THE ENCODING OF UTILITY AND ITS MODULATION BY PSYCHOSTIMULANTS IN RAT ANTERIOR CINGULATE CORTEX

SAEEDEH HASHEMNIA
Master of Science, **Isfahan University of Technology, 2013**

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

DOCTOR OF PHILOSOPHY

in

NEUROSCIENCE

Department of Neuroscience
University of Lethbridge
LETHBRIDGE, ALBERTA, CANADA

THE ENCODING OF UTILITY AND
ITS MODULATION BY PSYCHOSTIMULANTS
IN RAT ANTERIOR CINGULATE CORTEX

SAEEDEH HASHEMNIA

Date of Defence: August 15, 2019

Dr. Aaron J. Gruber	Associate Professor	Ph.D.
Dr. David R. Euston	Associate Professor	Ph.D.
Thesis Co-Supervisors		
Dr. Masami Tatsuno	Associate Professor	Ph.D.
Thesis Examination Committee Member		
Dr. Javid Sadr	Associate Professor	Ph.D.
Internal External Examiner		
Department of Psychology		
Faculty of Arts and Science		
Dr. Stephen L. Cowen	Associate Professor	Ph.D.
External Examiner		
University of Arizona		
Tucson, Arizona		
Dr. Ian Whishaw	Professor	Ph.D.
Chair, Thesis Examination Committee		

DEDICATION

To Mom and Dad, whose support and patience has shaped who I am today.

تقدیم به پدر و مادرم، که صبر و حمایتشان مرا آنگونه ساخت که امروز هستم.

ABSTRACT

The Anterior Cingulate Cortex (ACC) encodes many decision variables, including effort and reward. This region contains dense dopamine innervations and communicates with several brain regions known to participate in motivation. Lesion of the rat ACC and administration of d-amphetamine (AMPH), which increases extracellular catecholamines, affect choice behaviours. Here, I first investigated neural encoding in rat ACC in a binary choice task with varying amounts of effort and reward. Second, I examined the effects of AMPH on the neural encoding of task-relevant information. In addition to path, effort, and reward, I found that ACC neurons encode the utility. AMPH appeared to significantly disturb this information by decreasing reward encoding and utility signaling. Moreover, low-dose AMPH decreases population state-space volume and variations in neural trajectories; whereas, higher concentrations increase both. Our results thus demonstrate that the ACC is involved in processing the value of the ongoing action and AMPH impairs this function.

ACKNOWLEDGMENTS

First and foremost, I would like to thank my amazing supervisors, Dr. Aaron Gruber and Dr. David Euston, for mentoring, encouraging, and shaping the bigger picture with me. Throughout the past five years, they offered their support and knowledgeable insight which would direct me to find solutions to my research questions, while at the same time provided me with the space I required to develop into an independent researcher. I would also like to thank other members of my supervisory committee, Dr. Masami Tatsuno and Dr. Artur Luczak, for contributing their invaluable expertise and time.

A special thank you goes to Victorita Ivan, Ali Mashhoori, and other past and present members of the Gruber lab for all discussions that helped me through many parts of my research. We spent countless hours working next to each other, shared many happy and sad moments, and grew a friendship that lasts forever. I would also like to further extend my gratitude to all my colleagues at the Canadian Centre for Behavioural Neuroscience (CCBN) that have directly and indirectly contributed to my research. A special thanks to Amanda Mauthe-Kaddoura and Naomi Cramer for all their supports which made our graduate life much easier.

Special appreciation goes to my loving husband, Navvab Afrashteh, for having ultimate faith in me and being there at every step. My being truly loves yours! I am deeply thankful and blessed for my beloved family, Dad, Mom, Samaneh, Farzaneh, Mohsen, Hamed, Samin, Sina, Zeinab, Sormeh, Samad, Masoud, Pouria, Kazem, Saeedeh, and Mehrad. Much appreciation goes to my special, unique, and amazing friends: Mahsa, Robab, Hossein, Behroo, Mojtaba, Megan, Soroush, Parastoo, Milad,

Nahal, Reihaneh, Azam, Alireza, Donya, Bahareh, Amin, Afrooz, Zahra, Elaheh, Ali, Sara, Elham, Sajedeh, Atefeh, Michael, Sorina, Scott, Sienna, Raj, Cecilia, Sergey, Sam, Qandeel, Shaghayegh, Nima, Nasrin, Jamshid, Valerie, Adam, Maurice, Surjeet, Leo, and Kartik. They all have supported me, uplifted me and brought joy to my soul in the journey of graduate study.

I would also like to thank Dr. Uri Eden (from Boston University), Dr. Zhe Chen (from New York University), and Dr. Greg Vilk for being willing to share their insights, suggestions, and time with me. A great thank you goes to the Natural Sciences and Engineering Research Council of Canada (NSERC) CREATE Research Training Program for generously funding this research project. Another big thank you goes to my supervisors, the University of Lethbridge, and Graduate Student Association for their financial support of my education.

Saeedeh Hashemnia

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF ABREVIATIONS

ACC	Anterior Cingulate Cortex
AMPH	d-Amphetamine
AP	Antero-Posterior
ANOVA	Analysis of Variance
DA	Dopamine
FBN	Fischer Brown Norway
GPFA	Gaussian Process Factor Analysis
LE	Long_Evans
MCC	Mid-Cingulate Cortex
ML	Medio-Lateral
mPFC	Medial PreFrontal Cortex
NAc	Nucleus Accumbens
PBS	Phosphate-Buffered Saline
PCA	Principal Component Analysis
PC	Principal Component
PFA	Paraformaldehyde
RM- ANOVA	Repeated Measure Analysis of Variance
SMA	Supplementary Motor Area
VS	Ventral Striatum

1 CHAPTER1: General Introduction

The Anterior Cingulate Cortex (ACC) has been implicated in several aspects of mentation including emotional regulation, memory retrieval, and decision-making (Euston *et al.*, 2012; Heilbronner & Hayden, 2016). In this thesis, I focus on its role in making a decision. Early human neuroimaging studies revealed ACC activation when participants engaged in voluntary actions based on internal states and not external cues. Hence, these studies suggested that the ACC is responsible for internally generated actions (Deiber *et al.*, 1991; Frith *et al.*, 1991; Passingham *et al.*, 2010). Parallel to these studies, another group of researchers found ACC involvement in error processing (Holroyd & Coles, 2002). However, this function was later challenged by a finding showing that ACC activity increases in response to the need to change strategy, even when no error has yet been experienced (Amiez *et al.*, 2005). Subsequent lines of research have proposed that the ACC signals the occurrence of conflicts in information processing, or it may signal other areas to increase attention or other processing resources when a previously successful response strategy is no longer optimal (Botvinick *et al.*, 2001; Kerns *et al.*, 2004; Durstewitz *et al.*, 2010). The later proposal, in particular, indicates that ACC activity does not simply reflect conflict, but rather specifies the appropriate amount of control that should be invested at a specific time and location (Shenhav *et al.*, 2013). Another line of research has revealed that the ACC plays a key role in effort-reward based decision-making (From now on, effort refers to any physical activity, such as barrier climbing or lever pressing, that the subject is required to exert in order to receive some reward which refers to food). This research follows from the observation that damage to the ACC in rats induces a significant shift in their choice preference among options with different utilities (i.e. net value). Intact rats typically

prefer to climb a barrier to obtain a larger reward, indicating that it has higher utility to them, rather than collecting an easier-to-obtain small reward. Rats with ACC lesions however prefer the easy option (Walton *et al.*, 2002; Walton *et al.*, 2003; Schweimer & Hauber, 2005).

Despite the wealth of data on ACC function, an overarching theory for its specific role in choice remains elusive. Nonetheless, one proposed theory posits that its general function is to serve as a link between events, context, and corresponding adaptive responses (Euston *et al.*, 2012). Along similar lines, it has also been proposed that the ACC not only encodes the task state (context), but also guides behaviour (strategy) (Heilbronner & Hayden, 2016). Several studies have provided supportive evidence to some aspects of these proposals. For instance, the ACC has been shown to encode the rats' progression toward goals (Ma *et al.*, 2014). Our recent findings demonstrated that spatiotemporal information about the rat's current state on a well-learned track could be extracted with high accuracy from ACC spiking activity (Mashhoori *et al.*, 2018). Also, ACC neurons exhibit context-dependent information such as sensory cues in the environment, time, and movement patterns (Hyman *et al.*, 2012). All of these findings suggest that the ACC encodes a map of task contextual space (Lapish *et al.*, 2008; Mashhoori *et al.*, 2018). However, more research is required to elucidate whether the ACC is both necessary and sufficient for the proposed roles.

The foundation of many of these theories is the flexibility of PFC neurons responding to a broad range of task-relevant features (Duncan, 2001). In decision-making tasks, ACC neurons of both primates and rodents exhibit selective activity to the manipulation of costs and benefits of a choice. Electrophysiological studies revealed that a considerable fraction of ACC neurons encode reward (Kennerley & Wallis, 2009b;

Amemori & Graybiel, 2012; Cowen *et al.*, 2012; Blanchard *et al.*, 2015) and the cost of receiving reward in terms of aversive stimuli (Amemori & Graybiel, 2012), physical effort (Kennerley & Wallis, 2009b; Cowen *et al.*, 2012; Hosokawa *et al.*, 2013), or competitive effort (Hillman & Bilkey, 2012). Moreover, the expected utility (Hillman & Bilkey, 2010; Amemori & Graybiel, 2012), and some other related variables such as fictive reward, chosen path, and movement direction (Hayden *et al.*, 2009; Kennerley & Wallis, 2009b; Cai & Padoa-Schioppa, 2012; Cowen *et al.*, 2012; Blanchard *et al.*, 2015) were shown to be encoded by the ACC. However, there is an ongoing question as to whether this information is directly used in making a choice. Indeed, most evidence implies that ACC activation occurs primarily after the decision, rather than before. For instance, its response to the chosen reward probability, the amount of reward, and its associated cost has been reported either after the behavioural indications of the decision become evident (Cowen *et al.*, 2012; Blanchard & Hayden, 2014) or before that, but it is still argued to be post-decisional (Kennerley *et al.*, 2009, Cai and Paoda 2012). Some studies have reported pre-decisional choice-related signals in nearby regions, such as supplementary motor areas, while medial prefrontal activity was observed later in decision process (Wunderlich *et al.*, 2009; Sul *et al.*, 2010; Sul *et al.*, 2011).

One computation emerging prior to outcomes involves utility. The term utility is used in economics and game theory as a measure of pleasure or satisfaction. Also, in a decision-making problem, a rational subject prefers the option with greater utility value among all provided alternatives (Von Neumann & Morgenstern, 1947; Fishburn, 1970; Tversky & Kahneman, 1981). Utility functions are thus used to model the decision-maker preference in a formal mathematical sense. These functions represent a positive relationship with provided reward and negative relationship with response cost; they thus

assign higher utility values to the options that provide more reward but require less amount of work or risk (Phillips *et al.*, 2007). Utility-related signals have been observed in the ACC of humans (Apps & Ramnani, 2014; Klein-Flugge *et al.*, 2016), non-human primates (Kennerley *et al.*, 2009; Kennerley & Wallis, 2009b), and rodents (Hillman & Bilkey, 2010; 2012). Such a role requires the region to assess several parameters involved in the decision, and combine them into an integrative format (Kennerley *et al.*, 2011). However, some studies failed to find strong evidence of multiplexing signals related to task variables in this structure (Kennerley & Wallis, 2009a; Skvortsova *et al.*, 2014). Furthermore, if the ACC contains utility information, its activity should correlate with benefit (reward), but should be anticorrelated with cost (effort). However, there are some uncertainties about these signals. Some studies have reported positive correlations between ACC activity and effort (Prevost *et al.*, 2010; Kurniawan *et al.*, 2013) or reward (Croxson *et al.*, 2009; Klein-Flugge *et al.*, 2016), while others have observed the inverse - higher activity for smaller amounts of effort (Croxson *et al.*, 2009; Klein-Flugge *et al.*, 2016) or reward (Hayden *et al.*, 2011b; Chong *et al.*, 2017). Still some other studies have even observed a mixture of relationships (Kennerley *et al.*, 2011). These data beg several key questions: Does the ACC encode characteristics of upcoming choices and associated outcomes, and if yes, how? Does the ACC encode task variables separately (i.e., among separate sub-populations of cells) and send them downstream for a comparison, or does it integrate all required information to compute the utility of each option itself? If the ACC is involved in the assessment of the choice utility, does it occur prior to a decision? The answer to these and other similar questions can provide a better understanding of how the ACC contributes to information processing in a larger network of interconnected brain structures. Thus, in this thesis study I aim to discuss them.

The ACC and other nearby regions in the medial PFC have extensive connections with other cortical and subcortical regions. These include primary motor cortex (Strick *et al.*, 1998), amygdala (Van Hoesen *et al.*, 1993), striatum (Haber *et al.*, 2006), and brainstem dopaminergic nuclei (Porrino & Goldman-Rakic, 1982). These structures are well-known for their contributions to motor output, decision-making, motivation, and some other reward-based cognitive behaviours (Salamone, 1994; Floresco & Ghods-Sharifi, 2007; Kim *et al.*, 2009; Sul *et al.*, 2011). Hence, their connections with the ACC are consistent with many of the roles proposed for this region (Paus, 2001; Jones *et al.*, 2005; Rushworth *et al.*, 2011; Kolling *et al.*, 2016). For example, recent rat studies have indicated reciprocal communication between ACC and Ventral Tegmental Areal (VTA) in goal-oriented behaviours (Elston & Bilkey, 2017; Elston *et al.*, 2018; Elston *et al.*, 2019). Specifically, an increase in the ACC signaling to VTA has been observed during low-effort trials in an effort-reward task, which can be interpreted as an elevation in rat motivation (Elston & Bilkey, 2017). In the opposite direction, the VTA-to-ACC signaling observed in a cost-benefit decision-making task has been interpreted to be related to error-feedback and behavioural adaptation (Elston *et al.*, 2018).

Many of the computations ascribed to the above-mentioned structures are thought to involve catecholamines, particularly dopamine (DA). Psychostimulant drugs affecting these neuromodulators may therefore alter functions of these structures. For example, D-amphetamine (AMPH) has been shown to strongly affect the catecholamine system and behaviour (Seiden *et al.*, 1993; Leyton *et al.*, 2002; Ren *et al.*, 2009). Acute AMPH increases the extracellular concentration of DA and other monoamines (Pontieri *et al.*, 1995), and its administration prior to tasks has been shown to change effort-reward based decisions in rats (Floresco *et al.*, 2008a; Floresco *et al.*, 2008b; Bardgett *et al.*, 2009) and

to modulate task-related ACC neural activity (Lapish *et al.*, 2015). However, how this drug affects the neural activity underlying cognitive processes, particularly decision-making, is not well understood.

In this thesis, I aim to investigate rat ACC spiking activity in an effort-reward decision-making task to address two main questions: first, how does the rat ACC encode cost, benefit, and utility of an effort-reward choice? And second, how does systemic AMPH treatment affect the encoding of these parameters in the ACC? For the first question, I expect the rat ACC to dynamically encode the level of effort, reward, and utility of the upcoming choice. As discussed above, there are disparate findings about how ACC activity correlates with effort and reward and if they signal utility. Most of the previous reports are from human neuroimaging studies. Hence, our single cell data may shed light on the human literature. Moreover, as discussed above, there is no consensus on the integration of effort-reward information as a form of utility in the ACC. Thus, I will investigate neural responses not only to single parameters, but also the response to multiple variables in combination.

In addition to expected utility, there are several other parameters which might be encoded in the ACC. One is the path taken by the animal. I hypothesize that ACC neurons will differentiate between left and right paths, regardless of the choice conditions, because of previous data indicating that the ACC encodes location and other task-related states that the animal is currently experiencing (Lindsay *et al.*, 2018; Mashhoori *et al.*, 2018).

Studying the timing of variable encoding can also provide important insights into the function of the ACC region. For example if the ACC is not involved in making a decision, but instead is a critical component in keeping the animal motivated to pursue the goal, I would expect the ACC activity to be greatest after selecting a path and not before.

Another variable which might be encoded is freedom-of-choice. I would also expect no significant difference between ACC activity in the trials where the rat is forced into one of the two options and the ones where he is free to choose. However, observing the converse would support the proposed role of the ACC in volitional behaviours.

For the second question of this study, I expect that AMPH would significantly affect both effort- and reward-related behaviours as well as their corresponding ACC activity. Early studies have shown that AMPH induces motoric hyperactivity as well as stereotyped behaviours, defined as repetitive, purposeless movements (Randrup & Munkvad, 1967; Groves & Rebec, 1976; Sams-Dodd, 1998). Also, AMPH-treated animals have shown increased food seeking (Foltin, 2001), but less interest in consuming food (Foltin, 2001; Cannon *et al.*, 2004; Wellman *et al.*, 2009). Although these observations suggest an alteration of effort and reward information in the brain, how the encoding of these parameters in the ACC is affected by AMPH is not well studied. On the other hand, this stimulant has been shown to induce dose-dependent changes in rats' choice preference (Floresco *et al.*, 2008b). Specifically, low AMPH concentration increases rats' willingness to exert more effort for more reward (Floresco *et al.*, 2008b), while its high-dose decreases such willingness. Amphetamine also influences ACC neural dynamics in decision-making tasks (Lapish *et al.*, 2015). Low-dose AMPH facilitates separation between the neural states of different phases of a task, but the opposite effect is observed under high doses. Therefore, I will investigate the modulation of both behaviour and ACC ensemble activity by different doses of AMPH when rats face decisions about different combinations of effort and reward. Both effort and reward are constituent elements of utility, and thus I expect AMPH to alter the encoding of utility.

To achieve these objectives, I implanted rats' brains with arrays of fine-wire electrodes, called tetrodes, and recorded the spiking activity from the ACC while animals performed a decision-making task. The task is an automated figure-8 maze with two available effort-reward options, each of which requires the animal to exert a specific amount of effort in order to receive a specific amount of reward. In each experimental trial, the rat is either forced to select one option, or is allowed to freely choose among the two options. The rat then has to climb a barrier (the height of which determines the level of effort) to reach some food (the amount of which represent the level of reward). Further details of the maze and experimental design are provided in the method sections of the following chapters.

I manipulated the levels of effort and reward to investigate neural signaling in the rat ACC in response to these variables at each time and location of the task. The analysis and the results of this experiment are provided in the second chapter. Following that, I modified the experimental design on the same figure-8 maze in order to address the second question (the effect of AMPH on encoding in ACC). In this design, the rat performs a certain number of trials before and after injecting AMPH (i.p.) of a particular dose during each session. Then, I compared the animal's pre- and post-injection behaviour and ACC neural activity. The results were also compared with other sessions with different AMPH doses, as well as control sessions in which saline was injected instead of AMPH. The analysis and the results of this experiment are provided in the third chapter. Lastly, I present a general discussion in the fourth chapter.

2 CHAPTER2: Multiplexing of Effort and Reward Coding in Rat Anterior Cingulate

Abstract

Estimating costs and benefits of a given action is critical in coordinating goal-directed behaviours. The Anterior Cingulate Cortex (ACC) appears to be involved in this process; however, the underlying neural mechanisms are not well understood. Here we examined rats' ACC neuronal activity in response to varying amounts of effort and reward in a binary choice task. Single neurons in the ACC conveyed signals for effort, reward, and path, among which path appeared to be the strongest modulator of ACC neurons in most spatial locations, consistent with previous evidence that the ACC encodes the animal's location and features of the task. We also found that the activity of a sizeable sub-population of neurons correlates positively with reward and negatively with effort. In other words, they encoded the net utility of an option. The second most frequently observed signaling type had the opposite relationship with these two parameters, negative correlation with reward and positive correlation with effort, suggesting negative utility. The signaling balance of these two groups of neurons dynamically changed according to the value of the provided effort and reward as well as the phase of the task (e.g. the positive utility signal outweighs the negative utility signal for more valuable options, especially on approach to reward zones). The intermixing of positive and negative utility encoding neurons may hence explain the disparate findings observed in effort-reward studies using functional magnetic resonance imaging in humans. We also observed some prospective activity related to the remotely located reward amount when the animal was receiving reinforcement at the central feeder, long before the rat had made a choice. This finding is consistent with our previous report that the location represented by ACC

ensemble briefly deviates from encoding the present location to encode an alternative reward location during the pause after reward and suggests that, as rats experience one reward, their brains process information about upcoming or previous rewards.

2.1 Introduction

Numerous studies in rats suggest a key role for the anterior cingulate cortex (ACC) in effort-reward decision making. Given a choice between a high-reward attainable by climbing a barrier and a low-reward attainable without effort, rats with ACC lesions are biased towards the easier/lower reward (Walton *et al.*, 2002; Walton *et al.*, 2003). Lesioned rats are still capable of discriminating among reward amounts, suggesting a specific role for ACC in weighing effort against reward. This basic finding has been replicated many times with effort consisting of ramp climbing, attention to brief visual cues, or multiple lever presses (Schweimer & Hauber, 2005; Rudebeck *et al.*, 2006; Walton *et al.*, 2009; Holec *et al.*, 2014; Hosking *et al.*, 2014). Rats with a disconnection of ACC from either basal lateral amygdala or nucleus accumbens (NAc) show similar impairments (Floresco & Ghods-Sharifi, 2007; Hauber & Sommer, 2009). One caveat is that the impairment is apparently task specific. ACC lesions do not affect the breakpoint when lever press requirements are gradually increased (progressive ratio schedule), nor preference for a high-force/high-reward lever (Schweimer & Hauber, 2005; Holec *et al.*, 2014).

Human imaging studies also implicate parts of the medial prefrontal cortex (mPFC) in effort-reward decision making, most notably the mid-cingulate cortex (MCC) and adjacent supplementary and pre-supplementary motor areas (SMA and pre-SMA). During decisions to exert either cognitive or physical effort, activation is seen in MCC which codes for the degree of anticipated effort (Croxson *et al.*, 2009; Prevost *et*

al., 2010; Kurniawan *et al.*, 2013; Apps & Ramnani, 2014; Skvortsova *et al.*, 2014; Scholl *et al.*, 2015; Klein-Flugge *et al.*, 2016; Chong *et al.*, 2017). In cases where reward is varied, the MCC response is usually also sensitive to reward amount (Croxson *et al.*, 2009; Prevost *et al.*, 2010; Kurniawan *et al.*, 2013; Skvortsova *et al.*, 2014; Klein-Flugge *et al.*, 2016; Chong *et al.*, 2017). Most commonly, MCC activity is positively correlated with effort and negatively correlated with reward, suggesting a negative correlation with effort-discounted net reward (Prevost *et al.*, 2010; Skvortsova *et al.*, 2014; Chong *et al.*, 2017); however, the opposite pattern is reported in some studies (Croxson *et al.*, 2009; Apps & Ramnani, 2014; Klein-Flugge *et al.*, 2016). The difference might be attributable to regional variation. For example, one study reported a negative correlation with effort-discounted value in anterior MCC (and SMA) and a positive correlation in posterior MCC (Scholl *et al.*, 2015). Interestingly, the MCC is not typically activated during force exertion, per se. Instead, exertion force appears to activate SMA and sensory-motor cortex (Pessiglione *et al.*, 2007; Schmidt *et al.*, 2012). Activity in ACC related to both effort and reward shows up when subjects are given a choice about effort-reward payoffs as well as during outcome evaluation (Scholl *et al.*, 2015).

Consistent with rodent and human studies, ACC in non-human primates has been shown to be involved in making a decision. Lesion findings in monkeys suggest that ACC is required for sustaining effective choice behaviour (Rushworth *et al.*, 2003; Kennerley *et al.*, 2006; Chudasama *et al.*, 2013). Electrophysiological studies reveal that neurons encoding decision parameters are more prevalent in the ACC (Kennerley *et al.*, 2009; Kennerley & Wallis, 2009b) and signal more robustly than some other fronto-cortical regions in decision-making paradigms (Hosokawa *et al.*, 2013). In particular, this region exhibits discriminative firing patterns for the amount of reward, reward contingency, and

cost of receiving reward in terms of physical effort, aversive stimulus or delay (Seo & Lee, 2007; Hayden *et al.*, 2009; Kennerley & Wallis, 2009b; Hayden *et al.*, 2011a; Amemori & Graybiel, 2012; Hosokawa *et al.*, 2013). Several studies have emphasized that ACC neurons are integrative, encoding the net value or jointly encoding multiple parameters required for decision making (Kennerley *et al.*, 2009; Luk & Wallis, 2009; Hayden *et al.*, 2011b; Kennerley *et al.*, 2011; Hosokawa *et al.*, 2013). However, there is a debate about the stage of decision-making at which the ACC exhibits the greatest discriminative activity. Some studies have shown ACC activation before motor execution (Matsumoto *et al.*, 2003), but after the choice is selected (Cai & Padoa-Schioppa, 2012; Blanchard & Hayden, 2014); thus, suggested a post-decisional regulatory role for ACC. Whereas others have demonstrated the involvement of ACC in the evaluation of choice outcomes based on the animal's choice history (Seo & Lee, 2007). These results have led to the proposal that the role of ACC is to link actions to their outcomes in light of the past history of reward and punishment (Rushworth *et al.*, 2004).

Taken together, the evidence suggests that the ACC is a critical region in the decision-making process (Rushworth *et al.*, 2004; Seamans *et al.*, 2008; Euston *et al.*, 2012). However, the neural activity underlying the encoding of decision parameters is still not clear. There are a limited number of rodent studies investigated ACC ensemble activity relevant to reward and physical effort. Hillman and Bilkey (2010) provided animals with two options varying in ramp height and reward (Hillman & Bilkey, 2010). The majority of neurons showed higher firing for the effort/reward combination preferred by the rat, but only when a ramp was present on at least one of the arms. A follow-up study replaced ramp climbing effort with a competition with another rat for food pellets (Hillman & Bilkey, 2012). The results were largely the same, with firing rates bias

towards rats preferred choices over a range of effort/reward combinations. Again, the vast majority of cells failed to show reward selectivity when no effort was required on either arm. Taken together, these results suggest that ACC firing is biased towards the economically more advantageous arm, but only when the decision involves both effort and reward. Cowen et al (2012) also examined the effects of effort and reward on ACC activity (Cowen *et al.*, 2012). Rats made a choice between two alleys which then resulted in the presence/absence of a barrier followed, at a distance, by a low- or high-reward. They found that ACC neurons anticipated the forthcoming amount of effort and reward, however, this activity occurred largely after action selection, proximal to the barrier and reward location, respectively. Subpopulations of neurons encoded both effort and, to a lesser extent, reward amount, but many of the cells were modulated by the specific path the animal took. There was little interaction between effort and reward coding, with only 3% of cells sensitive to both variables and, contrary to Hillman and Bilkey (Hillman & Bilkey, 2010; 2012), the population did not respond selectively to the preferred choice. While these studies provide important insights into ACC function, they leave several questions unanswered. Hillman and Bilkey (2010) contrasted primarily low-effort/low-reward against high-effort/high-reward conditions, and offer only limited data on the separate effects of effort and reward (Hillman & Bilkey, 2010). Cowen et al. (2012) , on the other hand, provided an in-depth analysis, but the task structure was sufficiently different than the T-maze used in the original Walton studies (Walton *et al.*, 2002; Walton *et al.*, 2003; Cowen *et al.*, 2012), such that inferences about computational processes underlying the lesion effects remain open to question.

To fill this gap, we aimed to investigate the rat ACC spiking activity in response to changes in the physical effort, reward and path using a decision task closely modeled

after the original Walton effort/reward T-maze task (Walton *et al.*, 2002; Walton *et al.*, 2003). Using automated, dynamically variable ramps, we were able to examine the interaction of effort, reward, and path on ACC activity using a crossed design. We also included both forced and free trials, as numerous studies have suggested a role for the ACC in volitional behaviours (Shima *et al.*, 1991; Jueptner *et al.*, 1997; Passingham & Lau, 2006; Mueller *et al.*, 2007; Passingham *et al.*, 2010). Of particular interest was the activity on the central arm leading up to the choice. Do cells primarily encode effort, reward, or an interaction of the two?

2.2 Material and Methods

2.2.1 Animals

Four mature male rats in the age range of 4 to 10 months were used in this study, three of them Fischer Brown Norway (FBN) hybrid and one Long-Evans (LE). LE rats were prone to jumping off the sleeping platform and did not hold still for turning of electrodes, so we switched after the first animal. Since our analysis revealed very similar proportion of responsive neurons to task parameters for both species, we pooled their data together. After the surgery, rats were housed individually in a 12h-12h reverse light cycle. Experiments were performed in dim light during the animal's waking phase. After recovering from surgery, rats were food restricted and were monitored to keep their weight near 85% of their free-feeding weight for the duration of the experiment. All procedures were performed in accordance with the Canadian Council of Animal Care and the Animal Welfare Committee at the University of Lethbridge.

2.2.2 Surgery and electrode placement

A "hyperdrive" was used to record action-potentials from groups of 20-100 neurons simultaneously. The hyperdrive has 12 independently moveable 4-channel

electrodes ("tetrodes"). Each tetrode was constructed of four polyimide-coated, nichrome wires (14 μm diameter; Kanthal Palm Coast, Palm Coast, FL) twisted together (McNaughton *et al.*, 1983; Wilson & McNaughton, 1993). Drive construction has been described in detail by Gothard et al. (1996), except that our drive used sleeves of silica tubing around each tetrode to improve rigidity (Gothard *et al.*, 1996). Tetrodes were inserted into silica tubing (65 μm inner diameter, 125 μm outer diameter; part number TSP 065125; Polymicro Technologies, Phoenix, AZ). Cyanoacrylate glue was used to secure the electrodes in the silica tube. Tetrodes and two reference electrodes were inserted into a bundle of 14 tightly-packed 30 gauge stainless steel cannula (Microgroup, Inc., Medway, MA) which guided their entry into the brain.

Each rat was anesthetized with isoflurane (1-1.5 % by volume in oxygen at a flow rate of 1.5 L/min) and placed in a stereotaxic holder. The skull was cleared of skin and fascia and craniotomies were opened for the hyperdrive. The hyperdrive was centered over the left mPFC at 3.0 mm Antero-Posterior (AP), 1.3 mm Medio-Lateral (ML), and angled at 9.5 degrees toward the midline. It should be noted that the bundle of cannulae holding the electrodes as they entered the brain, and therefore the electrodes themselves, had a diameter of 1.4 mm so that cells were sampled from a relatively wide region of the medial frontal cortex. Rats were returned to ad libitum feeding and allowed to recover for 3-4 days after surgery. Metacam (1 mg/kg s. c. for 3 days) and Baytril (10 mg/kg s. c. for 5 days) were administrated every 24 hours starting pre-surgery. If pain was evidence after surgery, Buprenorphine (0.03 mg/kg s. c.) was administrated.

