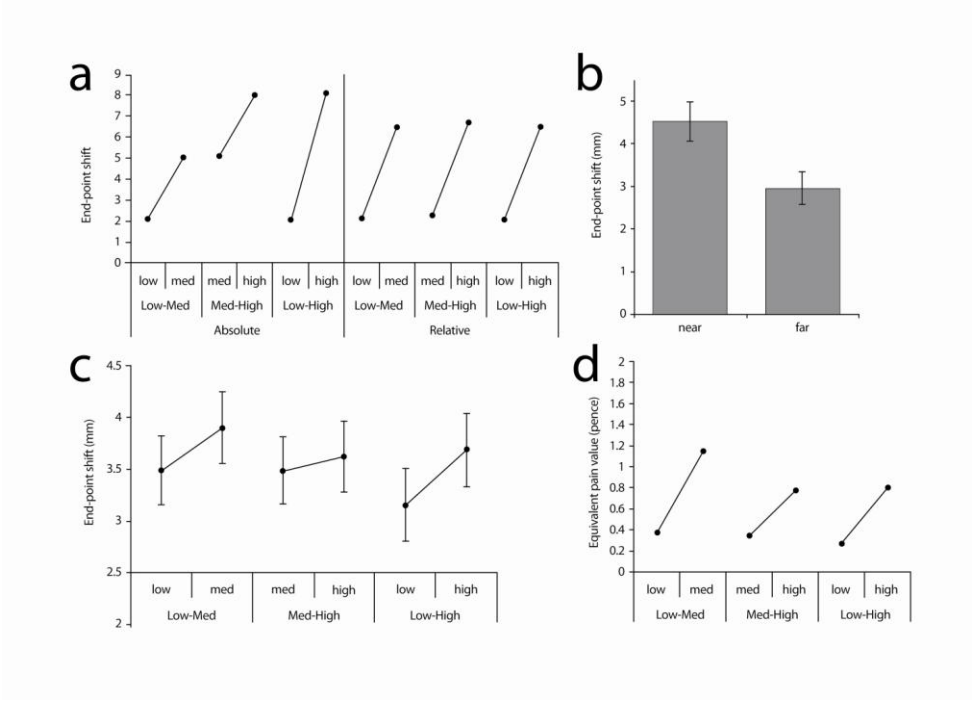

Figure 6-2 Context dependency in motor control for pain avoidance. The graphs in (a) illustrate the end-point shift predicted in the three experimental blocks according to an absolute-coding model (left) and a relative-coding model (right). Increases in end-point shift would be expected to reflect increases in pain magnitude in the absolute-coding model, but in the relative-coding model, increases in end-point shift would be expected to remain identical across experimental blocks. The observed pooled-participant mean end-point shifts are shown as a function of (b) target-penalty distance and (c) pain intensity in the low-medium, medium-high, and low-high blocks. Error bars represent standard errors of the mean. The graph in (d) shows estimated equivalent monetary value of low, medium, and high pain within the low-medium, medium-high, and low-high blocks.

I assumed that each response would not reflect a summarised coding of the two pain intensities within the block. Rather, I assumed that participants' motor systems would distinguish the two pain intensities consistently, such that a higher pain level would always be avoided by a greater distance than its lower-level counterpart. The crucial distinction between an absolute and a relative model of pain is that this higher-versus-lower pain-response pattern applies only within blocks in the case of relative coding, but applies both within and across blocks in the case of absolute coding.

Put differently, according to an absolute-coding model, end-point shifts should depend purely on the absolute pain intensity presented at each trial, and should be independent of the other shock intensity presented in that block. In

contrast, according to a relative-coding model, end-point shifts should vary according to a pain's intensity relative to the other pain stimulus occurring in the same block. For instance, a medium-intensity stimulus should repel end points to a greater degree if it is the higher of the two intensities in a block (i.e., in a low-medium block) than if it is the lower of two intensities (i.e., in a medium-high block). FIGURE 6-2A illustrates the predictions of these hypothesised absolute and relative models.

Stimuli

The visual stimulus presented on each trial consisted of a target and a penalty circle, each of which had a 9.24-mm radius. The target was always an open yellow circle. The penalty was always a filled circle.

Each of the three experimental blocks had two shock levels, which were indicated visually by different colours; different colours were also used in different blocks. For each participant, I randomly chose six penalty colours from among seven colours (excluding four colour pairs that could not be visually discriminated easily). This variability in colour coding was made clear to participants; they were able to visually distinguish the target circle from the penalty circle and expected two penalty colours representing different shock levels in each block. Colour coding allowed participants to identify which penalties had a higher pain level within an experimental block; this use of colour coding also ensured that the colour-pain association did not carry over to other blocks. For example, a blue circle might represent low pain throughout the first block, but in the next block, low pain would be associated with a different colour, such as pink. Penalty colours in practice blocks were different from the penalty colours in experimental blocks.

At the start of each trial, a cross (8 mm × 8 mm) appeared at the centre of the screen. When participants touched the cross, the stimulus appeared for 650 ms; its location was randomly selected to be 9.9 cm to the left of, to the right of, above, or below the cross.

Procedure

Appropriate shock levels for each participant were calibrated in advance of the trials. Two silver-chloride electrodes were placed on the back of the left hand. A brief current was delivered through the electrodes to cause a transitory aversive

sensation, which became increasingly painful as the current was increased. I administered shocks, starting at extremely low intensities and ascending in small steps, until participants reached their maximum tolerance. No shocks above a participant's stated tolerance level were administered. Participants rated each shock on a visual analogue scale from 0, *no pain at all*, to 10, *the worst possible pain*. Their ratings allowed us to determine the appropriate range of current amplitudes to use during the actual experiment and to assign pain levels (low, medium, and high) that were subjectively comparable across participants.

Once their maximum tolerance was reached, participants received fourteen random subtolerance shocks that removed expectancy effects created by the incremental procedures. A Weibull (sigmoid) function was statistically fitted to participants' ratings for the fourteen shocks and the intensities of current that related to three levels of pain (mild: 4; moderate: 6; strong: 8) were estimated; and subsequently used for the three shock levels (low, medium, and high) in the experiment. Participants were unaware that only three specific amplitudes of current were used during the experimental task. The participants rated the same set of fourteen subtolerance shocks in a random order at the end of experiment. A one-sample t test showed that the sum of the difference between participants' first and second ratings was not significantly different from zero, $t(16) = 1.25$, $p = .22$, which suggests that there was no systematic change between participants' first and second ratings.

To investigate the possibility of adaptation more precisely, I compared the second ratings made by participants who completed the low-medium, medium-high, or low-high block as their final block in the experiment. If participants had adapted after their final block, ratings made by participants whose final block included low intensities (e.g., the low-medium block) should have been higher than ratings made by participants whose final block included high intensities (e.g., the medium-high block). A Kruskal Wallis (nonparametric) test showed no evidence of such adaptation: The mean rating differences were the same among participants who had just completed the low-medium, medium-high, or low-high blocks, $\chi^2(2, N = 17) = 0.40$, $p = .81$. These results suggest that there was no significant habituation or sensitisation during the experiment.

Participants completed three practice phases and three experimental blocks. During the first practice phase, which had 64 trials (eight repeats of eight stimulus locations), participants learned to point within 650 ms. The penalty area appeared

randomly at a middle distance (9.24 mm) to the left or right of the target's center point. Participants then completed the second phase, which was the same as the first phase except that there were 72 trials and participants received a mild shock when they hit the penalty area. In the third phase, the penalty circle was randomly presented either near (6.6 mm) or far from (10.56 mm) the target (FIGURE 6-1B). Participants completed 112 trials (seven repeats of sixteen stimulus locations). The shock level was the same in Phases 2 and 3, but this level was different from the shock levels in the experimental blocks. Because of the time limit for responding, the task was difficult, and these three practice phases allowed participants to achieve adequate accuracy rates without learning the pain magnitudes to which they would be exposed in the experimental blocks.

There were 128 trials (four repeats of sixteen stimulus locations at two pain levels) in each of the three experimental blocks. The order of the experimental blocks was determined randomly for each participant. The experimental blocks represent three pairs of pain magnitudes, which allowed us to test whether finger-pointing shifts reflected relativistic or absolute coding of pain magnitudes.

Data analysis

I conducted repeated measures analyses of variance (ANOVAs) with three independent variables: distance (near or far), block (low-medium, medium-high, or low-high), and relative pain (lower or higher within each block). The dependent variables were average end-point shifts from the center of the target (FIGURE 6-1C) and reaction times (RTs). Responses on 14% (*SD* = 2%) of the trials were late (equally distributed across blocks), and these trials were excluded from all analyses. All trials during which participants responded within 650 ms (including trials with end-points outside the circles) were included in the analyses.

In principle, stimulus intensity (measured in milliamps) could have been added into the general linear model, although any significant association between stimulus intensity and end-point shifts would vary widely according to factors such as skin temperature, sweating, hydration, sex, and skin thickness. Therefore, in line with normal practice in the pain literature, it was not included.

To determine trade-offs between reward and pain, I compared the shifts I observed in participants' response to changes in pain intensity with the strategies of an optimal movement planner maximising gain. The only free parameter in this

comparison was alpha, which represented the pain-pence exchange rate for each shock level. This comparison yielded an estimate for the monetary value of the penalty that corresponded to the movement shift I observed in response to changes in pain intensity. The method for computing this equivalent monetary value is described in APPENDICES.

6.2.3 **Results**

As FIGURE 6-3 shows, participants hit the target-only area significantly more often than they hit the penalty-only area or the overlapping region, $F(1, 17.46) = 171.65$, $p < .00001$, $\eta_p^2 = .91$. I tested whether participants adjusted their end-points according to pain intensity and target-penalty proximity. To do this, I computed pooled-participant mean end-point shifts by computing median values for each participant's horizontal end-point shift in each condition, and then averaging these median values across all participants. This value served as an index of how far participants deviated from optimal pointing (Trommershäuser et al., 2005).

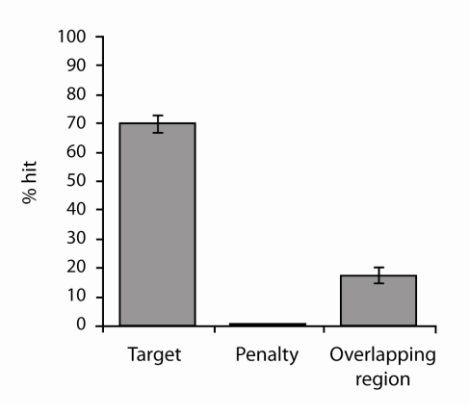

Figure 6-3 Mean percentages of participants' end points that hit the target-only area, penalty-only area, and overlapping region of the target and penalty area. Error bars represent standard errors of the mean.

