Modulation of Active Exploratory Behaviors in Humans

by

Nathaniel John Clement

Department of Psychology and Neuroscience
Duke University

Date:_____

Approved:

_____
R. Alison Adcock, Supervisor

_____
Scott Huettel

_____
Tobias Egner

_____
Timothy Strauman

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Psychology and Neuroscience in the Graduate School
of Duke University

2016

ABSTRACT

Modulation of Active Exploratory Behaviors in Humans

by

Nathaniel John Clement

Department of Psychology and Neuroscience
Duke University

Date:_____

Approved:

_____
R. Alison Adcock, Supervisor

_____
Scott Huettel

_____
Tobias Egner

_____
Timothy Strauman

An abstract of a dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Psychology and Neuroscience in the Graduate School
of Duke University

2016

Abstract

Human learning and memory relies on a broad network of neural substrates, and is sensitive to a range of environmental factors and behaviors. The present studies are designed to investigate the modulation of active exploration behaviors in humans, which can serve to regulate learning in response to environmental challenge. In the current work, we operationalize exploration in two ways: participants' spatial navigation (using a computer mouse) in environments containing rewarding and informative stimuli, and participants' eyegaze activity while viewing images on a computer screen. The study described in Study 1 investigates the relationship between spatial exploration and reward, using participants' reported anxiety levels to predict between-subject variability in vigor and information-seeking. The study described in Study 2 investigates the relationship between cue-outcome predictive validity and eyegaze behavior during learning; additionally, we test the extent to which differing states of expectation drive differences in eyegaze behavior to novel images. The study described in Study 3 expands on the questions raised in Study 2: using functional imaging and eyetracking, we investigate the relationship between predictive validity, gaze, and the neural systems supporting active exploration. Taken together, the findings in the present study suggest that emerging certainty in reward outcomes, rather than

uncertainty, drives exploration and associative learning for events and their outcomes as well as encoding into long-term memory.

Dedication

This dissertation is dedicated to my parents, Kurt and Sue, and to all of my teachers.

Contents

List of Figures

Acknowledgements

I would like to thank my advisor, Alison Adcock, as well as my committee for their help and advice. I would like to thank the members of the Adcock Lab for their invaluable contributions. Additionally, I would like to thank my family and friends for their love and support.

1. General Introduction

Humans (and other animals) are not passive vessels into which information is poured; instead, we must explore our environments to observe and experience events, to make and test predictions, and ultimately to form coherent theories about the structure of the world and our ability to pursue goals. The mere presence of reward within an animal's environment does not guarantee reward receipt. In order to successfully gain rewards, animals must be able to actively learn about their environments in order to predict where and when rewards will be available, and what actions they must perform to gain them.

Numerous studies in cognitive neuroscience have illustrated the relationship between motivation, active learning, and adaptive behavior: active exploration of spatial environments has been shown to enhance learning relative to passive presentation of information (Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011; Voss, Warren, et al., 2011), information associated with cues predicting high-magnitude or high-probability reward delivery is more reliably encoded relative to information associated with cues associated with a lower magnitude or probability of reward (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Wittmann et al., 2005), and choices between competing rewarding options can closely track the shifting learned values of available choices (Kovach et al., 2012; Li & Daw, 2011; Schonberg, Daw, Joel, & O'Doherty, 2007). In all these cases, encoding of information is enhanced when paired

with motivationally significant stimuli or actions, and within environments that contain the promise of potential rewards.

Not all motivation is reward-oriented; indeed, in naturalistic environments, a novel stimulus (be it an undiscovered kind of plant, an unexplored location, or a heretofore unknown individual) are at least as likely to pose a threat as to present an opportunity for reward. Studies examining the relationship between negative motivation (such as learning under the threat of aversive outcomes, such as monetary loss or shock) have shown that memory for the motivationally-relevant stimuli can be enhanced, but that the neural mechanisms supporting learning under potentially aversive conditions differs from those recruited when learning is modulated by the possibility for reward (Dunsmoor, Murty, Davachi, & Phelps, 2015; Murty, LaBar, Hamilton, & Adcock, 2011; Phelps, 2004). The systems involved show considerable overlap: for example, the dopaminergic midbrain has been shown to respond both to motivationally positive and aversive events (Schultz, 2007). Thus, to guide optimal behavior, both potentially rewarding stimuli and potentially aversive (or neutral) stimuli must be represented and compared, in a manner that may potentially be supported by the same re regulatory mechanisms.

However, "more" learning is not always "better" learning: though a barrage of events may occur within an individual's scope of reference, the events that an individual detects may or not be motivationally relevant and may or may not impact aspects of the

environment that an organism can influence. Additionally, learning comes at a cost: stimuli must be attended to, at the exclusion of other, potentially more beneficial or detrimental events; and, once attended to, must be maintained and successfully encoded into memory. This requires the expenditure of physical energy in processing, as well as the opportunity costs entailed in selecting one stimulus at the expense of others.

Exploration does not occur without a cost: the actions required to pursue potential rewards require energy, increase exposure to unknown risks in the environment, and present an opportunity cost in potential gains. Thus, animals must regulate the active exploration behaviors they employ.

Thus, the expenditure of effort in active exploration must be controlled, and modulated as a result of both external events (for example, the degree to which an event's outcome was predicted, as well as its positive, negative, or neutral motivational significance) as well as internal states (for example, physiological states of hunger or thirst will enhance the motivational relevance of food- or water-related stimuli.)

Computational and theoretical models of human learning in large part rest on the assumption that reward maximization is a primary goal of learning and reward-oriented behaviors (Moustafa, Cohen, Sherman, & Frank, 2008). While maximizing "reward," loosely defined to include not simply monetary or hedonically-relevant stimuli such as money, food, or other fundamental reinforcers, but also to include the information signaling shifts in the availability of these reinforcers. Thus, exploratory

behaviors may be regulated through signals relating to information present in the environment.

One potential driver of exploratory behaviors may be environmental volatility, or the extent to which outcomes can be predicted (Berlyne, 1960). Expectation-violations may generate an internal signal that more learning is needed in order to better predict outcomes, potentially driving shifts in active exploratory behaviors (Esber & Haselgrove, 2011; Hall, 1991; P. C. Holland, Thornton, & Ciali, 2000; Pearce & Hall, 1980; Sara & Bouret, 2012). However, an alternative driver of exploratory behavior may be effective prediction; that is, the extent to which outcomes can reliably be predicted from cues or events (Esber & Haselgrove, 2011; Haselgrove, Esber, Pearce, & Jones, 2010). As an individual comes to learn that a given stimulus strongly predicts motivationally relevant outcomes, its inherent salience may be increased relative to less-predictive stimuli, encouraging greater manipulation and active learning.

Both uncertainty signals and effective history could promote exploration, but both processes must be regulated. Discovering gold amongst railroad ties may be a surprising event that captures attention and increases our likelihood of exploring further; additionally, having reliably found scattered change at a railroad crossing may increase an individual's tendency to search that environment for coins. However, an enhanced willingness to explore the train tracks that outweighs our sense of caution to monitor the environment for oncoming trains is unlikely to prove adaptive in the long

term. Thus, exploration behaviors must be regulated dynamically, with motivationally salient stimuli impacting behavior within their behaviorally-relevant contexts.

The mechanisms by which information-seeking is instigated, maintained, and directed, however, remain incompletely characterized. The drivers of this state may be internal or external, but may reflect an individual's recognition of either an "information gap" (a mismatch or gap between its learned representation of the world and events in its environment) or the detection of a shift in this information gap, as new information confirms prior predictions and diminishes the gap between what is predicted and what occurs.

1.1 Associative Models of Learning

Learning to predict events is critical for an individual to make adaptive choices. Although, in many instances, following a choice or action, an individual can receive direct reinforcement which increases or decreases the likelihood of repeating that action, in many cases, participants must make decisions without immediate reinforcement. Adaptive learning to maximize reward relies on choices which have never directly been reinforced, but where past experience with relevant stimuli may provide information regarding potential outcomes. A number of computational models have attempted to characterize potential processes by which animals may learn to bind stimuli to each other and to their consequences, a cognitive domain known as associative learning (Eichenbaum, 2004). The most influential of the early generation of models was put

forward by Rescorla and Wagner in 1972 (Miller, Barnet, & Grahame, 1995; Rescorla, 1972).

According to the Rescorla-Wagner model, the link (formally, the associative strength) between each stimulus and a consequent changes incrementally with each repeated pairing, approaching an asymptote over extended training (Rescorla, 1972). The size of this change is determined according to the following equation:

$$\Delta Vx=\alpha x\beta(\lambda-V)$$

The change in associative strength (ΔV) is modulated by a set of parameters: α, determined by properties of the conditioned stimulus (CS) X, including its salience (defined by its physical and sensory properties); β, determined by, the rate at which of the unconditioned stimulus (US) can impact learning (Rescorla, 1972); λ, the maximum associative strength possible for the US; and, critically, the degree to which the US is already predicted by existing associations between it and other CSs. If a given CS is of low salience, and the US with which it is paired is likewise insufficient to produce a response, the associative strength between the two will increase slowly. If, however, either or both factors were to increase (if the CS were to become more salient, or if the US were to increase in effectiveness) the change in associative strength would increase accordingly, provided that the US is not already completely predicted by other stimuli.

This dual strength—providing a stepwise mechanism describing how the association between stimuli and events could change over time to predict learning, and

designing constraints to permit that learning to specifically target relevant aspects of an environment and determine which stimuli are reliable predictors—allowed researchers to apply Rescorla-Wagner models of learning to behavioral and neural data to understand aspects of learning. Studies of conditioning in rats supported aspects of the basic structure of the Rescorla-Wagner model, with variability in parameters such as CS salience and US strength predicting learning rates and behavior in rats (Miller et al., 1995). Of particular significance were the model's ability to accurately describe blocking effects: the tendency for an animal, having learned that CS A predicts US X, will show no learning towards CS B following repeated pairings of A and B followed by X (Rescorla, 1972). That is, learning that B also predicts X is "blocked" by the existing strong association between A and X. Thus, two main driving forces behind learning were suggested by Rescorla-Wagner's model: learning is determined by the structure of the environment and the strength of the US.

1.2 Salience in Associative Learning Models

The learning mechanisms put forth by Rescorla and Wagner describe how animals may learn to associate particular stimuli with one another through manipulations of these key variables: primarily, the effectiveness of a US, the properties of a CS as perceived by the animal, and the interaction of these properties with the complexity and existing structure of an animal's environment. The unique contributions of the Rescorla-Wagner were primarily in its delta-updating rule, where by the

associative strength between stimuli changes according to the difference between predicted events and experienced events. The implications of this and related error signals are discussed further in Section IV. Subsequent models of associative learning, however, including both the Pearce-Hall model and Mackintosh model, describe the effects of changes in the animal's perception of the CS as a potential mechanism for driving learning.

Current models of associative learning make opposing predictions about the variability associated with a stimulus and the quality and strength of learning that an organism may produce. This property—the ease with which a stimulus can become associated with another stimulus or events—is referred to as a stimulus' "associability" (Esber & Haselgrove, 2011; Mackintosh.J., 1975; Pearce & Hall, 1980). In both the Pearce-Hall and Mackintosh/Esber-Haselgrove models of learning, a stimulus' associability changes in response to the extent to which it can reliably serve as a predictor.

According to models put forward by Mackintosh (Mackintosh.J., 1975) and further developed by Esber and Haselgrove (Esber & Haselgrove, 2011), if stimulus A reliably predicts stimulus B, its associability will increase, increasing its effective salience to the organism, making it a likelier target for subsequent attention and additional learning. In a sense, the animal has learned that stimulus A is a reliable indicator that an event is about to occur: thus, attending to it will increase the odds of adaptively responding to its consequent. Thus, stimuli with higher effective salience, which serve as

more reliable predictors of outcomes, should capture more attention and potentially serve to increase information-seeking.

The Pearce-Hall model, however, offers a conflicting mechanism. According to their framework, if a stimulus reliably predicts a stimulus or event, its associability is diminished (Hall, 1991; Pearce & Hall, 1980). Instead, stimuli whose outcomes remain uncertain or ambiguous will show an increased associability, as animals devote additional attention towards components of their environment that are incompletely understood (Hall, 1991). Within this framework, the assocability of a given stimulus is defined not as its effective salience, but as the sum of its past prediction errors.

Data from humans and animals support aspects of both models: considerable evidence from rat models of learning suggest that shifts in behaviors in response to unexpected outcomes fits the predictions of the Pearce-Hall model (P. C. Holland & Gallagher, 1999; P. C. Holland et al., 2000), while neural data from animal models and functional imaging in humans shows evidence that neural activity tracks shifts in associability in response to unexpected shifts in contingency; however, considerable eyetracking and learning data in humans suggests that overt attention to predictive cues is greater than overt attention to less predictive cues (Beesley, Nguyen, Pearson, & Le Pelley, 2015; Failing, Nissens, Pearson, Le Pelley, & Theeuwes, 2015; Le Pelley, Beesley, & Griffiths, 2011; Le Pelley, Mitchell, Beesley, George, & Wills, 2016; Luque, Vadillo, Le

Pelley, & Beesley, 2016). Thus, aspects of both models are likely to be reflected in driving learning behaviors and neural signals.

1.3 Neural Systems Supporting Exploration Behaviors

Exploration behaviors allow an individual to probe complex, and shifting, environmental contexts and embedded stimuli. The behaviors that support adaptive learning and choice rely on a set of neural systems, described below.

1.3.1 Learning and Memory: Hippocampus and Neuromodulatory Inputs

The hippocampus is necessary for the formation and manipulation of long-term memory and for spatial navigation (Eichenbaum, 2004; Gabrieli, 1998). However, the nature and strength of the representations that are formed depend on the integration of the hippocampus within functional networks.

One prominent modulator of the hippocampal memory system is the neurotransmitter dopamine (DA): dopamine is critical for long-term potentiation (Frey, Schroeder, & Matthies, 1990; Otmakhova & Lisman, 1998) and enhanced connectivity between the hippocampus and the dopaminergic midbrain facilitates encoding success (Adcock et al., 2006). Outside its role in hippocampal plasticity, dopamine is critical both in learning to predict the appearance of reward and in facilitating the appropriate response in order to acquire it (Guitart-Masip, Fuentemilla, et al., 2011; Schultz, 1998)

Additionally, the neuromodulator norepinephrine (NE) has been shown to impact a variety of functions directly relevant to both attending to stimuli and

responding to environmental variability: phasic increases in noradrenergic activity are associated with identifying relevant stimuli and inducing switches of attention, as an animal disengages with a stimulus and re-engages with a different target (Aston-Jones, Chiang, & Alexinsky, 1991; Aston-Jones, Rajkowski, & Cohen, 1999; Aston-Jones, Rajkowski, Kubiak, & Alexinsky, 1994; Grant, Aston-Jones, & Redmond, 1988). Additionally, increased baseline or tonic levels of activity within the LC, producing an increased baseline level of NE in its efferent targets, including prefrontal cortex (Florin-Lechner, Druhan, Aston-Jones, & Valentino, 1996) and the amygdala (Asan, 1998) is associated with a broadening of attention (Berridge & Waterhouse, 2003).

Both dopamine and norepinephrine have been shown to influence aspects of ocular physiological function, including blink rate and pupil diameter (Samuels & Szabadi, 2008; Spiers & Calne, 1969). Investigations into eyetracking responses have demonstrated that pupil diameter shifts within a reward-driven bandit task (Jepma & Nieuwenhuis, 2011) and in response to surprising events in humans (Preuschoff, t Hart, & Einhauser, 2011), suggesting a role for these catecholamines in driving salience-driven shifts in active exploration.

1.3.2 Vigor

Dopaminergic projections to the nucleus accumbens drive locomotor activity (Wu, Brudzynski, & Mogenson, 1993), and accumbens activity in response to a reward-predictive cue predicts subsequent vigor (Catanese & van der Meer, 2013; McGinty,

Lardeux, Taha, Kim, & Nicola, 2013). Drug manipulations in humans also suggest that increased dopamine levels drive enhanced response vigor (Beierholm et al., 2013). However, disruption of inputs from either the dopaminergic midbrain or the basolateral amygdala impair this invigoration (Ambroggi, Ishikawa, Fields, & Nicola, 2008). The role of dopamine in information-seeking and vigor is further discussed in Chapter 2.

1.3.3 Salience

The amygdala has been consistently shown to respond to salient events, and functionally intact connections between the amygdala and the dopaminergic midbrain are necessary for generating orienting responses to unexpected events and updating behavior in response to violations of expectations (El-Amamy & Holland, 2006; Esber et al., 2012). Thus, adaptive regulation of active exploration may rely on the interaction of the hippocampus with a functional network including the dopaminergic midbrain, its striatal targets, and the amygdala. Of critical interest for the present study is how associative salience is determined, signaled within our critical systems of interest, and the behavioral relevance of that signal.

1.4 Implications for Learning

The hippocampus is critical for encoding information into long-term memory (Eichenbaum, 2004); however, data from humans and animals suggest that the strength and character of memory formation relies on system-level interactions between the hippocampus and associated neural structures supporting motivated behaviors,

primarily including the dopaminergic midbrain and amygdala (Frey et al., 1990; Kumaran & Maguire, 2009; Lisman & Grace, 2005; Phelps, 2004). In the present studies, a critical question is the sensitivity of these regulatory networks to environmental signals of reward certainty or uncertainty, and, especially, their effects in up- or down-regulating hippocampally-dependent processing of items into long-term memory.

1.5 Specific Hypotheses

The present studies are designed to investigate the modulation of active exploration behaviors in humans. In the current work, we operationalize exploration in two ways: participants' spatial navigation (using a computer mouse) in environments containing rewarding and informative stimuli, and participants' eyegaze activity while viewing images on a computer screen. The study described in Aim 1 investigates the relationship between spatial exploration and reward, using participants' reported anxiety levels to predict between-subject variability in vigor and information-seeking. The study described in Aim 2 investigates the relationship between cue-outcome predictive validity and eyegaze behavior during learning; additionally, we test the extent to which differing states of expectation drive differences in eyegaze behavior to novel images. The study described in Aim 3 expands on the questions raised in Aim 2: using functional imaging and eyetracking, we investigate the relationship between predictive validity, gaze, and the neural systems supporting active exploration.

In the following chapters, I will present three experiments designed to examine the relationship between participants' experience of reward, their active exploration, and the encoding of information into long-term memory.

We will describe tests of the following hypotheses:

1. Information-seeking behaviors will be facilitated in the presence of reward: participants will generate more explicit vigor in pursuing rewards of greater (vs. lesser) magnitude, and greater (vs. lesser) history of reward delivery will promote enhanced vigor in information-seeking.

2. Overt attention, as measured through eyetracking, will be greater to stimuli that serve as reliable predictors of events in their environment, relative to stimuli that serve as unreliable predictors.

3. The neural systems that support active learning--including the hippocampus/medial lobe cortices, amygdala, and dopaminergic midbrain--will show a greater response to stimuli that serve as reliable predictors of outcomes than to stimuli that serve as unreliable predictors.

4. Active learning about incidentally-presented stimuli will be enhanced for stimuli presented during states of high certainty, rather than states of high uncertainty.

The study presented in Chapter 2 will address the first hypothesis. The study presented in Chapter 3 will address the second and fourth hypotheses, and the study presented in Chapter 4 will address the third and fourth hypotheses.

2. Active Exploration and Information-Seeking

2.1 Introduction

Individuals are driven to explore their environments by a variety of factors: immediate threats, imminent rewards, and aversive physiological states (e.g., hunger) can all prompt active responses. The optimal means by which animals choose how to allocate their time spent seeking reinforcers is the subject of a considerable body of literature: models of animal foraging, for example, seek to determine how differences between the values of immediate rewards may be outweighed by the potential values of future rewards (Charnov, 1976; Pyke, 1977; Stephen, 1986). In contrast to the pursuit of primary reinforcers (such as the maximization of rewards in foraging) tied to one goal (e.g., the collection of food or the avoidance of harm), the regulation of information-seeking behaviors reflects a more open-ended state that can be impacted by arousal mediated by a number of sources.

Abundant evidence from the experimental, as well as public-health, literature, indicate that humans and other animals consistently act to place themselves in novel and stimulating environments with potential risk in addition to potential rewards (Beckmann, Marusich, Gipson, & Bardo, 2011; Harden & Tucker-Drob, 2011; Myrseth, Tvera, Hagatun, & Lindgren, 2012). Within these cases, the benefit that animals may gain is not limited simply to the chance of reward or risk of punishment, but to the potential gain in information.

Evidence from primates suggests that this is not merely a human phenomenon: rhesus macaques show a distinct preference, and willingness to work, for information about upcoming rewards, even if this information has no impact on the timing of reward delivery (Bromberg-Martin & Hikosaka, 2009). Thus, information that helps an organism improve its understanding of its environment is preferentially targeted, facilitating future choices (Guitart-Masip, Bunzeck, Stephan, Dolan, & Duzel, 2010; Järvelin, 2003; Shohamy & Adcock, 2010).

The relationship between information-seeking and reward, in particular, has significant implications for our understanding of the mechanistic link between learning, memory, and exploration, primarily through the known functions of the dopaminergic system. Considerable evidence suggests that dopamine is associated with exploratory action, in large part through its actions on the striatum, a region that receives significant innervation from the dopaminergic midbrain. Expectation (vs. receipt) of reward has been associated with the release of dopamine in the ventral striatum (VS) (de la Fuente-Fernandez et al., 2002); additionally, simultaneous PET-fMRI evidence suggests that increased dopamine levels in VS correlate with BOLD signal increases during the anticipation of reward, both in VS and in the dopaminergic midbrain (Schott et al., 2008). The disruption of dopaminergic projections to the VS, either through local antagonism or by GABAergic blockade of the VTA, disrupts the active approach of reward in rats (Ikemoto & Panksepp, 1996). Further, evidence from dopamine-depletion studies

suggest that dopamine's effect on the VS modulates a variety of functions including, approach-behavior initiation (Nicola, 2010), effort-based choice (Mai, Sommer, & Hauber, 2012; Salamone, Correa, Nunes, Randall, & Pardo, 2012), exploratory locomotion (Baldo, Sadeghian, Basso, & Kelley, 2002), and behavioral activation in general (Koob, Riley, Smith, & Robbins, 1978; Robbins & Everitt, 2007). Thus, striatal dopamine is integral to the initiation and maintenance of behaviors necessary for an animal to explore and interact with its environment, in addition to its potential role in reward anticipation.

