Different varieties of uncertainty in human decision-making

Amy R. Bland[1]* and Alexandre Schaefer[2]

[1] Neuroscience and Psychiatry Unit, University of Manchester, Manchester, UK
[2] Psychology Department, University of Durham, Durham, UK

The study of uncertainty in decision-making is receiving greater attention in the fields of cognitive and computational neuroscience. Several lines of evidence are beginning to elucidate different variants of uncertainty. Particularly, risk, ambiguity, and expected and unexpected forms of uncertainty are well articulated in the literature. In this article we review both empirical and theoretical evidence arguing for the potential distinction between three forms of uncertainty; expected uncertainty, unexpected uncertainty, and volatility. Particular attention will be devoted to exploring the distinction between unexpected uncertainty and volatility which has been less appreciated in the literature. This includes evidence mainly from neuroimaging, neuromodulation, and electrophysiological studies. We further address the possible differentiation of cognitive control mechanisms used to deal with these forms of uncertainty. Finally, we explore whether the dual modes of control theory provides a theoretical framework for understanding the distinction between unexpected uncertainty and volatility.

Keywords: uncertainty, unexpected uncertainty, volatility, decision-making

INTRODUCTION

Uncertainty is a common feature of many every day decisions. Uncertainty typically arises in a situation that has limited or incalculable information about the predicted outcomes of behavior (Huettel et al., 2005). Successfully detecting, processing and resolving uncertainty is important to successful adaptive behavior. Recent years have seen a growing body of research dedicated to exploring the brain mechanisms which underlie our choices during conditions of uncertainty. However, it is becoming clear that "*uncertainty*" is not comprised of a single dimension. More recent evidence is beginning to differentiate neural correlates involved in estimating, representing, and resolving different forms of uncertainty. For example, studies have demonstrated separable neural correlates of *reward expectancy* and *variance* (Preuschoff et al., 2006; Tobler et al., 2007), *reward probability* and *magnitude* (Knutson et al., 2005), and *ambiguity* and *risk* (Hsu et al., 2005; Huettel et al., 2006). A major contribution of this work has been a better understanding of how uncertainty can be induced by different variables in the decision-making environment. However, an important form of uncertainty which has received less attention is uncertainty induced by unexpected changes in learned Stimulus-Response-Outcome (S-R-O) contingences, often referred to as "unexpected uncertainty" or "volatility." However, as we will discuss below, unexpected uncertainty and volatility do not necessarily refer to the same phenomenon. Therefore, we review theoretical and empirical arguments supporting a potential distinction between three different forms of uncertainty: *expected uncertainty, unexpected uncertainty*, and *volatility*.

DISTINCT VARIETIES OF UNCERTAINTY

Successful decision-making relies on one's ability to form a stable representation of the underlying S-R-O rules learned from previous experience of gains and losses (e.g., Sutton and Barto, 1998; Ridderinkhof et al., 2004; Seymour et al., 2007). As such, agents can learn that a specific association between a stimulus (S) and a response (R) is linked with a positive or negative outcome (O). For instance, we may choose to enter (R) a particular restaurant (S) if we have previously found that it serves our preferred dish (O). Therefore through learning these associations between a *Stimulus* (restaurant), a *Response* (enter), and its positive or negative *Outcome* (preferred dish) we can guide future decision-making in order to choose the *Response* which will most likely lead to a rewarding *Outcome*. When faced with this kind of decision, an agent has a prediction or expectation of the probability of an outcome. This is derived from the recent history of outcomes of that choice (Sutton and Barto, 1998). Therefore an agent must have the ability to learn these S-R-O relationships and the likelihood to which they occur in order to make the most optimal choices. If we take the example above, our behavioral choice may be caused by previous experiences in which we learned that our preferred dish is available 8 out of 10 visits to that restaurant.

One of the most frequent methods used to manipulate uncertainty usually involves the systematic variation of the probability of these learned S-R-O contingencies. Using the example above, if we begin to learn that our preferred dish is available only 6 out of 10 visits, this increases uncertainty about the potential outcome (i.e., preferred dish) if we choose to enter this particular restaurant. In other words, when an agent is faced with a two-options choice, uncertainty is often said to be maximal when the probability of obtaining a reward linked to any of the two options is $p = 0.5$ but absent at the two extremes (probability = 1 and probability = 0; e.g., Fiorillo et al., 2003). Many studies have therefore used variations of a 75/25% S-R-O probability to create certain environments and 50/50% probability to create uncertain environments (Volz et al., 2003; Paulus et al., 2004; Huettel et al., 2005; Krain et al., 2006; Cohen et al., 2007b; Polezzi et al., 2008). It is also

possible to explore varying degrees of uncertainty (e.g., Volz et al., 2003; Huettel et al., 2005). Typically in these studies, participants are shown cues which are probabilistic predictors of a given outcome (e.g., a red triangle that predicts the occurrence of a reward on 80% of trials). Uncertainty in these paradigms is induced by lowering the predictability of the learned stimulus-response (S-R) association being rewarded (O). For instance, varying degrees of uncertainty may include 100, 90, 80, 70, 60, or 50% whereby 50% is the most uncertain and 100% being the least uncertain. Furthermore, if the predictability goes below 50% then uncertainty will decrease again, i.e., 40, 30, 20, 10, 0%.

A wealth of literature has begun to elucidate how the brain estimates, represents, and resolves this form of uncertainty which is induced by varying levels of probability. Neuroimaging evidence indicates that the DLPFC (Paulus et al., 2002; Huettel et al., 2005), posterior parietal cortex (Volz et al., 2003; Huettel et al., 2005), anterior cingulate cortex (ACC; Elliott and Dolan, 1998; Critchley et al., 2001; Stern et al., 2010), orbito-frontal cortex (OFC; Goel and Dolan, 2000; Critchley et al., 2001; Hsu et al., 2005; Tobler et al., 2007), and amygdala (Hsu et al., 2005) are involved in processing uncertainty. Electrophysiological evidence points to modulation of the P3, a positive going potential peaking around 300 ms post stimulus onset suggesting greater positivities are associated with greater uncertainty (Duncan-Johnson and Donchin, 1977; Donchin and Coles, 1988; Polich, 1990).

Importantly however, uncertainty can also be induced by unexpected changes in S-R-O contingencies, above and beyond the current S-R-O probability levels. For instance, using the example above, we may choose to enter a particular restaurant if we have previously found that our preferred dish is available 8 out of 10 visits to a particular restaurant. However, uncertainty could be induced if this S-R-O contingency suddenly changes because the usual kitchen chef was fired and replaced by another chef with different menu preferences which would take the "preferred dish probability" to 0.2 on that week. In this case, the choice of available dishes in that restaurant right after the replacement of the chef will be uncertain because it can no longer be predicted by past experience. Therefore uncertainty can be induced not only by lowering the probability of S-R-O contingencies, but also by fundamental changes in these contingencies that forces a modification of our previous beliefs.

More recent approaches to uncertainty have begun to establish that the two forms of uncertainty illustrated above refer to two distinct processes. Particularly, uncertainty can arise from (a) the stochasticity inherent in the decision-making environment (e.g., the stable probability of reward where an agent can learn that a stimulus predicts rewards on 80% of trials is less uncertain than a situation where this probability is set at 50%), and (b) from unexpected and fundamental changes in the S-R-O contingencies of the environment that invalidate prediction based on previous experience (Yu and Dayan, 2005; Courville et al., 2006; Behrens et al., 2007; Doya, 2008; Rushworth and Behrens, 2008; Krugel et al., 2009; Nassar et al., 2010; Payzan-LeNestour and Bossaerts, 2011). The former is usually referred to as *expected uncertainty* (Yu and Dayan, 2005) or *Feedback Validity* (e.g., Bland and Schaefer, 2011), and the latter is often referred to as *unexpected uncertainty* (Yu and Dayan, 2005).