An ANOVA revealed that participants displaced their end-point much farther when the penalty was near the target than when it was far from the target (FIGURE 6-2B), $F(1, 14) = 66.60$, $p < .00001$, $\eta_p^2 = .82$. This finding is consistent with the hypothesis (Trommershäuser et al., 2008) that movement execution incorporates information relating to judged movement variability (noise). Displacement from the target's centre also depended on relative pain magnitudes; that is, end-point shift was larger when pain was stronger than when pain was milder, $F(1, 14) = 4.84$, $p = .045$, $\eta_p^2 = .25$. End-point shift was not affected by

absolute pain intensities. The Block × Relative Pain interaction was not significant; the difference between lower and higher pain was similar across the three experimental blocks (FIGURE 6-2C). These results suggest that end-point shift was influenced by relative pain intensities. Other effects on end-point shifts were nonsignificant, Fs(2, 13) < 3.59, ps > .057, and F(1, 14) < 2.46, p > 0.13.

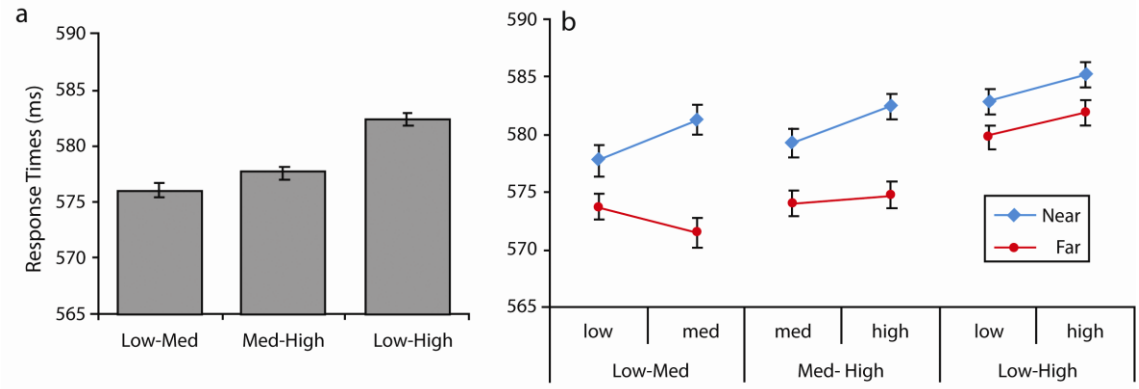

Figure 6-4 Reaction times (RTs). a) Pooled-participant mean RT for each of the three experimental blocks (low-medium, medium-high, and low-high). b) Pooled-participant mean RT as a function of target-penalty distance (near vs. far) and pain intensity within each experimental blocks. Bars show mean ± SEM.

I also examined participants' RTs (calculated from when they touched the fixation cross to when they touched the stimulus compound). Participants responded more slowly when the penalty circle was near than when it was far (see FIGURE 6-4B), F(1, 14) = 12.32, p = .003, η_p^2 = .46. RTs were also influenced by block, F(2, 28) = 5.2, p = .012, η_p^2 = .27 (FIGURE 6-4A). RTs in the low-high block were significantly slower than RTs in other blocks—low-medium block: t(14) = 2.7, p = .017; medium-high block: t(15) = 2.23, p = .041. Participants responded with equal quickness in the low-medium and medium-high conditions (p > .05). Although FIGURE 6-4B suggests that there may be a trend for an interaction, all interaction effects were nonsignificant, Fs < 0.12, ps > .80. See APPENDICES for complete descriptions of ANOVA results for end-points and RTs.

Under the assumption that end-point displacements corresponded to an optimal pointing strategy that maximised gain (Trommershäuser et al., 2008), I estimated the equivalent monetary value of each shock level to assess participants' trade-offs between reward and pain. Overall, participants consistently experienced higher shocks to be more painful and unpleasant than lower shocks. When converted into a hypothetical equivalent monetary value of pain for an optimal movement planner maximising gain, the shift in mean motor response to higher

shocks corresponded to higher equivalent monetary values than the shift in mean motor response to lower shocks did (details about the computation of the monetary values of pain can be found in APPENDICES). These results demonstrate that pain can be measured in equivalent monetary values. The results for this implicit measure correspond with those for my explicit measure (pain avoidance in end-point shifts), which suggests that pain is encoded relatively in guiding motor movement. FIGURE 6-2D depicts the context dependency of the estimated equivalent monetary values of pain.

6.2.4 Discussion

The data show that previous painful outcomes exert a pervasive influence on future movement control. First, I have shown that higher-intensity pain generally has a stronger influence on biasing future movement in a direction away from pain. Second, I have shown that the likelihood of pain, inferred by the proximity of pain to the goal target, biases movement in a similar way. This suggests that movement execution incorporates the consideration of both the magnitude and the probability of pain, as predicted by an optimal account of motor control. This study helps build a richer picture of the motivational dimension of pain because it shows that pain not only influences decisions about whether to perform an action (i.e., escape and avoidance behaviour), it also informs the actual execution of that action.

My results indicate that the influence of pain is more relative than absolute. That is, relatively intense pain that has been recently experienced has a greater effect on movement control than relatively mild pain that has been recently experienced. In addition, these findings suggest that noxious events are represented in relative terms at the level of basic motor control, which is putatively a much more fundamental index of the mental representation of such events than subjective ratings are. My results correspond nicely with the relativistic valuation of pain Vlaev and co-workers (Vlaev et al., 2009) demonstrated in an economic bidding game (borrowed from behavioural economics). The correspondence between explicit and implicit pain valuation in my study also resembles the correspondence between risk perception as examined via a classical economic decision-making task and an equivalent motor task (Wu, Delgado, & Maloney, 2009).

My implicit analysis of monetary values of pain implies that its context effect on movement control could be explained by differential economic values of pain. It is conceivable that people will tradeoff the amount of pain they will choose to suffer against the amount of money they are willing to pay to relieve that pain (Vlaev et al., 2009). Thus, the relative end-point shifts I found in this study could partially be explained by the fact that participants' monetary valuation of pain was sensitive to the relative context of that pain.

Two caveats should be noted in relation to the interpretation of my findings. First, it is difficult to rule out the possibility that short-term habituation to pain might have contributed to the relative coding I observed. Although I did not find evidence for habituation over the course of the experiment, it is possible that higher-intensity stimuli caused a relative diminution of pain through habituation effects that operated over the course of each block. Second, according to some accounts of relativity effects, participants use recent experiences to inform expectancies about forthcoming pain (Seymour & McClure, 2008). That is, participants infer distributions of anticipated pain and incorporate these distributions as priors in representational inference about inherently uncertain ascending afferent inputs. Thus, apparent relative effects might emerge not due to a fundamental limitation in people's ability to encode intensity, but because of uncertainty in the ascending input.

My results have implications for pain in clinical environments. A number of conditions and disorders cause pain that is exacerbated by movement; examples include conditions arising out of peripheral injury (e.g., post trauma), neuropathic conditions (e.g., complex regional pain syndrome), and central nervous system disorders (e.g., post stroke pain). Behaviours such as limb guarding (protecting a limb after recent trauma) are pervasive during recuperation and are essentially physiological. In other clinical situations, pain acts as a barrier to optimal functional recovery for the affected limb. Accordingly, an understanding of the exact ways in which pain modulates movement planning and execution can inform therapeutic strategies, particularly in poorly understood (but critically important) areas such as upper-limb physiotherapy. Furthermore, the existence of relative coding might inspire strategies that exploit context effects to improve movement recovery when pain experience is a recognised obstacle.

Chapter 7 General Discussion

7.1 Summary of findings

The decision to act requires a complex integration and anticipation of physical costs. There is currently little understanding regarding key questions, including: how are such costs computed in the brain, how are they integrated with value, and how do they influence the neural sensitivity to outcomes? Here I consider effort and pain as physical costs in a range of contexts; 1) effort choice, 2) pavlovian influence on effort learning, 3) basal ganglia-prefrontal sensitivity to effort anticipation and outcomes and 4) relativity influence in pain avoidance.

In the individual chapters I report that:

- Anticipated effort influences the likelihood to choose an action, its subjective likeability, and the time taken to decide. Persistence as a personality trait is also associated with likelihood of accepting an effortful action.
- Learning about effort actions is influenced by affective context (i.e. reward or punishment). When anticipating reward, we tend be active and expend effort, whereas when anticipating punishment, we tend to withdraw from expending effort. Computationally, this is accounted for by a pavlovian influence which specifies a 'spillover' from stimulus value associated with reward/punishment into an action value for exerting effort.
- Anterior cingulate cortex (ACC) is sensitive to the effort requirement of an upcoming action, but not to affective contexts namely rewarding or punishing outcomes.
- Over and above a nonspecific role for anticipating movement, the dorsal striatum (putamen) plays a crucial role in effort computations both in choice and non-choice contexts.
- Ventromedial prefrontal cortex (vmPFC) and ventral striatum are sensitive to the hedonic aspect of an action, by directly integrating its expended disvalue (effort) and expected value (outcome).
- Insula is sensitive to an action's displeasure, by responding most strongly to an unexpected punishment and most weakly to an expected reward.

- The context provided by relative pain magnitudes influence how the motor system implements pain avoidance (another exemplar of cost-driven behaviour). The same level of pain is avoided farther when paired with a milder pain than when it is paired with a stronger pain.

7.2 Apathy, persistence and compulsion

As alluded to earlier in this thesis, my behavioural and neuroimaging work in healthy participants may have implications for a more fine-grained analysis of neurological cases of apathy, using brain and behavioural evidence. First, several distinct types of brain insult are associated with apathy in humans. For example, bilateral ACC lesions can present with akinetic mutism, a wakeful state characterised by prominent apathy, indifference to painful stimulation, lack of motor and psychological initiative (Tekin & Cummings, 2002). Apathy is also often present in patients with subcortical brain lesions (involving BG), but is more commonly found in those with prefrontal, mainly ACC, lesions (van Reekum, Stuss, & Ostrander, 2005). More recently apathy in Alzheimer's disease patients has been associated with weaker ACC white matter integrity (J. W. Kim et al., 2011), whereas apathy in frontotemporal dementia population has no association with basal ganglia grey matter volume (Links et al., 2009).