What is the behavioral significance of these signals? Data from humans and animals suggests that there may be a strong relationship between reward, dopamine, and effortful motivation. Response vigor—the effort cost, or work performed, in generating a response—has been shown to relate to an individual's history of reward, and potentially to tonic dopamine levels (Beierholm et al., 2013; Guitart-Masip, Beierholm, Dolan, Duzel, & Dayan, 2011; Kurniawan, Guitart-Masip, & Dolan, 2011; Niv, Daw, Joel, & Dayan, 2007). Tonic dopamine has been characterized as regulating the "thrift" of an individual's response: elevated tonic dopamine signaling, associated with a higher history of mean rewards, predicts more active exploration, while diminished tonic dopamine signaling, associated with a reduced history of reward delivery, predicts less exploration in favor of "thrifty" capitalization of rewards and minimal effort expenditure (Beeler, 2012). Additionally, data from humans suggests that

this relationship between tonic dopamine and response vigor may vary with an individual's perceived sense of self-efficacy: a more external locus of control (which suggests that an individual views the events occurring within their life as determined not by their own actions, but as determined by the actions of others and of the world around them) is associated with variability in exploration behaviors during pharmacological manipulation of tonic dopamine (Kayser, Mitchell, Weinstein, & Frank, 2015).

In the present study, we gave participants the opportunity to gain monetary rewards by generating motor responses. Additionally, on critical trials we extended participants the opportunity to perform effort to gain information about upcoming rewards. Thus the present study investigates the relationships between varying magnitudes of available rewards, varied information available to participants, and effort generated to gain rewards—both directly, in the case of reward pursuit, and indirectly, as participants may expend varying amounts of effort to gain information about upcoming potential rewards.

2.2 Hypotheses

Active exploration relies upon systems supporting reward anticipation and motivated behaviors, including the expenditure of effort or vigor. In Study I, we monitored reward-approach and effortful information-seeking behavior in humans

while manipulating levels of reward magnitude and reward conflict. We anticipate that variability in reward magnitude and conflict will predict invigoration towards rewards. Further, we tested how history of rewards impacted participants' self-generated effort for information relating to future rewards; consistent with literature suggesting that increased dopamine levels in the ventral striatum predicts greater invigoration and effortful choice, we hypothesized that increasing rewards would drive increasing vigor in information-seeking.

2.3 Methods

2.3.A Participants

30 participants (mean age=24, 17 female) were recruited through email. All participants provided informed consent as approved by the Duke University Institutional Review Board.

2.3.B Materials

Stimuli consisted of simple shapes presented on a computer screen. Participants were asked to use a computer mouse to move the cursor toward onscreen locations in order to earn rewards. Additionally, participants were asked to use the keyboard to provide motor responses.

2.3.C Experimental Design and Procedure

The experimental task was designed to measure vigor and motivation in two domains: reward-seeking and information-seeking behavior. To this end, we developed a task in which participants were asked to make simple choices and motor movements in order to gain monetary rewards, while periodically being presented with the opportunity to perform effort to gain information about upcoming rewards.

To gain a reward, participants needed to move the mouse cursor into a target stimulus (a colored rectangle) in either the top-left or top-right corner of the computer screen, as shown in Figures 1 and 2.

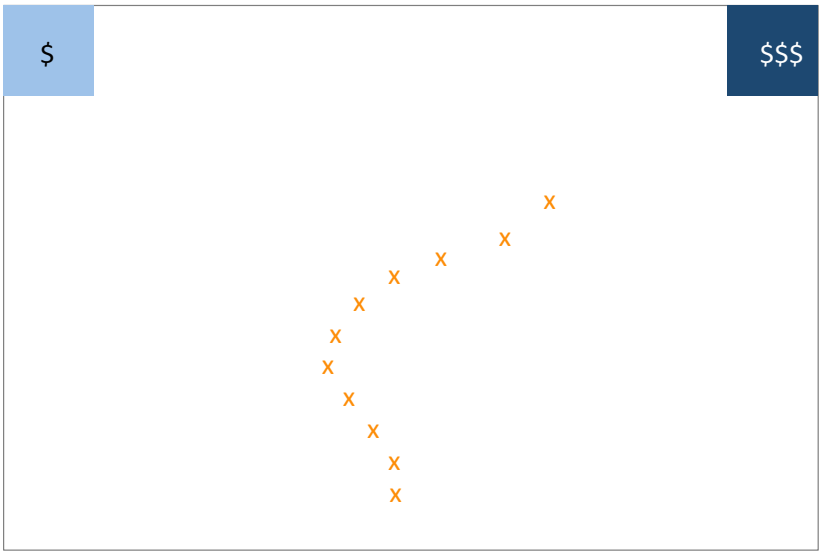

Figure 1: Task schematic for Study 1. Shown here is one example of a "Dual" trial, in which two competing rewards are available. The lesser reward, in the upper left, is a lighter color intensity, while the darker target in the upper right corresponds to a greater monetary reward. Orange hatch marks indicate one (hypothetical) trajectory in reward-seeking.

Figure 2: Apparatus used for Study 2. Participants used a cordless mouse to move up a ramp and toward their targeted stimulus of interest. Following trial conclusion, they moved the mouse to the bottom to begin the next trial. Keypresses were generated using the index finger of their dominant hand on the keyboard.

Trials were split among three conditions: "Single", "Dual", and "Press". "Single" trials presented participants with explicit information about the value of a single potential reward, which was available in either the top-left or top-right corner. "Dual" trials presented participants with explicit information about the value of a two potential rewards, which served as exclusive and competitive choice (participants could not gain both rewards, but needed to choose one.) In both "Single" and "Dual" trials, the value of potential rewards was indicated by the color intensity of the relevant target images, which ranged from fully saturated blue (indicating a value of $1) to black (indicating a value of $0.)

"Press" trials differed critically from the above conditions in that, by default, the explicit value of potential choices on each trial was not shown to participants. Instead, prior to trial start, participants had the opportunity to repeatedly press a specified button on the keyboard within a limited period of time (2 s). If they opted not to generate keypresses, or responded at an insufficiently slow rate, the trial proceeded and participants were asked to choose between two potentially rewarding targets, whose value was not signaled to them. However, if they generated button-presses at or above the response threshold, the values of the potential targets were revealed at trial onset (and were maintained onscreen throughout the trial.) Response thresholds for button-pressing were drawn from a random distribution (bounded between 7 and 12), such that an explicit minimum number of button-presses required to succeed on each trial was not stable and participants thus could not simply learn to generate a minimum response. Instead, the rate, or vigor, of their button-presses corresponds to their internally generated motivation to gain information in pursuit of reward.

Thus, the task allows us to characterize motivated motor behaviors in three domains: simple reward-seeking ("Single"), reward choice conflict ("Double"), and information-seeking ("Press").

Participants were allotted 5 seconds in which to move their mouse from the starting position into one of the onscreen targets in order to gain rewards on each trial. On "Press" trials, the pre-trial period (during which participants were allowed to expend

effort in order to gain information) lasted 2 seconds. There were 50 "Single" trials (25 with the sole available rewarded target on the left, 25 on the right), 25 "Double" trials, and 100 "Press" trials, for 175 trials in total.

The task took less than one hour to complete. Following task conclusion, participants were asked to complete a battery of psychological questionnaires

2.3.D Analysis

As the magnitude of potential rewards and associated reward conflict varied across trials, we used mixed-effect linear regression models to relate trial-by-trial variability in potential rewards, reward conflict and reward-seeking success or failure to three critical dependent variables: movement time (the time elapsed between trial start and the participant reaching a goal), spatial sampling (the ratio of total movement in the X direction relative to the Y direction), and button-press rate in "Press" trials.

To test these relationships, we ran three regression models, detailed as follows.

Model 1 (Reward Pursuit) included the summed magnitude of potential rewards on each trial, the difference in available rewards, and an interaction term as fixed effects and treated participants as a random effect. The dependent variable in this model was movement time.

Model 2 (Reward Conflict) included the summed magnitude of potential rewards on each trial, the difference in available rewards, and an interaction term as

fixed effects and treated participants as a random effect. The dependent variable in this model was the ratio of movement in the X dimension to movement in the Y dimension.

Model 3 (Information-Seeking) tested, for "Press" trials, the relationship between the magnitude of rewards earned on the immediately preceding trial (predictor variable) and the rate of button-presses on each trial (dependent variable).

Additionally, we were curious whether variability in reward-seeking behavior could be shown to relate to individual differences in psychological measures of interest. We included individual scores on the STAI and LOC as group-level regressors to test the relationship between individual variability and the behaviors of interest.

2.4 Results

2.4.A Reward Pursuit

The model's overall adjusted R-squared was 0.3284. The summed magnitude of reward available on each trial showed a significant positive relationship with movement speed ($p=0.041$). The difference in potential reward did not predict movement speed ($p=0.77$), and the reward x conflict interaction was not significant ($p=0.70$). This effect did not vary with individual differences.

2.4.B Reward Conflict

The summed magnitude of reward available on each trial significantly did not predict increased movement in the X dimension relative to Y ($p=0.85$). The difference in

potential reward showed a highly significant positive relationship with the X:Y ratio ($p<0.005$), and the reward x conflict interaction was not significant ($p=0.12$). Additionally, the relationship between choice conflict and X:Y movement showed a significant negative relationship with individuals' reported self-efficacy, as measured by their responses on the Internal Locus of Control ($p<0.05$).

2.4.C Information-Seeking

The model's overall adjusted R-squared was 0.6487. The proportion of reward earned on the immediately preceding trial showed a significant positive relationship with rates of motor response in button-pressing ($p=0.013$). This effect did not vary with individual differences.

2.5 Discussion

2.4.A Summary of Specific Hypotheses and Results

In Study 1, we sought to explore the factors driving allocation of effort, both in direct reward-seeking behavior and in indirect reward-seeking behavior in the form of effort recruited to gain information about upcoming rewards. We operationalized vigor in two ways: spatial motor responses to gain immediate monetary rewards, and response rates for button-pressing in order to gain information. Our findings are consistent with findings in humans and animals relating dopaminergic signaling with reward motivation and active exploration. Increasing magnitudes of potential rewards

encourages vigorous responding on the part of individuals attempting to gain those rewards; however, the presence of multiple, competing rewards diminishes this effect and predicts increasing spatial uncertainty in participant's paths. The critical manipulation in this study regarding information-seeking, that relating reward history with vigor in active information-seeking, is strongly consistent with data from human neuroimaging experiments as well as the hypothesized mechanisms. Additionally, the finding that individuals' self-reported sense of self-efficacy, as measured through their internal vs. external locus of control, predicted shifts in exploratory movements was consistent with our hypotheses and with findings reported in the extant literature. However, the lack of a relationship between the observed effects and participants' reported anxiety levels was surprising: we anticipated that anxiety would serve to moderate the observed effects, but did not find such relationships.

2.4.B Relationship to Prior Literature

Data collection in Study 1 was limited to behavioral measures and self-reported responses to psychological questionnaires; thus, our ability to relate the conclusions of the present study to extant findings in the neurobiological literature are limited. However, we can compare our findings with outcomes predicted by models of dopaminergic function and the psychological literature.

Models of reward history and response vigor, such as that of Niv and colleauges (Niv et al., 2007), suggest that the tonic levels of dopamine, contrasted by (and

potentially reflecting temporally and spatially smoothed input by) phasic spikes in activity by dopaminergic neurons in response to singular events. In this framework, tonic dopamine levels within ventral striatum are thought to reflect the average reward rate, in contrast to the cue- or outcome-specific prediction-error signals encoded by phasic activity. This is supported by evidence from subsequent empirical work, further confirming that tonic dopamine signals the value of expected actions (Zenon, Devesse, & Olivier, 2016), while dopamine depletion (in unmedicated patients with Parkinson's disease) diminished effort recruitment and value estimation (Le Bouc et al., 2016). Our findings add to this body of evidence, both in the confirmatory result of expected reward value driving pursuit vigor (tracking speed), and in the novel finding of participants' willingness to expend effort to gain information about upcoming rewards and the variable relationship of reward history in predicting information-seeking vigor. Of particular note is the direction of this relationship: although we had conjectured that one potential driver for enhanced information-seeking may by a signal corresponding to reward prediction errors, the evidence here suggests that the magnitude of reward, not the errors in predicting it, may serve to up-regulate information-seeking.

Our findings are enriched by a consideration of the psychological concept of self-efficacy. As put forth by Bandura and others (Bandura, 1977, 1982), theories of self-efficacy state that the value and motivation an individual ascribes to a given action reflects not only their context and history of reward or punishment, but critically their

sense of the strength (or weakness) of their actions to influence their experience and the events impacting themselves and the world around them (Bandura, 1982). This construct can be operationalized by scaling individuals according to their perceived locus of control (LOC), reflecting the extent to which they believe the events that happen their reflect their actions and intentions (Internal), those of those forces around them beyond their ability to control (External) (Rotter, 1990). Data from computational models and genetic testing have suggested that a more external LOC is associated with a greater likelihood that an individual's choices in reward pursuit would be biased towards exploration (greater likelihood of novel or risky choices) than towards exploitation (repetition of a known or "safe" response) only when dopamine levels were pharmacologically enhanced. Under placebo, participants with high internal LOC were more exploratory than were participants with low LOC (Kayser et al., 2015). Our data confirm and expand these findings: action selection was impacted by potential reward conflict, but the strength of this effect varied across participants as a function of their reported internal LOC. This finding, in the context of Kayser et al.'s illustration of the sensitivity of individuals' bias towards exploration vs exploitation as a function of dopamine and self-efficacy, underscores the relationship between an individual's experienced history of reward and the value they place upon their own actions in discriminating motivationally relevant events, selecting optimal outcomes, and gaining information to inform future choice.

Additional neurobiological investigations of dopamine, self-efficacy, and response vigor have further clarified the mechanisms relating effort and reward choice. Work by Michael Frank and others have described distinct "No" and "No-go" pathways within the basal ganglia, which differentially support appetitive/approach-oriented and aversive/avoidance-oriented reinforcement and rely on distinct direct (Go) and indirect (NoGo) pathways to excite, or inhibit, outputs from striatum to cortex and ultimately drive choice (Frank, Seeberger, & O'Reilly R, 2004). Work in animals has suggested that motivated, effortful action can be impacted through manipulation of dopaminergic circuitry. In mice with well-characterized motivational deficits, de-activation of the no-go pathway selectively enhanced motivated action selection, but impaired efficiency in maintenance of a sustained response (Carvalho Poyraz et al., 2016). In the context of our data, we may expect to observe similar effects if participants were sorted according to their observed bias in reward learning, potentially relating to variability in go/nogo baseline action and their perceived self-efficacy.

2.4.C Caveats and Limitations

Our interpretation of the results of Study 1 is bounded by limitations in study design and analysis. A principle limitation is the question of uncertainty: in the present study, to simplify choice behaviors and enhance the potential value of information about upcoming rewards, participants made choices whose outcomes were deterministically known. The magnitude of rewards participants could earn on each trial was overtly

indicated, and was not associated with any variable probability of reward. Thus, information about upcoming rewards was reliable and questions regarding participants' willingness to work to gain information was not complicated by differential interpretations of, and tolerance for, ambiguity in participants' awareness of contingencies between choices and outcomes.

2.4.D Outstanding questions

Consideration of the structural and functional dissociations of basal ganglia function offers a clear set of future considerations. The present study manipulated reward and approach motivations; given evidence relating variability in go/no-go function and approach motivation and choice, an open question stands as to whether the observed relationships would persist in the presence of punishments or aversive outcomes, either in whole (replacing the reward-pursuit motivation with loss-avoidance entirely) or in part (presenting rewards and punishments intermixed.) Given the evidence linking tonic dopamine with average history of reward (suggesting a state-specific effect rather than one time-locked to particular events or outcomes), as opposed to the well-characterized relationship between stressors and sympathetic "fight-or-flight" behavioral activation, the presence of punishing or aversive outcomes may serve to alter the relationships observed here, in which trials varied only in the magnitude of reward that could be expected, and not the qualitative shift from the expectation of reward to that of punishment.

3. Associative Learning and Eyetracking

3.1 Introduction

Models of associative learning (e.g., (Esber & Haselgrove, 2011; Mackintosh.J., 1975; Pearce & Hall, 1980)) make strongly differing predictions on the role of outcome uncertainty in driving learning, and provide useful tools in describing signals that may drive learning. However, these models do not attempt to describe the relationship between environmental factors (such as the extent to which cues can, or cannot, be used to reliably predict an outcome) and the modulation of active exploratory behaviors. A question that remains outstanding is: how does variability in outcome certainty drive differences in participants' interaction with their environment?

Data from humans and animals provide a conflicting background. Evidence that both predictive validity and uncertainty drive learning is found both in humans and in animal manipulations.

Behavioral observations of cue-outcome learning in animals suggest that, in many cases, learning proceeds in a manner consistent with the Pearce-Hall model, as reversal-learning is facilitated for stimuli with a high level of uncertainty salience, compared to stimuli with lower levels of uncertainty salience (Dickinson, 2012; Gallagher & Holland, 1994; Hall, 1991). Ablation studies and drug manipulations suggest that disruption of the amygdala or dopaminergic midbrain (or, critically,

disrupting the integrity of amygdalar-midbrain connectivity) disrupt these effects, additionally supporting the functional role for the amygdala and dopamine signaling in supporting shifts in associative salience in response to surprise (P. C. Holland & Gallagher, 1993b; H. J. Lee, Gallagher, & Holland, 2010; H. J. Lee et al., 2005; H. J. Lee, Wheeler, & Holland, 2011). Imaging studies have also demonstrated evidence that, during reversal learning of cues paired with shock, observed BOLD signals in the amygdala is consistent with shifts in associative salience reflecting the putative mechanisms of the Pearce-Hall model (Li, Schiller, Schoenbaum, Phelps, & Daw, 2011).

However, extensive evidence in humans counters this approach, instead finding evidence in behavior and learning consistent with the role of effective predictive validity in driving shifts in salience, consistent with the predictions of the Mackintosh or Esber-Haselgrove models. Eyetracking data suggest that overt gaze during a category-learning task is greater to consistent exemplars than inconsistent stimuli (Le Pelley et al., 2011). (However, data from eyetracking reflect conflicting findings: (Hogarth, Dickinson, Austin, Brown, & Duka, 2008) found that gaze was more consistent with the predictions of the Pearce-Hall model.)

In comparing the attention captured by streams of stimuli that differ according to their predictability (i.e., one sequence of stimuli repeats according to a pattern, while three other sequences are randomly ordered), work in Nicholas Turk-Browne's laboratory suggests that, rather than responding to consistently unpredictable

information, participants' attention is instead preferentially captured by the predictable stream. Participants were faster and more accurate at identifying stimuli sharing properties with or occurring in the same visual space as the predictable (vs. unpredictable) stimuli (Zhao, Al-Aidroos, & Turk-Browne, 2013). Additionally, predictive stimuli were associated with an increased BOLD response in the hippocampus, which was related to an enhanced BOLD signal in occipital and parietal cortices that predicted increases in memory (Turk-Browne, Scholl, Johnson, & Chun, 2010).

As noted, current evidence provides conflicting accounts of the relationship between cue-outcome predictive validity, associative salience, and learning. Both schools of models, critically, provide a potential mechanism whereby environmental variability may induce exploratory behavior: the salience of a stimulus is determined, in large part, by its history in predicting reward. These models differ in their proposed mechanisms for determining and updating this salience value; however, in both cases the associative salience of a stimulus, whether reflecting its history of reliably predicting reward (its effective salience) or its history of unreliably predicting reward delivery (its uncertainty salience) can serve to increase arousal regarding the stimulus in question, enhancing exploratory behavior and increasing the recruitment of activity within the functional networks supporting active learning of cue-outcome relationships and subsequent long-term memory success.

The present study was designed to examine whether, and how, the acquired salience of reward-predictive cues reflected the certainty or uncertainty of the relationship between cues and their outcomes. Further, we intended to test the effects of this acquired salience on environmental exploration: how does the experience of certainty or uncertainty impact how individuals interact with other information presented incidentally in their environments? The presence of uncertainty may provoke individuals to more actively explore their environments to maximize information gains; conversely, the emergence of strong predictions, and not the experience of uncertainty, may bias individuals towards a more active relationship with incidentally-presented stimuli.

3.2. Hypotheses

In the present study, we hypothesized that attention to reward-associated cues, as measured through eyetracking, will vary as a function of those cues' predictive value, consistent with the predictions of the Mackintosh and Esber-Haselgrove models. Further, we hypothesize that participants will show preferential gaze to cues who serve as strong predictors of their outcomes, while showing less preferential gaze for less strongly predictive cues. Further, we hypothesized that exploration at encoding of incidentally-presented information will vary as a function of cue-outcome predictiveness.

3.3 Methods

3.3.A Participants

42 participants were recruited through email. All participants provided informed consent as approved by the Duke University Institutional Review Board. For 11 participants, satisfactory eyetracker calibration could not be achieved and their behavioral data are not included in the described analyses. Thus, 31 participants were included in the present analyses. (Mean age=23, 17 female.)

3.3.B Materials

Stimuli consisted of images presented on a computer screen. Images were of three categories: (1) five grayscale symbols composed of dots superimposed over squares; (2) scenes, color images of indoor and outdoor spaces collected from the MIT-CSAIL Database of Objects and Scenes; and (3) outcome images, which consisted of a one-dollar bill and a spatially-scrambled version of the dollar image. Scene images were selected to exclude humans, animals, famous landmarks, or written words or symbols. Image sets were counterbalanced across conditions and across participants.

3.3.C Experimental Design and Procedure

Participants completed the task over two days. On Day 1, participants were asked to complete the Encoding and Cue Probe tasks. Encoding was broken into five blocks, with six blocks of the Cue Probe task interspersed before and after each block.

Eyetracking data were collected during the Encoding task. Following a 24-hour delay, participants returned to the same testing location to complete subsequent-memory testing.

3.3.C.1 Session I: Encoding and Eyetracking

Participants were presented with images on a computer screen while eyetracking data were recorded. The task is summarized in Figure 3. On each trial, participants were presented with a cue, which was followed (after a delay) by either the delivery or omission of a fixed monetary reward. Participants were not required to generate an overt response in order to gain the reward; instead, the procedure relies of Pavlovian associative learning, or trace conditioning.

Four cues were used in this experiment, as shown in Figure 4. Each cue consisted of a compound dot-image: a control cue, which was repeated on every trial in each condition, and a condition-specific cue, which was unique to each condition. Control and condition-specific cues were displayed on the left and right of the screen, and were counterbalanced for left/right presentation.

