DISSERTATION

CATEGORICAL EVIDENCE, CONFIDENCE AND URGENCY DURING THE INTEGRATION OF MULTI-FEATURE INFORMATION

Submitted by

Kurt Braunlich

Department of Psychology

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2015

Doctoral Committee:

Advisor: Carol Seger

Charles Anderson
Matthew Rhodes
Lucy Troup

ABSTRACT

CATEGORICAL EVIDENCE, CONFIDENCE AND URGENCY DURING THE INTEGRATION OF MULTI-FEATURE INFORMATION

The present experiment utilized a temporally-extended categorization task to investigate the neural substrates underlying our ability to integrate information over time and across multiple stimulus features. Importantly, the design allowed differentiation of three important decision functions: 1) categorical evidence, 2) decisional confidence (the choice-inpendent probability that a decision will lead to a desirable state), and 3) urgency (a hypothetical signal representing a growing pressure to produce a behavioral response within each trial). In conjunction with model-based fMRI, the temporal evolution of these variables were tracked as participants deliberated about impending choices. The approach allowed investigation of the independent effects of urgency across the brain, and also the investigation of how urgency might modulate representations of categorical evidence and confidence. Representations associated with prediction errors during feedback were also investigated. Many cortical and striatal somatomotor regions tracked the dynamical evolution of categorical evidence, while many regions of the dorsal and ventral attention networks (Corbetta and Shulman, 2002) tracked decisional confidence and uncertainty. Urgency influenced activity in regions known to be associated with flexible control of the speed-accuracy trade-off (particularly the pre-SMA and striatum), and additionally modulated representations of categorical evidence and confidence. The results, therefore, link the urgency signal to two hypothetical mechanisms underling flexible control of decision thresholding (Bogacz et al., 2010): gain modulation of the striatal thresholding circuitry, and gain modulation of the integrated categorical evidence.

ACKNOWLEDGEMENTS

During my time in graduate school, I have received support and encouragement from a great number of people. In particular, I would like to thank my graduate advisors: Kristi Lemm and Carol Seger. Kristi you were encouraging and helpful, and my time working with you provided a foundation for this dissertation. Carol, thank you for taking me on as a graduate student, for giving me the freedom to pursue my interests, and allowing me the time to develop the skills needed to pursue my interests. Erik, I appreciated your always opinionated conversations and feedback. Hillary, thank you for being a great labmate. To my parents, thank you for the constant support and encouragement. Pop, I hope you would have enjoyed reading this. Finally, I would like to thank Leif and Natalie; this dissertation was typeset using Leif's document class for LaTeX.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Introduction

Decision making can be characterized as a deliberative process that involves weighing noisy samples of evidence for competing hypotheses until committing to a choice. For instance, when deciding about the direction in which a field of dots is moving, a decision maker must integrate information over time (Figure 1.1.A). Many real-world decisions additionally require decision makers to integrate information from different information sources according to their estimated reliability. For instance, after training participants to categorize circular sine-wave gratings into the two categories shown in Figure 1.1.B, we would expect that they would place heavy weight on the frequency dimension while tending to ignore orientation. Real-world decisions, of course, tend to be more complex. First, they tend to be higher-dimensional (e.g., categorical decisions such as "car" vs. "bike" sometimes require decision makers to consider more than "number of wheels;" e.g., Figure 1.1.C). Second, many features provide only probabilistic information. For instance, while dogs tend to be larger than cats, this is not always the case. To be able to reliably classify these kinds of complex stimuli, decision-makers must generally adopt strategies where they average informational content across multiple stimulus dimensions and across time.

It can often be advantageous, however, to reduce the dimensionality of complex problems by considering only a subset of the information. This has the well-known effect of reducing computational demands, and can sometimes improve decisional accuracy (Gigerenzer and Brighton, 2009; Gigerenzer and Gaissmaier, 2011; Guyon and Elisseeff, 2003; Tversky and Kahneman, 1974). Real-world demands often additionally require decision-makers to compromise accuracy in order to respond quickly, a phenomenon known as the *speed-accuracy*

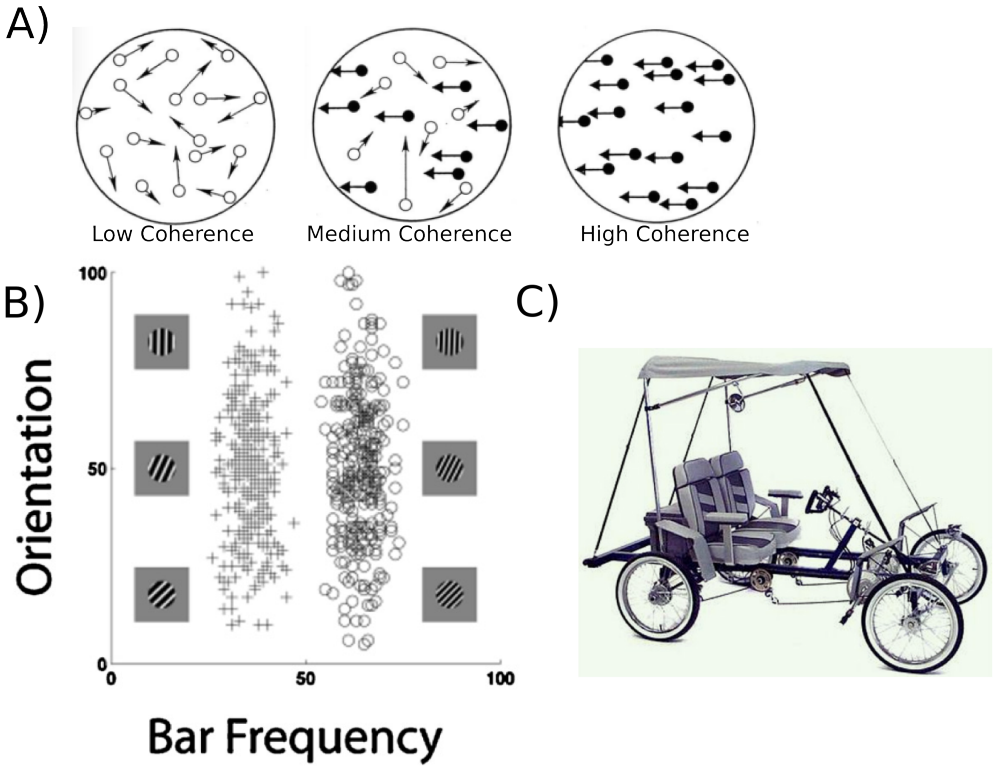

FIGURE 1.1. Categorization Task Examples. A) In the random dot motion task (Downing and Movshon, 1989), participants categorize fields of moving dots based on the overall direction of motion in the array. The difficulty of individual trials is inversely proportional to proportion of dots moving in a coherent direction. B) An example of a "rule-based" categorization task, where categories are defined according the frequency of sine-wave stimuli, and where the orientation of the individual exemplars does not provide reliable categorical information. Example stimuli are presented to the right and left of the plot illustrating the bivariate category distributions. C) An ambiguous bike-car stimulus. (A) is adapted from Downing and Movshon 1989, and B) is adapted from Smith et al. 2012)

trade-off (SAT). For instance, although the consequences of either choice might be unknown, a skier heading for a tree must commit to turn left or right within a limited period of time. In order to make advantageous decisions within such environments, decision makers must keep track of the probability that each choice will lead to a desirable state, and they must also track the time available to respond. Finally, in complex environments, decision-makers must also track their decisional confidence, the choice-independent probability of reaching

a desirable state. This allows decision-makers to plan sequences of choices in the absence of feedback, and to determine the value of accumulating additional information (when confidence is low, the value of considering additional information tends to be high). Finally, confidence also provides a basis for the calculation of reward prediction error, an important learning signal in the brain (Schultz et al., 1997).

Several lines of investigation suggest that *"urgency,"* a monotonically-increasing signal representing an increasing pressure to respond within each trial, may allow decision makers to commit to choices with successively less evidence as costs associated with deliberation accrue, or as opportunities for reward decrease (Cisek et al., 2009; Niyogi and Wong-Lin, 2013; Reddi and Carpenter, 2000; Schultz et al., 1997; Standage et al., 2011; Thura et al., 2012, 2014; Thura and Cisek, 2014). In the present experiment a temporally-extended version of the "weather-prediction" categorization task (Knowlton et al., 1994, 1996), was used to track the evolution of categorical evidence and decisional confidence as they evolved prior to commitment, and to investigate how these functions were modulated by urgency. Before our task and predictions are described in greater detail, several relevant categorization models are briefly reviewed. These models provide a framework for considering the integration of information across multiple stimulus dimensions. Particular focus is given to the COVIS model (Ashby et al., 2011), which currently provides the most biologically-detailed account of human classification performance. The sequential sampling and attractor-network classes of decision-making model are then reviewed. These model classes represent powerful frameworks for understanding how we integrate information over time, and how we threshold our decision processes to determine how much information we should accumulate before committing to a choice.

1.1. Categorization and Attention

The idea of modeling stimuli as a collection of attributes in multidimensional space has a long history in cognitive psychology. For instance, according to the Shepard-Luce choice rule (Luce, 1959; Shepard, 1957), a stimulus is represented as a set of points, x, and the distance between two K-dimensional stimuli (e.g., x_i and x_j) can be estimated by summing their unsigned differences across dimensions:

$$d_{ij} = \sum_{k=1}^{K} |x_{ik} - x_{jk}| \tag{1}$$

While this modeling approach provides a good account of the confusion data obtained from tasks requiring the identification of specific stimuli (where individual stimuli are associated with individual responses), it provides a poor account of confusion data obtained from categorization tasks (where groups of stimuli are associated with a common response). Shepard (1961) interpreted this finding as providing evidence that identification and categorization are based upon fundamentally different cognitive principles. This agrees with the logical assertion that the identification of individual stimuli must require the consideration of many stimulus dimensions (exemplars must be differentiated from all other stimuli), while categorization only requires the consideration of those dimensions which differentiate category members (exemplars associated with a common response) from non-category members (exemplars associated with other available response options). Nosofsky (1986), however, demonstrated that by modifying Equation 1 so that each psychological dimension was modulated by an attentional weight, w_k, the approach could also provide a reasonable account for confusion data obtained from categorization tasks:

$$d_{ij} = \sqrt{\sum_{k=1}^{K} w_k |x_{ik} - x_{jk}|^2} \tag{2}$$

This simple modification captures another long-standing idea in cognitive science: that rule-based strategies are equivalent to assigning exclusive attentional weights to a limited number of dimensions according to a strategy that is easily verbalizable (Bruner et al., 1956). This is also the case for many contemporary theories (e.g., Ashby et al., 2011; Johansen and Palmeri, 2002; Nosofsky et al., 1994; Paul and Ashby, 2013). For instance in the COVIS model (Ashby et al., 2011; Paul and Ashby, 2013), unidimensional rule-based strategies can be learned through a hypothesis-testing strategy wherein participants develop simple, verbalizable hypotheses for attentional weighting schemes, and then test them through trial and error. Conjunctive rules can be constructed by combining unidimensional weights through verbalizable Boolean logic. For instance, to categorize the stimuli shown in Figure 1.2.A, participants might begin by testing unidimensional strategies (e.g, by placing attentional weights on only the orientation dimension), but then through trial and error, learn to use the optimal rule: *"when the orientation is less than X and the frequency is greater than Y, choose category A; otherwise, choose B."*

1.1.1. Dissociable Categorization Systems. Whereas single-system approaches emphasize the flexibility and parsimony of Equation 2 to account for a variety of psychological phenomena, recent research suggests that different categorization strategies may place dissociable demands on neural systems (Ashby and Maddox, 2005, 2011; Maddox et al., 2004b,a; Maddox and Ashby, 2004). $Rule-based$ categorization tasks are thought to rely on working memory to maintain the candidate hypotheses, and to recruit neural regions associated with

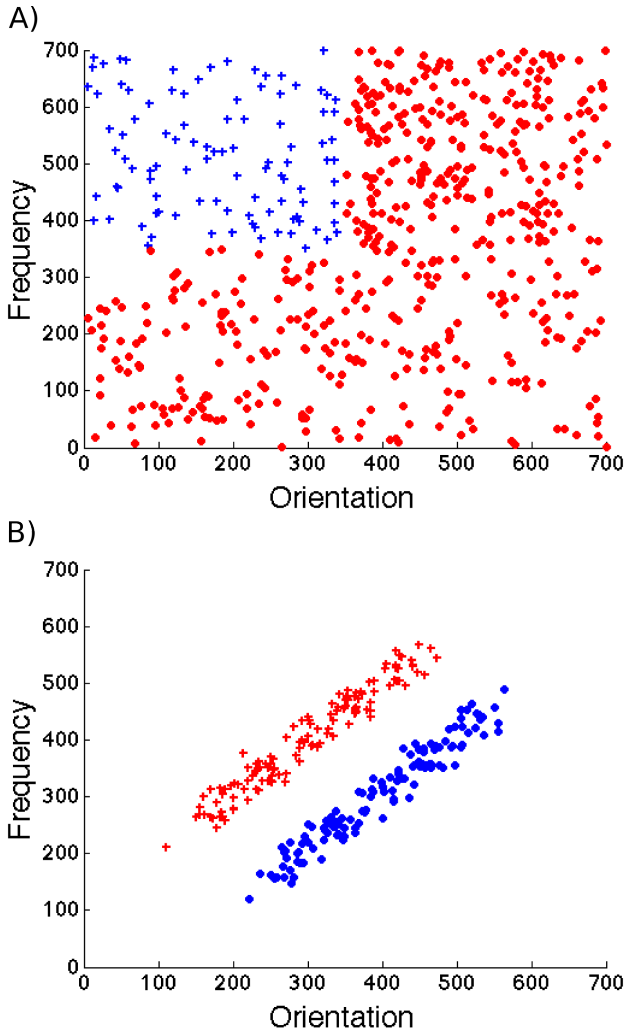

FIGURE 1.2. Rule-Based and Information-Information Categorization Tasks: Examples of bivariate distributions for stimuli that vary on two dimensions: frequency and orientation (e.g., the sine-wave gratings illustrated in Figure 1.1.B). A) A conjunctive rule-based categorization task. Rule-based categorization tasks are thought to rely on explicit hypothesis testing strategies. B) An example of an information-integration categorization task, which are thought to recruit the procedural learning system.

working memory storage (e.g., DLPFC) and with rule-selection (e.g., anterior cingulate and caudate head). The salience of a specific rule depends on its past success; on each trial, the salience of a particular rule is increased if its application allows the observer to transition to a desirable state, and its salience is decreased if its application results in a transition to an undesirable state. Switching between rules is thought to be mediated by an attentional shift

from the active rule to another via a corticostriatal loop involving the anterior cingulate and head of the caudate nucleus (Ashby and Maddox, 2005; Ashby et al., 1998; Helie et al., 2010). *Information-integration* tasks (Figure 1.2.B), however, are thought to rely on the procedural learning system. These tasks typically require participants to integrate information from multiple incommensurable stimulus dimensions. They typically require more time to learn than rule-based tasks, and, once learned, participants typically have difficulty verbalizing the strategy underlying their behavioral performance (Maddox et al., 2003, 2008, 2010). Posterior regions of the striatum (particularly the tail of the caudate and posterior putamen, which are reciprocally connected with high-level visual and motor areas) are thought to play an important role in information-integration category learning (Cantwell et al., 2015; Waldschmidt and Ashby, 2011).

An attractive feature of the COVIS model is that it provides a biologically-plausible framework underlying the key computation outlined in Equation 2: that of integrating information across multiple stimulus dimensions according to their estimated reliability. Like many other models of the procedural learning system, it posits that the the key site of reward-mediated learning is the corticostriatal synapse. Specifically, it posits that the learning of category structure is mediated by the visual corticostriatal loop (and specifically the corticostriatal synapses linking visual cortical regions to the body and tail of the caudate). Formally, COVIS posits that striatal units, S, integrate information from a set of perceptual units, K, each of which is maximally sensitive to a particular stimulus. The activation of a particular striatal unit, S_j, on trial n is determined by the activity of these inputs, I_K, and by the strength of the corticostriatal synapses w_{Kj} connecting them. This is a compelling

idea, as each medium spiny projection neuron of the striatum is known to receive converging glutamatergic input from nearly 20,000 cortical axons (Kincaid et al., 1998; Zheng and Wilson, 2002), and are thus anatomically well-situated to integrate and classify patterns of cortical input. For simplicity, COVIS reduces the number of perceptual neurons projecting to each striatal unit to 625, each of which is maximally sensitive to a particular stimulus in a 25 x 25 two-dimensional stimulus space (e.g., the 700 x 700 stimulus space illustrated in Figure 1.2).

$$S_j(n) = \sum_{k=1}^{625} w_{kj}(n) I_k(n) \tag{3}$$

The synaptic weights, w_{kj}, are updated after each trial according to the phasic dopaminergic reward prediction error signal (Schultz et al., 1997), which is proportional to the difference between expected and observed rewards:

$$RewardPredictionError = ObservedReward - PredictedReward \tag{4}$$

Thus positive reward prediction errors, which occur when reward is greater than expected, strengthen the synaptic weights, w, at synapses characterized by pre and post-synaptic excitation, while negative prediction errors, which occur when reward is smaller than expected, have the opposite effect. Thus, COVIS provides a framework where, through trial and error, the procedural learning system incrementally learns to parcellate multidimensional perceptual space into regions associated with specific behavioral or cognitive responses. It does so by strengthening the synaptic connections that have lead to unexpected reward, and by

weakening synaptic connections that lead to unexpected loss (or lack of reward). Recent works supports the assertion that dimensional weights can be learned without reference to a specialized representation learning system (i.e., model-based processes), and can be learned directly through this hypothesized corticostriatal mechanism (Niv et al., 2015).

The distinction between rule-based and information-integration categorization strategies mirrors the more widely-known distinction between *goal-directed* (or model-based) and *habitual* (or model-free) strategies underlying performance in instrumental conditioning paradigms. In goal-directed learning, animals use knowledge about the contingencies between actions and outcomes to select appropriate behaviors (Yin and Knowlton, 2006; Dickinson, 1985). Goal-directed behaviors can be formalized using model-based reinforcement learning algorithms that potentially implement forward (or "tree") searching (Gläscher et al., 2010) or mental-simulation, (Hassabis and Maguire, 2007) to maximize rewards and minimize punishments. Functional neuroimaging studies have found caudate activity consistent with these model-based computations (Gläscher et al., 2010; Wunderlich et al., 2012). Habitual behavior, however, is elicited without consideration of outcome or reward expectancy, based on stimuli in the environment (Gläscher et al., 2010; Wunderlich et al., 2012). Accordingly, dissociations between systems are often assessed via reward devaluation (e.g., lever pressing is considered habitual when a rat will continue to press the lever for food even after it has become satiated). Other researchers have suggested that habitual behavior should be resistant to interference from a simultaneously performed dual task (Ashby et al., 2010), or require extensive training (Mishkin et al., 1984; Waldschmidt and Ashby, 2011). Functional neuroimaging studies have found sensorimotor putamen activity consistent with these model-free computations (Miyachi et al., 2002).

1.1.2. The Weather Prediction Categorization Task. A popular paradigm for investigating how participants integrate probabilistic information across different information sources is the weather prediction categorization task (Knowlton et al., 1994, 1996), which requires participants to make predictions (e.g., "rain" vs. "sun") based on the information provided by multiple partially-informative cues (Figure 1.3). Early work by Knowlton and colleagues (Knowlton et al., 1994, 1996) suggested that, relative to healthy controls, amnesic patients were unimpaired during the early stages of learning, but showed impairments later in training (despite impaired declarative memory for the training session). Patients with Parkinson's disease, however, were unable to learn the task (despite intact declarative memory for the training session). Researchers at the time interpreted these results as evidence that probabilistic categorization may not rely on declarative knowledge, but may instead rely on the procedural learning system.

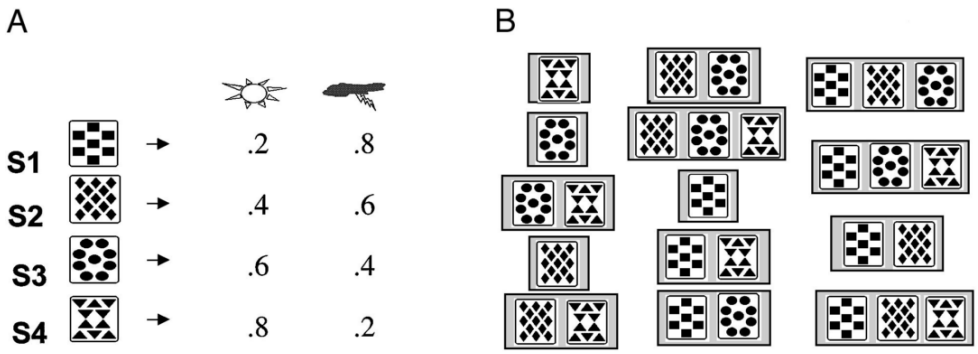

Figure 1.3. The Weather Prediction Task. A) Each of four cards provides a different amount of probabilistic evidence towards either a "sun" or a "rain" response. B) On each trial, participants are shown a set of cards, and are asked to make a prediction about the correct response. Although a range of suboptimal strategies are possible (see: Gluck et al., 2002; Meeter et al., 2008), accurate categorization generally requires integration of probabilistic information across cues. (Image from Shohamy et al., 2004)

Subsequent research, however, has suggested that the story is more complicated. For instance, task-knowledge tends to at least be partially declarative (Lagdano et al., 2006;

Meeter et al., 2006). In support of this conjecture, Poldrack et al. (2001) noted that the anterior hippocampus, a region that has commonly been associated with declarative memory, was recruited during the early stages of training. However, they also noted that the task recruited regions of the basal ganglia, which tended to negatively covary with hippocampal activity. Thus, one possibility is that the task recruits both declarative and procedural components; the anterior hippocampus may be required to set up the initial representations that may then be utilized by other memory systems (Meeter et al., 2008).