2.2.3 Electrode Positions

Tetrodes were lowered 950 μm from skull surface right after the surgery and gradually lowered over the next 3-4 weeks to a depth between 1200 μm and 2600 μm,

corresponding to the ACC, area A32D in the Paxinos and Watson MRI atlas (Paxinos *et al.*, 2015). Once in the target region, electrodes were not moved except when cells were lost and it was necessary to acquire new ones. If electrodes were moved, it was done the day before the experiment and electrodes remained untouched overnight. When repositioning tetrodes which had not been moved for an extended period, the electrode was first raised by 40-320 μm before lowering. This procedure, which we arrived at via trial and error, presumably loosens the tetrode if it is stuck in the guide cannula or has adhered to brain tissue.

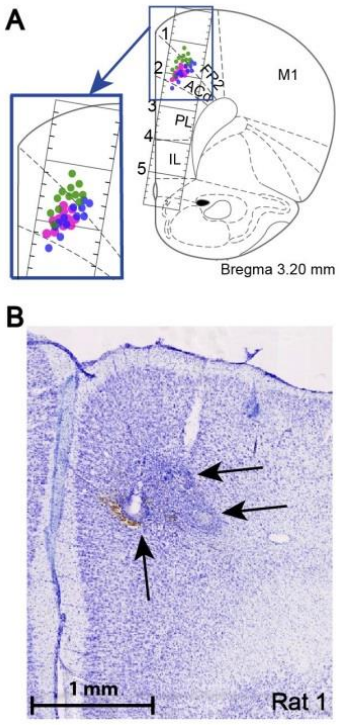

Figure 2.1. Electrode marks and estimation of electrode locations in the brain. A, Recording locations from all electrodes, as inferred based on stereotaxic coordinates. Each colored dot shows the location of one tetrode on one recording session, color coded by rat. As electrodes were moved infrequently, most dots represent 4 recording sessions. A magnified view of the area is provided in the inset. Figure altered from (Paxinos & Watson, 2014) FR2: frontal area 2; ACd: dorsal anterior cingulate; PL: prelimbic cortex. B, Electrode tracks are visible in a coronal brain slice from one of our experimental rats. The slice has been stained with cresyl violet. The arrows mark the locations of the three electrode tips, visible due to the necrosis in their vicinity.

Electrode positions were inferred for all rats using stereotaxic implant coordinates for medial-lateral and anterior-posterior position and the number of turns of the electrode positioning screw for depth (317 μm/full turn). Final electrode position was histologically verified in 2 of our 4 rats and found to lie in the ACC as expected from our electrode positioning logs (see Figure 2.1B). The other two brains were accidentally destroyed during histology; the estimated locations of the electrodes were determined by the daily logs of electrode depth, and are indicated in the Figure 2.1A. For histological verification, the location of electrodes was marked by passing 5 μA direct current for 10 sec through each electrode in each tetrode, positive to the electrode. After a week the rats were transcardially perfused with Phosphate-Buffered Saline (PBS) followed by 4% Paraformaldehyde (PFA) while they were deeply anesthetized with isoflurane. Coronal sections of paraformaldehyde-fixed brains were sliced at 40 microns and Nissl stained using cresyl violet and scanned with a whole-slide microscope (NanoZoomer, Hamamatsu Corporation, Bridgewater NJ).

2.2.4 Experimental Design

2.2.4.1 Apparatus

The effort-reward maze is the same as has been previously used in our lab (Holec *et al.*, 2014). As show in Figure 2.2A, the maze is similar to the T-maze used in previous effort-reward studies in rats (Salamone *et al.*, 1994a; Walton *et al.*, 2002) but adds a return path to each arm, forming a "figure-8". Rats receive food reward from a "central feeder" on the T-stem and then, at the T intersection, turn either left or right to reach a feeder at the end of each arm. These two "side feeders" were located on moveable platforms (15 by 25 cm) which could be elevated up to 51 cm above the maze surface, providing a vertical wire mesh climbing surface ("barrier") on the approach side and a

gradual down “ramp” on the return side. The wire-mesh was made of 1.6 mm thick galvanized steel wire with a 1.25 cm square spacing. Two drawbridge-style “gates” were located on both the approach and departure points from the central arm and were used to enforce forced choices on some trials as well as to ensure that rats did not run backwards. The maze itself was 102 cm long by 114 cm wide and 60 cm height from the floor. The animal could freely run on a 15 cm-wide wooden surface painted in light gray with 36 cm-high black plastic walls on both sides. To make the maze totally automated, each reward platform, fashioned from aluminum, was attached to a carriage which slid on a guiderail, allowing height adjustment using a rack and pinion gear system driven by a

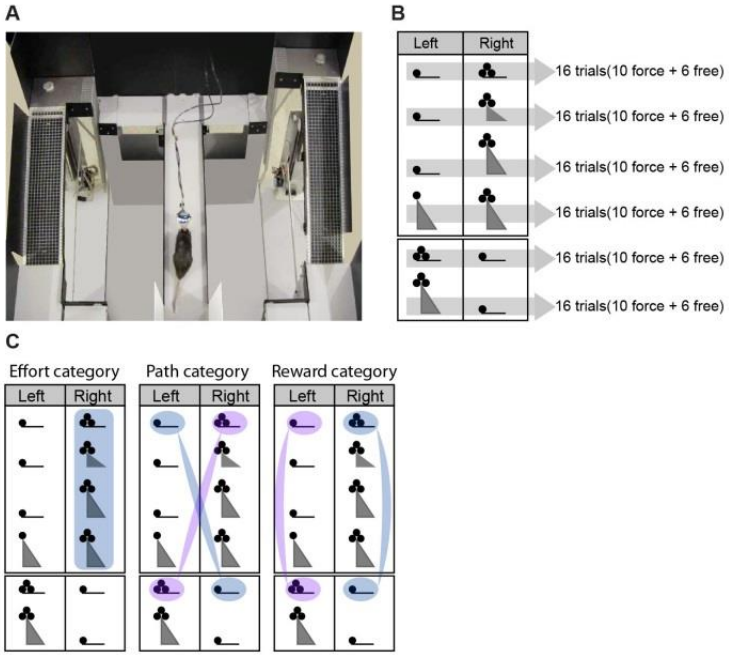

Figure 2.2. Design of the experiment and analysis. A, top view of the figure-8 maze. B, Task schematic showing task parameters for each set of trials, in the order they were presented. Height of triangles shows ramp height while number of black circles shows relative reward amount. C, The subset of trials used for effort, path, and reward encoding tests. The highlighted groups of trials were chosen so that only the parameter of interest changes and the other two remain constant. For example, only the barrier height changes in the effort group, while the amount of reward is the same and all of those trials occur on the same side. For reward and path, two groups of trials were used (blue and pink), such that each of the two groups was analyzed separately and then their results were combined.

stepper motor (Model 23Y9, Anaheim Automation, Anaheim, CA) and powered by a stepper motor controller (Model G251X, Gecko drive, Tustin, CA). The stepper motors, gates, and feeders were all controlled via a digital interface board (described below).

For reward, a high-calorie, chocolate-flavored liquid (Ensure®, Abbott laboratories, Brockville, ON) was delivered through a silicone tube (inner diameter:1.98mm; outer diameter: 3.18mm, wall thickness: 0.61mm, VWR international, Mississauga, ON), one end of which was attached to a 50 ml syringe hung at the feeder location but behind the maze wall and 51 cm above the maze surface. The other end was attached to a port at the bottom of a conical well through a hole on the track. The conical well, which played the role of a feeder, was 24 mm in diameter and was glued to a washer which was magnetically attached to the track floorboard. The silicone tube was not accessible by the rat, preventing the rat from biting it. Reward flow was controlled via a solenoid pinch valve (Asco Scientific, model SCH284B004, Florham Park, NJ). The flow rate when the valve was open was approximately 0.0126 ml of liquid per 100 ms.

2.2.4.2 Behavioural Procedures

Each trial was initiated by the rat reaching the feeder on the stem of the T, referred to as the central feeder, at which point reward was delivered and two automated gates, located at the entry of T stem, were raised to prevent the rat from going in a backwards direction. The central feeder delivered 0.03 ml (200 ms of valve open time) of Ensure on every visit. At the T intersection, rats turned left or right, choosing between arms that varied in the barrier height (our proxy for effort) and reward amount. Low-reward was 0.03 ml of Ensure (200 ms of valve open time) and high-reward was 0.12 ml (600ms of valve open time). Gates located just after the T intersection ("choice gates") were used for forced choices and were also raised if the rat tried to return to the central feeder by

backtracking once they had obtained a reward at one of the side feeders. Once the rat had obtained reward, he returned by descending a ramp and returning to the feeder on the T-stem segment of the maze. The gates at the T-stem entrance were raised so as to guide the rat into the T-stem and prevent him from going backwards into the other half of the maze.

All animals were both trained and tested while they were implanted and food-restricted. Testing occurred daily for up to 60 minutes per session over the course of 8 to 12 weeks. We started with a barrier height of zero and both choice gates down, allowing free choice on every trial. Rats were manually guided, as needed, during the first 5 to 10 training sessions until they could successfully complete 192 trials without external aid. Next, rats were trained using forced alternation and a 7.6 cm barrier at the high-reward and no barrier at the low-reward feeder location. Once the rat was able to complete 192 trails per session they were next tested on the full experimental protocol in which both barrier height and reward amount were varied dynamically according to a partial factorial design with three levels of effort and two levels of reward on both left and right arms of the maze as shown in Figure 2.2B. Each block of 96 trials consisted of six different effort-reward conditions. Within a condition, the first 10 trials were forced (choice gate blocked one arm alternatively) and the next 6 were free. Once the block was completed, the block of trials was repeated in the same order until the total elapsed time of 60 minutes had been reached or 256 trials. The medium barrier height was initially set to 7.6 cm with high-effort set to 15 cm until rats could successfully complete 192 trials per session (2 full blocks of trials). Barrier heights were gradually raised to 10 cm medium/20 cm high, 12.7 cm medium/25.4 cm high, and ultimately, 15 cm medium/30 cm high, using successful completion of 192 trials/session as a success criterion. In two of the sessions reported here for one rat, barrier heights were 23 cm medium and 46 cm for high.

The partial factorial design meant that one side had more high-reward trials than the other on any given session. To reduce the likelihood that rats would develop a side bias, the trial order was reversed every other session (i.e., the right arm started as high-reward/high-effort one day, then left arm started as high-reward/high-effort the second and so on).

2.2.4.3 Data Acquisition and Experiment Control

The experiment was performed in a dimly lit room and was recorded by a standard video camera mounted to the ceiling directly above the maze. Neural spiking activity was recorded by a 48-channel unity-gain headstage, HS-54 (Neuralynx, Bozeman, MT), and transferred to the Cheetah neurophysiology data acquisition system via two multi-wire tether cables (TETH-HS-27-3M, Neuralynx, Bozeman, MT) and a 128 channel commutator (PSR-36-4, Neuralynx, Bozeman, MT). All signals were amplified, digitized and bandpass filtered between 600 and 6000 Hz via a Digital Lynx system (Neuralynx, Bozeman, MT). A counter-weight was used to reduce the weight of the headstage and tethers, allowing the rats to climb. The counterweight was attached to the tether just above the headstage using fishing line running through a plastic pulley attached to the bottom of the commutator. Red and green LEDs installed on the headstage marked the front and back of the rat's head, respectively, and were used by the Cheetah software to track the animal's position on-line. The tracker used circular zones centered on each of the feeders as well as one on the far corner of the return path to coordinate maze events with the rats position. For example, entering the tracking zone around the central feeder triggered food delivery and the raising of the gates at the entrance to the T stem which

prevented the rat from going backwards. If the trial was a forced trial, the appropriate gate on the exit of the T stem was also raised to direct the rat to the opposite arm.

Two standard computers running Microsoft Windows 7 were used: one to control the maze, using a National Instruments PCIe-7841R programmable digital input/output board and custom software written in Microsoft Visual Basic and Labview (National Instruments, Toronto, ON), and the second one to run the Cheetah (Neuralynx, Bozeman, MT) data acquisition software. Synchronizing the two system's clocks was done by connecting Cheetah acquisition clock output to a digital input on the PCIe-7841R board. Experimental events such as the times of gate openings/closings, times of feeder onsets/offsets, and trial numbers were automatically logged on the experiment control computer, along with the Cheetah clock input. Spiking activity, local field potentials signals (not reported here), and video files were collected via the data acquisition system. The PCIe-7841R board was also programmed to send a digital pulse to the Digital Lynx system every time a feeder opened. These feeder events were then used to synchronize the timestamps of neural and behavioural events.

2.2.4.4 Spike Sorting

Recorded spike events were sorted off-line automatically (KlustaKwik, author: K. D. Harris, Rutgers-Newark) and then manually refined using a second specialized software package MClust (MClust; David Redish, University of Minnesota, Minneapolis) with the purpose of isolating individual cells based on the clustering of similar spike waveforms. In both steps, the properties of action potential waveforms such as peak amplitude, waveform energy and match to canonical waveforms on different channels of each tetrode were utilized for clustering (McNaughton et al., 1983; Wilson and

McNaughton, 1994). The output of this procedure is a train of spike times for each putative neuron.

Neurons were also manually classified into pyramidal and interneurons using half-amplitude duration and trough-to-peak duration, following the method of Bartho et al. (Bartho *et al.*, 2004), as shown in Figure 2.3. Each of the four average waveforms from a tetrode was classified separately and the cell classified as pyramidal or interneuron based on majority rule with ties being classified as pyramidal. Given the low number of interneurons within each session (typically less than 8 units), the analysis was restricted only to putative pyramidal neurons.

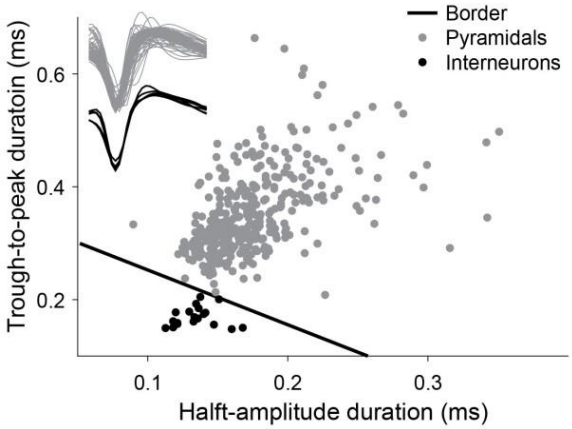

Figure 2.3. Separation of pyramidal cells from interneurons. Shown are data from one sample recording session. Each dot represents the waveform from one cell (after spike sorting) on one channel. Because cells were recorded using tetrodes, each cell is therefore represented by 4 dots. Two clusters are apparent and the line used to separate them is shown. Cells were classified by a weighted sum of the category membership of each of its individual waveforms, with those directly in the middle assumed as pyramidal neurons. Sample waveforms from each of the clusters is shown in inset at upper left.

2.2.4.5 Video Tracking

Each recorded video frame was also preprocessed to extract the rat's position on the maze. A threshold was used to extract the position of all pixels corresponding to the headstage tracking LEDs. We then took the center of mass of these, corresponding roughly to the center of the rats head, as the rat's position.

2.2.5 Analysis

2.2.5.1 Behavioural Analysis

Four different sessions from each of four animals were selected for analysis. These sessions were selected to be at least two weeks into training and include both left- and right-biased configurations (i.e., more high-reward on left or right trials) for each animal. All trials in which the rat paused for more than 1 minute were manually removed from consideration, which amounted to no more than 15 trials in any analyzed session. For behavioural analysis, repeated measures ANOVAs and pair-wise t-tests were used to compare the percentage of high-reward choice (Fig. 2.4) in eight pairs of effort conditions: (LL,LM), (LL,HL), (LL,LH), (LL,HH), (LL,LL(R)), (LM,LH), (LH,HH), and (LH,HL), where L, M, and H indicate low-, medium-, and high-effort, respectively; (R) following similar conditions indicates the reversed high-reward feeder location. Data were considered significant for p value less than Bonferroni corrected alpha value (0.05/8 = 0.00625).

2.2.5.2 Mapping 2D position on the maze to linear axis

We divided the maze into 36 spatial bins starting at the central feeder and looping around one half of the maze (i.e., the bin containing the central feeder was bin 1 and the bin just before arrival at the central feeder was 36). The stem of the T and the parallel returning paths were binned linearly, but to capture figure-8 curve on choice and

returning arms, we binned the turns using radial spokes. The bin numbers in the right loop of the maze were the mirror image of the bin numbers in the left loop. In this way, continuous trajectories on the maze were divided into 36 segments, which were then analyzed independently.

2.2.5.3 Calculating event-related neural activity

To explore how neural firing rates related to task events, we looked at spike-rasters and peri-event time histograms with a bin size of 100 msec. Spikes from each trial were time-warped so that the following events aligned: arrival at the central feeder, departure from the central feeder, onset of climbing, and arrival at the side feeder. As shown in the inset of Figure 2.5, the behavioural epochs defined by these events were color coded: 1) consuming reward at base feeder (blue), 2) leaving from base feeder to climbing (pink), 3) climbing zone (green), and 4) receiving side feeder (peach). The time of arrival and departure from the central feeder was determined by velocity. The rat's velocity was calculated by measuring the distance the animal had covered from one video frame to the next (i.e., every 33.6 msec). On each trial, reward arrival time was determined to occur when the rat's velocity dropped below 25 pixels/second (~ 4 cm/sec). Similarly, reward departure was the time at which velocity exceeded 25 pixels/second. The barrier was located in the center of the 12^{th} spatial bin (Fig. 2.6A-inset), so entry into this bin was used as the start of the climb epoch. The end of the climb epoch and start of the side feeder reward consumption was judged to occur when the rat's velocity dropped below 50 pixels/second (~8 cm/sec) while on the 14^{th} or 15^{th} bin. This side feeder reward consumption ended when the rat's velocity again exceeded the threshold or left the 15^{th} bin. For time warping, we used averaged duration of the four epochs (central reward, 1.86 sec; climb approach, 2.51 sec; climb, 0.84 sec; and side reward, 5.07 sec) as a reference

and linearly expanded or contracted the spikes so that each epoch lasted the same time as the corresponding reference epoch. Although warping time obscures some trial-by-trial variability, it provides a clearer picture of how an individual cell's activity relates to each behavioural event.

2.2.5.4 Categorizing task selective neurons in each spatial bin

Note that the time the rat stays in each bin is different. For instance, rats spent much longer in the bins around the feeder because it took them time to eat. This creates a confound when comparing discriminability of task parameters such as effort and reward across bins, because bins with higher occupancy have a higher signal-to-noise ratio. To ensure that spike rates in each bin were calculated over an equivalent time period, we limited our estimate of spike rate in each bin to a 0.3 second window centered on the time at which the rat was in the center of each bin (determined by averaging the entry and exit times). The 0.3 second window was used because it is close to the median time spent in each bin. Note that, using this method, the spike rates in bins where the rat was moving rapidly may actually include spikes from the adjacent bins.

The major objective of this study was to show how ACC neurons integrate the information about costs and benefits needed to make a decision. Hence, the first step was to investigate if the ACC neurons were selectively responding to task features such as effort, reward, and path (i.e., the trajectory, left or right, that the animal took). To address this question, neural spike rates were evaluated on a subgroup of trials on which only the feature under investigation was changing and the other features were constant (Fig. 2.1C). For instance, to study effort, we compared trials (highlighted in blue in the left table of Fig. 2.1C) in which the path and reward were constant and only the barrier height changed. An ANOVA was then used to assess whether spike rate varied significantly with

the barrier height. As the criteria for statistical significance, $p < 0.05$ is used but only those significant results with a large effect size (> 0.1379) were considered. The effect size was computed by the partial eta-square formula: $\eta_p^2 = \frac{SS(effect)}{SS(total)}$ in which $SS(effect)$ is the ANOVA sum of the squares of the effect and $SS(total)$ corresponds to the total ANOVA sum of the squares (Cohen, 1988). For path and reward, we had two pairings of trials, as shown in Figure 2.1C, each of which was tested separately. To assess path-related activity, for example, we first limited our comparison to trials in which both barriers were equal and low. We then used an ANOVA to compare the right/low-reward trials with the left/low-reward trials. A separate ANOVA was used to compare the right/high-reward to left/high-reward trials. A cell was judged to be path-modulated if it had a low p value ($p<.05$) and large effect size ($\eta_p^2 > 0.1379$) in at least one of these two comparisons. For reward, we used a similar method, comparing high-reward and low-reward on each side independently, as indicated in Figure 2.1C. Assuming a rat did 192 trials (two blocks) per session and his choice in free trials was always high-reward, 88 trials would be available to evaluate effort selectivity, 20 and 44 trails to evaluate path selectivity, and 32 and 32 trails to evaluate reward selectivity. The rational for considering effect size is that p-value depends on sample size, whereas effect size is much less sensitive to sample size. Since sample size is not equal for each of our parameter comparisons, p-values between groups are not comparable. Effect size, on the other hand, quantifies how much a given parameter affects the cell's firing relative to noise levels and is directly comparable across conditions.

We limited our analysis to those cells that were "task-related", defined as those cells whose firing rate was modulated by one of our three primary parameters (effort,

reward, and path) in at least two bins, using the criteria of low p and high eta (as above). Within this subset, the percentage of feature selective cells were obtained for each session and averaged across all sessions from all rats at the end.

To find levels of false alarms expected by chance, data of each session was shuffled for 100 random permutations and was analyzed each time as above. Then the percentage of feature selective cells were obtained for each session and averaged on all sessions at the end. Each shuffled data was obtained by uniformly randomized spiking time points of each cell keeping its first and last spiking time point fixed.

2.2.5.5 Decoding task features by ensemble firing rates

For each session, we created a data set of neural firing rate vectors, consisting of the firing rate of each neuron when the rat was in each bin during each trial. This process resulted in $B \times N \times T$ activity vectors, where B is the number of bins, N is the number of cells, and T is the number of trials.

Depending on the properties of the trial corresponding to each activity vector, three target values were assigned to that vector: the direction that the rat ran in that trial (right or left), the amount of reward that the rat received (high or low), and the level of effort in that trial (high or low). Note that medium effort trials were ignored for this analysis.

For each task feature (path, reward, and effort), we created a balanced set by randomly selecting trials from the trial category that represented that feature (keeping the other two features constant). By balanced, we mean that the set contained an equal number of instances from each class of the task feature (e.g. equal number of high-reward and low-reward instances). Note that in the selected trials for each task feature, the other two features are the same for all the instances.

In the next step, we performed a 10-fold cross validation on the instances of each set. In this approach, the dataset is divided into 10 subsets with equal size. In each fold, we use the data from 9 subsets to train a classifier and test the model on the other remaining subset. Each subset is used exactly once as the test set, and the final accuracy is the average of the 10 accuracies that we obtain on different test sets. As for the classifier, we used a support vector machine (SVM) with a linear kernel.

2.2.5.6 Stability of neural encoding during preparatory and approach phases

To address the question of whether ACC neuron tuning changes during a single trial, we compared the encoding of effort and reward in two phases of the trajectory from the central to side feeder. The "preparatory phase" was defined as the period after leaving the central feeder up to the T-intersection, spatial bins # 4-8, and the "approach phase" was the period from the T-intersection to the barrier, spatial bins # 9-13. The reason we started from fourth spatial bin for the preparatory phase is that we wanted to have the same number of spatial bins for the two phases; also, the animal may still be eating in the first two spatial bins. To compare effort encoding, we found the bin in both phases where a cell's average firing rate was the most discriminatory of the three different levels of effort (i.e., its maximum value situated at the greatest distance from its minimum value). Then the firing rate on the low-, medium-, high-effort trials were arranged into two vectors, one for the maximum discriminable bin on the preparatory phase and another for the maximum discriminable bin for the approach phase. These two bins were then correlated, yielding a measure of the degree to which the cell's effort coding was stable across the preparatory and approach phases. A positive correlation would mean the cell's selectivity for effort (e.g., high firing rate for high-effort) was the same in both phases. A negative correlation would mean that the cell switched its coding (e.g., high firing rate for

high-effort during the preparatory phase but high firing rate for low-effort during the approach phase). A histogram of "encoding stability" correlations was created for each session and these values averaged to present population data across all animals. The stability of reward coding was computed similarly except that there were two subsets of trials used to compute rewards selectivity, one for the right and the other for the left path (see Fig. 2.1C). Correlations for each cell were computed separately for the two paths and treated as two cells in the final histogram of correlation values. The histogram counts for reward encoding stability were then divided by two to reflect the true number of cells.

To determine the range of encoding stability correlation scores expected by chance, we have repeated the analysis for 100 times per session on the shuffled data with random spiking time series, keeping the first, last, and average spiking activity fixed for each neuron. Then the histogram was found for each random permutation and then averaged over all at the end.

2.2.5.7 Effort-reward joint encoding

In order to judge the relative strength of effort and reward coding for each cell, we computed the Pearson's linear correlation coefficient for each cell's firing rate as it relates to either effort or reward levels. For each cell, these correlations were computed for each of the spatial bins from the central feeder to the side feeder. Reward trials are chosen to be on the same side as effort trials to eliminate the effect of path in effort-reward joint coding. When the correlations of every cell with effort and reward are plotted as a scatterplot with reward on the horizontal axis, they fall into four quadrants. We numbered these quadrants similar to quadrants of a Cartesian coordinate system in plane geometry, from the upper right as I through IV in a counter clockwise fashion (Fig. 2.10A inset). Cells in the fourth quadrant have firing rates that are positively correlated with reward

and negatively correlated with effort, essentially coding net utility (E-/R+). Conversely, the cells in the second quadrant seem to encode negative net utility (E+/R-). Although this analysis of the effort-reward joint coding provides information about the number of cells encoding positive and negative utility, it doesn't provide any information about the strength of their signals. Hence, we sought to examine how a pre-selected group of positive and negative utility cells change their firing rate, as a population, throughout the task. We first chose a spatial bin between central and side feeder (bin # 6) that showed clear evidence of positive and negative net utility selectivity (see Fig. 2.10B) yet was remote from the effort and reward events (i.e., climb and food consumption) which strongly drive cellular activity. We then plotted the average firing rate of all cells coding positive net utility (cells in Q_{IV}) for every position on the maze, separated by effort and reward levels (see Fig. 2.12A). We did the same for the average firing rate of all cells encoding negative net utility (cells in Q_{II}). Finally, we examined the net utility signal available from these cells by taking the sum of the positive net utility firing rates minus the negative ones (see Fig. 2.12B). The sum was used, instead of making an average, as the number of positive utility cells outnumbered the negative utility cells and so would be expected to have a stronger influence on downstream structures.

2.3 Results

Four rats performed the effort-reward decision making task while single-neuron firing rates were recorded using a tetrode array located in the dorsal part of the medial prefrontal cortex. Four sessions per animal were selected for the analysis. Because of the low cell counts in one animal, neural analysis are based on three animals while behavioural results are based on all four.

2.3.1 Climbing effort influences decision making

A total of 113 to 256 trials per session (mean of 221trials/session) were performed by the animals in the selected sessions for the analysis. Each of the four rats performed 10 alternating forced-choice trials for every effort-reward combination tested followed by 6 free-choice trials. Figure 2.4 shows the mean percentage of high-reward choices for all free-choice trials, averaged across sessions and animals but separated by effort-reward condition. As indicated in the figure legend, the amount of reward available for the two choice arms was always unequal but the effort varied. The last two conditions in the figure correspond to the second phase of testing where the high-reward feeder location was swapped.

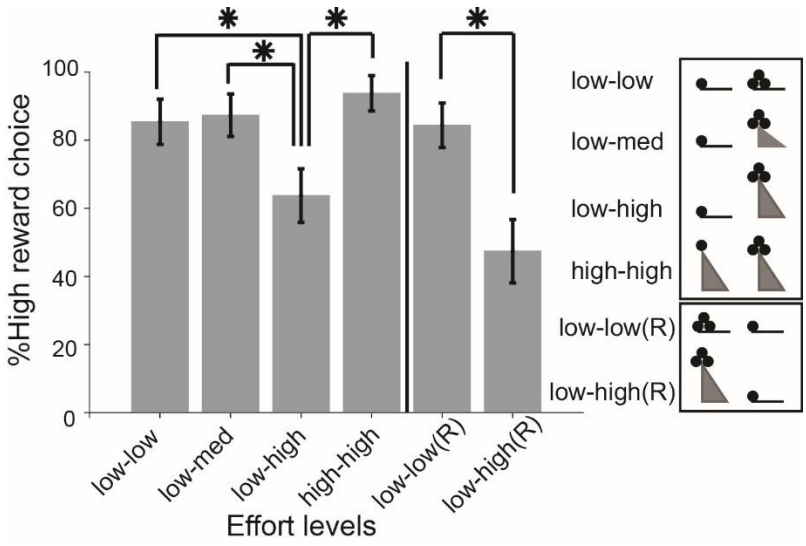

Figure 2.4. Mean choice of the high-reward arm on free trials. Results are averaged across 16 sessions. The x-axis shows the effort level for the low- and high-reward arms. Note that the position of the high-reward switches for the two conditions on the right. The legend at the right provides a graphic summary of the different choices in each condition with reward amounts indicated by small black circles and ramp height depicted by the height of the triangle. Error bars show the standard error of the mean (*: Bonferroni corrected significant at $p<.00625$, (R) indicates reversed reward conditions at the side feeders).

Rats chose the high-reward feeder location roughly 85-90 percent of the time, but this choice was influenced by effort (Fig. 2.4). Planned pair-wise comparisons between a group of trial types with high-reward on the same feeder location showed that rats chose the high-reward arm significantly less when it had a high-barrier compared to low or medium (see Fig. 2.4; low-low vs low-high: $t(15) = 3.568$, $p = 0.003$; low-medium vs low-high: $t(15) = 3.890$, $p = 0.001$). Similar results were found when high-reward was on the opposite location (low-low vs high-low: $t(15) = 4.366$, $p = 0.001$). However, when the choice was between a high-barrier on both arms, rats again preferred the high-reward significantly more than when the low-reward option did not involve a climb (low-high vs high-high: $t(15) = -4.495$, $p<0.001$). In sum, the results show that choices were definitely influenced by both reward and effort.