Recent developments suggest that *volatility,* has also to be considered (Behrens et al., 2007; Bland and Schaefer, 2011). Volatility can be defined as a variation in the frequency of changes in existing S-R-O contingencies across time. In our example above, a stable situation (low volatility) is attained when our preferred dish is served in our chosen restaurant 8 days out of 10 during an entire year. However, a volatile situation can arise if the manager of the restaurant decides to dynamically change the menu several times during the year. In such case, the "preferred dish probability" will frequently change (e.g., 0.9 in the first week, 0.2 in the second week, 0.7 in the third week, etc.). In this case, the dynamic changes in S-R-O contingencies will constrain agents to continually update their representation of the environment in order to obtain accurate prediction levels. Therefore volatility and unexpected uncertainty can be distinguished by the frequency of contingency changes. Unexpected uncertainty is characterized by rare unpredicted changes in underlying S-R-O rules, whereas high volatility is characterized by frequent occurrences of fundamental changes in S-R-O rules. In addition, it is important to note that a high frequency of changes may potentially cause agents to learn that changes occur rapidly. Therefore, volatility can be *expected* by decision-making agents.

In summary, three distinct forms of uncertainty can be identified: (1) Expected uncertainty: S-R-O rules learned from past events are weak predictors of the outcomes of future actions, and this unreliability is known and stable. (2) Unexpected uncertainty: a rare fundamental change in the environment which invalidates existing S-R-O rules that are no longer able to accurately predict the outcomes of our actions. (3) Volatility: frequent changes in the environment which require a constant updating of S-R-O rules[1].

THE MODEL OF YU AND DAYAN (2005)

Yu and Dayan (2005) have proposed a distinction between *expected* and *unexpected* forms of uncertainty. Yu and Dayan employed a task involving a set of arrows pointing to the left or right hand side of a screen. The directions of the colored arrows are randomized independently of each other on every trial, but one of them, the *cue*, specified by its color, predicts the location of the subsequent target (a light bulb) with a significant probability whilst the rest of the arrows are irrelevant distracters. The color of the cue arrow (i.e., the "relevant" color which generally predicts the location of the light bulb) persists over many trials, defining a relatively stable context. However, the relevant cue color can suddenly change without informing the subject. According to Yu and Dayan's (2005)

---

[1]Although the main focus of this article was on human decision-making, these three forms of uncertainty can also potentially be encountered in animals. For instance, an animal may learn that pressing (response) a blue lever (stimulus) is paired with food (outcome) for 7 out of 10 lever presses. It is reasonable to expect that through learning, the animal might form a representation of the expected amount of error (30%) in this S-R-O contingency (expected uncertainty). If the blue lever predicts food delivery for only 2 out of 10 lever presses, there will be a fundamental change in the contingencies that were previously guiding behavior (unexpected uncertainty), and the animal must adapt to this new situation (probably through the exploration of other levers present in the environment). However, if the association between the lever press and the food reward is constantly changing then the environment becomes volatile and in order to adapt, an optimal solution would be for the animal to form a representation of the fact that these S-R-O contingencies are likely to frequently change.

influential theory, expected uncertainty arises from known unreliability of predictive relationships within a familiar environment (e.g., learning that the relevant color predicts the location of the light bulb on 80% of trials) whereas unexpected uncertainty is induced by fundamental changes in the environment that produce sensory observations strongly violating expectations (e.g., the previously relevant color no longer predicts the location of the target). The former has been equated with environmental stochasticity in an otherwise stable S-R-O relationship (Nassar et al., 2010). This stochasticity is analogous to uncertainty induced by manipulating the predictive value of decision cues. So an agent can learn that a cue predicts a reward on 80% of trials and so on 20% of trials the outcome is not a valid predictor of the S-R-O relationship. This creates a level of expected uncertainty in a familiar environment which can be thought of as the expected amount of error. Indeed, an agent learns to expect that there will be a certain amount of uncertainty when making their decision through sampling the environment. In other words, expected uncertainty remains the same as long as the 20–80% contingencies are maintained, but unexpected uncertainty increases temporarily during an uncued reversal from 80 to 20%.

Unexpected uncertainty arising from fundamental changes in learned predictive relationships should signal for a revision of an agent's belief about the best course of action. Unexpected uncertainty must therefore require a mechanism for suppressing potentially outdated expectations and encouraging faster adaptation to new S-R-O contingencies (Dayan and Yu, 2002). Indeed, learning rate parameters tend to increase during periods of unexpected uncertainty (Yu and Dayan, 2005) and volatility (Behrens et al., 2007; Nassar et al., 2010). In this way, fundamental changes in S-R-O contingencies increase uncertainty, and speed up subsequent learning, by making historical outcomes irrelevant and new outcomes influencing beliefs strongly (Courville et al., 2006; Nassar et al., 2010). Furthermore, surprise induced by changes in S-R-O contingencies can enhance the speed of learning whereas random variation under constant probabilities (as with a sequence of coin flips) will not be surprising (Courville et al., 2006).

Taken together, these findings highlight the importance of considering different forms of uncertainty and how they interact to produce adaptive behavior. Importantly, an agent must possess the neural and cognitive mechanisms to detect if an S-R-O contingency has changed by representing the probabilistic chance that an error is caused by inherent stochasticity. This parameter is crucial for determining a contingency change. For instance, during everyday decision-making, there is often only a probabilistic chance (rather than a certainty) of success therefore the lack of reward on a particular occasion may not necessarily signal the need to switch to an alternative course of action (Kennerley et al., 2006). Therefore, when a participant responds according to the learned S-R-O rule and receives negative feedback, they must possess the ability to infer whether the erroneous response is due to the inherent stochasticity of the task or whether the S-R-O rule has fundamentally changed. Therefore a changing world requires a mechanism which will allow the successful detection and adaptation to both forms of uncertainty. We will discuss in the remainder of this article some mechanisms potentially involved in this adaptation. Most research to date has conceptualized uncertainty as variations in expected uncertainty (as well as slightly different forms of uncertainty, such as ambiguity and risk). Although unexpected uncertainty has received less attention, there is now a growing body of research that has tackled this phenomenon. In addition, the potential distinction between unexpected uncertainty and volatility has not received much attention, and both concepts tend to be somewhat confounded in the literature. We will therefore review existing evidence on unexpected uncertainty, and we will also review the possibility that specific cognitive strategies might be employed for volatility which are not necessarily employed for unexpected uncertainty.

COMPUTATIONAL MODELING OF UNEXPECTED UNCERTAINTY AND VOLATILITY

Modeling human behavior using computational approaches has provided some important insights into the potential mechanisms involved in decision-making under uncertainty. Such behavior can be modeled by Bayesian algorithms (Behrens et al., 2007; Nassar et al., 2010; Mathys et al., 2011). Indeed, Bayesian statistical theory formalizes the notion that optimal inference and learning depend critically on representing and processing the various sorts of uncertainty associated with a behavioral context (Yu and Dayan, 2005). In the specific case of volatility, it has been suggested that humans adapt to a volatile decision-making environment following Bayesian rules (Behrens et al., 2007; Nassar et al., 2010). Particularly, Behrens et al. (2007) showed that using an ideal Bayesian model, human participants can optimally assess volatility and adjust decision-making accordingly to produce the most advantageous future outcomes. In Behrens et al.'s (2007) study, subjects carried out a one-armed bandit task in which they had to choose between blue and green stimuli. Subjects underwent trials where the probability of a blue outcome was 75% (a certain/stable environment) and trials where reward probabilities switched between 80% blue and 80% green every 30 or 40 trials (an uncertain/volatile environment). This study illustrated how human participants repeatedly combine prior and subsequent information as data accumulates, even when faced with a rapidly changing environment by continually tracking the statistics of the environment to assess the salience of every new piece of information. Behrens et al.'s fMRI data suggests the BOLD activity in the ACC might reflect a Bayesian estimate of the environment's volatility during a monitoring stage, i.e., when outcomes are being evaluated in order to regulate current beliefs about the underlying S-R contingencies of the environment. This model also suggests that the ACC might encode how much influence feedback should give to subsequent decisions, with more recent outcomes being more salient in volatile contexts (Rushworth and Behrens, 2008).