Second, effort is a salient variable in individuals with apathy who lack the ability to initiate simple day-to-day activities with excessive reliance on external control (a spectrum that incorporates abulia) (Lévy & Dubois, 2006; van Reekum et al., 2005). This lack of internally generated actions may stem from impaired incentive motivation: the ability to convert basic valuation of reward into action execution (Schmidt et al., 2008). Patients with auto-activation deficit (AAD), the most severe form of apathy, are characterised by lack of self-initiated action (van Reekum et al., 2005) or a quantitative reduction in self-generated voluntary behaviours (Lévy & Dubois, 2006). Thus, the key feature in AAD is an inability to internally generate goal-based actions, a deficit that may variously reflect an ability to (1) encode that the consequence of an action as pleasurable or as having hedonic value (e.g., to attain reward, 'liking') (2) execute the action; and (3) represent the association between action and reward. I now discuss a proposal that

the behavioural and neural mechanisms underlying AAD are mostly intimately linked to the third sub-process.

AAD is not associated with impaired 'liking' as patients with AAD have a normal skin conductance response to receipt of rewards and verbally distinguish between different magnitudes of monetary reward (Schmidt et al., 2008). In addition, the most prominent damage in AAD pertains to BG and the dopaminergic system. Secondly, AAD is probably not linked to specific impairments of action execution. Schmidt and colleagues (2008) tested patients with bilateral BG lesions with the history of AAD and found that, compared to normal and Parkinson's disease control groups, patients with AAD are worse when generating voluntary vigourous actions based on contingent reward, but are equally able to generate the same motor response if based on external instructions. This provides evidence against AAD being explicable in terms of an impairment in pure motor action execution.

I suggest that AAD reflects an impairment in linking reward anticipation to action. Damage to BG in AAD most commonly involves a focal bilateral insult to the internal portion of pallidum (Lévy & Dubois, 2006). Pessiglione and colleagues investigated the role of ventral pallidum in incentive motivation employing a task where individuals voluntarily squeezed a handgrip device in response to different reward magnitudes (Pessiglione et al., 2007). Notably, the amount of voluntary force during squeezing was proportional to reward magnitude, suggesting that participants were able to identify a reward context where it was advantageous to produce more physical effort. Furthermore, ventral pallidal activity correlated with outcome context, providing a neural basis for enhanced effort as a response to increased payoff. Similarly, damage to BG in AAD may have caused a failure to recognise an advantageous context to make an adaptive action (Lévy & Dubois, 2006; Walton et al., 2004). These data suggest that bilateral BG damage, at least in AAD, produces a syndrome that arises out of a deficit in translating reward cues into appropriate action selection and execution.

In light of Schmidt and co-workers' (2008) findings that AAD patients were mostly impaired in the execution of actions, when an internal link between a reward and action is required, it is noteworthy that AAD may cause impairments beyond simple abstract action-reward association. In other words, AAD may cause impairments in the actual execution of reward-based actions. This highlights the importance of BG in energising individuals to act with perseverance, a deficit

commonly found in patients with Parkinson's disease (which is largely associated with a dysfunction in BG). Schneider tested Parkinson's disease patients in solving a difficult cognitive task, and found that the patients were making significantly fewer attempts to solve the task than normal controls, pointing to a deficit in mental persistence in such patients (Schneider, 2007). It may well be that persistence is linked to a higher tendency to generate internal motivation or arousal which then energises individuals to persevere (Gusnard et al., 2003), or perhaps lessens a tendency to distraction (Nicola, 2010).

Taken together, apathy, as a manifestation of impaired motivation to overcome the cost of an action, is associated with damage to a cortico-subcortical network (either lesions in the ACC or BG) that generates internal association between action and its consequences. This highlights a key involvement of the ACC and BG in the anticipation and execution of effortful actions.

I have discussed how apathy could be an instance where effortful, motivated behaviour, to gain reward is impaired. At first glance, persistence seems to be the opposite of apathy which provides an optimal behaviour where persistent individuals are capable of exerting effort while still maximising gain. However, a further opposite of apathy which highlights another state of impaired behaviour might be the case of compulsive behaviour. Using the orthogonalisation of action and outcome valence (Boureau & Dayan, 2011; Guitart-Masip et al., 2011; Huys et al., 2011) described in CHAPTER 4 we can see that, on the one hand, apathy is a marker of an impairment in invigoration for reward (FIGURE 7-1), while on the other hand, compulsion could be a marker for a failure to suppress habitual responses to evade punishments. Compulsion has been shown to be a behavioural manifestation of trait impulsivity (Belin, Mar, Dalley, Robbins, & Everitt, 2008). Early evidence for failure in response suppression comes from a rodent experiment which shows that while rats with limited cocaine exposure are able to suppress lever pressing to self-administer cocaine when the lever is now associated with shock delivery, rats with extended cocaine exposure develop a compulsive cocaine-taking response and fail to suppress this behaviour and perseverate in self-administering cocaine even when this simultaneously delivers shock (Vanderschuren & Everitt, 2004). Two fundamental points to clarify here are i) whether compulsive behaviour is specific to failure in this 'NoGo-to avoid punishment' quadrant, and ii) whether this precise failure takes place during learning or performance.

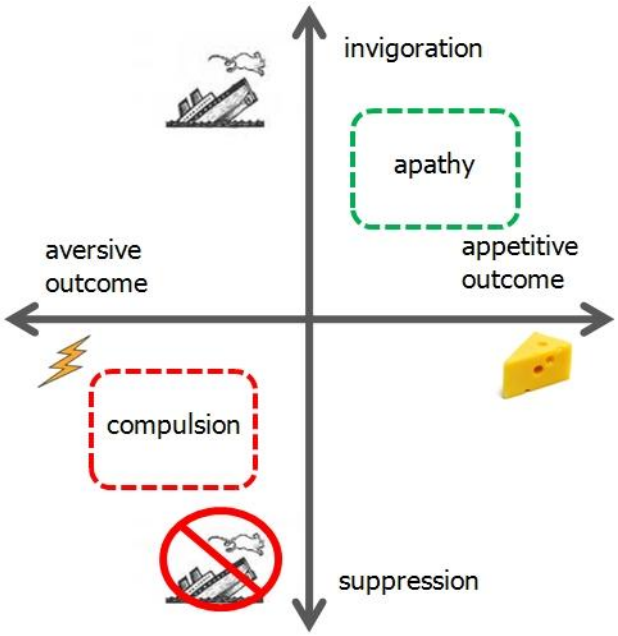

Figure 7-1 The affect-effect plot (Boureau & Dayan, 2011) depicted in chapter 4. By using this framework, we could potentially characterise apathy as having impairments in vigour for rewards and compulsion as having impairments in suppressing actions for punishment avoidance

This novel opponency between apathy and compulsion, depicted in the affect-effect plot in FIGURE 7-1, may provide a fruitful framework for investigating aberrations in the forms of lack, and excess, of effortful actions and in understanding the intimate relationship between action and outcome. Future work should further exploit this orthogonal paradigm akin to studies in CHAPTER 4 and in Guitart-Masip et al. (2011) and Huys et al. (2011) to test clinical populations such as individuals with apathy and obsessive compulsive disorder (OCD).

7.3 Outstanding issues and future directions

Repetitive responding

In the research reported here, I utilised one means of measuring and manipulating physical effort through force production. It is worth considering a different form of effort such as that in repetitive responding. Like force production, repetitive responding is differentially influenced by a dopaminergic manipulation (Ishiwari et al. 2004). Both forms of effort may be associated with a behavioural trait of persistence which characterises a human tendency to exert self-regulatory effort (Segerstrom and Nes 2007) in order to achieve long term goals (Duckworth et al. 2007). Reported in CHAPTER 3, persistence is associated with activity in dorsal ACC

when participants rejected an option with low effort. This provides provisional support for an extensive neurological literature that links circuitry damage involving the ACC to various motivational impairments, as for example seen in apathetic patients (van Reekum et al. 2005; Eslinger and Damasio 1985). An important future research avenue would be to examine if repetitive responding interacts with force production in influencing action choices and how this relates to a persistence trait, apathetic syndromes and OCD.

Serotonin and dopamine

Emerging work using the affect-effect orthogonalisation often highlighted the opponency between 5HT and DA (Boureau & Dayan, 2011; Cools, O. J. Robinson, & Sahakian, 2008; Daw, 2002; Huys et al., 2011). While DA has been much robustly implicated in vigour, as well as reward, some weaker evidence points to an association of 5HT with inhibition and punishment. In relation to our novel attempt to extend an action spectrum into vigour (CHAPTER 4**)**, it is difficult to explain our behavioural data using this opponency. A nascent opponency literature is extremely complicated and underexplored in the typical approach-avoidance spectrum (Huys et al., 2011), and our paradigm does not provide a straightforward extension from existing paradigms, nor does it allow a clear inference about the likely neurotransmitters involved in such processes.

Effort integration into outcome value

In addition, I have shown in CHAPTER 5 evidence of effort integration in the prefrontal-striatal sensitivity to hedonic outcomes. I found that vmPFC and ventral striatal response to expected, desirable, outcomes is stronger after expending less effort and weaker after large effort. Of course, this conceptualisation of effort integration into hedonic outcomes (more effort, less hedonic) is somewhat informal. Thus, a stronger analysis should utilise computational approaches to formalise how just-expended effort could be integrated into outcome valuation.

Effort may boost value

Throughout this thesis, I base my research on the assumption and empirical observation that effort is costly. In contrast, under some conditions, effort may boost preference or value (briefly listed in CHAPTER 1). This may arise in the

context of relief; that having done a hard work may yield a sense of relief which yields appetitive value. This is loosely associated with cognitive dissonance (Festinger, 1957) where an individual needs to justify his/her effortful action by concluding that s/he must like the outcome that action brings and therefore has a high preference for that action. Another suggestion is that this boosting effect could be related to social contexts attached to the actions (Heyman & Ariely, 2004), such that depending on the context, monetary or social, the effort people are willing to exert may depend (or may not depend, respectively) on the 'benefits' that they receive, to the extent that in a social context they are willing to expend more effort for no monetary return (Heyman & Ariely, 2004). One viable paradigm to explore this would be to manipulate goals as either benevolent (e.g, charity) or self-interested. Anecdotally, people are more willing to suffer (e.g., run greater distances) for 'good causes', and recent work show that once they commit to financially contribute for good causes such as a charity, their contribution increases more when the contribution process involves pain and effort than when it is enjoyable (Olivola & Shafir, under review). Future work could explore the neural circuitry which determines effort's discounting or boosting effects.