2.3.2 ACC neurons encode the path, effort, and reward

A total of 760 putative single-unit were recorded and, after screening out putative interneurons and those which showed no modulation by any task variable, 647 of them were chosen for the further analysis. The total number of neurons analyzed per session was: 82, 82, 94, and 99 (for rat 1); 41, 42, 46, and 47 (for rat 2); 23, 24, 30, and 37 (for rat 3).

Electrode positions were inferred based on the stereotaxic coordinates of the implanted microdrive. As shown in Figure 2.1A, they were largely localized to the ACC with a few electrodes in FR2. Histology was performed on two rats and confirmed our stereotaxic-based estimates (Fig. 2.1B). Histology of the one rat that had low-quality units showed that the electrodes were within the target region, and helped to confirm that our stereological coordinates are correct for the other two rats.

Some typical examples of single-unit responses to effort, reward and path are shown in Figure 2.5. As illustrated, single ACC cells showed stronger discriminant activity for effort and path than for reward. Discriminant activity was observed in all phases of the task, including the period while the rat was at the central feeder, on the stem of the T, climbing, and consuming reward at the side locations. However, most cells showed discriminant activity in only a limited region.

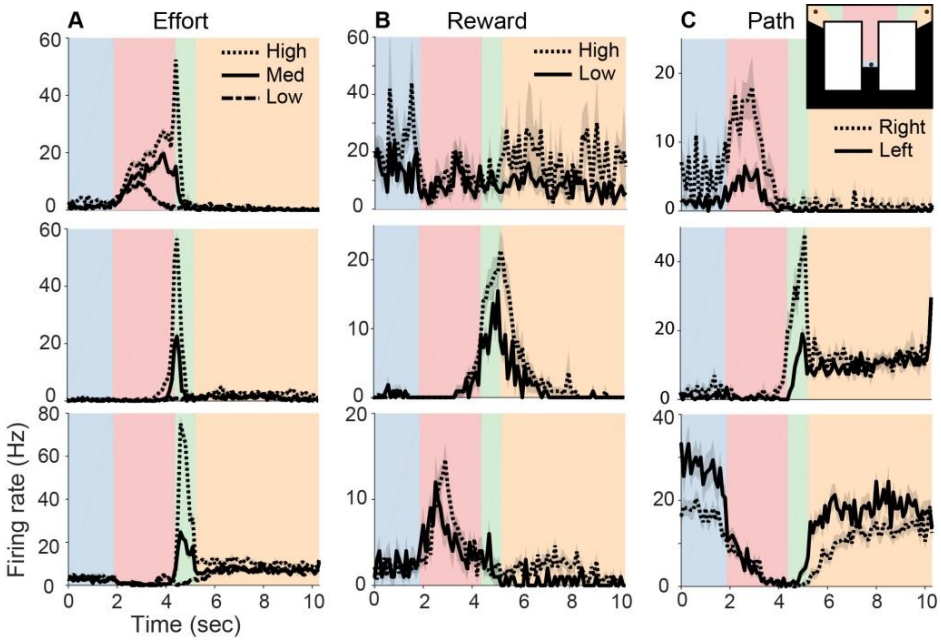

Figure 2.5. Responses to task features by nine example cells. Colors correspond to the maze regions shown in the inset, upper right. Light gray/blue corresponds to the period spent at the central feeder, pink to the approach to the ramp, green to the actual climb and peach to the time spent at the side feeder. A, Effort: Shown are three examples of neurons which had discriminant activity for the three ramp heights while on approach to the ramp (top), right before climbing (middle), and while climbing (bottom). B, Reward: Three neurons showing differential activity to high- and low-reward on the base and side feeders (top), during the climb and approach to side reward (middle), and on the stem of the T (bottom). C, Path: Three examples of cells which show differential activity on the left and right arms of the maze. Differential activity is seen just after leaving the central reward (top), while climbing (middle), and during consumption of food at both the central feeder and side feeders. Note that the rat's position was the same while at the central feeder and beginning of the T stem. These peri-event histograms have been time-warped on a trial-by-trial basis before averaging so that the entry and exit from each analysis zone were matched. Standard error of the mean is shown in light gray behind each track.

Unsurprisingly, some cells showed differential encoding of barrier height during the climb (Fig. 2.5A bottom) while other cells encoded effort just as the climb was commencing (middle). Still other cells show anticipatory encoding of effort well before rats reached the barrier (top). Effort encoding typically dropped sharply at the side feeder. Just as effort-related activity was strongest during the climb, activity discriminant of reward magnitude was common at the side feeders, where the rat received either the high- or low-reward. However, reward discriminant activity was occasionally observed on central feeder (Fig. 2.5B, top) as well as on approach to the side feeders (Fig. 2.5B, bottom). Path discriminant cells were common on all parts of the maze. Surprisingly, this included the central feeder (Fig. 2.5C, bottom) and stem of the T (top), where the rats' running trajectories were nearly identical across trials. However, rats' movement trajectories for left and right turns start to diverge not later than the second spatial bin; hence, they significantly differentiate between these two paths even on this segment of the maze (ANOVA, $P<0.001$) (e.g. Fig. 2.6A). This activity might be a reflection of the rat's tendency to lean left or right before the paths diverge. However, another possibility is that this activity corresponds either to the upcoming turn or previous trajectory. Unfortunately, we did not have enough trials with different combinations of previous/upcoming choice to discriminate between retrospective and prospective coding.

As described below, population data from all cells in all sessions confirmed the results illustrated in our single-cell examples. To perform these analyses, data were analyzed in 36 bins covering the entire extent of the rat's path and discriminant activity was judged based on the results of a one-way ANOVA, comparing firing rates at each level of a given parameter (e.g., effort, reward, path). A cell was judged to encode a

parameter in a particular bin if it had both a large effect size ($\eta_p^2 > 0.1379$) and low p-value ($p<.05$). Chance levels were computed by randomly shuffling spiking activity of each cell for 100 times per session. Under random shuffling, the average number of cells classified as tuned to Path was less than 5.2%, Reward, less than 4.4%, Effort, less than 0.6 %, Free-Forced Choice, less than 0.0025%.

Because previous studies have implicated the ACC in volitional behaviour (Jueptner *et al.*, 1997; Passingham *et al.*, 2010) we expected to find differential activity between free- and forced-choice trials. During forced trials, one gate at the T intersection was raised, forcing the rat to choose the opposite arm, whereas in free-choice trials, both arms were open. The raised gate thus provides a salient visual cue that differed between forced and free trials. Given the perceptual and decision-related differences, we expected at least some cells to differentiate free and forced trials, but this was not the case. As shown in Figure 2.6B, the average number of cells showing differential activity related to free/forced choice never exceeded 1.1% and was close to zero in most bins. Because of the lack of effect, subsequent analyses combined data from both free- and forced-choice trials.

Path was by far the most robustly encoded variable, with roughly a quarter of cells discriminating path in the majority of spatial bins (Fig. 2.6B). As expected, the percentage of path-encoding cells was reduced along the T-stem which is a common segment of the track between left and right loops (Fig. 2.6B); however, even in places where trajectories for left and right turns were completely overlapping, the number of cells encoding path was still above chance.

Effort was the next most influential task parameter. The proportion of cells encoding the barrier height peaked at 23% at the time of the climb (bin 12) with the number of cells increasing monotonically during the journey from the central feeder to the ramp. The number of effort-encoding cells dipped while rats were at the side feeder and then reached a low plateau around 8% from bins 16 to 23, corresponding to the downward ramp (see Fig. 2.6B, 3-level effort comparison). This result indicates that ACC cells were also sensitive to the slope of the down ramp which, of course, was tied to

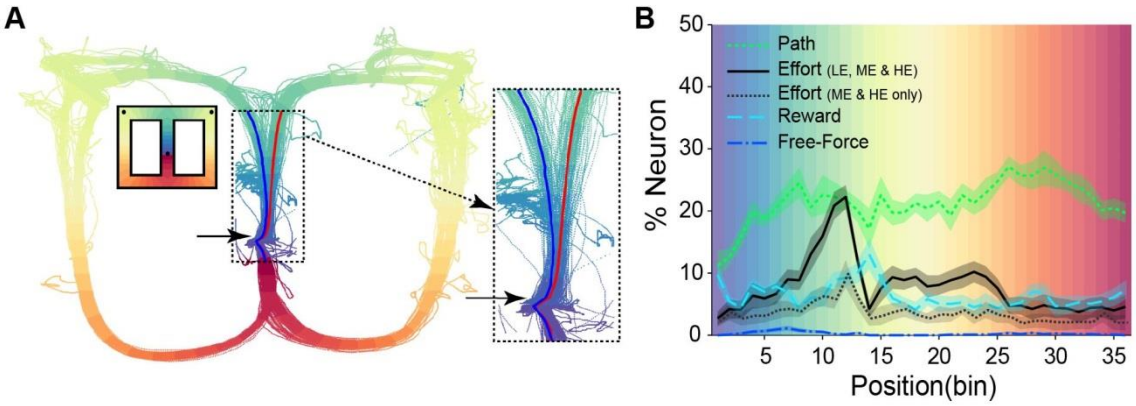

Figure 2.6. An example of rat's movement trajectories and population selectivity for task parameters. A, The route trajectory that one of the animals followed on the maze in a session is shown. The magnified section illustrates the point at which the animal's path diverged. Colors correspond to each of the 36 spatial bins of the maze as shown in the inset. Blue and red lines indicate the average of left and right trajectories, respectively. The arrow shows where these trajectories start to significantly differ (ANOVA, $p < 0.001$). B, Averaged percentage of path, effort, reward and free-force selective neurons in each spatial bin. Shaded area around each plot indicates the standard error of the mean, computed based on 12 session averages. Chance level, computed by randomly shuffling data, is about 5% for path and reward and close to zero for Effort and Free-force parameters. Background colors correspond to bin positions illustrated in the inset of panel A. (ANOVA, $p < 0.05$ and effect size>0.1379)

barrier height. Rats often slowed on the down ramp, consistent with the idea that this portion of the ramp was also effortful. Effort encoding was lowest during the return to the central feeder. While these results indicate a strong effort effect, the large number of cells encoding effort is partly due to the inclusion of three levels of effort, one of which is zero.

Many of the effort-sensitive cells responded strongly to the contrast between some ramp and zero ramp (Figure supplement 2.2). When only the medium- and high-barrier conditions were used to determine the proportion of effort-encoding cells, the proportions were much lower (see Fig. 2.6B, black dashed line). Thus, the ACC appears to encode the barrier presence or absence, rather than producing a graded response of effort amount (see Figure supplement 2.2 and 2.3 for the effort encoding and decoding results using pairwise comparisons between the barrier heights).

Similar to effort, the proportion of reward-related cells varies based on task events. The number of cells discriminant of reward amount peaked at 15% while the rat was at the side feeder. This is, of course, where the rat received either a large or small reward and so this finding is not unexpected. There was a small ramp up in the percentage of reward encoding cells during the approach to the feeder location, but this started later and was less robust than that seen for effort (Fig. 2.6B). The percentage of reward encoding cells was in the vicinity of 5% for the rest of the trajectory, with the exception of the region around the central feeder. At this location, the proportion of reward-encoding cells was near 10%. This activity is remarkable because the reward magnitude at the central feeder does not vary, and the rat's behaviour (i.e., position) at this point is in no way related to the upcoming reward choice. Being at the central feeder hence evokes activity in ACC which reflects upcoming (or perhaps preceding) reward amount. Examples of two cells showing reward discriminant activity at the central feeder are shown in Figure 2.7. This activity was observed both on arrival at the central feeder, while the rat is decelerating (Fig. 2.7 left), and while the rat is at the central feeder and, presumably, eating (Fig. 2.7 right).

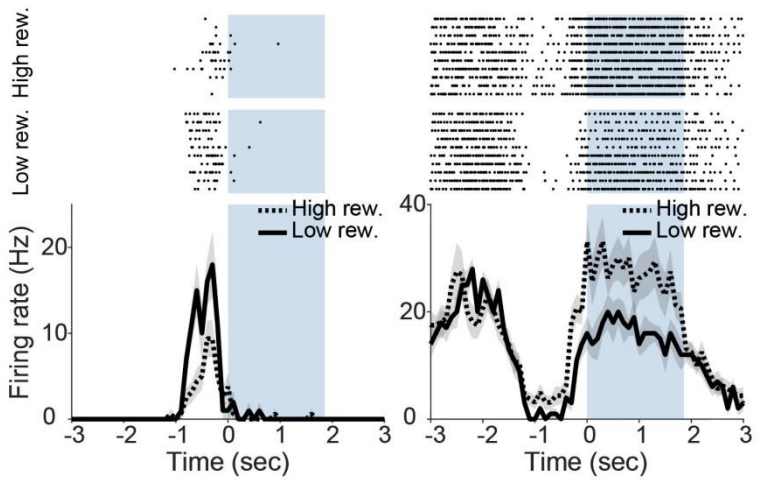

Figure 2.7. Single units encoding the amount of reward at side feeders while the rat was at the central feeder. Shown are two separate neurons which discriminate between low- and high-reward right before or while consuming a fixed amount of reward at the central feeder. Top and bottom plots show corresponding raster and time-warped peri-event time histograms (bin-size: 100 msec) of neural spikes respectively. Blue background highlights the duration of reward consumption at the central feeder. Start and stop of the feeder period was determined based on the rat's velocity, which drops below a certain threshold while feeding.

2.3.3 ACC populations provide high fidelity prediction of path, effort and reward

Our single-unit counts tell us what fraction of the population of ACC cells is involved in coding task parameters. To assess the strength of the signal, we used ensemble activity within each session to predict upcoming parameter values.

For all 12 sessions with sufficient neurophysiological data, we trained a SVM classifier with linear kernel to predict either path, effort, or reward based on the instantaneous firing rates of all cells within a spatial bin. Despite the limited number of cells used in the decoding (23-99 cells), average prediction accuracy for all sessions was well above the 50% chance level for all parameters in almost every bin (Fig. 2.8). This population decoding data has strong parallels with the overall number of cells encoding parameters (Fig. 2.6A). Path again is predicted with the highest fidelity followed by effort and then reward. Path prediction accuracy was strongly influenced by the rat's position, being very

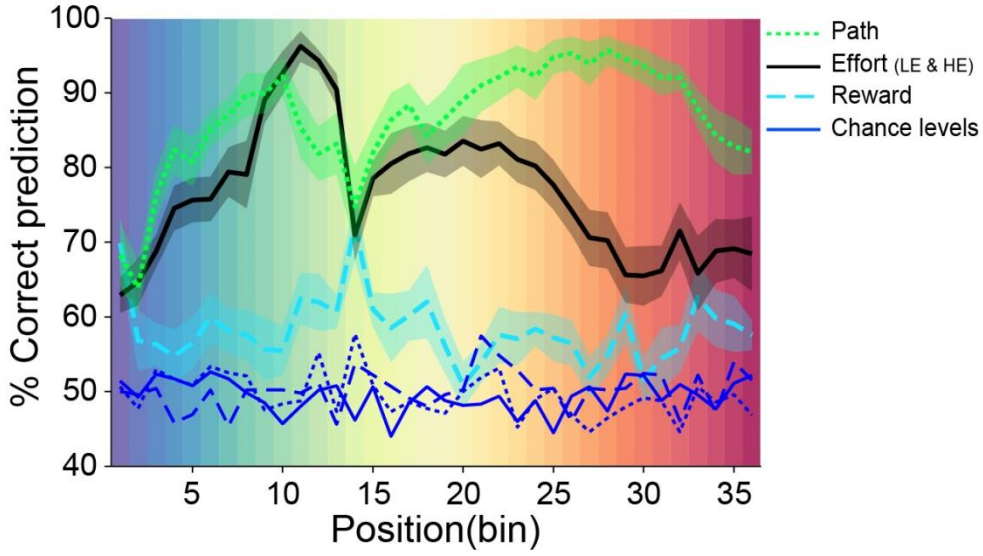

Figure 2.8. Accuracy of decoding effort level, reward amount, and path by ACC ensemble activity. The average percentage of correct prediction of all 12 sessions, 10 times each, for effort, reward and path is illustrated. Blue lines illustrate corresponding chance levels to each parameter. Standard error of mean is represented by shaded area around the mean. Color map of the figure background is provided in the inset of Figure 2.6A.

high whenever the paths were separate and relatively low when paths overlapped (bin 2-4).

A similar pattern was seen with effort and reward decoding accuracy, though not to the same extent. To keep the number of classes the same across all variables, we used only low- and high-effort trials for effort decoding. Decoding of the existence or absence of the barrier was extremely accurate during the climb, even better than that for path, suggesting that encoding of effort is particularly robust at this point. Effort decoding accuracy dips during the time at the side feeder and then reaches a plateau level just above 80% during the down ramp (bins 16-23). Reward decoding shows increasing accuracy as the rat moves from the central feeder to side feeder and, though it peaks slightly at the side feeder, still shows moderate prediction accuracy (~60%) throughout the rest of the trajectory back to the central feeder. A small peak in the decoding of all

three parameters (higher for path and reward) is apparent at central feeder, consistent with the pattern observed in the number of cells encoding reward.

2.3.4 Effort and reward encodings are stable during approach

Given recent evidence that the encoding of motion parameters in motor cortex can shift dynamically within a reach trajectory (Churchland *et al.*, 2010), we wondered whether ACC cells tuned to effort and reward might show a similar shift of tuning properties during a trial. We separated each trial into preparatory (stem of the T) and approach phases (arms of the T) and compared the parameter tuning between these two phases for each cell. Figure 2.9A shows specific examples of cells with shifting (top row) and stable (bottom row) tuning properties. To quantify this effect, we found the bins in the preparatory and approach phases with the strongest discriminatory activity and then computed the correlation of the firing rate between these two bins across trials (details in the Methods section). If the cell's tuning was stable across the trial (e.g., it preferred high-effort on both preparatory and approach phases), the correlation would be positive. Conversely, a cell whose tuning switched would have a negative correlation. A correlation of zero would mean the cell was not tuned to the parameter in at least one task phase. Histograms of the across-phase correlation coefficients are shown in Figure 2.9B. The histograms are notably skewed to positive values, indicating that the majority of cells showed stable tuning to both effort and reward across the trial. Cells with negative correlations were present (Fig. 2.9A top) but were no more frequent than that seen when we shuffled the order of trials. We hence conclude that effort and reward coding in ACC are stable across the trial.

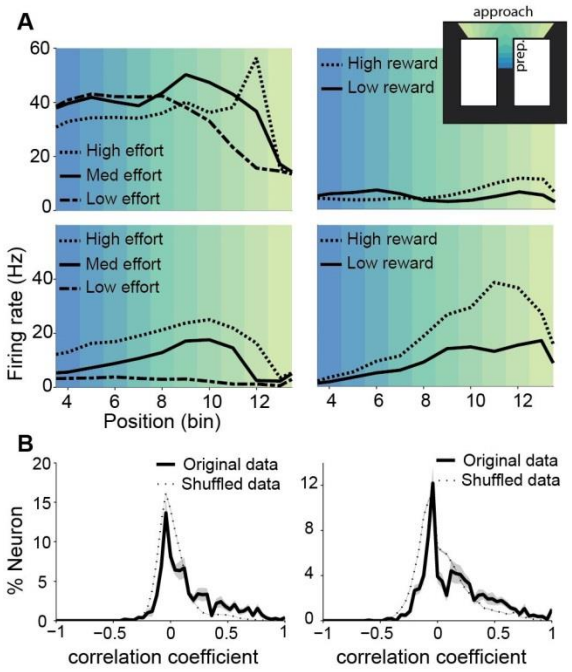

Figure 2.9. Stability of neural encoding during preparatory and approach phases. A, Average firing rate of four example neurons in effort (left) and reward (right) conditions are illustrated. Cells in the top two plots exhibit unstable encoding in that their preference for effort and reward switches between the preparatory phase (T-stem, bin # 4-8) and the approach phase (T-arms, bin # 9:13). The cell in upper right, for example, prefers low-reward during preparation and high reward during approach. The cells on the bottom row, in contrast, show stable encoding across the trajectory. Inset (upper right) shows regions classified as preparatory and approach as well as the color coding of bins. B, Histogram showing the averaged percentage of neurons having stable (positive correlations) or unstable (negative correlations) tuning for groups of trials that varied in effort (left) or reward (right). As described in the main text, correlation coefficients were computed for firing rates of cells between preparatory and approach phases. Dashed lines show chance levels of correlations, estimated by shuffling data across trials 1000 times. Standard error of mean is indicated by shaded area around the plots.

2.3.5 ACC neurons which jointly encode effort and reward are biased towards positive net utility

We next examined the joint encoding of effort and reward. A simple Pearson correlation coefficient was used to quantify the degree to which a given cell's firing co-varied with effort or reward. Scatterplots showing the relationship between tuning to effort and reward are provided in Figure 2.10A for all analyzed cells (647 in total). If neurons

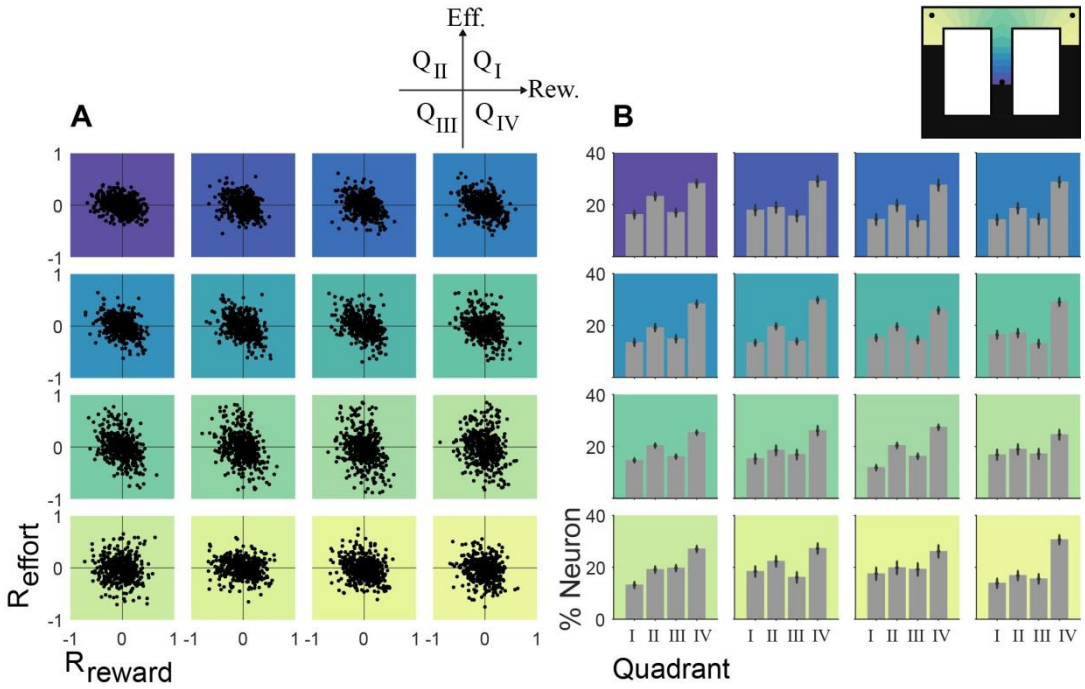

Figure 2.10. Joint encoding of effort and reward. A, Data shown is the correlation coefficient of firing rate and effort levels versus the correlation coefficient of firing rate and reward amounts for each neuron (a total of 647). The plane of R values is divided into four quadrants which are numbered as shown (upper right inset). B, In the first 16 spatial bins, the average percentage of units falling in each quadrant (over 12 sessions) is shown. Error bars show standard error of the mean. Inset (upper right) shows the maze color map corresponding to square color bars on the upper right corners of figures.

encode effort and reward separately, the data should be distributed either vertically (on y axes) or horizontally (on x axes). Instead, we see an over-representation of cells in the upper left and lower right quadrants (Q_{II} and Q_{IV}, respectively), as is quantified in Figure 2.10B. However, this Q_{II}–Q_{IV} tuning shifts toward effort and reward axis at the barrier and reward zones, respectively. Across all bins from the central feeder to the side feeder, the most common cell type was that showing a negative correlation with effort and positive correlation with reward (E-/R+). These cells hence encode the positive net utility of the chosen path. The next most common type of cell was found in Q_{II}. These cells respond most robustly to high-effort and low-reward (E+/R-), thus signalling negative net utility. These negative net utility cells are more prevalent than Q_I and Q_{III} cells only

around central feeder and during the approach to the climb but not during the climb (bins # 12 and 13). Peri-event time histograms showing specific examples of positive and negative net utility cells are provided in Figure 2.11.

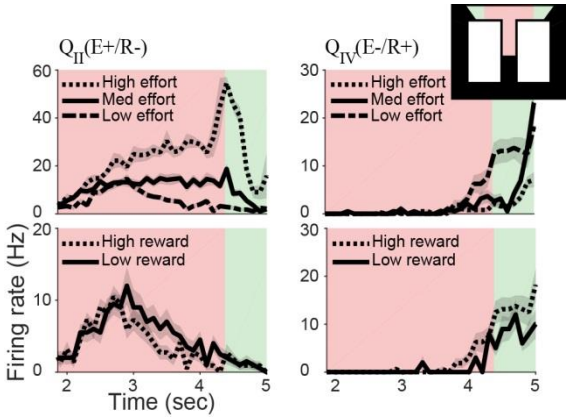

Figure 2.11. Examples of single units that encode positive and negative overall net utility. Time-warped peri-event time histograms (bin-size: 100 msec) for two neurons that encode positive (right column) and negative (left column) utility between central and side feeder are shown. The time at which the rat starts consuming food at the central feeder is time zero (not shown). Background color corresponds to the color-coded sections of the maze shown in the inset. Green indicates the climb. The positive utility cell (Q_{IV}) shows higher firing rate for low-effort and high-reward conditions, whereas the firing rate of the negative utility cell (Q_{II}) is higher for high-effort and low-reward conditions.

2.3.6 The balance of negative and positive utility signals changes on approach to effort or reward

The preceding section showed that, in ACC, there is an over-representation of cells representing positive and negative net utility. To provide a more detailed picture of how the signal from these sub-populations changed through time, we chose positive and negative utility cells, determined in one spatial bin (bin#6), and tracked their averaged firing rate over the entire maze. This spatial bin was specifically chosen to be well before the climb and reward so that proximity to these events would have less influence on our estimates of cell tuning preference (see Fig. 2.10). As illustrated in Figure 2.12A, the locations at which these cells provide discriminatory activity are almost exclusively

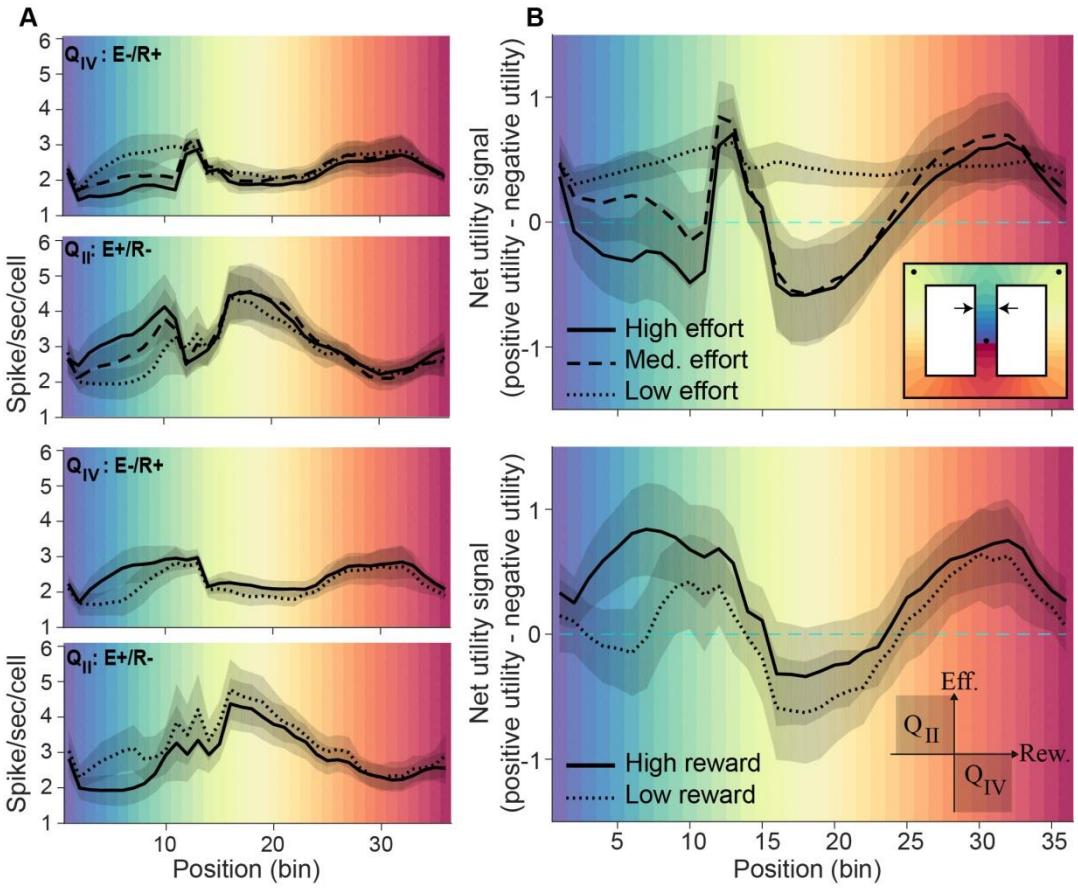

Figure 2.12. Differences between positive and negative net utility signals. A, Average spike rate of all neurons in fourth quadrant (cells encoding positive net utility, E-/R+) and second quadrant (cells encoding negative net utility, E+/R-) when the animal is in spatial bin#6, separated by effort (top two) and reward (bottom two). Legends are the same as the corresponding figure in panel B. B, The difference between the overall firing rates of neurons in fourth quadrant and second quadrant as the animal was in spatial bin#6 and tracking them all over the maze in effort (top) and reward (bottom) conditions. The average of both panels is on 12 sessions and the shaded area around each plot indicates standard error of mean. (top inset: maze color map, the bin#6 is indicated by arrows; bottom inset: highlighted quadrants corresponding to positive and negative utility cells).

located between the central feeder and arrival at one of the side reward feeders. This observation is consistent with the idea that this activity is related to the net utility of the choice. We next computed the total net utility signal available from both positive and negative utility cells by performing a weighted average of the activity of all cells in Q_{IV} and subtracting the weighted averaged activity of all cells in Q_{II}. As illustrated in Figure 2.12B (top), when no barrier is presented, cells encoding positive net utility (Q_{IV}) provide

a stronger signal than negative utility cells (Q_{II}) in all spatial bins. However, in the trials with medium- and high-effort, negative utility cells almost outbalance positive utility signal when approaching the climb as well as descending the ramp. Figure 2.12B (bottom) shows that, when approaching the feeders, positive utility cells (Q_{IV}) fire more strongly than negative utility cells (Q_{II}). This effect becomes reversed after leaving the side feeder. Although the effect is similar for both high- and low-reward, the strength of the positive utility cells outweighs the negative utility cells more in the high-reward condition than the low-reward condition specifically over the T-stem, as one might expect. Another peak is observed on the returning path toward the central feeder. Overall, these data suggest that, for more valuable conditions (LE and HR), ACC encodes a signal which is biased toward positive utility and this effect is reversed for more effortful or less rewardable trails particularly during the approach to the barrier and over the ramp.