1.2. Decision Thresholding

Real-world decisions require decision-makers to not only deliberate about the best possible choice, but to adjust the temporal dynamics of this process to minimize costs and maximize gains. In this section, the well-known sequential sampling class of decision-making model (SSM) is reviewed. Then, the attractor-network class of model is discussed. This model class instantiates the algorithmic processes outlined in the SSM class within a biologically-plausible framework. Both model classes provide useful frameworks for considering the temporal dynamics of decision-making, and for considering how we flexibly adjust decision thresholds to maximize transient opportunities to transition to desirable states.

1.2.1. Sequential Sampling Models.

Sequential sampling models have been particularly successful in accounting for the temporal dynamics of decision-making behavior. Although they differ in terms of how information is accumulated, they typically posit that information is accumulated from some baseline level, until a decision threshold is crossed. In the Drift Diffusion model (DDM; Figure 1.4; Palmer et al., 2005; Ratcliff, 1978; Smith and Ratcliff, 2004) the strength of decision evidence modulates the speed with which a single decision variable (representing accumulated evidence) diffuses towards a decision boundary,

and the speed with which decision-makers commit to a decision is modulated both by the strength of this evidence, and the distance between the starting point of the accumulation process and a decision threshold. In race models, two separate accumulators "race" to a threshold. Race models have the advantage that they can account for multiple (> 2) response options and can more closely mirror known biological principles (i.e., many regions of the brain are known to accumulate information for specific choices; see Gold and Shadlen, 2007, for review).

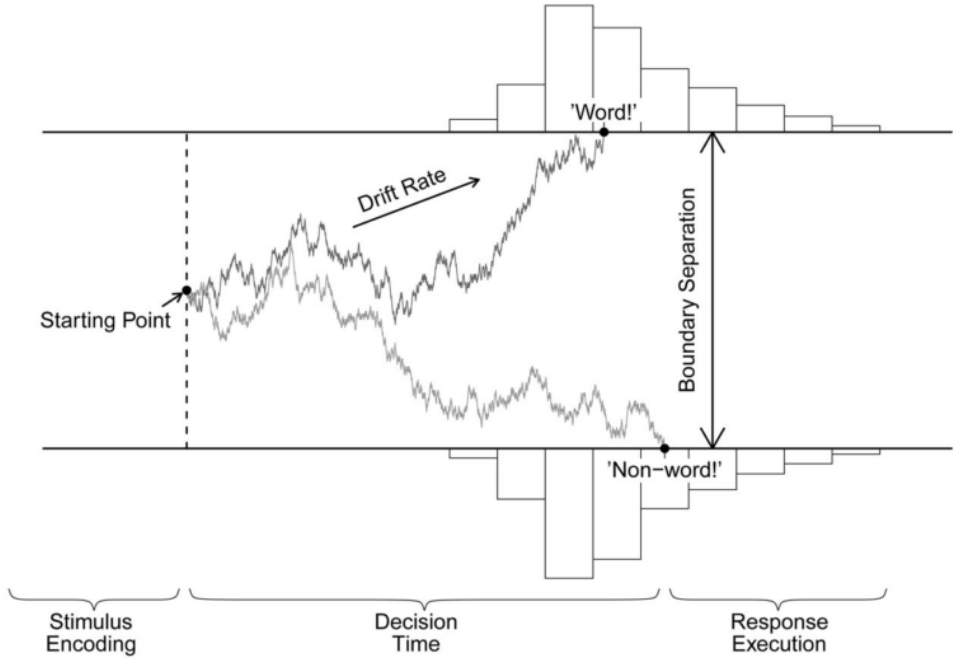

FIGURE 1.4. The Drift Diffusion Model. In this example, participants make decisions about whether a stimulus is a "word" or a "non-word." The DDM decomposes the decision process into *drift rate* (the speed with which information is integrated), *boundary separation* (i.e. decisional "caution"), the *starting point* of the accumulation process (which can account for decisional bias) and *non-decisional time* (e.g., stimulus encoding and response execution). (Image from van Ravenzwaaij et al., 2012)

One reason the SSM framework is attractive is that it provides a compelling characterization of neural activity observed in many regions of the brain. One particularly influential body of research examined perceptual decision making in the domain of motion direction

judgments. In the random dot motion task (RDM; Figure 1.1.A), participants must judge the direction of motion of a field of dots, and the difficulty of the task is manipulated by varying the percent of dots moving in a coherent direction. A well-known finding is that error rates and reaction times are negatively correlated with motion coherence. Interestingly, activity in many regions is well characterized by this accumulate-to-bound framework. For instance, the activity of neurons in the primate lateral intraparietal area (LIP) is known to track evidence for specific behavioral responses. It also tends to show a stereotyped pattern at the time of behavioral response. These findings suggest that these neurons may represent an important component in a biological framework where evidence is accumulated until a decision threshold is crossed (for review, see: Gold and Shadlen, 2007). Specifically, in the random to motion task, these "accumulator" LIP neurons seems to integrate information across multiple motion-sensitive neuronal pools in the middle temporal area (area MT). Each of these sensory pools is tuned to a particular motion direction (Britten et al., 1992), and in binary decisions, the activity of LIP accumulator neurons is well characterized as the unsigned difference between opposing MT neural populations (Roitman and Shadlen, 2002; Shadlen and Newsome, 1996). More recent studies have shown that LIP neural activity reflects *learned* categorical structure (Fitzgerald et al., 2011; Freedman and Assad, 2006; Swaminathan and Freedman, 2012), and human studies have found that accumulator-related activity is not limited to the human homologue of LIP (the intraparietal sulcus) but exists in a variety of cortical regions including adjoining parietal regions (inferior parietal cortex), pre and primary motor cortices, insula, and lateral frontal regions (Bernier et al., 2012; Gluth et al., 2012; Ho et al., 2009; Kayser et al., 2010b; Ploran et al., 2007, 2011).

Interestingly, regions which track categorical evidence appear to depend on specific task demands. For instance, when the mapping between category and response is known during deliberation, regions associated with motor-preparation (e.g., premotor and primary motor cortices, LIP/IPS) tend to track categorical evidence (Gluth et al., 2012; Thura et al., 2012; Thura and Cisek, 2014). However, these regions appear to shift based on manipulations of effector modality (Tosoni et al., 2008), while other regions (e.g., anterior insula) tend to accumulate information regardless of specific effector modalities (Ho et al., 2009).

SSM models have also allowed researchers to overcome the temporal limitations of fMRI to investigate distinct components of the decision-making process. For instance, Kayser et al. (2010a) used an accumulator model (Palmer et al., 2005) to investigate whether any regions of the brain might covary with the speed with which perceptual information accumulates in a random dot motion task (see Figure 1.1.A). Their results confirmed that the intraparietal sulcus (IPS; the homologue of the primate LIP), like the LIP, is sensitive to the rate at which information accumulates during motion judgments. Interestingly, they additionally found that the strength of functional connectivity (beta-series correlation; Rissman et al., 2004) between MT (a regions strongly associated with motion perception) and the IPS negatively covaried with the normative strength of perceptual evidence (i.e., motion coherence). Conversely, the connectivity between area MT and a more inferior region of the posterior parietal cortex positively covaried with motion strength. This finding suggests that both regions of the parietal lobe are sensitive to decisional evidence, but may play opposing roles during deliberation.

Similarly, Forstmann et al. (2008) used fMRI in conjunction with a race model (the linear ballistic accumulator; Brown and Heathcote, 2008) to estimate trial-by-trial adjustments of

response caution. Interestingly, they found that the activity of the striatum and the pre-SMA negatively covaried with estimates of decisional caution. Using a similar experimental paradigm and modeling approach, Forstmann and colleagues (2010) found that individual differences in the strength of anatomical connectivity between the striatum and the pre-SMA predicted individual differences in the capacity to flexibly-adjust the decision threshold to match environmental incentives for speed vs. accuracy.

1.2.2. Cortical Attractor Dynamics. Although the SSM framework successfully models choice behavior and provides a compelling account of the computations involved in decision making, it provides an incomplete account of the underlying biological mechanisms. Many of the attractive qualities of the SSM framework, however, have been successfully instantiated (for formal comparison, see: Bogacz et al., 2006) in biophysically-based models that characterize the decision process according to the attractor dynamics of competitive category-selective neuronal pools (e.g., Deco and Rolls, 2006; Furman and Wang, 2008; Standage et al., 2011; Usher and McClelland, 2001; Wong and Wang, 2006). Within this framework, slow recurrent excitatory connections within neuronal pools support the integration of transient sensory information over time, and mutual inhibition between pools assures categorical selectivity (i.e., "winner take all" dynamics; for reviews, see: Deco et al., 2012; Wang, 2008).

The duration of information representation within each category-selective pool emerges naturally from the balance between recurrent excitation within pools and mutual inhibition between pools (Figure 1.5). When inactive, recurrent excitatory effects are weak, and the network is characterized by spontaneous activity. When excitatory and inhibitory effects are appropriately balanced, the network can integrate information over time (and can do

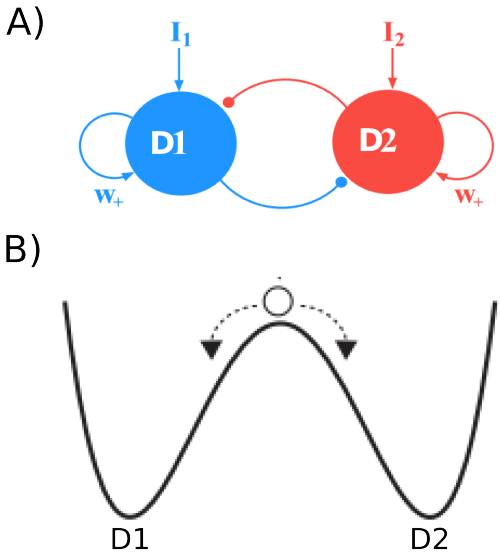

FIGURE 1.5. Attractor Dynamics. A) A schematic of a 2-variable attractor network (Wong and Wang, 2006), in which two populations of category-selective neuronal pools ($D1$ and $D2$) receive excitatory sensory innervation from different information sources (I_1 and I_2). The dynamics of the system are influenced by recurrent excitation (W_+), inhibition between pools, and by the sensory inputs. B) The network dynamics can be characterized as an energy landscape where a the motion of a particle across the landscape characterizes the deliberative process. $D1$ and $D2$ compete through inhibitory connections until one wins, at which point the system falls into a basin of attraction. When excitatory and inhibitory influences are balanced, the network can integrate information over time. (A) adapted from Wong and Wang, 2006; B) adapted from Standage et al. 2011).

so optimally, given constraints; Bogacz et al., 2006; van Ravenzwaaij et al., 2012). When recurrent excitation is sufficiently strong, the network falls into a basin of attraction, and the decision process tends towards commitment; new information is not integrated, but information related to the choice can be stored until a response is required (Cisek et al., 2009; Standage et al., 2011, 2013). Gain modulation of this recurrent excitation, therefore, is a simple biological mechanism that may underlie flexible adjustments of the speed-accuracy trade-off, and one that is common to several models that account for the effect (Furman and

Wang, 2008; Standage et al., 2011, 2013). In support of this framework, several groups have observed, within regions associated with the integration of categorical evidence, effects associated with a monotonically-increasing "urgency" signal, representing a growing pressure to produce a behavioral response within each trial (Churchland et al., 2008; Thura and Cisek, 2014; Thura et al., 2014).

While several groups have modeled the effect of urgency on decisional evidence as a multiplicative interaction between urgency and decision evidence (where urgency modulates the slope of activity in category-selective pools; Cisek et al., 2009; Ditterich, 2006; Thura et al., 2012), others have modeled it as an additive effect (where urgency increases baseline activity across all category-selective neuronal pools; Churchland et al., 2008; Gluth et al., 2012; Hanks et al., 2014). It is unclear whether the distinction between these functions is theoretically meaningful at the algorithmic level; both mechanisms may drive the system towards commitment. However, the distinction implies underlying mechanistic differences; an additive signal increases activity across all category-selective pools, while the multiplicative signal implies competitive inhibitory influences between pools.

1.2.3. The Role of the Basal Ganglia. Gain modulation within cortical accumulator circuits is likely only one component within a system underlying our capacity to flexibly adjust decision strategy; modulation of the cortico-basal ganglia circuit likely plays an important role as well (for review, see: Bogacz et al., 2010). Several features of the basal ganglia make it an attractive candidate for playing an important role in action selection and decision-making. First, its nuclei are somatotopically organized, such that different regions are associated with different body parts, and are subdivided such that different channels can represent different movements (Alexander et al., 1986). Second, the basal ganglia are known

to disinhibit salient cortical representations (Chevalier and Deniau, 1990), and have been associated with the selection of both motor responses (Redgrave et al., 1999; Mink, 1996) and higher level representations (McNab and Klingberg, 2008).

There are three primary pathways through which the basal ganglia interacts with cerebral cortex. These pathways exert their effects though opposing influences on thalamic excitation. The first two pathways to be identified were the *direct* and *indirect* pathways (Albin et al., 1989; Alexander and Crutcher, 1990). The GABAergic medium spiny projection neurons (MSN's) in the direct pathway (also sometimes referred to as the "Go" pathway) inhibit activity in the internal segment of the globus pallidus (GPi), attenuating its tonic inhibitory influence on the thalamus, and ultimately lowering the threshold required for the thalamus to become excited in response to other inputs (Chevalier and Deniau, 1990). The MSNs in the indirect pathway (also sometimes referred to as the "NoGo" pathway) inhibit activity in the external segment of the globus pallidus (GPe), thereby attenuating its tonic inhibitory effect on the GPi. The GPis tonic inhibitory influence on the thalamus is thereby increased, raising the threshold required to elicit thalamic activity. The direct and indirect pathways influence the threshold required for thalamic activity, but thalamic excitation seems to be driven primarily via direct projections from other regions. The medial dorsal thalamic nucleus, for instance, is primarily innervated by the prefrontal cortex (PFC), and the ventral lateral nucleus is primarily innervated by premotor regions (Haber and Calzavara, 2009). It should be noted that some thalamocortical cells, however, show post-inhibitory rebound bursting (Ulrich and Huguenard, 1997), suggesting that the basal ganglia output nuclei may have at least some direct excitatory influence on thalamic cells.

A third pathway, the *hyperdirect* pathway (Nambu et al., 2000), projects to the subthalamic nucleus (STN) and then to the internal segment of the globus pallidus (GPi), bypassing the striatum. This pathway is glutamatergic rather than GABAergic, and so cortically-driven excitation of the STN excites the GPi, ultimately inhibiting the thalamus. Projections from the STN to the thalamus are diffuse, and a great deal of research (e.g., Aron and Poldrack, 2006; Frank, 2006; Mink, 1996) suggests that while the direct and indirect pathways are responsible for thresholding specific input/output associations, excitation of the hyperdirect pathway instead raises the global thalamic excitation threshold such that stronger cortical innervation is required to produce any thalamic output. This global inhibition would likely be useful in situations in which it is desirable to withhold fast responding in order to improve accuracy. Excitation of the striatum may thus serve to lower decision thresholds by decreasing the tonic inhibitory influence of the basal ganglia output nuclei on specific salient cortical representations (Simen et al., 2006; Chevalier and Deniau, 1990), and excitation of the STN may conversely serve to increase global inhibitory effects (Ratcliff and Frank, 2012; Wiecki and Frank, 2013).

1.3. Confidence and Uncertainty

Even in the controlled experimental context provided by the random dot motion task (Figure 1.1.A), the decision maker must correctly make inferences across many forms of uncertainty. For instance, a decision maker must correctly filter out the noise associated with the visual stimulus, uncertainty about the appropriate response, and uncertainty about the consequences of the final choice. Our understanding of the brain suggests that, in many ways, it mirrors this functional hierarchy. For instance, neurons in area MT are known to be sensitive to the proportion of dots moving in a coherent direction (Britten et al., 1992),

neurons in the, parietal lobe (Newsome et al., 1989) and frontal eye fields (FEF) (Gold and Shadlen, 2000) are thought to integrate the uncertainty expressed across MT pools in the service of selecting an appropriate response (see Gold and Shadlen, 2007, for review), and neurons in the OFC may integrate uncertainty across neuronal pools to form an estimate concerning the probability that the selected option will be correct (Rolls et al., 2010b,a).

Neural computations are well-approximated by the rules of probabilistic inference, where the brain represents information in the form of probability distributions (Doya et al., 2007; Knill and Pouget, 2004; Pouget et al., 2013). Within this framework, the assignment of degrees of confidence to beliefs, whether they be beliefs about the causes of sensory input (e.g., MT), the consequences of actions (e.g., LIP and FEF), or the probability of reaching a desirable state (e.g., orbitofrontal cortex (OFC)), is a fundamental computational principle. In the present paper, however, "confidence" is reserved for *decisional* confidence, the choice-independent probability of reaching a desirable state (rather than perceptual confidence or categorical confidence). This matches its current use in the literature (Gherman and Philiastides, 2014; Hebart et al., 2014; Kiani and Shadlen, 2009; Rolls et al., 2010b,a).

Decisional confidence is thought to support several important functions. First, it motivates information seeking (Daw et al., 2006; Friston et al., 2012), such that uncertainty (the inverse of confidence) tends to be negatively correlated with the value of obtaining additional information. Second, in complex environments, it allows decision-makers to prepare for the consequences of their choices; this is particularly important in environments when sequences of actions must be chained without feedback. Thirdly, as confidence represents the degree of belief associated with reaching a desirable state, it plays an important role in the calculation of reward prediction error (eq. 4).

In the present paper, confidence was calculated from normative characteristics of the experimental design and from idiosyncratic patterns of participant behavior. This differs from the way confidence is typically operationalized in the metacognitive literature, where a goal is to dissociate subjective confidence from normative confidence. Thus, in the metacognitive literature, confidence is typically assessed through measures that are distinct from behavioral accuracy. For instance, in post-decision wagering tasks (Persaud et al., 2007; Seth, 2008), decision-makers are asked to indicate how much they would be willing to wager that were correct after making a choice. In "uncertain option" tasks (Hampton, 2001), decision-makers are given a low-valued, but certain option that can be an advantageous choice under conditions of low confidence. In "delayed incentives" tasks (Kepecs et al., 2008), decision makers make a choice, but must then wait before receiving a reward (the amount of time that they are willing to wait is assumed to be positively correlated with confidence). Thus, an important question in the domain of metacognitive research is how various biases mediate differences between subjective estimates of decisional confidence (i.e., subjective estimates of decisional accuracy or "type 2" sensitivity) and actual decisional accuracy (i.e., "type 1" sensitivity; Clarke et al., 1959; Fleming and Lau, 2014; Maniscalco and Lau, 2012).

Although this is an important question, before investigating the neural substrates underlying fine-grained differences between subjective and objective estimates of confidence, it is important to understand how objective confidence is represented in the brain. Following previous work (Gherman and Philiastides, 2014; Hebart et al., 2014; Rolls et al., 2010b,a), the current study used a model-based approach to identify latent representations of normative confidence in the absence of explicit behavioral measures. The approach is similar to Bayesian decision theory and signal detection theory, where a normatively optimal decision

variable, D (an estimate of the evidence for a particular choice), for option, A, given a stimulus, s, can be calculated as the log likelihood ratio:

$$D(s) = log\frac{P(s|A)}{P(s|B)} \tag{5}$$

where $P(s|B) = 1 - P(s|A)$.

Within this framework, an optimal choice would be to choose option A when $D(s)$ is greater than 0, and an estimate of confidence emerges as the unsigned decision variable: $|D(s)|$. Confidence thus represents the strength of decision evidence, without regard to a particular choice. negat different groups have used similar approaches to model confidence. Vickers and Packer (1982), for instance, suggested that, within a race model framework, confidence could be estimated as the unsigned difference between the winning and losing races (which they termed the "balance of evidence"). Rolls et al. (2010; 2010a; 2010b) found that regions of the human medial prefrontal cortex and cingulate cortex covaried with confidence and hypothesized that this representation might emerge from the unsigned difference in activity between category-selective attractor networks. In task requiring discriminations of faces vs. houses, Philiastides et al. (2010) found that activity in the ventromedial prefrontal cortex (VMPFC) covaried both with the unsigned difference in evidence between categories, and with the unsigned difference in activity between regions associated with perceptual evidence for the choice options "face" (posterior fusiform gyrus) vs. "house" (parahippacampal gyrus). Similarly, using direct cellular measures in behaving rats, Kepecs et al. (2008) found that that neurons in the orbitofrontal cortex were sensitive to the distance from an arbitrary category boundary in an odor classification task, and that this representation predicted the

amount of time that rats were willing to wait for reward. Additionally, using the random dot motion task with non-human primates, Ding and Gold (2010) found that distinct neural populations in the primate caudate nucleus tracked confidence and categorical evidence.

