



Published in final edited form as:

Top Cogn Sci. 2010 October ; 2(4): 658–677. doi:10.1111/j.1756-8765.2010.01085.x.

Computational models of performance monitoring and cognitive control

William H. Alexander and Joshua W. Brown

Dept. of Psychological and Brain Sciences, Indiana University, Bloomington IN

Abstract

The medial prefrontal cortex (mPFC) has been the subject of intense interest as a locus of cognitive control. Several computational models have been proposed to account for a range of effects including error detection, conflict monitoring, error likelihood prediction, and numerous other effects observed with single-unit neurophysiology, fMRI, and lesion studies. Here we review the state of computational models of cognitive control and offer a new theoretical synthesis of the mPFC as signaling response-outcome predictions. This new synthesis has two interacting components. The first component learns to predict the various possible outcomes of a planned action, and the second component detects discrepancies between the actual and intended responses; the detected discrepancies in turn update the outcome predictions. This single construct is consistent with a wide array of performance monitoring effects in mPFC and suggests a unifying account of the cognitive role of medial PFC in performance monitoring.

1. Introduction

Models of cognitive or executive control came of age with a seminal qualitative framework proposed by Norman and Shallice (1986). In this framework, “schemas” mapped stimuli to corresponding responses. When two schemas were activated simultaneously and conflicted with each other or otherwise needed to be coordinated, a “contention scheduling” mechanism ensured that only one schema could carry out its stimulus-response mapping at a time. Sometimes, however, a schema might not only need to be coordinated with another stimulus-driven schema, but also coordinated with larger behavioral goals. According to Norman and Shallice, this required a “supervisory attentional system” to impose top-down goals that might deactivate some schemas and activate others in service of a higher level goal. The organization of component behaviors into coordinated, goal-directed actions is arguably the essence of cognitive control. Computational theories of cognitive control have generally focused on two aspects of Norman and Shallice’s supervisory attentional system, namely attentional biasing and performance monitoring. We treat each of these in turn with a view toward a novel synthesis of computational theories of cognitive control.

2. Attentional biasing and task switching

Attentional biasing (Desimone & Duncan, 1995; Norman & Shallice, 1986; Posner & DiGirolamo, 1998) refers to the top-down modulation of attention toward cues that will drive behavior consistent with higher-level goals. This concept of attentional biasing forms the basis of the *biased competition model* (Cohen, Dunbar, & McClelland, 1990; Miller & Cohen, 2001), as shown in Figure 1, in which the usefulness of attentional biasing is exemplified with the Stroop task (Stroop, 1935). In the Stroop task, color words such as

“green” or “red” are presented to a subject in various ink colors, and subjects must ignore the meaning of the word while naming the ink color. A top-down attentional bias signal focuses attention on the ink color, thus enabling subjects to name the ink color instead of (incorrectly) reading the word. *This top-down biasing that directs performance based on a current task set or goal is an essential component of cognitive control.* Several computational models have been developed from the biased competition model, and these have been used to argue that neural network models can learn rule categories and implement higher level goals, arguably without requiring symbolic goal representations (Herd, Banich, & O’Reilly, 2006; Rey, Lew, & Zanutto, 2007; Rougier, Noelle, Braver, Cohen, & O’Reilly, 2005).

2.1. Task switching and cognitive control

The biased competition model has been especially studied in the case of two or more tasks that must be performed in alternation, as the task rules switch. In the Stroop task (Stroop, 1935) for example, a cue preceding each trial can specify whether the color naming or word reading task must be performed. Response times are generally longer following a task switch as compared to repeated performance of the same task, an effect known as the *switch cost* (Jersild, 1927). A renaissance of work on task switching beginning in the early 1990’s has generally proposed two different computational accounts of the source of the switch cost: one as due to the time needed for an executive controller to reconfigure the task set (Rogers & Monsell, 1995), and the other as due to the need to overcome priming of the task that is no longer relevant after a switch (Allport, Styles, & Hsieh, 1994). Computational models have largely simulated task switch costs as due to priming effects (Altmann & Gray, 2008; Badre & Wagner, 2006; Brown, Reynolds, & Braver, 2007; Gilbert & Shallice, 2002), although one model suggests that switch costs may be due to a combination of priming and executive control effects (Brown et al., 2007). Overall, the extent to which switch costs reflect a top-down executive controller versus a bottom-up priming effect remains in dispute (Altmann, 2003).

2.2. Priming and cognitive control

The controversy over task-switch costs has often cast priming effects as antithetical to postulated cognitive control mechanisms, but such is not necessarily the case. Theoretically, priming effects are not simply a liability that leads to slower performance on task switch trials. Priming may improve performance by increasing the connection between task cues and the current task rule. This in turn allows task cues to strengthen the representation of the current task set. In other words, priming effects may directly activate cognitive control mechanisms rather than replace them. In terms of reaction time, this may speed up repeated task performance. Priming may also reduce errors by allowing the task rule to be reactivated more easily if it decays from working memory in the framework of the biased competition model. In support of this possibility, the *failure-to-engage* hypothesis (de Jong, 2000) suggests that task representations in working memory do occasionally decay, in which case the task stimuli initiate a retrieval of the most recent task cue into working memory, so that the task can be performed correctly. A computational model of the failure-to-engage hypothesis (Reynolds, Braver, Brown, & Stigchel, 2006) has predicted that switch costs may be increased when task set activity decays prior to target appearance (i.e., a failure to engage), but priming allows the task set to be retrieved so that a correct response can be generated, albeit more slowly. Furthermore, the probability of failure to engage may be reduced with increased motivation, as manipulated by performance incentives (Nieuwenhuis & Monsell, 2002).

2.3. Dopaminergic tone and cognitive control

Motivational effects on cognitive control, and especially the stability of task set representations, may be mediated in part by dopamine. At the neural level, dopamine plays a critical role in stabilizing working memory representations, which include task set (Muly, Szigeti, & Goldman-Rakic, 1998), so that they are less likely to decay. Several attempts have been made to simulate this effect computationally, using rate-coded and spiking computational neural models (Braver & Cohen, 2000; Brunel & Wang, 2001; Durstewitz, Kelc, & Gunturkun, 1999; Durstewitz, Seamans, & Sejnowski, 2000). One of these models (Brunel & Wang, 2001) has been extended to simulate more directly the role of dopamine in task set engagement (Deco & Rolls, 2003, 2005). A basic property of these computational models is that there is an optimal level of dopamine in prefrontal regions associated with task set maintenance, and too little or too much dopamine concentration impairs performance (Chadderdon & Sporns, 2006; Muly et al., 1998).

3. Performance monitoring

A fundamental question is how the brain knows when to implement cognitive control in the first place, versus when to simply allow control by prepotent stimulus-response associations (“schemas” according to Norman and Shallice (1986)). What is needed is a mechanism to monitor ongoing performance and determine when additional control signals are required, then in turn increase the activity of cognitive control signals such as task set representations. Some models have been construed to argue that executive control mechanisms are not fundamentally different from schemas, in that both implement control whenever a set of underlying conditions is satisfied (e.g., Meyer & Kieras, 1997). In other words, cognitive control might be viewed as arbitrarily complex stimulus-response associations. In contrast, other computational neural models of cognitive control incorporate relatively simple stimulus-response associations, and behavioral complexity is introduced through performance monitoring signals which modulate the expression of goal-directed behavior (Botvinick et al., 2001; Brown et al., 2007; Jones, Cho, Nystrom, Cohen, & Braver, 2002). In practice, these perspectives may be reconciled in that goals must be activated in appropriate circumstances, but their level of activity may be further increased as difficult conditions warrant. Several distinct computational mechanisms of control have been proposed as satisfying this performance monitoring requirement, and these can be broken down into *reactive vs. proactive monitoring* (Braver, Gray, & Burgess, 2007).