2.4 Discussion

We investigated the spiking activity of simultaneously recorded ACC ensembles in response to effort, reward, and path while the animal was performing a figure-8 cost-benefit decision-making task where the cost was the effort of climbing a barrier. Our results show that left and right paths were discriminated by about one-fifth of neurons at all maze locations, although the effect was reduced where the paths overlap. The second and third most influential parameters were, in turn, effort and reward. While the dominance of effort over reward was partially accounted for by the neurons responding to the emergence of the barrier (Figure supplement 2.2), we also found that effort encoding (even medium- versus high-effort) was generally more robust than reward coding at the single cell level (e.g., Fig. 2.5). This finding is consistent with our decoding results in which we could predict the level of effort with a higher accuracy when using barrier-

present and barrier-absent trials, rather than using high- and medium-barrier trials (Figure supplement 2.3). However, the prediction accuracy for the later condition is still higher than reward decoding (Figure 2.8 and Figure supplement 2.3). The proportion of cells encoding effort started to ramp up as the rat left the central feeder and approached the barrier, peaked at the barrier with more than 20% of the population, then dropped greatly after climbing. However, the reward discrimination peaked at the reward feeders with about 15% of neurons encoding this parameter. Using decoding methods on simultaneously recorded neural activity patterns, we showed that discriminability for effort, reward, and path was generally quite good everywhere except on the common stem of the maze before the rat turned left or right. Further, these patterns showed the same general pattern observed with the proportion of encoding cells in that path was the most strongly encoded feature, followed by effort and then reward.

It is also worth noting that the effort encoding is not strongly influenced by animals' running velocity. Supplemental figure 2.1 shows some examples of one of the rats' velocity in each condition. The velocity profile appears to stay fairly similar between the conditions; only at the climb zone it is decreased when the barrier exists. If the effort selectivity observed in our cells was due solely running speed, one would expected elevated firing at other points on the maze where the animal slowed down, such as on approach to feeders. This was not the case (data not shown).

2.4.1 Decision making versus action maintenance

Effort-coding in ACC started to dominate as the animals' left and right path trajectories diverged, while the effort-reward decision would have had to have been made much earlier, on the common segment. The proportion of neurons responding to reward started to ramp up even later and peaked at the location of target feeder. On the common

segment of the T-maze, where decision demands are presumably greatest, the ACC demonstrated the lowest discriminant activity for effort, reward, and path of anywhere on the maze. While our decoding results still show above-chance predictions of upcoming choices on the common segment, the most parsimonious interpretation of these results is that the ACC doesn't drives decision processes in this task. Some primate studies also have reported gradually increasing ACC activity prior to reward delivery but after a decision has been made (Shidara & Richmond, 2002; Cai & Padoa-Schioppa, 2012; Blanchard & Hayden, 2014; Blanchard *et al.*, 2015). If ACC doesn't make decisions, which region does? Prior research has shown that secondary motor cortex signals choice ahead of behavioural manifestation of the animal's choice (Sul *et al.*, 2011), while the various medial prefrontal structures, including the ACC, starts to signal choice later, when the animal's behaviour actually diverges (Sul *et al.*, 2010). In another single-unit study of rat's ACC on a different effort-reward decision-making task, path encoding appeared higher than chance level only where paths diverged and the encoding of reward and effort before the decision were negligible (Cowen *et al.*, 2012). In a decoding study, our group has also shown that populations of ACC neurons provide high-resolution decoding of the rat's present location, not where he is going. In short, the ACC seems to encode what is, not what shall be.

Cowen (2012) has suggested that the ACC may be necessary for persistence on a plan of action once chosen (Cowen *et al.*, 2012). This is consistent with our single-cell results, as encoding of effort and reward is quite robust after the choice until the ramp is climbed and reward achieved. It is also consistent with an observation from a behaviour study using the same task where ACC-lesioned rats turned into the high-reward arm but then backed away just before climbing the barrier and eventually chose the alternative

arm (Holec *et al.*, 2014). The ACC event horizon is apparently too short for it to play a role in the decision, but long enough for it to anticipation immediate outcomes, such as expected effort and reward, which helps the rat maintain the course of action until goals are achieved.

On the other hand, a small but significant fraction of neurons predicted upcoming reward magnitude on the side feeders while the rat was pausing at the central feeder, long before a choice was actually made. Although the small peak at this location, which appeared in the encoding results (Fig 2.6B), doesn't seem compelling at the first glance, we found many examples of cells whose activity robustly differentiating between low and high levels of the side feeder reward (two of which are presented in Fig. 2.7). Our decoding results were also consistent with this finding, showing a slight peak in prediction accuracy near the central feeder (Figure 2.8). Similar prospective activity was found for path in both encoding and decoding results (See also Fig 2.5C for two example neurons). However, comparing with the rest of the maze, path encoding at the central feeder is weak. Such an effect was not observed for effort; thus, it seems to be specific to reward and side. Although the prospective activity of ACC neurons in response to reward expectation has been observed in another study, it seems to be different in pattern and character. In a study by Blanchard et al. (Blanchard *et al.*, 2015) with monkeys, the ACC ensemble activity observed after choice was reported to gradually increase and come to resemble the post-reward pattern, while, we observed that the reward-related prospective activity diminishes as the animal leaves the central feeder and then ramps up as the animal is approaching the side feeder. Also, they found this activity during a delay in anticipation of the reward, but we observed it when the animal paused around reward zone. We have previously shown that population activity at side feeders often

momentarily switches to a pattern seen on the alternate side feeder, as though the rat was imagining himself at that alternate location (i.e., the path not taken) (Mashhoori *et al.*, 2018). This mental excursion to the remote reward location was shown to occur more after disfavoured reward (low-reward). Hence, the idea that the pause at reward zone is a time when the brain reflects on alternate course of action is not new. However, this is the first report of such predictive activity in the ACC. Several groups have reported that while rats pause after reward, cells in several regions show prospective encoding of trajectories. This phenomenon has been reported in orbitofrontal cortex (Steiner & Redish, 2014) and hippocampus (Foster & Wilson, 2006; Carr *et al.*, 2011). It has been speculated that this activity may relate to the switch of perspective ascribed to the "default mode" network in humans, of which the ACC in rats appears to be a member (Buckner & Carroll, 2007; Buckner *et al.*, 2008; Lu *et al.*, 2012).

2.4.2 So, what does ACC care about?

Although it has been suggested that the dominant role of ACC is in the evaluation of the amount of required effort to achieve a reward (Walton *et al.*, 2002; Walton *et al.*, 2003), the results of the present study indicate that path information in ACC is even stronger. The ACC neurons differentiate between the chosen left and right paths more robustly than effort and reward at every location on the maze. This finding implies that spatial coding might take precedence over effort-reward calculation in ACC. This finding is not an artifact of our task design. Both high- and low-reward and high- and low-effort were encountered on both side feeder locations. Hence, our task design allows us to separate the influence of side from effort and reward. We have recently shown that activity in ACC is completely sufficient to decode the position of the animal on this task, further strengthening the argument that ACC plays an important role in tracking the

animal spatial position (Mashhoori *et al.*, 2018). One might argue that the position is not independent of existing manifestations of effort and reward in the task. Information about the proximity of barriers and reward wells certainly also conveys information about the animal's position on the maze. However, the decoding algorithm was blind to the amount of effort and reward on any given trial and, despite the varying levels of these variables, still decoded position with extremely high accuracy (Mashhoori *et al.*, 2018). While not frequently studied, there are similar reports of strong spatial encoding in ACC and other mPFC structures in both rats (Euston & McNaughton, 2006; Cowen & McNaughton, 2007; Fujisawa *et al.*, 2008) and primates (Strait *et al.*, 2016). Also, other primate studies who monitored animals eye movements have found spatial selective neurons in ACC (Kennerley *et al.*, 2009) and more specifically in dorsal part of ACC (Cai & Padoa-Schioppa, 2012).

In a previous study, Cowen et al has shown that the number of ACC neurons responsive to path and effort is much higher than those responsive to reward (Cowen *et al.*, 2012). At the first glance, our results also show that after path, effort has a much stronger effect on ACC neurons than reward; however, part of this was due to a design limitation: we had three levels of effort (zero, medium, and high) and only two levels of reward (low or high). One can imagine that effort coding would have been much more prominent had we contrasted some reward against no reward. Hence, ramp, which included a no-ramp condition may have had a structural advantage in driving ACC activity. Further, the use of three levels increased our statistical power to detect effort-encoding cells. When we looked at only the top two levels of effort, we found that the proportion of cells encoding effort was similar to that encoding reward (Fig. 2.6). The fact that the detected proportion of effort-encoding cells dropped so dramatically when

we excluded the no-ramp condition suggests that the contrast between no ramp and some ramp is more important to these cells than the contrast between a medium- and high-barrier (Figure supplement 2.1). On the other hand, reward itself still had a discriminatory effect on neural firing rates, even when the required physical effort was minimal. This finding contrasts with the lack of reward coding in the nearly identical study of Hillman and Bilkey (2010) when ramps were zero on both arms (Hillman & Bilkey, 2010), but consistent with a large literature showing reward-magnitude encoding in medial prefrontal regions (Pratt & Mizumori, 2001; Burton *et al.*, 2009; Gruber *et al.*, 2010; Euston *et al.*, 2012). Perhaps the lack of reward coding without effort was simply a result of the low number of cells in the earlier Hillman and Bilkey's study.

Our analysis didn't show any significant difference between ACC neural activity on free versus forced trials. This is particularly remarkable because there were significant physical differences between the two conditions, namely the physical gates that are either raised to block access to one of the arms or lowered to allow free choice. Despite this difference, the ACC activity was nearly identical whether the animal made a choice of their own free will or were directed by the gates. This suggests that the ACC activity is not influenced by the volitional state of behaviour. Another reason this finding is particularly surprising is that the cingulate cortex is often implicated in the initiation of self-generated actions, especially in human/primate studies (Mueller *et al.*, 2007; Passingham *et al.*, 2010). Internally guided actions are referred to those decisions in which no external cue tells the animal which action to make; instead, the animal decides based on internal states (Passingham & Lau, 2006). The neural activity in these volitional behaviours was first compared with externally triggered actions in early studies (Okano & Tanji, 1987; Romo & Schultz, 1987) and supplementary motor area (SMA and pre-SMA)

and ACC were found to be involved (Romo & Schultz, 1987; Shima & Tanji, 1998a; b; Mueller *et al.*, 2007). Human research has shown discriminative activity in these regions well before the subjects' awareness of their internal states, decision, or intention to make an action (Fried *et al.*, 2011). A finer resolution study in monkeys has shown that the dorsal ACC is most active before movement initiation while the ventral bank of ACC encodes reward delivery, but neither of them seem to encode pre-decision variables (Cai & Padoa-Schioppa, 2012). It may be that the structure of our task, perhaps the extensive pre-training, minimized the salience of gates or obviated the need to attend to the free/force trial status. It is also worth noting that the evidence for ACC involvement in movement initiation is largely based on primate and human studies, so there may be a species difference or confusion about homologous structures (Heilbronner *et al.*, 2016).

2.4.3 Net utility of a choice

In the effort-reward decision-making paradigm, selecting the optimum option ultimately depends on weighing both effort and reward. Our behavioural results confirm that both variables influence the animal choice. If effort and reward are separately encoded in ACC, this region might be simply a relay to another area which integrates the necessary information to eventually make the decision. By contrast, if the encoding of these variables is integrated, it might suggest that ACC is involved in combining these variables to compute net utility. Certainly, primate studies suggest that multiplexing is a signature of ACC encoding (Kennerley *et al.*, 2009). By investigating how effort and reward are jointly encoded, we aimed to address this issue. We have observed that some ACC neurons encode only effort and others only reward. However, the majority of cells encoded a combination of these two task parameters. It should be noted that this contrasts with the findings of Cowen et al (2012), who found few cells which jointly encoded both

effort and reward (Cowen *et al.*, 2012). This may be because effort and reward occurred at separate locations on the track compared to our study, where reward was delivered at the top of the ramp. The firing rate of the most of these neurons in our study correlated with effort and reward in opposite directions, consistent with the idea that they encode net utility. A significant portion of neurons fired most strongly for low-effort and high-reward, suggesting that they encode cost-discounted potential gain, typically referred to as positive net utility (E-/R+). Others fired most strongly for high-effort and low-reward, perhaps signalling negative net utility (E+/R-). These utility-like signals peak just before the ramp (and reward) are encountered, consistent with the results we obtained when analyzing each variable separately (compare Figures 2.6 and 2.12). Of note, these signals are also largely stable throughout a trial (Fig. 2.9). Unlike cells in the motor cortex, which have been shown to shift their tuning to preferred direction during a movement (Churchland *et al.*, 2010), we found the tuning of cells to effort and reward to be stable across the preparatory and approach phases (e.g. a reward cell whose firing is positively corrected with level of reward is unlikely to show a negative relationship between firing and reward in another location).

The majority of the evidence for ACC encoding comes from human neuroimaging studies, and here the picture is far from clear. ACC activity has been shown to correlates with effort, reward, as well as the net value of the choice in both positive and negative directions (Croxson *et al.*, 2009; Prevost *et al.*, 2010; Kurniawan *et al.*, 2013; Klein-Flugge *et al.*, 2016; Chong *et al.*, 2017). For instance, ACC activity was negatively correlated with subjective value in one study (Prevost *et al.*, 2010), and positively correlated with the same variable in another study (Klein-Flugge *et al.*, 2016). When the cost of the decision is a physical activity of squeezing a bulb, ACC activity was shown to

increase with the higher level of effort (Prevost *et al.*, 2010; Kurniawan *et al.*, 2013), while other human studies have found a negative correlation between ACC activity and the level of effort (Croxson *et al.*, 2009; Klein-Flugge *et al.*, 2016). Also, the level of reward had a negative correlation with ACC activity in some fMRI studies (Chong *et al.*, 2017) and positive in others (Croxson *et al.*, 2009; Klein-Flugge *et al.*, 2016). Similarly, a mixture of different relationships between ACC firing rate and the amount of reward or required effort has been reported by animal studies (Heilbronner & Hayden, 2016) . For instance, some studies have observed ACC neurons negatively responded to the level of effort (Porter *et al.*, 2019), or reward (Luk & Wallis, 2009; Hayden *et al.*, 2011b; Kennerley *et al.*, 2011); whereas, others have reported ACC neurons positively tuned to the reward (Hayden *et al.*, 2009; Hillman & Bilkey, 2010) and still other studies have found both negative and positive encoding of reward in ACC (Hayden *et al.*, 2011a; Blanchard & Hayden, 2014). In light of these findings, it is worthwhile to note that we also see evidence that the ACC population is biased towards positive and negative net utility (see Figure 2.10, where more cells are in Q_{IV} and Q_{II} than elsewhere) and these signals are encoded simultaneously by separate subpopulations of intermixed cells. Similar sub-populations have been reported in monkey pregenual ACC using a cost-benefit task involving aversive air-puffs and reward (Amemori & Graybiel, 2012). In their study, both positive and negative net utility cells were evenly distributed except in one anatomical region where the negative cells where more numerous. Whether such a region exists in rats remains to be determined.

When we examined the net firing rate of negative and positive utility cells combined, we find that the balance between these two signals changes in different phases of the task and as the value of the outcomes change (Fig. 2.12B). For low-effort and high-

reward trials in which the outcomes possess more value, positive utility signal outweighs the negative utility in most spatial bins, except after leaving the feeder with high-reward condition. This effect becomes less or even reversed as the value of outcomes decreases (higher-effort or low-reward). These results suggest that ACC is an element of a computational network in the brain involved in assessing the utility of an ongoing action. These data may thus be critical in keeping the animal motivated on the task. It is consistent with recent findings of the communication between ACC and VTA, a substantial component of motivational network, in effort-reward decision processes (Elston & Bilkey, 2017; Elston *et al.*, 2018; Elston *et al.*, 2019). Specifically, in a task which requires the rat to climb a barrier in only some random trials in order to receive a reward, the ACC is shown to signal the VTA when the rat faces no barrier (Elston & Bilkey, 2017). Also, ACC-to-VTA signaling was shown to be related to the anticipation of task reversal, whereas VTA-to-ACC signals were observed in response to errors in outcomes (Elston *et al.*, 2019). Taken together, our results are consistent with the idea that ACC integrates both effort and reward information to convey the net utility of a condition.

2.5 Figure supplement

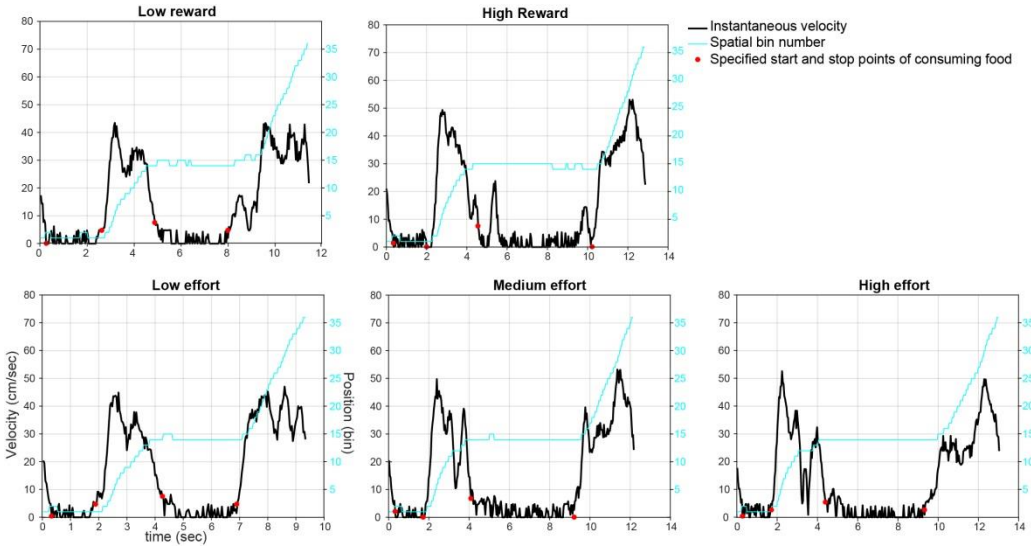

Figure supplement 2.1. Examples of an animal's running velocity in five different trials. The instantaneous animal's running velocity is illustrated in five different trials, each correspond to a specific condition.

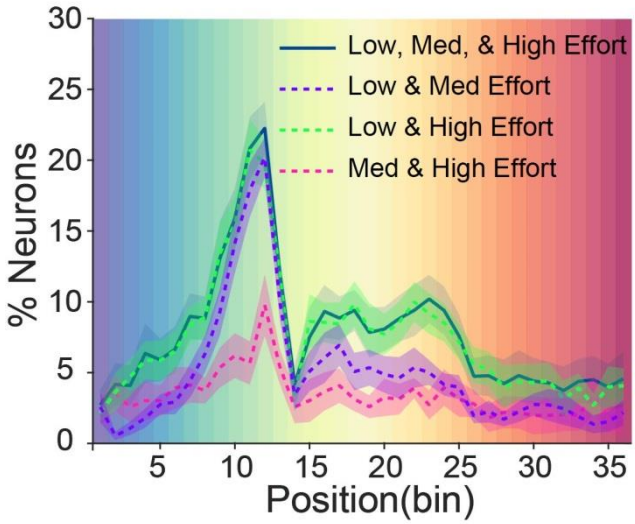

Figure supplement 2.2. Population data on effort selectivity. The average percentage of neurons discriminate different pairs of barrier levels in each spatial bin of the maze are plotted along with the average percentage of neurons differentiate at least one of the three levels of the barrier. Shaded area around each plot indicates the standard error of the mean. Background colors correspond to bin positions illustrated in the inset of panel A of Figure 2.6. (ANOVA, $p < 0.05$ and effect size>0.1379)

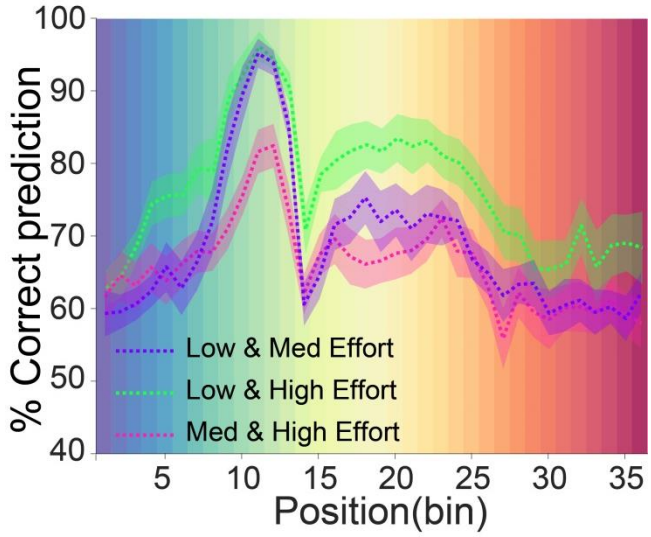

Figure supplement 2.3. Accuracy of decoding effort level using pairwise comparisons of the barrier height. The average percentage of correct prediction of all 12 sessions, 10 times each, for effort amount using trials with two levels of the barrier is illustrated. Shaded area around each plot indicates the standard error of the mean. Background colors correspond to bin positions illustrated in the inset of panel A of Figure 2.6.

3 CHAPTER 3: Amphetamine-induced reduction of effort-reward utility encoded by rat anterior cingulate cortex

Abstract

D-amphetamine (AMPH) is well-known to affect catecholamine transmission in the brain and behavioural choice during tasks involving effort and reward. A primary cortical target of catecholamines, particularly dopamine, is the anterior cingulate cortex (ACC), which is implicated in effort-reward choice. Thus, the effect of AMPH on behaviour may involve the modulation of information processing in the ACC. We tested this hypothesis here by quantifying the effect of vehicle or AMPH (0.5, 1, 1.5 mg/kg) on the encoding of effort, reward, and utility in the ACC of freely-moving rats. Rats performed a two-alternative task in which they received either a large or small reward by exerting either high- or low-effort (barrier climb). Consistent with previous data, we found that individual ACC neurons encoded effort, reward, and utility, among other variables. AMPH had dose-dependent effects on this encoding as well as behaviour. Low doses increased task engagement and had only mild effects on neural encoding, whereas high doses disrupted utility signaling and decreased animals' engagement in the task. The disruption was primarily mediated by a reduction in reward encoding and compression of the joint effort-reward encoding of utility signaling cells. Interestingly, animals spent less time at feeders and increased the number of times they bypassed feeders with increasing AMPH, suggesting that the reduction in reward encoding may be manifested as reduced interest in the reward. Multivariate analyses of simultaneously recorded ensemble activity revealed that low-dose AMPH caused a contraction of neural trajectories in state space and an increase in the neural trajectory velocity within the space, which mirrored the increase in running speed in rats. At high doses, the state-space occupancy expanded and

trajectories became highly variable, which was concordant with frequent disengagement in the task. Besides its attenuation of value signaling, we suggest that AMPH has two effects on population dynamics: it stabilizes the dynamics underlying learned neural trajectories in ACC, and increases the excitability of the system. At low doses, these balance such that the systems evolves more quickly (because of the increased energy of increased excitability), but stays within the bounds of learned trajectories. At high doses, the additional excitation overwhelms the stability effects such that the neural activity deviates from the attractor basin of task-related ensemble patterns, and thereby breaks behaviour.

3.1 Introduction

The control of volitional behaviour involves complex interactions among several brain areas, each of which appears to have some specialization of processing information about the many factors pertinent to action selection (Balleine & Dickinson, 1998; O'Doherty, 2004; Daw *et al.*, 2006; Doya, 2008; Pennartz *et al.*, 2009; Gruber & McDonald, 2012). Empirical studies of choice among options with varying costs and benefits indicate involvement of several brain structures, including the medial prefrontal cortex (mPFC), medial striatum, and midbrain dopamine (DA) neurons in rodent (Walton *et al.*, 2002; Salamone *et al.*, 2003; Walton *et al.*, 2003; Schweimer & Hauber, 2005; 2006; Lapish *et al.*, 2008; Hauber & Sommer, 2009; Kim *et al.*, 2009; Walton *et al.*, 2009; Hillman & Bilkey, 2010; Sul *et al.*, 2010; Braun & Hauber, 2011; Pasquereau & Turner, 2013; Holec *et al.*, 2014; Elston & Bilkey, 2017). This is consistent with data in primates (Samejima *et al.*, 2005; Kennerley *et al.*, 2006; Seo & Lee, 2007; Amemori & Graybiel, 2012), suggesting at least partial conservation of function (Murray *et al.*, 2017). The results of these studies suggest that cortico-striatal networks and their

neuromodulation underlie evaluation of choice options and execution of the subsequent action (Seo & Lee, 2007; Doya, 2008; Floresco *et al.*, 2008a; Assadi *et al.*, 2009; Kurniawan *et al.*, 2011; Euston *et al.*, 2012).

Rodents readily learn to exert effort in order to acquire food rewards. Intact rats typically are willing to exert more effort, such as barrier climbing or lever pressing, if the associated reward is of considerably higher value than that of lower-effort options. This preference is reduced or eliminated by lesions of the Anterior Cingulate Cortex (ACC) (Walton *et al.*, 2002; Walton *et al.*, 2003; Schweimer & Hauber, 2005; Holec *et al.*, 2014), by lesions of ventral striatum (VS) (Hauber & Sommer, 2009), or by disruption of DA transmission in these areas (Salamone *et al.*, 1994b; Cousins *et al.*, 1996; Schweimer *et al.*, 2005; Schweimer & Hauber, 2006; Floresco *et al.*, 2008b; Bardgett *et al.*, 2009; Mai *et al.*, 2012). Systemic administration of D1-type and/or D2-type DA receptors antagonists (Denk *et al.*, 2005; Floresco *et al.*, 2008b; Bardgett *et al.*, 2009; Hosking *et al.*, 2015), as well as intra-ACC D1 receptor blockage (Schweimer & Hauber, 2006), show a similar effect, whereas intra-ACC D2 receptor antagonists (Schweimer & Hauber, 2006) or systemic administration of D3 agonists/antagonist (Bardgett *et al.*, 2009) do not affect effort-reward preference. These data suggest that drugs affecting DA transmission via D1/D2 receptors will influence choices involving effort and reward. One such drug is d-amphetamine (AMPH). A prominent effect of AMPH is the extracellular increase of DA and norepinephrine in the striatum (Besson *et al.*, 1969a; Besson *et al.*, 1969b; Chiueh & Moore, 1973; Zetterstrom *et al.*, 1983; Hernandez *et al.*, 1987; Di Chiara & Imperato, 1988; Avelar *et al.*, 2013) due to its action on uptake transporters for these transmitters (Seiden *et al.*, 1993; Easton *et al.*, 2007; Avelar *et al.*, 2013; Covey *et al.*, 2013). AMPH also increases extracellular concentrations of these transmitters in the

neocortex by interfering with their enzymatic degradation in synapses (Shoblock *et al.*, 2003; Pum *et al.*, 2007; Ren *et al.*, 2009). This is particularly important in the dorsal mPFC, including the ACC, because of its dense DA innervation relative to other cortical regions (Berger *et al.*, 1991; Ash *et al.*, 2014). Although the mechanism is not well understood, evidence strongly indicates that choice preference is biased toward high-effort, high-reward options when AMPH administered systemically (Floresco *et al.*, 2008b; Bardgett *et al.*, 2009). It remains unclear what aspects of information processing are affected by AMPH to produce this effect.

Effort, reward, and other features pertinent for costly economic decisions are often formalized by the concept of utility (Phillips *et al.*, 2007). Options requiring low effort but yielding a large food reward to a hungry animal have high utility, whereas options requiring high effort or yielding little/unwanted food have low utility. The bias toward high-effort, high-reward options under AMPH could reflect a perceived discounting of effort and/or an increase in perceived reward value when the drug is on board. Behavioural evidence for this is mixed. Animals are willing to exert more physical effort under AMPH (Floresco *et al.*, 2008b; Bardgett *et al.*, 2009), but it is unclear if this reflects changes in utility (e.g. devaluation of effort) or is a by-product of generally increased motoric output (Shin *et al.*, 2010; Wardle *et al.*, 2011). The behavioural data for effects on reward value is enigmatic. AMPH-treated rats appear less motivated to eat; their latency to first consumption is longer, they consume less than usual, and they spend less time on eating their meals than when no AMPH is administered (Blundell *et al.*, 1976; Blundell *et al.*, 1979; Leibowitz *et al.*, 1986). Suppression of food-intake by AMPH has been reported in other types of rodents (Cannon *et al.*, 2004; Wellman *et al.*, 2009), as well as in primates (Foltin & Fischman, 1988; Foltin, 2001). These findings suggest that

the perceived reward value is decreased. On the other hand, AMPH induces rodents and primates to become more engaged in food-seeking activity (Evans & Foltin, 1997; Foltin & Evans, 1999; Foltin, 2001; Odum & Shahan, 2004). This suggests that they are more willing to do work for the reward, as if the food had higher utility, but then show less interest in consuming the reward, as if the food had lower utility. The behavioural data, therefore, do not support a clear prediction about how AMPH may influence the neural encoding of utility.