As decisional confidence negatively covaries with the value of seeking additional information, we would expect that decisional uncertainty would slow response times. In fact, it is well-known that response times negatively covary with decision difficulty (Vickers and Packer, 1982). Thus, as both confidence and urgency are thought to influence the dynamics of the deliberative process (shifting the system toward decisional commitment; Figure 1.5), we might expect to effects of these functions in overlapping regions of the brain. Additionally, we might expect that urgency multiplicatively modulates the slope of the confidence function if, as Rolls et al. have suggested (2010a; 2010b), confidence is calculated based on the difference in activity between category selective neuronal pools, and if, as experimental and computational work suggests, activity in these regions is (multiplicatively) modulated by urgency (Churchland et al., 2008; Cisek et al., 2009; Ditterich, 2006; Gluth et al., 2012; Hanks et al., 2014; Thura et al., 2012).

1.4. Overview of the Current Experiment

Many previous studies have used perceptual decision making tasks in conjunction with direct cellular recordings to track neural representations associated with decisional evidence. While these studies have provided important insights into how we integrate information over time, perceptual decision-making tasks are generally poorly suited to dissociate representations associated with evidence accumulation from those associated with decisional confidence and urgency; in these tasks, each of these functions increases monotonically with

time. To disentangle their neural representations, researchers must therefore employ experimental paradigms in which the strength of evidence for and against particular categories is modulated within single trials.

The present experiment uses a temporally-extended version of the weather-prediction categorization task (Knowlton et al., 1994, 1996) in which the features of an abstract "amoeba" stimulus were added to the display one-by-one over four discrete steps (Figure 3.1). Each feature was associated with a different amount of probabilistic evidence for a response with fingers of the left vs. right hands, and participants had to integrate this information across time and across distinct stimulus features (cf., Wheeler et al., 2014; Yang and Shadlen, 2007). This allowed the use of a model-based approach, wherein both the normative probabilistic characteristics of the experimental design, and measures based on participant behavior, were used to track the dynamical evolution of categorical evidence and confidence as participants deliberated about impending choices.

Categorical evidence was defined as the cumulative log likelihood ratio (CmLogLR) that a particular categorical response would lead to reward. Decisional confidence was modeled as the unsigned difference in evidence for each of the two categories (De Martino et al., 2013; Hebart et al., 2014; Kepecs and Mainen, 2012). Finally, urgency was modeled as a monotonically-increasing signal, which peaked at the last opportunity to earn reward (i.e., the onset of the fourth stimulus feature). In addition to tracking the independent evolution of urgency, the present research also investigated the possibility that urgency modulates the slopes of the neural representations of categorical evidence and confidence. In line with previous research, it was hypothesized that the integration of categorical evidence would recruit regions of the parietal lobe (Ploran et al., 2007, 2011), as well as primary and premotor

regions (Gluth et al., 2012; Pastor-Bernier and Cisek, 2011; Thura and Cisek, 2014; Wheeler et al., 2014). It was predicted that frontal regions, such as the VMPFC, as well as the OFC and dorsolateral prefrontal cortex would represent effector-independent decisional confidence (Cohen et al., 2007; De Martino et al., 2013; Heekeren et al., 2006; Padoa-Schioppa, 2011; Rolls et al., 2010a,b). It was also predicted that urgency would be represented in regions associated with modulation of the SAT (e.g., DLPFC, pre-SMA and striatum; Forstmann et al., 2008, 2010; van Veen et al., 2008). Urgency may additionally modulate representations of categorical evidence, either increasing activity across all neuronal pools (Cisek et al., 2009; Ditterich, 2006), or by modulating the slope of decisional evidence (Churchland et al., 2008; Gluth et al., 2012; Hanks et al., 2014). To investigate the neural substrates associated with these signals, an approach was adopted which allowed the investigation of the additive and multiplicative effects separately and in conjunction.

1.5. Overview of Remaining Sections

The remaining sections of the paper describe the pilot study and the primary study. The pilot utilized a similar experimental paradigm to estimate the number of participants needed to achieve sufficient statistical power. Additionally, based on the results from the pilot, temporal jitter was added to the final experimental paradigm, and the log-likelihood ratios were slightly modified to improve behavioral performance. After describing the behavioral and neuroimaing results from the primary experiment, future directions and limitations will be discussed.

CHAPTER 2

Pilot Study

The pilot experiment employed a temporally-extended categorization task in which seven different stimuli accumulated on the screen, each of which was associated with a different amount of probabilistic evidence for a response with the right or left hands. The design was similar to that used in the final experiment, however, participants were required to integrate probabilistic information across different stimuli, rather than across the features of a single stimulus (Figure 2.3 vs. Figure 3.1). Additionally, the inter-feature intervals were not jittered, and thus the model-based approach adopted in the final study was not used. Rather, in order to provide useful estimates of the number of participants required for the final study, the cumulative categorical evidence available on the last step was used for power analyses.

2.0.1. Participants. Three participants were recruited from the CSU community. All participants were healthy, right handed adults (2 female) with an average age of 24 (range 22-30). Participants were screened for history of psychiatric or neurological disorders, for current use of psychoactive medications, and for exclusionary MRI criteria. Based on low behavioral accuracy, one participant's data was excluded from the final analyses.

2.0.2. Stimuli. Stimuli were abstract shapes, each of which was associated with a different weight of probabilistic evidence towards responses with fingers of the left and right hands (the “d” and “k” keyboard locations, respectively). The logLR's assigned to the individual stimuli were 1.12, -1.12, .71, -.71, .32, -.32 and 0 (where positive weights arbitrarily

indicate evidence for a left response). To mitigate confounds associated with visual salience, the mapping between logLR and visual stimulus was randomized for each participant.

2.1. Training

At least a day in advance of scanning, participants trained on the individual stimuli in three blocks. On each trial, a single stimulus was presented in the center of the screen, and participants made a guess about whether it belonged in category "A" or category "B" (Figure 2.1). Feedback was disbursed according to Equation 7, and participants learned through trial and error. In the first block they trained on the stimuli with logLR equal to -1.12, 1.12, -.32 and .32. In the second block, they trained on the stimuli with logLR equal to -.71, .71 and 0. In the third block, they trained on all the stimuli at once. For each block, stimuli were presented in random order, and participants trained until they reach an 80% accuracy criterion. The criterion was polled every 150 trials, and was reset if participants were unsuccessful in reaching it.

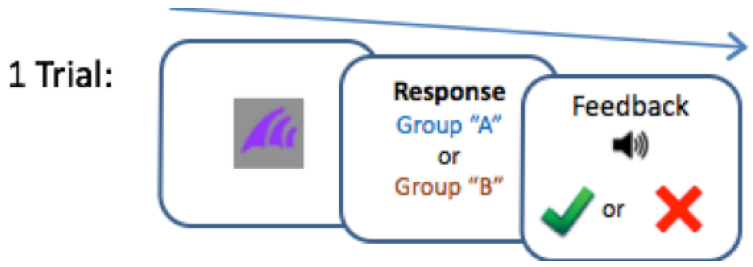

Figure 2.1. Training Task. On each trial, participants saw an abstract stimulus presented in the center of the computer screen. They made a guess about it's category membership, and then were given corrective feedback. Positive feedback consisted of a green checkmark and a pleasant tone, while negative feedback consisted of a red "X" and an unpleasant tone.

2.1.1. Post Decisional Wagering.

To allow estimation of the subjective weights of evidence that each participant placed on the individual stimuli, after training, participants

performed a post-decision wagering task (Persaud et al., 2007; Seth, 2008). In this task, participants categorized a stimulus, and then indicated how much they would be willing to bet that they were correct (Figure 2.2). Feedback was similar to the training task, however, participants also received information about the (hypothetical) money won or lost on each trial. Stimuli were presented in random order, and participants rated each exemplar three times. After performing the post-decisional wagering task, participants practiced the temporally-extended task that they would later perform in the scanner by completing one block of 120 trials.

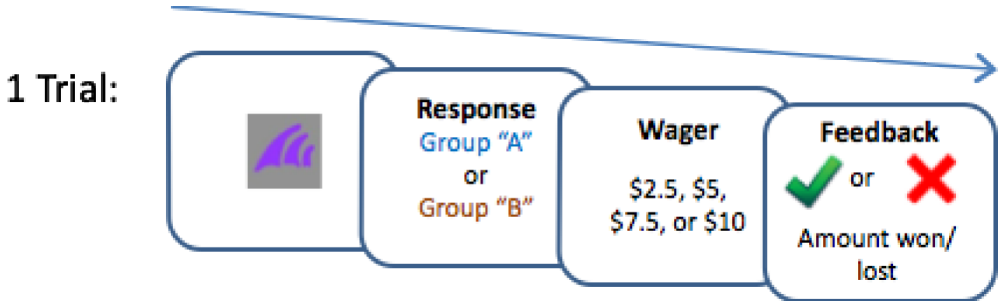

FIGURE 2.2. Post Decisional Wagering Task. On each trial, participants saw a stimulus, made a categorization decision, and then were asked to indicate how much they would like to bet that they were correct (options were $2.50, $5, $7.50 or $10). They were then given feedback as in the training task (figure 2.1), and were also told how much (hypothetical) money they had won or lost.

2.2. SCANNING

2.2.1. TASK. In the scanner, participants performed a temporally-extended version of the task wherein the stimuli accumulated on the screen in random (non-overlapping) locations, one-by-one, over seven discrete steps (Figure 2.3). Stimulus onsets coincided with the TR and were not jittered (for discussion of jitter, see section 3.2.2). Participants were free to respond at any time during the trial, and so the number of trials per session varied based on

behavioral performance. The optimal response for each trial was determined by the sign of the sum of the logLRs associated with the individual stimuli:

$$CmLogLR_{(step7)} = log_7 \frac{P(Left|S_{1:7})}{P(Right|S_{1:7})} = \sum_{i=1}^{7} w_i \tag{6}$$

where $S_{1:7}$ indicate the specific stimuli included in a particular trial, and w, the logLR assigned to each stimulus. Feedback was disbursed according to the corresponding probability:

$$P(Left|S_{1:7}) = \frac{7^{\sum_{i=1}^{7} w_i}}{1 + 7^{\sum_{i=1}^{7} w_i}} \tag{7}$$

where $P(Right)$ = 1 - $P(Left)$.

Through behavioral piloting, a points system was developed to balance demands for speed and accuracy. For each step of "revelation" (Ploran et al., 2007), the points available reduced by 2 points (22 points were available to win on the first step, while only 10 were available on the last step). To discourage participants from using strategies based on guessing, participants lost double the number of points they would have won if they were correct. The points were adjusted to maximize variability in response timing.

2.2.2. Image Acquisition. Data were acquired using a 3-tesla Siemens Magnetom Trio-Tim scanner. Stimuli were presented using the psychtoolbox for MATLAB (Brainard, 1997). Functional volumes were acquired in three sessions. Each session involved the acquisition of 560 whole brain volumes with an interleaved 2-dimensionsal echo planar imaging (EPI) sequence (28 slices, TR = 1,500 ms, TE = 25 ms. Flip angle=75 degrees). The first 4 volumes

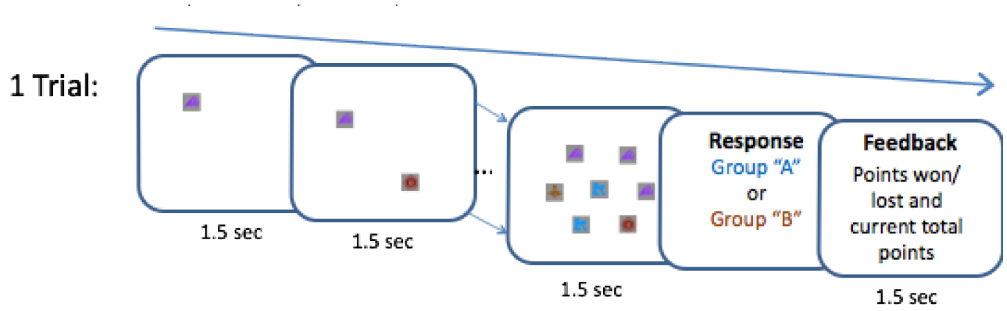

FIGURE 2.3. Scanner Task. In the scanner, participants performed a temporally-extended paradigm in which the stimuli accumulated one-by-one, until a maximum of seven stimuli were on the screen. Participants were free to respond at any time. The opportunity to earn a decreasing number of points on each step provided incentive for participants to respond earlier in the trial.

were discarded to allow for T1 equilibration effects. Anatomical images were collected using an magnetization-prepared rapid gradient-echo (MP-RAGE) sequence.

2.3. Analyses

2.3.1. Preprocessing. Preprocessing involved slice-timing correction to the tenth slice, motion correction, coregistration, segmentation of the anatomical images, and normalization of the structural and functional images to the Montreal Neuroligical Institute (MNI) template. Low frequency drifts in the timeseries were removed via a high pass filter with a cutoff frequency of 1/ 128 Hz. The normalized images were smoothed with a 6-mm full-width-at-half-maximum Gaussian kernel.

2.4. Results

2.4.1. Behavior. Average accuracy, determined relative to an optimal observer making decisions based on the sign of the cumulative logLR available on the last step was 77% correct (SD = 3.4%). Results from the post-decisional wagering task suggested that participants

had difficulty learning the stimuli with logLR's of -.36 and .36 (Figure 2.2). Accordingly, in the final experiment, these weights were increased to .56.

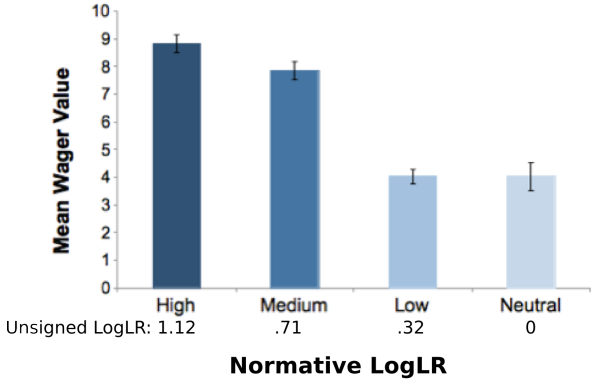

FIGURE 2.4. Post Decisional Wagering Results. The amount of money participants were willing to bet reflected the logLR's used to determine the probabilistic disbursement of feedback. Error bars represent the standard error of the mean.

2.4.2. POWER ANALYSES. The Powermap statistical software developed by Joyce and colleagues (2012) was used to estimate the number of participants needed in the final study. This software allows voxelwise estimation of the number of subjects needed to achieve a specified power level (>= 80%) while correcting for multiple comparisons (False Discovery Rate $q < .05$). The method is based on random field theory, which is also used by SPM (Wellcome Trust Centre for Neuroimaging; the software used for analyses in the final study). Notably, this analysis indicated that with fewer than 20 participants, it would be possible to achieve power of 80% for effects related to categorical evidence (for a response with the right hand) in the left SMA, left postcentral motor region, left caudate and right cerebellum. Conversely, activity in left cortical motor regions, and right superior parietal lobe exceeded this criterion for evidence towards a response with the left hand (Table A.1). Based partially

on these results and past research (Friston, 2012; Kayser et al., 2010a; Ploran et al., 2011), 20 participants were included in the primary experiment.

As it is known to inflate type 1 error rate, power analyses on the data from the primary data-set were not performed. For instance, if one were to run a power-analysis based on the first 2-3 participants of the primary experiment, the number of subjects in the final experiment becomes a free parameter that depends on the sample data. As standard statistic tests do not account for this effect, this procedure can inflate the type 1 error rate, and the magnitude of the effect depends on the proportion of final data included in the power analyses (Mumford, 2012). Posthoc power analyses were also not performed on data from the primary experiment, as these are considered to be uninformative (see Goodman and Berline, 1994; Hoenig and Heisey, 2001; Levine and Ensom, 2001; Mumford, 2012; Poldrack et al., 2011).

CHAPTER 3

Primary Experiment

3.1. Participants

Twenty right-handed participants (mean age = 24, SD = 4; 12 Female) were recruited from the Colorado State University and University of Colorado at Boulder communities. Participants were screened for history of psychiatric and neurological disorders, for current use of psychoactive medications and for exclusionary MR criteria. They were compensated at a rate of $20 per hour.

3.2. Task

Participants performed a temporally-extended, binary, probabilistic categorization task (Wheeler et al., 2014; Yang and Shadlen, 2007) wherein they categorized different "amoeba" (Figure 3.1) into one of two categories: "A" (indicated with a left hand response) and "B" (indicated by a right hand response). For mnemonic purposes, throughout the paper, the categories will be labeled according to the response with which they were associated (Category *"Left"* and Category *"Right"*). The amoeba stimuli consisted of a black outline upon which seven different features could appear. Features could be repeated, and could appear at four different locations. Presenting the cues as features of an amoeba was chosen for several regions. First, it was expected that presenting cues as features of single objects would encourage integration and representation as aspects of a single object. Second, it provided a suitable cover story for subjects, as biological kinds are often characterized by features that may occur probabilistically.

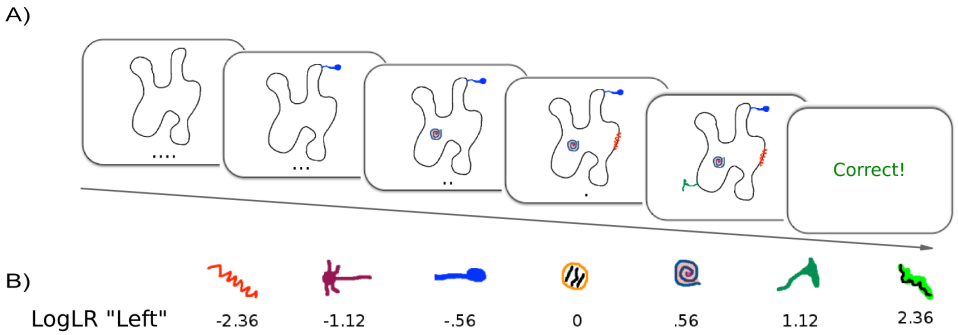

FIGURE 3.1. A) Trial Format. In the scanner, each trial began with the presentation of the blank amoeba profile. The features (abstract organelles, nuclei and flagella, each representing different amounts of probabilistic evidence towards a right or left response) then accumulated over four steps. Feature onsets were separated by a jittered interval. Participants were free to respond at any time during the trial. After a response, positive feedback was disbursed according to the probability corresponding to the sum of the logLR's associated with the four individual features (Equation 9). A series of small dots, presented below each amoeba, indicated the number of steps remaining in the trial. B) Although the association between visual feature and logLR was randomized for each participant, one possible mapping is illustrated. Positive logLR's (arbitrarily) represent evidence for a response made with the left hand.

At the beginning of each trial, the black outline, without any features, served as an ad hoc fixation point, and cued the beginning of each trial. The features then accumulated on the amoeba, one-by-one, over four discrete steps. Each feature provided different probabilistic evidence towards a specific response. The logLRs associated with the individual features were -2.36, -1.12, -.56, 0, .56, 1.12, and 2.36, where positive weights (arbitrarily) indicates evidence towards Category Left, and where 0 indicates no evidence towards either response. Participants were free to respond at any time during the trial, and were instructed to simply wait until they had enough information before doing so. To avoid possible confounds associated with visual salience, the mapping between logLR and visual feature was randomized for each participant. The optimal response for each trial was determined by the sign of the sum of the logLRs associated with the individual features:

$$CmLogLR_{(step4)} = log_7 \frac{P(Left|F_1, F_2, F_3, F_4)}{P(Right|F_1, F_2, F_3, F_4)} = \sum_{i=1}^{4} w_i \tag{8}$$

where $F_{(1:4)}$ indicate the specific features of a particular amoeba and w, the weights assigned to each feature. Feedback was disbursed according to the corresponding probability:

$$P(Left|F_1, F_2, F_3, F_4) = \frac{7^{\sum_{i=1}^{4} w_i}}{1 + 7^{\sum_{i=1}^{4} w_i}} \tag{9}$$

where $P(Right)$ = 1 - $P(Left)$.

3.2.1. Training. Participants performed two training sessions occurring on separate days. The three goals of the first training session were 1) to teach participants about the logLRs associated with the individual features, 2) to teach them how to integrate information across features, and 3) to give them experience with the temporally-extended task that they would perform in the scanner. To allow participants to quickly gain experience with a large number of experimental trials, in the first training task, the temporally-extended paradigm described above was not used, instead all features were presented at the same time. Through trial and error, participants learned to categorize amoeba with one, two, three, and then all four features. On each trial, an amoeba and its feature(s) were presented, and participants made a decision about whether it belonged in the Left or Right categories. Following correct responses, positive feedback was disbursed according to the probability corresponding to sum of the four individual logLRs (eq. 8; the word "Correct" was presented in green for .75 seconds, and was accompanied by a pleasant tone). Following incorrect responses (or trials

in which the probabilistic feedback schedule mandated it), negative feedback was disbursed (the word "Wrong" was shown for .75 seconds in red, and was accompanied by an unpleasant tone).