3.1. Reactive monitoring

Reactive monitoring detects a problem *ex post facto*, and this principally involves error detection, although rewards may be monitored as well (Ito, Stuphorn, Brown, & Schall, 2003). Empirically, error commission is commonly followed by a lower error rate and longer reaction times in subsequent trials (Laming, 1979; Rabbitt, 1966). Thus, error signals can provide an indication that a greater level of cognitive control is needed. Error signals have been found in single units of the medial prefrontal cortex (mPFC) of monkeys (Gemba, Sasaki, & Brooks, 1986; Niki & Watanabe, 1979). Similar signals were later found with event-related potential methods in the medial prefrontal cortex of humans. These error effects were referred to as the error-related negativity (ERN) (Falkenstein, Hohnsbein, Hoorman, & Blanke, 1991; Gehring, Coles, Meyer, & Donchin, 1990).

3.1.1 Reactive error detection—Several models have proposed how errors are detected and lead to mPFC activation (Table 1). First, error effects may be due to a comparison of actual vs. desired outcomes (Ito et al., 2003; Scheffers & Coles, 2000). This computation may be carried out in midbrain dopamine cells (Brown, Bullock, & Grossberg, 1999; Schultz, Dayan, & Montague, 1997), which pause their firing in response to an error.

This has been proposed to cause subsequent activation of medial PFC (Holroyd & Coles, 2002).

Alternatively, an error may be detected as a second response that follows an immediately preceding and presumably incorrect first response (Steinhauser et al., 2008). In this framework, a response is generated when accumulated evidence for the response passes some threshold, and this accumulation of evidence continues even after the response has been given. In error trials, continued accumulation of evidence may lead to a second response. An error is inferred if a discrepancy exists between the first and second response.

3.2 Proactive performance monitoring

It would be very useful for an individual to learn to predict the need for increased cognitive control at the earliest possible time, so that cognitive control can be implemented proactively to meet the increased demand to prevent an error in the first place. Several studies have suggested that mPFC, and especially the anterior cingulate cortex (ACC), serves not only to detect errors, but also to anticipate the need for control (Aarts, Roelofs, & van Turenout, 2008; De Pisapia & Braver, 2006; Fan et al., 2007; Sohn, Albert, Jung, Carter, & Anderson, 2007).

3.2.1. Conflict monitoring—As the importance of error monitoring by mPFC was gaining new appreciation, a series of high-profile papers (Botvinick, Nystrom, Fissel, Carter, & Cohen, 1999; Carter et al., 1998; MacDonald, Cohen, Stenger, & Carter, 2000) proposed that mPFC error effects may reflect a signal of *response conflict* between correct vs. incorrect response processes. In its simplest form, the conflict model specifies the conflict signal as the *multiplicative product of two mutually incompatible response processes* (Botvinick et al., 2001). When an incorrect response representation is more active alongside the correct response representation, then a state of response conflict exists by definition. To the extent that mPFC signals the presence of response conflict, mPFC activity will be greater even if no error actually occurs on that trial. The response conflict signal, as shown in Figure 1, was proposed to drive proactive increases in task set-related activity in the dorsolateral prefrontal cortex (DLPFC) leading to greater cognitive control before an error might occur. This theory was then formalized in several computational models of performance monitoring and cognitive control (Botvinick et al., 2001; Jones et al., 2002).

3.2.2. Error likelihood prediction—Brown & Braver proposed that apparent response conflict effects due to incongruent response cues may actually reflect an increase in the perceived likelihood of an error (Brown & Braver, 2005). This proposal was supported by combined computational modeling and fMRI. In the error likelihood computational model, mPFC cells in a self-organizing map (Kohonen, 1982) are recruited and trained by error signals to respond to stimuli that preceded an error. When similar circumstances arise later, the response cues associated with higher error likelihood drive activation of a greater number of model mPFC cells, even when no response conflict exists and no error is committed. This computational model prediction was supported by fMRI findings (Brown & Braver, 2005). The report generated some controversy, as a subsequent study did not find error likelihood effects as expected (Nieuwenhuis, Schweizer, Mars, Botvinick, & Hajcak, 2007). One of the issues raised by Nieuwenhuis et al. was the distinction between error likelihood effects driven by the predictive information in the cue, vs. the difficulty of the task at the time of response. Nonetheless, further studies have replicated the error likelihood effect and suggest an explanation for its absence in some individuals (Brown & Braver, 2007), namely that error likelihood effects drive risk avoidance (Magno, Foxe, Molholm, Robertson, & Garavan, 2006; Paulus & Frank, 2006). Thus, error likelihood effects may be strongest in more risk-averse individuals but relatively absent in risk-seeking individuals.

Further computational modeling work with the error likelihood model showed that the magnitude of the error likelihood effect can be simulated by manipulating how efficiently individuals learn from past errors (Brown & Braver, 2008), which may account for individual differences in error and reward sensitivity and ERN magnitude (Frank, Worocho, & Curran, 2005; Hewig, Hagemann, Seifert, Naumann, & Bartussek, 2006; Klein et al., 2007).

Incongruent response cues (Eriksen & Eriksen, 1974; Stroop, 1935) are often associated with higher error rates. Therefore error likelihood effects and the incongruity effects often attributed to response conflict computation (Botvinick et al., 2001) in mPFC might be expected to show a positive correlation across subjects. To the contrary, the error likelihood computational model surprisingly predicted an inverse relationship between incongruity and error likelihood effects, which was subsequently supported by fMRI findings (Brown & Braver, 2008). On the surface, this seems to be at odds with the hypothesis that incongruity effects reflect an underlying prediction of error likelihood. In the model, incongruent response cues indicate that in addition to a correct response, an error is also likely. In the case of an incongruent stimulus, there are two competing cues that drive two mutually incompatible response processes. Initially, before error likelihood associations are learned, this leads to the broad mPFC activation that is measured as the incongruity effect. These incongruity effects resulting from diffuse excitation of the model mPFC become smaller as error likelihood learning proceeds and sharpens the error likelihood representations (Brown & Braver, 2008). Thus, as error likelihood representations are learned, the more widespread activation associated with incongruity effects diminishes. This accounts for the observed inverse relationship between error likelihood and incongruity effects.

A striking further prediction of the error likelihood model is that incongruity effects may be found when a greater number of responses are cued, *even when those responses do not conflict with each other*. Incongruent stimuli cue two possible responses, but congruent stimuli cue only a single response. Typically, a simultaneous response to both the relevant and distractor components of the incongruent cue constitutes an error, but what if instead the task were manipulated so that a simultaneous response to both relevant and distractor components were required? In that case, an incongruent stimulus would elicit multiple responses but not entail response conflict. According to the error likelihood model, an incongruent response cue, but not an incongruent stimulus cue (van Veen, Cohen, Botvinick, Stenger, & Carter, 2001), should lead to greater mPFC activation regardless of whether or not the response to the distractor is required. This is due to the fact that as the number of response cues increases, the number of activated response and outcome representations in mPFC increases, which leads to greater overall activation of mPFC. Thus, the error likelihood model predicts that incongruity effects do not depend on the presence of actual response conflict. This hypothesis has now been tested and is consistent with fMRI findings (Brown, 2009).