Electrophysiological recordings indicate that rat ACC and nearby regions in mPFC encode a variety of signals related to choice. These include the position of the animal (Euston & McNaughton, 2006; Fujisawa *et al.*, 2008; Mashhoori *et al.*, 2018), task phase (Lapish *et al.*, 2008; Balaguer-Ballester *et al.*, 2011), reward (Gruber *et al.*, 2010; Cowen *et al.*, 2012), choice (Cowen *et al.*, 2012), effort (Cowen *et al.*, 2012; Hashemniayetorshizi *et al.*, 2015), and other task-related features (Cowen & McNaughton, 2007; Gruber *et al.*, 2009; Durstewitz *et al.*, 2010; Sul *et al.*, 2010). These are consistent with findings in monkeys and humans (Isomura *et al.*, 2003; Kennerley *et al.*, 2006; Croxson *et al.*, 2009; Kennerley *et al.*, 2009; Kennerley & Wallis, 2009b; Cai & Padoa-Schioppa, 2012; Apps & Ramnani, 2014; Skvortsova *et al.*, 2014; Blanchard *et al.*, 2015; Klein-Flugge *et al.*, 2016; Chong *et al.*, 2017). Circuits involving mPFC and striatum are proposed to be involved in cost-benefit computations (Hillman & Bilkey, 2010; Cowen *et al.*, 2012; Hillman & Bilkey, 2012), and we have shown encoding of net utility by single ACC units (Hashemniayetorshizi *et al.*, 2015). It has been suggested that AMPH modulates the ACC neural dynamics (Lapish *et al.*, 2015), but this has not been explicitly shown in the effort-reward encoding of neurons. Here, we use high-density

neural recordings in freely-moving rats to assess how systemic AMPH affects the encoding of effort and reward by single units and ensembles in ACC. (Sul *et al.*, 2010).

3.2 Material and Methods

3.2.1 Subjects and surgical procedure

In this study, four adult male Fischer Brown Norway (FBN) hybrid aged 6 to 10 months were used. Rats were born and raised on-site, housed individually in a 12h-12h reverse light cycle, and habituated to handling for two weeks prior to surgery. The drive and implantation surgeries were conducted as described previously (Euston & McNaughton, 2006), but is explained briefly here. Surgeries were carried out prior to any training. Animals were deeply anesthetized with isoflurane throughout the procedure (1-1.5 % by volume in oxygen at a flow rate of 1.5 L/min). Each animal was implanted with a "hyperdrive" consisting of 12 independently-movable tetrodes (McNaughton *et al.*, 1983; Wilson & McNaughton, 1993) and 2 reference electrodes. The hyperdrive bundle was centered at 3.00 mm AP and 1.3 mm ML of left mPFC and angled 9.5 degrees toward the midline. A craniotomy was made around the electrode exit site of the drive, and the dive bundle was lowered to the brain surface. The dura was retracted, and the hyperdrive body was secured to the skull with anchor screws embedded in dental acrylic (Lang Dental, Wheeling, US). Following surgery, rats were administered daily injections of 1mg/kg Metacam (analgesic) for 3 days and 10 mg/kg Baytril (antibiotic) for 5 days. Tetrodes were lowered 950 μm from the skull surface after the surgery and then gradually lowered over the next 2-3 weeks to reach the target depth. Food restriction started after the animal recovered from the surgery (7 days) and was monitored daily to ensure the weight was at least 85% of the free-feeding weight for the duration of the experiment. The experiments were performed in dim light during animal's waking phase. All

procedures were performed in accordance with the Canadian Council of Animal Care and the Animal Welfare Committee at the University of Lethbridge.

3.2.2 Experiment

The behavioural apparatus and the method of data collection used in this study were described previously (Mashhoori *et al.*, 2018). Briefly, we used an automated figure-8 maze (Fig. 3.1A), which is a modified version of the classic T-maze frequently

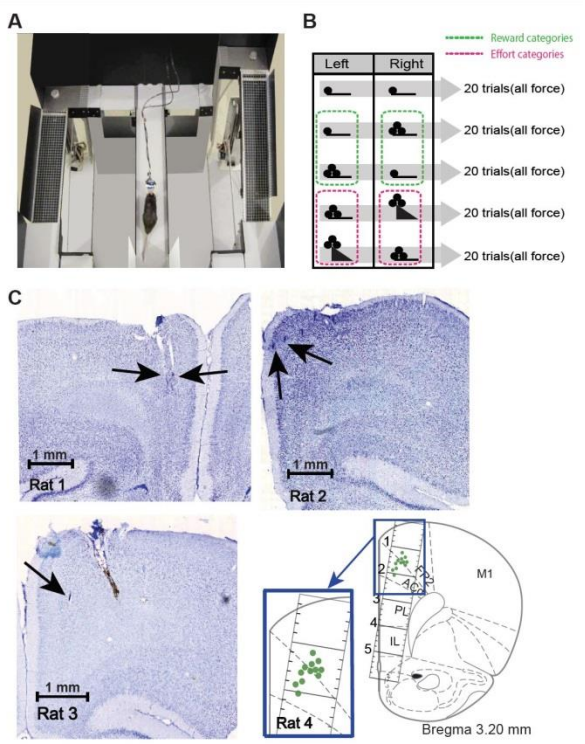

Figure 3.1. Experimental apparatus and task. A, Top view photograph of the figure-8 maze. The rat in the picture is consuming reward at the central feeder. One barrier is elevated for the purpose of illustration. Side feeder wells (white circles) are located on top of the platforms, and wire-mesh ramps. Lead back to the starting feeder. B, Task schematic illustrating one full block of the task. The table shows the order of grouped trials in each effort-reward condition. The height of triangles is equivalent to the barrier height and the number of filled circles represents the amount of food reward at the side feeder. The same trial sequence is repeated two times per session. Dashed lines show the subset of trials used for testing the effort (red) and reward (green) encoding. C, Histological representation of electrode locations. The electrode marks pointed by arrows for three rats are shown on stained coronal brain slices. The estimated electrode location for another rat is also shown on a figure altered from (Paxinos & Watson, 2014).

used in studies of effort-reward decision-making (Salamone *et al.*, 1994a; Walton *et al.*, 2002). The track of the maze was 15 cm wide, and configured into a rectangular pathway measuring 102 x 114cm. The maze contains a central feeder on the T-stem from which a trial was initiated. Two side feeder wells were located on two platforms located in the upper corners of the maze. The platforms could move vertically to require a variable height climb to reach the feeder. Animals descend from the platform by a ramp to return to the starting feeder. The elevation of the platform was 0 cm for the low-effort condition (it was level with the track), and was 23 cm for the high-effort condition. The reward was Ensure beverage (chocolate flavoured), and the volume was 0.03ml as a low-reward, and 0.12 ml as high-reward. The small volume was delivered at the center feeder in all trials. Four gates were located on the entry points of the T-stem and T-arms to prevent animals going backward. These were also used on some trials to force animals to select one target feeder. On other trials, animals were free to choose either target feeder. Over a course of roughly three weeks, animals were trained on the maze with a mixture of force and free trials. Animals were identified as well-trained when they choose high-reward options over the low-reward ones more than 80% of the trials in which the barrier height was the same on both sides. In the following sessions, the animals were required to perform two full blocks of forced trials as shown in Figure 3.1B. Each block of trials consisted of five groups of 20 forced trials where the effort-reward conditions were constant. Each of these trial groups was designed to investigate one of the three decision features (effort, reward, path), which was conducted by manipulating only one of these and keeping the other two constant. However, only effort and reward groups were used in this study. Trials were arranged alternatively left and right throughout the task. One of the recorded rats used a different task design, consisting of either four or six groups of 16 trials (10 alternate-side

force followed by 6 free trials) in which another level of effort was also added (9 inch: medium-effort and 18 inch: high-effort).

All rats were administered with saline or AMPH in a 5 to 10 minutes gap after finishing one block of the experiment and before starting the second one. Three doses of AMPH (0.5, 1, 1.5 mg/kg) were used in this study. Only one dose (or saline) was given per daily session. The order of drug/saline injection across sessions was: saline, AMPH 0.5 mg/kg, AMPH 1 mg/kg, saline, AMPH 1.5 mg/kg. The rat performing the alternate task design received the saline and drugs in different orders: saline, AMPH 0.5 mg/kg, AMPH 1.5 mg/kg, saline, AMPH 1 mg/kg, saline, AMPH 0.5 mg/kg.

3.2.3 Histology

Following completion of the study and 2-5 days before transcardial perfusion, recoding sites were marked by passing 10 uA direct current for 10 second through one electrode of each tetrode. Then, rats received lethal injections of sodium pentobarbital (100 mg/kg i.p.) and were perfused with PBS and 4% paraformaldehyde (PFA). The brains were post-fixed for 24 hours in 4% PFA and then transferred and stored in 30% sucrose and PBS solution with sodium azide (0.02%). After at least 24 hours, the brains were coronally sectioned at 40 μm thickness using a CM3050 S freezing Cryostat (Leica, Germany) and mounted on glass microscope slides, then stained by cresyl violet. Digital images of the prepared brain sections were produced with a Nano-Zoomer slide scanner (Hamamatsu, Japan) and visually inspected to determine the location of marking lesions. As shown in Figure 3.1C, electrode marks were identified for three brains. The 4^{th} brain was destroyed during histology; the estimated locations of the electrodes were determined by the daily logs of electrode depth, and are indicated in the figure.

3.2.4 Analysis

A total of 22 sessions were included in the analysis. The data include at least two sessions with saline treatment and one session of each drug dosage for each animal. In total, 1266 cells were recorded, out of which 1209 were assumed as pyramidal neurons and included in the analysis. Each session is partitioned into two phases, pre- and post-injection, and the analysis contrasts the relative change in the post-injection phase with respect to the pre-injection phase.

3.2.4.1 Preprocessing

Camera-based tracking of LEDs located on preamplifiers attached to the recording drives were used to estimate the position of animals on the maze. We used image registration methods in MATLAB to correct any maze or camera shifts or rotations between the recorded sessions in order to normalize maze position across sessions.

Recorded spikes waveforms were first automatically clustered using KlustaKwik (author: K. D. Harris, Rutgers-Newark) and then manually sorted using MClust (David Redish, University of Minnesota, Minneapolis). Putative individual neurons were then manually classified to pyramidal and interneurons based on half-amplitude and trough-to-peak durations (Bartho *et al.*, 2004). Neural analysis was restricted to only putative pyramidal neurons unless otherwise stated.

3.2.4.2 Behavioural analysis

The effect of AMPH on task performance was investigated using four behavioural measures: the smoothness of rat locomotion trajectories, the median velocity of each trial, the number of off-task distractions, and time spent on reward consumption. Finally, the smoothness was measured by the Hausdorff (fractal) dimension (Hausdorff, 1919; Gneiting *et al.*, 2012), which quantifies the roughness of an object. The object in our case

is the rat's trajectory of movement. Hausdorff dimension increases for the trajectories with movement deviating from the linear axis of the track. We first made images of the running trajectory for each trial of the task. The images were 1 for any pixel the rat traversed, and 0 otherwise. These images were then superimposed to create a grand image of all pre-injection trials and a grand image of all post-injection trials. The dimension was next obtained for each grand image separately. In case the number of trials was not the same for the pre- and post-injection phases, we randomly sub-sampled the trials (100 times), computed Hausdorff dimension for each, and then took an average so as to ensure that this is not an effect of heterogeneous sample sizes among conditions. Secondly, running velocity was calculated by the change in rat's position (in pixels) between every 111 msec time bin (1 pixel/sec is equivalent to 2.36 mm/sec). Thirdly, trials with off-task behaviours were found by manually scoring the video. Off-task behaviours included long pauses on the track, continuous repetition of circling, or backtracking. Since these stereotyped behaviours are known to be induced by AMPH (Randrup *et al.*, 1963; Randrup & Munkvad, 1967), we expected to observe them more often in the post-inj. phase. Finally, the reward consumption time was measured by the time the animal's velocity dropped below a certain threshold (59 and 118 mm/sec on central and side feeders, respectively) at each of the three reward zones. The relative change in all these four measures were found with respect to the pre-inj. phase, and the mean were compared statistically between saline and drug sessions (ANOVA for mean and Kruskal-Wallis for median tests)

3.2.4.3 Modulation of ACC single neuron activity by effort or reward

We first linearized the track by dividing each loop of the figure-8 maze into 36 spatial 2D bins, starting at the central feeder (Fig. 3.3-inset). The firing rate of individual

neurons was computed as the number of spikes in a 0.3 sec time bin centered on the animal's occupancy of each of these spatial bins. In each of these spatio-temporal bins, the change in firing rate was investigated between trials in which only one parameter was different, and the two others were the same (organization of trials is shown in Figure 3.1B). For instance, to compute the effect of effort on neural signaling, we made conditional means from trials in which reward size and reward location were the same. A neuron was considered to be responsive to effort (or reward) if it significantly discriminated low- and high-conditions in at least one of the two effort (or reward) groups shown in Figure 3.1B (t-test or ANOVA, significant at $p<0.05$). The percentage of effort- (or reward-) responsive neurons was found in each of the 36 bins after removing the cells with no discrimination of effort, reward, or feeder location from the population. Grouping the sessions with similar treatments, the average portion of responsive neurons were obtained for pre- and post-inj. phases separately.

The main aim of this study was to investigate the effect of AMPH on the encoding of cost and benefit, including the joint encoding in cells signaling 'utility'. To address this, we first measured the bias strength of individual neurons toward effort and reward, and then compared population variability in pre- and post-inj. phases. Bias was computed by the Pearson correlation coefficient between the spiking rate of each neuron and the two levels of effort or reward in each of the first 16 spatiotemporal bins (from the central feeder to the bin after the side feeder). These coefficients were then collapsed across the sessions with similar treatment. They were then plotted in a symmetric plane with axes of correlation for effort and for reward. In order to identify where the majority of neurons lie in the joint-encoding plane of effort (E) and reward (R), a polar histogram with three equal bins in each quadrant was used. Note that the plane makes four quadrants (Fig.

3.4B-inset), which we numbered as in plane geometry, Q_I (0-90°): E+/R+, Q_{II} (90-180°): E+/R-, Q_{III} (180-270°): E-/R-, Q_{IV} (270-360°): E-/R+. The '+' indicates positive correlation with either reward (R) or effort (E), while the '-' indicates negative correlation. Next, quantified the direction in which neurons are tuned in E-R space. Using principal component analysis (PCA), the direction of the population with maximum variability was then found. We bootstrapped data by randomly selecting the same number of neurons for pre and post-injection phases for 100 times per spatial bin and found principal components (PCs) to be stable, indicating that this is a reliable method. Next, we statistically tested the difference of principal component coefficients of the two phases for each treatment (Kuiper two-sample test which is the circular analogue of the Kolmogrove-Smirnov test). Also, we tested the maximum variance explained by the first principal component and plotted the relative difference between pre- and post-injection of this parameter in each treatment condition. This plot is again based on 100-times bootstrapped data points.

3.2.4.4 Analysis of ACC population activity

Spikes from each trial were time-warped in each of the four important task epochs: reward-consumption at central feeder, interval from the central feeder to the barrier, climbing the barrier, and reward-consumption at target feeder. In other words, spikes were linearly expanded or contracted to have the same time-length epochs as the reference. The reference was obtained by making an average over the time spent in each of these epochs by the animals. This reference is then used for all animals. Having the same time-length enables us to align the epochs and compare them either trial by trial or on average between groups of trials. Moreover, it controls for variability that may occur

due to altered movement speed or pauses in some trials. The aforementioned task epochs were found using animal position as well as velocity with certain thresholds as explained in behavioural section. These epochs were color coded as in Figure 3.6, 3.7, and 3.8-insets.

In order to investigate the modulation of ACC population activity by effort and reward, we used Gaussian Process Factor Analysis (GPFA) (Yu *et al.*, 2009). Using this technique, the latent structure embedded in the recorded population was extracted. In each experimental session, neural trajectories were first made by the firing rate of individual neurons in each 100 msec time-bin in the high-dimension space of all recorded neurons. The first 8 most important latent factors were extracted from the pre-inj. data. These factors in turn represent the directions in the high-dimensional space along which projections show largest variance. We removed trials in which it took the animal more than 20 sec to return back to the central from the side feeder. The post-inj. trajectories were projected on the same low-dimension space in order to compare them with the pre-inj. phase. Finally, we computed the state space occupancy of neural trajectories in 3D space made by the first three GPFA factors, as well as the 2D representation, of each factor versus time. Note that trials with off-task behaviours are removed from the illustrations. The trajectories of the first and second latent factors were averaged over each task epoch and statistically compared between low- and high-effort or reward trials (ANOVA). Following this test, the fraction of sessions with similar treatments that significantly discriminate these features in each epoch were found for pre- and post-inj. phase.

In the next step, we computed the volume of state space trajectories, as well as the velocity of population decorrelation function. For the former, the boundary of the 3D

space was obtained by the convex hull of the trajectory from the central to target feeders; then the enclosed volume was calculated. By finding the change in the volume of each post-inj. low- (or high-) reward trial relative to the mean volume of the same group of trials in pre-inj., we statistically investigated the effect of AMPH on the space size (RM-ANOVA). Next, we aimed to investigate the effect of AMPH on how quickly the population correlation of spatial firing rate decay with distance. Thus, smaller spatial bins were used to improve resolution. Dividing each loop of the figure-8 into 200 2D bins, the place field profile was obtained by averaging the firing rate of each cell in each spatial bin. Next, the population vector cross-correlation matrix was built by calculating the matrix of correlation coefficients of all neurons as a function of spatial lag. By making an average over its diagonal and all parallel elements, the symmetric vector of decorrelation function was calculated afterward (McNaughton, 1998; Battaglia *et al.*, 2004). To quantify how quickly the population decay over the distance, we found the spatial distance where this function first drops below the half-amplitude (i.e. threshold = 0.5). The relative change in the half-amplitude crossing distance of post-inj. phase with respect to the first phase was tested between saline and different drug doses (ANOVA).

3.3 Results

We recorded single neuron activity from the ACC of well-trained rats (n=4) performing a forced-alternation task in which the reward volume and the effort (barrier climbing) to enter either of two reward zones were manipulated independently. We examined the effect of d-amphetamine (AMPH) on behaviour and its neural correlates by means of a pre/post design in which the drug or vehicle was administered (i.p.) roughly halfway through the session.

3.3.1 AMPH increases running speed, but decreases reward consumption and running trajectory smoothness

The running trajectory of the rats, estimated by video tracking of head-mounted LEDs, showed decreasing path smoothness as the dose of AMPH increased (Fig. 3.2A). We quantified this by the Hausdorff fractal dimension, which assesses the distance between spatial data points in order to provide a measure of path roughness (Hausdorff, 1919; Gneiting *et al.*, 2012) (Fig. 3.2B; ANOVA, Main effect of dose on fractal dimension $F_{0.014,0.024} = 3.47$; $p = 0.0378$). The median running velocity increased following administration of higher doses (1.0 & 1.5 mg/kg) as compared to saline (Fig. 3.2C; Kruskal-Wallis, all pairwise comparisons except for saline and dose 0.5 mg/kg are

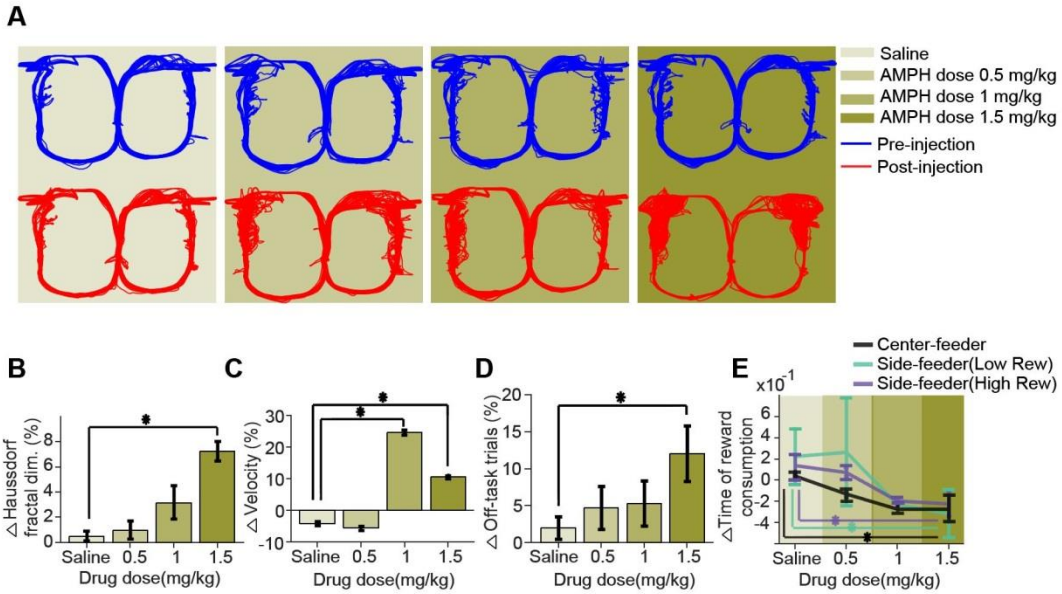

Figure 3.2. Effects of AMPH on task performance. A, A representative example of locomotion trajectories before and after injection of saline or AMPH. B, Mean change of Hausdorff fractal dimension from pre-drug, showing an increase in path variance with increasing AMPH. (1.397 < fractal dimensions < 1.564) C, Median change of velocity. D, Mean change in the relative proportion of trials with off-task behaviour. E, Median change in the time of reward consumption. Errorbars show the Standard Error of the Mean (SEM) or Median (SEMd) (*: statistically significant - but only comparisons with saline are illustrated).

significant, Chi-sq = 109.15; p = 1.675×10^{-23}). The effect appeared to peak at the 1.0mg/kg dose. The amount of off-task behaviour (circling, pausing, backtracking) increased with AMPH dose (Fig 3.2 D; ANOVA, $F_{262.9,469.8} = 3.36$; p = 0.0419), whereas the amount of time at the three reward feeders strongly decreased with increasing dose (Fig 3.2 E; Kruskal-Wallis median test, Chi-sq = 208.9; p = 5.04×10^{-45}; df = 3 for central feeder, Chi-sq = 148.76, p = 4.87×10^{-32}; df = 3 for low-reward feeder, Chi-sq = 136.24; p = 2.45×10^{-29}; df = 3 for high-reward feeder). All pairwise comparisons but one (1 and 1.5 mg/kg) at the central feeder were significantly different (p<0.00001), as well as all except two (1 and 1.5 mg/kg) and (saline and 0.5 mg/kg) at side feeders. Engagement in the task thus appeared to decrease with higher doses of AMPH, while lower concentrations do not appear to affect it significantly. Moreover, rats did not complete a sufficient numbers of trials for analysis following administration of 2.0mg/kg (not shown), further supporting this observation. These data indicate that AMPH increases motoric output, but decreases feeder engagement, consistent with previous reports (Randrup *et al.*, 1963; Randrup & Munkvad, 1967; Blundell *et al.*, 1976; Blundell *et al.*, 1979; Leibowitz *et al.*, 1986; Floresco *et al.*, 2008b).

3.3.2 AMPH does not affect the proportion of effort- or reward-responsive ACC neurons

We next examined if AMPH administration affected single unit encoding of effort or reward in a manner that may explain its behavioural effects. We first linearized the maze by partitioning the track into 36 two-dimensional bins in order to facilitate analysis of neural signaling during specific epochs of the task (Fig. 3.3A inset). The percentage of recorded neurons discriminating the ramp height (i.e. encoding effort) gradually elevates from about 10% at the central feeder to a peak near 35% at the barrier, then sharply drops

after climbing (Fig. 3.3A). The percentage of cells discriminating effort again increases to about 25% while the animal is descending the ramp back to the central feeder. A smaller proportion of ACC neurons encode reward volume (Fig. 3.3B), consistent with previous reports (Cowen *et al.*, 2012). Approximately 10-20 % neurons encode reward on the maze, as shown in Figure 3.3B. The apparent change in the smoothness of these plots at high doses is most likely an artifact of sampling due to fewer neurons recorded in these sessions. The relative difference between the proportion of effort (and reward) selective neurons before and after the injection is not significantly different across drug doses and saline (Fig. 3.3; ANOVA $p>0.05$ in all spatial bins). Thus, AMPH treatment does not appear to affect the proportion of cells encoding effort and reward. It is possible, however, that AMPH affects the firing rate of these cells. Therefore, we next investigated how the encoding of these variables by single units is affected by AMPH.

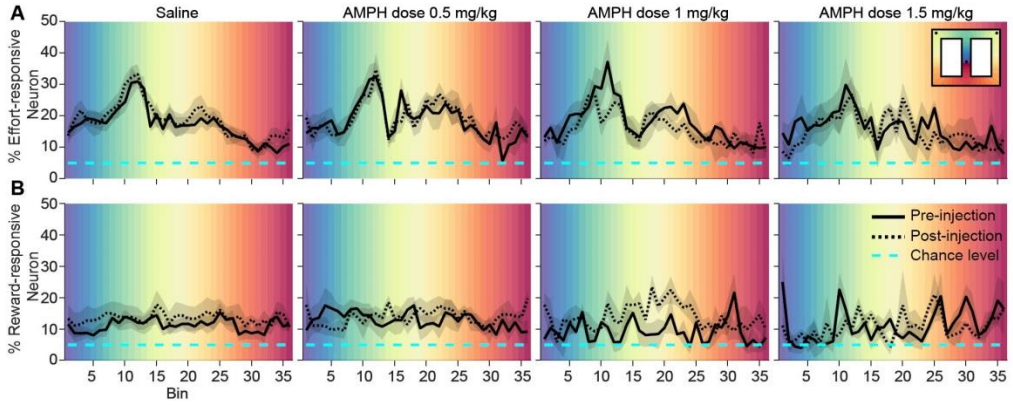

Figure 3.3. The proportion of ACC neurons encoding effort or reward. A, Mean proportion of recorded neurons that discriminate barrier height in pre- and post-injection conditions. B, Mean proportion of recorded neurons responding to the reward in pre- and post-injection conditions. The shaded region surrounding the curves indicates SEM. Background colors correspond to each of the 36 spatial bins of the maze shown in the inset.

3.3.3 AMPH compresses utility encoding in single-unit firing

The analysis of costs and benefits is typically formalized through the concept of utility, which can include many features (Kennerley *et al.*, 2006; Hayden *et al.*, 2009; Kennerley *et al.*, 2009; Kennerley & Wallis, 2009b; Hillman & Bilkey, 2010; Amemori & Graybiel, 2012; Porter *et al.*, 2019). Here, we focus on the joint encoding of effort and reward by ACC neurons. The analysis described above revealed that the largest proportion of effort-encoding neurons modulate activity in the several bins prior to the barrier climb, so we analyzed neural activity during this task phase. We computed the correlation of each neuron's firing rate with effort, and independently computed its correlation with reward in each spatial bin from the central to the side feeders. A scatter plot of these correlation coefficients before drug administration reveals a negative slope of the first principal component of data, indicating that the maximum variance of data is explained by neurons located in the second and fourth quadrants. Figure 3.4A provides an example of these plots, and Figure 3.4B shows the distribution of neurons in effort-reward coding plane at four leading locations of the maze (Fig. 3.4A, B). Neurons of quadrants Q_{II} and Q_{IV} have opposing signaling of utility. Neurons of Q_{IV} tend to generate more action potentials for high-value conditions (i.e. when the effort is low or the reward volume is high), and fire less in less valuable conditions (i.e. high-effort or low-reward). Neurons in Q_{II} exhibit the inverse relationship among firing and value. The slope of principal component becomes steeper with increasing AMPH (Fig. 3.4C), indicating a loss of neural firing correlation with reward. Circular statistical analysis reveals that the post-injection PCs are significantly different than the pre-drug condition for AMPH, whereas saline injection has no effect (Kuiper two-sample test, test statistic $k = 160$;

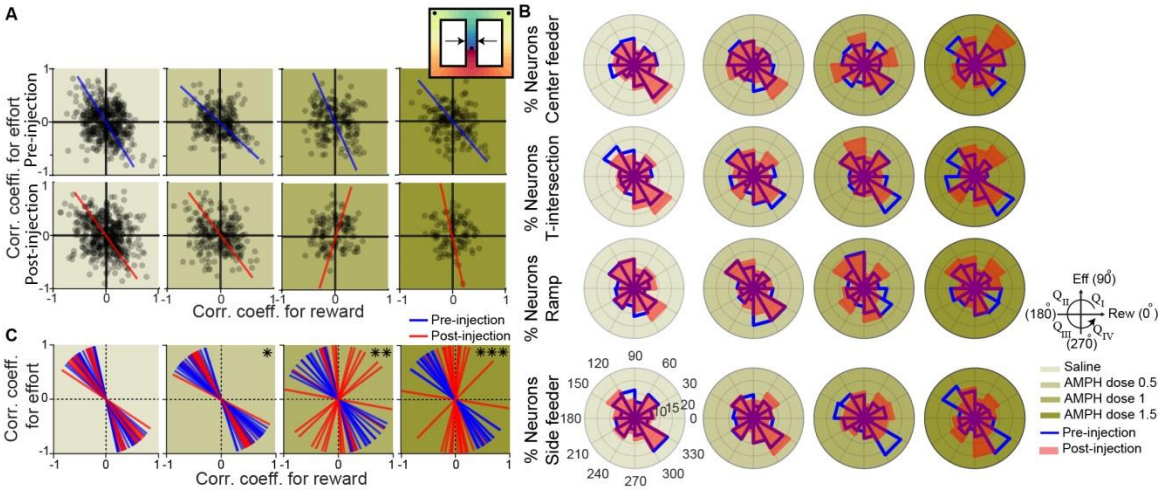

Figure 3.4. Joint encoding of reward and effort by single neurons. A, Representative example of effort-reward joint encoding at spatial bin #3 (right after leaving the central feeder as indicated by arrows in the inset). Each black circle shows the Pearson correlation coefficient between the amount of effort and firing rate of one neuron versus the correlation value for reward and firing rate of the same neuron. The blue and red lines indicate the first principal component coefficient of pre- (blue) and post- (red) injection. B, Distribution of angles of neurons in the effort-reward joint-coding plane. The percentage of neurons in each angle of the effort-reward plane is shown in four spatial bins: Central feeder, T-intersection, Barrier, and Side feeder. Each of four quadrants is divided into three equal-size polar bins where then the proportion of the population is obtained. Angles (0° to 360°) and quadrants (Q_I to Q_{IV}) are defined as in the plane geometry. The order of quadrants and angles are shown at the right subpanel. C, First principal component coefficient of the first 16 spatial bins, from central feeder to a bin after the side feeder. (Kuiper two-sample test for pre- vs post-injection, *: significant at $p<0.05$, **: $p<0.01$, ***: $p<0.001$).