Participants trained until they reached an 80% accuracy criterion twice (accuracy was determined in relation to an optimal classifier making decisions based on the sign of the sum of the logLR's associated with the individual features; the criterion was polled every 30 correct trials, and was reset after 35). After completing this initial training task, participants then practiced the temporally-extended paradigm that they would later perform in the scanner (described above). The second training session, which occurred on the same day as scanning, was identical to the first, but participants began the session by training with four features that were presented simultaneously. After reaching the 80% accuracy criterion (twice), they again practiced with the temporally-extended paradigm.

3.2.2. Scanning. In the scanner, the task was similar to the temporally-extended task included during the training session, but both the inter-feature and inter-trial intervals were jittered. The interval between each feature (and between the last step and feedback) was jittered according to a uniform distribution with a minimum of 2 seconds (to mitigate non-linear effects associated with shorter intervals; cf., Buckner, 1998; Friston et al., 2000), and a maximum of 4 seconds, in steps of 0.5 seconds (Gluth et al., 2012). The interval between each trial was jittered according to a positively-skewed geometric distribution ranging from 2-9 seconds.

The idea behind incorporating temporal jitter between events in fMRI is that it allows characterization of the shape of the hemodynamic response function (HRF) while maximizing the efficiency of the experimental design. The most popular HRF consists of the sum of two

gamma functions (Figure 3.2), which models the peak of the BOLD response as occuring roughly 4-6 seconds after the onset of a stimulus, an "undershoot" occuring roughly 12-16 seconds after stimulus onset, and a return to baseline shortly thereafter (Friston et al., 1999). Thus, as our inferences about neural activity are mediated by these slow hemodynamics, in order to recover the shape of the HRF, a researcher might decide to place stimulus onsets roughly 16 seconds apart (thereby allowing the BOLD response to return to baseline between successive events). Such a design would be highly ineffecient (as it would be characterized by a low ratio of number of trials to units of time). Perhaps more importantly, such a design would constrain the kinds of experimental questions researchers would be able to investigate. A more sophisticated approach, however, involves the use of "fast event" related designs, where the onsets of trials are placed more closely together. By jittering the onsets of the stimulus events in time, the researcher can sample different points of the HRF shape, and gain a more accurate representation of the true HRF associated with each experimental condition.

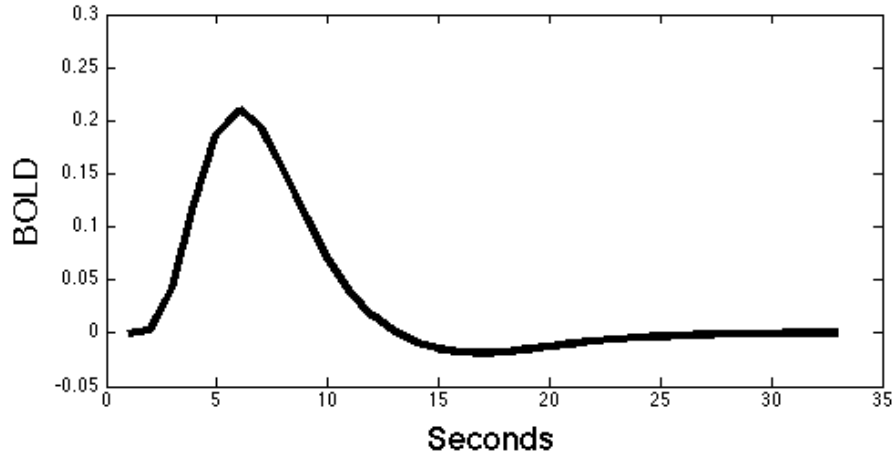

FIGURE 3.2. The canonical, double gamma, hemodynamic response function (HRF).

The fast event related approach depends crucially on the superposition principle that underlies all linear systems: that the HRF associated with multiple events is well-approximated by the sum of the HRF's associated with individual events. This generally holds when the intertrial interval is greater than 1 second (Dale and Buckner, 1997; Friston et al., 2000; Vazquez and Noll, 1998). However, when the interval trial is shorter, the BOLD response is known to display nonlinear dynamics (Hinrichs et al., 2000; Huettel and McCarthy, 2000; Wager et al., 2005). Specifically, the amplitude of the BOLD response tends to be attenuated, relative to the model predictions.

Participants performed 3 scanner sessions, each of which required 14 minutes. As participants were free to respond at any time during the trial, the number of trials per session varied based on participant behavior.

3.2.2.1. *Image Acquisition.* Data were acquired using a 3-tesla Siemens Magnetom TrioTim scanner with a 32 channel-head coil. Stimuli were presented using the psychtoolbox for MATLAB (Brainard, 1997). Each session involved the acquisition of 420 whole brain volumes with an interleaved EPI-2D sequence (TR = 2,000 ms, TE = 25 ms, voxel size: 2.3 x 2.3 x 3.5, flip angle = 75, GRAPPA acceleration factor 2). The first three volumes were discarded to allow for T1 equilibration effects. Anatomical images were collected using an MP-RAGE sequence.

3.3. Analyses

3.3.1. Preprocessing. Preprocessing involved slice-timing correction, motion correction, coregistration, segmentation of the anatomical images, and normalization of the structural and functional images to the MNI template. For the classical analyses, the normalized images were smoothed with a 6-mm full-width-at-half-maximum Gaussian kernel. For the

Random Effects Bayesian Model Selection (RFX-BMS) procedure (described below), the normalized unsmoothed functional images (down sampled to 3-mm isotropic voxels, to improve computational efficiency) were used to calculate the log-evidence maps for each model and each participant. The log-evidence maps were smoothed with an 8-mm full-width-at-half-maximum Gaussian kernel prior to group-level analyses. Preprocessing and classical (restricted maximum likelihood) data analyses were performed using Statistical Parametric Mapping (SPM8).

3.3.2. Model-Based Analyses.

3.3.2.1. *Overview.* A model-based approach was used to track the evolution of categorical evidence and confidence, and to investigate how these representations were modulated by urgency. As some of these models were correlated, and yielded highly similar statistical maps, in order to determine which models best accounted for the data, compared their relative fit using a random-effects Bayesian model selection procedure (RFX-BMS; Rosa et al., 2010; Stephan et al., 2009). This involves three steps. First, a classical, contrastive approach was used to identify voxels surviving standard statistical thresholds. Second, maps representing the log evidence ($logP(y|m)$, the log probability of the data, y, if it were generated by the model, m) were generated for each model and each participant using a variational Bayesian approach to integrate over model parameters. Third, a random effects analysis was used at the group level (Stephan et al., 2009) to calculate the protected exceedance probability (PXP; Rigoux et al., 2014), the probability that a given model is more frequent than any of the other models tested, above and beyond chance. Although calculation of the log evidence maps is computationally intensive, the RFX-BMS approach is attractive, as it provides an intuitive metric of model fit, allows comparison of non-nested models, and has some

favorable properties when compared to the Akaike Information Criterion and the Bayesian Information Criterion (Penny, 2012; Rigoux et al., 2014). For this analyses, the spm_BMS function included in SPM12 was used.

For both the classical and Bayesian analyses, it was assumed that neural activity associated with deliberation would continue across each jittered inter-feature interval, and so each feature step was modeled with a duration equal to the difference between its onset and the onset of the following feature or the behavioral response. Such variable-duration epoch models tend to be more sensitive for paradigms involving cognitive events of variable duration than constant epoch or variable amplitude impulse models (Grinband et al., 2008). The canonical SPM hemodynamic response function was used for all models. The topological false discovery rate (Chumbley and Friston, 2009) was used to correct for multiple comparisons (q ¡ .05, with an initial cluster forming threshold of p ¡ .001).

3.3.2.2. *Definition of Parametric Modulators (PMs).* As positive feedback was disbursed based on all evidence available at the end of the trial, it was not possible to accurately calculate cumulative logLR for steps 1-3 by summing the logLRs for these features (as in equation 8). Instead, these values were determined computationally, by tabulating the full (2401 x 4) permutation matrix and calculating the proportion that each partial feature sequence would be optimally categorized as Left. As evidence for the two categories was perfectly anticorrelated, confidence was calculated as the unsigned difference in evidence for the two categories (unsigned CmLogLR; cf., Fetsch et al., 2014; Hebart et al., 2014; Rolls et al., 2010a). Uncertainty was perfectly anticorrelated with confidence. Urgency was modeled as a linearly increasing signal that peaked on the fourth step (the last possible step to earn positive feedback). A linear function was used, as preliminary RFX-BMS analyses provided

evidence that it provided a better fit than logistic or exponential functions. To model the multiplicative effect of urgency on categorical evidence and confidence, the mean-centered evidence PM (CmLogLR or confidence) was multiplied by urgency. The multiplicative effect of urgency therefore modulated the slope of the evidence variables (categorical evidence and confidence) such that the difference between high and low values increased with urgency.

To fit the models to the strategies used by individual participants, the subjective weights of evidence (sWOE) that each participant placed on the individual features were estimated (Figure 3.3.B). A Bayesian logistic regression analysis (prior mean = 0 for all coefficients; Gelman et al., 2008) was employed to do so in a manner robust to separation. To calculate the cumulative subjective logLR, a similar procedure was used to calculate cumulative normative logLR; the full permutation matrix of beta coefficients was tabulated, and then the proportion that each sequence of features would be optimally categorized as Left was calculated. Calculation of the other subjective regressors also followed the process outlined for the construction of the normative regressors.

3.3.2.3. *Classical Analyses.* Four models (Table 3.1) were used to investigate categorical evidence and confidence as well as both additive and multiplicative effects of urgency. The first parametric modulator (PM1) tracked decisional evidence (categorical evidence or confidence) while the second PM (PM2) tracked urgency. As the second PM was orthogonalized with respect to the first, it represents an *additive* effect of urgency. The interaction term included in the first in the second and fourth models represent multiplicative effects of urgency (where urgency modulates the slope of the decisional evidence function). As positive logLR's (arbitrarily) indicate evidence for a "Left" response, and as evidence for the two categories was perfectly anticorrelated, the positive contrast revealed regions tracking evidence

for this response, while the negative contrast revealed regions tracking evidence for category "Right." Similarly, as confidence and uncertainty were also perfectly anticorrelated, it was possible able to track their representations via positive and negative contrasts, respectively.

TABLE 3.1. Classical Models. "CmLogLR": Categorical evidence. "Conf.": Confidence. "Urg.": Urgency. PM1: the first parametric modulator. PM2: the second parametric modulator.

Model Name	PM1	PM2	Positive Contrast	Negative Contrast
CmLogLR	CmLogLR	Urg.	"Left"	"Right"
CmLogLR* Urg.	CmLogLR* Urg.	Urg.	"Left" * Urg.	"Right" * Urg.
Conf.	Conf.	Urg.	Conf.	Uncertainty
Conf.* Urg.	Conf.* Urg.	Urg.	Conf.*Urg.	Uncertainty*Urg.

The default orthogonalization of regressors in SPM was left on (thus shared variance was assigned to the first PM while urgency was orthogonalized with respect to it). It was thus possible to investigate multiplicative effects of urgency via the first PM, and orthogonal, additive effects via the second. Parametric effects were considered only for correct trials, but incorrect trials were included in the design matrix as variables of no interest.

3.3.2.4. *RFX-BMS.* As there was considerable overlap between the statistical maps derived from the classical analyses, three RFX-BMS analyses were conducted. The first compared categorical evidence to categorical evidence (multiplicatively) modulated by urgency. The second compared confidence to confidence (multiplicatively) modulated by urgency. The third compared uncertainty to uncertainty (multiplicatively) modulated by urgency. For each model, one PM per mean regressor was included. Whereas the classical analyses supported use of a contrastive approach (which differentiates positive and negative correlations with particular regressors), the RFX-BMS analyses provide information about the ability of each *complete* model to account for variance in the BOLD signal. These analyses were limited

to voxels within a binary mask that included all significant voxels from the corresponding classical statistical maps (Gluth et al., 2014). This allowed me to compare these models within a single framework, to use these analyses to inform our interpretations of the classical statistical maps, and to limit inferences to voxels surpassing standard statistical thresholds. In Figure 3.11, each PXP map was thresholded to include only voxels where the model provided the best fit. These thresholded probability maps allow the reader to quickly identify these voxels, and to make inferences about the strength of evidence relative to the other models tested.

3.4. Behavioral Results

Average accuracy (determined relative to an optimal classifier making decisions based on the sign of the cumulative normative logLR on the last step) was 85% correct (SD = 6.34%). Participants were fairly conservative, and showed a strong tendency to wait until they saw all features before making a behavioral response (average response step = 3.77, SD = 0.25). A logistic regression analyses (described above) was conducted to confirm that participants used a strategy wherein they weighted the features according to their reliability rather than using a simpler strategy (e.g., counting the features belonging to each category, or considering only the most informative features; cf., Gluck et al., 2002; Meeter et al., 2006). In Figure 3.3.B, the normalized beta weights for each participant are plotted against the normalized normative logLR. The linear pattern and close correspondence between the two sets of estimates provides evidence that participants applied weights closely resembling the normative logLR used to determine the disbursement of feedback. A second logistic regression analysis, which was conducted to investigate whether participants placed greater

weights on particular steps (Figure 3.3.C), found no evidence of such an effect, $F(3,76) = 0.48$, $p = 0.7$; Bayes factor in favor of the null hypothesis = 8.87.

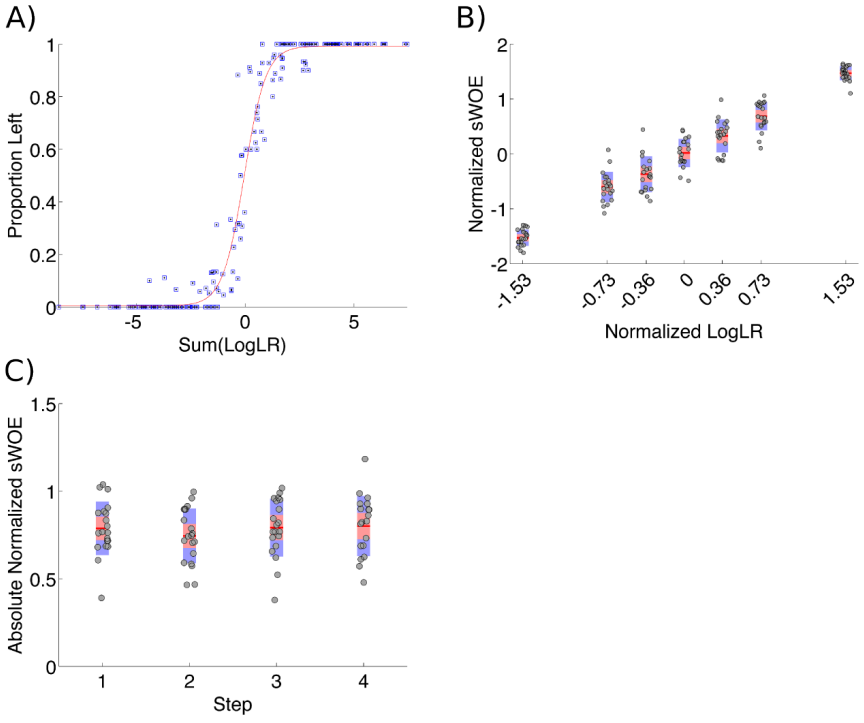

FIGURE 3.3. Behavioral performance. A) For each participant, the proportion of Category Left responses is plotted across 10 bins of cumulative normative LogLR (equation 8). B) Normalized subjective weights of evidence (sWOE) vs. the normalized normative logLR used to determine the probabilistic dispersement of feedback. X-axis: normalized logLR. Y-axis: normalized sWOE derived from each participant's pattern of behavioral performance. (Grey dots represent individual subject estimates, Dark red = mean, light red = 95% confidence interval, blue = 1 SD). C) Influence of each step on choice. Y-axis: absolute normalized mean sWOE for each step. These analyses provide evidence that subjects tended to integrate information across features according to a weighting scheme that closely approximated the normative characteristics of the probabilistic task environment.

3.5. Neuroimaging Results

3.5.1. Cumulative Normative LogLR vs. Cumulative Subjective LogLR.

To reduce the number of models investigated, the normative (which was used to defined the probabilistic disbursement of feedback) and subjective (which was derived from the logistic regression coefficients) CmLogLR models were compared using the RFX-BMS procedure. As described above, significant clusters were first identified by considering the positive and negative contrasts of the CmLogLR regressor over implicit baseline. The supra-threshold voxels from both maps were then combined into a single mask. Considering only voxels within this mask, the relative model evidence was then compared using the protected exceedance probability (PXP) maps derived from the RFX-BMS procedure. Although the classical analyses yielded similar maps for the normative and subjective models, the model comparison procedure provided stronger evidence for the subjective model in the majority of clusters (67%). Based on this weak evidence, the subjective CmLogLR was considered for all remaining neuroimaging analyses.

3.5.2. Classical Analyses.

3.5.2.1. *Mean Effect of Features.* A contrast of mean features > implicit baseline was used to investigate regions broadly involved in processing the stimulus features. As the resulting map yielded widespread activation, in order to provide meaningful interpretation, instead of correcting for multiple comparisons using the topological false discovery rate (Chumbley and Friston, 2009), the more conservative family-wise error rate was used. The bilateral intraparietal sulcus, bilateral insula, middle cingulate, superior and orbitofrontal gyri were involved in processing the stimulus features (Table A.2; Figure 3.4).

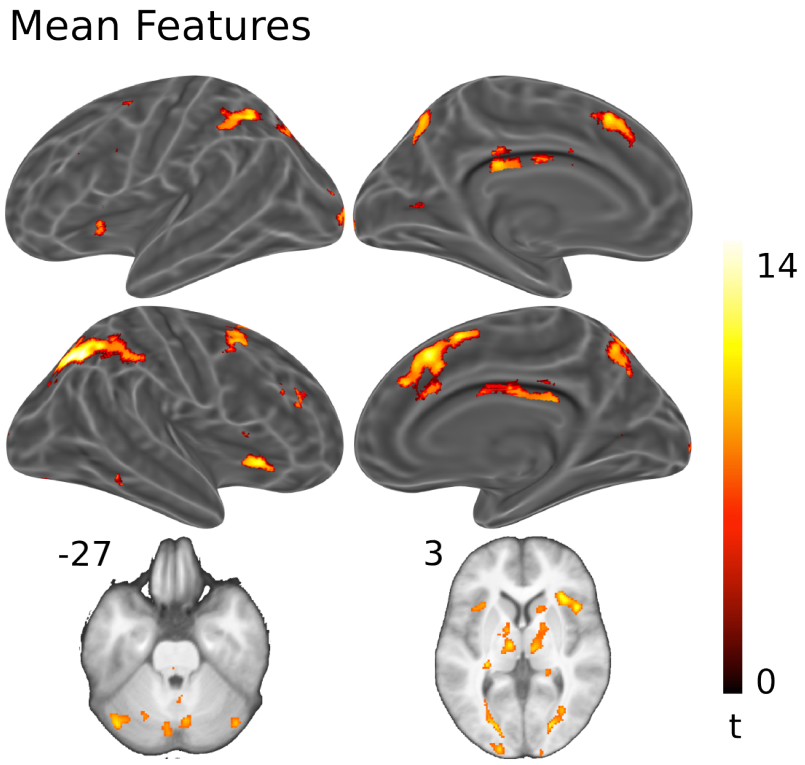

FIGURE 3.4. Mean Effect of Features. Regions associated with this contrast were generally involved in processing the stimulus features. (Contrast: Mean Features > implicit baseline).

3.5.2.2. *Task Negative Effects.* A contrast of implicit baseline > mean features was used to investigate regions negatively correlated with feature processing. This analysis revealed regions known to be involved in the default mode network, including inferior medial frontal cortex, posterior cingulate, as well as bilateral hippocampus (Table A.3); Figure 3.5).