Under a Bayesian framework, unexpected observations increase uncertainty whereby a sustained level of such uncertainty results in a high estimate for volatility, which in turn leads to a high learning rate. Indeed, Behrens et al. (2007) showed that the learning rate for human participants was adjusted depending on the estimate of volatility. In situations where the S-R-O rules are changing, new information has more influence. This is because looking too far back in the history of rewarded outcomes is of little use if there has been a recent fundamental change in S-R-O contingencies. This can make prediction more difficult and thus new outcomes

have a large impact on future expectations either because they are surprising (inducing a large prediction error) or because of uncertainty about current expectations (inducing a large learning rate; Rushworth and Behrens, 2008). Indeed, learning is enhanced when outcomes occur that are not fully predicted, then slows down as outcomes become increasingly predicted and ends when outcomes are fully predicted (Hollerman and Schultz, 1998).

Other studies have emphasized the idea that learning rates are flexibly adapted to best suit environmental statistics. In fast-changing or volatile situations subjects learn quickly from new outcomes thus a faster learning rate is required (Courville et al., 2006). Indeed, Nassar et al. (2010) accurately modeled subjects' behavior with a Bayesian model finding that the model adjusts the influence of newly experienced outcomes according to on-going estimates of uncertainty and the probability of a fundamental change in the process by which outcomes are generated. Thus outcomes that are unexpected because of a fundamental change in the environment carry more influence than outcomes that are unexpected because of persistent environmental stochasticity.

Together, evidence from computational models suggests that agents can act in a Bayesian fashion in order to track S-R-O contingencies and update these accordingly. In doing so, agents can represent the level of expected uncertainty and use this to detect unexpected changes in the decision-making environment. Importantly however a distinction between unexpected uncertainty and volatility has not been explicitly addressed in this literature. Indeed, there appears to be differences in how these two forms of uncertainty are computed. For instance, during unexpected uncertainty the agent must detect and adapt to the specific change in contingency. However in volatile contexts the agent must also represent the frequency in which S-R-O contingencies are changing. This is what Behrens et al. (2007) refer to as tracking volatility as a high order statistic of the environment.

NEUROMODULATORS ASSOCIATED WITH UNCERTAINTY

Acetylcholine (ACh) and Noradrenaline (NA) may be critical neurotransmitters involved in signaling expected and unexpected sources of uncertainty (Phillips et al., 2000; Bouret and Sara, 2005; Yu and Dayan, 2005; Preuschoff et al., 2011; Avery et al., 2012). Particularly, ACh is said to signal expected uncertainty due to known unreliability in the behavioral context whereas NA is said to signal unexpected uncertainty arising from fundamental changes in the S-R-O contingencies. Evidence that ACh is crucial in expected uncertainty comes from data that ACh varies inversely with the level of estimated cue validity (Witte et al., 1997; Phillips et al., 2000; Sarter and Parikh, 2005; Yu and Dayan, 2005). This cue validity represents the probability of the cue being correct, e.g. the cue is a valid predictor of the S-R-O rule on 80% of trials. This is typically constant over a whole experimental session and thus measures the stochasticity of the task. This suggests that ACh reports a form of expected uncertainty which can be learned through past experience of S-R-O relationships. Studies suggest that ACh increases in a sustained fashion for expected unreliability of the environment when attention needs to be maintained (Dalley et al., 2001). This implies that in order to grasp the predictive relationships of an environment, an agent must utilize a temporally sustained mechanism for estimating uncertainty.

It has been suggested that NA may signal unexpected uncertainty (Bouret and Sara, 2005; Yu and Dayan, 2005; Preuschoff et al., 2011; Avery et al., 2012). There is some empirical evidence supporting this notion. For instance, the prefrontal NA system, unlike the ACh system, is engaged by *novel* S-R-O contingencies, which is compatible with a role in mechanisms of plasticity and new learning (Dalley et al., 2001). Next, available evidence suggests that NA originates in the locus coeruleus (LC) where LC neurons fire phasically (opposed to tonically) and robustly to unpredicted changes in stimulus properties or reversal of S-R-O contingencies (Aston-Jones et al., 1997; Yu and Dayan, 2003; Bouret and Sara, 2004). More recent evidence has shown that NA signals unexpected uncertainty as measured by pupil dilation (Preuschoff et al., 2011). Indeed, Preuschoff et al. (2011) have shown that unexpected uncertainty is closely linked with pupil size and is dissociated from expected uncertainty. Pupil size is thought to correlate remarkably with NA in both animal and human studies (Rajkowski et al., 1993; Gilzenrat et al., 2010). Taken together, these observations suggest that the LC-NA system facilitates attentional and cognitive shifts in behavioral adaptation in changing environments (see Sara, 2009). NA levels, could therefore signal when expectations about our world need to be revised (Cohen et al., 2007a).

Although phasic bursts of NA activity are likely to signal unexpected uncertainty, volatility characterized by a high frequency of fundamental S-R-O changes may be signaled by tonically high levels of NA (Yu, 2007). Indeed, McClure et al. (2006) propose that increased long-term response conflict (induced by frequent changes in S-R-O contingencies) biases the LC toward a tonic NA firing mode to increase exploratory behavior. These authors suggest that increased tonic firing reflects increased environmental uncertainty. This tonic mode of LC functioning may therefore reflect volatility in the environment triggered by frequent changes in the underlying rules guiding behavior.

Taken together, psychopharmacological evidence suggests that unexpected uncertainty is linked with phasic bursts of NA which signal changes in S-R-O contingencies. However expected uncertainty shows a more tonic mode of ACh in order to temporally sustain past S-R-O contingencies and hence the expected level of stochasticity. Further, volatility could be signaled by tonic levels of NA as opposed to phasic bursts (McClure et al., 2006; Yu, 2011). Therefore an important distinction could be made between unexpected uncertainty and volatility in terms of their temporal characteristics of neuromodulation.

EXPLOITATION VERSUS EXPLORATION DILEMMA

Some authors suggest that the distinction between expected and unexpected forms of uncertainty may be an important element in behavioral adaptation i.e., in choosing whether to *explore* or *exploit* the decision-making environment (Cohen et al., 2007a). The *exploitation* versus *exploration* dilemma suggests a trade-off between persisting in our current behavior (exploit) or selecting alternative options (explore) in reinforcement learning. For example, if we experience a poor quality meal at our preferred restaurant then we could choose to persist in our current behavior and continue to visit the restaurant on the assumption that the restaurant is still the best option given its good past record (exploitation). Alternatively, we may decide to explore other restaurants in search

of a better dining experience (exploration). Indeed, the *exploitation* versus *exploration* trade-off is a fundamental challenge for the adaptive control of behavior (Cohen et al., 2007a).