3.3 Reactive-Proactive performance monitoring

Although we distinguish between reactive and proactive control, it is likely that both kinds of signals are used in performance monitoring. The conflict model for example provides reactive as well as proactive control signals (Yeung et al., 2004). Specifically, the possibility of error is inferred from the simultaneous activation of two potential responses (proactive), while actual errors are detected by continued co-activation of response plans following an error (reactive). Stimuli that cue conflicting responses initiate a competition between response units, one representing the correct response and the other representing an error. According to the conflict theory, ACC reports the amount of conflict as the product of the activation of all response units. On correct trials, the unit representing the correct response is

predominantly activated; mutual inhibition between response units ensures that conflict remains low. On incorrect trials, random noise in the model pushes activity in the incorrect response unit above a response threshold, resulting in an error. However, because the error was due to processing noise, attentionally-biased processing of stimuli ensures that the correct unit continues to be active, resulting in a high degree of conflict. This high conflict situation occurs after the commission of the error, consistent with the observed timing of the error-related negativity (ERN) (Yeung et al., 2004).

Another model (Holroyd et al., 2005) implements error detection in the framework of reinforcement learning algorithms. Errors are detected as deviations from temporal difference predictions arising from a learned conjunction of state information (reactive), while the predictions themselves are used to exert control on response generation (proactive).

3.3.1. The Response Selection Model—An alternative account of the ERN, the *response selection model of ACC* (Holroyd & Coles, 2002), casts ACC as a control filter that learns which of several motor controllers should be given authority. In contrast to conflict theory, which suggests that ACC is a fairly static calculator of conflict, ACC in the response-selection model is adaptive: prior to the beginning of an experiment, ACC has no information about the nature of the task or appropriate responses, and must learn which of several potential responses is required (although task relevant information may be encoded in the model weights before learning – e.g., (Holroyd et al., 2005).

The mechanism of learning in the response selection model is a temporal difference error which updates weights between candidate motor responses and the control filter. Initially, the ERN results from magnitude of the negative temporal difference error driven by unexpected error-related feedback. As learning in the model proceeds, however, the model generates a prediction error following incorrect responses as well as in response to stimuli which predict errors. Physiologically, ascending projections to ACC from midbrain dopamine neurons may actively inhibit firing in ACC; transient depressions below baseline firing rates of midbrain neurons may disinhibit ACC, giving rise to the ERN. While this hypothesis has been implemented in other models of cognitive control (Brown & Braver, 2007; Brown & Braver, 2005), and evidence for the interaction between ACC and midbrain has been observed in imaging studies (Alexander & Brown, In Press) it is possible that ACC instead inhibits dopamine neurons via descending striatal projections (Frank, D’Lauro, & Curran, 2007). If dopamine trains the ACC to perform monitoring and control functions, then this leaves open the question of how dopamine cells might learn to signal errors.

3.3.2. Actor-critic models—The response selection model is one of a class of models using an actor-critic architecture, which is intimately linked to reinforcement learning and dopamine activity (Brown, Bullock, & Grossberg, 2004; Frank, Loughry, & O’Reilly R, 2001; Montague, Dayan, Person, & Sejnowski, 1995; Schultz et al., 1997; Suri & Schultz, 1999; Sutton, 1988). In these models, one component, the critic, implements a version of temporal difference learning. The critic monitors the task state and calculates both a prediction of future reward and subsequent deviations (better or worse) from these predictions. The temporal difference signal is used to update both the critic’s prediction of reward and the weights between stimulus units to response units in the actor component. The actor component learns a mapping between inputs and responses that lead to rewarding states. The basal ganglia are frequently referred to by these models as the loci of action selection in the brain, and the activity of dopamine cells of the basal ganglia is thought to correspond roughly with the prediction error generated by the critic component (Schultz et al., 1997), while the direct and indirect pathways are thought to correspond roughly with the actor component (Brown et al., 2004). While actor-critic models typically learn a simple

mapping between stimulus and response for a single task, and thus seem to involve only a limited form of cognitive control, various models have proposed mechanisms by which cognitive control might be implemented in an actor-critic model, either through maintenance of working memory (Beiser & Houk, 1998; Frank et al., 2001), or through control of top-down attentional processes (Alexander, 2007), or by adjustment of meta-parameters governing response selection (Doya, 2002).

Implicit in the actor-critic architecture is a learned expectation of future events, and especially what will be the most rewarding course of action. The Posner attentional orienting task (Posner, 1980) provides an early example of learned expectation. In the Posner task, a prior cue predicts with probability 0.8 (for instance) which of two locations will shortly reveal a movement target. A valid cue that correctly predicts subsequent target appearance leads to improved performance. More recently, models of performance monitoring have been proposed which examine behavior and potential brain function during shifts in the underlying contingencies of a task. Yu & Dayan (2005), using a Bayesian statistical framework, show how signals corresponding with *expected* and *unexpected uncertainty* can be used to drive shifts in behavior in a stochastic environment. Their model suggests that the neuromodulators acetylcholine and noradrenaline reflect, respectively, uncertainty about the validity of a cue, and uncertainty about the identity of a cue in an extended version of the Posner task. Similarly, Behrens et al. (2007) implemented a Bayesian learning algorithm which tracked volatility in a decision making task in which the underlying probability of a positive outcome shifted over the course of an experimental session. A model-based analysis of fMRI data showed that increased activity in ACC corresponded with higher estimated volatility in the Bayesian model. Furthermore, the learning rate for human participants (estimated by fitting a delta rule model to behavioral data) also tracked environmental shifts. These computational model results suggest that the need for increased control may be signaled by violations of learned reinforcement contingencies.

4. A synthesis: Response-outcome prediction as a basis of performance monitoring

The monitoring and control models reviewed above collectively account for a large corpus of empirical results, but it remains unclear whether or how these disparate models may form a unified whole. Toward this end, we propose a new theoretical synthesis, the PRO (Prediction of Response-Outcome) theory, which casts performance monitoring in an actor-critic framework. According to this theory, one component of the mPFC predicts the likelihood of the various possible (i.e., previously experienced) outcomes of an action, both good and bad. The other component of the mPFC compares the actual vs. the expected outcomes, generating activity that signals when a discrepancy occurs (Figure 2). Discrepancy signals in turn train the actor to generate better predictions of the likely outcomes of an action. The PRO model reinterprets mPFC effects as reflecting a prediction of how likely an adverse outcome is, but more generally, anticipated outcomes may be desirable as well as undesirable. There is evidence that these functions may be carried out at least in part in mPFC, particularly in dorsal ACC and pre-supplementary motor area. There may be multiple possible outcomes of a given action, each predicted in mPFC with some probability (Quintana, Wong, Ortiz-Portillo, Marder, & Mazziotta, 2004). One proposed function of mPFC is to weigh the potential benefits against the potential costs of an action (Kennerley, Walton, Behrens, Buckley, & Rushworth, 2006), biasing action toward the greatest payoff with the least amount of effort (Botvinick, 2007; Botvinick & Rosen, 2008; Croxson, Walton, O'Reilly, Behrens, & Rushworth, 2009; Kennerley, Dahmubed, Lara, & Wallis, 2009) or risk (Brown & Braver, 2007). Those with greater sensitivity to actual reward or punishment may be correspondingly sensitive to predictions of rewarding or aversive outcomes, respectively (Brown & Braver, 2008; Frank et al., 2005; Hewig et al.,

2006; Klein et al., 2007). Neurophysiological studies of awake, behaving monkeys suggest a variety of mPFC signals related to responses, response cues, and their predicted outcomes (Isomura, Ito, Akazawa, Nambu, & Takada, 2003; Matsumoto, Suzuki, & Tanaka, 2003; Olson & Gettner, 2002), as well as unexpected reward and error signals (Ito et al., 2003).