$p= 0.0285$ for 0.5mg/kg, $k = 176$; $p = 0.0068$ for 1mg/kg, $k = 240$; $p < 10^{-3}$ for 1.5mg/kg AMPH). To ensure that this is not an effect of heterogeneous sample sizes among conditions, we ran a bootstrap analysis in which we randomly sub-sampled the data (100 times), so as to ensure the same number of neurons among conditions. The PCs were highly stable, and the results did not change with downsampling. Moreover, the percentage of total variance explained by the first PCs remains above 60% for all conditions which further indicates that the results of the PCA are reliable. The shift in neural tuning indicates that the encoding of reward is being compressed more than is the

encoding of effort. Furthermore, the explained variance by the first principal component is significantly decreased, which further indicates an overall compression in utility signaling at higher doses of AMPH, as opposed to saline or low-dose AMPH (Fig. 3.5; ANOVA, $F_{3,60}$ = 17.3454; p = 2.7154×10^{-8}, bootstrapped 100 times). This effect might be through a breakdown in the encoding of effort and reward or an increased dispersion in effort-reward coding by shifting some neurons to the first quadrant in which the neurons respond positively to both effort and reward. As a result, these neurons fail to encode utility.

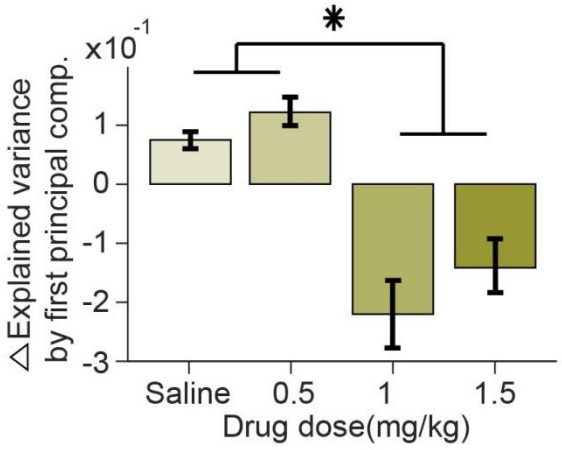

Figure 3.5. Dose-dependent effect of AMPH on the explained variance of effort-reward encoding by first principal component. The mean relative changes of explained variance of data by first principal component of effort-reward jointly encoding neurons, as well as its SEM. (ANOVA, $F_{3,60}$ = 17.5434; p = 2.7154×10^{-8}, *: significant at p<0.005 - pairwise comparisons).

In sum, high concentrations of AMPH induce a significant impairment in the encoding of utility by single units. This effect appears to be caused predominantly by a reduction in reward rather than effort encoding, which is consistent with our behavioural observations that AMPH-treated rats are less interested in consuming the food. These

results therefore suggest that reward is devalued by AMPH to a greater extent than is effort.

3.3.4 AMPH contracts ensemble state space

The analysis above suggests that the encoding of utility by single units is compressed by AMPH. This presumes that the primary carrier of information is firing rate, and does not take into account coordination of firing among neurons. We therefore conducted a state-space analysis of simultaneously recorded neurons to assess how AMPH affects information encoded by temporally-evolving patterns of neural ensembles. We used a method termed Gaussian Process Factor Analysis (GPFA) to reduce the dimensionality of the data. This algorithm is particularly advantageous for producing smooth trajectories in low dimensional space from processes with discrete events, such as action potentials (Yu *et al.*, 2009). Neural trajectories obtained by GPFA are modulated by multiple features of the task, including position, effort, and reward. In the 3D space of the first three GPFA factors, neural trajectories are similar for similar trial types, such as those with the same conditions of effort, reward, and path. Figure 3.6 shows an example of these trajectories for trials of low/high effort in two sessions, and are color coded based on the animal's location in the maze (Fig. 3.6 inset). The return path, from side feeder to the central feeder, is removed from the Figure to facilitate visualization. Trajectories therefore start at the central feeder (blue) and end at one of two side feeders (orange). Trajectories deviate among low- and high-effort trials, particularly while the rat is approaching and climbing the barrier. This shows, as expected from the single neuron analysis (Fig. 3.3), that effort modulates the activity of neural population.

Injection of saline had little effect on the trajectories in the reduced-dimension state-space (Fig. 3.6A), whereas AMPH appears to cause a modest contraction in the space (Fig. 3.6B). To better visualize the encoding of task features and the effects of

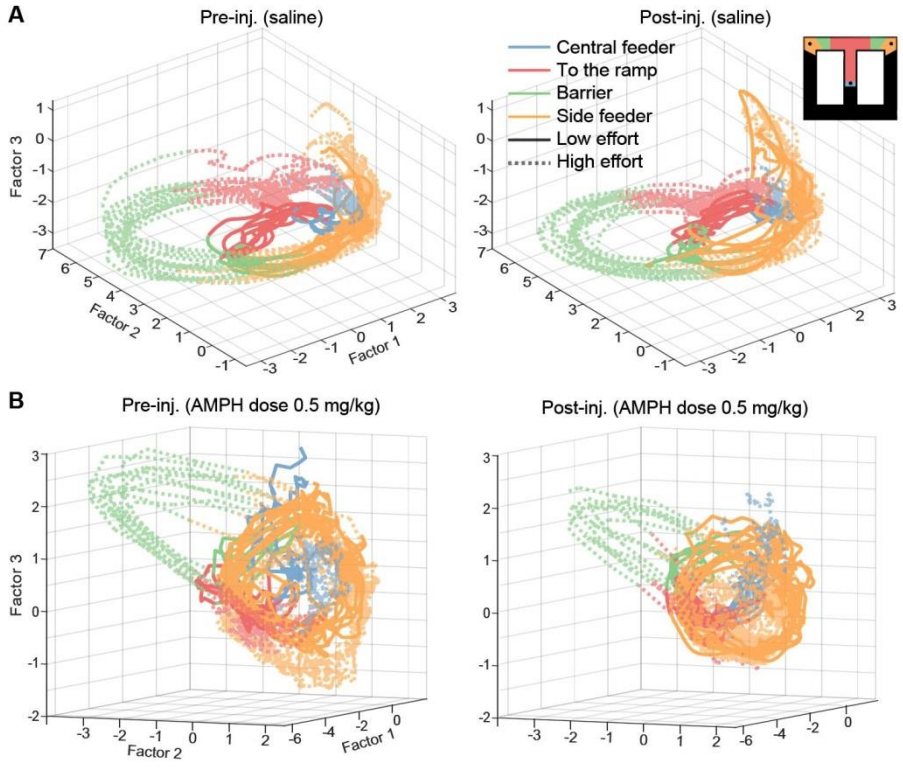

Figure 3.6. Representative example of effort discrimination via neural trajectories. Neural trajectories of effort trials from the central to side feeders are shown in the 3D state-space before and after the injection of saline (panel A) and AMPH dose 0.5 mg/kg (panel B). The trajectories in A and B are for the same animal but two different sessions, and thus, two different state-spaces. The figure inset shows the color-map of the task epochs. High-effort trials are shaded with transparent colors for a better visualization.

AMPH, we next independently plotted the first four GPFA factors on trials with high- vs low-reward (and constant effort; Figure 3.7). The first factor in each session shows a clear discrimination of reward at the feeder zones. In some cases, the second factor does as well. AMPH does not appear to strongly affect the discrimination by the ensemble, but does appear to reduce the overall variance of factors (peak to peak amplitude). Note that the roughness of some trajectories (e.g. Fig. 3.7B and C as compared to A and D) might

be due to the lower number of cells in those sessions or increased population correlation and is not the effect of drug conditions as they appear even before the injections.

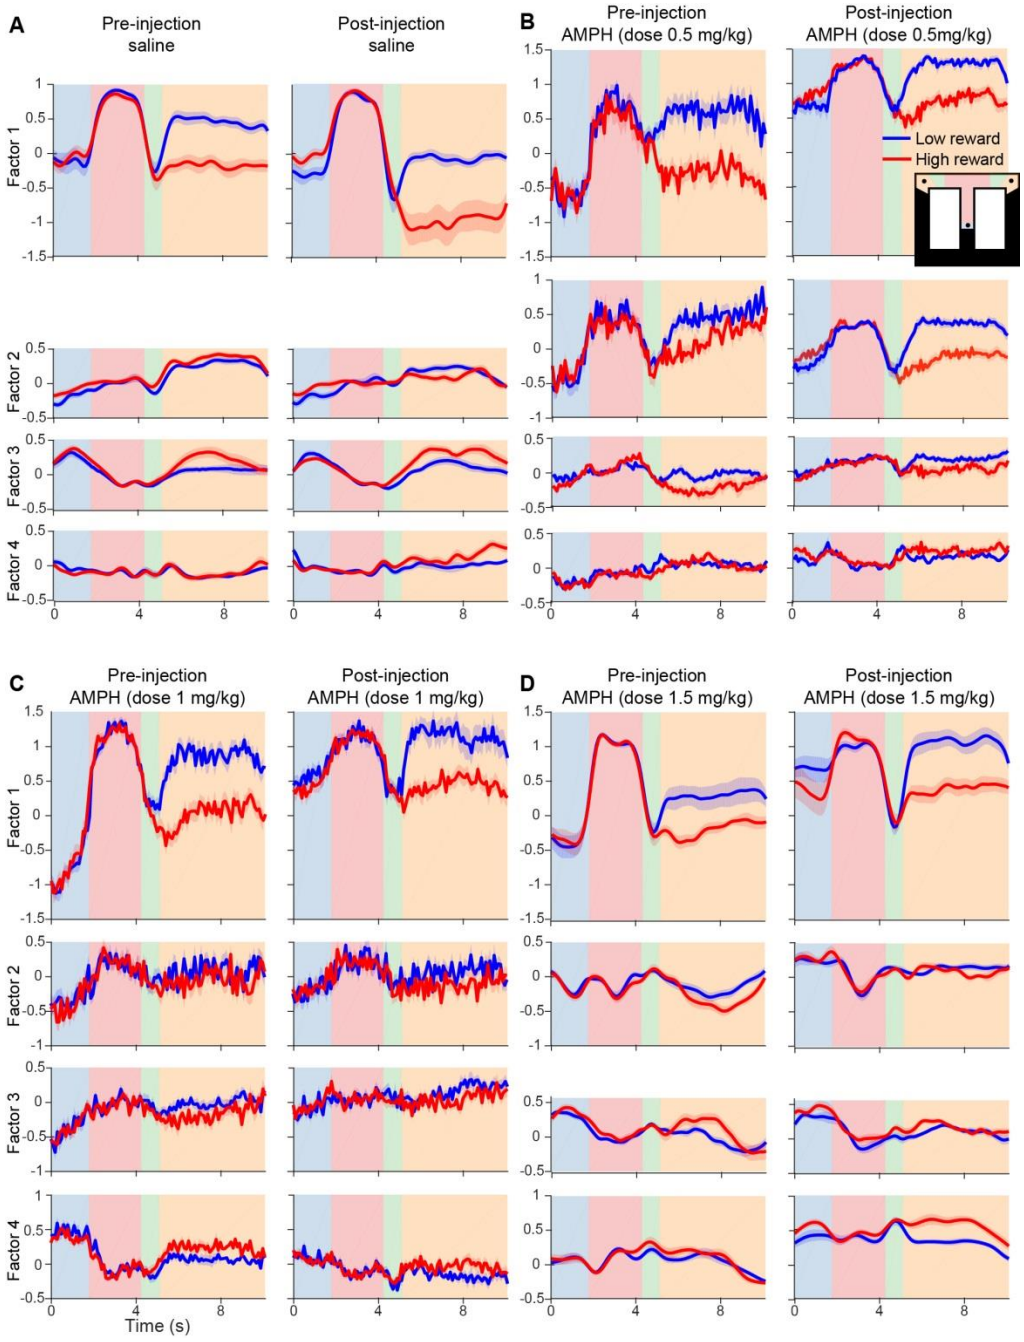

Figure 3.7. Representative example of reward discrimination via population neural dynamics. A, For both pre- and post-injection of saline, the first four latent factors of population activity in low- and high-reward trials are shown as the rat traverses from central to side feeders. The solid lines indicate the mean, and the shaded area represents the SEM. Background colours indicate the task epoch as shown in the inset of panel B.

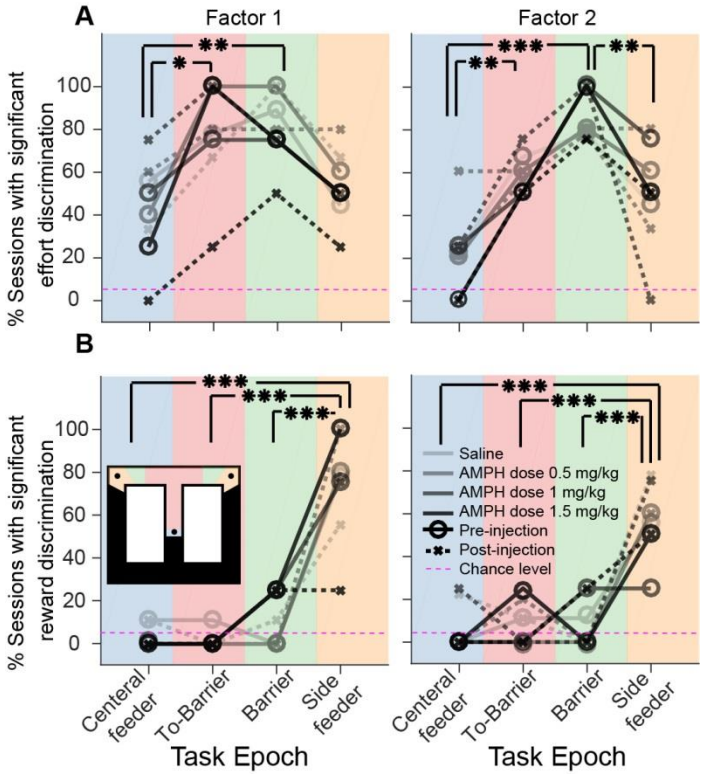

Figure 3.8. Effort and reward encoding by the neural population. Mean proportion of sessions in which the conditional mean of the first or second GPFA factors can discriminate effort (panel A) or reward (panel B) before or after injection. Each drug or saline condition is distinguished by a specific opacity of black color. Solid lines and circular markers indicate the percentage of sessions in which the given parameter (effort or reward) could significantly modulate pre-injection GPFA factors; dotted lines of the same color with cross markers indicates the same variable, but for post-injection phase of the same condition). The background color indicates the four task epochs analyzed as shown in the inset: central feeder, approach to the barrier, climbing the barrier, and side feeder. (ANOVA, *: significant at $p<0.01$, **: $p<0.005$, ***: $p<0.0001$).

The duration of each trial is normalized prior to producing group statistics. B, C, and D, are data from the same rat for AMPH doses of 0.5, 1, and 1.5 mg/kg, respectively.

We next sought to quantitatively test these observations. Using only data from bins between the central feeder to target feeder, we tested if the effort or reward could be discriminated from the average population activity of either the first or second GPFA factor (ANOVA, significant at $p<0.05$). We then compute the fraction of sessions with significant effort or reward encoding. The effort is well discriminated at all epochs, but

peaks significantly near the barrier (Fig. 3.8A) (ANOVA, $F_{3,28} = 6.6$; $p = 0.0016$ for first, and $F_{3,28} = 19.01$; $p = 6.36\times10^{-7}$ for second factor). The reward, on the other hand, is only discriminated well at the target reward feeders (Fig. 3.8B) (ANOVA, $F_{3,28} = 47.55$; $p = 4.09\times10^{-11}$ for first, and $F_{3,28} = 31.41$; $p = 4.22\times10^{-9}$ for second factor). AMPH does not affect the discrimination of effort or reward (Fisher's exact test, $p>0.05$ for pre- vs post-injection of saline or drug in each epoch).

To test if AMPH constricted the trajectories, we computed the change by AMPH of the volume occupied by the GPFA space trajectories in high- and low-reward trials (Fig. 3.9). As compared to saline, the volume is significantly reduced by 0.5 and 1.5 mg/kg of AMPH in low-reward condition (RM-ANOVA, $F_{2.465,93.68} = 12.267$, $p = 4.012\times10^{-6}$) and by 0.5 and 1.0 mg/kg in high-reward condition (RM-ANOVA,

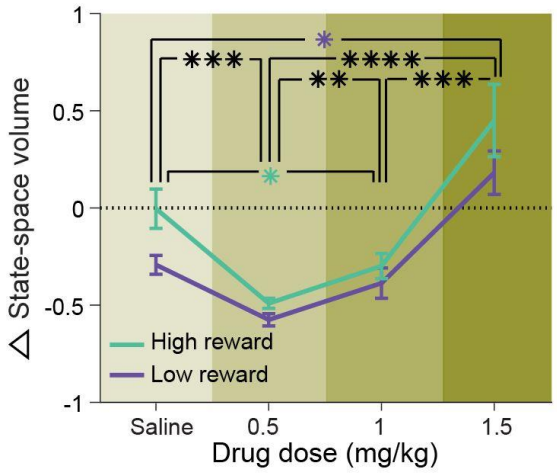

Figure 3.9. Dose-dependent effect of AMPH on the state-space volume. The mean relative volume changes in high- and low-reward trials, as well as its SEM. The volume is calculated as the state-space enclosed by the neural trajectory from central to side feeders in the 3D space formed by the first three latent factors. The volumes of all post-injection trials were then compared with the mean volume of the same group of trials in pre-inj. phase. This comparison occurs in the same space, which is specific to the session, and then the relative change is compared across the sessions as shown here. Negative and positive values in relative change demonstrate contraction and expansion in the state-space volume, respectively. (RM-ANOVA, *: significant at $p<0.05$, **: $p<0.01$, ***: $p<0.005$ and ****: $p<10^{-6}$, for both low- and high-reward conditions if in black.)

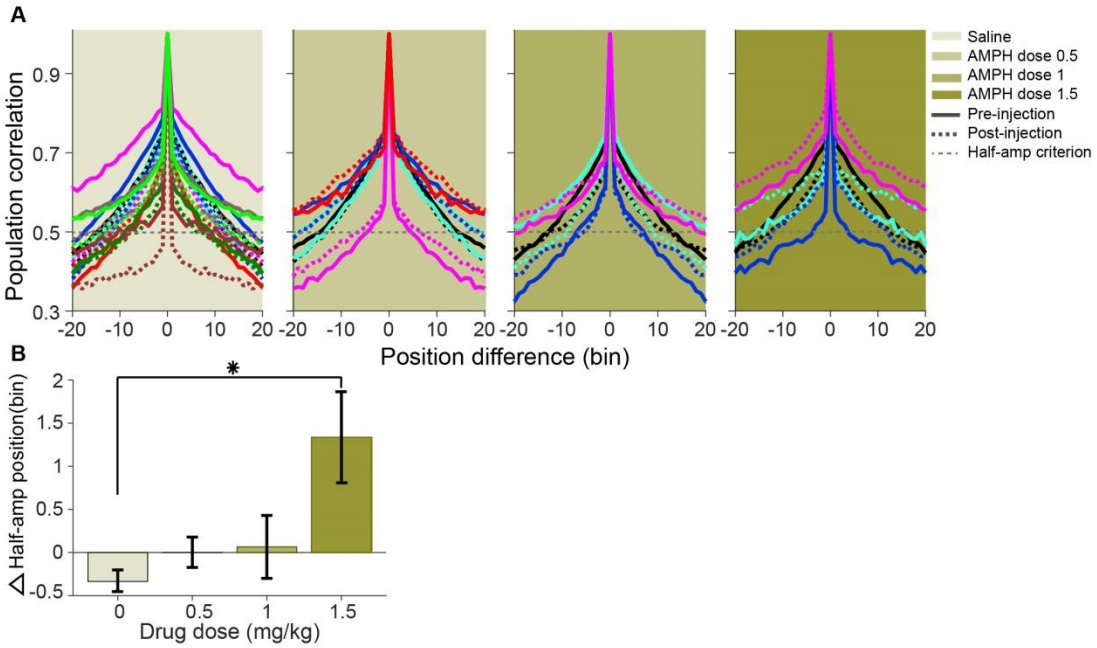

Figure 3.10. Effect of AMPH on population vector decorrelation. A, Population vector correlation of all task epochs of pre- and post-injection for each session (distinguished by color; such that each solid line indicates pre-injection population correlation and dotted line of the same color indicates the post-injection correlation function of the same session). Grey dotted horizontal lines represent the half-amplitude criteria. The crossing of this line indicates how fast the population decorrelates. B, Mean relative change in the half-amplitude crossing point following injections. (Errorbars: SEM and *: significant at $p<0.005$, ANOVA).

$F_{1.729,67.437} = 10.073$, $p = 2.98\times10^{-4}$). It is possible that the increased variance in the rat's movement path and/or increased propensity for off-task behaviour in high-dose may obscure the direct pharmacological effects in the ACC by increasing superfluous variance.

We lastly analyzed the dynamics of the ensemble by computing the state-space decorrelation, which describes how the correlation of vectors of binned activity change in time and space (McNaughton, 1998; Battaglia *et al.*, 2004). The state vector correlations are shifted upward, indicating that the patterns are more correlated across the entire maze (Fig. 3.10A). Only the 1.5 mg/kg dose showed a statistically significant difference from

the saline condition (ANOVA, $F_{3,18} = 6.97$; $p = 0.0026$). For this high dose, the half-amplitude of the decorrelation is higher, meaning that the rate of change of correlation is lower (Fig. 3.10B).

In sum, the first and second latent GPFA factors indicate that the neural ensemble is strongly modulated prior to the effort (approaching to the barrier), while exerting effort (climbing the barrier), and while at the feeder. Although AMPH did not affect the mean difference of the first two GPFA factors based on effort or reward, intermediate doses did decrease the state space volume.

3.4 Discussion

In this study, we investigated the effect of AMPH on ACC neural activity in rats performing a task with variable effort and reward. Our results demonstrate that the proportion of neurons that significantly discriminate low and high levels of effort or reward are not influenced strongly by AMPH, whereas, the joint encoding of effort and reward is significantly affected by this drug. In a four-quadrant joint encoding plane, we investigated how the firing rate of individual neurons is correlated with effort and reward levels. In the pre-injection phase of the experiment, the recorded population is biased toward the second and fourth quadrants (Q_{II}: E+/R-, and Q_{IV}: E-/R+), with most units positively correlated with the reward amount and negatively with the effort level (Q_{IV}). The abundance of neurons in these two quadrants supports the conjecture that the ACC primarily encodes utility of the choice (either positively (Q_{IV}) or negatively (Q_{II})) (Kennerley *et al.*, 2009; Kennerley & Wallis, 2009b; Blanchard & Hayden, 2014). Although the population is slightly biased toward the effort and reward axes as the animal reaches the barrier and feeders, respectively, the first principal component coefficient remains tuned to these two quadrants the entire duration from the central feeder to side

feeders, indicating that this distribution of utility encoding explains the maximum variance of the population throughout several phases of the task. AMPH injection increased dispersion of effort-reward coding, and sometimes biased the distribution to the first quadrant (Q_I: E+/R+), wherein the neurons fire more as the level of either effort or reward increases. This quadrant is inconsistent with the concept of utility encoding because neural activity increases with both benefit and cost. The first principal component is shifted toward this quadrant as a result of AMPH administration. Furthermore, higher doses of AMPH significantly decrease the data variance explained by the first component (Fig. 3.5), indicating a contraction in the encoding of utility in terms of both reward and effort. Nonetheless, the rotation of the first principle reveals that the ability of AMPH to decrease the utility encoding at the level of single units in the ACC is likely more related to a compression of correlation among firing and reward, rather than disturbance of effort encoding.

The population analysis revealed a clear event-specific discrimination of effort and reward levels. Consistent with the single-unit analysis, the effort discrimination was maximal in the proximity of the barrier, whereas reward discrimination was largest at the side feeders. Although the discrimination for effort and reward was not attenuated by AMPH, the amplitude of signaling within the reduced space was attenuated by intermediate doses of AMPH, leading to a reduced volume of neural trajectories in the latent space. Moreover, the spatial decorrelation of neural patterns was slower under increasing doses of AMPH. These data reveal that AMPH causes a contraction of signaling and an increase in the self-similarity over a longer distance. This is consistent with the analysis of single units, which indicate that the contraction may be more

pronounced in the reward domain. A loss of reward-related variance without increase in variance related to other task features will manifest as increased self-similarity.

3.4.1 AMPH and the perception of choice value

How might the above-described effects of AMPH affect the function of ACC? The systemic injection of AMPH used in the present experiment will have widespread effects in the brain. Catecholamines have long been postulated to participate in action generation and reward learning (Berridge, 2007). They are critical components of neural systems that generate effortful responses for goal-directed behaviours (Salamone & Correa, 2002). Part of the effect may be purely locomotor. AMPH increases motoric activity (Randrup & Munkvad, 1967; Groves & Rebec, 1976; Robinson & Berridge, 1993; Wilkinson *et al.*, 1993; Sams-Dodd, 1998), which we observed here. Moreover, AMPH decreases feeding (Foltin, 2001; Shoblock *et al.*, 2003; Cannon *et al.*, 2004; Wellman *et al.*, 2009) and rats' sensitivity to reward omissions (Wong *et al.*, 2017). Consistent with these data, we found that animals spent less time at the feeders, and were more likely to skip the reward, as the dose of AMPH increased. Thus, the behavioural effects appear dichotomous; animals ran faster on the maze, but were less interested in the reward. This pattern of behavioural effects may reflect discounting of both the physical effort and reward value. This may cause a change in the perception of a choice value (positive utility). It remains unclear why an animal would engage in the operant response for a highly discounted outcome. Indeed, animals did not engage in the task at higher doses, and instead engaged in off-task behaviours such as grooming, sniffing, and exploring the edges of the track. It may be that at high doses, the outcome has devalued to a point in which the animal no longer engages in the task. Similarly, a previous study has reported an inverted U-shape dose-dependent function of AMPH on rats' performance in

a working memory task (Shoblock *et al.*, 2003). Specifically, low concentrations of AMPH significantly improved working memory and increased locomotor activity, whereas higher doses didn't improve task performance. Moreover, this study also found that rats consumed less reward after AMPH treatment. We therefore speculate that at low doses, the outcome still has value while the effort has discounted substantially, leading to increased task engagement but relative disinterest in the reward.

We cannot infer where in the brain AMPH may be affecting decision-related processing. We can, however, use neural activity in ACC as a window into the network processing. We found neural correlates of effort and reward at appropriate task phases, consistent with past reports (Sul *et al.*, 2010; Cowen *et al.*, 2012; Blanchard & Hayden, 2014). The proportion of single units encoding reward peaked in the proximity of the barrier climb, whereas the proportion of reward-related units peaked at the side feeders. Previous studies also observed ACC activity in anticipation of effort and reward (Sul *et al.*, 2010; Cowen *et al.*, 2012; Hashemniayetorshizi *et al.*, 2015). Our results show that most neurons primarily encode utility in the portion of the maze from the center feeder to side feeders. It is possible that the animals anticipate the effort and reward of the upcoming trial because of the task design; the animal was directed to alternate between left and right options, and the effort and rewards were static over blocks of n = 20 trials. It is therefore possible that neural activity tracks expectations or keeps the animal engaged in the task, rather than functioning to generate optimal cost-benefit decisions. The present task lacks choice of feeder by the animal, so the 'choice' is between performing and not performing the task. AMPH may either change the value of the choice to engage in the task, or could disrupt the mechanisms by which the animal maintains engagement. This function is more akin to attention and vigilance, which are both increased by AMPH in a

variety of task settings and species (Sostek *et al.*, 1980; Ridley *et al.*, 1982; Koelega, 1993; Solanto, 1998; Grilly, 2000; Silber *et al.*, 2006; Sagvolden, 2011).

3.4.2 AMPH contracts populations neural dynamics

The population ACC activity appears to undergo stereotyped changes during task performance such that each task phase is encoded by a specific encoding state and/or cognitive function (Lapish *et al.*, 2008; Balaguer-Ballester *et al.*, 2011). Computational models have suggested that dopamine stabilizes the dynamics in the cortex by both direct neuromodulatory effects on cortical neurons (Gruber *et al.*, 2006; Durstewitz & Seamans, 2008), as well as 'locking' the input to the cortex from the basal ganglia (Gruber *et al.*, 2006). This provides a mechanistic explanation by which AMPH could produce the observed effects on the present task. Moderately increased dopamine and other catecholamines by low-dose AMPH stabilizes task-related cortical and striatal patterns, thus increasing task engagement and velocity. Note that this mechanism requires no notion of reward value or utility; rather, it is a property of changing the neural dynamics in the network. Furthermore, AMPH has been reported to increase the separation between distinct ensemble activity patterns during distinct task states at a moderate dose (1 mg/kg), whereas it reduces the distance between such distinct neural activity states at high doses (3.3 mg/kg) (Lapish *et al.*, 2015). We also found dose-dependent effects on population neural activity in similar range of AMPH doses. At low doses (0.5-1.0 mg/kg), we found a slight contraction of state space. This became an expansion of state-space at high doses (1.5 mg/kg). Moreover, the decorrelation time increased at high doses. These observations are consistent with that of Lapish and colleagues. Although the state space contracts under low doses, the variance from trial to trial also appears to reduce, which can increase pattern separation from samples at different points of the trajectory (different

task phases). At high doses, the state space expands, but so does the variance across trials. The increased variance and increased correlation across location both increase the similarity (overlap) of pattern samples at different task phases. In other words, the decrease in uninformative variance ('noise') with low-dose AMPH is proportionally greater than the decrease in state space contraction of ensemble patterns across phases of the task, thus leading to a greater separation of patterns (in the sense of discrimination analyses, such as D prime in signal detection theory) of distinct task phases. This breaks down under higher doses of AMPH such that the encoding becomes plagued with variance that can knock the state of ACC encoding off of its learned task-related trajectory, thus breaking engagement in the task in lieu of off-task behaviours. We further speculate that some brain states are more stable than others, and that this is likely a function of experience (amount of training) and genetics/neurodevelopment (i.e. dominant innate behaviours such as grooming). It may be that at high doses, only very robust behaviours can withstand the additional variance, such that the behavioural output consists of a repertoire of innate behaviours and 'over-learned' operant responses. Other less stable responses may begin to be initiated, but are disrupted by the high variance. Indeed, the behavioural output under high doses of AMPH and other psychostimulants are typically highly stereotyped actions reflecting parts of grooming, sniffing, and licking with repetitive head movement (Randrup & Munkvad, 1967; Fog, 1970; Schiorring, 1971). The data presented in this study are entirely consistent with the many reports of AMPH effects on behaviour, and the few papers on the modulation of neural activity in the medial prefrontal cortex. Although our data are purely correlational, they suggest novel linkages between the effects of AMPH on neural activity and behaviour. We cannot presently discern if it is the effects of AMPH to compress utility representations, or to

introduce variance into the state dynamics, that provide a more compelling explanation for the behavioural effects.