Particularly relevant is that uncertainty may precede the decision to explore an alternative option or exploit the current situation (Daw et al., 2006; Cohen et al., 2007a; Frank et al., 2009). For example, the detection of unexpected uncertainty can be an important signal of the need to promote exploration and has a central role in the acquisition of adaptive behavior in environments that change (Daw et al., 2006; Cohen et al., 2007a). For instance, in a familiar, reliable environment with a stable level of expected uncertainty, there is no need for exploration (i.e., the restaurant chef works 8 out of every 10 days so we are likely to gain our preferred dish on 80% of visits, thus we are just *exploiting* knowledge learned from previous experiences). If we experience a bad meal which is a consequence of a brief absence of the chef then we may continue to visit this restaurant (exploit). In contrast, during unexpected changes in the environment that lead to a durable invalidity of our previous representations, one needs to take exploratory actions (Doya, 2008). For instance, if we experience a poor meal because the previous chef was fired and replaced by a less experienced chef, this unexpected uncertainty about future visits to the restaurant might promote our exploration of other restaurants. Therefore uncertainty-driven exploration is a potentially important facet of decision-making and adaptive behavior (Cavanagh et al., 2011).

Research has begun to show that trial-to-trial variations in response-locked frontal theta are related to unexpected uncertainty and are larger in individuals who use uncertainty to guide exploration (Cavanagh et al., 2011). In addition, empirical studies have begun to reveal mechanisms that animals may use to adapt to changes in the environment, by regulating the balance between exploitation and exploration. These studies appear to be converging on the view that neuromodulatory systems; in particular, ACh and NA, interacting with DA-mediated reinforcement learning mechanisms may play a critical role in unexpected uncertainty induced exploration (Cohen et al., 2007a). Indeed, recent studies find that shifts between task engagement (exploitation) and disengagement (exploration) affect the pupil response which is thought to index NA neurotransmission (Preuschoff et al., 2011). This is consistent with Yu and Dayan (2005) theory of unexpected uncertainty and the adaptive gain theory of LC-NA (noradrenaline) mediated explore/exploit behavior (Aston-Jones and Cohen, 2005).

Together this evidence suggests a close relationship between uncertainty and the adaptive control of behavior. Indeed, it appears likely that uncertainty, and particularly unexpected uncertainty signals a contextual change which promotes exploratory adaptive behavior. Conversely, by tracking past representations of S-R-O rules and measuring the stochasticity of the environment, one can represent a form of expected uncertainty which promotes exploitative behavior. The interaction of expected and unexpected forms of uncertainty is likely to drive behavior in an optimal manner. Therefore it may be the case that successfully adapting to uncertainty could depend upon the levels of expected uncertainty and the frequency of changes in S-R-O contingencies.

To our knowledge, the distinction between volatility and unexpected uncertainty has not been explicitly articulated from the perspective of exploitation/exploration behaviors. However, it is reasonable to think that volatility should be characterized by a state in which the need for sustained exploration is anticipated. Indeed, if volatility leads to the formation of a representation that an underlying S-R-O rule can frequently change, then this should enable decision-making agents to be prepared to engage in exploration in this type of contexts. A possible prediction is that exploratory behaviors following an S-R-O rule change would be more rapidly engaged in volatile contexts compared to situations where S-R-O changes are rare because the need for exploration is anticipated. Further research will be needed to examine this question.

COGNITIVE CONTROL

As we have outlined, an emerging body of literature is beginning to demonstrate how different forms of uncertainty are processed. One aspect that has yet to be adequately addressed is the potential involvement of cognitive control processes in the resolution of uncertainty (Mushtaq et al., 2011). Indeed, the ability to rapidly and flexibly adjust behavior to changing environmental demands is a defining characteristic of *cognitive control* (Braver et al., 2003). Therefore successful adaptation to unexpected uncertainty may require the involvement of the dynamic and flexible engagement of cognitive control functions. Interestingly, different cognitive control strategies may be utilized to deal with different forms of uncertainty (i.e., expected uncertainty, unexpected uncertainty, and volatility). Particularly conflict monitoring mechanisms and working memory (WM) are two canonical instances of cognitive control processes that appear to be likely candidates for successful adaption to various forms of uncertainty (Bland and Schaefer, 2011; Mushtaq et al., 2011).

CONFLICT MONITORING AND WORKING MEMORY

The conflict hypothesis (Botvinick et al., 2001; van Veen and Carter, 2002; Kerns et al., 2004) provides a theoretical framework that can be used to understand some of the interactions between uncertainty and cognitive control. According to the conflict hypothesis, adjustments in cognitive control are likely to occur during a high degree of response conflict (Botvinick et al., 2001). According to this hypothesis, response conflict occurs whenever two or more incompatible response tendencies are simultaneously active. For example, response conflict is high when a response must be withheld in contexts in which there is a pre-potent tendency to make an overt response (Nieuwenhuis et al., 2003). Therefore, a change in learned S-R-O contingencies might require inhibiting habitual behavior (e.g., learned from the previous S-R-O rule) following a negative outcome, and overriding it with new behavior adapted to the new rule. This type of behavioral adaptation is likely to rely on *conflict processing*, that is, the ability to efficiently arbitrate between two conflicting behavioral responses (usually a habitual response that needs to be overridden by a new response). Conflict processing is thought to be a key mode of cognitive control (Botvinick et al., 2001; Yeung and Cohen, 2006), and it is more often observed in tasks with a habitual context interrupted by rare high-conflict trials (Botvinick et al., 1999). Indeed, changes in learned S-R-O contingencies and hence unexpected uncertainty are likely to produce conflict and so unexpected uncertainty may

be important in signaling the need for increased cognitive control in order to successfully adapt behavior (Mushtaq et al., 2011).

In addition to conflict monitoring mechanism, WM may also play an important role in successfully adapting to varying forms of uncertainty. WM is defined as a system providing temporary storage, manipulation and processing of information (Baddeley, 1992) and is kept on-line or available for immediate access by other cognitive processes (Awh and Jonides, 2001). WM has a key role in active maintenance and updating of information in order to allow task-relevant information to be utilized in a manner that directly biases on-going processing. This makes WM a likely candidate in decision-making whereby adaptive choices in an uncertain environment relies on tracking S-R-O contingencies and the ability to monitor and update for any changes in S-R-O associations. WM is particularly important in many tasks that require the active maintenance and updating of information in order to facilitate goal directed behavior (Owen et al., 2005). Therefore the concepts of WM and cognitive control may be closely linked with decision-making in situations where S-R-O changes might occur such as unexpected uncertainty or volatility. We will next review the link between cognitive control and different varieties of uncertainty from three perspectives: Neuroimaging studies (fMRI and ERP), models suggesting the existence of distinct modes of cognitive control (Koechlin et al., 2003; Braver et al., 2007) and neuromodulation studies.

NEURAL CORRELATES OF COGNITIVE CONTROL IN UNCERTAIN ENVIRONMENTS

Neuroimaging evidence has demonstrated greater ACC activation in studies examining conflict and conflict monitoring (Carter et al., 1998; Botvinick et al., 2001). The error-related negativity (ERN), a negative deflection in the ERP waveform at the time of an erroneous response (e.g., Gehring et al., 1990) which also originates in the ACC (Dehaene et al., 1994) is thought to be an electrophysiological marker which underlies a conflict monitoring mechanism (Carter et al., 1998; Botvinick et al., 1999, 2001; Yeung and Cohen, 2006). In addition, the anterior N2, an ERP thought to be generated in the ACC, has also been shown to reflect the monitoring of response conflict (Nieuwenhuis et al., 2003; Yeung et al., 2004). Importantly, the N2 has been associated with volatility in a habitual environment (Bland and Schaefer, 2011). Bland and Schaefer (2011) presented participants with either a blue or red triangle which was associated with two possible responses. Participants had to learn the correct S-R-O rule (red triangle – response 1 = reward; blue triangle – response 2 = reward). In this task two contextual determinants of decision uncertainty were independently manipulated: *Volatility* (i.e., the frequency of changes in the S-R-O rules) and *Feedback validity* (i.e., the extent to which an S-R-O rule accurately predicts outcomes, synonymous with expected uncertainty). Bland and Schaefer (2011) demonstrated that frequent S-R-O rule changes in an otherwise predictable environment (where *Feedback validity* is high) was associated with a frontally based N2 component. This perhaps reflects the implementation of cognitive control through a mechanism suited to detecting conflict in learned S-R-O contingencies.