The second component of the PRO model signals discrepancies between actual and expected outcomes, as has been suggested by findings in monkey ACC (Ito et al., 2003). The outcome predictions are timed using an adaptive timing mechanism similar to a tapped delay line structure, as has been proposed in models of reinforcement learning and dopamine signaling (Brown et al., 1999; Schultz et al., 1997). The timed nature of the outcome expectancies allows discrepancies to be detected at the moment when an outcome is actually expected to occur. In this manner, if a correct outcome is most likely but an error occurs, the familiar ERN or feedback ERN is obtained at the earliest time an error can be detected based on available information (Falkenstein et al., 1991; Gehring et al., 1990; Holroyd & Coles, 2002). But what if an error is the most likely outcome, and a correct response unexpectedly occurs? More recent ERP evidence suggests that the mPFC may not be limited to the detection of errors, but it may also signal surprisingly positive outcomes, as suggested by a recent ERP study (Oliveira, McDonald, & Goodman, 2007). In particular, there are several distinct ERP components localized to medial frontal sources that may signal unexpected positive or negative events, including the P300 and N200, as well as a separate feedback correct-related positivity (Holroyd & Krigolson, 2007; Holroyd, Pakzad-Vaezi, & Krigolson, 2008; Yeung, Holroyd, & Cohen, 2005). More recently, ACC activity has also been found with fMRI and unexpected wins in a gambling task (Jessup, Busemeyer, & Brown, In Press). This is consistent with the possibility that ACC may not only predict and detect errors, but it may also predict and detect desirable outcomes. ACC has also been observed to become active in conjunction with low-frequency responses (Braver, Barch, Gray, Molfese, & Snyder, 2001), suggesting that ACC may signal unexpected or unusual behaviors. This casts the ACC in a more general role of detecting discrepancies between anticipated and actual outcomes (Ito et al., 2003), whether desirable or undesirable.

Notably, the PRO model learns associations between the outcome of responses rather than the outcome of stimuli. Other studies have demonstrated that mPFC is most activated when responses must be chosen and their outcome monitored (Walton, Devlin, & Rushworth, 2004), consistent with an instrumental rather than classical conditioning paradigm. On the other hand, orbitofrontal regions are more active when the outcome is cued by a stimulus but cannot be controlled by instrumental responding (Schoenbaum, Setlow, Saddoris, & Gallagher, 2003; Walton et al., 2004).

4.1. PRO theory compared with previous models

Many aspects of the PRO theory are indebted to previous modeling efforts in reinforcement learning and motor control. The nature of the interaction between the outcome prediction and discrepancy detection is akin to formulations of temporal difference learning (Sutton, 1988; Sutton & Barto, 1981). The outcome predictions correspond to an actor (albeit at the cognitive rather than motor level) and can signal the value or risk of a planned action, while the discrepancy detection corresponds to a critic and provides a training signal that can update the outcome predictions as needed. Actor-critic architectures (Sutton, 1988; Sutton & Barto, 1981) provide a convenient framework through which predictive and discrepancy signals in the model can be used for driving approach or avoidance behavior (prediction), or by updating action plans (discrepancy). Nonetheless, the similarity ends there as actor-critic architectures learn to generate a response given a stimulus (S-R learning), but the PRO model learns to predict the outcome given a planned response (R-O learning). Furthermore, models of reinforcement learning are typically identified with the basal ganglia (Schultz et

al., 1997), but the PRO theory pertains to the mPFC. The PRO theory is consistent with the mPFC as influencing the initial preparation of cognitive and motor plan-related activity as well as their continued evolution toward response execution. In contrast, the basal ganglia may provide a selection among the already activated plan representations, i.e. to determine which of the already planned actions will be allowed to execute, based on reinforcement history (Brown et al., 2004).

This framework also diverges from typical formulations of reinforcement learning in which a scalar value signal is learned for each state or state-action pair. In such formulations, the value signal is a composite prediction of all future (positive and/or negative) outcomes, with a positive value signal indicating a prediction of reward, and a negative value signal indicating a prediction of punishment. The PRO model, in contrast, learns a prediction of distinct outcomes, regardless of valence, and indicates discrepancies between actual and *predicted* outcomes. This differentiates the PRO model from other models (e.g., (Scheffers & Coles, 2000)) which signal discrepancies between actual and *intended* outcomes. However, the PRO model is consistent with reinforcement models (e.g., Daw, Kakade, & Dayan, 2002) which suggest that aversive and appetitive events are represented independently. Furthermore, it generalizes typical reinforcement learning algorithms in that the discrepancy signal (i.e. prediction error) is not restricted to a single scalar value but may encode multiple prediction errors that follow from the corresponding predictions.

There is also a strong similarity between the PRO model in the cognitive domain and forward/inverse models of control in the motor domain (Shadmehr & Wise, 2004; Wolpert & Ghahramani, 2000). The prediction component of the PRO model may be considered a forward model in the cognitive rather than the motor domain, functioning to map a movement command to its predicted consequences with respect to higher goals instead of low level actions. The model is consistent with recent proposals of a hierarchy of low-level and high-level learned forward models of outcome prediction at the motor and higher cognitive levels (Krigolson & Holroyd, 2007; Krigolson, Holroyd, Van Gyn, & Heath, 2008), each of which learns to predict outcomes and detects discrepancies between actual and predicted outcomes.

In many ways, the PRO theory is an extension and generalization of the error likelihood model of mPFC (Brown & Braver, 2005). Similar to the error likelihood model, the PRO theory suggests that the preparation of multiple actions will lead to a greater mPFC activity, regardless of whether the prepared actions are in conflict with each other (Brown, 2009). The PRO theory suggests the ability to predict multiple outcomes including favorable ones, where the error likelihood model only predicts errors. Also, the error likelihood model assumes the existence of an error signal, presumably dopaminergic. The PRO theory includes a comparator component that computes discrepancies between actual and intended outcomes which includes the special case of an error. Therefore the PRO theory casts the mPFC as not necessarily dependent on dopaminergic error signals to train outcome likelihood representations, consistent with some reports that have called into question the nature of signals from the dopaminergic midbrain to the mPFC (Frank et al., 2007). The PRO theory suggests several predictions beyond those of the error likelihood model. First, it suggests that unexpectedly correct as well as unexpectedly incorrect outcomes will lead to mPFC activation (Jessup et al., In Press). Second, it suggests that mPFC represents the timing as well as the nature of anticipated outcomes, and therefore that an unexpected delay of the expected outcome will yield a discrepancy signal.