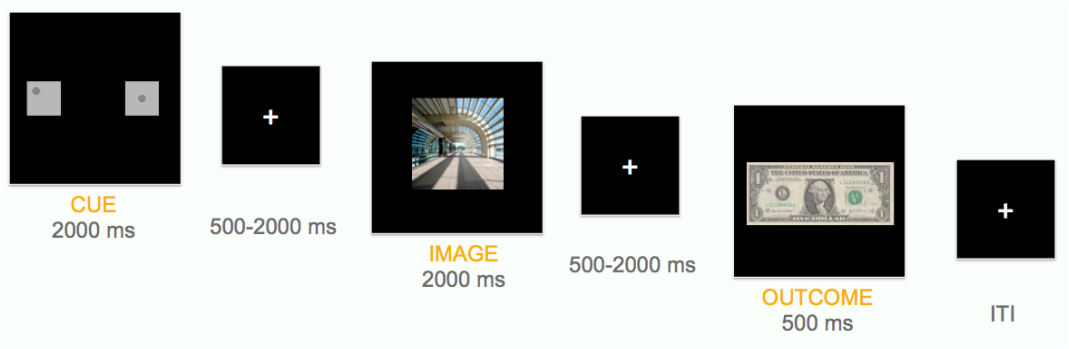

Figure 3: Trial structure for Study 2 (Eyetracking task.) Cues were associated with reward delivery at varying predictive strengths. During the delay period between cue offset and outcome onset, a trial-unique image of an indoor or outdoor scene was presented. Each cue was presented 8 times per block over 5 blocks (40 trials total.) Cues were counterbalanced for left/right presentation. No overt response was required for participants to receive monetary rewards; cue-outcome contingency awareness was probed at the end of each block using a Likert scale. Following a twenty-four hour delay, participants returned to the laboratory to complete a subsequent-memory recognition test for the scene images.

Critically, cues varied in the extent to which they consistently predicted reward delivery: they were either "Predictive" and were deterministically and consistently paired with reward delivery or omission, or were "Nonpredictive" and were not reliably paired with reward delivery or omission over the course of the experiment. The Predictive class of cues was split into two conditions: Predictive-Plus (P+), in which each appearance of the cue was paired with reward delivery; and Predictive-Minus (P-), in which each appearance of the cue was paired with reward omission. The nonpredictive class of cues was also split into two conditions: Nonpredictive-Flat (NPflat), in which 50% of cue appearances were paired with reward delivery, which 50% were paired with reward omission; and Nonpredictive-Walk (NPwalk), like the NPflat condition, featured an overall reinforcement rate of 50%, but, unlike the NPflat condition, was locally

predictive, shifting its reward-predictive strength over the course of the experiment. For the first four blocks, the NPwalk cue alternated between 75% and 25% predictive validity in reward outcomes; for the final block, it predicted reward delivery with a 50% reinforcement rate.

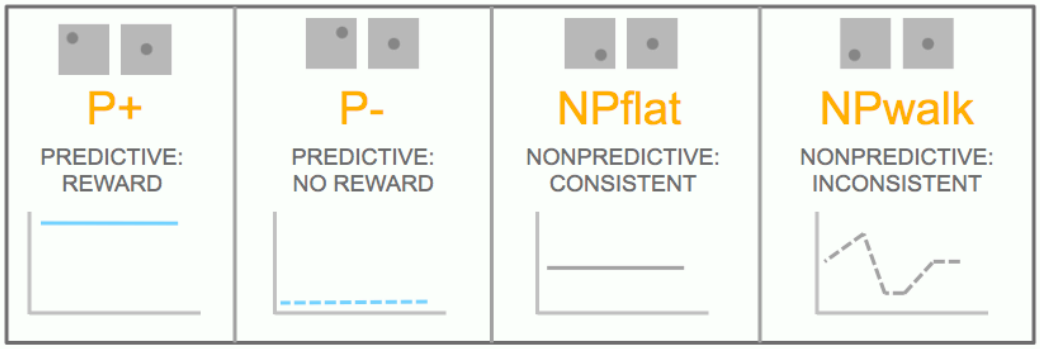

Figure 4: Cue-outcome contingencies for Study 2. Predictive-plus (P+) cues were paired with reward delivery 100% of the time; Predictive-minus (P-) cues were paired with reward omission 100% of the time. Nonpredictive-flat cues were reinforced at a constant 50% rate, while Nonpredictive-walk cues were reinforced an alternating 75%/25% rate over the first four blocks, and at a rate of 50% reward delivery in the final block. Cues were counterbalanced for left/right display.

To isolate effects of reward to encoding, monetary rewards generated through presentation of the cues was provided to participants at the end of Session I.

Participants were instructed that their goal in the task was to learn to predict whether, and to what extent, each cue predicted the delivery or omission of reward. Scenes were presented incidentally, and participants were not instructed to deliberately encode them.

Each cue was presented 40 times, over 5 blocks (8 times per condition per block), for a total of 160 trials. In subsequent memory testing, participants were presented with an equal number of “old” and “new” images, for a total of 320 scene images.

To gauge participants' ability to associate each cue with its outcome, we repeatedly probed participants' awareness of cue-outcome contingency. To do so, we displayed each cue (in isolation) with a Likert scale. Using the mouse, we asked participants to indicate, using the Likert scale, the likelihood of reward delivery. The minimum value of the scale corresponded to "Certain of reward omission" (and was labeled correspondingly) while the maximum value corresponded to "Certain of reward delivery." Participants' responses to the cue-outcome contingency probes were untimed. The probes were presented five times: once following each block.

In each trial, cues were displayed onscreen for 2000 ms. The outcome (dollar/scrambled dollar) was displayed for 1000 ms. Scenes were presented for 2000 ms; onset of these scene images was pseudo-randomly jittered within the fixed cue-outcome delay. During the delay period following each cue, but prior to reward delivery, a trial-unique scene image was presented (2000 ms). Cue presentation, image presentation, and reward delivery were each separated by brief periods of fixation: a variable 500-2000 ms delay period separated both the cue presentation from the scene image, and the scene image from the feedback delivery; for each trial, total intra-trial fixation time summed to 2500 ms.

Eyetracking data were collected using an Eyetribe eyetracker. Eyetracking data were collected at 60Hz. Prior to each run of the encoding task, the eyetracker was calibrated to within 0.5 degrees of visual angle.

3.3.C.2 Session II: Subsequent Memory Testing

Approximately 24 hours after encoding, participants were asked to return to the laboratory to complete a test of recognition memory and of source memory for the images they had encountered. These tests were structured as follows. First, we probed participants' recognition memory for the incidentally presented scenes they had viewed on the previous day. Participants were serially presented with each of the images they had viewed on the previous day ("old" images) alongside an equal number of novel foils ("new" images), and were asked to identify each image as old or new. Following their old/new categorization of each image, they were asked to provide an estimate for their confidence in their memory choice, using a scale of 1 ("Guess") to 3 ("Highly Certain"). Following this recognition-memory test, participants completed a source-memory test: each image was presented onscreen above a set of cues, and participants were asked to indicate which cue accompanied the given target image.

After completing memory testing, participants were paid for their time and debriefed regarding the goals of the experiment.

3.4 Analysis

3.4.A Behavioral Analysis

For each participant, Likert responses to each cue in each round of the Cue Probe task were scaled to generate estimates of reward likelihood. These values were then used to fit the Pearce-Kaye-Hall model for each participant, such that optimal parameters for each participant were selected to minimize the squared error between model-derived values of net associative strength and that participant's reported estimates of reward likelihood for each cue.

Memory was analyzed using standard subsequent-memory methods: for each participant, a Hit Rate (HR) and False Alarm Rate (FA) were calculated, according to their tendency to accurately identify old targets and mis-identify foils as targets. For each participant, a general FA rate was calculated, and then subtracted from each within-condition HR to generate a corrected HR for each condition. Participants whose average corrected HR was less than zero (indicating that they could not identify old vs. new images at a rate greater than chance) were removed from analysis.

Using participants' reported estimates of cue contingency, subject-specific parameters were calculated for the Pearce-Kaye-Hall (1982) model of associative learning. Of particular interest to the present study was the cue-specific learning rate, alpha, and the net associative strength (V_net). Uncertainty salience ("alpha") was determined according to the following equation:

$$\alpha = \gamma * abs(\lambda - \Sigma) + (1 - \gamma) * \alpha$$

where λ=magnitude of reward, γ=weighting factor for recent vs. more distant events, Σ= summed associative strength of present cues.

Following the outcome of each trial (reward delivery or omission), the shift in learning for each cue present at that trial is determined by the following equation:

$$\Delta V = \alpha * \beta * RPE\_abs$$

where α=uncertainty salience for given cue, β=intrinsic salience for given cue, RPE_abs= unsigned reward prediction error.

The "effective salience" of a cue, in contrast, was calculated not as a reflection of its uncertainty, but as the summed positive and negative associative strength for each cue.

3.4.B Eyetracking Analysis

Eyetracking data were analyzed based on the methods described in (Hannula et al., 2010). Blinks and other unusable timepoints were removed from consideration. Each timepoint was categorized as fixation or saccade (a timepoint was considered a "fixation" if its velocity was less than a visual angle of 1 degree, and this sub-threshold velocity had been maintained for at least 100 ms.) Mid-saccade timepoints were removed, and only fixations were considered in subsequent analysis. Sorted data were spatially smoothed prior to subsequent dwell-time or path-length analyses.

Dwell time measurements were calculated by expressing the amount of time a given criterion was met (e.g., that participants' gaze was inside or outside of an onscreen region of interest) as a proportion of total viable (non-saccade) timepoints for that trial.

To characterize the role of predictiveness in driving gaze behavior, we calculated dwell times to each cue for every trial. Regions of interest were defined based on the dimensions of the images displayed onscreen at each timepoint. For cue gaze analysis, relative dwell times were computed by subtracting the proportion of time spent looking at the cue-specific element vs. the repeated control element of each cue.

Gaze to scene images was characterized as an exploration rate. This was calculated as a path length--the total distance (in pixels) that each participant's gaze traversed over each image of interest--divided by the duration of viable (e.g., non-blink, non-saccade) timepoints for that trial. Exploration rates were compared across predictive conditions (P vs. NP.)

In addition to mean-based analyses, two mixed-effects linear regression models were used to test the relationship between behavior (including gaze behavior and subsequent memory), model-derived parameters of interest (e.g., uncertainty salience, effective salience, and prediction errors), and reward history over time. Parameters of interest, within-participant, were treated as fixed effects, while participants were treated as random effects. In model 1, we compared the relationship of uncertainty salience, effective salience, and prediction errors as fixed-effect predictors, with preferential cue

gaze as the outcome. In model 2, we compared the relationship of uncertainty salience, effective salience, and prediction errors, and exploration rate as fixed-effect predictors, with subsequent-memory confidence as the outcome.

3.5 Results

3.5.A Behavioral Results

Participants were able to distinguish the cues and to predict reward outcomes. Figure 5 illustrates the model-derived estimate of reward delivery for each cue over time.

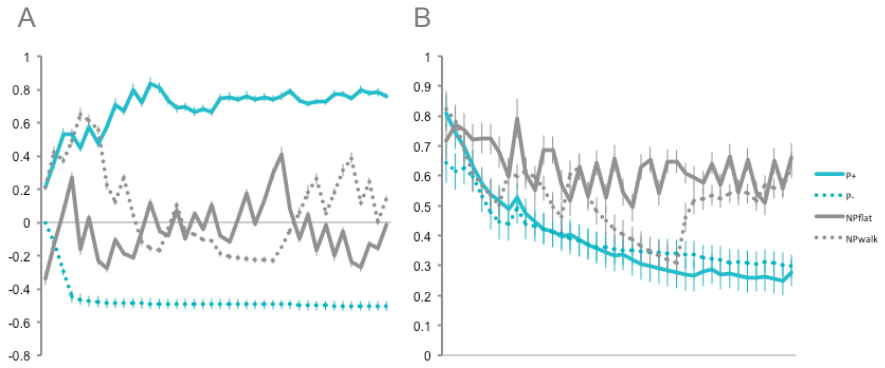

Figure 5: Group means for PKH model-derived estimates of (A) net associative strength for each cue over time, and (B) uncertainty salience for each cue over time. (Error bars correspond to +/- 1 standard error of the mean.)

As shown in Figure 6, subsequent memory for target images did not differ by condition (1-way ANOVA, p=0.66).

3.5.B Eyetracking Results

3.5.B.1 Gaze to Cues: Mean Comparisons

As shown in Figure 7, participants' dwell time to informative (e.g., condition-specific) cues was greater than to non-informative (control) cues. Further, this preferential dwell time for predictive cues (P+ and P-) (M=0.16, SD=0.15) was greater than for nonpredictive cues (NPflat and NPwalk) (M=0.12, SD=0.14); t(30)=2.24, p=0.03). Further, gaze to the P+ (M=0.17, SD=0.15) showed a trend-level significance toward greater than gaze to the NPf cue (M=0.13, SD=0.16); t(30)=2.03, p=0.051) and was significantly greater than gaze to the NPw cue (M=0.14, SD=0.13); t(30)=2.17; p=0.038.

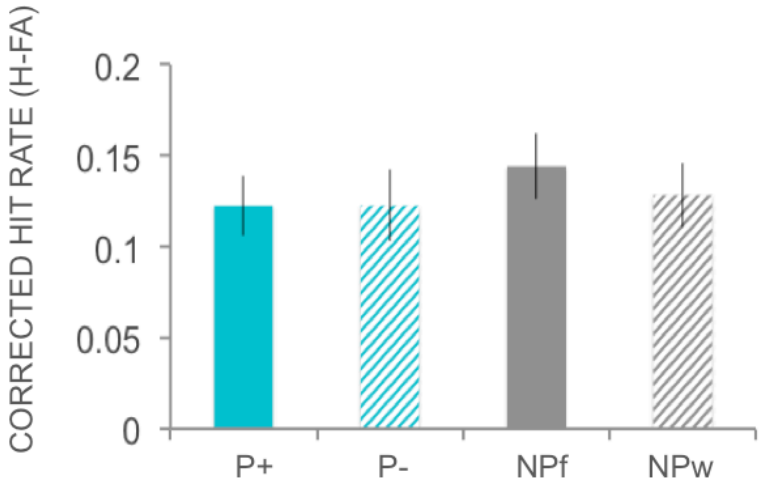

Figure 6: Subsequent memory for Study 2. Corrected hit rates (Hits – False Alarms) are plotted. These data included medium- and high-confidence responses, excluding guesses. Memory did not differ by condition.

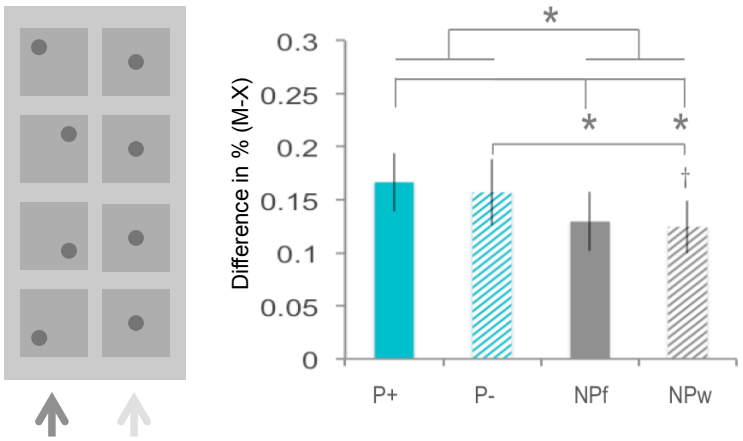

Figure 7: Gaze to cues. Plotted values are calculated as difference scores in proportion of time spent viewing the condition-specific cue minus the proportion of time spent viewing the control cue. Participants spent more time looking at Predictive vs. Non-Predictive cues (p=0.03).

3.5.B.2 Gaze to Cues: Regression Models

To investigate what signals may drive trial-to-trial shifts in gaze allocation, we tested whether model-derived estimates of parameters of interest predicted preferential dwell time to cues, calculated as the gaze difference score (dwell time to meaningful cues - dwell time to control cue). We treated uncertainty-salience, cue-outcome predictive validity or effective salience (the summed associative strength for each cue), or recent history of reward expectation violations (unsigned reward-prediction error for the preceding trial) as fixed effects, and participants as random effects. The model's adjusted R-squared was 0.1152. As shown in Figure 8, uncertainty salience was not a

significant predictor of gaze (p=0.89), while preceding-trial RPE significantly predicted gaze (p=0.03), and effective validity (unsigned net associative strength) predicted gaze at the trend level (p=0.07). Thus, these data suggest that gaze to cues is not sensitive to the uncertainty salience of a cue, but to its effective salience and to temporally local signals that expectations have been violated.

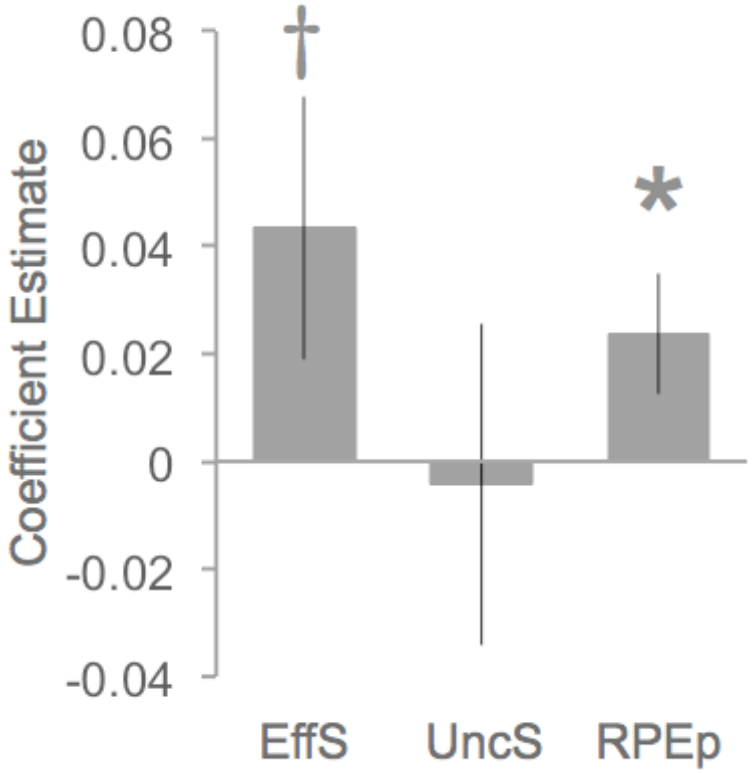

Figure 8: Coefficient estimates from linear mixed-effects model relating, for each trial, effective salience, uncertainty salience, and the preceding trial's signed reward prediction error with the dwell time to cue for the corresponding trial. Uncertainty salience did not predict gaze to cues, while effective salience was a significant predictor at the trend level (p=0.07) and the preceding trial's reward prediction error significant predicted gaze (p=0.03).

3.5.B.3 Gaze to Images: Hits vs. Misses

As shown in Figure 9, within-participants comparisons demonstrate that participants show more active exploration for images that were subsequently remembered ("Hits") than for images that were subsequently forgotten ("Misses"). Hits: M=4.51, SD=1.95; Misses: M=5.06, SD=1.70. t(30)=2.33, p=0.027.

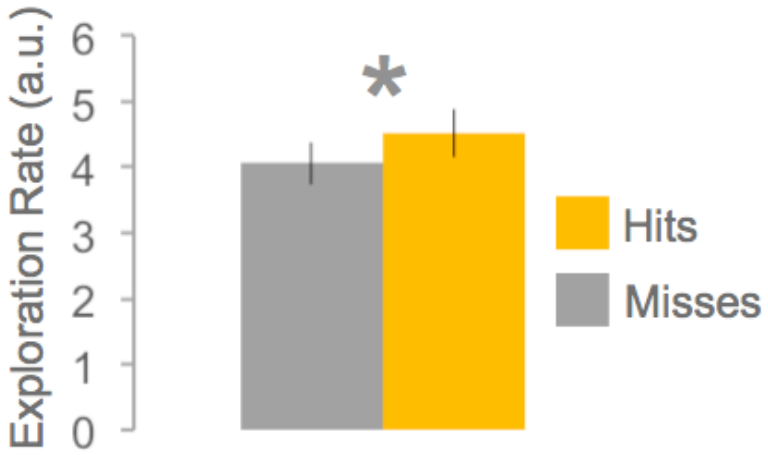

Figure 9: Exploration rate differs by memory outcome. Participants are significantly more active in viewing scene images that are subsequently remembered ("Hits") than images that are subsequently forgotten ("Misses") ($p<0.05$)

3.5.B.4 Gaze to Images: Subsequent Memory Confidence

A mixed-effects linear regression model tested, with memory confidence as the outcome variable, participants treated as a random effect, and effective salience, uncertainty salience, immediately-preceding RPE, and exploration rate as fixed effects, the relationship between variability in parameters of interest and encoding success. As shown in Figure 10, only exploration rate was found to significantly predict memory confidence ($p<0.05$).

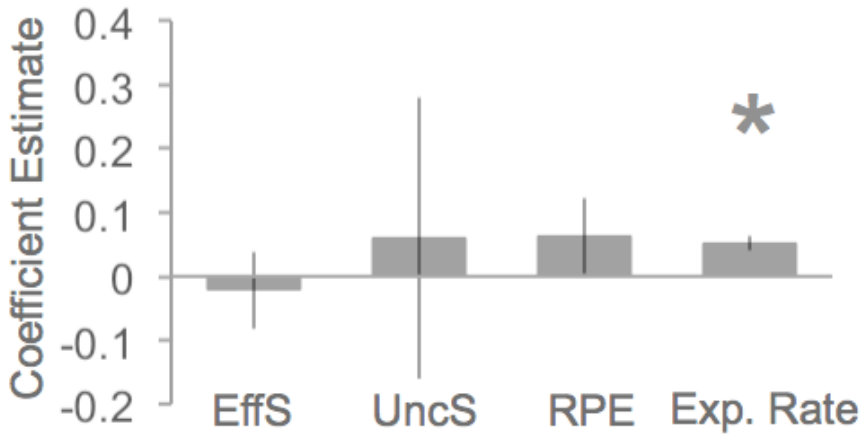

Figure 10: Coefficient estimates from linear mixed-effects model relating, for each trial, effective salience, uncertainty salience, the preceding trial's signed reward prediction error, and the rate at which each individual explored each scene image, with the subsequent-memory confidence for that image. Only exploration rate significantly predicted subsequent memory ($p<0.05$).