In relation to the conflict hypothesis, it has been suggested that the detection of conflict by the ACC leads to the delivery of trigger signals to systems specialized in implementing control (e.g., the prefrontal cortex, PFC). Support for this idea comes from evidence suggesting that conflict-related activity in ACC predicts a subsequent increase in PFC activity and corresponding adjustments in performance (Kerns et al., 2004). Specifically, the ACC is thought to play an essential role in the adjustment of executive control mechanisms governed by the PFC (Botvinick et al., 2001; Kerns et al., 2004; Brown and Braver, 2005; Egner and Hirsch, 2005; di Pellegrino et al., 2007; Mansouri et al., 2009). Given that unexpected uncertainty and volatility are characterized by environmental changes requiring the suppression or adjustment of existing S-R-O representations, these forms of uncertainty could then be seen as states that trigger conflict and therefore the cascade of processes leading to the implementation of cognitive control processes. In other words, these forms of uncertainty can be perceived as a summary of the contextual antecedents of the implementation of cognitive control processes (Mushtaq et al., 2011).

Another theoretical interpretation proposes a link between unexpected uncertainty and specific mechanisms of cognitive control (Nieuwenhuis, 2011). An interesting review by Nieuwenhuis (2011) addresses the relationship between the LC system and the P3 ERP. By bringing together Yu and Dayan's (2005) theory and the prominent theory of the P3 proposed by Donchin (1981), Nieuwenhuis (2011) explores how unexpected uncertainty requires agents to update their representation of the environment. Indeed, a surprising and unexpected outcome must call for revision of an agent's mental model of the decision-making environment. This is indexed by the P3 amplitude which is strongly thought to be generated by the LC and NA signaling. An increased phasic release of NA may have direct enhancing effects on task-specific control representations in PFC contributing to the compensatory increase in control following a transient decrease in performance and/or reward (Aston-Jones and Cohen, 2005). Global changes in the external environment thus serves as an alarm system for contextual switches. Indeed, empirical studies are beginning to show that the variants of the P3 and late positive complex (LPC) are associated with changing S-R-O contingencies (Bland and Schaefer, 2011). Bland and Schaefer (2011) also demonstrated that frequent S-R-O rule changes in a challenging environment (where *Feedback validity* is low) was associated with a frontally based LPC component. This perhaps reflects a mechanism for integrating past outcomes in order to update a mental model of the current S-R-O contingency and the frequency in which it occurs. An S-R-O rule change is likely to signal for a revision in one's mental model which is likely to be reflected in the enhanced amplitude of the P3/LPC complex. This is in contrast to a rule change in an otherwise fairly habitual context (high *Feedback validity*) where volatility is indexed by an N2 component and likely reflects conflict monitoring (Bland and Schaefer, 2011). Therefore, there is some evidence to point to different forms of cognitive control depending on the interaction of expected uncertainty and volatility. However, unexpected uncertainty and volatility are yet to be explicitly dissociated in neuroimaging and electrophysiological studies.

SEPARABLE MODES OF COGNITIVE CONTROL

Recent studies are beginning to explore differential modes of cognitive control which may have important overlaps with the computational and neurobiological evidence outlined above. The dual modes of control (DMC) theory (Braver et al., 2007, 2009) suggests that cognitive flexibility can be achieved by modulating the manner in which a particular control mechanism is deployed in response to changing task demands or internal goal states. Specifically, this theory proposes a distinction between *proactive* and *reactive* modes of cognitive control (Braver et al., 2007). The proactive control is the early selection of goal-relevant information which is actively maintained in a sustained/anticipatory manner, before the occurrence of cognitively demanding events, to optimally bias attention, perception, and action systems in a goal-driven manner. In contrast, the reactive mode is a late correction mechanism whereby cognitive control is recruited only as needed, such as after a high-interference event is detected. Thus, proactive control relies on the anticipation and prevention of interference before it occurs, whereas reactive control relies on the *post hoc* detection and resolution of interference after its onset (Braver et al., 2009).

A clear prediction of this hypothesis is that proactive and reactive control can be distinguished in terms of lateral PFC activity. For instance, proactive control should be associated with sustained and/or anticipatory activation of PFC, which reflects the active maintenance of task goals. In contrast, reactive control should be reflected in transient activation of lateral PFC, along with a wider network of additional brain regions including the ACC (Braver et al., 2007, 2009). In addition, the DMC theory has been related to distinct ERP components. Particularly, it has been claimed that the P3 and late positivities are linked to proactive control and N2 to reactive control (van Wouwe et al., 2011). Interestingly, the P3 has been linked to WM and sustained maintenance of information in WM (Duncan-Johnson and Donchin, 1982) and the N2 has been linked with conflict monitoring and error detection (van Veen and Carter, 2002).

Importantly, proactive and reactive modes of control may be useful in successfully adapting to different forms of uncertainty. The DMC theory suggests that the temporal dynamics of neural activity can differ between a transient to a predominantly tonic mode. For instance, expected uncertainty may involve a more proactive mode of control in order to implement sustained attentional resources to facilitate internal representations of S-R-O contingencies (however, it might also be argued that automatic processes might be sufficient in a situation with learned and stable levels of expected uncertainty).

Separable modes of control have also been proposed by Koechlin and colleagues using a hierarchical framework. Koechlin et al. suggest two forms of control; *contextual and episodic control* (Koechlin et al., 2003; Koechlin and Summerfield, 2007). *Contextual control* refers to the use of a current cue (context) for selecting task appropriate behavior whereas *episodic control*, refers to the use of past cues that determine, for an extended period of time the way that current stimuli and contextual cues are interpreted (Egner, 2009). The modes of control are arranged hierarchically whereby episodic control affects contextual control, but not vice versa. According to Kouneiher et al. (2009) transient posterior-lateral PFC regions subserve contextual control whilst sustained mid-lateral PFC regions are associated with episodic control. Importantly, these two modes of control may also play a role in adapting to different forms of uncertainty. For instance, episodic control refers to temporally extended information over a behavioral episode. This requires a sustained mechanism to integrate past representations and form a mental model of the environment. This mode of control may therefore be particularly important to integrating past S-R-O occurrences and representing expected forms of uncertainty. Conversely, contextual control as indicated by transient neural activity in the PFC may be useful in detecting contextual shifts such as a change in underlying S-R-O contingencies.

Together the theories outlined above suggest that there are separable modes of cognitive control. Here, we suggest that these may be particularly relevant to estimating and resolving different forms of uncertainty. As suggested above, expected forms of uncertainty may be estimated by sustained episodic control (Kouneiher et al., 2009) or proactive control (Braver et al., 2009) whilst unexpected forms of uncertainty may be detected by transient contextual control (Kouneiher et al., 2009) or reactive control (Braver et al., 2009).

Importantly however, a reactive mode of control may not necessarily be the most optimal mode in volatile environments in which a high frequency of S-R-O changes occur. Indeed, an agent may learn that the environment is frequently changing and thus these unexpected changes may become anticipated. Therefore a proactive mode of control may be ideal in this type of environment for two reasons. First, it would allow a sustained activation of a representation of the frequency of changes in the environment and hence the potential need for constant exploratory behaviors. Second, a proactive mode of control would allow the maintenance and integration of temporally extended information about past S-R-O contingencies in order to dynamically update current mental models. A parallel could be drawn from the theory of Koechlin et al. (2003), Koechlin and Summerfield (2007), Kouneiher et al. (2009) from which it could be speculated that episodic control could also be useful in order to integrate temporally extended information needed to successfully adapt to volatile situations.