3.5.2.3. *Categorical Evidence (CmLogLR) and its Interaction with Urgency.* Categorical evidence tracked the evidence for responses with the left vs. right hands (and was operationalized as the cumulative logLR for each response). Evidence for a left response was represented in a large contralateral region of the right pre- and post-central gyri (overlapping BA4 and BA6) extending into the right superior parietal lobe, and in the subcortical

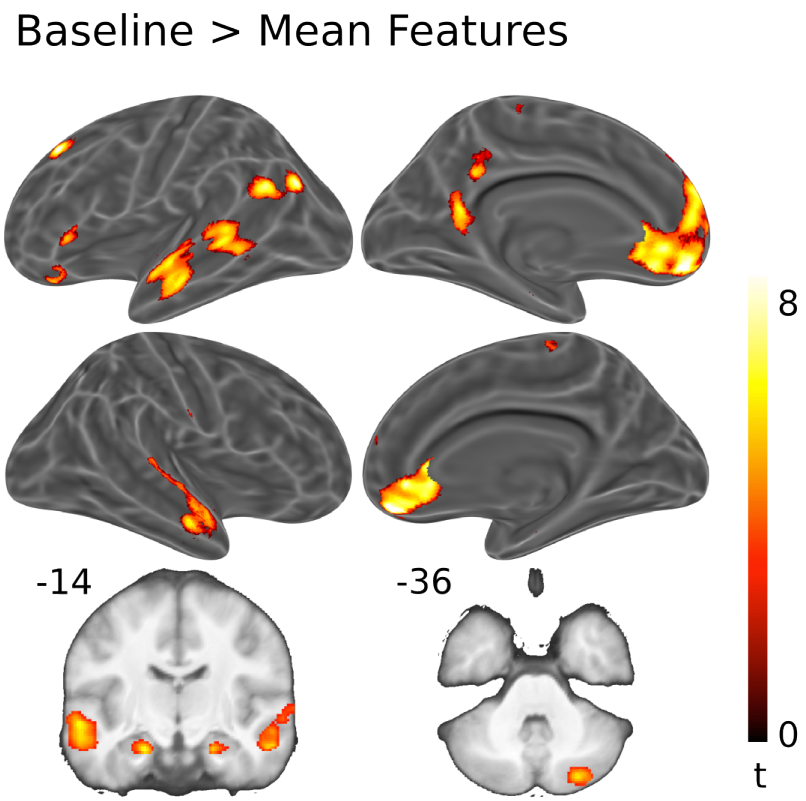

FIGURE 3.5. Task-Negative Effects. Regions associated with this contrast generally showed greater activity during baseline (contrast: implicit baseline > Mean Features).

activity of left cerebellar lobules IV and V, as well as in the right posterior putamen / globus pallidus and motor regions of the thalamus (Figure 3.6). Conversely, the regressors tracking evidence for a right response elicited activity in the left pre and post-central gyri. The model tracking the multiplicative interaction between categorical evidence and urgency yielded a highly similar statistical map (Table A.4; Figure 3.6.B).

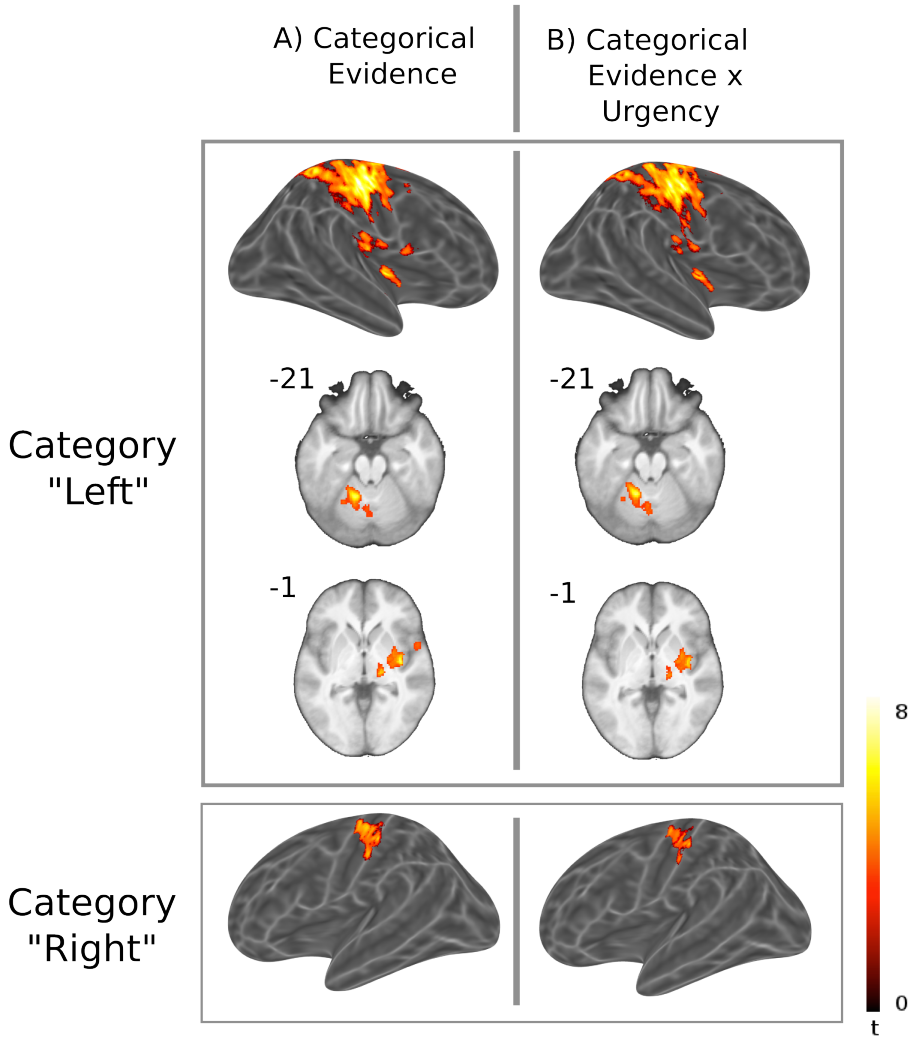

FIGURE 3.6. Cortical (bilateral pre- and post-central and primary motor regions), subcortical (right posterior putamen and posterior putamen) and cerebellar (Lobules V and VI) somatomotor regions tracked the evolution of categorical evidence. Left and Right indicate evidence towards the respective categories. A) Categorical evidence. B) The model tracking the (multiplicative) interaction between categorical evidence and urgency. The model tracking categorical evidence and the model tracking its interaction with urgency yielded similar statistical maps (contrast: PM > implicit baseline).

3.5.2.4. *Classical Results: Confidence and Uncertainty (Unsigned CmLogLR) and its Interaction with Urgency.* As evidence for each of the two categories were perfectly anticorrelated, confidence was modeled as unsigned CmLogLR. Whereas categorical evidence represents evidence for specific categorical responses (i.e., the probability of positive feedback given a specific action), confidence tracks the probability of positive feedback irrespective of particular categories. Therefore, whereas categorical evidence can be used to select the best possible choice, confidence allows the decision maker to determine whether it might be advantageous to opt-out of certain trials (e.g., if the costs of an incorrect response outweigh the potential benefits of guessing), and to learn efficiently from feedback (the unsigned CmLogLR can also be interpreted as a reward prediction signal).

Confidence (Table A.5; Figure 3.7.A) was represented in the activity of the bilateral SMA and striatum (bilateral putamen and right caudate body), bilateral temporoparietal junction, bilateral insula, bilateral middle frontal regions, right middle orbital gyrus, anterior and middle cingulate, bilateral cuneus and Crus I and II regions of the left cerebellum. The (multiplicative) interaction between confidence and urgency was associated with a highly similar statistical map, but was not associated with regions of the right medial and lateral frontal cortex. This finding implies that the slope of activity associated with confidence did not vary according to urgency, however it does not preclude the possibility that activity in this region was additionally sensitive to an additive urgency signal (i.e., that activity in this region tended to increase with urgency). Activity in bilateral occipital cortex, fusiform and superior parietal lobe, as well as the left precentral motor cortex and right superior frontal lobe covaried with decisional uncertainty (Table A.6; Figure 3.8).

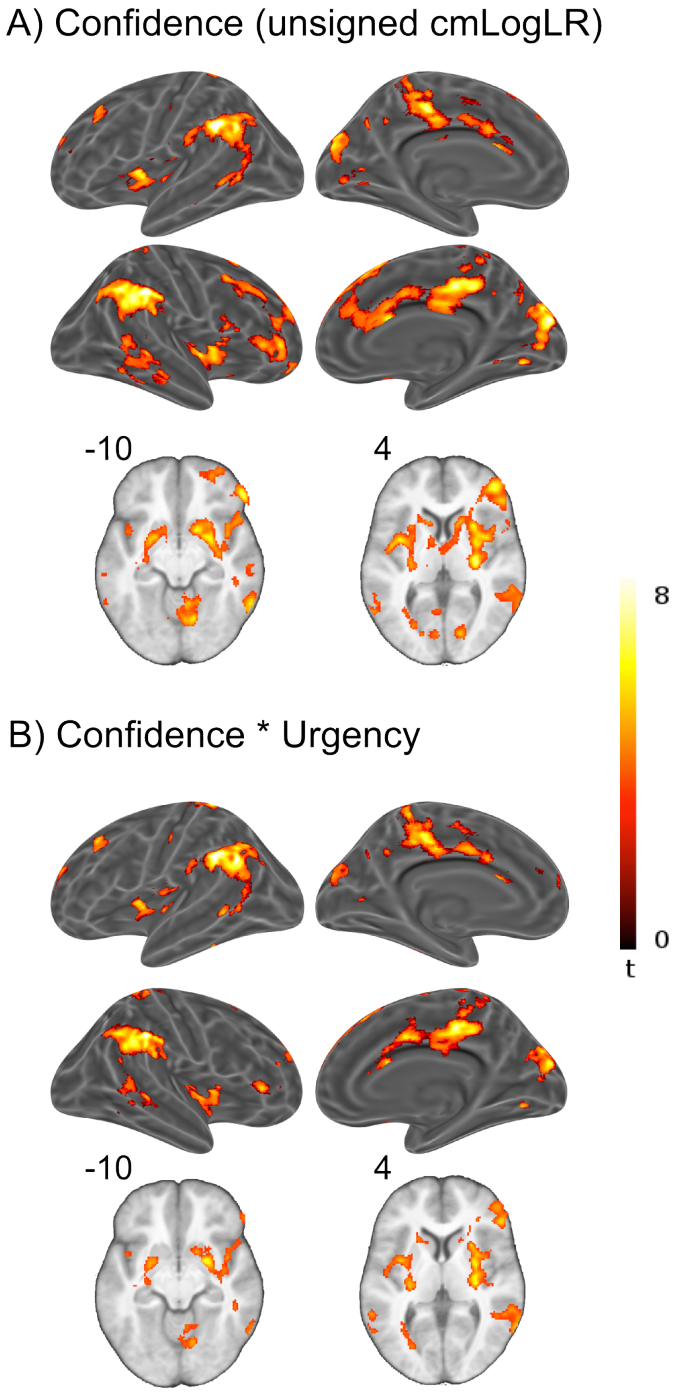

FIGURE 3.7. Classical Results: Confidence. A) Confidence (unsigned CmLogLR). B) The multiplicative interaction between confidence and urgency.

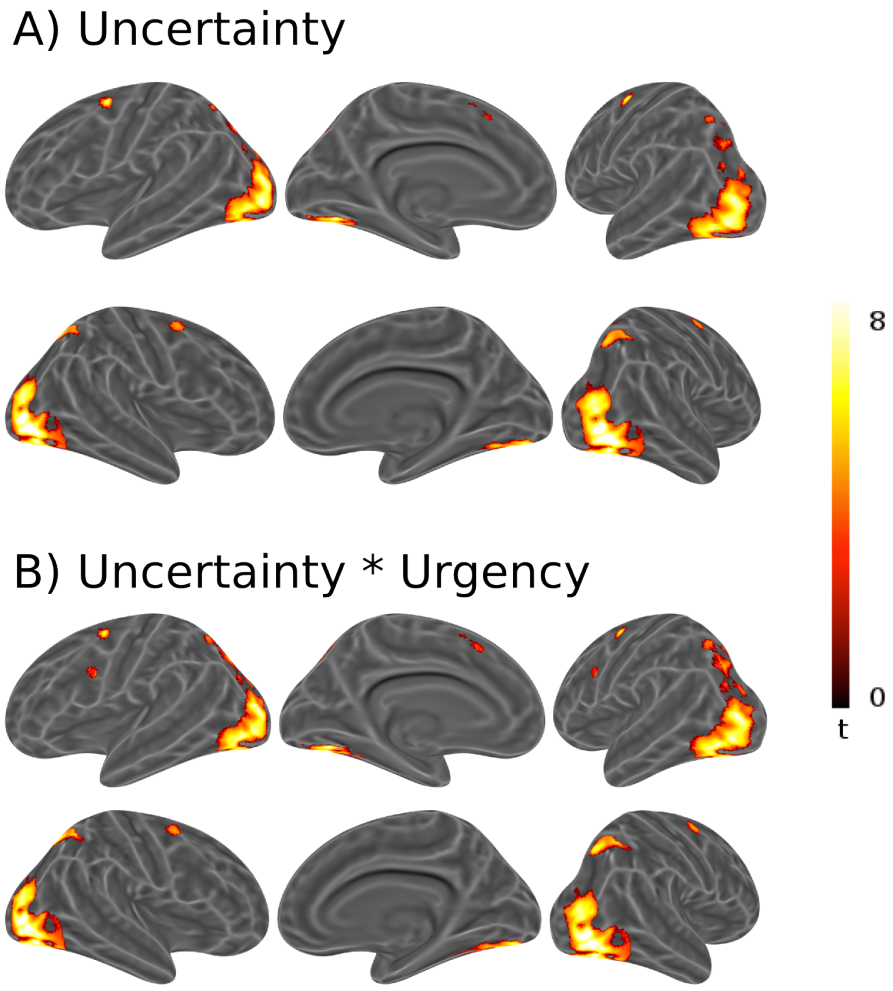

FIGURE 3.8. Classical Results: Uncertainty. Activity in bilateral occipital lobe, fusiform, and superior parietal lobe was negatively correlated with confidence. Highly similar statistical maps were obtained for the model tracking the interaction with urgency.

3.5.2.5. *Urgency.* As described above, for each of the four classical models, 2 PMs were included per mean regressor. The first PM tracked an evidence variable (categorical evidence or confidence), and the second PM tracked urgency. Whereas the first PM allowed the investigation of multiplicative interactions between the evidence variable and urgency (i.e., whether urgency modulates the slope of evidence variable), the second PM allowed the investigation of additive effects. As this regressor was orthogonalized with respect to the first PM, overlap between the statistical maps associated with the first and second PMs provides evidence that activity within these regions tracked the first PM and also tended to increase with urgency. To identify regions associated with this orthogonalized representation, a conjunction analyses was conducted for the contrast of urgency greater than implicit baseline across each of the four classical models. The average t-values across these maps are reported in Table A.7 and Figure 3.9.

Urgency was notably associated with the pre-SMA, bilateral cerebellum, right precuneus, anterior and middle cingulate, right middle frontal gyrus, right inferior frontal gyrus, and a large region of the right inferior parietal lobe and supramarginal gyrus. Urgency was also represented in widespread striatal regions, including the bilateral body of the caudate and the ventral striatum. Right inferior parietal and medial frontal regions also notably tracked decisional confidence (Table A.5; Figure 3.7), providing evidence that activity in these regions tracked not only confidence, but also tended to increase with urgency.

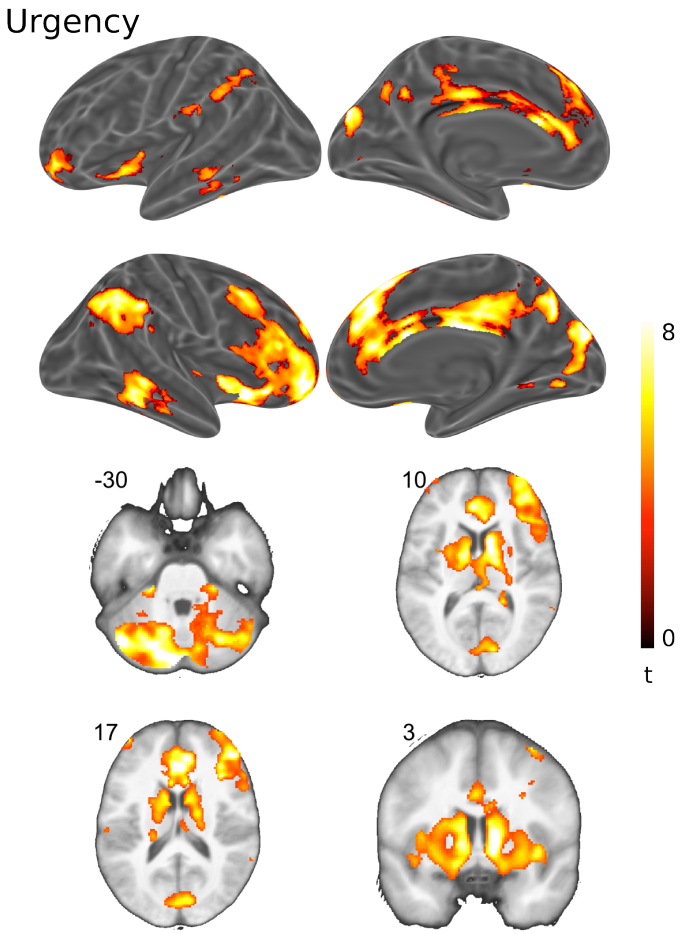

FIGURE 3.9. Urgency was modeled as a linearly-increasing signal that peaked on the last step. It was included as the second PM in each model, and was orthogonalized with respect to the first PM. This statistical map is the result of a conjunction analysis across each of the four models (Table 3.1), and the values reflect averaged t-values across these maps. Many of these regions overlapped with those associated with the first PM of each model, providing evidence of additive gain modulation of their representations.

3.5.2.6. *Feedback.* During feedback several regions were sensitive to prediction errors related to the categorical evidence and confidence functions. Given a limited number of incorrect trials, these analyses were limited to trials where participants made the normatively correct choice. Given a limited number of trials, we constrained these analyses to trials where participants made the normatively correct choice. Reward prediction errors was calculated as:

$$PredictionError = R - P(correct) \tag{10}$$

where $R = 1$ for positive feedback and $R = 0$ for negative feedback and P indicates the probability of a correct choice. To investigate effector-specific prediction errors, P represented the CmLogLR on the response step. To investigate effector-independent prediction errors, P represented the confidence (unsigned CmLogLR) on the last step.

It was predicted that regions involved in learning the contingencies between features and specific motor responses would show positive prediction errors associated with categorical evidence (Gershman et al., 2009; Samejima et al., 2005; Stalnaker et al., 2012), while regions involved in learning the contingencies between features and rewarded outcomes, without regard to specific effectors, would show prediction errors related to decisional confidence (Fiorillo et al., 2003).

Regions neighboring the central sulcus of right hemisphere showed effector-specific prediction errors (Table A.8; Figure 3.10.B). Mirroring the generally weaker results for the Right category, no regions were found to track positive prediction error for the Left category. Activity within regions in the right putamen (overlapping regions associated with confidence

and its interaction with urgency), left caudate body, left insula, right frontal mid-orbital gyrus, and the left crus I and bilateral crus II regions of the cerebellum, was consistent with an effector-independent reward prediction signal .

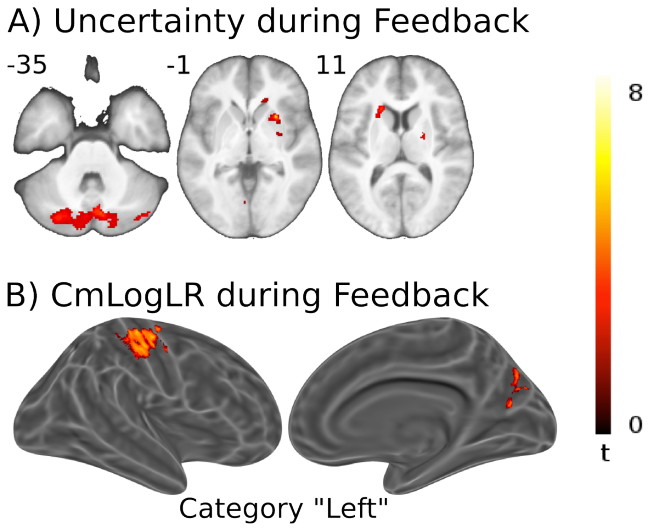

FIGURE 3.10. A) Effector-independent reward prediction error in the bilateral head of the caudate, right putamen and cerebellum. B) Effector-specific prediction error corresponding to the left response in left-cortical motor regions and right precuneus. No regions sensitive to prediction errors corresponding to the right response were identified.

3.5.3. RFX-BMS. While all voxels included in the RFX-BMS analyses surpassed classical significance thresholds, only moderate and strong effects (PXP >= 0.8) will be discussed. As shown in Figure 3.11(and in Tables A.9, Table A.10, and Table A.11), each individual PM (categorical evidence, confidence, and uncertainty) was investigated individually with the individual PM multiplied by urgency. For categorical evidence, there was strong evidence that the slope of the representation of categorical evidence within bilateral cortical somatomotor regions (peak PXP > 0.9) is multiplicatively modulated by urgency. In cerebellar lobule VI, however, this representation is more likely to be unmodulated by urgency

(peak PXP = 0.8). In addition, there were trends (PXP > 0.8) for the putamen and premotor and superior parietal regions to be better fit by the function that was not modulated by urgency. For the confidence regressor, the right medial superior frontal gyrus was likely unmodulated by urgency (peak PXP = .08), while activity in the right superior parietal lobe (PXP > 0.8), right cuneus (peak PXP > 0.8) and right superior temporal pole (peak PXP > 0.8) were likely modulated urgency. For uncertainty, no regions exceeded the criterion, indicating no strong evidence about whether the slopes of the representations of uncertainty were modulated by urgency or not.