3.5.B.4 Gaze to Images: Exploration During Predictive vs. Non-Predictive Trials

As shown in Figure 11, although subsequent memory for images did not significantly differ across conditions, within-participants comparisons revealed that participants show more active exploration for images presented during Predictive (P+ and P-) (M=4.24, SD=1.73) than for Non-Predictive trials (NPwalk and NPflat) (M=3.99, SD=1.61); t(30)=3.08, p=0.004).

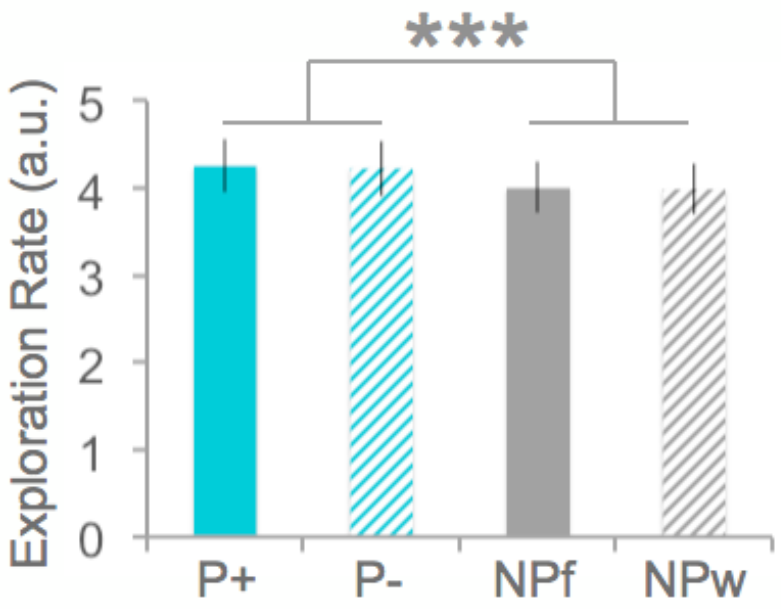

Figure 11: Participants show a greater exploration rate for scenes presented during Predictive trials vs. scenes presented during Non-Predictive trials (p<0.005).

3.6 Discussion

3.6.A Summary of Specific Hypotheses and Results

In Study 2, we investigated how individuals interacted with stimuli during states of variable expectancy, and how these interactions related to their learning and subsequent memory. By pairing symbols with rewards and varying rates and asking participants to learn whether, and how strongly, to predict reward delivery after viewing each cue, we attempted to induce differential levels of certainty regarding upcoming events; self-reported estimates of reward-delivery likelihood for each cue provide evidence that participants learned to distinguish the cues and to predict their outcomes. Additionally, and critically, we presented novel scenes during the delay

between cue offset and reward delivery or omission; we predicted that viewing of the scene images would differ as a function of participants' level of uncertainty regarding the outcome of the current trial. We found that participants generated more active exploration when viewing incidentally-presented scenes paired with predictive cues than incidentally-presented scenes paired with nonpredictive cues. These differences in exploration were predictive of subsequent memory success. However, we did not observe direct differences in subsequent memory between conditions. In a further consideration of eyetracking data during learning, we compared gaze to the reward-predictive (or non-predictive) cues during learning, in order to compare the predictions of associative models. Participants spent more time looking at all condition-specific cues than at the control cue, and, further, looked more at predictive cues than at non-predictive cues.

3.6.B Relationship to Prior Literature

Associative learning--binding one stimulus with another, or with a given response--is a core function of learning and memory, and a primary focus of cognitive science. Thus, a rich literature describing potential mechanisms supporting associative learning has developed, with sets of models making opposing, and overlapping, predictions regarding reward history, salience, and action. As discussed previously, the Pearce-Model (among others) posits that the associability, or degree of learning that a stimulus can attract (that is, the increment or decrement of the association between a

stimulus and an outcome, following each presentation of the stimulus and the delivery, or omission, of the outcome) is partly a function of its reward history. According to the Pearce-Kaye-Hall model, a stimulus whose outcome is not well known, or with which an animal has recently experienced a large amount of prediction errors, will have a higher associability than a stimulus whose outcome is known (and with which the animal has not experienced errors in reward prediction.) This prediction is contrasted by that of the Mackintosh and Esber-Haselgrove models, which instead relate associability (or associative salience) with the effective strength of a cue in predicting an outcome (whether positive, negative or neutral), potentially facilitating the generation of an action required to gain a reward (Esber & Haselgrove, 2011; Mackintosh.J., 1975).

The tension in these models reflects variable, competing, roles of attention in driving learning vs. regulating action (Maddux, Kerfoot, Chatterjee, & Holland, 2007). Our data are consistent with the predictions of the Mackintosh and Esber-Haselgrove models, which suggests that cue predictive validity is a strong driver of associative salience and supports encoding success. However, they do not serve to overturn the large body of evidence from animals (P. C. Holland & Gallagher, 2004; P. C. M. Holland, J.M., 2010) and humans (P. C. Holland, 2013; Li et al., 2011); one potential explanation for the lack of an observed role of uncertainty salience in driving behavior is that here, observed events are divorced from choice. If participants were required to generate a choice in order to gain reward, enhanced dopaminergic signaling of salience-related

prediction errors may enhance the motivational relevance of uncertainty and bias gaze towards non-predictive cues.

The eyetracking effects reported here relating to dwell time to predictive vs. non-predictive cues replicates the basic findings extant in the literature, primarily those of Le Pelley and colleagues (Le Pelley et al., 2011; Le Pelley et al., 2016) as well as the reported effects of regularities in attracting attention reported by Zhao and colleagues (Zhao et al., 2013). Consistency and predictiveness, not variability or unpredictiveness, seem overall to attract attention in measures of eyetracking during learning and serve as targets for changes in cue-outcome associations.

Our findings also address the nature of exploration behavior in learning as measured using eyetracking. Extant data have illustrated the strong relevance of eye gaze, as measured using fixations or dwell time, to active encoding. Such data strongly illustrate the value of eyetracking data in considering participants' behavior during learning and the factors that drive it: when variables such as familiarity are manipulated above and beyond task constraints, participants' eyegaze behavior serves to indicate their history of reinforcement even when such responses are not instructed and are counter to overt goals on each trial (Schwedes & Wentura, 2012). Gaze during encoding, specifically, is critically linked with memory success and with activity within the hippocampus and medial-temporal-lobe system which is critical for supporting the formation of long-term memories (Gabrieli, 1998; Hannula, Baym, Warren, & Cohen,

2012; Nickel, Henke, & Hannula, 2015). Our data confirm these findings, as participants' visual exploration of novel scenes predicted subsequent memory success; further, we extend our understanding of the relationship between visual exploration and learning success by noting that visual exploration differed as a function of predictive certainty experienced during encoding.

It is surprising to note that, although participants' dwell time when viewing predictive cues was greater than when viewing non-predictive cues, and exploration for scenes encountered during high-certainty prediction states was greater than for scenes encountered during low-certainty prediction states, we did not observe such differences when comparing reward-predictive cues vs. omission-predictive cues. Thus, the degree to which a cue served as a predictor of events, not the value of the events it predicted, served to drive gaze. This contrasts, in part, with findings suggesting that the magnitude of expected reward is a strong driver of the systems supporting long-term memory and dopaminergic function (Adcock et al., 2006; Wittmann et al., 2005). These effects are further considered in Study 3 and in the General Discussion.

3.6.C Caveats and Limitations

The present study investigated Pavlovian conditioning; participants were not asked to generate a particular response in order to gain a reward or to make a choice between competing outcomes. This simplified analysis and enhanced our ability to compare simple computational models of associative learning, but neglects to address a core

component of learning and choice in the real world. Adopting the experimental framework used here to employ an operant reinforcement task would expand the relevance to real-world contexts, and would potentially enhance the role of dopaminergic function in recruiting motivationally-relevant structures supporting choice and vigor for actions.

3.6.D Outstanding questions

The present study supports the perspective of the Mackintosh model with regards to the relationship between predictive validity and associative salience. However, the present study could not address the neural mechanisms supporting the observed effects. Multiple mechanisms could interact or compete to generate the set of learning and eyegaze behaviors observed in the present study. Therefore, in order to provide insight into neural mechanisms, we adapted the procedure in the present study to be performed using functional magnetic resonance imaging (fMRI). The imaging experiment and findings are discussed in the subsequent chapter.

4. Neural Systems Supporting Associative Learning

4.1 Introduction

In Experiment 2, we asked participants to learn to predict reward delivery when presented with cues with varying predictive strengths. We found that participants spent more time looking at cues with a stronger predictive relationship with their outcomes (regardless of whether that predicted outcome was the delivery or omission of reward), and that participants generated more active exploration when viewing scenes presented during predictive (vs. non-predictive) trials. In Experiment 3, we expand our scope of inquiry to include functional neuroimaging in order to characterize the neural systems involved in supporting these processes. Here, we test whether blood oxygen level dependent (BOLD) signal within the amygdala, hippocampus, and midbrain structures differs as a function of cue predictiveness, and whether these observed differences relate to learning in our participants.

4.1.1 Neural Systems of Interest

Data from animals and humans strongly implicate a set of neural structures in supporting associative learning, including the amygdala/extended amygdala, hippocampus and associated medial temporal lobe cortices, and the dopaminergic and noradrenergic concentrations of cells within the midbrain.

4.1.1.A Functional Role of Amygdala and Bed Nucleus of the Stria Terminalis

Evidence from human imaging and nonhuman experimental work has long suggested that the activity of neurons in the amygdala, though popularly associated with the processing of fear-related or otherwise overtly emotional stimuli (Phelps & LeDoux, 2005), is generally responsive to stimulus salience. PET imaging suggests that increased activity within the amygdala is related not just to negative images, but to positive images as well; additionally, activity increased following the presentation of non-emotional but highly visually engaging, or high-salience, images (S. B. Hamann, Ely, Hoffman, & Kilts, 2002). FMRI data confirm this interpretation, with increased salience in olfactory stimuli corresponding to increased BOLD response in the amygdala (Anderson et al., 2003). Additional functional imaging data from humans suggest that changes in the BOLD signal within the amygdala in response to changing salience are not related solely to the physical qualities of the stimuli involved (such as brightness or contrast), but to the interaction of the external qualities of the stimuli with an individual's current goals—that is, with a stimulus' motivational salience (Cunningham, 2012).

Both the Pearce-Hall and the Mackintosh models suggest that salience may be associated with cue-outcome learning via a cue-outcome contingency's rate of change—that is, when a stimulus becomes more (or less) salient reflects when an animal may learn about it (Mackintosh.J., 1975; Pearce & Hall, 1980). Evidence from work in rats

suggests that the activity of neurons within the amygdala relates not simply to the degree of sensory salience an animal accords a stimulus, but to mismatch between its predictions and its experience—that is, may represent a salience prediction error (H. J. Lee, Youn, O, Gallagher, & Holland, 2006). Lee et al. further found that severing the projections of the substantia nigra to the amygdala prevented learning, in a manner consistent with Pearce-Hall learning (H. J. Lee et al., 2006). However, data from animals and humans are also consistent with a role for the amygdala in signaling effective salience. Neurons within the rat basolateral amygdala show enhanced response rates to conditioned stimuli regardless of positive or negative outcome (S. C. Lee, Amir, Headley, Haufler, & Pare, 2016), while distinct neurons within the basolateral amygdala of rhesus macaques responded to cues associated both with positive and with aversive reinforcement (Paton, Belova, Morrison, & Salzman, 2006).

Additionally, data from humans and animals suggest that motivational signals transmitted by neurons in the amygdala may not solely be reflected in the activity or BOLD response within the amygdala, but also within the set of associated structures in the bed nucleus of the stria terminalis or BNST, also termed the "extended amygdala" (Avery, Clauss, & Blackford, 2016). Particularly, the relationship between individual differences and functionally relevant processes supported by the amygdala have been demonstrated to relate to activity within the extended amygdala: activity within the BNST has been found to vary with trait anxiety (Somerville, Whalen, & Kelley, 2010),

while differential responding to unpredictable vs. predictable threat stimuli has been localized to the BNST as well (Alvarez, Chen, Bodurka, Kaplan, & Grillon, 2011).

4.1.1.B Functional Role of Dopamine in Associative Learning

Dopamine is strongly associated with the anticipation of reward (Schultz, 1998): phasic activity which occurs initially in response to reward is transferred to the earliest predictive stimulus (Barto, 1998), while tonic firing is increased during the anticipation of a reward (Adcock et al., 2006). Animal models of associative learning have demonstrated that intact connectivity between the dopaminergic midbrain and the amygdala is necessary for updating representations in associative learning (H. J. Lee et al., 2010; H. J. Lee et al., 2005). Evidence from electrophysiological recording in cats demonstrates a strong relationship between the salient events and activity of neurons in the dopaminergic midbrain (Horvitz, Stewart, & Jacobs, 1997; Steinfels, Heym, Strecker, & Jacobs, 1983; Strecker & Jacobs, 1985). Dopamine neurons show a phasic increase to unexpected stimuli (Horvitz et al., 1997; Schultz, 1998), consistent with a potential role in broadcasting violations of expectancy.

One potential mechanism for dopamine may be in transmitting an SPE signal to the amygdala, facilitating the enhancement of stimulus salience in overt attention and behavior. Dopaminergic cells within both the SNc and VTA project directly to the amygdala (H. J. Lee et al., 2011; Loughlin & Fallon, 1983). The amygdala, in turn, may reciprocally impact the activity of the dopaminergic system through GABAergic

projections to the SNc (and to a smaller extent, the VTA, potentially through indirect projections) (Geisler & Zahm, 2005; H. J. Lee et al., 2011). Data from rats suggest that severing the dopaminergic projections to the amygdala disrupt associative learning (H. J. Lee et al., 2006). Additional studies indicate that dopaminergic signaling may impact the time rodents spend exploring novel objects, a potential index of motivational salience for novelty within an animal model (Powell, Paulus, Hartman, Godel, & Geyer, 2003).

4.1.1.C Relationship between Dopamine and Hippocampally-Dependent Learning

The expectation of reward can impact explicit learning in at least two ways, each of which can inform our understanding of the role of dopamine in curiosity. Memory is enhanced for stimuli encountered during the anticipation of a reward, and this enhancement of memory is associated with increased BOLD signal in both the midbrain and the hippocampus (Adcock et al., 2006; Wittmann et al., 2005). Electrophysiological evidence suggests that dopamine is necessary for the formation and maintenance of long-term potentiation (LTP) in the hippocampus (Frey et al., 1990; Otmakhova & Lisman, 1998), which receives projections from the VTA (Gasbarri, Packard, Campana, & Pacitti, 1994). Together, these findings suggest one way that dopamine directly impact learning: through direct modulation of hippocampal activity and synaptic strength.

Dopamine may also impact hippocampal memory indirectly, through its impact on motivation and exploration. Human learning in a spatial navigation task can differ

according to participants' motivational states: when learning under the anticipation of rewards, participants showed better spatial memory than when learning without reward anticipation (Murty et al., 2011). Striatal dopamine, in facilitating active exploration (Baldo et al., 2002) can thus potentially drive enhancements in hippocampus-dependent memory as well.

4.1.1.D Additional Role of Dopamine: Response to Novelty

When an animal is exposed to a novel, unexpected, or nonspecifically high-salience stimulus, its attention can be captured, as it automatically orients towards the cue (Sokolov, 1963). This orienting response classically includes overt foveation towards the relevant cue, as well as enhanced attention towards the cue across a variety of modalities (Sokolov, 1963). Evidence from rats indicates that the dopaminergic midbrain is necessary for exogenous stimuli to provoke an orienting response (Ljungberg & Ungerstedt, 1976). The hippocampus has also been shown to be highly responsive to novelty, potentially reflecting the function of a "novelty detection" circuit relying on connectivity between the hippocampus and dopaminergic midbrain (Lisman & Grace, 2005). However, recent neuroimaging findings have suggested that this hippocampal response to novelty can habituate over time: as increasing expectation of novelty emerges from repeated exposure to novel stimuli diminishes the novelty-related enhancement of the BOLD response (Murty, Ballard, Macduffie, Krebs, & Adcock, 2013). Thus, learned predictions of novelty may engage similar dopaminergic-hippocampal networks as in dopamine-mediated facilitation of long-term memory.

4.1.1.E Potential Role of Norepinephrine

The neuromodulator norepinephrine (NE) has been shown to impact a variety of functions directly relevant to both attending to stimuli and responding to environmental variability: phasic increases in noradrenergic activity are associated with identifying relevant stimuli and inducing switches of attention, as an animal disengages with a stimulus and re-engages with a different target (Aston-Jones et al., 1991; Aston-Jones et al., 1999; Aston-Jones et al., 1994; Grant et al., 1988). Additionally, increased baseline or tonic levels of activity within the LC, producing an increased baseline level of NE in its efferent targets, including prefrontal cortex (Florin-Lechner et al., 1996) and the amygdala and extended amygdala/BNST (Asan, 1998), is associated with a broadening of attention (Berridge & Waterhouse, 2003). With increased tonic firing associated with increased states of arousal, a particular stimulus is less likely to sustain attention, as attention is more likely to be globally distributed and biased towards rapid shifts (Aston-Jones et al., 1999).

One potential role for NE in the modulation of stimulus salience may be in diminishing the potential associability of objects under conditions—such as elevated arousal--that evoke high levels of NE. As NE levels increase, the ability for a given stimulus to capture attention will decrease, diminishing its associability and decreasing learning according to salience-driven associative learning models (Aston-Jones et al., 1999; Mackintosh.J., 1975; Pearce & Hall, 1980). This is not necessarily the case for highly

motivationally salient stimuli—e.g., explicit threat cues—which have been shown to increase in salience under conditions of high arousal according to theories of arousal-biased competition (Mather & Sutherland, 2011). However, the mechanisms underlying learning in response to highly salient cues and active exploration may differ: while active exploration serves to sample the relation between target stimuli and the environment as a whole, arousal-biased competition predicts the monopolization of attention with a single, highly salient cue which determines the coordination of a response (e.g., escape) (Mather & Sutherland, 2011). Indeed, consistent with (Aston-Jones et al., 1999), investigations into arousal-biased attention in human participants found that performance is impaired when choice is dependent on the monitoring of multiple stimuli distributed within an environment (Morelli, 2009). The ability to selectively modify a stimulus' learned predictiveness is dependent not just on the ability to focus attention on that target stimulus, but also to monitor the associative strength of other objects present (Mackintosh.J., 1975; Pearce & Hall, 1980; Rescorla, 1972).

4.1.2 Interpreting Potential Network Interactions

We focused analyses on the defined regions based on findings from humans and animals implicating amygdalar-midbrain connectivity in supporting associative learning, and numerous findings demonstrating that encoding-dependent activity within the hippocampus was sensitive to modulation via activity within the amygdala (Dougal, Phelps, & Davachi, 2007; Murty et al., 2011; Phelps, 2004) as well as the

dopaminergic midbrain (Adcock et al., 2006; Shohamy & Adcock, 2010; Wittmann et al., 2005). We anticipated that these regions would show consistent responses to learning during the task: as cues acquired salience through their presentation and subsequent delivery or omission of rewards, we anticipated that BOLD signal within the amygdala and dopaminergic midbrain would vary according to the nature of the acquired salience, up- or down-modulating hippocampal activity as determined by the trial-appropriate response. The critical question, then, is what form acquired salience may take: whether increasing predictive validity predicts increased acquired salience, consistent with the concept of "effective salience," or whether repeated presentation with non-predictive cues enhances their acquired salience, consistent with "uncertainty salience." Thus, we intended to characterize the nature of these systems' responses to learned predictiveness vs. learned non-predictiveness.

4.2 Hypotheses

We hypothesized that cue-outcome learning will be associated with increased signal in our ROIs, primarily amygdala and dopaminergic midbrain. We anticipated that this would be reflected both in main effects/contrasts and in heightened connectivity. Mechanistically, we anticipated that trial-by-trial shifts in salience will be reflected in trial-by-trial shifts in BOLD signal within our network of interest. Further, we hypothesized that these shifts in salience will reflect emerging expectancy regarding a

predictive cue-outcome contingency, and not its uncertainty or history of unpredictiveness.

4.3 Methods

4.3.A Participants

Twenty participants (15 female; mean age = 26; range = 19-35) were recruited from the local community through email and the Brain Imaging and Analysis Center at Duke University. All participants were right-handed, had normal or corrected-to-normal vision, and were screened for contraindications to MRI or history of psychological or neurological illness. All participants provided informed consent as approved by the Duke University Institutional Review Board.

4.3.B Materials

Stimuli consisted of images projected onto a screen or displayed on a computer screen. Images were of three categories: (1) five grayscale symbols composed of dots superimposed over squares; (2) scenes, color images of indoor and outdoor spaces collected from the MIT-CSAIL Database of Objects and Scenes; and (3) outcome images, which consisted of a one-dollar bill and a spatially-scrambled version of the dollar image. Scene images were selected to exclude humans, animals, or written words or symbols. Image sets were counterbalanced across conditions and across participants.

4.3.A Design and Procedure

Task design was highly similar to that described above in Chapter 3. Participants completed the task over two days. On Day 1, participants were asked to complete the Encoding and Cue Probe tasks, while undergoing functional magnetic resonance imaging (fMRI). Encoding was broken into five blocks, with six blocks of the Cue Probe task interspersed before and after each block. Imaging data were collected during the Encoding task and during rest. After scans were completed, they received compensation for rewards accrued during scanning and filled out psychological questionnaires measuring anxiety and to depression. Following a 24-hour delay, participants returned to complete subsequent-memory testing.

4.3.A.1 Session I: Imaging and Encoding

As shown in Figure 12, task design was highly similar to that described above in Chapter 3. Participants viewed cues that were associated with rewards with varying validity, and were asked to learn to predict reward delivery or omission in response to each cue. As above, cues were "Predictive" (P+ and P-, deterministically and consistently paired with reward delivery and omission, respectively) or "Non-Predictive," which were paired with reward delivery on 50% of trials and reward omission on 50% of trials. An image was presented during the cue-outcome delay, but participants were not explicitly instructed to encode this image or informed that their memory for these images would be tested.