A common theme across these theories is that the separable modes of control can be distinguished by sustained and transient neural activity. This may be particularly important for estimating different forms of uncertainty. Indeed, neurotransmitters thought to underlie expected and unexpected forms of uncertainty have been distinguished by tonic and phasic activity. For instance, ACh increases in a sustained fashion for expected unreliability of the environment (Dalley et al., 2001) and is involved in a prolonged state of readiness to respond to rarely and unpredictable occurring signals (Sarter et al., 2001). It could therefore be speculated that unexpected uncertainty would be associated to transient forms of neural activity related to cognitive control, whereas volatility would be associated to more sustained patterns of neural activity in cognitive control brain networks.

ADAPTIVE GAIN THEORY OF LC-NA FUNCTIONING

The adaptive gain theory of LC-NA functioning suggests that there are at least two distinguishable modes of LC function which drive behavior. In a phasic mode, bursts of LC activity are observed

in association with the outcome of decision processes and are closely coupled with behavioral responses that are generally highly accurate. In a tonic mode however, LC baseline activity is elevated but phasic bursts of activity are absent (Aston-Jones and Cohen, 2005). Interestingly it has been proposed that the OFC and ACC could drive this LC phasic activity directly which in turn promotes exploratory or exploitative behavior (Aston-Jones and Cohen, 2005). This may have important implications for the mode in which cognitive control is implemented. For instance, unexpected uncertainty arises from strong violations of predictions that are expected to be correct (Yu and Dayan, 2005). Phasic NA signals have been associated with novelty and changes in S-R-O contingencies (Aston-Jones et al., 1997; Aston-Jones and Cohen, 2005; Yu and Dayan, 2005; Avery et al., 2012). This would fit well with a reactive mode of control which arises as a consequence of high-conflict events (Braver et al., 2007) which could be cause by strong violations of predictions. Furthermore this is also linked with a view that unexpected uncertainty is induced by a mismatch between prediction and observation and is signaled phasically with rapid habituation (Yu and Dayan, 2005). Indeed, strong projections from the OFC and ACC to the LC may drive this phasic response where signals from OFC and ACC augment the LC phasic release of NA thus improving performance on subsequent trials (Aston-Jones et al., 2002; Aston-Jones and Cohen, 2005). According to the adaptive gain theory, this effect could further contribute to the compensatory increase in control following a transient decrease in performance and/or reward. Indeed, empirical evidence suggests that NA is specifically involved in performance monitoring (Riba et al., 2005). Furthermore, there is substantial evidence for the modulatory influence of NA on cognitive functions that depend on the frontal cortex, particularly selective attention and working-memory tasks (Sara, 2009).

The adaptive gain theory further suggests that signals from ACC to LC (indicating an adverse outcome), possibly complemented by signals from OFC to LC (indicating absence of an expected reward) may augment the LC phasic mode (Aston-Jones and Cohen, 2005). This, in turn, would improve performance on subsequent trials by enhancing the LC phasic release of NA thus having direct enhancing effects on task-specific control representations in PFC (Aston-Jones and Cohen, 2005). Thus conflict detection as reflected by the ACC response which then sends triggers for compensatory adjustments in cognitive control may be mediated by LC-NA functioning. This would be consistent with Yu and Dayan's (2005) theoretical framework of NA functioning as a signal for unexpected uncertainty.

Indeed unexpected uncertainty can be seen as a state signaling the potential need to suppress of previous S-R-O rules in order to override these with more adaptive S-R-O contingencies. This requires flexible adaption of behavior in environments that are changeable. Thus signaling of NA in response to unexpected uncertainty may be crucially involved in ACC-PFC implementation of cognitive control. Indeed, functional neuroimaging studies investigating uncertainty have uncovered a neural network that has a remarkable overlap with brain networks usually associated with cognitive control tasks. In particular, a network involving lateral PFC areas, parietal cortex and the ACC seems to be constantly activated for decision-making tasks in which volatility and

expected forms of uncertainty are manipulated and also in a wide range of classical cognitive control tasks (for a review of the neural correlates of uncertainty and cognitive control see Mushtaq et al., 2011). Therefore cognitive control and particularly a reactive mode as indexed by early negativities in the EEG and ACC fluctuations in the BOLD response as well as phasic bursts of NA may be particularly important for estimating, detecting, and resolving unexpected uncertainty. Alternatively, a proactive control mode characterized by sustained neural activity in the PFC and the P3/LPC complex may be important for successful integration of past outcomes in order to measures the stochasticity of the environment and deal with expected uncertainty. However, it is also likely that stable levels of stochasticity could be learned through automatic processes without the involvement of cognitive control processes. In addition, it is possible that proactive control might be also particularly useful in volatile contexts, where the temporally sustained maintenance and updating of past outcome information in WM might be useful to adapt to a context of frequent S-R-O changes.

In summary, it seems that reactive control could be used following a highly unexpected S-R-O change. However, a proactive mode can be very efficient at dealing with volatility. Therefore unexpected uncertainty and volatility should be differentiated: unexpected uncertainty occurs from a single or infrequent unpredicted fundamental changes in S-R-O contingency whereas volatility can be seen as a series of frequent fundamental changes in S-R-O frequencies, and this frequency of changes can itself become predictable. For our example above, a volatile situation is reached when our usual restaurant tends to hire a new chef very often during the year. If customers know this tendency, they will be able to use proactive strategies in order to detect if a change in the quality of the food is due to a transient change in a more stable pattern (e.g., the usual chef is absent 1 day every week) or if it reflects a more fundamental change, i.e., the previous chef was fired and replaced by a new one). Therefore how the brain estimates the relative frequency of changes on the environment is crucial. Behrens et al. (2007) suggest that this is reflected by ACC activity. Indeed, the ACC may be able to estimate the rate at which reward contingencies are changing and signal to the PFC to implement a reactive or more proactive mode of control. This likely reflects a highly sophisticated control mechanism which adjusts for suitable changes in the environment as possibly reflected by neuromodulation of ACh and NA mediated by the ACC-PFC.

SYNTHESIS AND CONCLUSIONS

We have reviewed existing empirical evidence and theoretical evidence in order to form a case for considering three distinct forms of uncertainty; expected uncertainty, unexpected uncertainty, and volatility. Whilst expected uncertainty has received much attention in the literature, the latter two forms of uncertainty are relatively less well explored. Nevertheless a growing body of literature is beginning to unravel how the brain deals with unexpected changes in the environment. This is an exciting line of research which is beginning to prove fruitful (Yu and Dayan, 2005; Behrens et al., 2007; Doya, 2008; Krugel et al., 2009; Nassar et al., 2010; Bland and Schaefer, 2011; Nieuwenhuis, 2011; Preuschoff et al., 2011).

However an explicit distinction between unexpected uncertainty and volatility has yet to be addressed. We have suggested

that computational modeling studies provide evidence of how we can deal with unexpected changes in S-R-O contingences and adjust the learning rate accordingly. However, volatility appears to promote a further computation by representing a "volatility" parameter as a high order statistic of the environment (Behrens et al., 2007). Next, the temporal activity of neuromodulators involved in signaling uncertainty may differentiate unexpected uncertainty and volatility. Particularly, unexpected uncertainty appears to be signaled by phasic bursts of NA activity whereas prolonged unexpected uncertainty i.e., volatility may recruit a more tonic mode. Finally these two forms of uncertainty may be differentiated in terms of the involvement of distinct cognitive control modes. It is possible that unexpected changes may be dealt with by a reactive mode of control recruiting conflict detection mechanisms to overcome competing responses in S-R-O contingencies. Alternatively successful adaptation to volatility may be associated with a proactive and sustained mode of control through the continual maintenance and updating of S-R-O contingencies in WM.