Effort and temporal discounting

Finding the exact trade-off point for rewards againts effort and time costs is what we strive for. Often we make seemingly imbalanced decisions which could be described in the following quote: "I consider that the sufferings of this present time are as nothing compared with the glory to be revealed to us." (Rom 8:18). Indeed, looking at each time point (now and later), we see cost-benefit imbalances where there is greater suffering than reward in the present, but a much greater reward than the suffering, in the future. Notwithstanding the temporal element, a cost-benefit imbalance could be seen as irrationality. However, a challenging research avenue would take into account the temporal aspect of this choice problem and examine how humans are able to, perhaps optimally, integrate effort, reward, and delay when making decisions. For example, it would be fruitful to create experimental situations where although the effort-reward trade-off *now* yields an action value that is incredibly low, people might continue persisting for the vision of a much better effort-reward trade-off *later*.

7.4 Contribution to the field

My doctoral work has contributed to the field of decision neuroscience in the following ways:

- I have developed laboratory paradigms to manipulate physical effort expenditure and control for temporal costs. These paradigms involve extensive training protocols which allow precise control on effort expenditure and representation. The trade-offs between effort and reward are easily tipped by a slight increase in reward levels, thus central to such experimental paradigms is finding the right tradeoff points. In addition, I described two fMRI paradigms that segregate BOLD signal of abstract representation of effort from signal related to motor anticipation.
- I provide converging evidence about the role of anterior cingulate cortex and striatum in effort processing in healthy humans, and these support previous findings from animal and clinical neuroscience concerning pathologies in effort-based behaviours.
- I show viability of a computational approach to capture the pavlovian relationship between affective outcomes and effort deployment.
- I report novel findings which highlight that effort just expended may have a modulatory influence on ventromedial prefrontal cortical and ventral striatal sensitivity to outcome delivery. This finding points to effort being integrated into outcome value.
- I extend the examination of cost-driven behaviour to studying pain avoidance. I show that as effort was sensitive to context manipulation such as outcome valence, pain avoidance is also sensitive to a context manipulation of relative pain magnitudes.

To conclude, the field of effort-based learning and decision making has contributed to knowledge about decisions, actions and their neural underpinnings. I have shown ways in which effort may influence choice and interact with outcomes, and how prefrontal-striatal circuitry is sensitive and influenced by the presence of effort demands. How pathologies might result from the interaction between vigour and rewards/ punishments is poorly understood. Future research programs investigating clinical populations such as those with apathy syndromes and OCD

might clarify our understanding of action and outcome. Such an eclectic and multidisciplinary approach, which takes into account non-human animal literature and healthy and clinical human research endeavours, is likely to be crucial paramount in providing an integrated framework in which to understand both healthy cognition and near-optimal actions on the one hand, and aberrant, suboptimal behaviour on the other.

Appendices

7.5 Study 1: Details of effort and reward parameters.

Here are listed effort levels in % of maximum force and reward levels in pence for each individual included in the analysis of Chapter 3, study 1. The last 10 subjects had the same experimental parameters and are analysed separately.

	GRIP										HOLD
	Effort levels % max force)					Reward Levels (pence)					
Sub no.	Eff1	Eff2	Eff3	Eff4	Eff5	Rew1	Rew2	Rew3	Rew4	Rew5	Rew
101	30	55	70	80	90	2	5	10	15	20	1
102	30	55	70	80	90	2	5	10	15	20	1
103	30	55	70	80	90	2	5	10	15	20	1
104	50	60	70	80	90	2	5	10	15	20	1
105	50	60	70	80	90	2	5	10	15	20	1
106	50	60	70	80	90	2	5	8	12	15	1
107	50	60	70	80	90	2	5	8	12	15	1
108	50	60	70	80	90	3	6	9	12	15	2
109	50	60	70	80	90	3	6	9	12	15	2
110	50	60	70	80	90	3	6	9	12	15	2
111	50	60	70	80	90	2	5	10	15	20	1
112	50	60	70	80	90	2	5	10	15	20	1
113	50	60	70	80	90	2	5	10	15	20	1
114	50	60	70	80	90	2	5	10	15	20	1
115	50	60	70	80	90	2	3	6	9	12	1
116	50	60	70	80	90	2	3	6	9	12	1
117	50	60	70	80	90	2	3	6	9	12	1
118	50	60	70	80	90	2	3	6	9	12	1
119	50	60	70	80	90	2	3	6	9	12	1
120	50	60	70	80	90	2	3	6	9	12	1
121	50	60	70	80	90	2	3	6	9	12	1
122	50	60	70	80	90	2	3	6	9	12	1
123	50	60	70	80	90	2	3	6	9	12	1
124	50	60	70	80	90	2	3	6	9	12	1
125	50	60	70	80	90	2	3	6	9	12	1

7.6 Study 2: Manipulation checks

1. For each of the four 'grip' and 1 'hold' stimuli, participants rated how much they liked a particular grip-for-money combination associated with that stimulus. They used a visual analog scale where they could slide the cursor on a bar to indicate their liking from 'I do not like it at all' to 'I like it very much'. Instructions: "You will now see the circles again. For each circle you see, please think about how much gripping and how much money associated with it and indicate HOW MUCH YOU LIKE that grip-for-money action."
2. For each low and high reward level stimuli, participants answered to the question: "How much money does the horizontal line on the circle mean?"
3. For each of the low and high effort indicated by the target line in the thermometer cue, participants answered to the question: "How much money do you think is considered a fair pay for gripping at the yellow line 10 times in a row?"

7.7 Study 2: Persistence Scale

The following are statements people might use to describe their attitudes, opinions, interests and other personal feelings. For each of the following questions, please write the number that best describes the way you generally act or feel, not just how you are feeling right now. Remember, there are no right or wrong answers, just describe your *own* personal opinions and feelings.

Response scale:

1	2	3	4	5
Definitely False	Mostly False	Neither True nor False	Mostly True	Definitely True

1. I like a challenge better than easy jobs.
2. I am usually eager to get going on any job I have to do.
3. I often give up a job if it takes much longer than I thought it would.
4. I am a very ambitious person.
5. When I fail at something at first, I become even more determined to do a better job.
6. I am usually so determined that I continue to work long after other people have given up.

7. I have often been called an "eager beaver" because of my enthusiasm for hard work.
8. I often drag my heels a while before starting any project.
9. I love to excel at everything I do.
10. I am more hard working than most people.
11. No matter how hard a job is, I like to get started quickly.
12. The harder a job is, the less I enjoy it.
13. I am eager to start work on any assigned duty.
14. I often accomplish more than people expect of me.
15. I usually push myself harder than most people do because I want to do as well as I possibly can.
16. I am never described as an overachiever.
17. If something doesn't work as I expected, I am more likely to quit than to keep going for a long time.
18. I like to strive for bigger and better things.
19. I am more of a perfectionist than most people.
20. No job is too hard for me to do my best.

7.8 Study 2: Brain activity during squeezing

Table 1 MNI coordinates of regions the activity of which is correlated with 'squeeze' period (thresholded at $p = 0.001$, unc., > 5 voxels)

Region	Nearest Brodmann Areas	Coordinates (mm)			Z value	No. of voxels	*P*
		x	y	z			
Contrast: Squeeze > Hold (Execute Period)							
Cerebellum Anterior Lobe	N/A	+15	-52	-23	4.58	27	.03 (corr.)
Primary motor cortex	4	-42	-22	+49	3.72	11	.0001 (unc.)
Caudate Nucleus	N/A	-15	+26	+1	3.54	6	.0001 (unc.)

7.9 Study 7: Additional methods

Computation for monetary trade-offs between reward and pain

I briefly summarise how to compute optimal movement strategies maximising expected gain in the context of an unspecified loss function, previous work has provided (Trommershäuser, Maloney, & Landy, 2003a, 2003b) details on how to compute gain functions in the presence of monetary rewards and penalties.
In the experiments by Trommershäuser and colleagues, participants win and lose small monetary rewards by touching a reward and penalty region on a plane before the timeout. Penalties and rewards depend only on the position of the end point in this plane, and a visuo-motor strategy S is identified with the mean end point on the plane (x,y) that results from adopting strategy S.

The model can be applied to the experiments reported in the present study as follows: the scene is divided into three regions: the circular target region (R_1) which carries a positive gain, the circular penalty region (R_2) which carries no gain or a negative gain, and the background (no gain). An optimal visuo-motor strategy S on any trial is one that maximizes the participant's expected gain

$$\Gamma(S) = \sum_{i=1}^{2} G_i \, P(R_i|S) \qquad (1)$$

Here G_i denoted the gain the participant receives if region R_i is touched within the time limit ($G_1 = 6p$ for hitting the target region R_1; $G_2 = -\alpha p$ for hitting penalty region R_2); $P(R_i|S)$ is the probability, given a particular choice of strategy S, of reaching region R_i before the time limit (t = timeout) has expired,

$$P(R_i|S) = \int_{R_i^{timeout}} P(\tau|S)\, d\tau \qquad (2)$$

and $R_i^{timeout}$ denotes the set of possible trajectories τ that pass through R_i after movement onset and before the timeout. Because the task requires a quick response (before the timeout), Eq. 1 contains a term for this timeout penalty. The probability that a visuo-motor strategy S leads to a timeout is $P(timeout|S)$.

Maximizing Eq. (1) requires knowledge of the probability of hitting each region R_i. In our experiments, movement end points are distributed around the

mean end point $(\bar{x},\bar{y})$ according to a bivariate Gaussian distribution with widths (σ_x,σ_y) (see also Trommershäuser et al., 2005, for more details about the shape of the end point distribution),

$$p(x,y \mid \bar{x},\bar{y},\sigma_x,\sigma_y) = \frac{1}{2\pi\sigma_x\sigma_y}\exp\left[-\left(x-\bar{x}\right)^2/2\sigma_x^2\right]\exp\left[-\left(y-\bar{y}\right)^2/2\sigma_y^2\right]. \quad (3)$$

The probability of hitting region R_i is then computed by integrating over region R_i,

$$P(R_i|\bar{x},\bar{y},\sigma_x,\sigma_y) = \int_{R_i} p(x,y \mid \bar{x},\bar{y},\sigma_x,\sigma_y)\,dx\,dy. \quad (4)$$

In the experiment, the probability of a timeout is effectively constant over the limited range of relevant screen locations (and effectively zero once participants are practiced in the task), so – for any given end point variance (σ_x,σ_y) – finding an optimal movement strategy corresponds to choosing a strategy with mean aim point $(\bar{x},\bar{y})$ that maximises,

$$\Gamma(x,y) = G_1 P(R_1|\bar{x},\bar{y},\sigma_x,\sigma_y) + \alpha P(R_2|\bar{x},\bar{y},\sigma_x,\sigma_y) \quad (5)$$

Under the assumption that the measured mean end points $(\bar{x},\bar{y})$ correspond to the optimal movement strategy maximizing expected gain, the solution of Eq. (5) yields an estimate of the fit parameter α. This parameter α corresponds to the penalty value (in pence) that would have resulted into the observed mean shifts in each of the spatial configurations.