4 CHAPTER4: General Discussion

In this thesis, I sought to elucidate the role of ACC in goal-directed behaviours, particularly how it encodes variables related to cost-benefit decisions. A second aim was to understand how the information encoded by ACC is affected by amphetamine, which has long been known to affect decisions involving effort and rewards. I simultaneously recorded the spiking activity of multiple single units of ACC of rats performing an effort-reward decision-making task with and without AMPH on board. Based on spiking characteristics such as waveform and energy, I separated single units from recorded activity. After that, I analyzed firing rates of putative single neurons in fine spatiotemporal resolution to address the following questions: Do they discriminate between the levels of effort, reward, and different paths (and if yes, how)? Does this discrimination remain stable or shift in preparation for the goal and approach to it? Are effort and reward encoded individually by distinct neurons, or are they jointly encoded (i.e., do ACC cells encode net utility)? And how is the encoding influenced by AMPH at low, medium, and high doses? In addition to addressing these questions at the level of single units, I investigated neural trajectories of the recorded population in a space created by the most influential latent factors extracted from the population activity. In this low-dimension space, I considered the ensemble activity in response to effort and reward. Then, I examined the effect of AMPH on the population state-space volume and how far the population remains self-similar on the maze using the population decorrelation function.

4.1 Effort, Reward, and utility encoding by ACC

My data indicates that a considerable fraction of ACC individual neurons discriminates between different levels of the barrier and reward. Consistent with previous

reports (Cowen *et al.*, 2012), this activity started to increase prior to the corresponding event, peaked during the event, and then dropped quickly. It is thus not surprising that this effort and reward selectivity is robustly reflected in the ACC population activity, such that we could decode them with a high accuracy (chapter 2). Furthermore, ACC neural trajectories significantly discriminate between low and high levels of effort when approaching and climbing the barrier, as well as low and high levels of reward at side feeders (chapter 3). I found that the firing rate of the majority of task-related neurons correlates with both effort and reward. This finding is in accordance with the multiplexing feature of frontal cortical neurons observed by others, meaning that cells respond to more than one variable in a variety of tasks (Jung *et al.*, 1998; Lapish *et al.*, 2008; Hayden *et al.*, 2009; Kennerley & Wallis, 2009b; Hillman & Bilkey, 2010; Balaguer-Ballester *et al.*, 2011; Kennerley *et al.*, 2011; Wallis & Kennerley, 2011; Horst & Laubach, 2012). I found that the correlation between the activity of a population of ACC neurons and reward is opposite to their relationship with effort, consistent with the concept of utility. These neurons fell into two groups. Neurons of the first group fired more as the height of the required barrier decreased or the level of forthcoming reward increased (E-/R+); thus, they are purported to encode positive utility of the choice. Conversely, the activity of neurons in the second group increased when more effort was required or less reward was presented (E+/R-); hence, they possibly encode negative utility. The first and second groups may thus signal impetus and disincentive or convey motivational-related information to the corresponding brain regions. This interpretation is in line with previous studies that suggest that the ACC is part of a brain network initializing a motivational process so that the subject is willing to exert more effort to receive reward (Walton *et al.*, 2006; Phillips *et al.*, 2007; Holroyd & Yeung, 2011; Spielberg *et al.*, 2012). Supporting

this idea, I found that more neurons encode positive rather than negative utility in most spatial bins from the central feeder to the side feeders. Also, the activity of these neurons discriminates between the levels of effort and reward almost exclusively in this part of the maze. This region of the task is of a particular importance as the animal has already initiated his movement toward the choice after leaving the central feeder and has to keep focus on the pursuit of his goal when travelling to the side feeder. Hence, the encoded information in this segment of the maze is likely providing the animal foresight into the upcoming outcomes. In addition, I compared the superimposed firing rate of positive and negative utility neurons over the entire maze. For the choices of a higher value (high-reward or low-effort), I found that the positive-utility signal is on average higher than the negative-utility signal in most spatial bins. Conversely, the overall activity of negative-utility cells approaches the magnitude, or even surpasses, the positive-utility signal when leaving reward zones in less valuable trials (high-effort or low-reward). This effect remains for about 5 to 10 spatial bins and then reverses. These dynamic changes in the balance of these two signals and their dependency on the levels of effort and reward suggest that these two groups of neurons are in competition (although another possibility is that they encode opposite parameter values rather than being in direct competition). These findings, together with the post-decisional appearance of effort and reward encoding (chapter 2) and the lack of significant differentiation between free and forced trials by ACC neurons, suggest that ACC may not be a site of comparison between available options prior to a decision. Instead, they suggest that the ACC is involved in the ongoing process of approaching the choice zone and maintaining concentration on what to expect, as suggested by previous studies (Kennerley *et al.*, 2006; Sul *et al.*, 2010; Cai & Padoa-Schioppa, 2012; Blanchard & Hayden, 2014; Holec *et al.*, 2014).

These results are also consistent with previously proposed roles of the ACC in controlling adaptive decisions (Walton *et al.*, 2007). Prior reports demonstrate that ACC neurons adapt their firing activity to the changes in the task context or in behavioural requirements. For instance, as the rule of a task is changed in the absence of explicit cues related to the rule, a transition is observed in the rat's ACC ensemble activity from an old state, corresponding to the familiar rule, to a new state, associated with the novel rule (Durstewitz *et al.*, 2010). Primate studies illustrate that ACC neurons progressively compute the value of a depleting source of food and signal a decision of when to leave it for a new food source (Hayden *et al.*, 2011b). They also fire more strongly to reward as its expectancy increases based on a pre-defined schedule (Shidara & Richmond, 2002). In summary, the ACC is suggested to be involved in situations where no single certain response is cued as correct or when there is a change in outcomes of the same action. My task design consisted of varying choice conditions in each experimental session. Also, it contained trials where the rat was free to choose between low-reward/low-effort and high-reward/high-effort alternatives, none of which was specified as correct choice by the experiment. The results presented in the second chapter demonstrate adaptive neural responses to effort and reward conditions, even when conditions change or no certain correct choice exists, thus support the previously suggested role of the ACC in adaptive decision making (Walton *et al.*, 2007).

4.2 Representation of the task map by ACC

In a previous study, we found that ACC ensemble activity encodes the rats' current position on the task (Mashhoori *et al.*, 2018). Using a deep artificial neural network, we decoded this spatial information. We found that about 40 randomly chosen neurons out of the recorded population are sufficient to provide a high spatiotemporal

resolution estimate of the rat's position on the maze. Spatial encoding in the ACC appeared to be different than in the hippocampus, as the neural spatial firing fields are much broader than that of place cells (Wilson & McNaughton, 1993). Also, ACC neurons do not fire only in a particular location of the environment, but instead specify several small regions of the maze or a large continuous part of it. Using several training and test sets of trials for the input of our decoder, we found this spatial information quite stable within each experimental session, regardless of the choice variables, such as presented amount of reward or the height of the barrier. We thus speculated that "the ACC may form a topographically-organized representational space, based on real space or action/events at particular locations on the maze, which can serve as a scaffold for the encoding of behaviourally-relevant events" (Mashhoori *et al.*, 2018). This speculation, although similar to some previous proposals indicating that the ACC links events, context and behaviours (Lapish *et al.*, 2008; Euston *et al.*, 2012), introduce the spatial representation in the ACC as a contextual map to which relevant events and behaviours are pinned.

The results of the present study provide supportive evidence for this proposal. The data presented in the second chapter indicate that the rat's future running path (left/right choice) modulates ACC activity even stronger than effort and reward. The number of neurons which discriminate these two paths are more than the other two factors (effort and reward) in almost all locations of the maze. Also, path was decoded from ACC population activity at more than 80% accuracy, except in a couple of spatial bins at reward zones. Considering that rats' left and right trajectories diverge as soon as the animal leaves the central feeder, even when he is still on the T-stem, path-selectivity might, in fact, be a reflection of rat's current location instead of upcoming choice. Either

way, this spatial selectivity is independent of effort and reward conditions and remains fairly stable over the entire maze. Effort and reward-encoding, on the other hand, seem to be a function of proximity to the effort or reward location. The fact that cells are responsive to changes in these parameters at the same locations, shows that these are additional dimensions of the rats experience encoded within ACC. The task space encoded by the ACC thus encompasses both place and behaviourally relevant parameters like effort levels and reward amounts. Taken together, the ACC might express a unique pattern of activity for each phase of the task (Cowen *et al.,* 2012) and thus provide a measure of the animal's progression through the task space, as suggested by prior studies (Lapish *et al.*, 2008).

Nevertheless, my experiment was designed to investigate the encoding of choice variables and not to specify if the spatial representation in the ACC exercises control over these types of putative cognitive signals. Indeed, task phase is confounded with location in my experiment. The multiplexing of a broad range of information by ACC neurons that has been reported by many researchers (Jung *et al.*, 1998; Lapish *et al.*, 2008; Horst & Laubach, 2012), makes it challenging to distinguish between the encoding of task phase and location in tasks where the same actions occur in the same location (Jung *et al.*, 1998; Euston & McNaughton, 2006; Cowen & McNaughton, 2007). A recent study confirms the encoding of spatial information in the ACC (Lindsay *et al.*, 2018). This study utilized a small, simple behavioural chamber to investigate the modulation of ACC ensemble by different movements and locations. The chamber, although originally designed for a Pavlovian conditioning task, was used with no conditional cues or outcomes to prevent their interference with the parameters that were under investigation. These parameters included rat's limb movements, instantaneous velocity, and spatial position. The study

demonstrated that the activity of ACC individual neurons is weakly influenced by these parameters, yet the information is widely distributed across the population and was decodable from the ensemble activity. For instance, about 27% of ACC neurons were selective to the animal's location but the R^2 value between the model factor and the firing rate of these neurons rarely passed 0.1 (GLM: generalized linear model). Although the percentage of these spatial-selective neurons is comparable to the percentage of effort- and reward-selective neurons I found in the ACC (Fig. 2.6), the relationships they reported between the spiking activity of the neurons and the location (R^2) is much weaker than what I observed for effort and reward (Fig 2.10). However, the limited effect of the location might be caused by relatively small size of the enclosure they used ($30 \times 25 \times 60$ cm). Taken together, future studies are needed to compare the spatial encoding with effort and reward computation or task-phase information in the ACC and clarify which one is more prominent.

4.3 Prospective reward processing by ACC

I found that a small group of ACC neurons encoded the amount of reward expected to be received at a remote location of the maze far before that. These neurons, two examples of which are provided in chapter 2, discriminate between low- and high-reward values at side-feeders specifically when the rat receives a fixed, low amount of reinforcement at the central feeder. This observation, although novel, is consistent with our previous report of post-reinforcement ACC excursions to remote not-taken reward location (Mashhoori *et al.*, 2018). Similar to what I found here, that ACC activity was observed specifically when disfavoured low reward was experienced by the rat. However, it has also some dissimilarity with those signals. In that study, we showed that the decoded position of the rat, obtained from ACC population activity, sometimes deviates

from his actual position. This deviation emerged almost exclusively at side-feeder sites, and formed a mental excursion originating in the animal's current position, connecting several consecutive spatial points to the alternate reward location, and returning to the animal's position. Such excursions did not appear at any other locations on the maze including the central feeder. Hence, we concluded that they are not signaling the general quality of reward or future actions (Mashhoori *et al.*, 2018). Since this phenomenon was found to occur more frequently after less-preferred reinforcement, we suggested it is involved in regret or the evaluation of a better outcome that could have happened. On the other hand, prospective reward processing I observed in the present study occurs at the central feeder and reflects future reward levels, not their locations. I thus speculated that they are involved in the evaluation of future outcomes. Thus, the difference between where this activity is observed in our previous study and here might be due to the difference in what they encode. Nevertheless, the presence prospective activity at reward zones related to either alternate reward location or forthcoming reward amount is consistent with the suggested role of the ACC in default mode network (discussed more in section 2.4) (Buckner *et al.*, 2008; Lu *et al.*, 2012).

4.4 AMPH-induced changes in the encoding of task variables in the ACC

AMPH significantly affected rat's behaviour in the effort-reward alternation task in a dose-dependent manner. Animals treated with high doses of AMPH were less engaged in the task (more off-task distractions and less reward consumption), as opposed to intermediate doses and saline. These results are consistent with prior reports of AMPH-induced stereotype behaviours (Randrup *et al.*, 1963; Randrup & Munkvad, 1967) as well as increased motoric activity (Wilkinson *et al.*, 1993; Sams-Dodd, 1998) and satiety (Foltin, 2001; Cannon *et al.*, 2004; Wellman *et al.*, 2009). Therefore, they imply potential

changes in underlying effort and reward neural encoding. I tested ACC activity in response to effort and reward before and after AMPH/Saline injections and found that the proportion of effort- and reward-selective neurons remained fairly similar for pre- and post-injection of either solution. Multivariate analysis also showed that ACC population activity significantly differentiates low- and high-levels of effort (while approaching and climbing the barrier) and of reward (while pausing at side reward zones) both with and without the drug on board. However, jointly responding to effort and reward by individual ACC neurons was found to be significantly influenced by AMPH. As explained before, the spiking activity of a large fraction of ACC neurons normally are tuned to positive and negative utilities. As shown in Figure 3.4, when each cell is represented on scatterplot according to their effort and reward selectivity, cells tend to be more numerous in the 2^{nd} and 4^{th} quadrants. Further, the maximum variance explained by the first principal component of data is aligned with the Q_{II}-Q_{IV} diagonal. However, AMPH was found to shift this neural tuning to a less negative slope thus indicating decreasing reward-selectivity. Also, AMPH of higher doses resulted in a general reduction in the explained variance by the first principal component (Figure 3.5). These effects are consistent with behavioural results and may explain why the rats were less interested in reward and showed more off-task distractions after receiving high doses of AMPH. Indeed, it is possible that these behaviours are induced by the lack of precise utility information normally represented by ACC.

I then sought to investigate whether the similar compression I observed in the population of utility neurons is reflected in the population neural activity. As a measure of population variance, I measured the volume of 3D state-space enclosed by neural trajectories. I found a significant contraction and expansion in the volume, respectively,

due to moderate and high doses of AMPH. However, the general pattern of neural trajectories from central to side feeder remained fairly similar between pre- and post-injection of this drug. In other words, the change in the volume was due to the change in the peak-to-peak activity, while the sequence of task-events related states remained intact. Although still sensitive to both effort and reward, this compression observed in both effort-reward joint encoding and population state-space volume suggests that AMPH restricts the ACC activity to a smaller range. On the other hand, AMPH of low-to-moderate doses decreases variations in trial-by-trial neural trajectories. This reduction appears to be greater than the reduction in the state-space volume, and thus improving the separation between task-related states, consistent with previous studies (Lapish *et al.*, 2015). High level of AMPH, however, induces opposite effects, expanding the state-space volume and increasing the variance of neural trajectories.

Considering that AMPH is known to increase extracellular DA concentration, these findings provide supportive evidence for the proposed role of DA in estimating the cost-benefit utility of a decision (Phillips *et al.*, 2007). This neuromodulator has been postulated to mediate instrumental behaviours by processing different aspects of reward (Phillips *et al.*, 2007) and motivating the animal to invest more physical effort in order to overcome behavioural constraints (Salamone *et al.*, 2006). For instance, midbrain DA neurons have been shown to encode the magnitude and the expectancy of future reward (Tobler *et al.*, 2005). On the other hand, accumbens DA depletion appears to shift the animals' choice preference from high-effort/high-reward to low-effort/low-reward choices (Cousins *et al.*, 1996; Mai *et al.*, 2012). A similar effect has also been reported after intra-ACC D1DR receptor blockade, but not after D2 antagonist (Schweimer & Hauber, 2006). DA is postulated to be involved not only in effort-reward computations,

but also in stabilizing cortical representations regardless of the task. Computational models suggest that prefrontal cortex DA concentration plays a biphasic role in modulating neural trajectories and stabilizing attractor dynamics (Durstewitz & Seamans, 2002; 2008). Also, AMPH itself, similar to many other psychostimulants that affect catecholamine transmission, induces dose-dependent effects on both behaviour and neural activity (Robinson & Berridge, 1993; Floresco *et al.*, 2008b; Lapish *et al.*, 2015). Our results also demonstrate dose-dependent effects of AMPH on ACC population states. As explained previously, low concentrations of this drug decrease the variance of neural trajectories across trials, suggesting more stabilized task-related states. Also, AMPH shrinks the state-space domain through which neural trajectories traverse, thus contracting ACC ensemble activity. However, these effects are reversed by higher concentrations. Nevertheless, an underlying mechanism whereby AMPH changes neural dynamics via DA transmission remain to be tested in future studies.

DA is released in brain structures regulating goal-directed behaviours, such as ACC, NAc, and VTA. Recent electrophysiological studies in rodents have revealed reciprocal communication between ACC and VTA in effort-reward decision-making tasks (Elston & Bilkey, 2017; Elston *et al.*, 2018; Elston *et al.*, 2019). They showed that the ACC signals to the VTA in the anticipation of the task-reversal as well as immediately after switching to the new condition (Elston *et al.*, 2019). This signal also increased when the barrier was absent, in an experiment where its presence was probabilistic (Elston & Bilkey, 2017). On the other hand, VTA-to-ACC signaling was observed in response to errors (i.e., when choosing the less preferred option after a task reversal) (Elston *et al.*, 2019). VTA-to-ACC signaling was also reported to increase when the rats' choices were incorrect in a rule switching task involving a predictable switch

(Elston *et al.*, 2018). These findings led the authors to speculate that bidirectional communications between these two areas provides the VTA-involved circuit of motivation access to cortically-derived information about the task while the VTA provides the ACC feedback about when to modify the representation of task space.

It is noteworthy that to better understand the role of DA in cost-benefit types of computations and underlying neural mechanisms, direct manipulations of these neurons along with simultaneous electrophysiological recording from target regions should be employed in future experiments. Also, it is important to note that the changes observed in neural activity may not be only through DA, as it is not the only neurotransmitter whose concentrations are affected by AMPH. Indeed, AMPH has a complex multifaceted effect on brain function (Fleckenstein *et al.*, 2007) and more research is required to clarify its specific effects on cognitive networks.

4.5 General Conclusion

The findings presented in this thesis are consistent with previous reports. In addition, I present here a novel finding regarding the ACC encoding of the utility of an available effort-reward option. Moreover, my data shows that this information is significantly disturbed by administration of AMPH in a dose-dependent manner. Utility appears to be encoded via two groups of ACC neurons, one of which encodes effort-discounted reward and the other one encodes the inverse. The balance of activation between these two neural populations changes through different phases of the task based on the value of the effort-reward condition. Moreover, some ACC neurons prospectively encode properties of remote feeders when the animal pauses at another reward location. AMPH treated rats spend less time at reward zones, which is concordant with the observed reduction in the ACC reward encoding. Furthermore, ACC neural trajectories

shrink at lower doses of AMPH, but the trial-by-trial variation in neural trajectories also decreases. These effects, which are reversed by high concentration of AMPH, are consistent with dose-dependent effects of AMPH on the rat's performance in the task. Together, these findings help to better understand ACC's function in a cost-benefit computation, and how this function is affected by AMPH. These results suggest some mechanistic explanations of previously reported changes in choice behaviour via ACC lesion or AMPH administration. Future experiments should determine how task-related information in the ACC might be used by other structures, as well as underlying mechanisms through which AMPH affects ACC neural encoding.

REFERENCES

Amemori, K. & Graybiel, A.M. (2012) Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nat Neurosci*, **15**, 776-785.

Amiez, C., Joseph, J.P. & Procyk, E. (2005) Anterior cingulate error-related activity is modulated by predicted reward. *European Journal of Neuroscience*, **21**, 3447-3452.

Apps, M.A.J. & Ramnani, N. (2014) The Anterior Cingulate Gyrus Signals the Net Value of Others' Rewards. *Journal of Neuroscience*, **34**, 6190-6200.

Ash, E.S., Heal, D.J. & Clare Stanford, S. (2014) Contrasting changes in extracellular dopamine and glutamate along the rostrocaudal axis of the anterior cingulate cortex of the rat following an acute d-amphetamine or dopamine challenge. *Neuropharmacology*, **87**, 180-187.

Assadi, S.M., Yucel, M. & Pantelis, C. (2009) Dopamine modulates neural networks involved in effort-based decision-making. *Neurosci Biobehav Rev*, **33**, 383-393.

Avelar, A.J., Juliano, S.A. & Garris, P.A. (2013) Amphetamine augments vesicular dopamine release in the dorsal and ventral striatum through different mechanisms. *J Neurochem*, **125**, 373-385.

Balaguer-Ballester, E., Lapish, C.C., Seamans, J.K. & Durstewitz, D. (2011) Attracting dynamics of frontal cortex ensembles during memory-guided decision-making. *PLoS Comput Biol*, **7**, e1002057.

Balleine, B.W. & Dickinson, A. (1998) Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, **37**, 407-419.

Bardgett, M.E., Depenbrock, M., Downs, N., Points, M. & Green, L. (2009) Dopamine modulates effort-based decision making in rats. *Behav Neurosci*, **123**, 242-251.

Bartho, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K.D. & Buzsaki, G. (2004) Characterization of neocortical principal cells and Interneurons by network interactions and extracellular features. *Journal of Neurophysiology*, **92**, 600-608.

Battaglia, F.P., Sutherland, G.R. & McNaughton, B.L. (2004) Local sensory cues and place cell directionality: Additional evidence of prospective coding in the hippocampus. *Journal of Neuroscience*, **24**, 4541-4550.

Berger, B., Gaspar, P. & Verney, C. (1991) Dopaminergic innervation of the cerebral cortex: unexpected differences between rodents and primates. *Trends Neurosci*, **14**, 21-27.

Berridge, K.C. (2007) The debate over dopamine's role in reward: the case for incentive salience. *Psychopharmacology*, **191**, 391-431.

Besson, M.J., Cheramy, A., Feltz, P. & Glowinski, J. (1969a) Release of newly synthesized dopamine from dopamine-containing terminals in the striatum of the rat. *Proc Natl Acad Sci U S A*, **62**, 741-748.

Besson, M.J., Cheramy, A. & Glowinski, J. (1969b) Effects of amphetamine and desmethylimipramine on amines synthesis and release in central catecholamine-containing neurons. *Eur J Pharmacol*, **7**, 111-114.

Blanchard, T.C. & Hayden, B.Y. (2014) Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *J Neurosci*, **34**, 646-655.

Blanchard, T.C., Strait, C.E. & Hayden, B.Y. (2015) Ramping ensemble activity in dorsal anterior cingulate neurons during persistent commitment to a decision. *J Neurophysiol*, **114**, 2439-2449.

Blundell, J.E., Latham, C.J. & Leshem, M.B. (1976) Differences between Anorexic Actions of Amphetamine and Fenfluramine-Possible Effects on Hunger and Satiety. *J Pharm Pharmacol*, **28**, 471-477.

Blundell, J.E., Latham, C.J., Moniz, E., Mcarthur, R.A. & Rogers, P.J. (1979) Structural-Analysis of the Actions of Amphetamine and Fenfluramine on Food-Intake and Feeding-Behavior in Animals and in Man. *Curr Med Res Opin*, **6**, 34-54.

Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S. & Cohen, J.D. (2001) Conflict monitoring and cognitive control. *Psychol Rev*, **108**, 624-652.

Braun, S. & Hauber, W. (2011) The dorsomedial striatum mediates flexible choice behavior in spatial tasks. *Behav Brain Res*, **220**, 288-293.

Buckner, R.L., Andrews-Hanna, J.R. & Schacter, D.L. (2008) The brain's default network - Anatomy, function, and relevance to disease. *Ann Ny Acad Sci*, **1124**, 1-38.

Buckner, R.L. & Carroll, D.C. (2007) Self-projection and the brain. *Trends Cogn Sci*, **11**, 49-57.

Burton, B.G., Hok, V., Save, E. & Poucet, B. (2009) Lesion of the ventral and intermediate hippocampus abolishes anticipatory activity in the medial prefrontal cortex of the rat. *Behavioural Brain Research*, **199**, 222-234.

Cai, X. & Padoa-Schioppa, C. (2012) Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *J Neurosci*, **32**, 3791-3808.

Cannon, C.M., Abdallah, L., Tecott, L.H., During, M.J. & Palmiter, R.D. (2004) Dysregulation of striatal dopamine signaling by amphetamine inhibits feeding by hungry mice. *Neuron*, **44**, 509-520.

Carr, M.F., Jadhav, S.P. & Frank, L.M. (2011) Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, **14**, 147-153.

Chiueh, C.C. & Moore, K.E. (1973) Release of Endogenously Synthesized Catechols from Caudate-Nucleus by Stimulation of Nigro-Striatal Pathway and by Administration of D-Amphetamine. *Brain Research*, **50**, 221-225.

Chong, T.T.J., Apps, M., Giehl, K., Sillence, A., Grima, L.L. & Husain, M. (2017) Neurocomputational mechanisms underlying subjective valuation of effort costs. *Plos Biol*, **15**.

Chudasama, Y., Daniels, T.E., Gorrin, D.P., Rhodes, S.E.V., Rudebeck, P.H. & Murray, E.A. (2013) The Role of the Anterior Cingulate Cortex in Choices based on Reward Value and Reward Contingency. *Cerebral Cortex*, **23**, 2884-2898.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I. & Shenoy, K.V. (2010) Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron*, **68**, 387-400.

Cohen, J. (1988) *Statistical power analysis for the behavioral sciences*. L. Erlbaum Associates, Hillsdale, N.J.

Cousins, M.S., Atherton, A., Turner, L. & Salamone, J.D. (1996) Nucleus accumbens dopamine depletions alter relative response allocation in a T-maze cost/benefit task. *Behav Brain Res*, **74**, 189-197.

Covey, D.P., Juliano, S.A. & Garris, P.A. (2013) Amphetamine elicits opposing actions on readily releasable and reserve pools for dopamine. *PLoS One*, **8**, e60763.

Cowen, S.L., Davis, G.A. & Nitz, D.A. (2012) Anterior cingulate neurons in the rat map anticipated effort and reward to their associated action sequences. *J Neurophysiol*, **107**, 2393-2407.

Cowen, S.L. & McNaughton, B.L. (2007) Selective delay activity in the medial prefrontal cortex of the rat: contribution of sensorimotor information and contingency. *J Neurophysiol*, **98**, 303-316.

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**, 4531-4541.

Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B. & Dolan, R.J. (2006) Cortical substrates for exploratory decisions in humans. *Nature*, **441**, 876-879.

Deiber, M.P., Passingham, R.E., Colebatch, J.G., Friston, K.J., Nixon, P.D. & Frackowiak, R.S.J. (1991) Cortical Areas and the Selection of Movement - a Study with Positron Emission Tomography. *Exp Brain Res*, **84**, 393-402.

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

Di Chiara, G. & Imperato, A. (1988) Drugs abused by humans preferentially increase synaptic dopamine concentrations in the mesolimbic system of freely moving rats. *Proc Natl Acad Sci U S A*, **85**, 5274-5278.

Doya, K. (2008) Modulators of decision making. *Nat Neurosci*, **11**, 410-416.

Duncan, J. (2001) An adaptive coding model of neural function in prefrontal cortex. *Nat Rev Neurosci*, **2**, 820-829.

Durstewitz, D. & Seamans, J.K. (2002) The computational role of dopamine D1 receptors in working memory. *Neural Networks*, **15**, 561-572.

Durstewitz, D. & Seamans, J.K. (2008) The Dual-State Theory of Prefrontal Cortex Dopamine Function with Relevance to Catechol-O-Methyltransferase Genotypes and Schizophrenia. *Biol Psychiat*, **64**, 739-749.

Durstewitz, D., Vittoz, N.M., Floresco, S.B. & Seamans, J.K. (2010) Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, **66**, 438-448.

Easton, N., Steward, C., Marshall, F., Fone, K. & Marsden, C. (2007) Effects of amphetamine isomers, methylphenidate and atomoxetine on synaptosomal and synaptic vesicle accumulation and release of dopamine and noradrenaline in vitro in the rat brain. *Neuropharmacology*, **52**, 405-414.

Elston, T.W. & Bilkey, D.K. (2017) Anterior Cingulate Cortex Modulation of the Ventral Tegmental Area in an Effort Task. *Cell Rep*, **19**, 2220-2230.

Elston, T.W., Croy, E. & Bilkey, D.K. (2019) Communication between the Anterior Cingulate Cortex and Ventral Tegmental Area during a Cost-Benefit Reversal Task. *Cell Rep*, **26**, 2353-+.

Elston, T.W., Kalhan, S. & Bilkey, D.K. (2018) Conflict and adaptation signals in the anterior cingulate cortex and ventral tegmental area. *Sci Rep-Uk*, **8**.

Euston, D.R., Gruber, A.J. & McNaughton, B.L. (2012) The role of medial prefrontal cortex in memory and decision making. *Neuron*, **76**, 1057-1070.

Euston, D.R. & McNaughton, B.L. (2006) Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. *J Neurosci*, **26**, 13143-13155.

Evans, S.M. & Foltin, R.W. (1997) The effects of d-amphetamine on the reinforcing effects of food and fluid using a novel procedure combining self-administration and location preference. *Behav Pharmacol*, **8**, 429-441.

Fishburn, P., C., (1970) *Utility Theory for Decision Making*. JOHN WILEY & SONS, INC., New York, London, Sydney, Toronto.