4.2. Role of the PRO theory in cognitive functions

The signals generated according to the PRO theory may provide two valuable sources of information to the rest of the brain. The first is an advance prediction of the potential cost or

risk of an already planned action, which may lead to a veto of an already-planned but not yet executed action (Brass & Haggard, 2007), thus providing a “final predictive check” and possible cancellation of a planned action (Haggard, 2008). Second, the ability to predict outcomes means that not only could the medial PFC detect surprising occurrences, but it might also detect surprising non-occurrences, such as the absence of an expected reward. Discrepancies such as the surprising absence of an expected reward may signal the need for an update of a strategy. Consistent with this, the ACC provides signals that are necessary (Shima & Tanji, 1998) and sufficient (Bush et al., 2002; Procyk, Tanaka, & Joseph, 2000; Shima & Tanji, 1998) to signal a need to switch strategies due to reward omission. According to the PRO theory, activation that represents an expected occurrence is inhibited by a signal representing the actual occurrence of the expected outcome. If the outcome does not occur, the activation representing the anticipated outcome is unopposed, leading to activity that reflects the unexpected non-occurrence of a predicted event (Amador, Schlag-Rey, & Schlag, 2000). This activity signals the detection of a task switch or signals changes in the validity of a cue. Thus, the PRO theory provides an account of how task switches occur when reward contingencies change without an explicit cue, and how mPFC may signal the unexpected *absence* of reward that indicates a task switch is required (Bush et al., 2002; Shima & Tanji, 1998).

How might the outcome prediction and discrepancy signals actually influence decisions in progress? There are two ways. First, the timed outcome prediction signals are exactly the kind of signals needed to suppress dopamine cells at the time of expected reward (Hollerman, Tremblay, & Schultz, 1998; Schultz, 1998). Furthermore, there is anatomical evidence that such mPFC signals may influence nigral dopamine cells via the striatal striosomes (Eblen & Graybiel, 1995). Therefore the timed prediction signals may lead to error-related activation of the mPFC but error-related suppression of nigral dopamine cells. These error modulations of dopamine cells are key components of reinforcement learning signals.

Second, the mPFC predictions of aversive outcomes (as postulated by the PRO model) must be learned through the experience of actual aversive outcomes. The corresponding pauses of dopamine cell activity that correspond to errors are exactly the kind of signals needed to train not only the predictions of aversive outcomes but also the impact of those aversive outcome signals in suppressing the action plan with a predicted aversive outcome. Consistent with this possibility, recent anatomical studies suggest that mPFC provides predominantly inhibitory signals to lateral prefrontal regions involved in action planning (Medalla & Barbas, 2009). In sum, actions with both high expected reward and high risk may be planned initially due to the anticipated reward, but the experience of periodic aversive outcomes may train both the prediction of the aversive outcome and the impact of that prediction on preventing the action plan from being executed.

Aside from direct mPFC to lateral PFC inhibitory projections, another possible target for top-down control signals is the ventral striatum, which is a known recipient of mPFC projections (e.g., Devinsky, Morrell, & Vogt, 1995). Ventral striatum has been shown to become active during anticipation of reward (Schultz, Apicella, Scarnati, & Ljungberg, 1992), anticipation of aversive events (Delgado, Li, Schiller, & Phelps, 2008), as well as to unpredicted rewarding and aversive events. Transient activation of striatal neurons by primary rewards and punishments may allow association of abstract information about the likelihood of an event with its behavioral relevance (i.e., ‘value’).

Overall, the PRO theory suggests a functional role that resembles but is complementary to existing reinforcement learning mechanisms. In general, the role of a reinforcement learning algorithm is to select an action with the highest expected reward for a given task context.

This amounts to a learned stimulus-response (S-R) association. Nonetheless, an action with the highest anticipated reward may not be the best option, because it may entail an unacceptably high risk. In general, action planning may proceed as follows. First, a reinforcement learning system may activate a response via a learned S-R association. Second, the outcome of that response may be predicted prior to execution via a learned response-outcome (R-O) association (Figure 2). Third, if the predicted costs or risks of the planned action outweigh the potential benefits, then the action may be cancelled, or another action substituted. Fourth, the final action is selected and allowed to execute. This kind of decision model is consistent with a broad class of diffusion models of decision-making (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006).

5. Conclusion

Since the seminal qualitative framework of Norman and Shallice (1986) was first proposed, a number of models have been proposed to flesh out the computational basis of executive function. The contention scheduler has been implemented as lateral inhibition in a network of units representing schemas (Cooper & Shallice, 2000). The supervisory attentional system has been divided into two main components. The first is a representation of goals or task rules that direct attention and action in accordance with higher level goals, as in the biased competition model. The second is a performance monitor that detects the need to strengthen or update goal representations. There are currently many proposals regarding the exact computational nature of the performance monitor (Table 1), and further combined empirical and computational modeling work will be required to discriminate among competing models. As a step toward this end, we have developed a new theory of performance monitoring, the PRO theory, that may account for a variety of effects in a single model. These effects include greater predicted activity for error, conflict, and error likelihood, as well as for environmental volatility (non-stationarity) and unexpected outcomes in general, including unexpectedly positive as well as negative outcomes. With regard to the monkey literature, the PRO theory is consistent with a variety of cell types representing combinations of responses and their expected outcomes, as well as the unexpectedness of an outcome. Furthermore, the theory leads to a number of testable predictions. We are actively investigating these predictions with human fMRI, and the results are encouraging. At the same time, we are currently developing a rigorous computational neural model that instantiates the PRO theory to further elucidate its predictions (Alexander & Brown, 2008).

Acknowledgments

Supported by a 2005 NARSAD Young Investigator Award, the Sidney R. Baer, Jr. Foundation, AFOSR FA9550-07-1-0454, R03 DA023462, and R01 DA026457 to JWB. The authors thank the three anonymous reviewers for helpful comments.

References

- Aarts E, Roelofs A, van Turenout M. Anticipatory activity in anterior cingulate cortex can be independent of conflict and error likelihood. *J Neurosci*. 2008; 28(18):4671–4678. [PubMed: 18448644]
- Alexander, W.; Brown, J. A computational neural model of learned response-outcome predictions by anterior cingulate cortex. Program No. 682.21. Washington DC: 2008 Neuroscience Meeting Planner; 2008.
- Alexander WH. Shifting Attention Using a Temporal Difference Prediction Error and High-Dimensional Input. *Adaptive Behavior*. 2007; 15(2):121–133.
- Alexander WH, Brown JW. Competition between learned reward and error outcome predictions in anterior cingulate cortex. *Neuroimage*. (In Press).