7.10 Study 7: Additional results

Late trials

Late trials were equally distributed across Low-Medium, Medium-High, and Low-High blocks: 11% (.7%), 14% (.8%), 13% (.7%), respectively.

Complete three-way ANOVA results for end-point shifts

A three-way repeated measures ANOVA with Distance (2: Near vs. Far), Block (3: Low-Medium, Medium-High, Low-High), and Relative Pain (2: Lower vs. Higher) yielded a significant main effect of Distance ($F(1,14) = 66.60$, $p < .00001$, partial eta squared .82): participants displaced their end-point much farther when penalty

was near, compared to far, from the target (FIGURE 6-2B), consistent with the hypothesis that movement execution incorporates information relating to judged movement variability (noise) (Trommershäuser et al., 2008).

I also found a significant main effect of Relative Pain (Lower vs. Higher), $F(1,14) = 4.84$, $p = .045$, partial eta squared .25. Across all blocks, displacement was larger to penalty with stronger than that to milder pain magnitude. This indicates that the displacement from penalty depends on relative pain magnitudes; whether pain was *milder* or *stronger*, not on absolute pain intensities. Finally, it should be noted that the Block x Relative Pain interaction was not significant. This means the difference between Lower and Higher pain found in the main effect of Relative Pain, is comparable across three experimental blocks (FIGURE 6-2C), which suggests that end-point shift is influenced by *relative* pain intensities. No other effects were significant, $F < 3$, $p > .068$.

Complete three-way ANOVA results for RTs

A three way repeated measures ANOVA with Distance (2: Near vs. Far), Block (3: Low-Medium, Medium-High, Low-High), and Relative Pain (2: Lower vs. Higher) yielded a significant main effect of Distance on RT ($F(1,14) = 12.32$, $p = .003$, partial eta squared .46). Participants responded slower when the target was Near than Far. There was also a significant effect of Block on reaction time, $F(2,28) = 5.2$, $p = 0.012$, partial eta squared .27. Follow-up paired-samples t-tests revealed that RT's in the context of Low-High pain was significantly slower than RT's in Low-Medium and Medium-High blocks ($t(14) = 2.7$, $p = 0.017$; $t(15) = 2.23$, $p = 0.041$, respectively). Participants responded equally fast in Low-Medium and Medium-High conditions, $p > .05$.

REFERENCES

Aberman, J. E., & Salamone, J. D. (1999). Nucleus accumbens dopamine depletions make rats more sensitive to high ratio requirements but do not impair primary food reinforcement. *Neuroscience*, *92*(2), 549-552.

Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in Neurosciences*, *13*(7), 266-71. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1695401.

Andersson, J. L., Hutton, C., Ashburner, John, Turner, R., & Friston, K. J. (2001). Modeling geometric deformations in EPI time series. *NeuroImage*, *13*(5), 903-19. doi:10.1006/nimg.2001.0746.

Aronson, E. (1961). The effect of effort on the attractiveness of rewarded and unrewarded stimuli. *Journal of Abnormal and Social Psychology*, *63*, 375-380.

Ashburner, John, & Friston, K. J. (2004). Image segmentation. In F. RSJ, K. J. Friston, C. D. Frith, R. J. Dolan, C. Price, S. Zeki, John Ashburner, et al. (Eds.), *Human Brain Mapping* (Second., pp. 695–706). London: Elsevier.

Bardgett, M. E., Depenbrock, M., Downs, N., Points, M., & Green, L. (2009). Dopamine modulates effort-based decision making in rats. *Behavioral Neuroscience*, *123*(2), 242-251. doi:10.1037/a0014625.

Bautista, L. M., Tinbergen, J., & Kacelnik, A. (2001). To walk or to fly? How birds choose among foraging modes. *Proc. Natl. Acad. Sci. U.SA.*, *98*, 1089-1094.

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, *47*(1), 129-41. doi:10.1016/j.neuron.2005.05.020.

Beckman, M., Johansen-Berg, H., & Rushworth, M. F. S. (2009). Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *Journal of Neuroscience*, *29*(4), 1175-1190.

Belin, D., Mar, A. C., Dalley, J. W., Robbins, T. W., & Everitt, B. J. (2008). High impulsivity predicts the switch to compulsive cocaine-taking. *Science*, *320*(5881), 1352-5. doi:10.1126/science.1158136.

Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*, *28*(3), 309-369.

Berridge, K. C., Venier, I. L., & Robinson, T. E. (1989). Taste reactivity analysis of 6-hydroxydopamine-induced aphagia: Implications for arousal and anhedonia hypotheses of dopamine function. *Behavioral Neuroscience*, *103*, 36-45.

Bitgood, S. (2006). Not another step! Economy of movement and pedestrian choice point behavior in shopping malls. *Environment and Behavior*, *38*(3), 394-405. doi:10.1177/0013916505280081.

Bolam, J. P., Magill, P. J., & Bevan, M. D. (2002). The functional organisation of the basal ganglia: new insights from anatomical and physiological analyses. In L. Nicholson & R. Faull (Eds.), *Basal Ganglia VII* (pp. 3711-3378.). New York: Kluwer Academic/Plenum Publishers.

Botvinick, M. M., Huffstetler, S., & McGuire, J. T. (2009). Effort discounting in human nucleus accumbens. *Cognitive, Affective & Behavioral Neuroscience*, *9*(1), 16-27. doi:10.3758/CABN.9.1.16.

Boureau, Y.-L., & Dayan, P. (2011). Opponency revisited: Competition and cooperation between dopamine and serotonin. *Neuropsychopharmacology*, *36*, 74-97. doi:10.1038/npp.2010.151.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 443–446.

Matthew Brett, Jean-Luc Anton, Romain Valabregue, Jean-Baptiste Poline. Region of interest analysis using an SPM toolbox [abstract] Presented at the 8th International Conference on Functional Mapping of the Human Brain, June 2-6, 2002, Sendai, Japan. Available on CD-ROM in NeuroImage, Vol 16, No 2.

Bush, G., Luu, P., & Posner, M. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences*, *4*(6), 215-222. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10827444.

Camerer, C. F., & Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*(4), 827-874.

Cannon, C. M., & Palmiter, R. D. (2003). Reward without dopamine. *Journal of Neuroscience*, *23*(34), 10827-10831.

Carroll, L. (1865). Alice's Adventures in Wonderland, 1994 Edn. London: Penguin Classics.

Collier, G., & Levitsky, D. A. (1968). Operant running as a function of deprivation and effort. *Journal of Comparative and Physiological Psychology*, *66*, 522-523. Retrieved from PM:5722068.

Collier, G., Hirsch, E., Levitsky, D. A., & Leshner, A. I. (1975). Effort as a dimension of spontaneous activity in rats. *Journal of Comparative and Physiological Psychology*, *88*, 89-96. Retrieved from PM:1120820.

Cook, D. B., Connor, P. J. O., Ray, C. A., Meeusen, R., Carter, J. R., Monahan, K. D., & Sauder, C. L. (2000). Muscle pain perception and sympathetic nerve activity to exercise during opioid modulation. *American Journal Of Physiology*, *279*, R1565-R1573.

Cools, R., Robinson, O. J., & Sahakian, B. J. (2008). Acute tryptophan depletion in healthy volunteers enhances punishment prediction but does not affect reward prediction. *Neuropsychopharmacology*, *33*(9), 2291-9. doi:10.1038/sj.npp.1301598.

Cousins, M. S., & Salamone, J. D. (1994). Nucleus accumbens dopamine depletions in rats affect relative response allocation in a novel cost / benefit procedure. *Pharmacology, Biochemistry, and Behavior*, *49*(1), 85-91.

Craig, A. D. (2003). A new view of pain as a homeostatic emotion. *Trends in Neurosciences*, *26*(6), 303-307. doi:10.1016/S0166-2236(03)00123-1.

Critchley, H. D., Mathias, C. J., Josephs, O, O'Doherty, J. P., Zanini, S., Dewar, B.-K., Cipolotti, L., et al. (2003). Human cingulate cortex and autonomic control: converging neuroimaging and clinical evidence. *Brain*, *126*, 2139-2152.

Croxson, P. L., Walton, M. E., O'Reilly, J. X., Behrens, T. E. J., & Rushworth, M. F. S. (2009). Effort-based cost-benefit valuation and the human brain. *Journal of Neuroscience*, *29*(14), 4531-41. doi:10.1523/JNEUROSCI.4515-08.2009.

Daw, N. D. (2002). Opponent interactions between serotonin and dopamine. *Neural Networks*, *15*(4-6), 603-616. doi:10.1016/S0893-6080(02)00052-7.

Daw, N. D. (2009). Trial-by-trial data analysis using computational models. *Attention & Performance XXIII* (pp. 1-26).

Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, *16*(2), 199-204. doi:10.1016/j.conb.2006.03.006.

Dayan, P., & Seymour, B. (2008). Values and actions in aversion. In P.W. Glimcher, C.F. Camerer, E. Fehr, & R. A. Poldrack (Eds.), *Neuroeconomics: Decision making and the brain* (pp. 175–192). London: Academic Press.

Deichmann, R., Schwarzbauer, C., & Turner, R. (2004). Optimisation of the 3D MDEFT sequence for anatomical brain imaging: technical implications at 1.5 and 3 T. *NeuroImage*, *21*(2), 757-67. doi:10.1016/j.neuroimage.2003.09.062.

Denk, F., Walton, M. E., Jennings, K. A., Sharp, T., Rushworth, M. F. S., & Bannerman, D. M. (2005). Differential involvement of serotonin and dopamine systems in cost-benefit decisions about delay or effort. *Psychopharmacology*, *179*, 587-596. doi:10.1007/s00213-004-2059-4.

Duckworth, A. L., Peterson, C., Matthews, M. D., Kelly, D. R., Farah, M., Latham, G., Rozin, P., et al. (2007). Grit□: Perseverance and passion for long-term goals. *Journal of Personality and Social Psychology*, *92*(6), 1087-1101. doi:10.1037/0022-3514.92.6.1087.

Eisenberger, R., Weier, F., Masterson, F. A., & Theis, L. Y. (1989). Fixed-ratio schedules increase generalized self-control□: Preference for large rewards despite high effort or punishment. *Journal of Experimental Psychology: Animal Behavior Processes*, *15*(4), 383-392.

Eslinger, P. J., & Damasio, A. R. (1985). Severe disturbance of higher cognition after bilateral frontal lobe ablation. *Neurology*, *35*, 1731-1741.

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford: Stanford UP.

Fields, H. (2004). State-dependent opioid control of pain. *Nature Reviews Neuroscience*, *5*(7), 565-75. doi:10.1038/nrn1431.