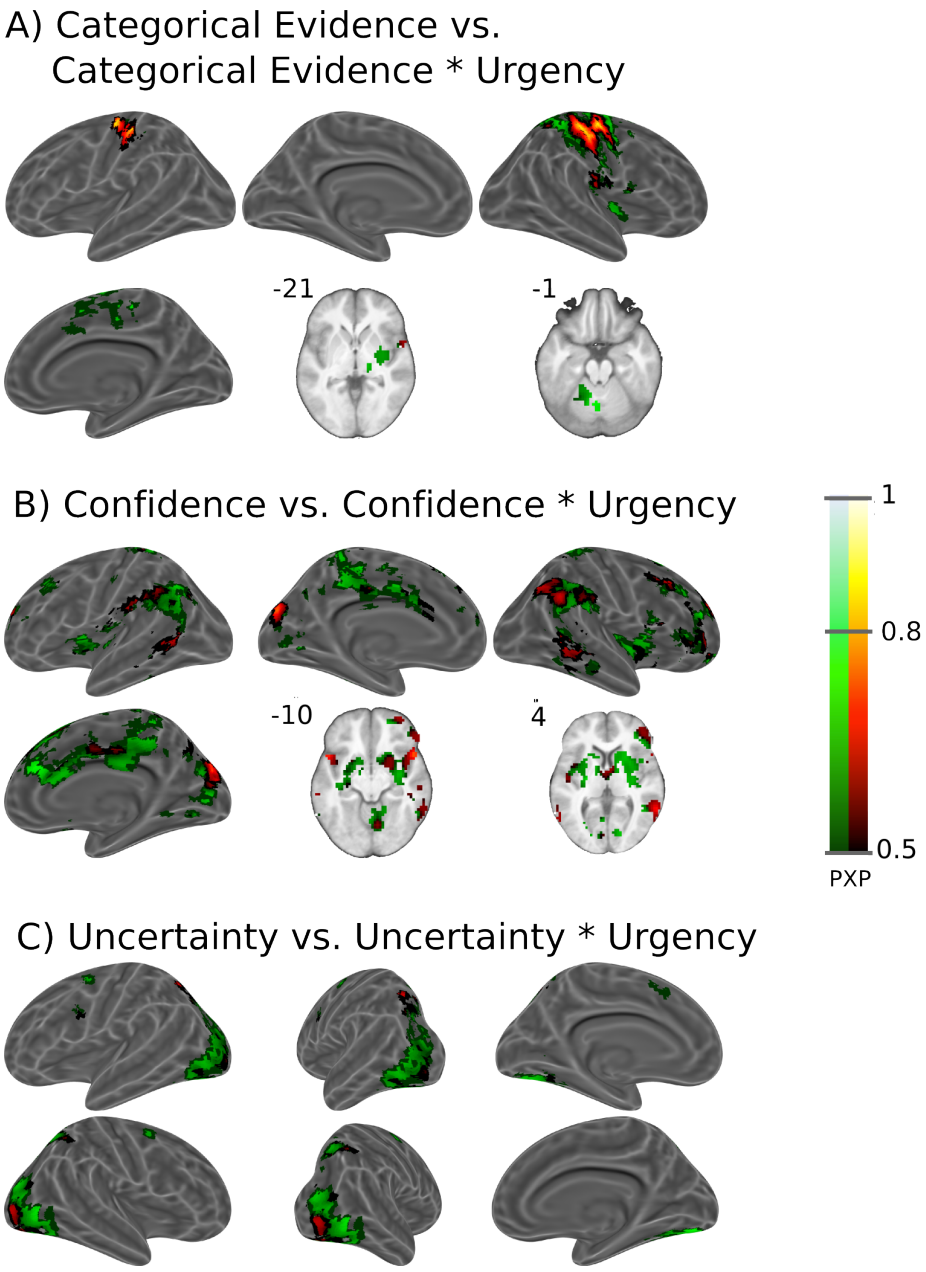

FIGURE 3.11. Model Comparison Results. Protected Exceedance Probability (PXP) maps for the three pairs of models tested. Each map was thresholded to include only voxels where the classical models provided the best fit. The PXP indicates the probability that the model provided the best fit, relative to the other models tested, above and beyond chance (Rigoux et al., 2014). A) Green: Categorical Evidence; Warm colors: Categorical Evidence weighted by urgency. Cortical somatomotor regions showed strong preference (PXP $>$ 0.9) for the model tracking the multiplicative interaction between categorical evidence and urgency. B) Green: Confidence; Warm colors: Confidence weighted by urgency. C) Uncertainty was perfectly anticorrelated with Confidence. Green: Uncertainty; Warm colors: Uncertainty modulated by urgency.

CHAPTER 4

Discussion

A temporally extended design was used to investigate how participants integrated probabilistic featural information across time in order to make categorization decisions. Importantly, although most of the features were informative, they were not perfectly predictive. This kind of probabilistic decision-making task is common in everyday life; environmental cues often provide only incomplete information about the course of action that will lead to a desired state. In order to interact advantageously with such environments, decision-makers must assign weights to information sources according to a weighting scheme that approximates the characteristics of the external environment (Craik, 1967). In the present experiment, our goal was to investigate the neural mechanisms underlying our capacity to integrate information across stimulus features of varying reliability, and to investigate how we modulate this process based upon the time available to respond. Through use of a model-based approach, the temporal evolution of three neural functions were differentiated: categorical evidence (the evidence for specific responses associated with the category), decisional confidence (an abstract representation proportional to the probability of reaching a desirable state), decisional uncertainty (the inverse of confidence), and urgency (a monotonically-increasing signal representing an increasing pressure to respond within each trial). The design also allowed investigation into how urgency might modulate the other functions.

4.1. Multi-Cue Integration and Categorization

A primary goal of our study was to examine how fundamental decision making mechanisms were adapted for categorical decision making. Through use of a probabilistic multiple cue categorization task, participants were encouraged to integrate information across multiple stimulus dimensions. Previous research has indicated that in non-temporally extended versions of the task, participants frequently use simpler heuristic strategies, where they consider only a subset probabilistic cues (e.g., considering only the cues that are most informative; Gluck et al., 2002; Meeter et al., 2006, 2008). A logistic regression analysis (Figures 3.3.B) provided evidence that participants in our tasks tended to use weighting schemes that closely matched the normative logLR's which determined the probabilistic disbursement of feedback (eq. 7). The temporally-extended nature of our paradigm, wherein participants were exposed not only to the uncertainty associated with the informational content of the features, but also to the uncertainty about what features would comprise the final stimulus, may have encouraged participants to consider each of the individual features during deliberation.

The goal of the current study was to investigate the neural bases underlying the integration of probabilistic information. Accordingly, our design was intended to encourage participants to integrate information across multiple stimulus dimensions. An important domain for future research will to investigate how different environmental influences may influence the use of different decision strategies. For instance, by providing incentive to respond quickly (e.g., a points system as use in the pilot study), it would be possible to encourage participants to consider only a subset of the stimulus features.

4.2. Categorical Evidence

Somatomotor regions (primary motor cortex, pre- and post-central motor cortex, posterior putamen, and cerebellum) tracked evidence (the cumulative logLR) towards specific categories and their associated responses. Activity appeared to be effector specific, with primary activity contralateral to the hand used to make the categorical response. Activity within cortical somatomotor regions, for instance, was positively correlated with evidence towards the contralateral response, while activity within somatomotor regions of the cerebellum was positively correlated with evidence for the ipsilateral response, consistent with the crossed nature of cortical-cerebellar projections. Previous studies have also found motor effector specific activity in decision making tasks (Gluth et al., 2012; Selen et al., 2012; Thura and Cisek, 2014; Tosoni et al., 2008; Wheeler et al., 2014). Of particular note are two studies in humans using relatively high level decision making tasks. Gluth et al. (2012) found that similar motor regions tracked decision evidence in a temporally-extended value-based decision task, and Wheeler et al. 2014 found that motor regions were sensitive to the rate at which information accumulated in a probabilistic feature task similar to ours. Previous groups (e.g., Cisek, 2007; Thura and Cisek, 2014) have argued that such activity is unlikely to be an epiphenomenal result of decision processes occurring upstream. Our results strengthen this argument by providing evidence that activity in the right central sulcus (a region strongly tracking evidence for a left response) was also re-instantiated as a positive, effector-specific prediction error during correct feedback (Table A.8; Figure 3.10); a signal which likely served to increase the probability of making the same response in the future, given similar stimuli.

4.3. Confidence and Uncertainty

Signals related to decisional confidence have been previously identified in the striatum (Ding and Gold, 2010), dopaminergic midbrain (Schwartenbeck et al., 2014), and in dorsolateral, rostrolateral, ventromedial and orbitofrontal cortical regions (Bowman et al., 2012; De Martino et al., 2013; Heekeren et al., 2004, 2006; Philiastides et al., 2011; Rolls et al., 2010a,b; Tobler et al., 2007). Our results mirrored these results and indicated that the DLPFC, RLPFC, and striatum tracked confidence. Additionally, regions of the ventral parietal lobe, which have been previously associated with *bottom-up* attentional control (Corbetta and Shulman, 2002), tracked the temporal evolution of confidence (Figures 3.7 and 3.11.B). Conversely, mirroring previous findings (White et al., 2012), the superior parietal lobe, which has been previously implicated in *top-down* attentional control (Cabeza et al., 2008; Corbetta and Shulman, 2002), tracked decisional uncertainty (Figures 3.8 and 3.11.C). Thus, our results provide evidence that dorsal and ventral parietal regions covary with decisional confidence, but may play opposing roles during deliberation.

Decisional confidence represents the choice-independent probability that a decision will be correct, and was calculated as the unsigned CmLogLR. Our confidence regressor, therefore, tracked the total amount of information accumulated from the amoeba stimulus, without regards to specific motor responses. As analyses were limited to correct trials only, it provides a normative estimate of the probability of transitioning to a desirable state (a state associated with positive feedback), without regard to the specific action needed to obtain it. Such predictions of reward provide the basis for calculation of reward prediction error (eq. 4), and so may play an important role in reward-mediated learning. As participants were free to respond at any time during each trial, such a function likely also provides an important

signal indicating the value of accumulating additional evidence; when confidence is high, the value of accumulating additional information is low, and vice versa. However, in the present experiment, no incentives were provided for participants to respond early, and most participants were fairly conservative in regards to their response timing (they showed a strong tendency to respond on the last trial).

Decisional confidence may also modulate the use of different decisional strategies. For instance, on difficult trials, confidence may also allow decision-makers to opt-out of difficult choices instead of gambling on uncertain information (Gherman and Philiastides, 2014; Kepecs et al., 2008). As only one information source (the amoeba features) provided a reliable source of information in our study, observed activity may have reflected the accumulation of information for this exogenous information channel only. If multiple sources of information were relevant, it is possible that other regions might have accumulated this information. Signals related to confidence in particular sources may allow decision-makers to shift between information channels based on their estimated reliability. For instance, when confidence in a mnemonic information source is low, decision makers may tend to place greater weight on perceptual channels, and vice versa (Hanks et al., 2011; Hutchinson et al., 2014). This kind of precision-weighted competition may also underly switches between different neural systems. In the COmpetition between Verbal and Implicit Systems model (COVIS; Ashby et al., 2011), the competition between the explicit and procedural systems is mediated by estimates of the reliability of each. An important goal for future research will be to understand how we flexibly recruit different neural systems based on their estimated reliability.

4.4. Urgency

There are a number of ways that flexible adjustments to the SAT may be implemented in the brain (for review, see Bogacz et al., 2010). One mechanism is analogous to manipulation of the decision boundary itself; through modulation of the cortico-basal ganglia circuit. For instance, excitation of the striatum may serve to lower thresholds by decreasing the tonic inhibitory influences of the basal ganglia output nuclei (Chevalier and Deniau, 1990; Simen et al., 2006). In support of this hypothesis, several previous studies have found that activity in the pre-SMA and striatum covaries with emphasis on speed vs. accuracy (Forstmann et al., 2008, 2010; Green et al., 2012; Ivanoff et al., 2008; van Maanen et al., 2011; van Veen et al., 2008). Additionally, individual differences in connectivity between these regions predict individual differences in the capacity to flexibly modulate the SAT Forstmann et al. (2010). Another hypothetical mechanism underlying flexible modulation of the SAT is analogous to modulating the decision variable through gain modulation of neural activity associated with the integration of categorical evidence (Furman and Wang, 2008; Salinas and Abbott, 1996; Standage et al., 2011, 2013; Thura and Cisek, 2014). Previous studies have provided evidence that activity within regions associated with the integration of decisional evidence shows neural gain well-accounted by this signal (e.g., Churchland et al., 2008; Thura and Cisek, 2014), and several successful biologically-plausible decision-making models (e.g., Cisek et al., 2009; Ditterich, 2006; Niyogi and Wong-Lin, 2013; Standage et al., 2011, 2013), have demonstrated that such a signal represents an effective way to modulate the SAT; one that closely matches observed patterns of human behavior.

In the present experiment, the SAT was not manipulated. Instead, its temporal evolution was tracked the evolution of this monotonically-increasing urgency signal. This signal

was strongly represented in the activity of the pre-SMA and striatum, regions previously implicated in flexible modulation of the SAT. In addition to the additive urgency signal, which increased the amplitude of the BOLD response across neuronal populations, urgency modulated the slope of the neural representations of categorical evidence and confidence. Therefore, urgency modulated two hypothetical mechanisms underlying flexible modulation of the SAT: gain modulation of the striatal thresholding circuitry, and gain modulation of the decisional evidence (see Bogacz et al., 2010, for review). Our results highlight similarities between contrastive modeling approaches based on manipulations of contextual demands for speed vs. accuracy (e.g., Forstmann et al., 2008), and the model-based approach applied here.

Specifically, our findings are in accordance with extant evidence that urgency can modulate the temporal dynamics of deliberation by moving the system towards commitment (e.g., Ditterich, 2006; Standage et al., 2011). It is important to note, however, that this may not hold for all tasks. Hawkins and colleagues (2015), for instance, recently provided evidence that across a range of perceptual decision-making tasks, behavioral performance tends to be best accounted for by models that include stationary, rather than collapsing decision boundaries (i.e., where the distance between the decision variable and the boundary is not modulated by time). However, it is unclear whether the insights gained from the drift diffusion model (which was originally intended for decisions with short deliberative periods; Ratcliff and Rouder, 1998) should be applied to paradigms where the information is presented in discrete epochs, where evidence for and against particular categories is modulated within single trials, and where the final commitment to a behavioral response might be based on several individual decisions.

4.5. Conclusions

Real world environments often require decisionmakers to integrate information across time and across multiple information sources. To do so effectively, they must weight different information channels according to their estimated reliability, and they must temper the timing of their behavioral responses to maximize transient opportunities to reach desirable states. In the present experiment, urgency modulated activity in regions known to be associated with flexible modulation of the speed-accuracy trade-off, and also representations of categorical evidence and confidence to match these demands on response timing. Emerging evidence suggests that somatomotor recruitment during deliberation depends on knowledge of the mapping between choice and response and that the specific regions recruited are sensitive to manipulations of effector modality (Ho et al., 2009; Tosoni et al., 2008). An important goal for future research, therefore, will be to understand how we recruit different neural populations during deliberation to support flexible cognitive control across variable task environments.

4.5.1. Limitations.

4.5.1.1. *Model Density.* The present experiment employed a Bayesian logistic regression analysis (Gelman et al., 2008) to characterize the strategies used by individual participants. This analysis employes an uninformative Cauchy prior (mean = 0 for all coefficients) which has the effect of shrinking the majority of regression coefficients towards zero while occasionally allowing large values. This general approach is applicable to a wide variety of regression problems and has the advantage of avoiding overfitting while also being robust to separation. However, particularly in high-dimensional real-world environments, the decision-making behavior of cognitively-limited organisms might be better accounted for by models

incorporating stronger priors for model sparsity. Such models would have the advantage of positing reduced mnemonic demands (i.e., requiring that only information from relevant dimensions would be considered) and so would more closely model the use of low-dimensional rule-based categorization strategies (e.g., Figure 1.1.B). A popular technique for encouraging model sparsity involves regularization via a L_1 norm. In a Bayesian framework, an L_1 penalty can be included in the model by using a Laplacian prior. In a frequentist framework, it can be included through use of the Least Absolute Shrinkage and Selection Operator (LASSO) penalty (Tibshirani, 1996). To avoid arbitrary decisions about the appropriate regularization strategy, future researchers might choose to use the elastic net (Zou and Hastie, 2005), a popular data-driven method to estimate the appopriated balance between L_1 and L_2 (e.g., ridge regression) norms. Among its attractive features, this approach is known to encourage the formation of sparse models (like the L_1 penalty), but also tends to yield better generalization performance than models including the L_1 penalty.

An alternative approach to characterize idiosyncratic strategies is to employ "strategy" analysis, which involves making strong assumptions about specific strategies that participants might use during the task. Previous researchers (e.g., Gluck et al., 2002; Meeter et al., 2006), for instance, have used this approach to compare evidence for singleton strategies (which involve consideration of only the most informative features), single-feature strategies (which involve consideration of only single features), "intermediate" integrative strategies (which involve integration of *some* information across features), and optimal strategies (which involve integrating information across all available features). A drawback of this approach (relative to the regression analyses described in the previous paragraph) is that it is

unable to identify strategies that were not predicted by the experimenter. This approach is also unable to capture subtle differences in the weights applied to different stimulus features.

4.5.1.2. *Strategy Changepoints.* When fit across entire experimental blocks, both regression and strategy analyses assume that participants have not switched between different strategies during the experiment. To identify possible strategy changepoints (points at which participants switch from one strategy to another), previous researchers have fit regression and strategy analyses to data within rolling windows of arbitrary size (typically 20-40 trials; Meeter et al., 2006, 2008). As this approach blurs information across trials within each window, a drawback of this approach is that it tends to provide imprecise estimates of strategy changepoints. Thus it also leads to imprecise characterization of the strategies employed on either side of the changepoints. An alternative approach, which has been used by Speekenbrink and colleagues (Speekenbrink et al., 2010), involves modeling strategy switches as hidden Markov models. This approach allows probabilistic inference about the precise locations of strategy changepoints, and thus allows estimation of the most likely patterns of subjective strategy use across the entire experiment.

4.5.1.3. *Manipulation of the SAT.* Urgency is thought to play the important functional role of modulating the temporal dynamics of the decision-making process to match contextual demands for speed vs. accuracy. A simple prediction based on this framework would be that the slope of the urgency signal should covary with demands for speed (i.e., a greater slope would be predicted when speed is emphasized over accuracy). In the present experiment, SAT demands were not manipulated, thus, to better establish the link between urgency and modulation of the dynamics of decision making, future researchers may wish to test this prediction. Additionally, as we found that urgency modulates representations of categorical

evidence and confidence through multiplicative and additive functions, it seems likely that urgency may serve to modulate other functions that influence the decision-making process. Urgency may, for instance, interact with prior knowledge to influence choice outcome (Hanks et al., 2011).

REFERENCES

Albin, R. L., Young, A. B., & Penney, J. B. (1989). The functional anatomy of basal ganglia disorders. *Trends in Neurosciences*, *12*(10), 366–375.

Alexander, G., DeLong, M., & Strick, P. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Reviews of Neuroscience*, *9*(1), 357:81.

Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: Neural substrates of parallel processing. *Trends in Neurosciences*, *13*(7), 266–71.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to Stop signal response inhibition: Role of the subthalamic nucleus. *The Journal of Neuroscience*, *26*(9), 2424–33.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*(3), 442–81.

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–78.

Ashby, F. G., & Maddox, W. T. (2011). Human category learning 2.0. *Annals of the New York Academy of Sciences*, *1224*(1), 147–61.

Ashby, F. G., Paul, E. J., & Maddox, W. T. (2011). COVIS. In E. Pothos & W. AJ (Eds.), *Formal approaches in categorization* (pp. 65–87). New York, New York, USA: Cambridge University Press.

Ashby, F., Turner, B., & Horvitz, J. (2010). Cortical and basal ganglia contributions to habit learning and automaticity. *Trends in Cognitive Sciences*, *14*(5), 208–215.

Bernier, P.-M., Cieslak, M., & Grafton, S. T. (2012). Effector selection precedes reach planning in the dorsal parietofrontal cortex. *Journal of Neurophysiology*, *108*(1), 57–68.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–65. doi:10.1037/0033-295X.113.4.700

Bogacz, R., Wagenmakers, E. J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends in Neurosciences*, *33*(1), 10–6.

Bowman, N. E., Kording, K. P., & Gottfried, J. A. (2012). Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*, *75*(5), 916–927. doi:10.1016/j.neuron.2012.06.035

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433–6.

Britten, K., Shadlen, M., Newsome, W., & Movshan, J. (1992). The analysis of visual Motion: A comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, *12*(12), 4745–4765.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–78.

Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. Oxford, England: John Wiley and Sons.

Buckner, R. L. (1998). Event-related fMRI and the hemodynamic response. *Human Brain Mapping*, *6*(5-6), 373–377.

Cabeza, R., Ciaramelli, E., Olson, I. R., & Moscovitch, M. (2008). The parietal cortex and episodic memory: An attentional account. *Nature Reviews Neuroscience*, *9*(8), 613–625.

Cantwell, G., Crossley, M. J., & Ashby, F. G. (2015). Multiple stages of learning in perceptual categorization: Evidence and neurocomputational theory. *Psychonomic Bulletin & Review*, 1–16.

Chevalier, G., & Deniau, J. M. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends in Neurosciences*, *13*(7), 277–80.