- Allport, DA.; Styles, EA.; Hsieh, S. Shifting intentional set: Exploring the dynamic control of tasks. In: Umiltà, MMC., editor. In Attention and Performance 15: Conscious and nonconscious information processing. Cambridge: MIT; 1994.
- Altmann EM. Task switching and the pied homunculus: where are we being led? *Trends Cogn Sci.* 2003; 7(8):340–341. [PubMed: 12907226]
- Altmann EM, Gray WD. An integrated model of cognitive control in task switching. *Psychol Rev.* 2008; 115(3):602–639. [PubMed: 18729594]
- Amador N, Schlag-Rey M, Schlag J. Reward-predicting and reward-detecting neuronal activity in the primate supplementary eye field. *J Neurophysiol.* 2000; 84(4):2166–2170. [PubMed: 11024104]
- Badre D, Wagner AD. Computational and neurobiological mechanisms underlying cognitive flexibility. *Proc Natl Acad Sci U S A.* 2006; 103(18):7186–7191. [PubMed: 16632612]
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat Neurosci.* 2007; 10(9):1214–1221. [PubMed: 17676057]
- Beiser DG, Houk JC. Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *J Neurophysiol.* 1998; 79(6):3168–3188. [PubMed: 9636117]
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev.* 2006; 113(4):700–765. [PubMed: 17014301]
- Botvinick MM. Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. *Cogn Affect Behav Neurosci.* 2007; 7(4):356–366. [PubMed: 18189009]
- Botvinick MM, Braver TS, Barch DM, Carter CS, Cohen JC. Conflict monitoring and cognitive control. *Psychological Review.* 2001; 108:624–652. [PubMed: 11488380]
- Botvinick MM, Nystrom L, Fissel K, Carter CS, Cohen JD. Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature.* 1999; 402(6758):179–181. [PubMed: 10647008]
- Botvinick MM, Rosen ZB. Anticipation of cognitive demand during decision-making. *Psychol Res.* 2008
- Brass M, Haggard P. To do or not to do: the neural signature of self-control. *J Neurosci.* 2007; 27(34):9141–9145. [PubMed: 17715350]
- Braver TS, Barch DM, Gray JR, Molfese DL, Snyder A. Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. *Cereb Cortex.* 2001; 11(9):825–836. [PubMed: 11532888]
- Braver, TS.; Cohen, JD. On the control of control: The role of dopamine in regulating prefrontal function and working memory. In: Monsell, S.; Driver, J., editors. Attention and Performance XVIII. Cambridge, MA: MIT Press; 2000. p. 713-738.
- Braver, TS.; Gray, JR.; Burgess, GC. Explaining the many varieties of working memory variation: Dual mechanisms of cognitive control. In: Conway, CJA.; Kane, M.; Miyake, A.; Towse, J., editors. Variation of working memory. Oxford: Oxford University Press; 2007.
- Brown J, Braver TS. Risk prediction and aversion by anterior cingulate cortex. *Cog Aff Behav Neurosci.* 2007; 7(4):266–277.
- Brown J, Bullock D, Grossberg S. How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *Journal of Neuroscience.* 1999; 19(23):10502–10511. [PubMed: 10575046]
- Brown J, Reynolds J, Braver TS. A Computational model of fractionated conflict-control mechanisms in task switching. *Cognitive Psychology.* 2007; 55:37–85. [PubMed: 17078941]
- Brown JW. Conflict effects without conflict in anterior cingulate cortex: multiple response effects and context specific representations. *Neuroimage.* 2009; 47(1):334–341. [PubMed: 19375509]
- Brown JW, Braver TS. Learned Predictions of Error Likelihood in the Anterior Cingulate Cortex. *Science.* 2005; 307(5712):1118–1121. [PubMed: 15718473]
- Brown JW, Braver TS. A computational model of risk, conflict, and individual difference effects in the anterior cingulate cortex. *Brain Res.* 2008; 1202:99–108. [PubMed: 17707352]
- Brown JW, Bullock D, Grossberg S. How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Netw.* 2004; 17(4):471–510. [PubMed: 15109680]

- Brunel N, Wang XJ. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J Comput Neurosci*. 2001; 11(1):63–85. [PubMed: 11524578]
- Bush G, Vogt BA, Holmes J, Dale AM, Greve D, Jenike MA. Dorsal anterior cingulate cortex: a role in reward-based decision making. *PNAS*. 2002; 99(1):507–512. [PubMed: 11752394]
- Carter CS, Braver TS, Barch DM, Botvinick MM, Noll DC, Cohen JD. Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*. 1998; 280:747–749. [PubMed: 9563953]
- Chadderton GL, Sporns O. A large-scale neurocomputational model of task-oriented behavior selection and working memory in prefrontal cortex. *J Cogn Neurosci*. 2006; 18(2):242–257. [PubMed: 16494684]
- Cohen JD, Dunbar K, McClelland JL. On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*. 1990; 97(3):332–361. [PubMed: 2200075]
- Cooper R, Shallice T. Contention scheduling and the control of routine activities. *Cognitive neural psychology*. 2000; 17(4):297–338.
- Crosson PL, Walton ME, O'Reilly JX, Behrens TE, Rushworth MF. Effort-based cost-benefit valuation and the human brain. *J Neurosci*. 2009; 29(14):4531–4541. [PubMed: 19357278]
- Daw ND, Kakade S, Dayan P. Opponent interactions between serotonin and dopamine. *Neural Netw*. 2002; 15(4–6):603–616. [PubMed: 12371515]
- de Jong, R. Task switching and multitask performance. In: Monsell, S.; Driver, J., editors. *Control of cognitive processes: Attention and performance XVIII*. Cambridge, MA: MIT Press; 2000.
- De Pisapia N, Braver TS. A model of dual control mechanisms through anterior cingulate and prefrontal cortex interactions. *Neurocomputing*. 2006; 69(10–12):1322–1326.
- Deco G, Rolls ET. Attention and working memory: a dynamical model of neuronal activity in the prefrontal cortex. *Eur J Neurosci*. 2003; 18(8):2374–2390. [PubMed: 14622200]
- Deco G, Rolls ET. Neurodynamics of biased competition and cooperation for attention: a model with spiking neurons. *J Neurophysiol*. 2005; 94(1):295–313. [PubMed: 15703227]
- Delgado MR, Li J, Schiller D, Phelps EA. The role of the striatum in aversive learning and aversive prediction errors. *Philos Trans R Soc Lond B Biol Sci*. 2008; 363(1511):3787–3800. [PubMed: 18829426]
- Desimone R, Duncan J. Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*. 1995; 18:193–222.
- Devinsky O, Morrell MJ, Vogt BA. Contributions of anterior cingulate cortex to behaviour. *Brain*. 1995; 118(Pt 1):279–306. [PubMed: 7895011]
- Doya K. Metalearning and neuromodulation. *Neural Netw*. 2002; 15(4–6):495–506. [PubMed: 12371507]
- Durstewitz D, Kelc M, Gunturkun O. A neurocomputational theory of the dopaminergic modulation of working memory functions. *Journal of Neuroscience*. 1999; 19(7):2807–2822. [PubMed: 10087092]
- Durstewitz D, Seamans JK, Sejnowski TJ. Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *Journal of Neurophysiology*. 2000; 83(3):1733–1750. [PubMed: 10712493]
- Eblen F, Graybiel AM. Highly restricted origin of prefrontal cortical inputs to striosomes in the macaque monkey. *J Neurosci*. 1995; 15(9):5999–6013. [PubMed: 7666184]
- Eriksen BA, Eriksen CW. Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*. 1974; 16(1):143–149.
- Falkenstein M, Hohnsbein J, Hoorman J, Blanke L. Effects of crossmodal divided attention on late ERP components: II. Error processing in choice reaction tasks. *Electroencephalography and Clinical Neurophysiology*. 1991; 78:447–455. [PubMed: 1712280]
- Fan J, Kolster R, Ghajar J, Suh M, Knight RT, Sarkar R, et al. Response anticipation and response conflict: an event-related potential and functional magnetic resonance imaging study. *J Neurosci*. 2007; 27(9):2272–2282. [PubMed: 17329424]