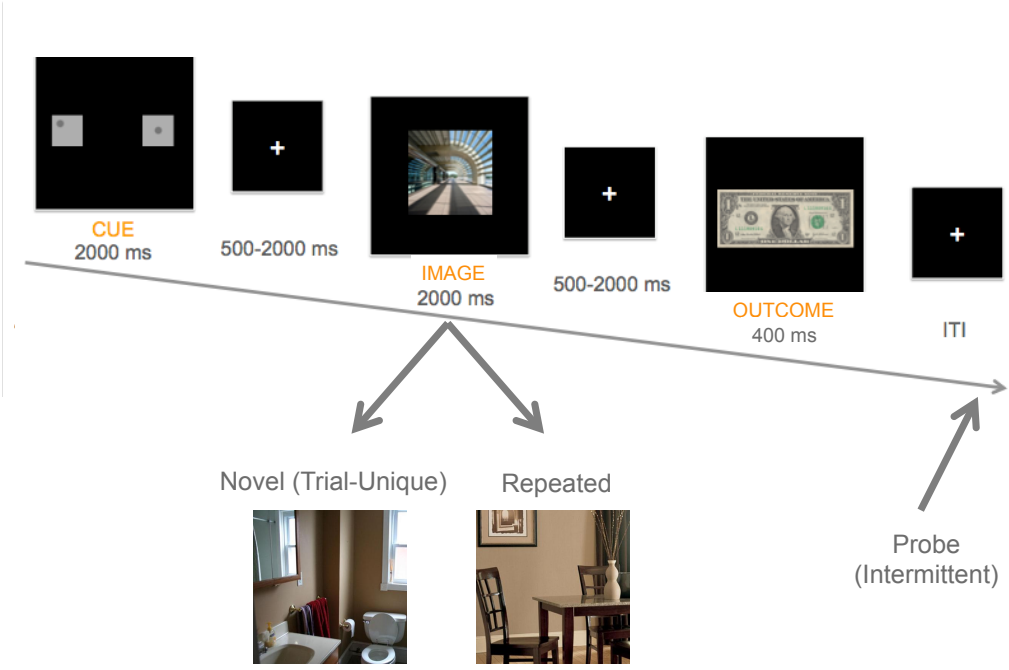

Figure 12: Trial schematic for Study 3. Cues were associated with reward delivery at varying predictive strengths. During the delay period between cue offset and outcome onset, an image of an indoor or outdoor scene was presented. For "Novel" trials (P+, P- and NP_novel), this image was a trial-unique indoor or outdoor scene. For NP_familiar trials, this image was a repeated, pre-familiarized scene. Each cue was presented 10 times per block over 5 blocks (50 trials total.) Cues were counterbalanced for left/right presentation. No overt response was required for participants to receive monetary rewards; cue-outcome contingency awareness was probed intermittently during the blocks, at the end of each block, using a Likert scale. Following a twenty-four hour delay, participants returned to the laboratory to complete a subsequent-memory recognition test for the scene images. Participants underwent functional imaging while completing the learning task on Day 1, but were not scanned during the memory task on Day 2.

The present task differs from that described in Chapter 3 with regards to novelty. In both the P+ and P- conditions, as in Chapter 3, a novel, trial-unique scene image was presented during the cue-outcome delay period. The two Non-Predictive cues were

paired equally with reward over time (monetary reward was delivered on 50% of trials) but differed with respect to the novelty of the scene images with which they were presented. In one non-predictive condition, NP_nov, a non-predictive cue was paired with a novel, trial-unique image on each trial. In the other non-predictive condition, NP_fam, a second non-predictive cue was paired with a repeated, familiar image on each trial. Thus, each Predictive cue (P+ and P-) consistently predicts reward delivery and novelty for that trial, while one NP cue (NP_nov) does not predict reward delivery but is predictive of novelty, while the final cue (NP_fam) does not predict reward delivery but is predictive of familiarity. Participants were pre-familiarized with this image prior to entering the scanner.

Participants were instructed that their goal in the task was to learn to predict whether, and to what extent, each cue predicted the delivery or omission of reward. Scenes were presented incidentally, and participants were not instructed to deliberately encode them.

Each cue was presented 50 times, over 5 blocks (10 times per condition per block), for a total of 200 trials. In subsequent memory testing, participants were presented with an equal number of “old” and “new” images, for a total of 300 scene images.

To gauge participants' ability to associate each cue with its outcome, we repeatedly probed participants' awareness of cue-outcome contingency. To do so, we

displayed each cue (in isolation) with a Likert scale. Using the button box, we asked participants to move a vertical hash mark to indicate, using the Likert scale, the likelihood of reward delivery. The minimum value of the scale corresponded to "Certain of reward omission" (and was labeled correspondingly) while the maximum value corresponded to "Certain of reward delivery." Participants' responses to the cue-outcome contingency probes were untimed. Reward contingency awareness for each cue was probed sixteen times: once at the beginning of the task, after each block, and twice within each block.

In each trial, cues were displayed onscreen for 2000 ms. The outcome (dollar/scrambled dollar) was displayed for 400 ms. Scenes were presented for 2000 ms; onset of these scene images was pseudo-randomly jittered within the fixed cue-outcome delay. During the delay period following each cue, but prior to reward delivery, a trial-unique scene image was presented (2000 ms). Cue presentation, image presentation, and reward delivery were each separated by brief periods of fixation: a variable 500-2000 ms delay period separated both the cue presentation from the scene image, and the scene image from the feedback delivery; for each trial, total intra-trial fixation time summed to 2500 ms.

The experiment used a fast event-related design; trial onset times were pseudorandomly jittered, optimized by OptSeq2. Participants' heartbeat and respiration were monitored while they were in the scanner in order to control for the effects of

physiological noise in subsequent imaging analysis. Following the completion of the task, participants completed the Beck Depression Inventory and the State-Trait Anxiety Inventory.

4.3.A.1 Session II: Memory Testing

Following completion of Session I, participants returned approximately twenty-four hours later for memory testing. On a computer screen, participants were shown each of the trial-unique scene images they had been shown during Session I, alongside an equal number of novel foils. Using the keyboard, participants were asked, first, to indicate whether they thought each image was "old" or "new." Following this response, they were asked to indicate their confidence in their memory judgment using a three-point scale. Memory testing was not timed, and was performed outside the scanner.

4.3.B Image Acquisition and Preprocessing

fMRI data were acquired on a 3.0-T GE Discovery MRI scanner using two sets of standard echo-planar imaging (EPI) sequences. The first set of EPI sequences were used during the five functional runs and included coverage across a section of the brain including temporal lobe and midbrain structures (TE = 27 ms, flip = 77 degrees, TR = 1 s, 19 contiguous slices, size = 3.75 mm x 3.75 mm x 3.8 mm), and the second set of EPI sequences were used during the single resting run and included coverage across the whole brain (TE = 27 ms, flip = 77 degrees, TR = 2 s, 34 contiguous slices, size = 3.75 mm x 3.75 mm x 3.8 mm). Each functional run consisted of 464 volumes, and the resting run

consisted of 150 volumes. Prior to the functional runs, we collected whole-brain high-resolution anatomical images (voxel size = 1mm isotropic) for use in spatial normalization and co-registration. The first 12 volumes of the functional runs (12 s) and the first 3 volumes of the resting run (6 s) were discarded to allow for signal saturation. Given that the functional runs were collected using a TR of 1 s, we did not collect data from the whole brain, but instead focused our acquisitions on slices that allowed us to collect data from midbrain and the medial temporal lobes, extending into the frontal lobe. Critically, this coverage did not extend to the entire dorsal striatum. The group coverage map is shown in Figure 13.

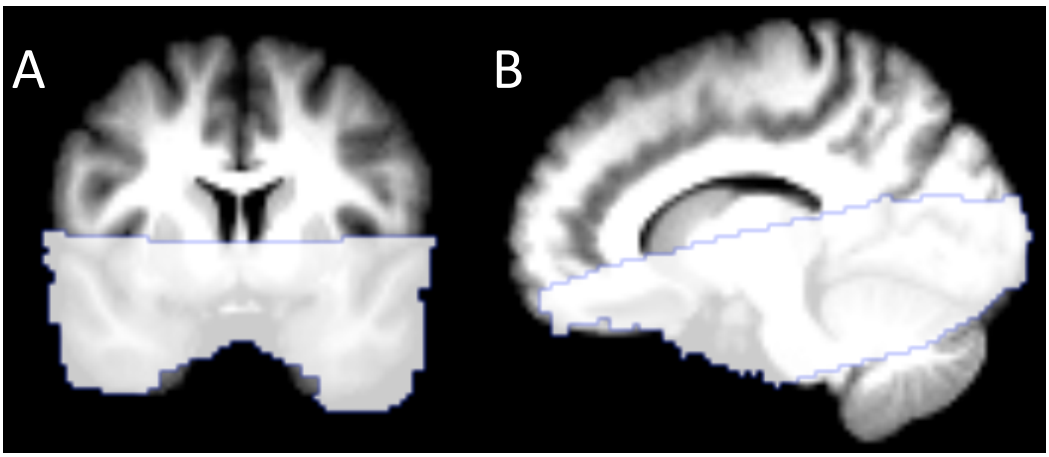

Figure 13: (A) Coronal and (B) Sagittal display of coverage for data collection in functional runs of Study 3. Note that while the midbrain, medial temporal lobes, and ventral striatum are included, much of dorsal striatum is excluded.

fMRI preprocessing was performed using fMRI Expert Analysis Tool (FEAT) version 6.0 as implemented in FSL 5.0.8. BOLD images were skull stripped using the

Brain Extraction Tool. Images were realigned within run, intensity-normalized by a single multiplicative factor, spatially smoothed with a 4 mm full width half maximum (FWHM) kernel, and subjected to a high-pass filter (100 s).

Spatial normalization was performed using a two-step procedure on fMRIb Linear Registration Tool (FLIRT). First, mean EPIs from each run were co-registered to the high-resolution anatomical image. Then the high-resolution anatomical image was normalized to the high-resolution standard space image in Montreal Neurological Institute (MNI) space using a nonlinear transformation with a 10-mm warp resolution, as implemented by fMRI Nonlinear Registration Tool (FNIRT). All coordinates are reported in MNI space.

In order to correct for potentially confounding physiological signals within the data, we analyzed peaks in heart rate and respiration recorded from participants during the task and included these as a regressor in our first level models.

4.3.C Behavioral Analysis

Behavioral analyses, both for subsequent-memory analyses and model fitting, were performed as described above in Chapter 3. Guesses were excluded from memory analyses, which prioritized medium- and high-confidence responses.

4.3.D Imaging Analysis

FMRI data were analyzed using FEAT version 6.0 as implemented in FSL 5.0.8. Three sets of analyses were conducted.

To investigate task-related activations, first-level (within-subject, within-run) general linear models (GLMs) included regressors of interest corresponding to each condition and baseline (ITI), and run-specific regressors corresponding to motion and physiological noise. In separate analyses, we modeled responses to the cue (duration: 2 s), image (duration: 2 s), outcome (duration: 0.4 s), and the whole trial (duration: 6.5 s). The events were modeled using an amplitude of 1, and were convolved with a double-gamma hemodynamic response function. Individual maps of parameter estimates were estimated for several contrasts of interest: task > baseline (at each epoch) P > NP (at cue), P > NP (at outcome), P+ > P- (at cue), P+ > P- (at outcome), Predictive (P+ and P-) > NP_nov) (whole trial), Remembered > Forgotten (at image), Remembered > Forgotten (at cue), Novel > Familiar (at image).

To test the relationship between BOLD signal within regions of interest and single-trial parameters of interest, trials were collapsed across conditions and first-level (within-subject, within-run) general linear models (GLMs) included parametric regressors corresponding to model-derive parameters of interest. In separate analyses, we modeled responses to the epoch of interest. To control for the main effect of activation related to an event and detect variability relating to our parameters of interest, each event was modeled using an amplitude of one as well as an amplitude corresponding to the demeaned model-estimated value for each trial. Individual maps of parameter estimates were estimated for parameters of interest:

Given that data from animals and functional imaging in humans have suggested that the neural mechanisms supporting associative learning may rely not simply on the magnitude of activation, but additionally on co-activation or intact signaling pathways between distinct regions, we examined connectivity during learning. Using psychophysiological interaction (PPI) analyses, we tested whether connectivity differed as a function of task conditions. We identified seed regions using conjunctions of functional activations with anatomically defined a priori regions of interest, including the bilateral amygdala, dopaminergic midbrain (including VTA and SN), orbitofrontal cortex, and the locus coeruleus. Seed ROIs were binarized, and timecourses were extracted from each run for each participant. To characterize the role of our ROIs in driving functional connectivity, we first ran base models on epochs of interest, as described above. We then ran secondary models including as regressors the timecourse, the main task regressors, as well as interaction terms between the timecourse and each of the conditions. Each condition in the base model was also included in the PPI. Second and third level analyses proceeded as described above.

Second level analyses for each contrast (combining across runs, within participant) were modeled using a fixed effects analysis. Third-level analyses (across-subjects) were modeled using FSL's mixed effects analyses (FLAME 1+2), which accounts for within-session/subject variance calculated at the first and second levels, on the parameter estimates for contrasts of interest derived from the second-level analysis.

To test whether, and to what extent, individual differences in anxiety may predict functional differences in neural activity or connectivity, we included individuals' (demeaned) trait anxiety scores as a group-level regressor. To investigate the main effects, two-way repeated measure ANOVAs were run in FSL.

4.4 Results

4.4.A Behavioral Results

4.4.A.1 Cue-Outcome Contingency Learning

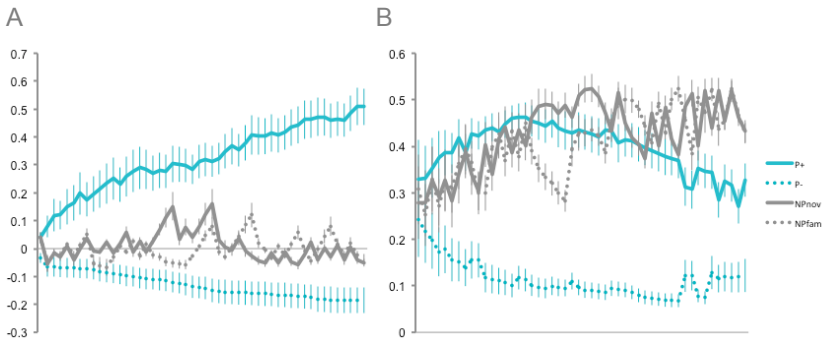

Figure 14: Group means for PKH model-derived estimates of (A) net associative strength for each cue over time, and (B) uncertainty salience for each cue over time. (Error bars correspond to +/- 1 standard error of the mean.)

As shown in Figure 14, participants learned to associate each cue with its approximate reward value, distinguishing P+ and P- cues from each other and from the Non-Predictive cues.

4.4.A.2 Subsequent Memory

When medium- and high-confidence responses are considered, memory for scenes paired with both P+ and P- cues is facilitated relative to scenes paired with Non-

Predictive cues. (P+: M=0.22, SD=.12; NP: M=0.19, SD=0.11; t(19)=2.3; p=0.03. P- : M=0.23, SD=0.14, t(19)=2.43, p=0.02. However, there is no difference in memory scores between P+ and P- conditions (t(19)=0.56, p=0.58. When restricting analyses only to high-confidence responses, memory is significantly greater for scenes presented during reward anticipation (e.g., during P+ trials) than either for scenes presented during P- and Non-Predictive trials. (P+: M=0.075, SD=0.049; P-: M=0.051, SD=0.062; NP: M=0.042, SD=0.041. P+ vs P-, t(19)=2.77; p=0.01. P+ vs. NP, t(19)=4.93, p<0.005.) There is no difference between recognition-memory performance for images presented during P- and NP trials (t(19)=0.80, p=0.43).

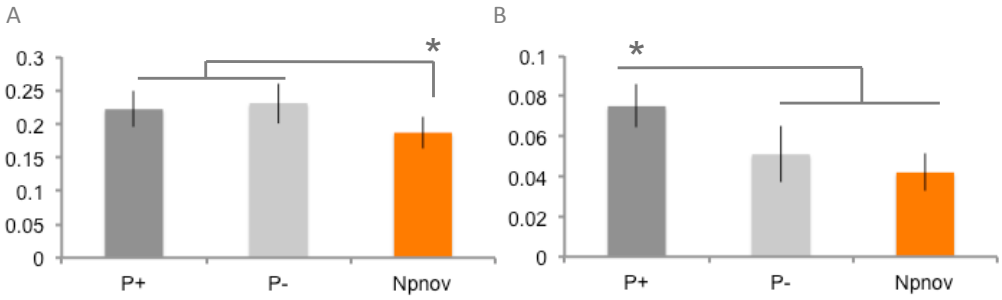

Figure 15: Subsequent memory (Corrected Hit Rate) by condition. (A) Medium- and High-Confidence responses (excluding Guesses.) Subsequent memory for scenes presented in both Predictive trials is facilitated relative to memory for scenes presented in Non-Predictive trials. (B) High-Confidence responses only. Memory is selectively facilitated for images presented during Predictive-Plus trials, during reward anticipation, relative to memory for images presented during Non-Predictive or Predictive-Minus trials.

4.4.B Imaging Results

We were interested in how our task, and the signals supporting associative learning over time, would be represented across the whole brain, as well as within our a

priori network of interest. A priori regions of interest were as follows: amygdala and its associated output regions within the extended amygdala/BNST; hippocampus; perirhinal, parahippocampal and entorhinal cortices, collectively referred to as MTL cortex; orbitofrontal cortex; dopaminergic midbrain (including both VTA and SN), as well as its efferent projections in the nucleus accumbens; and locus coeruleus. We had no a priori assumptions regarding laterality.

Our epochs of interest differed according to the specific questions we asked. In order to characterize BOLD responses consistent with forming and updating cue-outcome associative learning, we focused analyses primarily on the "cue" and "outcome" epochs, in order to identify activations related to reward anticipation and outcome-driven learning, respectively. To probe activations related to the successful (or unsuccessful) incidental encoding of scene images, we focused analyses on the "image" epoch (as well as probing the "cue" epoch to test for anticipatory effects on memory.)

4.4.B.1 Main Effects

4.4.B.1.1 BOLD Response to Cue

We identified a large number of voxels that showed greater activation in response to the cue than at baseline, as shown in Figure 16. These include activations within bilateral amygdala, hippocampus and MTL cortex, and orbitofrontal cortex, as well as extensive activations within regions supporting visual processing. However, no clusters within the midbrain or accumbens show a positive response. Instead, these

regions (among others, including the bilateral insula) show the opposite pattern, generating a greater response to fixation than that evoked by the cue.

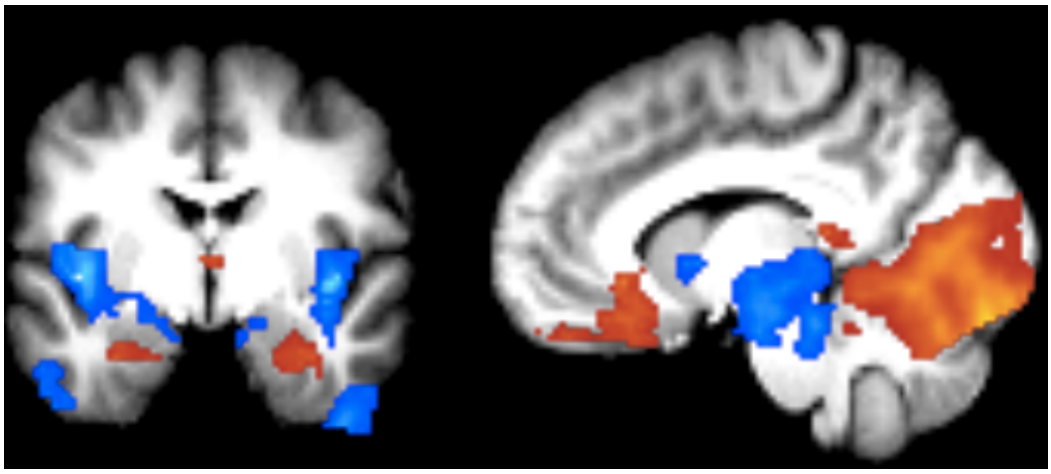

Figure 16: Cue-related activation vs. baseline. Cue > baseline depicted in red; Cue < baseline depicted in blue.

4.4.B.1.2 BOLD Response to Cue: Condition-Specificity

Having determined that our task engages our systems of interest, we wished to probe this response further to determine whether response to cue presentation differs across conditions. In response to the cue presentation, no clusters surviving correction for multiple comparisons showed a selectively greater response during Predictive trials than to Nonpredictive trials.

4.4.B.1.3 Main Effects of Predictive vs. Non-Predictive Cues

In order to identify whether structures of interest showed a greater response to predictable reward delivery/omission vs. unpredictable reward delivery, we compared activations at outcome for P (P+ and P- trials > NP (NP_nov and NP_fam). As shown in Figure 17, we found that BOLD signal at outcome was greater for P vs NP trials within

our network of interest: amygdala, orbitofrontal cortex, hippocampus and MTL cortex, as well as visual cortex. Interestingly, we found no significant activations within the dopaminergic midbrain, but did identify a functional activation consistent with the anatomical location of the locus coeruleus, consistent with a hypothesized modulatory role for noradrenaline in guiding attention for learning. We identified the local maxima within this cluster and generated spherical ROIs at these coordinates to probe connectivity.

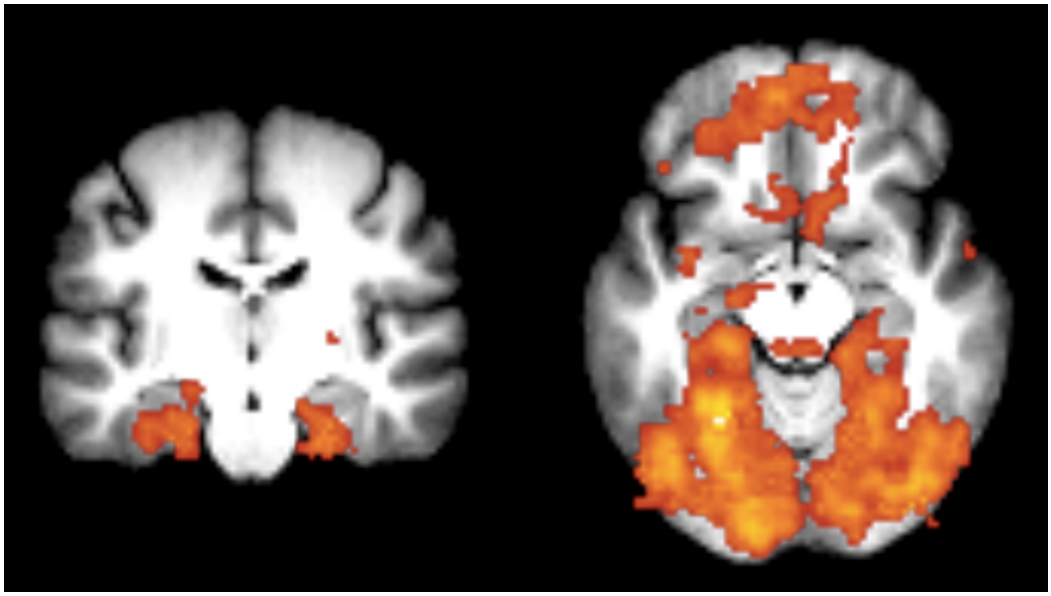

Figure 17: Activation at outcome, Predictive trials > Non-Predictive trials.