Chumbley, J. R., & Friston, K. J. (2009). False discovery rate revisited: FDR and topological inference using Gaussian random fields. *NeuroImage*, *44*(1), 62–70.

Churchland, A., Kiani, R., & Shadlen, M. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, *11*(6), 693–702.

Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *362*(1485), 1585–99.

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *The Journal of Neuroscience*, *29*(37), 11560–71.

Clarke, F. R., Birdsall, T. G., & Tanner, W. . (1959). Two types of ROC curves and definitions of parameters. *The Journal of the Acoustical Society of America*.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *362*(1481), 933–42.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–15.

Craik, K. (1967). *The nature of explanation*. (1st ed.). Chicago: Cambridge University Press.

Dale, A. M., & Buckner, R. L. (1997). Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping*, *5*(5), 329–40.

Daw, N., ODoherty, J., Dayan, P., & Seymour, B. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876–879. doi:10.1038/nature04766.Cortical

De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based choice. *Nature Neuroscience*, *16*(1), 105–10.

Deco, G., & Rolls, E. T. (2006). Decision-making and Weber's law: A neurophysiological model. *The European Journal of Neuroscience*, *24*(3), 901–16.

Deco, G., Rolls, E. T., Albantakis, L., & Romo, R. (2012). Brain mechanisms for perceptual and reward-related decision-making. *Progress in Neurobiology*, *103*, 194–213.

Dickinson, A. (1985). Actions and habits: The development of behavioral autonomy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *308*(1135), 67–78.

Ding, L., & Gold, J. I. (2010). Caudate encodes multiple computations for perceptual decisions. *The Journal of Neuroscience*, *30*(47), 15747–59.

Ditterich, J. (2006). Evidence for time-variant decision making. *The European Journal of Neuroscience*, *24*(12), 3628–41.

Downing, C. J., & Movshon, J. A. (1989). Spatial and temporal summation in the detection of motion in stochastic random dot displays. *Investigative Ophthalmology and Visual Science*, *30*.

Doya, K., Ishii, S., Pouget, A., & Rajesh P. N. Rao (Eds.). (2007). *Bayesian brain: Probabilistic approaches to neural coding*. MIT Press.

Fetsch, C. R., Kiani, R., Newsome, W. T., & Shadlen, M. N. (2014). Effects of cortical microstimulation on confidence in a perceptual decision. *Neuron*, *83*(4), 797–804.

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science (New York, N.Y.)*, *299*(5614), 1898–902.

Fitzgerald, J. K., Freedman, D. J., & Assad, J. A. (2011). Generalized associative representations in parietal cortex. *Nature Neuroscience*, *14*(8), 1075–9.

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*(July), 1–9.

Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., … Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(36), 15916–20.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(45), 17538–42.

Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, *19*(8), 1120–36.

Freedman, D. J., & Assad, J. A. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature*, *443*(7107), 85–8.

Friston, K. (2012). Ten ironic rules for non-statistical reviewers. *NeuroImage*, *61*(4), 1300–1310. doi:10.1016/j.neuroimage.2012.04.018

Friston, K. J., Mechelli, A., Turner, R., & Price, C. J. (2000). Nonlinear responses in fMRI: The Balloon model, Volterra kernels, and other hemodynamics. *NeuroImage*, *12*(4), 466–77.

Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., … Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, *8*(1), e1002327.

Friston, K. J., Zarahn, E., Josephs, O., Henson, R. N., & Dale, a M. (1999). Stochastic designs in event-related fMRI. *NeuroImage*, *10*(5), 607–19.

Furman, M., & Wang, X.-J. (2008). Similarity effect and optimal control of multiple-choice decision making. *Neuron*, *60*(6), 1153–68. doi:10.1016/j.neuron.2008.12.003

Gelman, A., Jakulin, A., Pittau, M. G., & Su, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, *2*(4), 1360–1383.

Gershman, S., Pesaran, B., & Daw, N. (2009). Human reinforcement learning subdivides structured action spaces by learning effector-specific values. *The Journal of Neuroscience*, *29*(43), 13524–13531.

Gherman, S., & Philiastides, M. G. (2014). Neural representations of confidence emerge from the process of decision formation during perceptual choices. *NeuroImage*, *106C*, 134–143.

Gigerenzer, G., & Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science*, *1*(1), 107–143.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–82.

Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, *66*(4), 585–595.

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the "weather prediction" task?: Individual variability in strategies for probabilistic category learning. *Learning & Memory (Cold Spring Harbor, N.Y.)*, *9*(6), 408–18.

Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding when to decide: Time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *The Journal of Neuroscience*, *32*(31), 10686–98.

Gluth, S., Rieskamp, J., & Büchel, C. (2014). Neural evidence for adaptive strategy selection in value-based decision-making. *Cerebral Cortex*, *24*, 2009–2021.

Gold, J. I., & Shadlen, M. N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, *404*(6776), 390–4.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Reviews of Neuroscience*, *30*, 535–74.

Goodman, S. N., & Berline, J. A. (1994). The use of predicted confidence intervals when planning experiments and the misuse of power when interpreting results. *Annals of Internal Medicine*, *121*, 200–206.

Green, N., Biele, G. P., & Heekeren, H. R. (2012). Changes in neural connectivity underlie decision threshold modulation for reward maximization. *Journal of Neuroscience*, *32*(43), 14942–14950.

Grinband, J., Wager, T. D., Lindquist, M., Ferrera, V. P., & Hirsch, J. (2008). Detection of time-varying signals in event-related fMRI designs. *NeuroImage*, *43*(3), 509–20.

Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, *3*, 1157–1182.

Haber, S. N., & Calzavara, R. (2009). The cortico-basal ganglia integrative network: The role of the thalamus. *Brain Besearch Bulletin*, *78*(2-3), 69–74.

Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(9), 5359–5362.

Hanks, T. D., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *eLife*, *2014*, 1–17.

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *The Journal of Neuroscience*, *31*(17), 6339–52.

Hassabis, D., & Maguire, E. A. (2007). Deconstructing episodic memory with construction. *Trends in Cognitive Sciences*, *11*(7), 299–306.

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *The Journal of Neuroscience*, *35*(6), 2476–2484.

Hebart, M. N., Schriever, Y., Donner, T. H., & Haynes, J.-D. (2014). The relationship between perceptual decision variables and confidence in the human brain. *Cerebral Cortex*.

Heekeren, H. R., Marrett, S., Bandettini, P. A., & Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature*, *431*(7010), 859–62.

Heekeren, H. R., Marrett, S., Ruff, D. A., Bandettini, P. A., & Ungerleider, L. G. (2006). Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(26), 10023–8.

Helie, S., Roeder, J. L., & Ashby, F. G. (2010). Evidence for cortical automaticity in rule-based categorization. *The Journal of Neuroscience*, *30*(42), 14225–34.

Hinrichs, H., Scholz, M., Tempelmann, C., Woldorff, M. G., Dale, a M., & Heinze, H. J. (2000). Deconvolution of event-related fMRI responses in fast-rate experimental designs: tracking amplitude variations. *Journal of Cognitive Neuroscience*, *12 Suppl 2*, 76–89.

Ho, T. C., Brown, S., & Serences, J. T. (2009). Domain general mechanisms of perceptual decision making in human cortex. *The Journal of Neuroscience*, *29*(27), 8675–87.

Hoenig, J. M., & Heisey, D. M. (2001). The abuse of power: The pervasive fallacy of power calculations for data analysis. *The American Statistician*, *55*(1), 1–6.

Huettel, S. A., & McCarthy, G. (2000). Evidence for a refractory period in the hemodynamic response to visual stimuli as measured by MRI. *NeuroImage*, *11*, 547–553.

Hutchinson, J. B., Uncapher, M. R., & Wagner, A. D. (2014). Increased functional connectivity between dorsal posterior parietal and ventral occipitotemporal cortex during uncertain memory decisions. *Neurobiology of Learning and Memory*, *117*, 71–83.

Insabato, A., Pannunzi, M., Rolls, E. T., & Deco, G. (2010). Confidence-related decision making. *Journal of Neurophysiology*, *104*, 539–547.

Ivanoff, J., Branning, P., & Marois, R. (2008). FMRI evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PloS One*, *3*(7), e2635.

Johansen, M. K., & Palmeri, T. J. (2002). Are there representational shifts during category learning? *Cognitive Psychology*, *45*(4), 482–553.

Joyce, K. E., & Hayasaka, S. (2012). Development of PowerMap: A software package for statistical power calculation in neuroimaging studies. *Neuroinformatics*, *10*(4), 351–65.

Kayser, A. S., Buchsbaum, B. R., Erickson, D. T., & D'Esposito, M. (2010). The functional anatomy of a perceptual decision in the human brain. *Journal of Neurophysiology*, *103*(3), 1179–94.

Kayser, A. S., Erickson, D. T., Buchsbaum, B. R., & D'Esposito, M. (2010). Neural representations of relevant and irrelevant features in perceptual decision making. *The Journal of Neuroscience*, *30*(47), 15778–89.

Kepecs, A., & Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *367*(1594), 1322–37.

Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, *455*(7210), 227–31.

Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science (New York, N.Y.)*, *324*(5928), 759–64.

Kincaid, A. E., Zheng, T., & Wilson, C. J. (1998). Connectivity and convergence of single corticostriatal axons. *The Journal of Neuroscience*, *18*(12), 4722–31.

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719.

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*, 106–120.

Lagdano, D. A., Newell, B., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple cue learning. *Journal of Experimental Psychology: General*, *135*(2), 162.

Levine, M., & Ensom, M. H. (2001). Post hoc power analysis: An idea whose time has passed? *Pharmacotherapy*, *21*(4), 405–409.

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. John Wiley and Sons.

Maddox, W. T., & Ashby, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes*, *66*(3), 309–332.

Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(4), 650–662.

Maddox, W. T., Ashby, F. G., Ing, A. D., & Pickering, A. D. (2004). Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition*, *32*(4), 582–91.

Maddox, W. T., Filoteo, J. V., Hejl, K. D., & Ing, A. D. (2004). Category number impacts rule-based but not information-integration category learning: further evidence for dissociable category-learning systems. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *30*(1), 227–45.

Maddox, W. T., Glass, B. D., O'Brien, J. B., Filoteo, J. V., & Ashby, F. G. (2010). Category label and response location shifts in category learning. *Psychological Research*, *74*(2), 219–36.

Maddox, W. T., Love, B., Glass, B., & Filoteo, J. V. (2008). When more is less: Feedback effects in perceptual category learning. *Cognition*, *108*(2), 578–589.

Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, *21*(1), 422–430.

McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*(1), 103–7.

Meeter, M., Myers, C. E., Shohamy, D., Hopkins, R. O., & Gluck, M. A. (2006). Strategies in probabilistic categorization: Results from a new way of analyzing performance. *Learning & Memory*, *13*, 230–239.

Meeter, M., Radics, G., Myers, C. E., Gluck, M. A., & Hopkins, R. O. (2008). Probabilistic categorization: How do normal participants and amnesic patients do it? *Neuroscience and Biobehavioral Reviews*, *32*(2), 237–48.

Mink, J. W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, *50*(4), 381–425.

Mishkin, M., Malamut, B., & Bachevalier, J. (1984). Memories and habits: Two neural systems. *Neurobiology of Human Learning and Memory*, 65–77.

Miyachi, S., Hikosaka, O., & Lu, X. (2002). Differential activation of monkey striatal neurons in the early and late stages of procedural learning. *Experimental Brain Research*, *146*(1), 122–6.

Mumford, J. A. (2012). A power calculation guide for fMRI studies. *Social Cognitive and Affective Neuroscience*, *7*(6), 738–42.

Nambu, A., Tokuno, H., & Hamada, I. (2000). Excitatory cortical inputs to pallidal neurons via the subthalamic nucleus in the monkey. *Journal of Neurophysiology*, *84*, 289–300.

Newsome, W. T., Britten, K. H., & Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. *Nature*, *341*(6237), 52–54.

Niv, Y., Daniel, R., Geana, a., Gershman, S. J., Leong, Y. C., Radulescu, a., & Wilson, R. C. (2015). Reinforcement Learning in Multidimensional Environments Relies on Attention Mechanisms. *Journal of Neuroscience*, *35*(21), 8145–8157.

Niyogi, R. K., & Wong-Lin, K. (2013). Dynamic excitatory and inhibitory gain modulation can produce flexible, robust and optimal decision-making. *PLoS Computational Biology*, *9*(6), e1003099.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology General*, *115*(1), 39–61.

Nosofsky, R. M., Gluck, M. A., Palmeri, T. J., McKinley, S. C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition*, *22*(3), 352–369.

Padoa-Schioppa, C. (2011). Neurobiology of economic choice: A good-based model. *Annual Review of Neuroscience*, *34*, 333–359.

Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 376–404. doi:10:1167/5.5.1

Pastor-Bernier, A., & Cisek, P. (2011). Neural correlates of biased competition in premotor cortex. *The Journal of Neuroscience*, *31*(19), 7083–8.

Paul, E. J., & Ashby, F. G. (2013). A neurocomputational theory of how explicit learning bootstraps early procedural learning. *Frontiers in Computational Neuroscience*, *7*(December), 177.

Penny, W. D. (2012). Comparing dynamic causal models using AIC, BIC and free energy. *NeuroImage*, *59*(1), 319–30.

Persaud, N., Mcleod, P., & Cowey, A. (2007). Post-decision wagering objectively measures awareness. *Nature Neuroscience*, *10*(2), 257–261.

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., & Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. *Current Biology*, *21*(11), 980–3.

Philiastides, M. G., Biele, G., & Heekeren, H. (2010). A mechanistic account of value computation in the human brain. *Proceedings of the National Academy of Sciences*, *107*(20), 9430–9435.

Ploran, E. J., Nelson, S. M., Velanova, K., Donaldson, D. I., Petersen, S. E., & Wheeler, M. E. (2007). Evidence accumulation and the moment of recognition: Dissociating perceptual recognition processes using fMRI. *The Journal of Neuroscience*, *27*(44), 11912–24.

Ploran, E. J., Tremel, J. J., Nelson, S. M., & Wheeler, M. E. (2011). High quality but limited quantity perceptual evidence produces neural accumulation in frontal and parietal cortex. *Cerebral Cortex*, *21*(11), 2650–62.

Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. a. (2001). Interactive memory systems in the human brain. *Nature*, *414*(6863), 546–50.

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 1–9.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59:108.

Ratcliff, R., & Frank, M. J. (2012). Reinforcement-based decision making in corticostriatal circuits: Mutual constraints by neurocomputational and diffusion models. *Neural Computation*, *24*(5), 1186–1229.

Ratcliff, R., & Rouder, J. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356.

Reddi, B. A., & Carpenter, R. H. (2000). The influence of urgency on decision time. *Nature Neuroscience*, *3*(8), 827–30.

Redgrave, P., Prescott, T., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*(4), 1009–1023.

Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies - revisited. *NeuroImage*, *84*, 971–85.

Rissman, J., Gazzaley, A., & D'Esposito, M. (2004). Measuring functional connectivity during distinct stages of a cognitive task. *NeuroImage*, *23*(2), 752–63.

Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, *22*(21), 9475–89.

Rolls, E. T., Grabenhorst, F., & Deco, G. (2010a). Choice, difficulty, and confidence in the brain. *NeuroImage*, *53*(2), 694–706.

Rolls, E. T., Grabenhorst, F., & Deco, G. (2010b). Decision-making, errors, and confidence in the brain. *Journal of Neurophysiology*, *104*(5), 2359–74.

Rosa, M. J., Bestmann, S., Harrison, L., & Penny, W. (2010). Bayesian model selection maps for group studies. *NeuroImage*, *49*(1), 217–24. doi:10.1016/j.neuroimage.2009.08.051

Salinas, E., & Abbott, L. F. (1996). A model of multiplicative neural responses in parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(21), 11956–61.

Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of action-specific reward values in the striatum. *Science (New York, N.Y.)*, *310*(5752), 1337–40.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–9.

Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., & Friston, K. (2014). The dopaminergic midbrain encodes the expected certainty about desired outcomes. *Cerebral Cortex*, 1–12.

Selen, L. P. J., Shadlen, M. N., & Wolpert, D. M. (2012). Deliberation in the motor system: Reflex gains track evolving evidence leading to a decision. *The Journal of Neuroscience*, *32*(7), 2276–86.

Seth, A. K. (2008). Post-decision wagering measures metacognitive content, not sensory consciousness. *Consciousness and Cognition*, *17*(3), 981–3.

Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: Seeing and deciding. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(2), 628–33.

Shepard, R. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, *75*(13).

Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, *22*(4), 325–345.

Shohamy, D., Myers, C. E., Onlaor, S., & Gluck, M. A. (2004). Role of the basal ganglia in category learning: How do patients with Parkinson's disease learn? *Behavioral Neuroscience*, *118*(4), 676–86.

Simen, P., Cohen, J., & Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks*, *19*(8), 1013–1026.

Smith, J. D., Berg, M. E., Cook, R. G., Murphy, M. S., Crossley, M. J., Boomer, J., … Grace, R. C. (2012). Implicit and explicit categorization: A tale of four species, *36*, 2355–2369.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–8.

Speekenbrink, M., Lagnado, D. a., Wilkinson, L., Jahanshahi, M., & Shanks, D. R. (2010). Models of probabilistic category learning in Parkinson's disease: Strategy use and the effects of L-dopa. *Journal of Mathematical Psychology*, *54*(1), 123–136.

Stalnaker, T. A., Calhoon, G. G., Ogawa, M., Roesch, M. R., & Schoenbaum, G. (2012). Reward prediction error signaling in posterior dorsomedial striatum is action specific. *The Journal of Neuroscience*, *32*(30), 10296–305.

Standage, D., You, H., Wang, D.-H., & Dorris, M. C. (2011). Gain modulation by an urgency signal controls the speed-accuracy trade-off in a network model of a cortical decision circuit. *Frontiers in Computational Neuroscience*, *5*, 7.

Standage, D., You, H., Wang, D.-H., & Dorris, M. C. (2013). Trading speed and accuracy by coding time: A coupled-circuit cortical model. *PLoS Computational Biology*, *9*(4), e1003021.

Stephan, K., Penny, W., & Daunizeau, J. (2009). Bayesian model selection for group studies. *Neuroimage*, *46*(4), 1004–1017.

Swaminathan, S. K., & Freedman, D. J. (2012). Preferential encoding of visual categories in parietal cortex compared with prefrontal cortex. *Nature Neuroscience*, *15*(2), 315–20.

Thura, D., Beauregard-Racine, J., Fradet, C.-W., & Cisek, P. (2012). Decision-making by urgency-gating: Theory and experimental support. *Journal of Neurophysiology*, *108*, 2912–2930.

Thura, D., & Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*, *81*(6), 1401–1416.

Thura, D., Cos, I., Trung, J., & Cisek, P. (2014). Context-dependent urgency influences speed-accuracy trade-offs in decision-making and movement execution. *Journal of Neuroscience*, *34*(49), 16442–16454.

Tibshirani, R. (1996). Regression shinkage and selection via the Lasso. *J. R. Statist. Soc. B*, *58*(1), 267–288.

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

Tosoni, A., Galati, G., Romani, G. L., & Corbetta, M. (2008). Sensory-motor mechanisms in human parietal cortex underlie arbitrary visual decisions. *Nature Neuroscience*, *11*(12), 1446–53.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science (New York, N.Y.)*, *185*(4157), 1124–31.

Ulrich, D., & Huguenard, J. R. (1997). GABA A -receptor-mediated rebound burst firing and burst shunting in thalamus. *Journal of Neurophysiology*, *78*, 1748–1751.

Usher, M., & McClelland, J. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592.

Van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E.-J., Ho, T., Serences, J., & Forstmann, B. U. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *The Journal of Neuroscience*, *31*(48), 17488–17495.

Van Ravenzwaaij, D., van der Maas, H. L. J., & Wagenmakers, E.-J. (2012). Optimal decision making in neural inhibition models. *Psychological Review*, *119*(1), 201–15.

Van Veen, V., Krug, M. K., & Carter, C. S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. *Journal of Cognitive Neuroscience*, *20*(11), 1952–65.

Vazquez, A. L., & Noll, D. C. (1998). Nonlinear aspects of the BOLD response in functional MRI. *NeuroImage*, *7*(2), 108–118.

Vickers, D., & Packer, J. (1982). Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica*, *50*, 179–197.

Wager, T. D., Vazquez, A., Hernandez, L., & Noll, D. C. (2005). Accounting for nonlinear BOLD effects in fMRI: Parameter estimates and a model for prediction in rapid event-related studies. *NeuroImage*, *25*(1), 206–18.

Waldschmidt, J. G., & Ashby, F. G. (2011). Cortical and striatal contributions to automaticity in information-integration categorization. *NeuroImage*, *56*(3), 1791–802.

Wang, X. (2008). Decision making in recurrent neuronal circuits. *Neuron*, *60*(2), 215–234.

Wheeler, M., Woo, S., Ansel, T., Tremel, J., Collier, A. L., Velanova, K., … Yang, T. (2014). The strength of gradually accruing probabilistic evidence modulates brain activity during a categorical decision. *Journal of Cognitive Neuroscience*, *27*(4), 705–719.

White, C. N., Mumford, J. A., & Poldrack, R. A. (2012). Perceptual criteria in the human brain. *The Journal of Neuroscience*, *32*(47), 16716–24.