- Frank M, Loughry B, O'Reilly RC. Interactions between the frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience*. 2001; 1(2):137–160.
- Frank MJ, D'Lauro C, Curran T. Cross-task individual differences in error processing: neural, electrophysiological, and genetic components. *Cogn Affect Behav Neurosci*. 2007; 7(4):297–308. [PubMed: 18189003]
- Frank MJ, Worocho BS, Curran T. Error-related negativity predicts reinforcement learning and conflict biases. *Neuron*. 2005; 47(4):495–501. [PubMed: 16102533]
- Gehring WJ, Coles MGH, Meyer DE, Donchin E. The error-related negativity: An event-related potential accompanying errors. *Psychophysiology*. 1990; 27:S34.
- Gemba H, Sasaki K, Brooks VB. 'Error' potentials in limbic cortex (anterior cingulate area 24) of monkeys during motor learning. *Neurosci Lett*. 1986; 70(2):223–227. [PubMed: 3774229]
- Gilbert SJ, Shallice T. Task switching: a PDP model. *Cognit Psychol*. 2002; 44(3):297–337. [PubMed: 11971634]
- Haggard P. Human volition: towards a neuroscience of will. *Nature Reviews Neurosci*. 2008; 9:934–946.
- Herd SA, Banich MT, O'Reilly RC. Neural mechanisms of cognitive control: an integrative model of stroop task performance and fMRI data. *J Cogn Neurosci*. 2006; 18(1):22–32. [PubMed: 16417680]
- Hewig J, Hagemann D, Seifert J, Naumann E, Bartussek D. The relation of cortical activity and BIS/BAS on the trait level. *Biol Psychol*. 2006; 71(1):42–53. [PubMed: 16360880]
- Hollerman JR, Tremblay L, Schultz W. Influence of reward expectation on behavior-related neuronal activity in primate striatum. *Journal of Neurophysiology*. 1998; 80(2):947–963. [PubMed: 9705481]
- Holroyd CB, Coles MG. The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psych Rev*. 2002; 109(4):679–709.
- Holroyd CB, Krigolson OE. Reward prediction error signals associated with a modified time estimation task. *Psychophysiology*. 2007; 44(6):913–917. [PubMed: 17640267]
- Holroyd CB, Pakzad-Vaezi KL, Krigolson OE. The feedback correct-related positivity: sensitivity of the event-related brain potential to unexpected positive feedback. *Psychophysiology*. 2008; 45(5):688–697. [PubMed: 18513364]
- Holroyd CB, Yeung N, Coles MG, Cohen JD. A mechanism for error detection in speeded response time tasks. *J Exp Psychol Gen*. 2005; 134(2):163–191. [PubMed: 15869344]
- Isomura Y, Ito Y, Akazawa T, Nambu A, Takada M. Neural coding of "attention for action" and "response selection" in primate anterior cingulate cortex. *J Neurosci*. 2003; 23(22):8002–8012. [PubMed: 12954861]
- Ito S, Stuphorn V, Brown J, Schall JD. Performance Monitoring by Anterior Cingulate Cortex During Saccade Countermanding. *Science*. 2003; 302:120–122. [PubMed: 14526085]
- Jersild AT. Mental set and shift. *Archives of Psychology*. 1927; (89):81.
- Jessup RK, Bussemeyer JR, Brown JW. Error effects in anterior cingulate cortex reverse when error likelihood is high. *J Neurosci*. (In Press).
- Jones AD, Cho R, Nystrom LE, Cohen JD, Braver TS. A computational model of anterior cingulate function in speeded response tasks: Effects of frequency, sequence, and conflict. *Cog Aff Behav Neurosci*. 2002; 2(4):300–317.
- Kennerley SW, Dahmubed AF, Lara AH, Wallis JD. Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci*. 2009; 21(6):1162–1178. [PubMed: 18752411]
- Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF. Optimal decision making and the anterior cingulate cortex. *Nat Neurosci*. 2006; 9(7):940–947. [PubMed: 16783368]
- Klein TA, Endrass T, Kathmann N, Neumann J, von Cramon DY, Ullsperger M. Neural correlates of error awareness. *Neuroimage*. 2007; 34(4):1774–1781. [PubMed: 17185003]
- Kohonen T. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*. 1982; 43:59–69.

- Krigolson OE, Holroyd CB. Hierarchical error processing: different errors, different systems. *Brain Res.* 2007; 1155:70–80. [PubMed: 17498670]
- Krigolson OE, Holroyd CB, Van Gyn G, Heath M. Electroencephalographic correlates of target and outcome errors. *Exp Brain Res.* 2008; 190(4):401–411. [PubMed: 18629483]
- Laming D. Choice reaction performance following an error. *Acta Psychologica.* 1979; 43:199–224.
- MacDonald AW, Cohen JD, Stenger VA, Carter CS. Dissociating the role of the dorsolateral prefrontal cortex and anterior cingulate cortex in cognitive control. *Science.* 2000; 288:1835–1838. [PubMed: 10846167]
- Magno E, Foxe JJ, Molholm S, Robertson IH, Garavan H. The anterior cingulate and error avoidance. *J Neurosci.* 2006; 26(18):4769–4773. [PubMed: 16672649]
- Matsumoto K, Suzuki W, Tanaka K. Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science.* 2003; 301(5630):229–232. [PubMed: 12855813]
- Medalla M, Barbas H. Synapses with inhibitory neurons differentiate anterior cingulate from dorsolateral prefrontal pathways associated with cognitive control. *Neuron.* 2009; 61(4):609–620. [PubMed: 19249280]
- Meyer DE, Kieras DE. A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review.* 1997; 104:3–65. [PubMed: 9009880]
- Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience.* 2001; 21:167–202.
- Montague PR, Dayan P, Person C, Sejnowski TJ. Bee foraging in uncertain environments using predictive hebbian learning. *Nature.* 1995; 377(6551):725–728. [PubMed: 7477260]
- Muly EC 3rd, Szigeti K, Goldman-Rakic PS. D1 receptor in interneurons of macaque prefrontal cortex: distribution and subcellular localization. *J Neurosci.* 1998; 18(24):10553–10565. [PubMed: 9852592]
- Nieuwenhuis S, Monsell S. Residual costs in task switching: testing the failure-to-engage hypothesis. *Psychon Bull Rev.* 2002; 9(1):86–92. [PubMed: 12026956]
- Nieuwenhuis S, Schweizer T, Mars RB, Botvinick MM, Hajcak G. Error-likelihood prediction in the medial frontal cortex: A critical evaluation. *Cereb Cortex.* 2007; 17:1570–1581. [PubMed: 16956979]
- Niki H, Watanabe M. Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res.* 1979; 171(2):213–224. [PubMed: 111772]
- Norman, D.; Shallice, T. Attention to action: Willed and automatic control of behavior. In: Davidson, R.; Schwartz, G.; Shapiro, D., editors. *Consciousness and Self Regulation: Advances in Research and Theory.* Vol. 4. New York: Plenum; 1986.
- Oliveira FT, McDonald JJ, Goodman D. Performance Monitoring in the Anterior Cingulate is Not All Error Related: Expectancy Deviation and the Representation of Action-Outcome Associations. *J Cogn Neurosci.* 2007
- Olson CR, Gettner SN. Neuronal activity related to rule and conflict in macaque supplementary eye field. *Physiol Behav.* 2002; 77(4–5):663–670. [PubMed: 12527016]
- Paulus MP, Frank LR. Anterior cingulate activity modulates nonlinear decision weight function of uncertain prospects. *Neuroimage.* 2006; 30(2):668–677. [PubMed: 16321546]
- Posner MI. Orienting of attention. *Quarterly Journal of Experimental Psychology.* 1980; 32:3–25. [PubMed: 7367577]
- Posner, MI.; DiGirolamo, G. Conflict, target detection and cognitive control. In: Parasuraman, R., editor. *The Attentive Brain.* Cambridge: MIT Press; 1998.
- Procyk E, Tanaka YL, Joseph JP. Anterior ingulate activity during routine and non-routine sequential behaviors in macaques. *Nature Neuroscience.* 2000; 3(5):502–508.
- Quintana J, Wong T, Ortiz-Portillo E, Marder SR, Mazzotta JC. Anterior cingulate dysfunction during choice anticipation in schizophrenia. *Psychiatry Res.* 2004; 132(2):117–130. [PubMed: 15598546]
- Rabbitt PMA. Errors and error correction in choice-response tasks. *Journal of Experimental Psychology.* 1966; 71(2):264–272. [PubMed: 5948188]