Floresco, S. B., & Ghods-sharifi, S. (2007). Amygdala-prefrontal cortical circuitry regulates effort-based decision making. *Cerebral Cortex*, *17*, 251 - 260. doi:10.1093/cercor/bhj143.

Floresco, S. B., St Onge, J. R., Ghods-Sharifi, S., & Winstanley, C. A. (2008). Cortico-limbic-striatal circuits subserving different forms of cost-benefit decision making. *Cognitive, Affective & Behavioral Neuroscience*, *8*(4), 375-89. doi:10.3758/CABN.8.4.375.

Floresco, S. B., Tse, M. T. L., & Ghods-Sharifi, S. (2008). Dopaminergic and glutamatergic regulation of effort- and delay-based decision making. *Neuropsychopharmacology*, *33*(8), 1966-79. doi:10.1038/sj.npp.1301565.

Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*, *17*(1), 51-72. doi:10.1162/0898929052880093.

Frank, M. J., & Fossella, J. A. (2011). Neurogenetics and pharmacology of learning, motivation, and cognition. *Neuropsychopharmacology*, *36*, 133-152. Nature Publishing Group.

Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: a computational model. *Cognitive, Affective & Behavioral Neuroscience*, *1*(2), 137-60. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12467110.

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*, *306*(5703), 1940-3. doi:10.1126/science.1102941.

Friedrich, A. M., & Zentall, T. R. (2004). Pigeons shift their preference toward locations of food that take more effort to obtain. *Behavioural Processes*, *67*(3), 405-15. doi:10.1016/j.beproc.2004.07.001.

Friston, K. J., Ashburner, J, Frith, C. D., Poline, J.-P., Heather, J. D., & Frackowiak, R. S. J. (1995). Spatial registration and normalization of images. *Human Brain Mapping*, *2*(3), 165-189. doi:10.1002/hbm.460030303.

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., & Frackowiak, R. S. J. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, *2*(4), 189-210. doi:10.1002/hbm.460020402.

Gan, J. O., Walton, M. E., & Phillips, P. E. M. (2010). Dissociable cost and benefit encoding of future rewards by mesolimbic dopamine. *Nature Neuroscience*, *13*(1), 25-7. Nature Publishing Group. doi:10.1038/nn.2460.

Garner, W. R. (1954). Context effects and the validity of loudness scales. *Journal of Experimental Psychology*, *48*, 218–224.

Gepshtein, S., Seydell, A., & Trommershäuser, J. (2007). Optimality of human movement under natural variations of visual – motor uncertainty. *Journal of Vision*, *7*(5), 1-18. doi:10.1167/7.5.13.Introduction.

Ghods-Sharifi, S., & Floresco, S. B. (2010). Differential effects on effort discounting induced by inactivations of the nucleus accumbens core or shell. *Behavioral Neuroscience*, *124*(2), 179-91. doi:10.1037/a0018932.

Guitart-Masip, M., Fuentemilla, L., Bach, D. R., Huys, Q. J. M., Dayan, P., Dolan, R. J., & Duzel, E. (2011). Action dominates valence in anticipatory representations in the

human striatum and dopaminergic midbrain. *Journal of Neuroscience*, *31*(21), 7867-7875. doi:10.1523/JNEUROSCI.6376-10.2011.

Guitart-Masip, M., Talmi, D., & Dolan, R. J. (2010). Conditioned associations and economic decision biases. *NeuroImage*, *53*(1), 206-214. Elsevier Inc. doi:10.1016/j.neuroimage.2010.06.021.

Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetics*, *84*(6), 401-10. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11417052.

Gusnard, D. A., Ollinger, J. M., Shulman, G. L., Cloninger, C. R., Price, J. L., Van Essen, D. C., & Raichle, M. E. (2003). Persistence and brain circuitry. *Proceedings of the National Academy of Sciences of the United States of America*, *100*(6), 3479-84. doi:10.1073/pnas.0538050100.

Haber, S. N., & Knutson, B. (2010). The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology*, *35*(1), 4-26. Nature Publishing Group. doi:10.1038/npp.2009.129.

Haber, S. N., Fudge, J. L., & McFarland, N. R. (2000). Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience*, *20*(6), 2369-82. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10704511.

Hadland, K. A., Rushworth, M. F. S., Gaffan, D., & Passingham, R. E. (2003). The anterior cingulate and reward-guided selection of actions. *Journal of Neurophysiology*, *89*(2), 1161-4. doi:10.1152/jn.00634.2002.

Hershberger, W. A. (1986). An approach through the looking-glass. *Animal Learning and Behavior*, *14*(4), 443-451.

Heyman, J., & Ariely, D. (2004). Effort for payment. A tale of two markets. *Psychological Science*, *15*(11), 787-93. doi:10.1111/j.0956-7976.2004.00757.

Hollander, D. B., Reeves, G. V., Clavier, J. D., Francois, M. R., Thomas, C., & Kraemer, R. R. (2010). Partial occlusion during resistance exercise alters effort sense and pain. *Journal of Strength and Conditioning Research*, *24*(1), 235-243.

Huettel, S. A., Song, A. W., & McCarthy, G. (2009). *Functional magnetic resonance imaging* (Second.). Sunderland, MA: Sinauer Associates, Inc.

Humphries, M. D., & Prescott, T. J. (2010). The ventral basal ganglia, a selection mechanism at the crossroads of space, strategy, and reward. *Progress in Neurobiology*, *90*(4), 385-417. doi:10.1016/j.pneurobio.2009.11.003.

Huys, Q. J. M., Cools, R., Gölzer, M., Friedel, E., Heinz, A., Dolan, R. J., & Dayan, P. (2011). Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding. (A. Rangel, Ed.)*PLoS Computational Biology*, *7*(4), 1-14. doi:10.1371/journal.pcbi.1002028.

Ikemoto, S. (2007). Dopamine reward circuitry: two projection systems from the ventral midbrain to the nucleus accumbens-olfactory tubercle complex. *Brain Research Reviews*, *56*(1), 27-78. doi:10.1016/j.brainresrev.2007.05.004.

Johnson, A. W., & Gallagher, M. (2010). Greater effort boosts the affective taste properties of food. *Proceedings of the Royal Society: Biological sciences*, (November). doi:10.1098/rspb.2010.1581.

Kable, J. W., & Glimcher, Paul W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625-1633. doi:10.1038/nn2007.

Kanarek, R. B., & Collier, G. (1973). Effort as a determinant of choice in rats. *Journal of Comparative and Physiological Psychology*, *84*, 332-338. Retrieved from PM:4723929.

Karason, K., Peltonen, M., Lindroos, A. K., Sjöström, L., Lönn, L., & Torgerson, J. S. (2005). Effort-related calf pain in the obese and long-term changes after surgical obesity treatment. *Obesity Research*, *13*(1), 137-45. doi:10.1038/oby.2005.18.

Kelley, A. E., Baldo, B. a, Pratt, W. E., & Will, M. J. (2005). Corticostriatal-hypothalamic circuitry and food motivation: integration of energy, action and reward. *Physiology & Behavior*, *86*(5), 773-95. doi:10.1016/j.physbeh.2005.08.066.

Kennerley, S. W., Dahmubed, A. F., Lara, A. H., & Wallis, J. D. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. *Journal of Cognitive Neuroscience*, *21*(6), 1162-1178.

Kennerley, S. W., Walton, M. E., Behrens, T. E. J., Buckley, M. J., & Rushworth, M. F. S. (2006). Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience*, *9*, 940-947.

Kim, J. W., Lee, D. Y., Choo, I. H., Seo, E. H., Kim, S. G., Park, S. Y., & Woo, J. I. (2011). Microstructural alteration of the anterior cingulum is associated with apathy in Alzheimer Disease. *American Journal of Geriatric Psychiatry*, *19*(7), 644-53. doi:10.1097/JGP.0b013e31820dcc73.

Kimball, W. R., Carwood, C. M., Chang, Y., McKenna, J. M., Peters, L. E., & Ballantyne, J. C. (2008). Effect of effort pain after upper abdominal surgery on two independent measures of respiratory function. *Journal of Clinical Anesthesia*, *20*(3), 200-5. doi:10.1016/j.jclinane.2007.10.009.

Knutson, B., Taylor, J., Kaufman, M., Peterson, R., & Glover, G. (2005). Distributed neural representation of expected value. *Journal of Neuroscience*, *25*(19), 4806-12. doi:10.1523/JNEUROSCI.0642-05.2005.

Kool, W., Mcguire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, *139*(4), 665- 682. doi:10.1037/a0020198.

Kunishio, K., & Haber, S. N. (1994). Primate cingulostriatal projection: limbic striatal versus sensorimotor striatal input. *Journal of Comparative Neurology*, *390*(3), 337-356.

Kurniawan, I. T., Seymour, B., Talmi, D., Yoshida, W., Chater, N., & Dolan, R. J. (2010). Choosing to make an effort: the role of striatum in signaling physical effort of a chosen action. *Journal of Neurophysiology*, *104*(1), 313-21. doi:10.1152/jn.00027.2010.

Laming, D. (1984). The relativity of "absolute" judgements. *British Journal of Mathematical and Statistical Psychology*, *37*, 152–183.

Laming, D. (1997). *The measurement of sensation*. New York: Oxford University Press.

Lau, B., & Glimcher, Paul W. (2007). Action and outcome encoding in the primate caudate nucleus. *Journal of Neuroscience*, *27*(52), 14502-14. doi:10.1523/JNEUROSCI.3060-07.2007.

Lau, B., & Glimcher, Paul W. (2008). Value representations in the primate striatum during matching behavior. *Neuron*, *58*(3), 451-63. doi:10.1016/j.neuron.2008.02.021.

Lewis, M. (1964). Effect of effort on value: An exploratory study of children. *Child Development*, *35*(4), 1337-1342.

Li, M., & Fleming, A. (2003). The nucleus accumbens shell is critical for normal expression of pup-retrieval in postpartum female rats. *Behavioural Brain Research*, *145*(1-2), 99-111. doi:10.1016/S0166-4328(03)00135-9.

Links, K. A., Chow, T. W., Binns, M., Freedman, M., Stuss, D. T., Scott, C. J. M., Ramirez, J., et al. (2009). Apathy is not associated with basal ganglia atrophy in frontotemporal dementia. *American Journal of Geriatric Psychiatry*, *19*(9), 819-821.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, *453*(7197), 869-78. doi:10.1038/nature06976.