In addition, a number of questions remain open. For instance, it is unclear at this stage whether volatility and unexpected uncertainty are associated with distinct brain networks. The evidence reviewed above about the potential involvement of distinct cognitive processes in these two forms of uncertainty suggests that they could be dissociated in terms of their neural correlates. Further research will be necessary to address this question. A more fundamental question regards the nature of the distinction between volatility and unexpected uncertainty. The main difference between them is the *frequency* of S-R-O changes in a given period of time. This frequency can be manipulated in a gradual, continuous way. However, it can be speculated that systems involved in processing uncertainty should be able to detect a threshold beyond which the processes implemented to deal with the environment will change (e.g., switching from a reactive toward a proactive mode of control). Further research will be needed to test this idea. Finally, although the theoretical avenues considered in this article suggest that volatility and unexpected uncertainty might lead to different modes of cognitive control, and to different neuromodulatory patterns, most of these ideas remain yet to be empirically tested.

In summary, this article has reviewed empirical and theoretical evidence for the distinction between three forms of uncertainty, and in particular, it highlighted a distinction between a rare unexpected change (unexpected uncertainty) and a frequently changing environment (volatility). Future research should therefore form a clear distinction between unexpected uncertainty and volatility in order to further explore how we successfully estimate, represent, and resolve these different forms of uncertainty.

ACKNOWLEDGMENTS

REFERENCES

Aston-Jones, G., and Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* 28, 403–450.

Aston-Jones, G., Rajkowski, J., and Kubiak, P. (1997). Conditioned responses of monkey locus coeruleus neurons anticipate acquisition of discriminative behavior in a vigilance task. *Neuroscience* 80, 697–715.

Aston-Jones, G., Rajkowski, J., Lu, W., Zhu, Y., Cohen, J. D., and Morecraft, R. J. (2002). Prominent projections from the orbital prefrontal cortex to the locus coeruleus in monkey. *Soc. Neurosci. Abstr.* 28, 86–89.

Avery, M. C., Nitz, D. A., Chiba, A. A., and Krichmar, J. L. (2012). Simulation of cholinergic and noradrenergic modulation of behavior in uncertain environments. *Front. Comput. Neurosci.* 6:5. doi:10.3389/fncom.2012.0000

Awh, E., and Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends Cogn. Sci. (Regul. Ed.)* 5, 119–126.

Baddeley, A. (1992). Working memory. *Science* 255, 556–559.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221.

Bland, A. R., and Schaefer, A. (2011). Electrophysiological correlates of decision making under varying levels of uncertainty. *Brain Res.* 1417, 55–66.

Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., and Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature* 402, 179–181.

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

Bouret, S., and Sara, S. J. (2004). Reward expectation, orientation of attention and locus coeruleus-medial frontal cortex interplay during learning. *Eur. J. Neurosci.* 20, 791–802.

Bouret, S., and Sara, S. J. (2005). Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends Neurosci.* 28, 574–582.

Braver, T. S., Gray, J. R., and Burgess, G. C. (2007). "Explaining the many varieties of working memory variation: dual mechanisms of cognitive control," in *Variation in Working Memory*, eds C. R. A. Andrew, J. Christopher, J. M. Kane, M. Akira, and T. N. John (New York: Oxford University Press), 76–106.

Braver, T. S., Paxton, J. L., Locke, H. S., and Barch, D. M. (2009). Flexible neural mechanisms of cognitive control within human prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* 106, 7351–7356.

Braver, T. S., Reynolds, J. R., and Donaldson, D. I. (2003). Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron* 39, 713–726.

Brown, J. W., and Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science* 307, 1118–1121.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., and Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science* 280, 747–749.

Cavanagh, J. F., Figueroa, C. M., Cohen, M. X., and Frank, M. J. (2011). Frontal theta reflects uncertainty and unexpectedness during exploration and exploitation. *Cereb. Cortex.* doi: 10.1093/cercor/bhr332

Cohen, J. D., McClure, S. M., and Yu, A. J. (2007a). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 362, 933–942.

Cohen, M. X., Elger, C. E., and Ranganath, C. (2007b). Reward expectation modulates feedback-related negativity and EEG spectra. *Neuroimage* 35, 968–978.

Courville, A. C., Daw, N. D., and Touretzky, D. S. (2006). Bayesian theories of conditioning in a changing world. *Trends Cogn. Sci. (Regul. Ed.)* 10, 294–300.

Critchley, H. D., Mathias, C. J., and Dolan, R. J. (2001). Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron* 29, 537–545.

Dalley, J. W., McGaughy, J., O'Connell, M. T., Cardinal, R. N., Levita, L., and Robbins, T. W. (2001). Distinct changes in cortical acetylcholine and noradrenaline efflux during contingent and noncontingent performance of a visual attentional task. *J. Neurosci.* 21, 4908–4914.

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

Dayan, P., and Yu, A. (2002). "ACh, uncertainty, and cortical inference," in *Advances in Neural Information Processing Systems 14,* Vols. 1 and 2, eds T. G. Dietterich, S. Becker, and Z. Ghahramani (MIT Press), 189–196.

Dehaene, S., Posner, M. I., and Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychol. Sci.* 5, 303–305.

di Pellegrino, G., Ciaramelli, E., and Làdavas, E. (2007). The regulation of cognitive control following rostral anterior cingulate cortex lesion in humans. *J. Cogn. Neurosci.* 19, 275–286.

Donchin, E. (1981). Surprise!... Surprise? *Psychophysiology* 18, 493–513.

Donchin, E., and Coles, M. G. (1988). Is the P300 component a manifestation of context updating? *Behav. Brain. Sci.* 11, 357–427.

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

Duncan-Johnson, C. C., and Donchin, E. (1977). On quantifying surprise: the variation of event-related potentials with subjective probability. *Psychophysiology* 14, 456–467.

Duncan-Johnson, C. C., and Donchin, E. (1982). The P300 component of the event-related brain potential as an index of information processing. *Biol. Psychol.* 14, 1–52.

Egner, T. (2009). Prefrontal cortex and cognitive control: motivating functional hierarchies. *Nat. Neurosci.* 12, 821–822.

Egner, T., and Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nat. Neurosci.* 8, 1784–1790.

Elliott, R., and Dolan, R. J. (1998). Activation of different anterior cingulate foci in association with hypothesis testing and response selection. *Neuroimage* 8, 17–29.

Fiorillo, C. D., Tobler, P. N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299, 1898–1902.

Frank, M. J., Doll, B. B., Oas-Terpstra, J., and Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nat. Neurosci.* 12, 1062–1068.

Gehring, W. J., Coles, M. G. H., Meyer, D. E., and Donchin, E. (1990). The error-related brain potential: an event-related brain potential accompanying errors. *Psychophysiology* 27, s34.

Gilzenrat, M., Nieuwenhuis, S., Jepma, M., and Cohen, J. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cogn. Affect. Behav. Neurosci.* 10, 252–269.

Goel, V., and Dolan, R. J. (2000). Anatomical segregation of component processes in an inductive inference task. *J. Cogn. Neurosci.* 12, 110–119.

Hollerman, J. R., and Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat. Neurosci.* 1, 304–309.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., and Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science* 310, 1680–1683.

Huettel, S., Stowe, J., Gordon, E., Warner, B., and Platt, M. (2006). Neural signatures of economic preferences for risk and ambiguity. *Neuron* 49, 765–775.

Huettel, S. A., Song, A. W., and McCarthy, G. (2005). Decisions under uncertainty: probabilistic context influences activation of prefrontal and parietal cortices. *J. Neurosci.* 25, 3304–3311.