4.4.B.1.4 Reward Delivery

To determine whether the effect described above was driven simply by reward delivery in the P+ condition, we modeled the outcome epoch and tested for differences between the P+ (reward delivered) and P- (reward omitted) conditions. No clusters within our regions of interest showed a significant effect; activations within visual

processing regions showed an effect of P+ > P-, and no clusters showed an effect of P- > P+.

4.4.B.1.5 Encoding Success vs. Failure

To identify voxels supporting encoding into long-term memory, we labeled trials as "hits" or "misses" according to each participants' subsequent memory, and contrasted BOLD signal in response to the images for Hits > Misses. As anticipated, we identified a cluster within the left medial temporal cortex/hippocampus that showed selectivity for encoding success (greater activity for Hits > Misses).

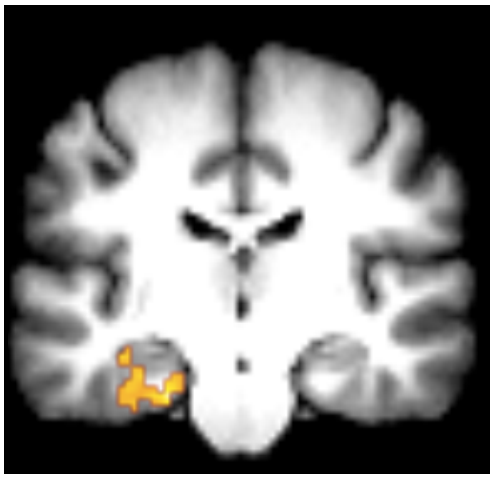

Figure 18: Response to images: Hits > Misses (Subsequently remembered > subsequently forgotten)

4.4.B.1.6 Anticipatory Effects on Encoding

To determine whether anticipatory effects on image encoding could be detected prior to image presentation, we modeled the "cue" epoch and contrasted Hits > Misses (sorting each trial according to the memory success for its associated image.) No clusters showed a significant effect, either for Hits > Misses or Misses > Hits.

4.4.B.2 Functional Tests of Model-Derived Parameters

The contrasts described above suggest that structures of interest, including the amygdala and medial temporal lobe memory systems, are implicated in our task and may support a role in associative learning. However, to characterize the precise mechanisms that they support, and whether their activation over time is consistent with theoretical learning signals, requires more computationally sophisticated analyses. To this end, we collapsed across different trial types and tested the significance of parametric regressors, corresponding to model-derived estimates of parameters of interest for each trial, in predicting BOLD signal within regions of interest.

4.4.B.2.1 Expected Value

To test whether the BOLD signal within our regions of interest tracked the (signed) magnitude of anticipated monetary rewards over the course of each trial, we included the net associative strength for cues present at each trial as a parametric regressor while modeling the "whole trial" epoch. No clusters within our regions of interest showed significant activation consistent with the magnitude of anticipated reward.

4.4.B.2.2 Effective Salience

To test whether cue-outcome certainty (or effective salience), rather than expected reward, was consistent with activity within our regions of interest, we modeled the "whole trial" epoch and included the summed associative strength between each cue and its outcome as a parametric regressor. Thus, the emerging certainty of

reward delivery for P+ trials was assigned a value comparable to the emerging certainty of reward omission for P- trials, in contrast to the consistently low certainty of reward for NP trials. We found significant activation associated with cue-outcome certainty within bilateral MTL cortex and visual processing regions.

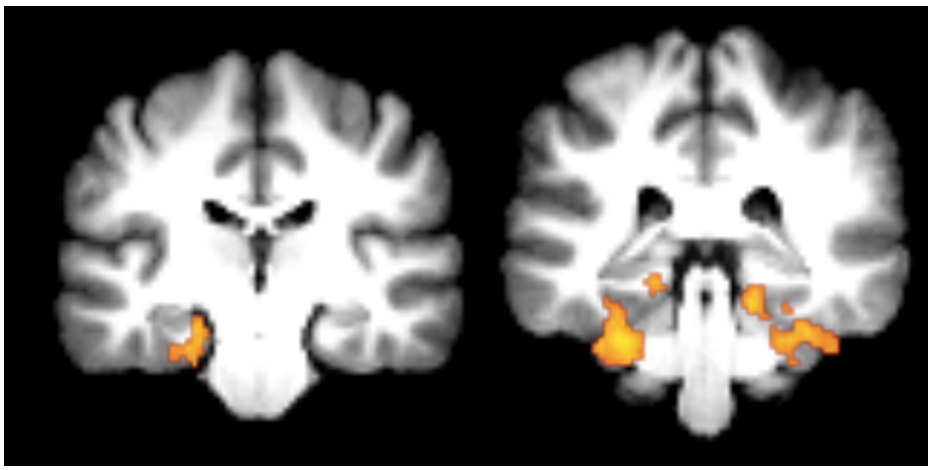

Figure 19: Significant clusters in which BOLD signal tracks trial-to-trial estimates of cue-outcome certainty, irrespective of magnitude of expected reward.

4.4.B.2.3 Trialwise Shifts in Associative Strength

Learning to predict a cue's outcome is represented within the framework of associative models as shifts (increments or decrements of varying magnitude) in the associative strength between relevant cues and outcomes. To test whether shifts in learning at outcome were associated with concomitant BOLD signal, we modeled outcome and included the (unsigned) shift in associate strength corresponding to each trial as a parametric regressor. No significant clusters showed a positive relationship with this regressor; however, a negative relationship was found, such that, trial by trial,

the magnitude of shifts in associative strength were associated with corresponding decreases in BOLD signal within the left hippocampus and MTL cortex, as well as visual processing regions.

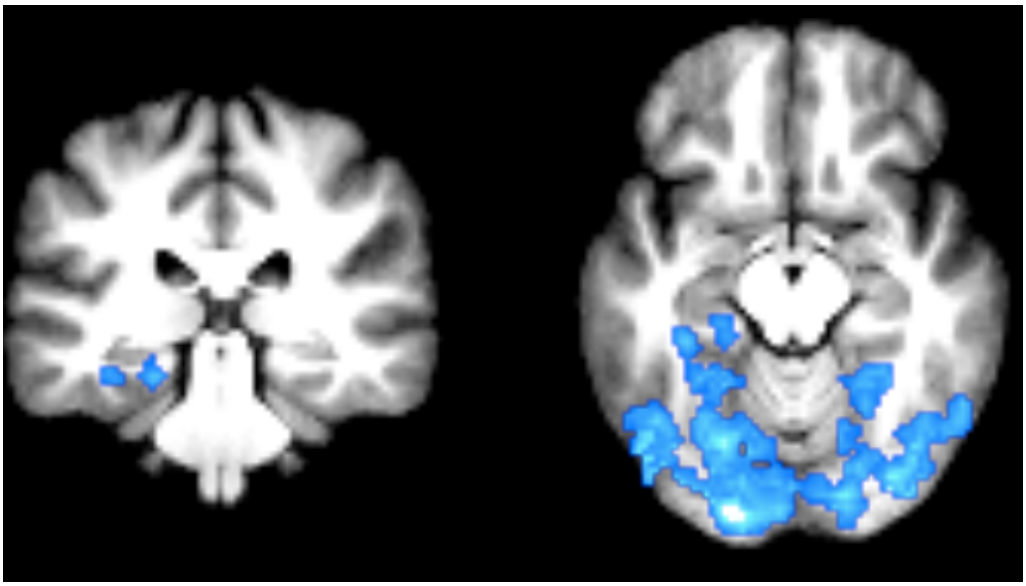

Figure 20: Response to outcome, showing a negative relationship with the magnitude of trial-by-trial shifts in associative strength between cues and their outcomes.

4.4.B.2.4 Uncertainty Salience

Uncertainty salience has been hypothesized to drive cue-outcome associative learning, and evidence from animals has suggested that this signal may be supported by the amygdala). We were curious whether we could detect trial-to-trial variability in BOLD signal that tracked trial-to-trial shifts in uncertainty salience within the amygdala or medial temporal lobes more broadly. We modeled the "whole trial" epoch and included each trial's estimated uncertainty salience as a parametric regressor. We did not identify and significant clusters that positively correlated with uncertainty salience; however, we did identify a set of clusters, including activation within the left medial

temporal lobe that showed a significant negative relationship with alpha. Interestingly, this cluster overlaps with the cluster identified as showing greater BOLD signal for Hits vs. Misses during image presentation, suggesting a common role for this region in cue-outcome learning and subsequent memory formation.

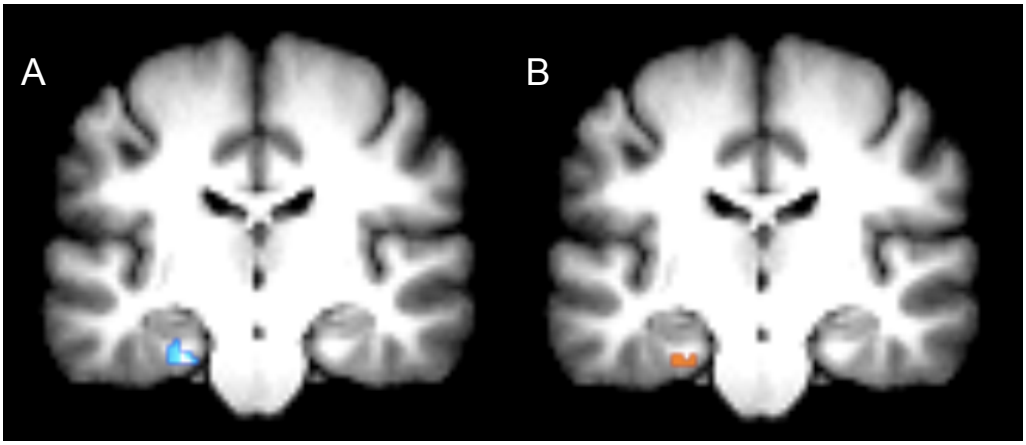

Figure 21: (A) Response to whole trial, showing a negative relationship with uncertainty salience. (B) Conjunction of cluster responding to uncertainty salience and cluster selective for memory encoding success.

4.4.B.3 Connectivity

To characterize connectivity between regions during learning, we focused on activity over the course of the whole trial.

4.4.B.3.1 Amygdala Connectivity: P>NP

If the amygdala plays a role in updating, or otherwise signaling, associative salience, then its functional connectivity may differ as a function of the salience, or motivational significance, of stimuli present in the environment. To test this, we performed a PPI analysis in order to compare amygdalar connectivity in Predictive (P+ and P-) vs. Non-Predictive Novel (NP_nov). Connectivity analyses revealed that greater

connectivity between the amygdala and dopaminergic midbrain, as well as extended amygdala/BNST and regions consistent with the location of the locus coeruleus, was seen during Predictive trials than during Non-Predictive trials.

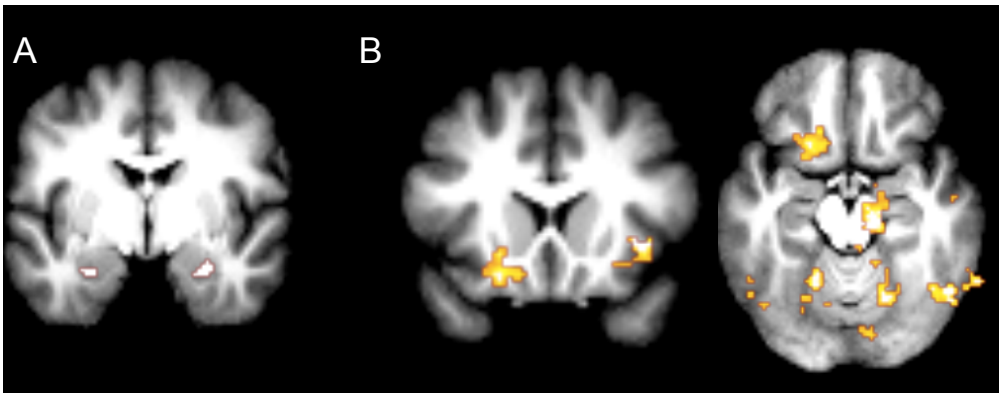

Figure 22: (A) Seed regions for connectivity in bilateral basolateral amygdala. These seed regions were defined using a conjunction between (1) the cluster-corrected whole-brain map of P-NP at cues, and (2) an anatomical mask generated from the Harvard-Oxford Probabilistic Atlas, binarized at 50% probability. (B) Regions showing significantly greater connectivity with the basolateral amygdala for P>NP, over the course of the whole trial.

4.4.B.3.2 Locus Coeruleus Connectivity

We were curious whether we could identify a functional role for the cluster, showing greater response for P>NP at outcome, that we have putatively identified as the locus coeruleus (see Figure 17). Using 2-mm radius spherical masks centered on the local maxima of this cluster, we extracted timecourses and performed a PPI analysis, modeling the whole trial vs. fixation, to characterize whether regions within our network of interest showed significant connectivity. We found that a cluster within the extended amygdala/BNST showed greater connectivity over the course of the whole trial

vs. connectivity during fixation, potentially consistent with the hypothesized neuromodulatory role of norepinephrine in driving shifts in attention during learning.

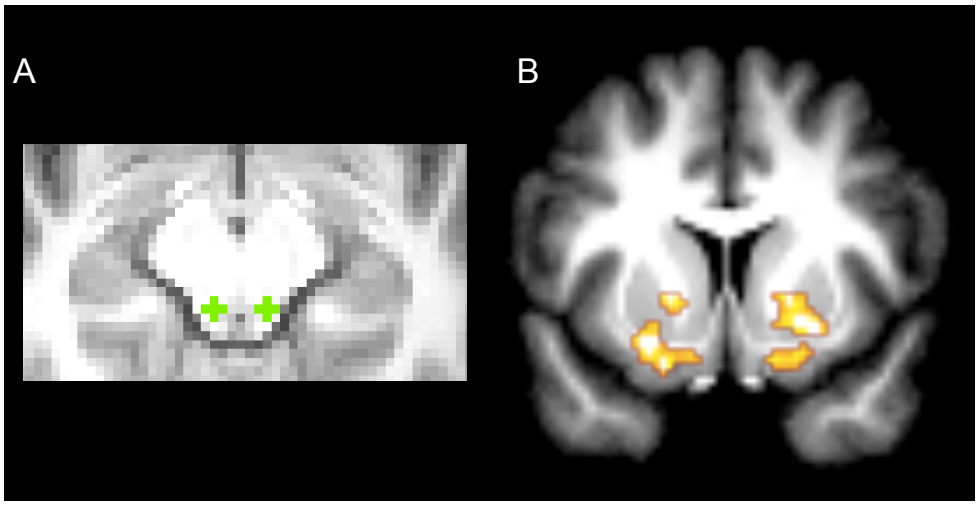

Figure 23: (A) 2mm masks defined by local maxima in contrast of P-NP at outcome, consistent with the anatomical location of the locus coeruleus. (B) Clusters showing significantly greater connectivity with these seed regions during the task period vs. baseline.

After identifying clusters that showed significant connectivity with the putatively-identified locus coeruleus during task vs. baseline, we questioned whether connectivity with this seed may differ as a function of condition. We compared connectivity during Predictive (P+ and P-) vs. Non-Predictive (alternatively, NP_nov or combining NP_nov and NP_flat) but did not identify significant clusters whose connectivity with the LC varied accordingly.

4.4.B.3.3 Connectivity in Response to Novelty

Given the role of novelty in driving responses within our regions of interest, we were interested to see whether connectivity differed as a function of novelty. We found that connectivity with an seed in the medial orbitofrontal cortex (defined anatomically

using the WFT Pickatlas) differed as a function of novelty: the bilateral amygdala and hippocampus showed greater connectivity with the orbitofrontal cortex when viewing novel images than when viewing familiar images.

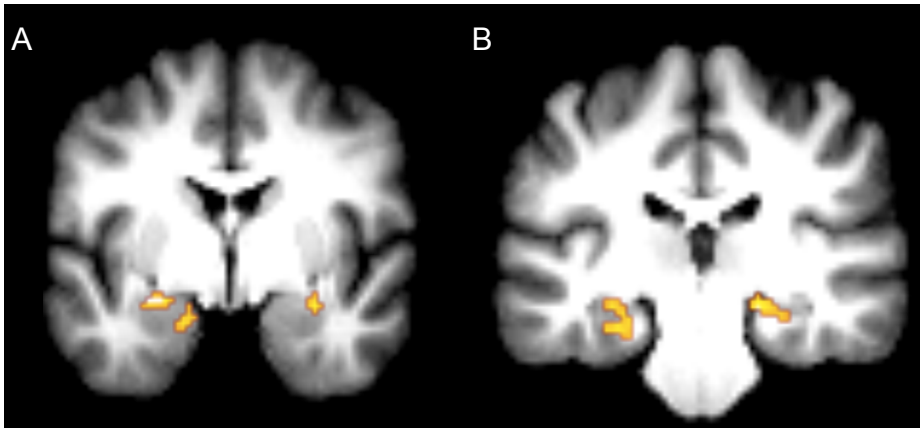

Figure 24: Regions showing greater connectivity with an anatomically-defined OFC seed in response to novel scenes vs. familiar scenes.

4.4.B.4 Individual Differences: Anxiety

In order to characterize the relationship between individual differences in trait anxiety and BOLD response to cue, we ran a group-level analysis, modeling deactivations at Cue relative to baseline with (demeaned) trait anxiety scores as a group-level covariate of interest. We identified a cluster within the amygdala that showed a significant positive relationship with trait anxiety.

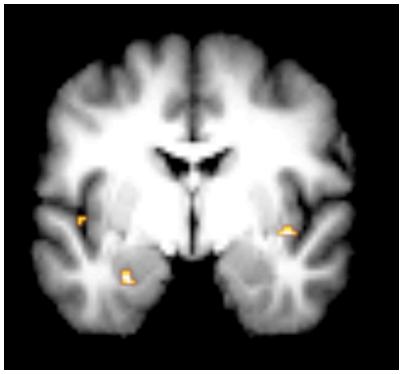

Figure 25: A significant cluster within the left amygdala, in which deactivations in response to Cue presentation shows a positive relationship with individual differences in trait anxiety.

4.5 Discussion

4.5.A Summary of Specific Hypotheses and Results

The present experiment suggests that cue-outcome predictiveness plays a strong role in driving learning, and that our neural systems of interest are strongly responsive to, and differ in their activity as a function of, learned predictiveness. We hypothesized that, as seen in Study 2, learning about cues that differed in their reward-predictive strength would reflect greater salience for predictive than non-predictive cues. Further, we predicted that regions of the brain sensitive to context-specific salience, motivation, and encoding success would all show significant activations with critical components of our task, and would differ over time consistent with the predictions of associative models of learning.

We identified a network of regions, including the amygdala and medial temporal lobes, which responded to our task. Additionally, these regions showed a selective effect of our manipulation: clusters within the medial temporal lobes showed a greater BOLD response to cues that served as reliable predictors of their outcomes, relative to cues that served as unreliable predictors. On a trial-by-trial basis, we showed that our regions of interest shifted in response to new information, and that these shifts tracked cue-outcome contingency, akin to effective salience, rather than uncertainty salience or simple reward anticipation. Subsequent memory for incidentally-presented scene images showed a significant effect of cue predictiveness: scenes presented during states of high outcome-certainty were better remembered than were scenes presented during high outcome-uncertainty.

These data, combined with the eyetracking findings discussed in Study 2, strongly implicate predictiveness, rather than uncertainty, as a driving factor in determining associative salience and driving associative learning (as well as explicit encoding for other information present in the environment, as indexed by memory for scenes.) Thus, these data suggest that motivation to learn in contexts of potentially ambiguous outcomes may be sensitive to the degree of certainty, rather than the magnitude of expected reward.

4.5.B Relationship to Prior Literature

The role of critical neuromodulators, including both dopamine and norepinephrine, is supported by the imaging findings, which suggest that BOLD signal within regions of the brain that support these neurotransmitters is responsive to our task (in the case of norepinephrine) and show altered connectivity with the amygdala as a function of cue predictiveness (in the case of dopamine). BOLD signal cannot be taken as strong evidence of the specific action of neurotransmitter systems; however, these findings are consistent with animal models and extant functional imaging findings in humans (e.g., (Rigoli, Chew, Dayan, & Dolan, 2016).

However, aspects of our findings conflict with some reported findings from humans and animals, primarily regarding the role of uncertainty salience. Imaging findings from humans have found that BOLD signal within the amygdala tracked trial-by-trial estimates of alpha, or uncertainty salience (Li et al., 2011). This is in direct contrast with our finding that BOLD signal within the amygdala tracks effective salience, and that the only region tracking uncertainty salience showed a negative relationship and was localized to the hippocampus and associated MTL cortex, not the amygdala.

However, differences in study design may account for aspects of the differing conclusions. In Li et al.'s study, participants learned to distinguish CS+ from CS- cues on the basis of aversive learning using electric shocks. Following an initial familiarization

phase, participants were re-exposed to the stimuli and reinforcements to undergo reversal learning, and it is in this context that BOLD signal in the amygdala tracked uncertainty salience. In our study, we attempted to relate BOLD activity with uncertainty salience in a context of reward learning, and in which participants lack prior reinforced associations with the relevant stimuli. The role of salience in driving learning in general, and in the role of the amygdala in signaling context-specific salience particular, is further considered in the General Discussion.

Our data also differ from animal models suggesting that uncertainty salience plays a role in driving behavior using rewarded outcomes, and that the amygdala and dopaminergic midbrain support these functions (Hall, 1991; P. C. Holland, 2008; P. C. M. Holland, J.M., 2010). However, these data were collected in contexts of operant reinforcement and feedback provided following overt choices, in contrast with the Pavlovian manipulations employed in the present study. Our data strongly suggest that predictive cues are treated as more salient and more deeply engaging of the neural structures supporting learning and memory when reward delivery is divorced from action. However, the changing task demands with the introduction of a choice, and the associated potential increase in a participant's efficacy in gaining information about an ambiguous choice (rather than, as in the present design, the efficacy necessary to merely allow that information to arrive) may prioritize the processing of uncertain predictors over certain ones. It is an open question, however, whether we would expect this

increase to enhance the salience of uncertain predictors relative to the low level observed in the present studies, or whether manipulations of context or trial demands could prioritize uncertain predictors relative to more certain predictors.