Wiecki, T. V, & Frank, M. J. (2013). A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychological Review*, *120*, 329–55.

Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience*, *26*(4), 1314–28.

Wunderlich, K., Dayan, P., & Dolan, R. J. (2012). Mapping value based planning and extensively trained choice in the human brain. *Nature Neuroscience*, *15*(5), 786–791.

Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, *447*(7148), 1075–80.

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, *7*(6), 464–76.

Zheng, T., & Wilson, C. J. (2002). Corticostriatal combinatorics: The implications of corticostriatal axonal arborizations. *Journal of Neurophysiology*, *87*(2), 1007–17.

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic-net. *Journal of the Royal Statistical Society*, *67*, 301–320.

APPENDIX A

Abbreviations

BA	Brodmann Area
BOLD	Blood Oxygen Level Dependent
CmLogLR	Cumulative Log Likelihood Ratio
COVIS	COmpetition between Verbal and Implicit Systems
DDM	Drift Diffusion Model
DLPFC	Dorsolateral Prefrontal Cortex
FDR	False Discovery Rate
FEF	Frontal Eye Fields
FWE	Familywise Error Rate
GPi	Internal Segment of the Globus Pallidus
LIP	Lateral Intraparietal Area
IPS	Intraparietal Sulcus
MNI	Montreal Neurological Institute
MT	Middle Temporal Area
OFC	Orbitofrontal Cortex
PM	Parametric Modulator
PXP	Protected Exceedance Probability
RFX-BMS	Random Effects Bayesian Model Selection
RLPFC	Rostrolateral Prefrontal Cortex
SAT	Speed Accuracy Trade-Off
SDT	Signal Detection Theory

SMA	Supplementary Motor Area
SNc	Substantia Nigra Pars Compacta
SPM	Statistical Parametric Mapping
SSM	Sequential Sampling Model
STN	Subthalamic Nucleus
sWOE	Subjective Weights of Evidence
TR	Repetition Time
VLPFC	Ventrolateral Prefrontal Cortex
VTA	Ventral Tegmental Area

TABLE A.1. Pilot Study: Power Analyses. "n" indicates the estimated number of participants required to reach acheive statistical power of 0.8, assuming correction for multiple comparisons with FDR ($q < 0.05$).

PM	x	y	z	n	Region	BA
Categorical						
Evidence	-7	-6	68	13	Supp_Motor_Area_L	6
(Left)	-30	-3	67	15	Frontal_Sup_L	6
	-40	-14	41	13	Postcentral_L	3
	-14	22	7	17	Caudate_L	-
	-14	0	24	14	Caudate_L	-
	4	-14	72	13	Supp_Motor_Area_R	6
	14	-3	72	15	Frontal_Sup_R	6
	28	-57	-32	15	Cerebelum_6_R	37
	50	-51	16	13	Temporal_Mid_R	21
Categorical						
Evidence	-35	6	32	12	Precentral_L	44
(Right)	21	20	62	14	Frontal_Sup_R	8
	-9	-24	69	29	Paracentral_Lobule_L	4
	38	-61	58	15	Parietal_Sup_R	7
	47	-71	27	16	Occipital_Mid_R	39

TABLE A.2. Classical Results: Mean Effect of Features. For other contrasts, the topological false discovery rate was used to account for multiple comparisons. However, this contrasts yielded more widespread activation, and so the more stringent family-wise error (FWE) rate was used instead. (Contrast: mean feature regressor > implicit baseline)

PM	Size	x	y	z	t	Region	BA
Mean							
Features	4528	34.5	-60.0	46.5	15.19	Angular_R	7
		-1.5	-78.0	49.5	11.05	Precuneus_L	7
	994	37.5	22.5	-7.5	13.66	Insula_R	47
	2700	4.5	24.0	40.5	12.95	Cingulum_Mid_R	32
	383	22.5	-79.5	1.5	12.63	Lingual_R	18
	44	-28.5	-30.0	1.5	12.09		37
	1039	-25.5	-75.0	-21.0	11.81	Cerebelum_6_L	18
	709	-21.0	-79.5	4.5	11.69	Calcarine_L	18
	1290	-28.5	-49.5	39.0	11.20	Parietal_Inf_L	40
	783	4.5	-1.5	27.0	10.88	Cingulum_Mid_R	-
	2411	10.5	-22.5	-4.5	10.75	Thalamus_R	-
	708	25.5	12.0	55.5	9.92	Frontal_Sup_R	8
	156	9.0	-69.0	-28.5	9.63	Vermis_7	-
	109	28.5	45.0	-21.0	9.30	Frontal_Mid_Orb_R	11
	55	43.5	-72.0	-27.0	9.15	Cerebelum_Crus1_R	19
	34	-13.5	12.0	-6.0	9.08	Putamen_L	25
	240	-6.0	-73.5	-28.5	8.81	Cerebelum_Crus1_L	-
	48	1.5	-36.0	-42.0	8.79		-
	186	-34.5	16.5	1.5	8.75	Insula_L	48
	39	36.0	-42.0	-36.0	8.65	Cerebelum_6_R	37
	43	55.5	-45.0	-13.5	8.51	Temporal_Inf_R	20
	82	3.0	-54.0	-22.5	8.51	Vermis_4_5	-
	11	-7.5	-39.0	19.5	8.19	Cingulum_Post_L	26
	51	13.5	15.0	3.0	8.17	Caudate_R	25
	33	16.5	-102.0	1.5	8.11	Calcarine_R	17
	16	-18.0	-37.5	6.0	8.11	Hippocampus_L	27
	24	-36.0	-43.5	-34.5	8.02	Cerebelum_Crus1_L	37
	27	-28.5	-1.5	45.0	8.00	Precentral_L	6
	18	37.5	-78.0	-19.5	7.93	Cerebelum_Crus1_R	19
	77	42.0	37.5	21.0	7.71	Frontal_Mid_R	45
	20	48.0	30.0	28.5	7.68	Frontal_Inf_Tri_R	45
	10	31.5	-58.5	-1.5	7.59	Fusiform_R	37

TABLE A.3. Classical Results: Task-Negative Activation. (Contrast: implicit baseline > mean feature regressor)

PM	Size	x	y	z	t	Region	BA
Task							
Negative	8859	-3.0	46.5	-7.5	10.86	Frontal_Mid_Orb_L	11
Activation		-19.5	30.0	36.0	9.57	Frontal_Mid_L	9
(Baseline ¿		1.5	4.5	-7.5	6.10	Olfactory_L	25
Features)	1718	-42.0	-75.0	31.5	8.20	Occipital_Mid_L	39
	518	-19.5	-13.5	-22.5	7.15	Hippocampus_L	35
	1939	55.5	-7.5	-18.0	7.09	Temporal_Mid_R	20
		69.0	-28.5	10.5	5.11	Temporal_Sup_R	22
	1834	-58.5	-12.0	-7.5	7.08	Temporal_Mid_L	22
	1269	-58.5	-42.0	4.5	6.79	Temporal_Mid_L	22
	255	-6.0	-54.0	10.5	6.64	Precuneus_L	30
	323	-52.5	27.0	0.0	6.55	Frontal_Inf_Tri_L	45
	258	27.0	-81.0	-36.0	6.52	Cerebelum_Crus2_R	-
	396	-6.0	-43.5	37.5	6.41	Precuneus_L	23
	190	10.5	46.5	51.0	6.21	Frontal_Sup_Medial_R	9
	418	21.0	-13.5	-21.0	6.02	ParaHippocampal_R	35
	268	67.5	-9.0	30.0	5.84	Postcentral_R	43
	362	-40.5	27.0	-18.0	5.41	Frontal_Inf_Orb_L	38
	213	-4.5	-27.0	60.0	5.35	Paracentral_Lobule_L	4

TABLE A.4. Classical Results: Categorical Evidence and Categorical Evidence (multiplicatively) modulated by Urgency.

PM	Size	x	y	z	t	Region	BA
Categorical Evidence (Left)	12316	40.5	-15.0	49.5	13.04	Precentral_R	6
		28.5	-42.0	61.5	7.06	Postcentral_R	2
		4.5	3.0	55.5	5.68	Supp_Motor_Area_R	6
	1416	-13.5	-52.5	-19.5	6.94	Cerebelum_4_5_L	19
	1299	37.5	-10.5	-1.5	6.70	Insula_R	48
	786	42.0	-21.0	24.0	5.95	Rolandic_Oper_R	48
	337	52.5	3.0	4.5	4.72	Rolandic_Oper_R	48
Categorical Evidence (Right)	1947	-36.0	-21.0	69.0	6.18	Precentral_L	6
		-57.0	-18.0	46.5	5.52	Postcentral_L	3
Categorical Evidence by Urgency (Left)	12461	40.5	-15.0	49.5	11.57	Precentral_R	6
		22.5	-55.5	67.5	6.69	Parietal_Sup_R	5
		4.5	3.0	55.5	5.40	Supp_Motor_Area_R	6
	1443	-12.0	-52.5	-19.5	7.02	Cerebelum_4_5_L	19
	1354	36.0	-9.0	-1.5	5.72	Putamen_R	-
Categorical Evidence by Urgency (Right)	1324	-36.0	-21.0	69.0	6.05	Precentral_L	6

TABLE A.5. Classical Results: Confidence and Confidence (multiplicatively) modulated by Urgency.

PM	Size	x	y	z	t	Region	BA
Confidence							
	5427	-58.5	-43.5	39.0	11.86	Parietal_Inf_L	40
		-58.5	-61.5	1.5	6.09	Temporal_Mid_L	37
	8107	55.5	-36.0	45.0	11.07	SupraMarginal_R	40
		66.0	-46.5	10.5	8.53	Temporal_Mid_R	22
		40.5	-69.0	40.5	5.48	Angular_R	7
	46026	9.0	-82.5	21.0	9.01	Cuneus_R	18
		-12.0	-30.0	36.0	8.10	Cingulum_Mid_L	23
		22.5	18.0	10.5	7.33	Caudate_R	-
		28.5	-18.0	3.0	7.32	Putamen_R	-
		-34.5	3.0	13.5	7.12	Rolandic_Oper_L	48
		46.5	46.5	1.5	7.08	Frontal_Mid_R	45
		-46.5	-48.0	-33.0	6.93	Cerebelum_Crus1_L	20
		12.0	21.0	57.0	6.31	Supp_Motor_Area_R	8
		24.0	58.5	27.0	5.98	Frontal_Mid_R	46
		6.0	-66.0	61.5	4.72	Precuneus_R	7
	376	-25.5	-43.5	28.5	6.26		-
	327	-27.0	25.5	28.5	6.11	Frontal_Mid_L	48
	227	64.5	-28.5	-15.0	5.27	Temporal_Mid_R	20
	341	-27.0	52.5	27.0	5.11	Frontal_Mid_L	46
Confidence by Urgency	6155	55.5	-37.5	45.0	9.32	SupraMarginal_R	40
		45.0	-40.5	-1.5	5.82	Temporal_Mid_R	21
	5621	-58.5	-42.0	39.0	9.24	Parietal_Inf_L	40
		-30.0	-64.5	9.0	6.06		19
	21686	10.5	-27.0	34.5	8.23	Cingulum_Mid_R	23
		22.5	18.0	10.5	7.44	Caudate_R	-
		0.0	7.5	36.0	7.15	Cingulum_Mid_L	24
		-22.5	-39.0	66.0	6.92	Postcentral_L	3
		-34.5	3.0	15.0	6.09	Rolandic_Oper_L	48
		21.0	-40.5	69.0	5.18	Postcentral_R	2
		4.5	-6.0	72.0	4.65	Supp_Motor_Area_R	6
		46.5	24.0	-18.0	4.44	Frontal_Inf_Orb_R	38
	3358	9.0	-82.5	21.0	7.31	Cuneus_R	18
	629	-46.5	-48.0	-33.0	6.87	Cerebelum_Crus1_L	20
	396	-27.0	24.0	28.5	6.41	Frontal_Mid_L	48
	1741	-3.0	-63.0	-15.0	6.08	Cerebelum_4_5_L	18
		-33.0	-78.0	-40.5	4.85	Cerebelum_Crus2_L	-
	1831	24.0	58.5	27.0	6.03	Frontal_Mid_R	46
		9.0	30.0	54.0	5.79	Frontal_Sup_Medial_R	8
	886	51.0	33.0	3.0	5.89	Frontal_Inf_Tri_R	45
	555	-25.5	51.0	25.5	4.83	Frontal_Mid_L	46

TABLE A.6. Classical Results: Uncertainty and Uncertainty (multiplicatively) modulated by Urgency

PM	Size	x	y	z	t	Region	BA
Uncertainty							
	4908	42.0	-81.0	-4.5	11.76	Occipital_Inf_R	19
	4628	-37.5	-82.5	-6.0	8.77	Occipital_Inf_L	19
	320	-27.0	-4.5	42.0	7.53	Precentral_L	6
	520	19.5	-63.0	51.0	5.98	Parietal_Sup_R	7
	186	-4.5	12.0	49.5	5.79	Supp_Motor_Area_L	32
	251	-22.5	-64.5	34.5	5.58	Occipital_Mid_L	7
	147	-18.0	-61.5	55.5	5.32	Parietal_Sup_L	7
	216	24.0	-1.5	48.0	5.29	Precentral_R	6
Uncertainty							
by Urgency	5172	42.0	-82.5	-4.5	10.70	Occipital_Inf_R	19
		33.0	-48.0	-15.0	5.17	Fusiform_R	37
	5737	-31.5	-70.5	-13.5	9.72	Fusiform_L	19
		-22.5	-64.5	34.5	6.21	Occipital_Mid_L	7
	329	-6.0	15.0	46.5	6.92	Supp_Motor_Area_L	32
	315	-27.0	-4.5	45.0	6.89	Precentral_L	6
	761	21.0	-61.5	49.5	6.83	Parietal_Sup_R	7
	200	-42.0	1.5	30.0	6.20	Precentral_L	44
	218	25.5	0.0	49.5	5.38	Frontal_Sup_R	6

TABLE A.7. Classical Results: Urgency. For the classical analyses, urgency was included as the second PM in each model. This map represents the results of a conjunction analysis across these models, and the t-values were averaged across models.

PM	Size	x	y	z	t	Region	BA
Urgency							
	79280	39.0	22.5	-10.5	12.92	Frontal_Inf_Orb_R	47
		4.5	-87.0	27.0	11.64	Cuneus_L	18
		4.5	33.0	34.5	11.14	Cingulum_Mid_R	32
		-45.0	-69.0	-31.5	10.25	Cerebelum_Crus1_L	-
		10.5	3.0	1.5	10.04	Pallidum_R	-
		3.0	-18.0	28.5	9.40	Cingulum_Mid_R	23
		43.5	13.5	36.0	9.29	Frontal_Inf_Oper_R	44
		1.5	-58.5	64.5	9.15	Precuneus_R	7
		-60.0	-42.0	45.0	7.63	Parietal_Inf_L	40
		27.0	-69.0	-30.0	7.57	Cerebelum_6_R	19
		-22.5	-31.5	-30.0	7.30	Cerebelum_4_5_L	37
		21.0	-42.0	9.0	6.87	Precuneus_R	37
	6903	52.5	-42.0	45.0	9.35	SupraMarginal_R	40
	2007	55.5	-40.5	-6.0	8.70	Temporal_Mid_R	21
	1881	-30.0	57.0	-12.0	7.94	Frontal_Mid_Orb_L	11

TABLE A.8. Classical Results: Feedback: Positive Prediction Errors Related to Confidence and Categorical Evidence.

PM	Size	x	y	z	t	Region	BA
Confidence							
	516	27.0	12.0	0.0	7.31	Putamen_R	-
	641	-34.5	10.5	16.5	6.25	Insula_L	48
	2404	-16.5	-73.5	-27.0	6.10	Cerebelum_Crus1_L	19
	268	21.0	37.5	-7.5	5.67	Frontal_Inf_Orb_R	11
	283	36.0	-76.5	-40.5	4.87	Cerebelum_Crus2_R	-
Categorical							
Evidence	2460	39.0	-18.0	52.5	6.54	Precentral_R	4
(Left)	364	12.0	-33.0	-40.5	5.66		-
	296	19.5	-69.0	21.0	4.83	Cuneus_R	18

TABLE A.9. RFX BMS Results: Categorical Evidence vs. Categorical Evidence (multiplicatively) modulated by Urgency. Each map was thresholded to include only voxels where the model provided the best fit.

Model	Size	x	y	z	PXP	Region	BA
Categorical							
Evidence	192	-3	-60	-21	0.80	Vermis_6	-
	792	15	-30	51	0.75	Paracentral_Lobule_R	-
		42	-6	45	0.69	Precentral_R	6
	33	15	-18	0	0.71	Thalamus_R	-
	141	27	-6	9	0.67	Putamen_R	48
	11	18	12	63	0.63	Frontal_Sup_R	6
	25	51	3	9	0.62	Rolandic_Oper_R	48
Categorical							
Evidence	817	42	-21	60	0.94	Precentral_R	4
by Urgency		12	-9	75	0.52	Supp_Motor_Area_R	6
	240	-39	-24	63	0.93	Precentral_L	4
	23	3	-3	72	0.67	Supp_Motor_Area_R	6
	69	60	-24	18	0.66	SupraMarginal_R	42
	20	57	9	0	0.64	Rolandic_Oper_R	48

TABLE A.10. RFX BMS Results: Confidence vs. Confidence (multiplicatively) modulated by Urgency. Each map was thresholded to include only voxels where the model provided the best fit.

Model	Size	x	y	z	PXP	Region	BA
Confidence							
	1991	12	30	42	0.80	Frontal_Sup_Medial_R	32
		-12	-27	51	0.76	Cingulum_Mid_L	4
		0	-60	-18	0.71	Vermis_6	-
		-45	-63	-39	0.64	Cerebelum_Crus2_L	-
	727	18	3	12	0.73		-
	544	-24	6	12	0.71	Putamen_L	48
	317	-42	-60	39	0.70	Angular_L	39
	279	45	-24	33	0.69	Postcentral_R	48
	65	36	42	0	0.69		47
	21	24	51	-12	0.67	Frontal_Mid_Orb_R	11
	57	30	30	33	0.66	Frontal_Mid_R	46
	22	-42	-6	18	0.64	Rolandic_Oper_L	48
	37	-36	-81	-42	0.63	Cerebelum_Crus2_L	-
	11	54	-51	-9	0.62	Temporal_Inf_R	37
	19	57	-30	-15	0.60	Temporal_Inf_R	20
Confidence by Urgency	741	33	-66	54	0.85	Parietal_Sup_R	7
		63	-48	6	0.73	Temporal_Mid_R	21
	461	6	-81	42	0.82		19
	182	51	15	-9	0.80		38
	508	21	57	30	0.79	Frontal_Sup_R	9
		-3	21	63	0.74	Supp_Motor_Area_L	8
		54	36	-15	0.65	Frontal_Inf_Orb_R	-
	100	-48	15	-9	0.75	Temporal_Pole_Sup_L	38
	44	-21	60	24	0.73	Frontal_Sup_L	10
	168	-30	-72	-27	0.68	Cerebelum_Crus1_L	19
	30	0	-6	69	0.67	Supp_Motor_Area_L	6
	100	30	21	54	0.67	Frontal_Mid_R	8
	375	-57	-60	9	0.66	Temporal_Mid_L	37
		-60	-24	15	0.66	SupraMarginal_L	42
	28	33	60	-6	0.65	Frontal_Sup_Orb_R	11
	57	-27	-39	72	0.64	Postcentral_L	1
	238	0	27	24	0.64	Cingulum_Ant_L	24
		3	-18	42	0.62	Cingulum_Mid_R	23
	12	-36	-69	51	0.63	Parietal_Sup_L	7
	78	24	3	-12	0.62	Amygdala_R	34
	24	33	-42	66	0.62	Postcentral_R	2
	42	9	-66	-12	0.62	Cerebelum_6_R	18
	46	6	-48	72	0.62	Precuneus_R	5

TABLE A.11. RFX BMS Results: Uncertainty vs. Uncertainty (multiplicatively) modulated by Urgency. Each map was thresholded to include only voxels where the model provided the best fit.

Model	Size	x	y	z	PXP	Region	BA
Uncertainty							
	508	18	-81	-9	0.77	Lingual_R	18
		36	-48	-15	0.62	Fusiform_R	37
	56	15	-60	54	0.72	Precuneus_R	5
	558	-33	-75	-9	0.71	Occipital_Inf_L	19
	35	-24	0	45	0.68		6
	41	-9	12	48	0.67	Supp_Motor_Area_L	32
	26	24	-3	45	0.65		6
Uncertainty by Urgency	54	-24	-63	54	0.70	Parietal_Sup_L	7
	108	33	-87	-15	0.67	Occipital_Inf_R	19
	28	30	-57	54	0.66	Parietal_Sup_R	7
	32	-27	-87	27	0.65	Occipital_Mid_L	19
	21	-48	3	36	0.63	Precentral_L	6
	12	51	-69	-3	0.63	Temporal_Inf_R	37
	29	30	-87	27	0.62	Occipital_Sup_R	19
	42	-33	-93	0	0.61	Occipital_Mid_L	18
	17	-45	-78	-12	0.60	Occipital_Inf_L	19