Lohman, E. B., Thorpe, D. L., Prior, M., George, J., & Kim, J. P. (2008). Can APB 2000 be used to discern sincerity of effort in unimpaired subjects from maximal performance in subjects with shoulder pain? *Journal of Forensic Sciences*, *53*(2), 392-6. doi:10.1111/j.1556-4029.2008.00670.x

Lévy, R., & Dubois, B. (2006). Apathy and the functional anatomy of the prefrontal cortex-basal ganglia circuits. *Cerebral Cortex*, *16*(7), 916-28. doi:10.1093/cercor/bhj043.

Maloney, L. T., Trommershäuser, J., & Landy, M. S. (2007). Questions without words□: A comparison between decision making under risk and movement planning under risk. In W. Gray (Ed.), *Integrated models of cognitive systems* (pp. 297-314). New York: Oxford University Press.

Marchand, W. R., Lee, J. N., Thatcher, J. W., Hsu, E. W., Rashkin, E., Suchy, Y., Chelune, G., et al. (2008). Putamen coactivation during motor task execution. *Neuroreport*, *19*(9), 957-60. doi:10.1097/WNR.0b013e328302c873.

Marsh, B., Schuck-paim, C., & Kacelnik, A. (2004). Energetic state during learning affects foraging choices in starlings. *Behavioral Ecology*, *15*(3), 396-399. doi:10.1093/beheco/arh034.

Matsumoto, M., & Hikosaka, O. (2009). Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature*, *459*(7248), 837-41. doi:10.1038/nature08028.

McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, *38*(2), 339-46. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12718866.

McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience*, *27*(21), 5796 -5804. doi:10.1523/JNEUROSCI.4246-06.2007.

McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(17), 7922-6. doi:10.1073/pnas.0910662107.

McMahon, S. B., & Koltzenburg, M. (2005). *Wall and Melzack's textbook of pain* (Fifth.). Philadelphia: Elsevier.

Mobbs, D., Petrovic, P., Marchant, J. L., Hassabis, D., Weiskopf, N., Seymour, B., Dolan, R. J., et al. (2007). When fear is near: threat imminence elicits prefrontal-periaqueductal gray shifts in humans. *Science*, *317*(5841), 1079-83. doi:10.1126/science.1144298.

Montague, P. R., & Berns, G. S. (2002). Neural economics and the biological substrates of valuation. *Neuron*, *36*(2), 265-84. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12383781.

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic predictive Hebbian learning. *Journal of Neuroscience*, *16*(5), 1936-1947.

Morris, G., Nevet, A., Arkadir, D., Vaadia, E., & Bergman, H. (2006). Midbrain dopamine neurons encode decisions for future action. *Nature Neuroscience*, *9*(8), 1057-63. doi:10.1038/nn1743.

Mosteller, F., & Nogee, P. (1951). An experimental measurement of utility. *Journal of Political Economy*, *59*(5), 371-404.

Naccache, L., Dehaene, S., Cohen, L., Habert, M.-O., Guichart-Gomez, E., Galanaud, D., & Willer, J.-C. (2005). Effortless control: executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia*, *43*(9), 1318-28. doi:10.1016/j.neuropsychologia.2004.11.024.

Nicola, S. M. (2007). The nucleus accumbens as part of a basal ganglia action selection circuit. *Psychopharmacology*, *191*(3), 521-50. doi:10.1007/s00213-006-0510-4.

Nicola, S. M. (2010). The flexible approach hypothesis: Unification of effort and cue-responding hypotheses for the role of nucleus accumbens dopamine in the activation of reward-seeking behavior. *Journal of Neuroscience*, *30*(49), 16585-16600. doi:10.1523/JNEUROSCI.3958-10.2010.

Niv, Y., Daw, N. D., & Dayan, P. (2005). How fast to work : Response vigor , motivation and tonic dopamine. In Y. Weiss, B. Scholkopf, & J. Platt (Eds.), *Neural Information Processing Systems* (pp. 1019-1026). MIT Press.

Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology*, *191*(3), 507-20. doi:10.1007/s00213-006-0502-4.

Niznikiewicz, M. a, & Delgado, M. R. (2011). Two sides of the same coin: Learning via positive and negative reinforcers in the human striatum. *Developmental Cognitive Neuroscience*, 1-12. doi:10.1016/j.dcn.2011.07.006.

Olivola, C. Y. (2010). When noble means hinder noble ends: The benefits and costs of a human preference for martyrdom in altruism. In D. M. Oppenheimer & C. Y. Olivola (Eds.), *The science of giving: Experimental approaches to the study of charity* (pp. 49-62). New York: Taylor and Francis.

O'Doherty, J. P., Dayan, P., Friston, K. J., Critchley, H. D., & Dolan, R. J. (2003). Temporal Difference Models and Reward-Related Learning in the Human Brain. *Neuron*, *38*(2), 329-337. doi:10.1016/S0896-6273(03)00169-7.

O'Doherty, J. P., Dayan, P., Schultz, J., Deichmann, R., Friston, K. J., & Dolan, R. J. (2004). Dissociable Roles of Ventral and Dorsal Striatum in Instrumental Conditioning. *Science*, *304*, 452 - 454. doi:10.1126/science.1094285.

Olivola, C. Y., & Shafir, E. (under review). The martyrdom effect: When pain and effort increase prosocial contributions.

Parkinson, J. A., Olmstead, M. C., Burns, L. H., Robbins, T. W., & Everitt, B. J. (1999). Dissociation in effects of lesions of the nucleus accumbens core and shell on appetitive pavlovian approach behavior and the potentiation of conditioned reinforcement and locomotor activity by D-amphetamine. *Journal of Neuroscience*, *19*(6), 2401-11. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10066290.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437-442.

Pennartz, C. M. A., Groenewegen, H. J., & Lopes da Silva, F. H. (1994). The nucleus accumbens as a complex of functionally distinct neuronal ensembles: an integration of behavioural, electrophysiological and anatomical data. *Progress in Neurobiology*, *42*(6), 719-61. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7938546.

Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., & Frith, C. D. (2007). How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science*, *316*, 904-906. doi:10.1126/science.1140459.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*(7106), 1042-5. doi:10.1038/nature05051.

Phillips, P. E. M., Walton, M. E., & Jhou, T. C. (2007). Calculating utility: preclinical evidence for cost-benefit analysis by mesolimbic dopamine. *Psychopharmacology*, *191*, 483-495.

Picard, N., & Strick, P. L. (2001). Imaging the premotor areas. *Current Opinion in Neurobiology*, *11*, 663-672.

Pierce, C. A., Block, R. A., & Aguinis, H. (2004). Cautionary note on reporting eta-squared values from multifactor ANOVA designs. *Educational and Psychological Measurement*, *64*, 916-924.

Pine, A., Shiner, T., Seymour, B., & Dolan, R. J. (2010). Dopamine, time, and impulsivity in humans. *Journal of Neuroscience*, *30*(26), 8888-96. doi:10.1523/JNEUROSCI.6028-09.2010.

Prevost, C., Pessiglione, M., Metereau, E., Clery-Melin, M.-L., & Dreher, J.-C. (2010). Separate valuation subsystems for delay and effort decision costs. *Journal of Neuroscience*, *30*(42), 14080 -14090. doi:10.1523/JNEUROSCI.2752-10.2010.

Price, D. D. (2000). Psychological and neural mechanisms of the affective dimension of pain. *Science*, *288*(5472), 1769-1772. doi:10.1126/science.288.5472.1769.

Prodoehl, J., Corcos, D. M., & Vaillancourt, D. E. (2009). Basal ganglia mechanisms underlying precision grip force control. *Neuroscience and Biobehavioral Reviews*, *33*(6), 900-8. doi:10.1016/j.neubiorev.2009.03.004.

Rangel, A., Camerer, Colin F, & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, *9*(June), 545-556. doi:10.1038/nrn2357.

Redgrave, P., & Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nature Reviews Neuroscience*, *7*(12), 967-75. doi:10.1038/nrn2022.

Redgrave, P., Prescott, T. J., & Gurney, K. (1999). Is the short-latency dopamine response too short to signal reward error□? *Trends in Neurosciences*, *2236*(98), 146-151.

van Reekum, R., Stuss, D. T., & Ostrander, L. (2005). Apathy: why care? *Journal of Neuropsychiatry and Clinical Neurosciences*, *17*(1), 7-19. doi:10.1176/appi.neuropsych.17.1.7.

Reynolds, J. N., Hyland, B. I., & Wickens, J. R. (2001). A cellular mechanism of reward-related learning. *Nature*, *413*(6851), 67-70. doi:10.1038/35092560.

Reynolds, S. M., & Berridge, K. C. (2002). Positive and negative motivation in nucleus accumbens shell: bivalent rostrocaudal gradients for GABA-elicited eating, taste "liking"/"disliking" reactions, place preference/avoidance, and fear. *Journal of Neuroscience*, *22*(16), 7308-20. doi:20026734.

Robinson, D. L., Venton, B. J., Heien, M. L. A. V., & Wightman, R. M. (2003). Detecting subsecond dopamine release with fast-scan cyclic voltammetry in vivo. *Clinical Chemistry*, *49*(10), 1763-73. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/14500617.

Roesch, M. R., Calu, D. J., & Schoenbaum, G. (2007). Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nature Neuroscience*, *10*(12), 1615-1624. doi:10.1038/nn2013.

Roitman, M. F., Stuber, G. D., Phillips, P. E. M., Wightman, R. M., & Carelli, R. M. (2004). Dopamine operates as a subsecond modulator of food seeking. *Journal of Neuroscience*, *24*(6), 1265-71. doi:10.1523/JNEUROSCI.3823-03.2004.

Rudebeck, P. H., Behrens, T. E. J., Kennerley, S. W., Baxter, M. G., Buckley, M. J., Walton, M. E., & Rushworth, M. F. S. (2008). Frontal cortex subregions play distinct roles in choices between actions and stimuli. *Journal of Neuroscience*, *28*(51), 13775-85. doi:10.1523/JNEUROSCI.3541-08.2008.

Rudebeck, P. H., Walton, M. E., Smyth, A. N., Bannerman, D. M., & Rushworth, M. F. S. (2006). Separate neural pathways process different decision costs. *Nature Neuroscience*, *9*(9), 1161-8. doi:10.1038/nn1756.

Rushworth, M. F. S., Behrens, T. E. J., Rudebeck, P. H., & Walton, M. E. (2007). Contrasting roles for cingulate and orbitofrontal cortex in decisions and social behaviour. *Trends in Cognitive Sciences*, *11*(4), 168-176. doi:10.1016/j.tics.2007.01.004.

Rushworth, M. F. S., Paus, T., & Sipila, P. K. (2001). Attention systems and the organization of the human parietal cortex. *Journal of Neuroscience*, *21*(14), 5262-71. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11438601.