4.5.C Caveats and Limitations

As in Study 2, the experimental design of Study 3 investigated reward learning using Pavlovian stimulus-outcome associative learning, and did not require participants to generate overt responses or make choices in order to gain

Additionally, the experiments described here and in Chapter 3 aimed to maximize the contrast between predictive and non-predictive cues, using deterministic and chance-level relationships between cues and their outcomes. However, this may reduce the role that ambiguity may play in driving attention and learning: in either predictive condition, participants may quickly adopt heuristics of an "all or nothing" learning approach. We may enhance participants' involvement with the stimuli, and potentially enhance the relevance of uncertainty as a meaningful signal for learning, by including cues with a probabilistic (but tractable) relationship in predicting reward delivery or omission.

Given the strong body of work linking reinforcement learning and stimulus-response pairings with activity within the dorsal striatum, we have strong predictions of the activity within dorsal striatum with relation to reward and learning over time; however, data collection in this study prioritized the midbrain, hippocampus and

amygdala and thus coverage did not extend to the dorsal striatum in most subjects. Future investigations would be enhanced by extending data collection to the whole brain, including dorsal striatum as well as a greater extent of frontal cortex.

4.5.D Outstanding questions

Future studies may develop these findings, testing, for example, the long-term resilience and flexibility of these learned representations. It may be, for example, that participants would be impeded in reversal learning should the cue-outcome contingencies change, while this effect may vary as a function of the initially learned values.

5. General Discussion

5.1 Motivated Information-Seeking: Summary of Findings

The mere presence of reward within an animal's environment does not guarantee reward receipt. In order to successfully gain rewards, animals must be able to actively learn about their environments in order to predict where and when rewards will be available, and what actions they must perform to gain them.

Motivated information-seeking behaviors are necessary to enable adaptive choices: an animal that is able to locate the rewards present in its environment and take the actions necessary to gain them faces a higher likelihood of adaptive success (Shohamy & Adcock, 2010). The processes driving our engagement of exploratory behaviors, and the effort with which we explore the environments we find ourselves in, thus directly shape our experience.

In the studies described above, we manipulated aspects of participants' environments and tested for differences in their resulting exploratory behavior. Through the completion of these studies, we demonstrated that the deployment of exploratory behavior can be experimentally manipulated in humans, and that the systems supporting active learning can influence exploratory behaviors across modalities.

Our exploration and reward-seeking findings confirm that individuals' vigor in reward pursuit is tightly linked with the magnitude of potential rewards, while overt information-seeking is likewise influenced by history of earned rewards. The

mechanisms supporting these effects could not be characterized within the behavioral study, but our findings are consistent with the hypothesized role of striatal dopamine in invigorating motor movements, particularly active exploration. Our findings regarding both vigor in pursuit of rewards and vigor in pursuit of information are consistent with the theories described here regarding dopamine's role in enhancing motivation and active exploration.

Our eyetracking findings suggest that, contrary to the predictions of the Pearce-Hall model and consistent with the predictions of Mackintosh, among others, increasing predictive validity drives increased associative salience. Participants spent more time looking at cues which served as reliable indicators, regardless of whether they indicated the delivery or absence of reward, relative to unreliable cues.

Our imaging findings confirm that the amygdala is responsive to cue-outcome contingencies, broadening its role in associative salience and beyond the realm of emotional valence and emotional learning. Here, the amygdala was found to play a strong role in responding to motivationally-relevant information and altered connectivity with key neuromodulatory (including, putatively, dopaminergic cells within the substantia nigra and ventral tegmental area.)

Across the three studies described here, we manipulated individuals' history of rewards and measured their willingness and vigor to seek, pursue, attend to, and encode new information. We found, across the studies considered here, that the

availability of reward in one's history induced a greater likelihood of active reward seeking and overt information-seeking. This may be related to observed effects relating tonic dopamine with exploratory actions; similarly, we found that when participants had a strong sense of the predicted outcome they were about to experience, they were more active in their exploration of novel information presented in the delay (in Study 2), and showed greater activation across regions of interest and better memory for the associated novel information (in Study 3.)

5.2 Models of Associative Learning

Models of associative learning have long been concerned with the relationship between a stimulus' predictive validity, its perceived salience, and its consequent potential as a target for new learning (its associability). However, families of models differ in their predictions about the nature of this relationship. According to the Pearce-Hall model, uncertainty regarding an event or stimulus' outcome (for example, the presence or absence of reward) can enhance the salience of relevant stimuli and facilitate learning relative to a stimulus whose outcome is known (Holland & Maddux 2010; Pearce, Kaye, & Hall 1982). In contrast, Mackintosh's (1975) model makes an opposing prediction: stimuli that consistently predict a given outcome (e.g., reward delivery or omission) should produce an increase in attention, and learning about that stimulus should be facilitated relative to a cue with a more ambiguous outcome.

Experimental evidence in humans has supported elements of both theories: Hogarth et al. (2008) found eyegaze behavior during passive conditioning to be consistent with Pearce-Hall models, while Le Pelley et al.'s investigation of category learning (2011) found evidence supporting the Mackintosh model. Thus, the nature of the relationship between stimulus predictiveness, salience (and consequent associability), and new learning remains uncharacterized. Our data show support for the mechanisms put forward by Mackintosh, as well as Esber and Haselgrove. However, as shown in Study 2, expectation violations (i.e., errors in reward prediction) also serve as significant drivers of attention in learning and gaze to motivationally relevant cues. Future studies combining varying levels of predictive certainty, active exploration and choice behaviors are necessary to fully characterize the relationships between uncertainty, predictions, and active exploration.

These data suggest that, in developing interventions to maximize exploration behaviors, ideal behavioral stimulation may not rely, as in the perspective of the Pearce-Hall model, on the presentation of information that violates expectations, but instead on exposing participants to stimuli and environments which allows them to develop, test, and confirm hypotheses. The greater overt gaze (in Study 2) and BOLD response within amygdala and hippocampus (in study 3) for predictive, vs. non-predictive, stimuli illustrate the critical significance, both in behavioral and neural function, of reliably predictive information in driving learning and exploration behaviors.

5.3 Salience and the Amygdala

The amygdala is a neural structure to which a large amount of inquiry has been focused. This large body of work has yielded a complex set of findings which reflect subtle degrees of variation in amygdala activity, reflecting an individual's internal state, history with relevant stimuli, and overall context. The "classical" framework regarding the amygdala, most commonly associated with the work of LeDoux and others (J. E. LeDoux, 2000), holds that the amygdala is primarily sensitive to negative emotions such as fear or anger, and that it prioritizes responses to aversive or generally negative stimuli in order to facilitate avoidance responses.

Subsequent work by LeDoux and others has complicated this theoretical link between the amygdala and negative emotion. The amygdala has been shown to respond to a wide class of stimuli, including stimuli directly inducing or indirectly associated with both negative and positive affect (Sander, Grafman, & Zalla, 2003). These include erotic imagery (Beauregard, Levesque, & Bourgouin, 2001), happy vs. neutral faces (Breiter et al., 1996) or words (S. Hamann & Mao, 2002), food during hungry vs. satiated conditions (LaBar et al., 2001), and positive feedback (the word "WIN") (Zalla et al., 2000), among a wide array of others. Thus, the amygdala, taken as a whole, is not sensitive to any one sensory domain or qualitative characteristic, but, potentially, to the context-dependent salience of the task an individual is engaged with. In naturalistic settings, it is highly likely that negative or potentially aversive stimuli demand a more

immediate response than do potential rewards, and are thus treated as more salient and preferentially processed by the amygdala. However, in contexts of reward-mediated learning as in the present study, the salient factor an individual confronts in the predictive strength of the cues in relating to reward delivery. This is consistent with our observed effects, as both greater overall BOLD signal in the amygdala and greater amygdalar-midbrain connectivity was observed during states of high predictive certainty relative to states of low predictive certainty.

Part of the complicating factors may be due to the underlying anatomy. The amygdala is not a single structure, but a set of interconnected nuclei who receive projections from, and project to, distinct cortical and subcortical structures. The central nucleus has been associated with Pearce-Hall-dependent learning signals (P. C. Holland & Gallagher, 1993a, 1993b, 1999), and with motor output structures and behaviors more broadly (Gentile, Jarrell, Teich, McCabe, & Schneiderman, 1986; J. LeDoux, 2003). In contast, the lateral and basolateral amygdala are thought to support the input of information into the amygdala, and receives heavy innervation from the hippocampus (Canteras & Swanson, 1992; J. LeDoux, 2003). This hippocampal-lateral amygdala connectivity is thought to support contextual conditioning (J. LeDoux, 2003). Fine-grain localization of signal within substructures of the amygdala at the spatial resolution of the functional data collected in Study 3 is difficult to validate; however, in Study 3 activations reported in both left and right amygdala corresponded to the lateral and

ventral portions of the amygdala, corresponding to the basolateral amygdala, consistent with the evidence linking this region with hippocampal neurons and supporting conditioning within a rewarded context.

5.4 Dopamine, Reward, Errors, and Effort

Most work relating dopamine signaling emphasizes the well-documented relationship between the magnitude of reward prediction errors following a given event and the associated phasic signal transmitted by populations of dopamine neurons (Schultz, 1998; Schultz, Dayan, & Montague, 1997; Schultz & Dickinson, 2000). A growing body of work, however, suggests that critical aspects of dopaminergic function may not be directly supported by high-temporal-resolution phasic responses, but by lower temporal- and spatial-resolution tonic signaling, as reward signals summed over time combine to drive tonic shifts in neuromodulatory state (Niv et al., 2007). This tonic signal has been characterized as associated with shifts in reward-pursuit vigor and mediating effortful choice. (Le Bouc et al., 2016; Niv et al., 2007; Zenon et al., 2016).

The sensitivity of behaviors to either trial-by-trial variations in outcome or in global trends in reward receipt may reflect aspects of these signals, dependent on the qualities substrates of the neural substrates involved. While ventral striatum is rich in dopamine receptors, allowing a fast and spatially localized response to temporally distinct phasic signals, the hippocampus supports a much sparser density of dopamine

receptors (Gangarossa et al., 2012). With such a dynamic in place, it is likely that dopamine's role in driving hippocampal learning reflects (at least) two indirect mechanisms: tonic signaling in the hippocampus itself, and the differential engagement with and pursuit of information acquired as a result of direct dopaminergic innervation of the striatum and associated behavioral activation.

In the data considered here, we find evidence for both forms of signaling: Study 1 is most explicitly designed to test for the relationship between earned rewards and response vigor, and we identified a relationship between reward availability, success and vigor both in reward pursuit and in information-seeking. However, the role of classical reward-prediction-error signaling was less than anticipated in Study 2 and 3: in contrast to our a prior expectations, the experience of uncertainty associated with reward-prediction errors was not associated with enhanced engagement, but instead predicted diminished eyegaze and engagement: greater connectivity with dopaminergic regions in Study 3 was seen not during states of low predictive certainty, but during high predictive certainty. However, in Study 2, prediction errors on immediately preceding trials predicted shifts in gaze to cues, suggesting that expectation-violations may be driving learning in ways not fully characterized by the associative-learning models considered here.

5.5 Foraging and Information-Seeking

Research into information-seeking behaviors has largely been concerned with learning from competing options, in a framework contrasted choices to known rewards ("exploit") vs sampling of potentially novel or differentially rewarding choices ("explore.") The studies considered here build on the findings reported from "exploration" but do not match directly: novelty-seeking, potentially to aberrant degrees, is strongly associated with this framework's conception of "exploration", which matches well with aspects of exploratory behaviors described here but falls short of the potential full complexity. Critically, "exploration" as indicated by novel choices in gambling contexts reflects novelty-seeking and value testing as concrete, temporally stimulus-locked choices ties to a discrete outcome (Frank et al., 2004). Here, exploration is investigated more broadly, as a set of behaviors associated with greater/more active interaction with a given stimulus (in a spatial context) or greater willingness to expend energy to gain information (in an effortful information-seeking context.)

5.6 Uncertainty and Risk

In the studies considered here, uncertainty reflects an individual's current lack of clarity between some stimulus and its true (if unknown) likelihood of reward delivery. These probabilities largely remain stable over time, and participants, in Studies 2 and 3, did not need to make a choice in order to gain them; in Study 1, choices were needed but

there was little ambiguity regarding the expected value of the choices being decided between. In none of the studies considered here, then, were cases in which the relationship between a necessary choice and its outcome was uncertain. This contrasts with most forms of deicion-making in daily life: we may not be certain of the value of two competing options, or may further lack certainty of a linkage between our preference for a given option and its ultimate delivery in response to our choice.

Neurophysiological studies suggest that this clarity may suppress, or diminish, the activity of dopaminergic neurons and the importance of dopaminergic signals in information-processing across the brain as a whole: while dopamine neurons respond to unpredicted rewards as well as to cues predicting rewards, evidence from cell recordings in animals suggests that they are responsive to ambiguities in reward-delivery likelihood, potentially generating a response "ramp" prior to weakly-anticipated rewards prior to delivery (Fiorillo, Tobler, & Schultz, 2003). This finding relates well to accounts of uncertainty or unpredicted outcomes in facilitating increments in associative salience, and to behavioral activation in response to novel or ambiguous events and contexts (Berlyne, 1960; Berlyne, Koenig, & Hirota, 1966).

5.7 General Conclusion

Individuals must make choices to gain rewards, but they must possess information in order to make optimal choices. Information is thus intrinsically valuable:

individuals who cannot generate active forms of learning behaviors, when contextually appropriate, are strongly disadvantaged in terms of their ability to gain available rewards now and to predict potentially available rewards in future. The present studies attempt to describe potential mechanisms by which such information-seeking behaviors are deployed and regulated in humans. Overall, a strong link between individuals' history of outcomes with given cues, and their ability to predict those outcomes when encountering the relevant stimuli, was shown both in spatial contexts of reward choice, and non-spatial contexts of reward delivery and passive observation. Future work should further explore the factors promoting and inhibiting individuals' tendency to actively seek information, primarily with regards to individual differences in psychological factors such as self-efficacy and anxiety.

References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: mesolimbic activation precedes memory formation. *Neuron, 50*(3), 507-517. doi: 10.1016/j.neuron.2006.03.036

Alvarez, R. P., Chen, G., Bodurka, J., Kaplan, R., & Grillon, C. (2011). Phasic and sustained fear in humans elicits distinct patterns of brain activity. *Neuroimage, 55*(1), 389-400. doi: 10.1016/j.neuroimage.2010.11.057

Ambroggi, F., Ishikawa, A., Fields, H. L., & Nicola, S. M. (2008). Basolateral amygdala neurons facilitate reward-seeking behavior by exciting nucleus accumbens neurons. *Neuron, 59*(4), 648-661. doi: 10.1016/j.neuron.2008.07.004

Anderson, A. K., Christoff, K., Stappen, I., Panitz, D., Ghahremani, D. G., Glover, G., . . . Sobel, N. (2003). Dissociated neural representations of intensity and valence in human olfaction. *Nat Neurosci, 6*(2), 196-202. doi: 10.1038/nn1001

Asan, E. (1998). The catecholaminergic innervation of the rat amygdala. *Adv Anat Embryol Cell Biol, 142*, 1-118.

Aston-Jones, G., Chiang, C., & Alexinsky, T. (1991). Discharge of noradrenergic locus coeruleus neurons in behaving rats and monkeys suggests a role in vigilance. *Prog Brain Res, 88*, 501-520.

Aston-Jones, G., Rajkowski, J., & Cohen, J. (1999). Role of locus coeruleus in attention and behavioral flexibility. *Biol Psychiatry, 46*(9), 1309-1320.

Aston-Jones, G., Rajkowski, J., Kubiak, P., & Alexinsky, T. (1994). Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *J Neurosci, 14*(7), 4467-4480.

Avery, S. N., Clauss, J. A., & Blackford, J. U. (2016). The Human BNST: Functional Role in Anxiety and Addiction. *Neuropsychopharmacology, 41*(1), 126-141. doi: 10.1038/npp.2015.185

Baldo, B. A., Sadeghian, K., Basso, A. M., & Kelley, A. E. (2002). Effects of selective dopamine D1 or D2 receptor blockade within nucleus accumbens subregions on ingestive behavior and associated motor activity. *Behav Brain Res, 137*(1-2), 165-177.

Bandura, A. (1977). Self-efficacy: toward a unifying theory of behavioral change. *Psychol Rev, 84*(2), 191-215.

Bandura, A. (1982). Self-efficacy mechanism in human agency. *American Psychologist, 37*(2), 122-147.

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

Beauregard, M., Levesque, J., & Bourgouin, P. (2001). Neural correlates of conscious self-regulation of emotion. *J Neurosci, 21*(18), RC165.

Beckmann, J. S., Marusich, J. A., Gipson, C. D., & Bardo, M. T. (2011). Novelty seeking, incentive salience and acquisition of cocaine self-administration in the rat. *Behav Brain Res, 216*(1), 159-165. doi: 10.1016/j.bbr.2010.07.022

Beeler, J. A. (2012). Thorndike's Law 2.0: Dopamine and the Regulation of Thrift. *Front Neurosci, 6*, 116. doi: 10.3389/fnins.2012.00116

Beesley, T., Nguyen, K. P., Pearson, D., & Le Pelley, M. E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *Q J Exp Psychol (Hove), 68*(11), 2175-2199. doi: 10.1080/17470218.2015.1009919

Beierholm, U., Guitart-Masip, M., Economides, M., Chowdhury, R., Duzel, E., Dolan, R., & Dayan, P. (2013). Dopamine modulates reward-related vigor. *Neuropsychopharmacology, 38*(8), 1495-1503. doi: 10.1038/npp.2013.48

Berlyne, D. E. (1960). Conflict, Arousal, and Curiosity.

Berlyne, D. E., Koenig, I. D., & Hirota, T. (1966). Novelty, arousal, and the reinforcement of diversive exploration in the rat. *J Comp Physiol Psychol, 62*(2), 222-226.

Berridge, C. W., & Waterhouse, B. D. (2003). The locus coeruleus-noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Res Brain Res Rev, 42*(1), 33-84.

Breiter, H. C., Etcoff, N. L., Whalen, P. J., Kennedy, W. A., Rauch, S. L., Buckner, R. L., . . . Rosen, B. R. (1996). Response and habituation of the human amygdala during visual processing of facial expression. *Neuron, 17*(5), 875-887.

Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron, 63*(1), 119-126. doi: 10.1016/j.neuron.2009.06.009

Canteras, N. S., & Swanson, L. W. (1992). Projections of the ventral subiculum to the amygdala, septum, and hypothalamus: a PHAL anterograde tract-tracing study in the rat. *J Comp Neurol, 324*(2), 180-194. doi: 10.1002/cne.903240204

Carvalho Poyraz, F., Holzner, E., Bailey, M. R., Meszaros, J., Kenney, L., Kheirbek, M. A., . . . Kellendonk, C. (2016). Decreasing Striatopallidal Pathway Function Enhances Motivation by Energizing the Initiation of Goal-Directed Action. *J Neurosci, 36*(22), 5988-6001. doi: 10.1523/JNEUROSCI.0444-16.2016

Catanese, J., & van der Meer, M. (2013). A network state linking motivation and action in the nucleus accumbens. *Neuron, 78*(5), 753-754. doi: 10.1016/j.neuron.2013.05.021

Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theor Popul Biol, 9*(2), 129-136.

Cunningham, W. A., & Brosch, T. (2012). Motivational Salience Amygdala Tuning From Traits, Needs, Values, and Goals. *Current Directions in Psychological Science, 21*(1), 54-59.

de la Fuente-Fernandez, R., Phillips, A. G., Zamburlini, M., Sossi, V., Calne, D. B., Ruth, T. J., & Stoessl, A. J. (2002). Dopamine release in human ventral striatum and expectation of reward. *Behav Brain Res, 136*(2), 359-363.

Dickinson, A. (2012). Associative learning and animal cognition. *Philos Trans R Soc Lond B Biol Sci, 367*(1603), 2733-2742. doi: 10.1098/rstb.2012.0220

Dougal, S., Phelps, E. A., & Davachi, L. (2007). The role of medial temporal lobe in item recognition and source recollection of emotional stimuli. *Cogn Affect Behav Neurosci, 7*(3), 233-242.

Dunsmoor, J. E., Murty, V. P., Davachi, L., & Phelps, E. A. (2015). Emotional learning selectively and retroactively strengthens memories for related events. *Nature, 520*(7547), 345-348. doi: 10.1038/nature14106

Eichenbaum, H. (2004). Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron, 44*(1), 109-120. doi: 10.1016/j.neuron.2004.08.028

El-Amamy, H., & Holland, P. C. (2006). Substantia nigra pars compacta is critical to both the acquisition and expression of learned orienting of rats. *Eur J Neurosci, 24*(1), 270-276. doi: 10.1111/j.1460-9568.2006.04896.x

Esber, G. R., & Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: a model of attention in associative learning. *Proc Biol Sci, 278*(1718), 2553-2561. doi: 10.1098/rspb.2011.0836

Esber, G. R., Roesch, M. R., Bali, S., Trageser, J., Bissonette, G. B., Puche, A. C., . . . Schoenbaum, G. (2012). Attention-related Pearce-Kaye-Hall signals in basolateral amygdala require the midbrain dopaminergic system. *Biol Psychiatry, 72*(12), 1012-1019. doi: 10.1016/j.biopsych.2012.05.023

Failing, M., Nissens, T., Pearson, D., Le Pelley, M., & Theeuwes, J. (2015). Oculomotor capture by stimuli that signal the availability of reward. *J Neurophysiol, 114*(4), 2316-2327. doi: 10.1152/jn.00441.2015

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science, 299*(5614), 1898-1902. doi: 10.1126/science.1077349

Florin-Lechner, S. M., Druhan, J. P., Aston-Jones, G., & Valentino, R. J. (1996). Enhanced norepinephrine release in prefrontal cortex with burst stimulation of the locus coeruleus. *Brain Res, 742*(1-2), 89-97.

Frank, M. J., Seeberger, L. C., & O'Reilly R, C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science, 306*(5703), 1940-1943. doi: 10.1126/science.1102941

Frey, U., Schroeder, H., & Matthies, H. (1990). Dopaminergic antagonists prevent long-term maintenance of posttetanic LTP in the CA1 region of rat hippocampal slices. *Brain Res, 522*(1), 69-75.