Fleckenstein, A.E., Volz, T.J., Riddle, E.L., Gibb, J.W. & Hanson, G.R. (2007) New insights into the mechanism of action of amphetamines. *Annu Rev Pharmacol*, **47**, 681-698.

Floresco, S.B. & Ghods-Sharifi, S. (2007) Amygdala-prefrontal cortical circuitry regulates effort-based decision making. *Cerebral Cortex*, **17**, 251-260.

Floresco, S.B., St Onge, J.R., Ghods-Sharifi, S. & Winstanley, C.A. (2008a) Cortico-limbic-striatal circuits subserving different forms of cost-benefit decision making. *Cogn Affect Behav Neurosci*, **8**, 375-389.

Floresco, S.B., Tse, M.T. & Ghods-Sharifi, S. (2008b) Dopaminergic and glutamatergic regulation of effort- and delay-based decision making. *Neuropsychopharmacology*, **33**, 1966-1979.

Fog, R. (1970) Behavioural effects in rats of morphine and amphetamine and of a combination of the two drugs. *Psychopharmacologia*, **16**, 305-312.

Foltin, R.W. (2001) Effects of amphetamine, dexfenfluramine, diazepam, and other pharmacological and dietary manipulations on food "seeking" and "taking" behavior in non-human primates. *Psychopharmacology*, **158**, 28-38.

Foltin, R.W. & Evans, S.M. (1999) The effects of d-amphetamine on intake of food and a sweet fluid containing cocaine. *Pharmacol Biochem Be*, **62**, 457-464.

Foltin, R.W. & Fischman, M.W. (1988) Food-Intake in Baboons - Effects of D-Amphetamine and Fenfluramine. *Pharmacol Biochem Be*, **31**, 585-592.

Foster, D.J. & Wilson, M.A. (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, **440**, 680-683.

Fried, I., Mukamel, R. & Kreiman, G. (2011) Internally Generated Preactivation of Single Neurons in Human Medial Frontal Cortex Predicts Volition. *Neuron*, **69**, 548-562.

Frith, C.D., Friston, K., Liddle, P.F. & Frackowiak, R.S.J. (1991) Willed Action and the Prefrontal Cortex in Man - a Study with Pet. *P Roy Soc B-Biol Sci*, **244**, 241-246.

Fujisawa, S., Amarasingham, A., Harrison, M.T. & Buzsaki, G. (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci*, **11**, 823-833.

Gneiting, T., Sevcikova, H. & Percival, D.B. (2012) Estimators of Fractal Dimension: Assessing the Roughness of Time Series and Spatial Data. *Stat Sci*, **27**, 247-277.

Gothard, K.M., Skaggs, W.E., Moore, K.M. & McNaughton, B.L. (1996) Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *Journal of Neuroscience*, **16**, 823-835.

Grilly, D.M. (2000) A verification of psychostimulant-induced improvement in sustained attention in rats: Effects of d-amphetamine, nicotine, and pemoline. *Exp Clin Psychopharm*, **8**, 14-21.

Groves, P.M. & Rebec, G.V. (1976) Biochemistry and behavior: some central actions of amphetamine and antipsychotic drugs. *Annu Rev Psychol*, **27**, 91-127.

Gruber, A.J., Calhoon, G.G., Shusterman, I., Schoenbaum, G., Roesch, M.R. & O'Donnell, P. (2010) More is less: a disinhibited prefrontal cortex impairs cognitive flexibility. *J Neurosci*, **30**, 17102-17110.

Gruber, A.J., Dayan, P., Gutkin, B.S. & Solla, S.A. (2006) Dopamine modulation in the basal ganglia locks the gate to working memory. *J Comput Neurosci*, **20**, 153-166.

Gruber, A.J., Hussain, R.J. & O'Donnell, P. (2009) The nucleus accumbens: a switchboard for goal-directed behaviors. *PLoS One*, **4**, e5062.

Gruber, A.J. & McDonald, R.J. (2012) Context, emotion, and the strategic pursuit of goals: interactions among multiple brain systems controlling motivated behavior. *Front Behav Neurosci*, **6**, 50.

Haber, S.N., Kim, K.S., Mailly, P. & Calzavara, R. (2006) Reward-related cortical inputs define a large striatal region in primates that interface with associative cortical connections, providing a substrate for incentive-based learning. *Journal of Neuroscience*, **26**, 8368-8376.

Hashemniayetorshizi, S., Sessford, D., Gruber, A.J. & Euston, D.R. (2015) Modulation of cellular activity by effort and reward in Anterior Cingulate Cortex. *Program No. 633.20. 2015 Neuroscience Meeting Planner. Washington, DC: Society for Neuroscience, 2015. Online.* .

Hauber, W. & Sommer, S. (2009) Prefrontostriatal circuitry regulates effort-related decision making. *Cereb Cortex*, **19**, 2240-2247.

Hausdorff, F. (1919) Dimension und äußeres Maß. *Mathematische Annalen*, **79**, 157-179.

Hayden, B.Y., Heilbronner, S.R., Pearson, J.M. & Platt, M.L. (2011a) Surprise Signals in Anterior Cingulate Cortex: Neuronal Encoding of Unsigned Reward Prediction Errors Driving Adjustment in Behavior. *Journal of Neuroscience*, **31**, 4178-4187.

Hayden, B.Y., Pearson, J.M. & Platt, M.L. (2009) Fictive Reward Signals in the Anterior Cingulate Cortex. *Science*, **324**, 948-950.

Hayden, B.Y., Pearson, J.M. & Platt, M.L. (2011b) Neuronal basis of sequential foraging decisions in a patchy environment. *Nature Neuroscience*, **14**, 933-U165.

Heilbronner, S.R. & Hayden, B.Y. (2016) Dorsal Anterior Cingulate Cortex: A Bottom-Up View. *Annu Rev Neurosci*, **39**, 149-170.

Heilbronner, S.R., Rodriguez-Romaguera, J., Quirk, G.J., Groenewegen, H.J. & Haber, S.N. (2016) Circuit-Based Corticostriatal Homologies Between Rat and Primate. *Biol Psychiat*, **80**, 509-521.

Hernandez, L., Lee, F. & Hoebel, B.G. (1987) Simultaneous microdialysis and amphetamine infusion in the nucleus accumbens and striatum of freely moving rats: increase in extracellular dopamine and serotonin. *Brain Res Bull*, **19**, 623-628.

Hillman, K.L. & Bilkey, D.K. (2010) Neurons in the rat anterior cingulate cortex dynamically encode cost-benefit in a spatial decision-making task. *J Neurosci*, **30**, 7705-7713.

Hillman, K.L. & Bilkey, D.K. (2012) Neural encoding of competitive effort in the anterior cingulate cortex. *Nat Neurosci*, **15**, 1290-1297.

Holec, V., Pirot, H.L. & Euston, D.R. (2014) Not all effort is equal: the role of the anterior cingulate cortex in different forms of effort-reward decisions. *Front Behav Neurosci*, **8**, 12.

Holroyd, C.B. & Coles, M.G.H. (2002) The neural basis. of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev*, **109**, 679-709.

Holroyd, C.B. & Yeung, N. (2011) An Integrative Theory of Anterior Cingulate Cortex Function: Option Selection in Hierarchical Reinforcement Learning. *Neural Basis of Motivational and Cognitive Control*, 333-349.

Horst, N.K. & Laubach, M. (2012) Working with memory: evidence for a role for the medial prefrontal cortex in performance monitoring during spatial delayed alternation. *Journal of Neurophysiology*, **108**, 3276-3288.

Hosking, J.G., Cocker, P.J. & Winstanley, C.A. (2014) Dissociable Contributions of Anterior Cingulate Cortex and Basolateral Amygdala on a Rodent Cost/Benefit Decision-Making Task of Cognitive Effort. *Neuropsychopharmacology*, **39**, 1558-1567.

Hosking, J.G., Floresco, S.B. & Winstanley, C.A. (2015) Dopamine antagonism decreases willingness to expend physical, but not cognitive, effort: a comparison of two rodent cost/benefit decision-making tasks. *Neuropsychopharmacology*, **40**, 1005-1015.

Hosokawa, T., Kennerley, S.W., Sloan, J. & Wallis, J.D. (2013) Single-Neuron Mechanisms Underlying Cost-Benefit Analysis in Frontal Cortex. *Journal of Neuroscience*, **33**, 17385-17397.

Hyman, J.M., Ma, L.Y., Balaguer-Ballester, E., Durstewitz, D. & Seamans, J.K. (2012) Contextual encoding by ensembles of medial prefrontal cortex neurons. *P Natl Acad Sci USA*, **109**, 5086-5091.

Isomura, Y., Ito, Y., Akazawa, T., Nambu, A. & Takada, M. (2003) Neural coding of "attention for action" and "response selection" in primate anterior cingulate cortex. *J Neurosci*, **23**, 8002-8012.

Jones, B.F., Groenewegen, H.J. & Witter, M.P. (2005) Intrinsic connections of the cingulate cortex in the rat suggest the existence of multiple functionally segregated networks (vol 133, pg 193, 2005). *Neuroscience*, **135**, 1011-1011.

Jueptner, M., Stephan, K.M., Frith, C.D., Brooks, D.J., Frackowiak, R.S.J. & Passingham, R.E. (1997) Anatomy of motor learning .1. Frontal cortex and attention to action. *Journal of Neurophysiology*, **77**, 1313-1324.

Jung, M.W., Qin, Y.L., McNaughton, B.L. & Barnes, C.A. (1998) Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebral Cortex*, **8**, 437-450.

Kennerley, S.W., Behrens, T.E.J. & Wallis, J.D. (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature Neuroscience*, **14**, 1581-U1119.

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

Kennerley, S.W. & Wallis, J.D. (2009a) Encoding of Reward and Space During a Working Memory Task in the Orbitofrontal Cortex and Anterior Cingulate Sulcus. *Journal of Neurophysiology*, **102**, 3352-3364.

Kennerley, S.W. & Wallis, J.D. (2009b) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci*, **29**, 2061-2073.

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

Kerns, J.G., Cohen, J.D., MacDonald, A.W., Cho, R.Y., Stenger, V.A. & Carter, C.S. (2004) Anterior Cingulate conflict monitoring and adjustments in control. *Science*, **303**, 1023-1026.

Kim, H., Sul, J.H., Huh, N., Lee, D. & Jung, M.W. (2009) Role of striatum in updating values of chosen actions. *J Neurosci*, **29**, 14701-14712.

Klein-Flugge, M.C., Kennerley, S.W., Friston, K. & Sestmann, S. (2016) Neural Signatures of Value Comparison in Human Cingulate Cortex during Decisions Requiring an Effort-Reward Trade-off. *Journal of Neuroscience*, **36**, 10002-10015.

Koelega, H.S. (1993) Stimulant-Drugs and Vigilance Performance - a Review. *Psychopharmacology*, **111**, 1-16.

Kolling, N., Wittmann, M.K., Behrens, T.E.J., Boorman, E.D., Mars, R.B. & Rushworth, M.F.S. (2016) Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience*, **19**, 1280-1285.

Kurniawan, I.T., Guitart-Masip, M., Dayan, P. & Dolan, R.J. (2013) Effort and Valuation in the Brain: The Effects of Anticipation and Execution. *Journal of Neuroscience*, **33**, 6160-6169.

Kurniawan, I.T., Guitart-Masip, M. & Dolan, R.J. (2011) Dopamine and effort-based decision making. *Front Neurosci*, **5**, 81.

Lapish, C.C., Balaguer-Ballester, E., Seamans, J.K., Phillips, A.G. & Durstewitz, D. (2015) Amphetamine Exerts Dose-Dependent Changes in Prefrontal Cortex Attractor Dynamics during Working Memory. *J Neurosci*, **35**, 10172-10187.

Lapish, C.C., Durstewitz, D., Chandler, L.J. & Seamans, J.K. (2008) Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proc Natl Acad Sci U S A*, **105**, 11963-11968.

Leibowitz, S.F., Shorposner, G., Maclow, C. & Grinker, J.A. (1986) Amphetamine - Effects on Meal Patterns and Macronutrient Selection. *Brain Research Bulletin*, **17**, 681-689.

Leyton, M., Boileau, I., Benkelfat, C., Diksic, M., Baker, G. & Dagher, A. (2002) Amphetamine-induced increases in extracellular dopamine, drug wanting, and novelty seeking: A PET/[C-11]raclopride study in healthy men. *Neuropsychopharmacology*, **27**, 1027-1035.

Lindsay, A.J., Caracheo, B.F., Grewal, J.J.S., Leibovitz, D. & Seamans, J.K. (2018) How Much Does Movement and Location Encoding Impact Prefrontal Cortex Activity? An Algorithmic Decoding Approach in Freely Moving Rats. *Eneuro*, **5**.

Lu, H.B., Zou, Q.H., Gu, H., Raichle, M.E., Stein, E.A. & Yang, Y.H. (2012) Rat brains also have a default mode network. *P Natl Acad Sci USA*, **109**, 3979-3984.

Luk, C.H. & Wallis, J.D. (2009) Dynamic Encoding of Responses and Outcomes by Neurons in Medial Prefrontal Cortex. *Journal of Neuroscience*, **29**, 7526-7539.

Ma, L.Y., Hyman, J.M., Phillips, A.G. & Seamans, J.K. (2014) Tracking Progress toward a Goal in Corticostriatal Ensembles. *Journal of Neuroscience*, **34**, 2244-2253.

Mai, B., Sommer, S. & Hauber, W. (2012) Motivational states influence effort-based decision making in rats: the role of dopamine in the nucleus accumbens. *Cogn Affect Behav Neurosci*, **12**, 74-84.

Mashhoori, A., Hashemnia, S., McNaughton, B.L., Euston, D.R. & Gruber, A.J. (2018) Rat anterior cingulate cortex recalls features of remote reward locations after disfavoured reinforcements. *Elife*, **7**.

Matsumoto, K., Suzuki, W. & Tanaka, K. (2003) Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science*, **301**, 229-232.

McNaughton, B.L. (1998) The neurophysiology of reminiscence. *Neurobiol Learn Mem*, **70**, 252-267.

McNaughton, B.L., O'Keefe, J. & Barnes, C.A. (1983) The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J Neurosci Methods*, **8**, 391-397.

Mueller, V.A., Brass, M., Waszak, F. & Prinz, W. (2007) The role of the preSMA and the rostral cingulate zone in internally selected actions. *Neuroimage*, **37**, 1354-1361.

Murray, E.A., Wise, S.P. & Graham, K.S. (2017) *The evolution of memory systems: ancestors, anatomy, and adaptations*, Oxford University Press.

O'Doherty, J.P. (2004) Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Curr Opin Neurobiol*, **14**, 769-776.

Odum, A.L. & Shahan, T.A. (2004) D-Amphetamine reinstates behavior previously maintained by food: importance of context. *Behav Pharmacol*, **15**, 513-516.

Okano, K. & Tanji, J. (1987) Neuronal Activities in the Primate Motor Fields of the Agranular Frontal-Cortex Preceding Visually Triggered and Self-Paced Movement. *Exp Brain Res*, **66**, 155-166.

Pasquereau, B. & Turner, R.S. (2013) Limited encoding of effort by dopamine neurons in a cost-benefit trade-off task. *J Neurosci*, **33**, 8288-8300.

Passingham, R.E., Bengtsson, S.L. & Lau, H.C. (2010) Medial frontal cortex: from self-generated action to reflection on one's own performance. *Trends Cogn Sci*, **14**, 16-21.

Passingham, R.E. & Lau, H.C. (2006) Free choice and the human brain. In Pockett, S., Banks, W.P., Gallagher, S. (eds) *Does consciousness cause behavior?* Mit Press, pp. 53-72.

Paus, T. (2001) Primate anterior cingulate cortex: Where motor control, drive and cognition interface. *Nat Rev Neurosci*, **2**, 417-424.

Paxinos, G. & Watson, C. (2014) *Paxino's and Watson's The rat brain in stereotaxic coordinates*. Elsevier/AP, Academic Press is an imprint of Elsevier, Amsterdam ; Boston.

Paxinos, G., Watson, C., Calabrese, E., Badea, A. & Johnson, G.A. (2015) *MRI / DTI atlas of the rat Brain*. Elsevier, Academic Press, London ; San Diego.

Pennartz, C.M., Berke, J.D., Graybiel, A.M., Ito, R., Lansink, C.S., van der Meer, M., Redish, A.D., Smith, K.S. & Voorn, P. (2009) Corticostriatal Interactions during Learning, Memory Processing, and Decision Making. *J Neurosci*, **29**, 12831-12838.

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.

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.

Pontieri, F.E., Tanda, G. & DiChiara, G. (1995) Intravenous cocaine, morphine, and amphetamine preferentially increase extracellular dopamine in the "shell" as compared with the "core" of the rat nucleus accumbens. *P Natl Acad Sci USA*, **92**, 12304-12308.

Porrino, L.J. & Goldman‐Rakic, P.S. (1982) Brainstem innervation of prefrontal and anterior cingulate cortex in the rhesus monkey revealed by retrograde transport of HRP. *Journal of Comparative Neurology*, **205**, 63-76.

Porter, B.S., Hillman, K.L. & Bilkey, D.K. (2019) Anterior cingulate cortex encoding of effortful behavior. *Journal of Neurophysiology*, **121**, 701-714.

Pratt, W.E. & Mizumori, S.J.Y. (2001) Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behavioural Brain Research*, **123**, 165-183.

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**, 14080-14090.

Pum, M., Carey, R.J., Huston, J.P. & Muller, C.P. (2007) Dissociating effects of cocaine and d-amphetamine on dopamine and serotonin in the perirhinal, entorhinal, and prefrontal cortex of freely moving rats. *Psychopharmacology (Berl)*, **193**, 375-390.

Randrup, A. & Munkvad, I. (1967) Stereotyped activities produced by amphetamine in several animal species and man. *Psychopharmacologia*, **11**, 300-310.

Randrup, A., Munkvad, I. & Udsen, P. (1963) Adrenergic Mechanisms and Amphetamine Induced Abnormal Behaviour. *Acta Pharmacol Toxicol (Copenh)*, **20**, 145-157.

Ren, J., Xu, H., Choi, J.K., Jenkins, B.G. & Chen, Y.I. (2009) Dopaminergic response to graded dopamine concentration elicited by four amphetamine doses. *Synapse*, **63**, 764-772.

Ridley, R.M., Baker, H.F., Owen, F., Cross, A.J. & Crow, T.J. (1982) Behavioral and Biochemical Effects of Chronic Amphetamine Treatment in the Vervet Monkey. *Psychopharmacology*, **78**, 245-251.

Robinson, T.E. & Berridge, K.C. (1993) The Neural Basis of Drug Craving - an Incentive-Sensitization Theory of Addiction. *Brain Res Rev*, **18**, 247-291.

Romo, R. & Schultz, W. (1987) Neuronal-Activity Preceding Self-Initiated or Externally Timed Arm Movements in Area-6 of Monkey Cortex. *Exp Brain Res*, **67**, 656-662.

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**, 1161-1168.

Rushworth, M.F.S., Hadland, K.A., Gaffan, D. & Passingham, R.E. (2003) The effect of cingulate cortex lesions on task switching and working memory. *J Cognitive Neurosci*, **15**, 338-353.

Rushworth, M.F.S., Noonan, M.P., Boorman, E.D., Walton, M.E. & Behrens, T.E. (2011) Frontal Cortex and Reward-Guided Learning and Decision-Making. *Neuron*, **70**, 1054-1069.

Rushworth, M.F.S., Walton, M.E., Kennerley, S.W. & Bannerman, D.M. (2004) Action sets and decisions in the medial frontal cortex. *Trends Cogn Sci*, **8**, 410-417.

Sagvolden, T. (2011) Impulsiveness, overactivity, and poorer sustained attention improve by chronic treatment with low doses of l-amphetamine in an animal model of Attention-Deficit/Hyperactivity Disorder (ADHD). *Behav Brain Funct*, **7**.

Salamone, J.D. (1994) The Involvement of Nucleus-Accumbens Dopamine in Appetitive and Aversive Motivation. *Behavioural Brain Research*, **61**, 117-133.

Salamone, J.D. & Correa, M. (2002) Motivational views of reinforcement: implications for understanding the behavioral functions of nucleus accumbens dopamine. *Behavioural Brain Research*, **137**, 3-25.

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, psychiatry, and drug abuse. *J Pharmacol Exp Ther*, **305**, 1-8.

Salamone, J.D., Correa, M., Mingote, S.M., Weber, S.M. & Farrar, A.M. (2006) Nucleus accumbens dopamine and the forebrain circuitry involved in behavioral activation and effort-related decision making: implications for understanding anergia and psychomotor slowing in depression. *Current Psychiatry Reviews*, **2**, 267-280.

Salamone, J.D., Cousins, M.S. & Bucher, S. (1994a) 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.

Salamone, J.D., Cousins, M.S., McCullough, L.D., Carriero, D.L. & Berkowitz, R.J. (1994b) Nucleus accumbens dopamine release increases during instrumental lever pressing for food but not free food consumption. *Pharmacol Biochem Behav*, **49**, 25-31.

Samejima, K., Ueda, Y., Doya, K. & Kimura, M. (2005) Representation of action-specific reward values in the striatum. *Science*, **310**, 1337-1340.

Sams-Dodd, F. (1998) Effects of continuous D-amphetamine and phencyclidine administration on social behaviour, stereotyped behaviour, and locomotor activity in rats. *Neuropsychopharmacology*, **19**, 18-25.

Schiorring, E. (1971) Amphetamine Induced Selective Stimulation of Certain Behaviour Items with Concurrent Inhibition of Others in an Open-Field Test with Rats. *Behaviour*, **39**, 1-+.

Schmidt, L., Lebreton, M., Clery-Melin, M.L., Daunizeau, J. & Pessiglione, M. (2012) Neural Mechanisms Underlying Motivation of Mental Versus Physical Effort. *Plos Biol*, **10**.

Scholl, J., Kolling, N., Nelissen, N., Wittmann, M.K., Harmer, C.J. & Rushworth, M.F.S. (2015) The Good, the Bad, and the Irrelevant: Neural Mechanisms of Learning Real and Hypothetical Rewards and Effort. *Journal of Neuroscience*, **35**, 11233-11251.

Schweimer, J. & Hauber, W. (2005) Involvement of the rat anterior cingulate cortex in control of instrumental responses guided by reward expectancy. *Learn Mem*, **12**, 334-342.

Schweimer, J. & Hauber, W. (2006) Dopamine D1 receptors in the anterior cingulate cortex regulate effort-based decision making. *Learn Mem*, **13**, 777-782.

Schweimer, J., Saft, S. & Hauber, W. (2005) Involvement of catecholamine neurotransmission in the rat anterior cingulate in effort-related decision making. *Behav Neurosci*, **119**, 1687-1692.

Seamans, J.K., Lapish, C.C. & Durstewitz, D. (2008) Comparing the Prefrontal Cortex of Rats and Primates: Insights from Electrophysiology. *Neurotox Res*, **14**, 249-262.

Seiden, L.S., Sabol, K.E. & Ricaurte, G.A. (1993) Amphetamine: effects on catecholamine systems and behavior. *Annu Rev Pharmacol Toxicol*, **33**, 639-677.

Seo, H. & Lee, D. (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci*, **27**, 8366-8377.

Shenhav, A., Botvinick, M.M. & Cohen, J.D. (2013) The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron*, **79**, 217-240.

Shidara, M. & Richmond, B.J. (2002) Anterior cingulate: Single neuronal signals related to degree of reward expectancy. *Science*, **296**, 1709-1711.

Shima, K., Aya, K., Mushiake, H., Inase, M., Aizawa, H. & Tanji, J. (1991) 2 Movement-Related Foci in the Primate Cingulate Cortex Observed in Signal-Triggered and Self-Paced Forelimb Movements. *Journal of Neurophysiology*, **65**, 188-202.

Shima, K. & Tanji, J. (1998a) Both supplementary and presupplementary motor areas are crucial for the temporal organization of multiple movements. *Journal of Neurophysiology*, **80**, 3247-3260.

Shima, K. & Tanji, J. (1998b) Role for cingulate motor area cells in voluntary movement selection based on reward. *Science*, **282**, 1335-1338.

Shin, R., Cao, J., Webb, S.M. & Ikemoto, S. (2010) Amphetamine administration into the ventral striatum facilitates behavioral interaction with unconditioned visual signals in rats. *PLoS One*, **5**, e8741.

Shoblock, J.R., Sullivan, E.B., Maisonneuve, I.M. & Glick, S.D. (2003) Neurochemical and behavioral differences between d-methamphetamine and d-amphetamine in rats. *Psychopharmacology (Berl)*, **165**, 359-369.

Silber, B.Y., Croft, R.J., Papafotiou, K. & Stough, C. (2006) The acute effects of d-amphetamine and methamphetamine on attention and psychomotor performance. *Psychopharmacology*, **187**, 154-169.

Skvortsova, V., Palminteri, S. & Pessiglione, M. (2014) Learning To Minimize Efforts versus Maximizing Rewards: Computational Principles and Neural Correlates. *Journal of Neuroscience*, **34**, 15621-15630.

Solanto, M.V. (1998) Neuropsychopharmacological mechanisms of stimulant drug action in attention-deficit hyperactivity disorder: a review and integration. *Behavioural Brain Research*, **94**, 127-152.

Sostek, A.J., Buchsbaum, M.S. & Rapoport, J.L. (1980) Effects of Amphetamine on Vigilance Performance in Normal and Hyperactive-Children. *J Abnorm Child Psych*, **8**, 491-500.

Spielberg, J.M., Miller, G.A., Warren, S.L., Engels, A.S., Crocker, L.D., Banich, M.T., Sutton, B.P. & Heller, W. (2012) A brain network instantiating approach and avoidance motivation. *Psychophysiology*, **49**, 1200-1214.

Steiner, A.P. & Redish, A.D. (2014) Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nature Neuroscience*, **17**, 995-1002.

Strait, C.E., Sleezer, B.J., Blanchard, T.C., Azab, H., Castagno, M.D. & Hayden, B.Y. (2016) Neuronal selectivity for spatial positions of offers and choices in five reward regions. *Journal of Neurophysiology*, **115**, 1098-1111.

Strick, P.L., Dum, R.P. & Picard, N. (1998) Motor areas on the medial wall of the hemisphere. *Novart Fdn Symp*, **218**, 64-80.

Sul, J.H., Jo, S., Lee, D. & Jung, M.W. (2011) Role of rodent secondary motor cortex in value-based action selection. *Nat Neurosci*, **14**, 1202-1208.

Sul, J.H., Kim, H., Huh, N., Lee, D. & Jung, M.W. (2010) Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron*, **66**, 449-460.

Tobler, P.N., Fiorillo, C.D. & Schultz, W. (2005) Adaptive coding of reward value by dopamine neurons. *Science*, **307**, 1642-1645.

Tversky, A. & Kahneman, D. (1981) The Framing of Decisions and the Psychology of Choice. *Science*, **211**, 453-458.

Van Hoesen, G.W., Morecraft, R.J. & Vogt, B.A. (1993) Connections of the monkey cingulate cortex. In Gabriel, M., Vogt, B.A. (eds) *Neurobiology of Cingulate Cortex and Limbic Thalamus: a comprehensive handbook*, Birkhauser, Boston, pp. 249-284.

Von Neumann, J. & Morgenstern, O. (1947) *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, N. J.,.

Wallis, J.D. & Kennerley, S.W. (2011) Contrasting reward signals in the orbitofrontal cortex and anterior cingulate cortex. *Critical Contributions of the Orbitofrontal Cortex to Behavior*, **1239**, 33-42.

Walton, M.E., Bannerman, D.M., Alterescu, K. & Rushworth, M.F.S. (2003) Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *Journal of Neuroscience*, **23**, 6475-6479.

Walton, M.E., Bannerman, D.M. & Rushworth, M.F. (2002) The role of rat medial frontal cortex in effort-based decision making. *J Neurosci*, **22**, 10996-11003.

Walton, M.E., Croxson, P.L., Behrens, T.E.J., Kennerley, S.W. & Rushworth, M.F.S. (2007) Adaptive decision making and value in the anterior cingulate cortex. *Neuroimage*, **36**, T142-T154.

Walton, M.E., Groves, J., Jennings, K.A., Croxson, P.L., Sharp, T., Rushworth, M.F. & Bannerman, D.M. (2009) Comparing the role of the anterior cingulate cortex and 6-hydroxydopamine nucleus accumbens lesions on operant effort-based decision making. *Eur J Neurosci*, **29**, 1678-1691.

Walton, M.E., Kennerley, S.W., Bannerman, D.M., Phillips, P.E.M. & Rushworth, M.F.S. (2006) Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, **19**, 1302-1314.

Wardle, M.C., Treadway, M.T., Mayo, L.M., Zald, D.H. & de Wit, H. (2011) Amping up effort: effects of d-amphetamine on human effort-based decision-making. *J Neurosci*, **31**, 16597-16602.

Wellman, P.J., Davis, K.W., Clifford, P.S., Rothman, R.B. & Blough, B.E. (2009) Changes in feeding and locomotion induced by amphetamine analogs in rats. *Drug Alcohol Depend*, **100**, 234-239.

Wilkinson, L.S., Mittleman, G., Torres, E., Humby, T., Hall, F.S. & Robbins, T.W. (1993) Enhancement of Amphetamine-Induced Locomotor-Activity and Dopamine Release in Nucleus-Accumbens Following Excitotoxic Lesions of the Hippocampus. *Behavioural Brain Research*, **55**, 143-150.

Wilson, M.A. & McNaughton, B.L. (1993) Dynamics of the hippocampal ensemble code for space. *Science*, **261**, 1055-1058.

Wong, S.A., Thapa, R., Badenhorst, C.A., Briggs, A.R., Sawada, J.A. & Gruber, A.J. (2017) Opposing Effects of Acute and Chronic D-Amphetamine on Decision-Making in Rats. *Neuroscience*, **345**, 218-228.

Wunderlich, K., Rangel, A. & O'Doherty, J.P. (2009) Neural computations underlying action-based decision making in the human brain. *P Natl Acad Sci USA*, **106**, 17199-17204.

Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V. & Sahani, M. (2009) Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity (vol 102, pg 614, 2009). *Journal of Neurophysiology*, **102**, 2008-2008.

Zetterstrom, T., Sharp, T., Marsden, C.A. & Ungerstedt, U. (1983) In vivo measurement of dopamine and its metabolites by intracerebral dialysis: changes after d-amphetamine. *J Neurochem*, **41**, 1769-1773.