Salamone, J. D., & Correa, M. (2002). Motivational views of reinforcement: implications for understanding the behavioral functions of nucleus accumbens dopamine. *Behavioural Brain Research*, *137*(1-2), 3-25. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12445713.

Salamone, J. D., Correa, M., Farrar, A., & Mingote, S. (2007). Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology*, *191*, 461-482.

Salamone, J. D., Correa, M., Mingote, S., & Weber, S. M. (2003). Nucleus accumbens dopamine and the regulation of effort in food-seeking behavior: implications for studies of natural motivation. *Journal of Pharmacology and Experimental Therapeutics*, *305*(1), 1-8.

Salamone, J. D., Cousins, M. S., & Bucher, S. (1994). Anhedonia or anergia? Effects of haloperidol and nucleus accumbens dopamine depletion on instrumental response selection in a T-maze cost/benefit procedure. *Behavioural Brain Research*, *65*, 221-229.

Salomon, L., Lanteri, C., Glowinski, J., & Tassin, J.-P. (2006). Behavioral sensitization to amphetamine results from an uncoupling between noradrenergic and serotonergic neurons. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(19), 7476-7481.

Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of action-specific reward values in the striatum. *Science*, *310*(5752), 1337-40. doi:10.1126/science.1115270.

Sarinopoulos, I., Dixon, G. E., Short, S. J., Davidson, R. J., & Nitschke, J. B. (2006). Brain mechanisms of expectation associated with insula and amygdala response to aversive taste: implications for placebo. *Brain, Behavior, and Immunity*, *20*(2), 120-32. doi:10.1016/j.bbi.2005.11.006.

Schmidt, L., D'Arc, B. F., Lafargue, G., Galanaud, D., Czernecki, V., Grabli, D., Schüpbach, M., et al. (2008). Disconnecting force from money: effects of basal ganglia damage on incentive motivation. *Brain*, *131*(Pt 5), 1303-10. doi:10.1093/brain/awn045.

Schneider, J. S. (2007). Behavioral persistence deficit in Parkinson's disease patients. *European Journal of Neurology*, *14*(3), 300-4. doi:10.1111/j.1468-1331.2006.01647.x

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593-9. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9054347.

Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, *6*(2), 461-464.

Seymour, B., & McClure, S. M. (2008). Anchors, scales and the relative coding of value in the brain. *Current Opinion in Neurobiology*, *18*, 1-6. doi:10.1016/j.conb.2008.07.010.

Seymour, B., Daw, N. D., Dayan, P., Singer, T., & Dolan, R. J. (2007). Differential encoding of losses and gains in the human striatum. *Journal of Neuroscience*, *27*(18), 4826-31. doi:10.1523/JNEUROSCI.0400-07.2007.

Seymour, B., Doherty, J. P. O., Dayan, P., Koltzenburg, M., Jones, A. K., Dolan, R. J., Friston, K. J., et al. (2004). Temporal difference models describe higher-order learning in humans. *Nature*, *429*(June), 664-667. doi:10.1038/nature02636.1.

Seymour, B., Singer, T., & Dolan, R. J. (2007). The neurobiology of punishment. *Nature Reviews Neuroscience*, *8*(4), 300-11. doi:10.1038/nrn2119.

Shackman, A. J., Salomons, T. V., Slagter, H. a, Fox, A. S., Winter, J. J., & Davidson, R. J. (2011). The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nature Reviews Neuroscience*, *12*(3), 154-167. doi:10.1038/nrn2994.

Shima, K., & Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science*, *282*(5392), 1335-1338. doi:10.1126/science.282.5392.1335.

Stalnaker, T. a, Calhoon, G. G., Ogawa, M., Roesch, M. R., & Schoenbaum, G. (2010). Neural correlates of stimulus-response and response-outcome associations in dorsolateral versus dorsomedial striatum. *Frontiers in Integrative Neuroscience*, *4*(May), 12. doi:10.3389/fnint.2010.00012.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, *46*(4), 1004-17. Elsevier Inc. doi:10.1016/j.neuroimage.2009.03.025.

Stevens, J. R., Rosati, A. G., Ross, K. R., & Hauser, M. D. (2005). Will travel for food: spatial discounting in two new world monkeys. *Current Biology*, *15*, 1855-1860.

Stevens, S. S. (1957). On the psychophysical law. *The Psychological Review*, *64*(3), 153-181.

Stroop, J. R. (1935). Studies of intereference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643-662.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Talairach, J., & Tournoux, P. (1988). *Co-Planar Stereotaxic Atlas of the Human Brain*. Stuttgart: Thieme.

Talmi, D., Dayan, P., Kiebel, S. J., Frith, C. D., & Dolan, R. J. (2009). How humans integrate the prospects of pain and reward during choice. *Journal of Neuroscience*, *29*(46), 14617-26. doi:10.1523/JNEUROSCI.2026-09.2009.

Talmi, D., Seymour, B., Dayan, P., & Dolan, R. J. (2008). Human Pavlovian – Instrumental Transfer. *Journal of Neuroscience*, *28*(2), 360 -368. doi:10.1523/JNEUROSCI.4028-07.2008.

Tam, G. Y. T., & Yeung, S. S. (2006). Perceived effort and low back pain in non-emergency ambulance workers: Implications for rehabilitation. *Journal of Occupational Rehabilitation*, *16*(2), 231-40. doi:10.1007/s10926-006-9019-2.

Tekin, S., & Cummings, J. L. (2002). Frontal-subcortical neuronal circuits and clinical neuropsychiatry: an update. *Journal of Psychosomatic Research*, *53*, 647-654.

Tobler, P. N., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *Journal of Neurophysiology*, *97*(2), 1621-32. doi:10.1152/jn.00745.2006.

Tracey, I., & Mantyh, P. W. (2007). The cerebral signature for pain perception and its modulation. *Neuron*, *55*(3), 377-91. doi:10.1016/j.neuron.2007.07.012.

Trommershäuser, J., Gepshtein, S., Maloney, L. T., Landy, M. S., & Banks, M. S. (2005). Optimal compensation for changes in task-relevant movement variability. *Journal of Neuroscience*, *25*(31), 7169 -7178. doi:10.1523/JNEUROSCI.1906-05.2005.

Trommershäuser, J., Landy, M. S., & Maloney, L. T. (2006). Humans rapidly estimate expected gain in movement planning. *Psychological Science*, *17*(11), 981-988.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America. A*, *20*(7), 1419-33. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12868646.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision*, *16*(3-4), 255-75. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12858951.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, *12*, 291-7. doi:10.1016/j.tics.2008.04.010.

Vanderschuren, L. J. M. J., & Everitt, B. J. (2004). Drug seeking becomes compulsive after prolonged cocaine self-administration. *Science*, *305*(5686), 1017-9. doi:10.1126/science.1098975.

Vitay, J., & Hamker, F. H. (2010). A computational model of Basal Ganglia and its role in memory retrieval in rewarded visual memory tasks. *Frontiers in Computational Neuroscience*, *4*(13), 1-18. doi:10.3389/fncom.2010.00013.

Vlaev, I., Seymour, B., Dolan, R. J., & Chater, N. (2009). The price of pain and the value of suffering. *Psychological Science*, *20*(3), 309-317.

Vogt, B. A. (2005). Pain and emotion interactions in subregions of the cingulate gyrus. *Nature Reviews Neuroscience*, *6*(7), 533-44. doi:10.1038/nrn1704.

Voorn, P., Vanderschuren, L. J. M. J., Groenewegen, H. J., Robbins, T. W., & Pennartz, C. M. A. (2004). Putting a spin on the dorsal – ventral divide of the striatum. *Trends in Neurosciences*, *27*(8), 468 - 474. doi:10.1016/j.tins.2004.06.006.

Walton, M. E., Bannerman, D. M., & Rushworth, M. F. S. (2002). The role of rat medial frontal cortex in effort-based decision making. *Journal of Neuroscience*, *22*(24), 10996-1003. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12486195.

Walton, M. E., Croxson, P. L., Rushworth, M. F. S., & Bannerman, D. M. (2005). The mesocortical dopamine projection to anterior cingulate cortex plays no role in guiding effort-related decisions. *Behavioral Neuroscience*, *119*(1), 323-328. doi:10.1037/0735-7044.119.1.323.

Walton, M. E., Devlin, J. T., & Rushworth, M. F. S. (2004). Interactions between decision making and performance monitoring within prefrontal cortex. . *Nature Neuroscience*, *7*(11), 1259-1265.

Walton, M. E., Groves, J., Jennings, K. A., Croxson, P. L., Sharp, T., Rushworth, M. F. S., & Bannerman, D. M. (2009). Comparing the role of the anterior cingulate cortex and 6-hydroxydopamine nucleus accumbens lesions on operant effort-based decision making. *European Journal of Neuroscience*, *29*, 1678-1691. doi:10.1111/j.1460-9568.2009.06726.x

Walton, M. E., Kennerley, S. W., Bannerman, D. M., Phillips, P. E. M., & Rushworth, M. F. S. (2006, October). Weighing up the benefits of work: behavioral and neural analyses of effort-related decision making. *Neural Networks*. doi:10.1016/j.neunet.2006.03.005.

Wasserman, E. a, Franklin, S. R., & Hearst, E. (1974). Pavlovian appetitive contingencies and approach versus withdrawal to conditioned stimuli in pigeons. *Journal of Comparative and Physiological Psychology*, *86*(4), 616-27. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/4823236.

Weiskopf, N., Hutton, C., Josephs, Oliver, & Deichmann, R. (2006). Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *Neuroimage*, *33*(2), 493-504. doi:10.1016/j.neuroimage.2006.07.029.

Wickens, J. R., Budd, C. S., Hyland, B. I., & Arbuthnott, G. W. (2007). Striatal contributions to reward and decision making: making sense of regional variations in a reiterated processing matrix. *Annals of the New York Academy of Sciences*, *1104*, 192-212. doi:10.1196/annals.1390.016.

Williams, D. R., & Williams, H. (1969). Auto-maintenance in the pigeon: Sustained pecking despite contingent non-reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*(4), 511-20. Retrieved from http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1338642&tool=pmcentrez&rendertype=abstract.

Wise, R. A. (1980). The dopamine synapse and the notion of ' pleasure centers ' in the brain. *Trends in Neurosciences*, *3*, 91-95.

Wojtyla, K. (1960). *Love and responsibility, 1981 English Translation*. London: William Collins Sons & Co.

Wu, S.-W., Delgado, M. R., & Maloney, L. T. (2009). Economic decision-making compared with an equivalent motor task. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(15), 6088-93. doi:10.1073/pnas.0900102106.