Gabrieli, J. D. (1998). Cognitive neuroscience of human memory. *Annu Rev Psychol, 49*, 87-115. doi: 10.1146/annurev.psych.49.1.87

Gallagher, M., & Holland, P. C. (1994). The amygdala complex: multiple roles in associative learning and attention. *Proc Natl Acad Sci U S A, 91*(25), 11771-11776.

Gangarossa, G., Longueville, S., De Bundel, D., Perroy, J., Herve, D., Girault, J. A., & Valjent, E. (2012). Characterization of dopamine D1 and D2 receptor-expressing

neurons in the mouse hippocampus. *Hippocampus, 22*(12), 2199-2207. doi: 10.1002/hipo.22044

Gasbarri, A., Packard, M. G., Campana, E., & Pacitti, C. (1994). Anterograde and retrograde tracing of projections from the ventral tegmental area to the hippocampal formation in the rat. *Brain Res Bull, 33*(4), 445-452.

Geisler, S., & Zahm, D. S. (2005). Afferents of the ventral tegmental area in the rat-anatomical substratum for integrative functions. *J Comp Neurol, 490*(3), 270-294. doi: 10.1002/cne.20668

Gentile, C. G., Jarrell, T. W., Teich, A., McCabe, P. M., & Schneiderman, N. (1986). The role of amygdaloid central nucleus in the retention of differential pavlovian conditioning of bradycardia in rabbits. *Behav Brain Res, 20*(3), 263-273.

Grant, S. J., Aston-Jones, G., & Redmond, D. E., Jr. (1988). Responses of primate locus coeruleus neurons to simple and complex sensory stimuli. *Brain Res Bull, 21*(3), 401-410.

Guitart-Masip, M., Beierholm, U. R., Dolan, R., Duzel, E., & Dayan, P. (2011). Vigor in the face of fluctuating rates of reward: an experimental examination. *J Cogn Neurosci, 23*(12), 3933-3938. doi: 10.1162/jocn_a_00090

Guitart-Masip, M., Bunzeck, N., Stephan, K. E., Dolan, R. J., & Duzel, E. (2010). Contextual novelty changes reward representations in the striatum. *J Neurosci, 30*(5), 1721-1726. doi: 10.1523/JNEUROSCI.5331-09.2010

Guitart-Masip, M., Fuentemilla, L., Bach, D. R., Huys, Q. J., Dayan, P., Dolan, R. J., & Duzel, E. (2011). Action dominates valence in anticipatory representations in the human striatum and dopaminergic midbrain. *J Neurosci, 31*(21), 7867-7875. doi: 10.1523/JNEUROSCI.6376-10.2011

Hall, G. (1991). *Perceptual and Associative Learning*. Oxford: Clarendon.

Hamann, S. B., Ely, T. D., Hoffman, J. M., & Kilts, C. D. (2002). Ecstasy and agony: activation of the human amygdala in positive and negative emotion. *Psychol Sci, 13*(2), 135-141.

Hamann, S., & Mao, H. (2002). Positive and negative emotional verbal stimuli elicit activity in the left amygdala. *Neuroreport, 13*(1), 15-19.

Hannula, D. E., Baym, C. L., Warren, D. E., & Cohen, N. J. (2012). The eyes know: eye movements as a veridical index of memory. *Psychol Sci, 23*(3), 278-287. doi: 10.1177/0956797611429799

Harden, K. P., & Tucker-Drob, E. M. (2011). Individual differences in the development of sensation seeking and impulsivity during adolescence: further evidence for a dual systems model. *Dev Psychol, 47*(3), 739-746. doi: 10.1037/a0023279

Haselgrove, M., Esber, G. R., Pearce, J. M., & Jones, P. M. (2010). Two kinds of attention in Pavlovian conditioning: evidence for a hybrid model of learning. *J Exp Psychol Anim Behav Process, 36*(4), 456-470. doi: 10.1037/a0018528

Hogarth, L., Dickinson, A., Austin, A., Brown, C., & Duka, T. (2008). Attention and expectation in human predictive learning: the role of uncertainty. *Q J Exp Psychol (Hove), 61*(11), 1658-1668. doi: 10.1080/17470210701643439

Holland, P. C. (2008). Cognitive versus stimulus-response theories of learning. *Learn Behav, 36*(3), 227-241.

Holland, P. C. (2013). Prediction and surprise in human fear conditioning (Commentary on Boll et al.). *Eur J Neurosci, 37*(5), 757. doi: 10.1111/ejn.12147

Holland, P. C., & Gallagher, M. (1993a). Amygdala central nucleus lesions disrupt increments, but not decrements, in conditioned stimulus processing. *Behav Neurosci, 107*(2), 246-253.

Holland, P. C., & Gallagher, M. (1993b). Effects of amygdala central nucleus lesions on blocking and unblocking. *Behav Neurosci, 107*(2), 235-245.

Holland, P. C., & Gallagher, M. (1999). Amygdala circuitry in attentional and representational processes. *Trends Cogn Sci, 3*(2), 65-73.

Holland, P. C., & Gallagher, M. (2004). Amygdala-frontal interactions and reward expectancy. *Curr Opin Neurobiol, 14*(2), 148-155. doi: 10.1016/j.conb.2004.03.007

Holland, P. C., Thornton, J. A., & Ciali, L. (2000). The influence of associability changes in negative patterning and other discriminations. *J Exp Psychol Anim Behav Process, 26*(4), 462-476.

Holland, P.C. & Maddux, J.M. (2010). Brain systems of attention in associative learning. In L. P. M. Mitchell CJ (Ed.), *Attention and associative learning: from brain to behaviour* (pp. 305-349).

Horvitz, J. C., Stewart, T., & Jacobs, B. L. (1997). Burst activity of ventral tegmental dopamine neurons is elicited by sensory stimuli in the awake cat. *Brain Res, 759*(2), 251-258.

Ikemoto, S., & Panksepp, J. (1996). Dissociations between appetitive and consummatory responses by pharmacological manipulations of reward-relevant brain regions. *Behav Neurosci, 110*(2), 331-345.

Järvelin, K., & Wilson, T. D., 9(1), 9-1. (2003). On conceptual models for information seeking and retrieval research. . *Information Research, 9*(1).

Jepma, M., & Nieuwenhuis, S. (2011). Pupil diameter predicts changes in the exploration-exploitation trade-off: evidence for the adaptive gain theory. *J Cogn Neurosci, 23*(7), 1587-1596. doi: 10.1162/jocn.2010.21548

Kayser, A. S., Mitchell, J. M., Weinstein, D., & Frank, M. J. (2015). Dopamine, locus of control, and the exploration-exploitation tradeoff. *Neuropsychopharmacology, 40*(2), 454-462. doi: 10.1038/npp.2014.193

Koob, G. F., Riley, S. J., Smith, S. C., & Robbins, T. W. (1978). Effects of 6-hydroxydopamine lesions of the nucleus accumbens septi and olfactory tubercle on feeding, locomotor activity, and amphetamine anorexia in the rat. *J Comp Physiol Psychol, 92*(5), 917-927.

Kovach, C. K., Daw, N. D., Rudrauf, D., Tranel, D., O'Doherty, J. P., & Adolphs, R. (2012). Anterior prefrontal cortex contributes to action selection through tracking of recent reward trends. *J Neurosci, 32*(25), 8434-8442. doi: 10.1523/jneurosci.5468-11.2012

Kumaran, D., & Maguire, E. A. (2009). Novelty signals: a window into hippocampal information processing. *Trends Cogn Sci, 13*(2), 47-54. doi: 10.1016/j.tics.2008.11.004

Kurniawan, I. T., Guitart-Masip, M., & Dolan, R. J. (2011). Dopamine and effort-based decision making. *Front Neurosci, 5*, 81. doi: 10.3389/fnins.2011.00081

LaBar, K. S., Gitelman, D. R., Parrish, T. B., Kim, Y. H., Nobre, A. C., & Mesulam, M. M. (2001). Hunger selectively modulates corticolimbic activation to food stimuli in humans. *Behav Neurosci, 115*(2), 493-500.

Le Bouc, R., Rigoux, L., Schmidt, L., Degos, B., Welter, M. L., Vidailhet, M., . . . Pessiglione, M. (2016). Computational Dissection of Dopamine Motor and

Motivational Functions in Humans. *J Neurosci, 36*(25), 6623-6633. doi: 10.1523/JNEUROSCI.3078-15.2016

Le Pelley, M. E., Beesley, T., & Griffiths, O. (2011). Overt attention and predictiveness in human contingency learning. *J Exp Psychol Anim Behav Process, 37*(2), 220-229. doi: 10.1037/a0021384

Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N., & Wills, A. J. (2016). Attention and associative learning in humans: An integrative review. *Psychol Bull, 142*(10), 1111-1140. doi: 10.1037/bul0000064

LeDoux, J. (2003). The emotional brain, fear, and the amygdala. *Cell Mol Neurobiol, 23*(4-5), 727-738.

LeDoux, J. E. (2000). Emotion circuits in the brain. *Annu Rev Neurosci, 23*, 155-184. doi: 10.1146/annurev.neuro.23.1.155

Lee, H. J., Gallagher, M., & Holland, P. C. (2010). The central amygdala projection to the substantia nigra reflects prediction error information in appetitive conditioning. *Learn Mem, 17*(10), 531-538. doi: 10.1101/lm.1889510

Lee, H. J., Groshek, F., Petrovich, G. D., Cantalini, J. P., Gallagher, M., & Holland, P. C. (2005). Role of amygdalo-nigral circuitry in conditioning of a visual stimulus paired with food. *J Neurosci, 25*(15), 3881-3888. doi: 10.1523/JNEUROSCI.0416-05.2005

Lee, H. J., Wheeler, D. S., & Holland, P. C. (2011). Interactions between amygdala central nucleus and the ventral tegmental area in the acquisition of conditioned cue-directed behavior in rats. *Eur J Neurosci, 33*(10), 1876-1884. doi: 10.1111/j.1460-9568.2011.07680.x

Lee, H. J., Youn, J. M., O, M. J., Gallagher, M., & Holland, P. C. (2006). Role of substantia nigra-amygdala connections in surprise-induced enhancement of attention. *J Neurosci, 26*(22), 6077-6081. doi: 10.1523/JNEUROSCI.1316-06.2006

Lee, S. C., Amir, A., Headley, D. B., Haufler, D., & Pare, D. (2016). Basolateral amygdala nucleus responses to appetitive conditioned stimuli correlate with variations in conditioned behaviour. *Nat Commun, 7*, 12275. doi: 10.1038/ncomms12275

Li, J., & Daw, N. D. (2011). Signals in human striatum are appropriate for policy update rather than value prediction. *J Neurosci, 31*(14), 5504-5511. doi: 10.1523/jneurosci.6316-10.2011

Li, J., Schiller, D., Schoenbaum, G., Phelps, E. A., & Daw, N. D. (2011). Differential roles of human striatum and amygdala in associative learning. *Nat Neurosci, 14*(10), 1250-1252. doi: 10.1038/nn.2904

Lisman, J. E., & Grace, A. A. (2005). The hippocampal-VTA loop: controlling the entry of information into long-term memory. *Neuron, 46*(5), 703-713. doi: 10.1016/j.neuron.2005.05.002

Ljungberg, T., & Ungerstedt, U. (1976). Sensory inattention produced by 6-hydroxydopamine-induced degeneration of ascending dopamine neurons in the brain. *Exp Neurol, 53*(3), 585-600.

Loughlin, S. E., & Fallon, J. H. (1983). Dopaminergic and non-dopaminergic projections to amygdala from substantia nigra and ventral tegmental area. *Brain Res, 262*(2), 334-338.

Luque, D., Vadillo, M. A., Le Pelley, M. E., & Beesley, T. (2016). Prediction and uncertainty in associative learning: examining controlled and automatic components of learned attentional biases. *Q J Exp Psychol (Hove)*, 1-19. doi: 10.1080/17470218.2016.1188407

Mackintosh.J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychol Rev, 82*(4), 276-298.

Maddux, J. M., Kerfoot, E. C., Chatterjee, S., & Holland, P. C. (2007). Dissociation of attention in learning and action: effects of lesions of the amygdala central nucleus, medial prefrontal cortex, and posterior parietal cortex. *Behav Neurosci, 121*(1), 63-79. doi: 10.1037/0735-7044.121.1.63

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*(1), 74-84. doi: 10.3758/s13415-011-0068-4

Mather, M., & Sutherland, M. R. (2011). Arousal-biased competition in perception and memory. *Perspect Psychol Sci, 6*(2), 114-133. doi: 10.1177/1745691611400234

McGinty, V. B., Lardeux, S., Taha, S. A., Kim, J. J., & Nicola, S. M. (2013). Invigoration of reward seeking by cue and proximity encoding in the nucleus accumbens. *Neuron, 78*(5), 910-922. doi: 10.1016/j.neuron.2013.04.010

Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla-Wagner model. *Psychol Bull, 117*(3), 363-386.

Morelli, F. & Burton, P.A. (2009). The impact of induced stress upon selective attention in multiple object tracking. *Military Psychology, 21*(1).

Moustafa, A. A., Cohen, M. X., Sherman, S. J., & Frank, M. J. (2008). A role for dopamine in temporal decision making and reward maximization in parkinsonism. *J Neurosci, 28*(47), 12294-12304. doi: 10.1523/jneurosci.3116-08.2008

Murty, V. P., Ballard, I. C., Macduffie, K. E., Krebs, R. M., & Adcock, R. A. (2013). Hippocampal networks habituate as novelty accumulates. *Learn Mem, 20*(4), 229-235. doi: 10.1101/lm.029728.112

Murty, V. P., LaBar, K. S., Hamilton, D. A., & Adcock, R. A. (2011). Is all motivation good for learning? Dissociable influences of approach and avoidance motivation in declarative memory. *Learn Mem, 18*(11), 712-717. doi: 10.1101/lm.023549.111

Myrseth, H., Tvera, R., Hagatun, S., & Lindgren, C. (2012). A comparison of impulsivity and sensation seeking in pathological gamblers and skydivers. *Scand J Psychol, 53*(4), 340-346. doi: 10.1111/j.1467-9450.2012.00944.x

Nickel, A. E., Henke, K., & Hannula, D. E. (2015). Relational Memory Is Evident in Eye Movement Behavior despite the Use of Subliminal Testing Methods. *PLoS One, 10*(10), e0141677. doi: 10.1371/journal.pone.0141677

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

Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl), 191*(3), 507-520. doi: 10.1007/s00213-006-0502-4

Otmakhova, N. A., & Lisman, J. E. (1998). D1/D5 dopamine receptors inhibit depotentiation at CA1 synapses via cAMP-dependent mechanism. *J Neurosci, 18*(4), 1270-1279.

Paton, J. J., Belova, M. A., Morrison, S. E., & Salzman, C. D. (2006). The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature, 439*(7078), 865-870. doi: 10.1038/nature04490

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol Rev, 87*(6), 532-552.

Phelps, E. A. (2004). Human emotion and memory: interactions of the amygdala and hippocampal complex. *Curr Opin Neurobiol, 14*(2), 198-202. doi: 10.1016/j.conb.2004.03.015

Phelps, E. A., & LeDoux, J. E. (2005). Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron, 48*(2), 175-187. doi: 10.1016/j.neuron.2005.09.025

Powell, S. B., Paulus, M. P., Hartman, D. S., Godel, T., & Geyer, M. A. (2003). RO-10-5824 is a selective dopamine D4 receptor agonist that increases novel object exploration in C57 mice. *Neuropharmacology, 44*(4), 473-481.

Preuschoff, K., t Hart, B. M., & Einhauser, W. (2011). Pupil Dilation Signals Surprise: Evidence for Noradrenaline's Role in Decision Making. *Front Neurosci, 5*, 115. doi: 10.3389/fnins.2011.00115

Pyke, G. H., Pulliam, H. R., & Charnov, E. L. . (1977). Optimal foraging: a selective review of theory and tests. *Quarterly Review of Biology*, 137-154.

Rescorla, RA; Wagner, AR. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement *Classical Conditioning II: Current Research and Theory* (pp. 96-99). New York: Appleton Century Crofts.

Rigoli, F., Chew, B., Dayan, P., & Dolan, R. J. (2016). The Dopaminergic Midbrain Mediates an Effect of Average Reward on Pavlovian Vigor. *J Cogn Neurosci, 28*(9), 1303-1317. doi: 10.1162/jocn_a_00972

Robbins, T. W., & Everitt, B. J. (2007). A role for mesencephalic dopamine in activation: commentary on Berridge (2006). *Psychopharmacology (Berl), 191*(3), 433-437. doi: 10.1007/s00213-006-0528-7

Rotter, Julian B. (1990). Internal versus external control of reinforcement: A case history of a variable. *American Psychologist, 45*(4), 489-493.

Salamone, J. D., Correa, M., Nunes, E. J., Randall, P. A., & Pardo, M. (2012). The behavioral pharmacology of effort-related choice behavior: dopamine, adenosine and beyond. *J Exp Anal Behav, 97*(1), 125-146. doi: 10.1901/jeab.2012.97-125

Samuels, E. R., & Szabadi, E. (2008). Functional neuroanatomy of the noradrenergic locus coeruleus: its roles in the regulation of arousal and autonomic function part II: physiological and pharmacological manipulations and pathological alterations of locus coeruleus activity in humans. *Curr Neuropharmacol, 6*(3), 254-285. doi: 10.2174/157015908785777193

Sander, D., Grafman, J., & Zalla, T. (2003). The human amygdala: an evolved system for relevance detection. *Rev Neurosci, 14*(4), 303-316.

Sara, S. J., & Bouret, S. (2012). Orienting and reorienting: the locus coeruleus mediates cognition through arousal. *Neuron, 76*(1), 130-141. doi: 10.1016/j.neuron.2012.09.011

Schonberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *J Neurosci, 27*(47), 12860-12867. doi: 10.1523/jneurosci.2496-07.2007

Schott, B. H., Minuzzi, L., Krebs, R. M., Elmenhorst, D., Lang, M., Winz, O. H., . . . Bauer, A. (2008). Mesolimbic functional magnetic resonance imaging activations during reward anticipation correlate with reward-related ventral striatal dopamine release. *J Neurosci, 28*(52), 14311-14319. doi: 10.1523/JNEUROSCI.2058-08.2008

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *J Neurophysiol, 80*(1), 1-27.

Schultz, W. (2007). Behavioral dopamine signals. *Trends Neurosci, 30*(5), 203-210. doi: 10.1016/j.tins.2007.03.007

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science, 275*(5306), 1593-1599.

Schultz, W., & Dickinson, A. (2000). Neuronal coding of prediction errors. *Annu Rev Neurosci, 23*, 473-500. doi: 10.1146/annurev.neuro.23.1.473

Schwedes, C., & Wentura, D. (2012). The revealing glance: eye gaze behavior to concealed information. *Mem Cognit, 40*(4), 642-651. doi: 10.3758/s13421-011-0173-1

Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends Cogn Sci, 14*(10), 464-472. doi: 10.1016/j.tics.2010.08.002

Sokolov, E. N. (1963). Higher nervous functions; the orienting reflex. *Annu Rev Physiol, 25*, 545-580. doi: 10.1146/annurev.ph.25.030163.002553

Somerville, L. H., Whalen, P. J., & Kelley, W. M. (2010). Human bed nucleus of the stria terminalis indexes hypervigilant threat monitoring. *Biol Psychiatry, 68*(5), 416-424. doi: 10.1016/j.biopsych.2010.04.002

Spiers, A. S., & Calne, D. B. (1969). Action of dopamine on the human iris. *Br Med J, 4*(5679), 333-335.

Steinfels, G. F., Heym, J., Strecker, R. E., & Jacobs, B. L. (1983). Behavioral correlates of dopaminergic unit activity in freely moving cats. *Brain Res, 258*(2), 217-228.

Stephen, D.W., and Krebs, J.R. (1986). *Foraging theory*: Princeton University Press.

Strecker, R. E., & Jacobs, B. L. (1985). Substantia nigra dopaminergic unit activity in behaving cats: effect of arousal on spontaneous discharge and sensory evoked activity. *Brain Res, 361*(1-2), 339-350.

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *J Neurosci, 30*(33), 11177-11187. doi: 10.1523/JNEUROSCI.0858-10.2010

Voss, J. L., Gonsalves, B. D., Federmeier, K. D., Tranel, D., & Cohen, N. J. (2011). Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nat Neurosci, 14*(1), 115-120. doi: 10.1038/nn.2693

Voss, J. L., Warren, D. E., Gonsalves, B. D., Federmeier, K. D., Tranel, D., & Cohen, N. J. (2011). Spontaneous revisitation during visual exploration as a link among strategic behavior, learning, and the hippocampus. *Proc Natl Acad Sci U S A, 108*(31), E402-409. doi: 10.1073/pnas.1100225108

Wittmann, B. C., Schott, B. H., Guderian, S., Frey, J. U., Heinze, H. J., & Duzel, E. (2005). Reward-related FMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron, 45*(3), 459-467. doi: 10.1016/j.neuron.2005.01.010

Wu, M., Brudzynski, S. M., & Mogenson, G. J. (1993). Functional interaction of dopamine and glutamate in the nucleus accumbens in the regulation of locomotion. *Can J Physiol Pharmacol, 71*(5-6), 407-413.

Zalla, T., Koechlin, E., Pietrini, P., Basso, G., Aquino, P., Sirigu, A., & Grafman, J. (2000). Differential amygdala responses to winning and losing: a functional magnetic resonance imaging study in humans. *Eur J Neurosci, 12*(5), 1764-1770.

Zenon, A., Devesse, S., & Olivier, E. (2016). Dopamine Manipulation Affects Response Vigor Independently of Opportunity Cost. *J Neurosci, 36*(37), 9516-9525. doi: 10.1523/JNEUROSCI.4467-15.2016

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention Is Spontaneously Biased Toward Regularities. *Psychol Sci*. doi: 10.1177/0956797612460407

Biography

Nathaniel John Clement was born in Illinois in 1985 to Sue Ann Fieldman Clement and Kurt Thomas Clement. He received his B.A. in Psychology from Stanford University in 2007, where he worked with Elizabeth Race and Anthony Wagner. After working with Daphna Shohamy at Columbia University, he joined the Adcock Laboratory at Duke University in 2010. He was awarded a James B. Duke Gradute Fellowship from 2010-2014 and a National Science Foundation Graduate Research Fellowship from 2012-2015.