- Rey, HG.; Lew, SE.; Zanutto, BS. Dopamine and norepinephrine modulation of cortical and subcortical dynamics during visuomotor learning. In: Tseng, K.; Atzori, M., editors. Monoaminergic modulation of cortical excitability. New York: Springer; 2007.
- Reynolds JR, Braver TS, Brown J, Stigchel S. Computational and Neural Mechanisms of Task Switching. *Neurocomputing*. 2006; 69(10):1332–1336.
- Rogers RD, Monsell S. Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*. 1995; 124(2):207–231.
- Rougier NP, Noelle DC, Braver TS, Cohen JD, O'Reilly RC. Prefrontal cortex and flexible cognitive control: rules without symbols. *Proc Natl Acad Sci U S A*. 2005; 102(20):7338–7343. [PubMed: 15883365]
- Scheffers MK, Coles MG. Performance monitoring in a confusing world: error-related brain activity, judgments of response accuracy, and types of errors. *J Exp Psychol Hum Percept Perform*. 2000; 26(1):141–151. [PubMed: 10696610]
- Schoenbaum G, Setlow B, Saddoris MP, Gallagher M. Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. *Neuron*. 2003; 39(5):855–867. [PubMed: 12948451]
- Schultz W. Predictive reward signal of dopamine neurons. *J Neurophysiol*. 1998; 80(1):1–27. [PubMed: 9658025]
- Schultz W, Apicella P, Scarnati E, Ljungberg T. Neuronal activity in monkey ventral striatum related to the expectation of reward. *J Neurosci*. 1992; 12(12):4595–4610. [PubMed: 1464759]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Shadmehr, R.; Wise, SP. Motor Learning and Memory for Reaching and Pointing. In: Gazzaniga, M., editor. *The Cognitive Neurosciences III*. 3. Cambridge: MIT Press; 2004.
- Shima K, Tanji J. Role of cingulate motor area cells in voluntary movement selection based on reward. *Science*. 1998; 282:1335–1338. [PubMed: 9812901]
- Sohn MH, Albert MV, Jung K, Carter CS, Anderson JR. Anticipation of conflict monitoring in the anterior cingulate cortex and the prefrontal cortex. *Proc Natl Acad Sci U S A*. 2007; 104(25):10330–10334. [PubMed: 17563353]
- Steinhauser M, Maier M, Hubner R. Modeling behavioral measures of error detection in choice tasks: response monitoring versus conflict monitoring. *J Exp Psychol Hum Percept Perform*. 2008; 34(1):158–176. [PubMed: 18248146]
- Stroop JR. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*. 1935; 18:643–662.
- Suri RE, Schultz W. A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience*. 1999; 91(3):871–890. [PubMed: 10391468]
- Sutton R. Learning to predict by the method of temporal difference. *Machine Learning*. 1988; 3:9–44.
- Sutton RS, Barto AG. Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*. 1981; 88(2):135–170. [PubMed: 7291377]
- van Veen V, Cohen JD, Botvinick MM, Stenger VA, Carter CS. Anterior cingulate cortex, conflict monitoring, and levels of processing. *Neuroimage*. 2001; 14(6):1302–1308. [PubMed: 11707086]
- Walton ME, Devlin JT, Rushworth MF. Interactions between decision making and performance monitoring within prefrontal cortex. *Nat Neurosci*. 2004; 7(11):1259–1265. [PubMed: 15494729]
- Wolpert DM, Ghahramani Z. Computational principles of movement neuroscience. *Nat Neurosci*. 2000; 3(Suppl):1212–1217. [PubMed: 11127840]
- Yeung N, Cohen JD, Botvinick MM. The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol Rev*. 2004; 111(4):931–959. [PubMed: 15482068]
- Yeung N, Holroyd CB, Cohen JD. ERP correlates of feedback and reward processing in the presence and absence of response choice. *Cereb Cortex*. 2005; 15(5):535–544. [PubMed: 15319308]
- Yu AJ, Dayan P. Uncertainty, neuromodulation, and attention. *Neuron*. 2005; 46(4):681–692. [PubMed: 15944135]




Figure 1. The Biased Competition model (Miller & Cohen, 2001) as applied to the Stroop task, with a conflict-monitoring loop (Botvinick et al., 2001) driving control.





Figure 2. ACC Response-Outcome model

Planned responses activate learned response-outcome predictions. These predicted outcome signals can feed back to amend or veto a planned action. Once an action is generated, the actual outcome (the movement itself or the feedback from the environment) is compared against the intended outcome, and any discrepancy leads to an update of the learned response-outcome predictions.

Table 1
Summary of error monitoring theories

Consider two representations of cued responses with continuous activities A and B in $[0, 1]$ that lead to binary 0 or 1 responses X and/or Y, representing correct and error responses, respectively. For concreteness, A may correspond to a valid cue in the Stroop task and drive correct response X, while B may represent an incongruent distractor that drives incorrect response Y.



Model	Method	Error Detection Pseudocode	References
Discrepancy	Subtractive (reactive)	$Y - X + 1$	(Scheffers & Coles, 2000)
Response Selection	Learned conjunction (reactive, proactive)	$V_t(A, Y) - V_{t-1}(A)$	(Holroyd, Yeung, Coles, & Cohen, 2005)
Correction detection	Sequence detection (reactive)	X after Y?	(Steinhauser, Maier, & Hubner, 2008)
Conflict	Multiplicative (reactive, proactive)	$A * B$	(Botvinick et al., 2001; Yeung, Cohen, & Botvinick, 2004)
Error likelihood	Learned Probability (proactive)	$P(Y A, B)$	(Brown & Braver, 2005)
PRO	Learned prediction, comparison (reactive, proactive)	$\text{abs}(Y - P(Y A))$	(Alexander & Brown, 2008)