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

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

Knutson, B., Taylor, J., Kaufman, M., Peterson, R., and Glover, G. (2005). Distributed neural representation of expected value. *J. Neurosci.* 25, 4806–4812.

Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science* 302, 1181–1185.

Koechlin, E., and Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends Cogn. Sci. (Regul. Ed.)* 11, 229–235.

Kouneiher, F., Charron, S., and Koechlin, E. (2009). Motivation and cognitive control in the human prefrontal cortex. *Nat. Neurosci.* 12, 939–945.

Krain, A. L., Wilson, A. M., Arbuckle, R., Castellanos, F. X., and Milham, M. P. (2006). Distinct neural mechanisms of risk and ambiguity: a meta-analysis of decision-making. *Neuroimage* 32, 477–484.

Krugel, L. K., Biele, G., Mohr, P. N. C., Li, S.-C., and Heekeren, H. R. (2009). Genetic variation in dopaminergic neuromodulation influences the ability to rapidly and flexibly adapt decisions. *Proc. Natl. Acad. Sci. U.S.A.* 106, 17951–17956.

Mansouri, F. A., Tanaka, K., and Buckley, M. J. (2009). Conflict-induced behavioural adjustment: a clue to the executive functions of the prefrontal cortex. *Nat. Rev. Neurosci.* 10, 141–152.

Mathys, C., Daunizeau, J., Friston, K. J., and Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* 5:39. doi:10.3389/fnhum.2011.00039

McClure, S. M., Gilzenrat, M. S., and Cohen, J. D. (2006). "An exploration – exploitation model based on norepinephrine and dopamine activity," in *Advances in Neural Information Processing Systems*, Vol. 18, eds Y. Weiss, B. Schölkopf, and J. Platt (Cambridge: MIT Press), 867–874.

Mushtaq, F., Bland, A. R., and Schaefer, A. (2011). Uncertainty and cognitive control. *Front. Psychol.* 2:249. doi:10.3389/fpsyg.2011.00249

Nassar, M. R., Wilson, R. C., Heasly, B., and Gold, J. I. (2010). An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* 30, 12366–12378.

Nieuwenhuis, S., Yeung, N., van den Wildenberg, W., and Ridderinkhof, K. (2003). Electrophysiological correlates of anterior cingulate function in a go/no-go task: effects of response conflict and trial type frequency. *Cogn. Affect. Behav. Neurosci.* 3, 17–26.

Nieuwenhuis, S. (2011). "Learning, the P3, and the locus coeruleus-norepinephrine system," in *Neural Basis of Motivational and Cognitive Control*, eds R. B. Mars, J. Sallet, M. F. S. Rushworth, and N. Yeung (Oxford University Press), 209–222.

Owen, A. M., McMillan, K. M., Laird, A. R., and Bullmore, E. (2005). N-back working memory paradigm: a meta-analysis of normative functional neuroimaging studies. *Hum. Brain Mapp.* 25, 46–59.

Paulus, M. P., Feinstein, J. S., Simmons, A., and Stein, M. B. (2004). Anterior cingulate activation in high trait anxious subjects is related to altered error processing during decision making. *Biol. Psychiatry* 55, 1179–1187.

Paulus, M. P., Hozack, N., Frank, L., and Brown, G. G. (2002). Error rate and outcome predictability affect neural activation in prefrontal cortex and anterior cingulate during decision-making. *Neuroimage* 15, 836–846.

Payzan-LeNestour, E., and Bossaerts, P. (2011). Risk, Unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Comput. Biol.* 7, e1001048. doi:10.1371/journal.pcbi.1001048

Phillips, J. M., McAlonan, K., Robb, W. G. K., and Brown, V. J. (2000). Cholinergic neurotransmission influences covert orientation of visuospatial attention in the rat. *Psychopharmacology* 150, 112–116.

Polezzi, D., Lotto, L., Daum, I., Sartori, G., and Rumiati, R. (2008). Predicting outcomes of decisions in the brain. *Behav. Brain Res.* 187, 116–122.

Polich, J. (1990). Probability and inter-stimulus interval effects on the P300 from auditory stimuli. *Int. J. Psychophysiol.* 10, 163–170.

Preuschoff, K., Bossaerts, P., and Quartz, S. R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51, 381–390.

Preuschoff, K., Hart, B. M., and Einhauser, W. (2011). Pupil dilation signals surprise: evidence for noradrenaline's role in decision making. *Front. Neurosci.* 5:115. doi:10.3389/fnins.2011.00115

Rajkowski, J., Kubiak, P., and Aston-Jones, G. (1993). Correlations between locus coeruleus (LC) neural activity, pupil diameter and behavior in monkey support a role of LC in attention. *Abstr. Soc. Neurosci.* 19, 974.

Riba, J., Rodríguez-Fornells, A., Morte, A., Münte, T. F., and Barbanoj, M. J. (2005). Noradrenergic stimulation enhances human action monitoring. *J. Neurosci.* 25, 4370–4374.

Ridderinkhof, K. R., van den Wildenberg, W. P. M., Segalowitz, S. J., and Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain Cogn.* 56, 129–140.

Rushworth, M. F., and Behrens, T. E. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat. Neurosci.* 11, 389–397.

Sara, S. J. (2009). The locus coeruleus and noradrenergic modulation of cognition. *Nat. Rev. Neurosci.* 10, 211–223.

Sarter, M., Givens, B., and Bruno, J. P. (2001). The cognitive neuroscience of sustained attention: where top-down meets bottom-up. *Brain Res. Rev.* 35, 146–160.

Sarter, M., and Parikh, V. (2005). Choline transporters, cholinergic transmission and cognition. *Nat. Rev. Neurosci.* 6, 48–56.

Seymour, B., Daw, N., Dayan, P., Singer, T., and Dolan, R. (2007). Differential encoding of losses and gains in the human striatum. *J. Neurosci.* 27, 4826–4831.

Stern, E. R., Gonzalez, R., Welsh, R. C., and Taylor, S. F. (2010). Updating beliefs for a decision: neural correlates of uncertainty and underconfidence. *J. Neurosci.* 30, 8032–8041.

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

Tobler, P. N., O'Doherty, J. P., Dolan, R. J., and Schultz, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *J. Neurophysiol.* 97, 1621–1632.

van Veen, V., and Carter, C. S. (2002). The anterior cingulate as a conflict monitor: fMRI and ERP studies. *Physiol. Behav.* 77, 477–482.

van Wouwe, N. C., Band, G. P. H., and Ridderinkhof, K. R. (2011). Positive affect modulates flexibility and evaluative control. *J. Cogn. Neurosci.* 23, 524–539.

Volz, K. G., Schubotz, R. I., and von Cramon, D. Y. (2003). Predicting events of varying probability: uncertainty investigated by fMRI. *Neuroimage* 19, 271–280.

Witte, E. A., Davidson, M. C., and Marrocco, R. T. (1997). Effects of altering brain cholinergic activity on covert orienting of attention: comparison of monkey and human performance. *Psychopharmacology (Berl.)* 132, 324–334.

Yeung, N., Botvinick, M., and Cohen, J. D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol. Rev.* 111, 931–959.

Yeung, N., and Cohen, J. D. (2006). The impact of cognitive deficits on conflict monitoring. *Psychol. Sci.* 17, 164–171.

Yu, A. J. (2007). Adaptive behavior: humans act as bayesian learners. *Curr. Biol.* 17, R977–R980.

Yu, A. J., and Dayan, P. (2003). Expected and unexpected uncertainty: ACh and NE in the neocortex. *Adv. Neural Inf. Proc. Syst.* 15, 157–164.

Yu, A. J., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692.

Yu, A. J. (ed). (2011). "Uncertainty and Neuromodulation," in *Neuroscience of Decision Making.* (New York: Psychology Press).

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.